

COVER TIME AND BROADCAST TIME

ROBERT ELSÄSSER¹ AND THOMAS SAUERWALD²

¹ Institute for Computer Science, University of Paderborn, 33102 Paderborn, Germany
E-mail address: `elsa@upb.de`

² International Computer Science Institute, 1947 Center Street, Berkeley, CA 94704, U.S.
E-mail address: `sauerwal@icsi.berkeley.edu`

ABSTRACT. We introduce a new technique for bounding the cover time of random walks by relating it to the runtime of randomized broadcast. In particular, we strongly confirm for dense graphs the intuition of Chandra et al. [8] that “the cover time of the graph is an appropriate metric for the performance of certain kinds of randomized broadcast algorithms”. In more detail, our results are as follows:

- For any graph $G = (V, E)$ of size n and minimum degree δ , we have $\mathcal{R}(G) = \mathcal{O}(\frac{|E|}{\delta} \cdot \log n)$, where $\mathcal{R}(G)$ denotes the quotient of the cover time and broadcast time. This bound is tight for binary trees and tight up to logarithmic factors for many graphs including hypercubes, expanders and lollipop graphs.
- For any δ -regular (or almost δ -regular) graph G it holds that $\mathcal{R}(G) = \Omega(\frac{\delta^2}{n} \cdot \frac{1}{\log n})$. Together with our upper bound on $\mathcal{R}(G)$, this lower bound strongly confirms the intuition of Chandra et al. for graphs with minimum degree $\Theta(n)$, since then the cover time equals the broadcast time multiplied by n (neglecting logarithmic factors).
- Conversely, for any δ we construct almost δ -regular graphs that satisfy $\mathcal{R}(G) = \mathcal{O}(\max\{\sqrt{n}, \delta\} \cdot \log^2 n)$. Since any regular expander satisfies $\mathcal{R}(G) = \Theta(n)$, the strong relationship given above does not hold if δ is polynomially smaller than n .

Our bounds also demonstrate that the relationship between cover time and broadcast time is much stronger than the known relationships between any of them and the mixing time (or the closely related spectral gap).

1. Introduction

Motivation. A *random walk* on a graph is the following process. Starting from a specified vertex, the walk proceeds at each step from its current position to an adjacent vertex chosen uniformly at random. The study of random walks has numerous applications in the design and analysis of algorithms (cf. [24] for a survey). Two of the most important

1998 ACM Subject Classification: G.3 Probability and Statistics [Probabilistic Algorithms, Stochastic Processes].

Key words and phrases: Random walk, randomized algorithms, parallel and distributed algorithms.

This work has been partially supported by the IST Program of the European Union under contract number 15964 (AEOLUS), by the German Science Foundation (DFG) Research Training Group GK-693 of the Paderborn Institute for Scientific Computation (PaSCo) and by the German Academic Exchange Service (DAAD).

parameters of random walks are its *mixing time* which is the time until the walk becomes close to the stationary distribution, and its *cover time* which is the expected time required for the random walk to visit all vertices.

Famous combinatorial problems solved by rapidly mixing random walks are, e.g., approximating the permanent and approximating the volume of convex bodies (cf. [24] for more details). The cover time comes naturally into play when the task is to explore a network, or to estimate the stationary distribution of a graph [31]. Moreover, the cover time is intimately related to combinatorial and algebraic properties such as the conductance and the spectral gap of the underlying graph [6] and thus, bounding the cover time may also lead to interesting combinatorial results.

In this paper, we are particularly interested in the relationship between the cover time of random walks and the runtime of *randomized broadcast* [16]. Broadcasting in large networks has various fields of application in distributed computing such as the maintenance of replicated databases or the spreading of information in networks [16, 21]. Furthermore it is closely related to certain mathematical models of epidemic diseases where infections are spread to some neighbours chosen uniformly at random with some probability. However, in most papers, spreaders are only active in a given time frame, and the question of interest is, whether on certain networks an epidemic outbreak occurs [22, 27]. Several threshold theorems involving the basic reproduction number, contact number, and the replacement number have been stated (see [19] for a collection of results).

