

Broadcasting vs. Mixing and Information Dissemination on Cayley Graphs [★]

Robert Elsässer and Thomas Sauerwald

University of Paderborn
Institute for Computer Science
33102 Paderborn, Germany
{elsa,sauerwal}@upb.de

Abstract. One frequently studied problem in the context of information dissemination in communication networks is the broadcasting problem. In this paper, we study the following randomized broadcasting protocol: At some time t an information r is placed at one of the nodes of a graph G . In the succeeding steps, each informed node chooses one neighbor, independently and uniformly at random, and informs this neighbor by sending a copy of r to it.

First, we consider the relationship between randomized broadcasting and random walks on graphs. In particular, we prove that the runtime of the algorithm described above is upper bounded by the corresponding mixing time, up to a logarithmic factor. One key ingredient of our proofs is the analysis of a continuous-type version of the afore mentioned algorithm, which might be of independent interest. Then, we introduce a general class of Cayley graphs, including (among others) Star graphs, Transposition graphs, and Pancake graphs. We show that randomized broadcasting has optimal runtime on all graphs belonging to this class. Finally, we develop a new proof technique by combining martingale tail estimates with combinatorial methods. Using this approach, we show the optimality of our algorithm on another Cayley graph and obtain new knowledge about the runtime distribution on several Cayley graphs.

1 Introduction

Models and Motivation: The study of information spreading in large networks has various fields of application in distributed computing. Consider for example the maintenance of replicated databases on name servers in a large network [9]. There are updates injected at various nodes, and these updates must be propagated to all the nodes in the network. In each step, a processor and its neighbors check whether their copies of the database agree, and if not, they perform the

[★] 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 EU within the 6th Framework Programme under contract 001907 “Dynamically Evolving, Large Scale Information Systems” (DELIS).

necessary updates. In order to be able to let all copies of the database converge to the same content, efficient broadcasting algorithms have to be developed.

There is an enormous amount of experimental and theoretical study of broadcasting algorithms in various models and on different networks. Several (deterministic and randomized) algorithms have been developed and analyzed. In this paper we only concentrate on the efficiency of randomized broadcasting and study the runtime of the *push model* [9] defined as follows. Place at some time t an information r on one of the nodes of a graph $G = (V, E)$. Then, in each succeeding time step, any *informed* vertex forwards a copy of r to a communication partner over an incident edge selected independently and uniformly at random. The advantage of randomized broadcasting is in its inherent robustness against several kinds of failures and dynamical changes compared to deterministic schemes that either need substantially more time [15] or can tolerate only a relatively small number of faults [19].

In this work we are particularly interested in the runtime of the push algorithm on *Cayley graphs*. A Cayley graph is given by a finite group \mathfrak{G} and a set of generators $S \subseteq \mathfrak{G}$. The vertices are the group elements and there is an edge from an element a to an element b iff if $a = bs$ in \mathfrak{G} for a generator $s \in S$.

This group theoretic model is often used for designing, analyzing, and improving symmetric interconnection networks. In designing interconnection networks, the objective is to construct large (vertex symmetric) graphs with small degree and diameter, high connectivity, and simple routing algorithms. Prominent examples that offer all these properties together are the Hypercube and the Star graph [1]. Other examples are explicit constructions of so called Ramanujan graphs, which are also obtained by using this group theoretic model [20].

An advantage of analyzing Cayley graphs is that properties can be proved for the class as a whole, instead of proving some property for each network independently. Moreover, even for specific networks we can often derive properties algebraically and interpret them graph theoretically.

Related Work: Most papers dealing with randomized broadcasting analyze the runtime of the push algorithm in different graph classes. Pittel [21] proved that with a certain probability an information is spread to all nodes by the push algorithm within $\log_2 N + \ln N + O(1)$ steps in a complete graph K_N . Feige et al [14] determined asymptotically optimal upper bounds for the runtime of this algorithm on random graphs, hypercubes and bounded degree graphs. In [12] we extended the results to Star graphs [2, 1], i.e. after $O(\log N)$ steps any information is spread to all of the N nodes with high probability¹.

We should also note that several broadcasting models have been analyzed in some scenarios that allow nodes and/or edges to fail during the algorithm is executed (e.g. [18]). Most of these papers deal with the worst case asymptotic behavior of broadcasting algorithms when the failures are governed by an adversary, however, in some papers the random failure scenario is also considered. In

¹ When we write “with high probability” or “w.h.p.” we mean with probability at least $1 - 1/N$. Accordingly, “with constant probability” or “w.c.p.” means with probability at least $1 - O(1) > 0$.

[13] we established a robustness result w.r.t. the push algorithm against random failures in general graphs.

Intuitively, rapid mixing implies fast broadcasting, but there is no (strong) bound on the runtime of the push algorithm, which uses mixing rates of Markov chains. In contrast to the push algorithm, mixing has been extensively studied in the past (e.g. [22, 7, 10, 4]). Thus, one of our goals is to derive efficient bounds on the runtime of the push algorithm by using mixing rates of Markov chains.

There is also a long history of the analysis of Markov chains on Cayley graphs. Consider for example the so called *card shuffling process*. The main question is how many times must a deck of cards be shuffled until it is close to random. Using different shuffling rules, the problem reduces to random walks on certain Cayley graphs. We will give examples for card shuffling procedures in Section 5.

Our Results: The next section contains the basic notations and definitions needed in our further analysis. In Section 3 we show that the runtime of the broadcasting algorithm is upper bounded by the mixing time of the corresponding Markov chain, up to a logarithmic factor. Section 4 contains the introduction of a new class of graphs which contains prominent examples of Cayley graphs. It is shown that the push algorithm has an optimal runtime on all these graphs. Finally, in Section 5 we develop a powerful approach which enables us to extend the optimality results mentioned before. This technique combines Azuma-Hoeffding type bounds with structural analysis of graphs. The last section contains our conclusions and points to some open problems. Due to space limitations, several proofs are omitted in this extended abstract.

2 Notation and Definitions

Let $G = (V(G), E(G))$ denote an unweighted, undirected, simple and connected graph, where $N := |V|$ denotes the size of the graph. In most cases, we will consider families of graphs $G(n) = (V_n, E_n)$, where $|V_n| \rightarrow \infty$ for $n \rightarrow \infty$. By $\text{diam}(G)$ we denote the diameter of G and $N(v)$ is the neighbourhood of some vertex $v \in V(G)$. For an arbitrary vertex $u \in V(G)$, we denote by $N_r(u) := \{v \in V(G) \mid \text{dist}(u, v) \leq r\}$ the r -neighborhood around u . Furthermore, let δ be the minimum and Δ be the maximum degree.

Definition 1. For any graph G and any integer $m \in \{1, \dots, \lfloor N/2 \rfloor\}$ define $E(m) = \min_{X \subseteq V(G), |X|=m} |E(X, X^c)|/|X|$. Here, $E(X, X^c)$ denotes the set of edges connecting X and its complement X^c .

As mentioned in the introduction, in this paper we mainly consider the following randomized broadcasting algorithm (known as the push model [9]): Place at time $t = 0$ an information r on one of the nodes of the graph G . In the succeeding time steps (or rounds) each *informed* vertex forwards a copy of r to a communication partner over an incident edge selected independently and uniformly at random.

This algorithm will be shortly abbreviated by RBA_d , where d indicates that the time steps are discrete numbers (In Section 3 and 5 we will introduce some

slightly modified versions of this algorithm). Throughout this paper, we denote by $I(t)$ the set of informed nodes at time t , and by $H(t)$ the set $V \setminus I(t)$.

Our main objective is to determine how many time steps are required to inform every node of G . Let $\text{RT}_d(G, p) := \min\{t \in \mathbb{N} \mid \Pr[I(t) = V] \geq p\}$ denote the runtime of RBA_d in G , i.e. the number of time steps needed by the push algorithm to inform all vertices of G with probability p . Since every broadcasting algorithm requires $\max\{\log_2 N, \text{diam}(G)\}$ rounds [14], we call RBA_d (asymptotically) optimal on G , if $\text{RT}_d(G, 1 - 1/N) = O(\log N + \text{diam}(G))$.

In the following, we will use basic notation of algebraic graph theory (cf. [5]).

Definition 2. A (directed) Cayley graph $G = (\mathfrak{G}, S)$ is given by a finite group \mathfrak{G} and a generating set $S = \{s_1, \dots, s_n\}$. The set of vertices consists of elements of \mathfrak{G} and there is an directed edge from u to v if and only if there exists a generator $s_i \in S$ such that $u = vs_i$.

If the set of generators is closed under inverses, i.e. $S^{-1} = S$ (which will be always the case in this paper), the resulting Cayley graph $G = (\mathfrak{G}, S)$ can be also viewed as an undirected graph. In the following, \mathfrak{G}_n will be always the symmetric group of n elements, denoted by \mathfrak{S}_n . For any distinct numbers $k_1, \dots, k_i \in [1, n]$ let $\mathfrak{S}_n(k_1, \dots, k_i) := \{\pi(n - i + j) = k_j, j \in \{1, \dots, i\} \mid \pi \in \mathfrak{S}_n\}$.

3 Broadcasting vs. Mixing

In this section we are going to show that rapid mixing implies fast broadcasting. It will be important to consider a slightly different broadcasting algorithm, called RBA_s (s for subimesteps) which is defined as follows.

In this model, the time axis is $\mathbb{T} = \mathbb{N} + \{i/N \mid i \in \{0, \dots, N - 1\}\}$. At such a (sub)timestep $t \in \mathbb{T}$, one node of $V(G)$ is chosen uniformly at random and this node, provided that it is already informed, sends the information to some neighbor, again chosen uniformly at random. This model has the advantage that the waiting times between the transmission of some informed node are geometrically distributed with mean 1 and thus are oblivious. Denote by $\text{RT}_s(G, 1 - 1/N)$ the runtime of this modified broadcasting algorithm. We say that a node $u \in V$ makes a transmission at time t , if node u is chosen by RBA_s at timestep t and sends the information to some neighbor.

The following theorem shows the equivalence of both introduced variants.

Theorem 1. For any G we have $\text{RT}_s(G, 1 - 1/N) = \Theta(\text{RT}_d(G, 1 - 1/N))$.

In order to derive a strong relationship between mixing and broadcasting, we first define the following Markov chain \mathfrak{M} on a graph $G = (V, E)$. \mathfrak{M} has state space $V(G)$, and its transition matrix is given by P where $p_{ii} = 1 - \alpha \deg(i)$, $p_{ij} = \alpha$ if $\{i, j\} \in E(G)$ and $p_{ij} = 0$ otherwise. Hereby, we set $\alpha = 1/(\Delta + 1)$ with Δ being the maximum degree in G . (P also corresponds to the diffusion matrix occurring in load balancing [11].) It is well-known that for our choice of α , the Markov chain \mathfrak{M} is ergodic and reversible [11, 22]. As usual, for any $k \in \mathbb{N}$, P^k denotes the k -step transition matrix.