Here, we consider the so-called *randomized broadcast algorithm* [16] (also known as *push algorithm*): at the beginning, a vertex s in a graph G knows of some rumor which has to be disseminated to all other vertices. Then, at each time-step every vertex that knows of the rumor chooses one of its neighbors uniformly at random and informs it of the rumor. The advantage of *randomized broadcast* is in its inherent robustness against several kinds of failures (e.g., [16]) and dynamical changes compared to deterministic schemes that either need substantially more time or can tolerate only a relatively small number of faults [21].

Related Work. There is a vast body of literature devoted to the cover time of random walks and we can only point to some results directly related to this paper. Aleliunas et al. [3] initiated the study of the cover time. Amongst other results, they proved that the cover time of any graph $G = (V, E)$ with n vertices is at most $\mathcal{O}(n \cdot |E|)$. To obtain this result they proved that the cover time is bounded by the weight of a spanning tree whose edges are weighted according to the commute times between the corresponding vertices. This approach was later refined by Feige [15] to obtain an upper bound of less than $2n^2$ for regular graphs. While the spanning tree technique is particularly useful for graphs that have a high cover time [15], it vastly overestimates the cover time of e.g., complete graphs.

The seminal work of Chandra et al. [8] established a close connection between the *electrical resistance* of a graph and its cover time. This correspondence allows the application of elegant methods from electrical network theory, e.g., the use of short-cut-principles or certain flow-based arguments. Nevertheless, for the computation of the resistance of a given graph other graph-theoretical parameters are often required, e.g., vertex-expansion, number of vertex-disjoint paths or the number of vertices within a certain distance [8].

A wide range of techniques to upper bound the cover time is based on the mixing time of a random walk or the closely related spectral gap. The technique of reducing the cover time to the coupon collector's problem on graphs with low mixing time traces back to Aldous [1] who derived tight bounds on the cover time of certain Cayley graphs. Later, Cooper and Frieze extended this technique to bound the cover time of several classes of random graphs,

e.g., [10]. The basic idea of this method is that after each mixing time steps, the random walk visits an (almost) randomly chosen vertex. The crux is to deal with the dependencies among the intermediate vertices. Hence, in addition to an upper bound on the mixing time of logarithmic [10] or at least sub-polynomial order [1], one has to bound the number of returns to the starting vertex within mixing time steps.

A related result was derived by Broder and Karlin [6] who bounded the cover time in terms of the *spectral gap* $1 - \lambda_2$, where λ_2 is the second largest eigenvalue of the transition matrix of the random walk.

Winkler and Zuckerman [31] introduced an interesting parameter called *blanket time* which is closely related to the cover time. Here, one asks for the first time-step at which the observed distribution of the visited vertices approximates the stationary distribution up to a constant factor. Winkler and Zuckerman conjectured that the blanket time is asymptotically the same as the cover time. In [20] Kahn et al. showed that the blanket time is upper bounded by the cover time multiplied by $\mathcal{O}((\ln \ln n)^2)$ for any graph.

Most papers dealing with randomized broadcast analyze the runtime on different graph classes. Pittel [28] proved that the runtime on complete graphs is $\log_2 n + \ln n \pm \mathcal{O}(1)$. Feige et al. [16] derived several upper bounds, in particular a bound of $\mathcal{O}(\log n)$ for hypercubes and random graphs. We extended the bound of $\mathcal{O}(\log n)$ to a certain class of Cayley graphs in [12]. Additionally, we proved that the broadcast time is upper bounded by the sum of the mixing time and an additional logarithmic factor [29] (a similar result for a related broadcast algorithm was derived by Boyd et al. [5]). However, the mixing time cannot be used for an appropriate lower bound on the broadcast time, as it may overestimate the broadcast time up to a factor of n on certain graphs (cf. Section 3.2).