For two given probability vectors $(\mu_i)_{i=1}^N$ and $(\nu_i)_{i=1}^N$ let

$$\|\mu - \nu\| = \frac{1}{2} \sum_{i=1}^N |\mu_i - \nu_i| = \max_{V' \subseteq V(G)} |\mu_{V'} - \nu_{V'}|$$

be the variation distance of these vectors [10]. Furthermore, we denote by

$$\text{MIX}_{\mathfrak{M}}(G, \epsilon) := \min\{t \in \mathbb{N} \mid \|P^t z - \pi\| \leq \epsilon \text{ for any probability vector } z\},$$

the mixing time (or mixing rate) of \mathfrak{M} . Observe that due to the proper choices of the p_{ii} 's, the vector $(1/N, \dots, 1/N)$ is the stationary distribution corresponding to P . \mathfrak{M} can be viewed as the Markov chain corresponding to a random walk on G , in which the transition probabilities are defined according to P .

Now we define the following random process on the graph G . Assume first that there are N indivisible unit size tokens x_1, \dots, x_N distributed somehow on the nodes of the graph. At each time $t \in \{i + k/N \mid i \in \mathbb{N}, k \in \{0, \dots, N-1\}\}$ we choose one node of the graph, uniformly at random, and one of the tokens on this node is allowed to perform a transition according to the matrix P . Hereby, each token of any node has a priority value, and when a node is chosen, then only the token with highest priority on this node is allowed to perform the transition described above. At the beginning, the tokens on any node u are assigned priority values in the range $[1, l(0, u)]$ arbitrarily, where $l(0, u)$ denotes the *load* (i.e., the number of tokens) on node u at time 0. When a token x_j performs a transition according to P from some node u to node v , then x_j is assigned, after the transition, the lowest priority among all tokens being on v (please note that v and u might coincide).

According to the description above, let $h(t, x_j)$ denote the host of token x_j at time t . Furthermore, let $l(t, u)$ denote the load of any node $u \in V$ at time t .

We are now ready to define another Markov chain \mathfrak{M}' based on the random process described above. \mathfrak{M}' has state space $S(\mathfrak{M}') = \{(l(1), \dots, l(N)) \mid 0 \leq l(i) \in \mathbb{N}, \sum_{i=1}^N l(i) = N\}$, and transition matrix P' , where $p'_{i,j} = \alpha/N$ if there are two states s and s' such that $s = (l(1), \dots, l(i), \dots, l(j), \dots, l(N))$, $s' = (l(1), \dots, l(i) + 1, \dots, l(j) - 1, \dots, l(N))$, where $l(j) \geq 1$, and $\{i, j\} \in E$. Obviously, the Markov chain \mathfrak{M}' simulates the random process described in the previous paragraphs. Since the transition matrix P' is symmetric, the stationary distribution equals the uniform distribution. Thus, the expected number of tokens equals 1 on each node in the stationary state.

Now we use the Markov chains introduced above to show the following.

Theorem 2. *For any graph $G = (V, E)$ it holds*

$$\text{RT}_s(G, 1 - \frac{1}{N}) \leq O\left(\text{MIX}_{\mathfrak{M}}(G, \frac{1}{2N}) \cdot \log N\right).$$

Proof. For simplicity, let $m := \text{MIX}_{\mathfrak{M}}(G, 1/(2N))$. First, we show that if there are $\log N \leq |I(t)| \leq N/2$ informed nodes at timestep t , then there will exist $(1 + \Omega(1))|I(t)|$ informed nodes at timestep $t + m$, w.c.p. In this proof we derive

a relationship between $\text{RT}_s(G, 1 - 1/N)$ and the Markov chain \mathfrak{M}' , and show that by using \mathfrak{M}' the information can be spread in time $O(m \cdot \log N)$ in G .

Now we consider the Markov chain \mathfrak{M} . We assume that there are N tokens distributed according to the stationary distribution of \mathfrak{M}' at some time t . Let I be a fixed, connected set of nodes in G with $\log N \leq |I| \leq N/2$. Let \mathcal{A} be set of tokens lying in I at timestep t , i.e., $\mathcal{A} := \{x_i \mid h(t, x_i) \in I\}$. Since $\mathbf{E}[|\mathcal{A}|] = |I|$, applying the Chernoff bounds [17], we get $\Pr[|\mathcal{A}| \geq |I|/2] \geq 1 - \exp(-\Omega(|I|))$.

We fix some token x_i on one of these nodes, and let this token perform a random walk according to P . Now we know that $\|P^m z - \pi\| \leq 1/(2N)$ for any probability vector z .

Let $D(i)$ denote the host of token x_i at time $t + m$ for some random instance \mathcal{I} of \mathfrak{M} . Then, define

$$\mathcal{B} := \{x_i \in \mathcal{A} \mid D(i) \in H\}, \quad \mathcal{C} := \left\{x_i \in \mathcal{A} \mid |\{D(i) = D(j) \mid j \neq i\}| \leq 32\right\},$$

where $H = V \setminus I$. Due to (1) we have $\Pr[x_i \in \mathcal{B}] \geq \frac{N-|I|}{N} - \frac{1}{2N} \geq \frac{7}{16}$, whenever $|I| \leq N/2$ and $\Pr[x_i \in \mathcal{C}] \leq \binom{N}{32} \left(\frac{3}{2N}\right)^{32} \leq \frac{1}{1024}$. Again, by the Markov inequality we obtain that $\Pr[|\mathcal{B}| \geq 1/4|\mathcal{A}|] \geq 1/4$ and $\Pr[|\mathcal{C}| \geq 31/32|\mathcal{A}|] \geq 31/32$.

Now we consider the walks performed by all tokens according to \mathfrak{M}' , and take into account the delays induced by other tokens. We assume that at time t these tokens are distributed according to the stationary state of \mathfrak{M}' . Let $u_{i,t}, \dots, u_{i,t+m}$ be the nodes visited by some fixed token x_i in steps $t, \dots, t + m$, respectively, according to \mathfrak{M} and instance \mathcal{I} . Let $f(u_{i,k})$ denote the number of time intervals $[j, j + 1]$ in which node $u_{i,k}$ is not chosen by the random process described above while x_i resides on $u_{i,k}$. Since a node is not chosen in time interval $[j, j + 1]$ with probability $(1 - 1/N)^N \approx 1/e$, the expected delay of token x_i is

$$\mathbf{E}[\Delta(i)] \leq \sum_{k=t}^{t+m} \mathbf{E}[l(u_{i,k}) + f(u_{i,k})] = \sum_{k=t}^{t+m} \mathbf{E}\left[\frac{e \cdot l(u_{i,k})}{e-1}\right] \leq \frac{e(m+1)}{e-1},$$

where $l(u_{i,k})$ is the load of node $u_{i,k}$ at the time when token x_i makes a transition to node $u_{i,k}$. Hence, token x_i reaches its destination after $32e(m+1)/(e-1) + m$ rounds, according to \mathfrak{M} and instance \mathcal{I} , with probability at least $31/32$.

Now let $\mathcal{D} := \{x_i \mid \Delta(i) \leq 32e(m+1)/(e-1) \text{ and } x_i \in \mathcal{A}\}$, i.e., the set of tokens of \mathcal{A} which reach their final destination after at most $32e \cdot (m+1)/(e-1) + m$ steps. Since $\mathbf{E}[|\mathcal{D}|] \geq |\mathcal{A}| \cdot 31/32$, the Markov inequality implies that $\Pr[|\mathcal{D}| \geq 13/16|\mathcal{A}|] \geq 5/6$. Putting all together, we get by the union bound

$$\Pr\left[|\mathcal{A}| \geq \frac{1}{2}|I| \wedge |\mathcal{B}| \geq \frac{3}{4}|\mathcal{A}| \wedge |\mathcal{C}| \geq \frac{31}{32}|\mathcal{A}| \wedge |\mathcal{D}| \geq \frac{13}{16}|\mathcal{A}| \right] \geq \frac{1}{32},$$

provided that N is large enough. Since \mathcal{B}, \mathcal{C} and \mathcal{D} are all subsets of \mathcal{A} we have $|\mathcal{B} \cap \mathcal{C}| \geq |\mathcal{A}| - |\mathcal{A} \setminus \mathcal{B}| - |\mathcal{A} \setminus \mathcal{C}| \geq |\mathcal{A}| - \frac{3}{4}|\mathcal{A}| - \frac{1}{32}|\mathcal{A}| - \frac{3}{16}|\mathcal{A}| = \frac{1}{32}|\mathcal{A}|$. Hence, at least $|\mathcal{A}|/32 \cdot 1/32 = |\mathcal{A}|/1024$ nodes of H will host a token of \mathcal{A} within the time interval $[t, t + m + 32e(m+1)/(e-1)]$, with probability $1/32$.

Now we consider $\text{RT}_s(G, 1 - 1/N)$. Since any node in the random process described by \mathfrak{M}' forwards a token (according to P) in some substep iff there is a token on this node, RT_s is able to spread an information faster than the tokens, which perform movements according to \mathfrak{M}' . Hence, $|I(t + O(m))| = (1 + \Omega(1))|I(t)|$, whenever $\log N \leq |I(t)| \leq \frac{N}{2}$.

Similar techniques imply that $|H(t + O(m))| \leq (1 - \Omega(1))|H(t)|$, whenever $\log N \leq |H(t)| \leq \frac{N}{2}$. If $|I(t)| \leq O(\log N)$ or $|I(t)| \geq N - O(\log N)$, then w.c.p. at least one single node becomes informed in some step $t + O(m)$. Applying now the Chernoff bounds [6, 17], we obtain the theorem. \square

However, for $G = K_{N/2} \times C_2$ we have $\text{MIX}(G, 1/(2N)) = \Omega(N)$, but $\text{RT}_d(G, 1 - 1/N) = \Theta(\log N)$ and thus a similar converse of Theorem 2 does not hold.

4 Broadcasting on Cayley Graphs

In this section we will prove that the RBA_d performs optimal on a certain class of Cayley Graphs which includes the Star Graph, Pancake Graph and Transposition Graph.

A vertex $v \in V$ in a graph $G = (V, E)$ is called α -approximated by the set $I(t)$, if $N_\alpha(v) \cap I(t) \neq \emptyset$. Furthermore, a vertex $v \in V$ is called contacted by a node $u \in V$ within some time interval $[a, b]$ (or conversely, u contacts v in time interval $[a, b]$) if there is a path $(u = u_1, u_2, \dots, u_{m-1}, u_m = v)$ in V such that

$$\exists t_1 < t_2 < \dots < t_{m-1} \in [a, b] \subseteq \mathbb{N} : \forall i \in [1, m-1] : u_i \text{ contacts } u_{i+1} \text{ in round } t_i.$$

Now we are ready to state the following theorem.

Theorem 3. *Assume that a family of Cayley graphs $G_n = (\mathfrak{S}_n, S_n)$ has the following properties:*

1. *for any $n \in \mathbb{N}$ it holds that $c_1 n^c \leq d(n) \leq c_2 n^c$, where $d(n)$ denotes the degree of G_n and $c_1, c_2, c \in \Theta(1)$,*
2. *$S_n \subseteq S_{n+1}$ for any $n \in \mathbb{N}$,*
3. *$\text{dist}(\tau, \mathfrak{S}_n(k)) := \min\{\text{dist}(\tau, \tau') \mid \tau' \in \mathfrak{S}_n(k)\} \leq c'$ for any $\tau \in \mathfrak{S}_n$, and $k \in [1, n]$, where c' is a constant,*
4. *$E(m) = \Omega(d(n))$ for any $m = O(n^{c \cdot c'})$.*

Then it holds that

$$\text{RT}_d(G_n, 1 - \frac{1}{N}) \leq O(\log N).$$

Proof. Since any Cayley graph is vertex-transitive [5], we may assume w.l.o.g. that the identity id is informed at the beginning. The proof is divided into two parts. In the first part, we will show that after $t = O(\log N)$ steps it holds for any vertex $w \in V$ that $N_{\alpha n}(w) \cap I(t) \neq \emptyset$, w.h.p., where α is a properly chosen constant. This approximation will consist of $\beta := (1 - \alpha)n$ disjoint phases $\mathcal{P}_1, \dots, \mathcal{P}_\beta$. To simplify notation let $\mathfrak{S}_n(i) := \mathfrak{S}_n(w_{n-i}, \dots, w_n)$. Phase \mathcal{P}_i , $i \in \{1, \dots, \beta\}$, begins when a node of $\mathfrak{S}_n(i - 1)$ becomes informed for the first time, and ends

when the information jumps from the set $\mathfrak{S}_n(i-1) \setminus \mathfrak{S}_n(i)$ to the set $\mathfrak{S}_n(i)$. Let X_i denote the random variable which represents the number of time steps needed by phase \mathcal{P}_i . Now, our goal is to derive an upper bound on X_i for an arbitrary fixed $i \in \{1, \dots, \beta\}$.

First we count the number of steps needed to inform $\Omega(d(n)^{c'})$ vertices of $\mathfrak{S}_n(i-1)$. Since $i \leq (1-\alpha)n$, $\mathfrak{S}_{n-i} \subset \mathfrak{S}_n$, and $d(n-i) \geq c_1(n-i)^c$, each node of $\mathfrak{S}_n(i-1)$ has $\Omega(d(n))$ neighbors in $\mathfrak{S}_n(i-1)$. Due to assumption (4), a constant fraction of these inner edges, incident to nodes of $\mathfrak{S}_n(i-1) \cap I(t)$, are connected to nodes of $\mathfrak{S}_n(i-1) \cap H(t)$. Now, let p_v denote the probability that some node $v \in \mathfrak{S}_n(i-1) \cap H(t)$ becomes informed in step $t+1$. Since $p_v = d_{I(t)}(v)/d(n)$, where $d_{I(t)}(v)$ denotes the number of neighbors of v in $I(t)$, it holds that

$$\mathbf{E}[|(I(t+1) \setminus I(t)) \cap \mathfrak{S}_n(i-1)|] = \sum_{v \in \mathfrak{S}_n(i-1) \cap H(t)} \frac{d_{I(t)}(v)}{d(n)},$$

which equals $\Omega(|I(t) \cap \mathfrak{S}_n(i-1)|)$. This implies that

$$|I(t+1) \cap \mathfrak{S}_n(i-1)| \geq (1+\rho)|I(t) \cap \mathfrak{S}_n(i-1)|,$$

where $\rho = \Theta(1)$, w.c.p.

Now we assume that at some proper time t' it holds that $|I(t') \cap \mathfrak{S}_n(i-1)| \geq \delta d(n)^{c'}$, where δ is a constant. Due to assumption (3), we know that for all $v \in \mathfrak{S}_n(i-1)$ the distance to $\mathfrak{S}_n(i)$ is at most c' .

Let us now consider the propagation of the information in $\mathfrak{S}_n(i-1)$ towards $\mathfrak{S}_n(i)$. Recall, that from each node $v \in I(t') \cap \mathfrak{S}_n(i-1)$ exists a path to some node in $\mathfrak{S}_n(i)$ of length at most c' .

Now define $\mathcal{L}_1 := I(t') \cap \mathfrak{S}_n(i-1)$, $\mathcal{L}_2 := \{w \in \mathfrak{S}_n(i-1) \mid \text{dist}(w, \mathfrak{S}_n(i)) = c' - 1\}$, \dots , $\mathcal{L}_{c'+1} := \mathfrak{S}_n(i)$. Assume w.l.o.g. that for each node $v \in I(t') \cap \mathfrak{S}_n(i-1)$ it holds $v \in \mathcal{L}_1$. Observe that $|\mathcal{L}_2| \geq |\mathcal{L}_1|/d(n)$, and generally $|\mathcal{L}_j| \geq \max\{1, |\mathcal{L}_1|/(d(n))^{j-1}\}$ for any j . Since any node of \mathcal{L}_j has a neighbor in \mathcal{L}_{j+1} , and a node v of \mathcal{L}_{j+1} becomes informed in some step $t''+1$ with probability $d_{I(t'')}(v)/d(n)$, it holds that

$$\mathbf{E}[|\mathcal{L}_{j+1} \cap I(t''+1)|] = \sum_{v \in \mathcal{L}_{j+1}} \frac{d_{I(t'')}(v)}{d(n)} \geq \frac{|\mathcal{L}_j \cap I(t'')|}{d(n)}$$

which implies $|\mathcal{L}_{j+1} \cap I(t'+O(1))| \geq \delta d(n)^{c'-j}$, w.c.p., provided that $|\mathcal{L}_j \cap I(t'+O(1))| \geq \delta d(n)^{c'-j+1}$. Summarizing, the time needed to complete \mathcal{P}_i can be modelled by a sum of $O(\log d(n)^{c'} + c') = O(\log d(n))$ independent geometrically distributed random variables with constant mean. Recall, that we have $O(n)$ phases. Thus, applying the Chernoff-Bound [6, 17] we conclude that some fixed vertex w is αn -approximated within $t_1 := O(n \log d(n)^{c'})$ steps with probability $1 - O(1/N^2)$. Using the Markov inequality we conclude that each vertex of G is αn -approximated at time t_1 , w.h.p.

Using the techniques of [14] we obtain

$$|I(t_1)| \geq \frac{n!}{d(n)^{\alpha n+1}} \geq \frac{n!}{n^{(\alpha n+1)c}} \geq n^{n-2\alpha cn}.$$

Furthermore, to obtain a subset of informed nodes $\mathcal{A} \subseteq I(t_1)$ which only contains vertices being at distance at least αn from each other we get by the same arguments that $|\mathcal{A}| \geq |I(t_1)|/(d(n)^{\alpha n+1}) \geq n^{n-4\alpha nc}$.

Using similar arguments as above, it can be shown that for any pair of vertices $v, w \in V$ and any time t_2 , there is a vertex $w' \in N_{\alpha n/2}(w)$ which contacts v within time interval $[t_2, t_2 + O(n \log d(n))]$, w.h.p.

In order to finish the proof we use similar techniques as in [12] which are omitted here due to space limitations. \square

It is now not too difficult to see that the class given in the previous theorem includes the following three well-known representatives of Cayley graphs.

Remark 1 *The Star graph, Pancake graph and Transposition graph [2, 1] satisfy the conditions of the Theorem 3.*

5 A New Martingale-Based Technique

Definition 3. [2] *The Bubble sort graph is defined as $B(n) = (\mathfrak{S}_n, S_n)$, where $S_n = \{(i, i+1) \mid i \in \{1, \dots, n-1\}\}$.*

Since the diameter of a Bubble sort graph is obviously $\Omega(n^2)$, Theorem 3 is not applicable and new techniques have to be developed. First, we briefly summarize the research history of related random processes on these graphs.

In spite of very refined techniques designed for the analysis of shuffling cards procedures, the mixing time of the Bubble sort graph has been an open question for almost two decades. Finally in 2002, Wilson proved the mixing time $\Theta(n^3 \log n)$ which is asymptotically tight up to a small constant factor [23].

Additionally, Diaconis and Ram considered the following generalization. First, fix some parameter $p \in (0, 1)$. In each step, choose uniformly at random one pair of adjacent cards and flip a coin that is heads with probability p . If the coin comes up heads, then arrange the cards in the correct order. Otherwise, arrange them in the reverse order.

For $1/2 < p \leq 1$ this shuffling card procedure models a randomized version of Bubble sort. In particular, the stationary distribution of this Markov chain is no longer uniform. Rather surprisingly, Benjami et.al. [4] proved very recently that the mixing time decreases to $O(n^2)$ if $p \neq \frac{1}{2}$ and thereby affirmed a conjecture of Diaconis and Ram. To follow the notation of Benjami et.al., denote by $\mathcal{DA}(n, p)$ the aforementioned card shuffling procedure. Then their result can be formally stated as follows. For any $p > 1/2$ it holds $\text{MIX}_{\mathcal{DA}(n,p)}(e^{-1}) \leq O(n^2)$. On the other hand, there is no cutoff [10] known yet. Thus, it is an open question of what magnitude is $\text{MIX}_{\mathcal{DA}(n,p)}(1 - o(1))$. However, by transferring the result of Benjami et.al. to RBA_s and using refined martingale techniques, we will prove a tight concentration of the distribution of the runtime $\text{RT}_s(G)$ around its mean.

Since RBA_s can simulate $\mathcal{DA}(n, p)$ we obtain the following result.

Lemma 1. *RBA_s informs some fixed node v within $O(n^2)$ rounds w.c.p.*

Our objective is to extend the Lemma above such that all nodes become informed after $O(n^2)$ rounds w.h.p.

To simplify the notation, we will analyze a slightly modified version of RBA_s , denoted by RBA_s' . Here, in each time step $t = 1, 2, \dots \in \mathbb{N}$ one node is chosen uniformly at random and sends the information to some randomly chosen neighbor provided that it is already informed. Obviously, this is just a scaling of the time axis by a factor of N compared to RBA_s .

In the following, we fix some node $v \in V$. We make use of the following doob martingale [3] (sometimes also called exposure martingale). Let $Z_0 := \mathbf{E}[\text{RT}_s'(v)]$, where $\text{RT}_s'(v)$ is the random variable representing the runtime required to inform v . Furthermore define $Z_t := \mathbf{E}[\text{RT}_s'(v) \mid I(0), \dots, I(t)] = \mathbf{E}[\text{RT}_s'(v) \mid I(t)]$. Thus, Z_t estimates the runtime conditioned on the set of informed nodes at time step t . Note that Z_t is a (random) function depending on $I(t)$. Moreover, if $Z_t \leq t$, then v has been informed and the sequence Z_t, Z_{t+1}, \dots becomes stationary. Additionally, for any two subsets $A \subseteq B \subseteq V$ we have

$$\mathbf{E}[\text{RT}_s'(v) \mid I(t) = A] \geq \mathbf{E}[\text{RT}_s'(v) \mid I(t) = B], \quad (1)$$

$$\mathbf{E}[\text{RT}_s'(v) \mid I(t-1) = A] + 1 = \mathbf{E}[\text{RT}_s'(v) \mid I(t) = A]. \quad (2)$$

Another building block will be the following concentration inequality.

Theorem 4. [8] *Let $Z_0 \dots, Z_t$ be a martingale w.r.t. the sequence $I(0), \dots, I(t)$ such that for $1 \leq k \leq t$ it holds $|Z_k - Z_{k-1}| \leq M$, $\mathbf{Var}[Z_k \mid I(0), \dots, I(k-1)] \leq \sigma_k^2$. Then for all $t \geq 0$ and $\lambda > 0$, $\Pr[|Z_t - Z_0| \geq \lambda] \leq 2e^{-\lambda^2 / (\sum_{k=1}^t \sigma_k^2 + M\lambda/3)}$.*

Let $\text{RT}_s'(u, v) := \min\{t \in \mathbb{N} \cup \{0\} \mid u \in I(t)\}$ conditioned on $I(0) = \{v\}$. and $\beta(G) := \max_{(u,v) \in E(G)} \mathbf{E}[\text{RT}_s'(u, v)]$. The following lemma improves the trivial bound $\beta(G) \leq \Delta(G) \cdot N$ for several graphs.

Lemma 2. *Let G be any d -regular graph. If for any two adjacent nodes $u, v \in V$ there exist $\Theta(d)$ node-disjoint paths of length at most 3, then $\beta(G) \leq O(d^{2/3}N)$.*

Note that the Transposition graph, Bubble sort graph and Hypercube satisfy the condition of this lemma. The following theorem relates the distribution of $Z_k - Z_{k-1}$ conditioned on Z_{k-1} to the combinatorial value $\beta(G)$.

Theorem 5. *For any graph $G = (V, E)$ we have for all $k \in \mathbb{N} \setminus \{0\}$*

$$-\beta(G) \leq Z_k - Z_{k-1} \leq 1 \quad \text{and} \quad \mathbf{Var}[Z_k \mid I(k-1)] \leq \beta(G).$$

Proof. Assume that $I(k-1) = I$ for a fixed I . We consider now two cases. In case of $I(k) = I$ we get $Z_k = Z_{k-1} + 1$ by (2). Secondly, if $I(k) = I \cup \{v\}$ for some $v \in N(u) \cap I^c$, $u \in I$, then

$$\begin{aligned} \mathbf{E}[\text{RT}_s'(v) \mid I(k-1) = I] &\stackrel{(1)}{\leq} \mathbf{E}\left[\min_{j \in \mathbb{N} \cup \{0\}} \{v \in I(k-1+j)\} \mid I(k-1) = \{u\}\right] \\ &\quad + \mathbf{E}[\text{RT}_s'(v) \mid I(k) = I \cup \{v\}] \\ &\stackrel{(2)}{=} \mathbf{E}\left[\min_{j \in \mathbb{N} \cup \{0\}} \{v \in I(j)\} \mid I(0) = \{u\}\right] \\ &\quad + \mathbf{E}[\text{RT}_s'(v) \mid I(k) = I \cup \{v\}] \end{aligned}$$

and thus $\mathbf{E} [\text{RT}_s'(v) \mid I(k) = I \cup \{v\}] - \mathbf{E} [\text{RT}_s'(v) \mid I(k-1) = I] \geq -\beta(G)$. By the first inequality we know that Z_k is a random variable whose values are all in the interval $[Z_{k-1} - \beta(G), Z_{k-1} + 1]$. Moreover, by the martingale property we have $\mathbf{E} [Z_k \mid I(k-1)] = Z_{k-1}$. Thus, by some standard upper bound on the variance [16] we finally obtain $\mathbf{Var} [Z_k] \leq |-\beta(G) \cdot 1| = \beta(G)$. \square

Note that in all previous results $\text{RT}_s'(v)$ can be replaced by $\text{RT}_s'(G)$.

Theorem 6. *It holds that $\text{RT}_d(B(n), 1 - \frac{1}{N}) = \Theta(n^2)$. Moreover, for any $x < 2$ we have that $\Pr [\text{RT}_s(B(n)) \leq \mathbf{E} [\text{RT}_s(B(n))] + n^x] \leq O(\exp(-n^{2x-8/3}))$.*

Proof. Fix some arbitrary node $w \in V(B(n))$. Due to Lemma 2 we have that $\beta(B(n)) = O(n^{2/3}N)$. Consequently by Theorem 5 it holds $\mathbf{Var} [Z_i - Z_{i-1/N}] = O(n^{2/3}N)$ and $|Z_i - Z_{i-1/N}| = O(n^{2/3}N)$. Then we apply Theorem 4 with $t = \mathbf{E} [\text{RT}_s'(v)]$, $\sigma_i^2 \leq O(n^{2/3}N)$, $\lambda := \mathbf{E} [\text{RT}_s'(v)] := \gamma n^2 N$, where $\gamma(n) = O(1)$ is some bounded function and obtain

$$\begin{aligned} \Pr [|Z_{2\mathbf{E}[\text{RT}_s'(v)]} - \mathbf{E} [\text{RT}_s'(v)] | \geq \lambda] &\leq 2e^{-\lambda^2 / (\sum_{k=1}^t \sigma_k^2 + M\lambda/3)}, \\ \Pr [|Z_{2\lambda} - \lambda| \geq \lambda] &\leq 2e^{\frac{-\gamma^2 N^2 n^4}{2N\gamma n^2(n^{2/3}N) + \gamma \cdot n^2 N \cdot n^{2/3} \cdot N}}, \\ \Pr [Z_{2\lambda} \leq 2\lambda] &\leq O(e^{\frac{-n^{12/3}}{n^{8/3}}}) \leq O(e^{-n^{4/3}}) \leq 1 - \frac{1}{N^2}. \end{aligned}$$

Thus after 2λ time steps, each single node of $B(n)$ has received the information with probability $1 - (1/N)^2$. Hence by Markov's inequality $\text{RT}_s'(B(n), 1 - \frac{1}{N}) = \Theta(Nn^2)$. The second claim is shown similarly. \square

It is worth mentioning that with the same techniques similar, but weaker tail estimates can be proven for Hypercubes, Star graphs and Pancake graphs.

6 Conclusions

In this paper we developed a new relationship between broadcasting and random walks, and proved that randomized broadcasting has optimal runtime on several classes of Cayley graphs. However, it would be still interesting whether the additional logarithmic factor in Theorem 2 can be reduced. It is also an open question on which graphs fast broadcasting implies fast mixing, though this has to be a more restricted class. Although the techniques introduced in Section 4 seem to be powerful, we could not apply it to all Cayley graphs considered in this paper. Our hope is that incorporating edge-expansion-based approaches would extend the applicability of this method.

7 Acknowledgments

We thank Peter B"urgisser for helpful suggestions concerning Section 5.

References

1. S. Akers, D. Harel, and B. Krishnamurthy. The star graph: An attractive alternative to the n -cube. In *Proc. of ICPP'87*, pages 393–400, 1987.
2. S. Akers and B. Krishnamurthy. A group-theoretic model for symmetric interconnection networks. In *Proc. of ICPP'86*, pages 555–565, 1986.
3. N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley-Interscience Series in Discrete Mathematics and Optimization, 2000.
4. I. Benjamini, N. Berger, C. Hoffmann, and E. Mossel. Mixing times of the biased card shuffling and the asymmetric exclusion process. *Transactions of the American Mathematical Society*, 357:3013–3029, 2005.
5. N. Biggs. *Algebraic Graph Theory*. Cambridge University Press, 1993.
6. H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Stat.*, 23:493–507, 1952.
7. F. Chung. *Spectral Graph Theory*, volume 92 of *CBMS Regional conference series in mathematics*. American Mathematical Society, 1997.
8. F. Chung and L. Lu. Concentration inequalities and martingale inequalities — a survey. *Internet Mathematics (to appear)*.
9. 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.
10. P. Diaconis. *Group Representations in Probability and Statistics*, volume 11. Lecture notes-Monograph Series, 1988.
11. R. Diekmann, A. Frommer, and B. Monien. Efficient schemes for nearest neighbor load balancing. *Parallel Computing*, 25(7):789–812, 1999.
12. R. Elsässer and T. Sauerwald. On randomized broadcasting in star graphs. In *Proc. of WG'05*, pages 307–318, 2005.
13. R. Elsässer and T. Sauerwald. On the runtime and robustness of randomized broadcasting. In *Proc. of ISAAC' 06*, pages 349–358, 2006.
14. U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithm*, I(4):447–460, 1990.
15. L. Gasieniec and A. Pelc. Adaptive broadcasting with faulty nodes. *Parallel Computing*, 22:903–912, 1996.
16. M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, and B. Reed. *Probabilistic Methods for Algorithmic Discrete Mathematics*. Algorithms and Combinatorics, 1991.
17. T. Hagerup and C. Rüb. A guided tour of chernoff bounds. *Information Processing Letters*, 36(6):305–308, 1990.
18. J. Hromkovič, R. Klasing, A. Pelc, P. Ruzicka, and W. Unger. *Dissemination of Information in Communication Networks*. Springer, 2005.
19. F. Leighton, B. Maggs, and R. Sitamaran. On the fault tolerance of some popular bounded-degree networks. In *Proc. of FOCS'92*, pages 542–552, 1992.
20. A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.
21. B. Pittel. On spreading rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.
22. A. Sinclair and M. Jerrum. Approximate counting, uniform generation, and rapidly mixing markov chains. *Inform. and Comput.*, 82:93–113, 1989.
23. D. Wilson. Mixing times of lozenge tiling and card shuffling markov chains. *Annals of Applied Probability*, 14:274–325, 2004.