Our Results. We present the first formal results relating the cover time to the broadcast time. In most of them, we will assume that the broadcast and the random walk both start from its respective worst-case initial vertex. Note that at a first look these processes seem not to be too closely related, since randomized broadcast is a *parallel* process where propagation occurs at *every* informed vertex simultaneously, while a random walk moves "only" from *one* vertex to another [16]. Nevertheless, Chandra et al. [8] mentioned that "The cover time of the graph is an appropriate metric for the performance of certain kinds of randomized broadcast algorithms". As a consequence of our main results, we obtain a fairly tight characterization of graph classes for which the cover time and broadcast time capture each other. On the positive side, for every graph with minimum degree $\Theta(n)$, the cover time equals the broadcast time multiplied by n , up to logarithmic factors (this kind of tightness (up to logarithmic factors) has been frequently considered in the study of random walks, e.g., when studying rapidly mixing Markov chains [30], or when bounding the cover time [8],[24, Theorem 2.7]). On the negative




Figure 1: All bounds on $\mathcal{R}(G)$ at a glance. \mathcal{G}_δ denotes the class of graphs with $\Delta = \mathcal{O}(\delta)$. The blue and red polygons indicate the gap between our lower and upper bounds on $\max_{G \in \mathcal{G}_\delta} \mathcal{R}(G)$ and $\min_{G \in \mathcal{G}_\delta} \mathcal{R}(G)$, resp.

side, when studying rapidly mixing Markov chains [30], or when bounding the cover time [8],[24, Theorem 2.7]). On the negative

side, this strong correspondence does not hold on almost regular graphs, when the degree is substantially smaller than n .

In more detail, our results illustrated in Figure 1 are as follows. First, we prove that the cover time of any graph with minimum degree δ is at most $\mathcal{O}(\frac{|E|}{\delta} \log n)$ multiplied by the (expected) broadcast time, that is, the quotient $\mathcal{R}(G)$ of the cover time and broadcast time is $\mathcal{O}(\frac{|E|}{\delta} \log n)$. This bound is tight up to a constant factor for binary trees and tight up to a logarithmic factor for various graphs including, e.g., expanders, hypercubes and lollipop-graphs. As an application, we use this result to upper bound the cover time of generalized random graphs that are used as a model for real world networks [9].

Conversely, we consider the question of lower bounding $\mathcal{R}(G)$. By showing that the commute time between two vertices u, v is at least $2 \cdot \text{dist}(u, v)^2$, we obtain that $\mathcal{R}(G) = \Omega(\frac{\sqrt{n \log n}}{\Delta})$ for any graph with maximum degree Δ . For constant Δ , this bound is tight for the two-dimensional $\sqrt{n} \times \sqrt{n}$ -torus up to logarithmic factors. We move on to improve this bound for denser graphs with $\Delta = \mathcal{O}(\delta)$ to $\mathcal{R}(G) = \Omega(\frac{\sqrt{n}}{\sqrt{\delta \log n}})$. More importantly, for any graph with $\Delta = \mathcal{O}(\delta)$ we establish that $\mathcal{R}(G) = \Omega(\frac{\delta^2}{n} \cdot \frac{1}{\log n})$. Together with our upper bound on $\mathcal{R}(G)$, this implies that on any graph with $\delta = \Theta(n)$, cover time and broadcast time (multiplied by n) capture each other up to logarithmic factors.

We complement these positive results by the construction of (almost) d -regular graphs for which $\mathcal{R}(G) = \mathcal{O}(\max\{\sqrt{n}, d\} \cdot \log n)$. Since for any d -regular expander (graphs for which the spectral gap satisfies $(1 - \lambda_2)^{-1} = \mathcal{O}(1)$), $\mathcal{R}(G) = \Theta(n)$, the cover time does *not* always capture the performance of randomized broadcast for the class of almost d -regular graphs when d is polynomially smaller than n .

All of our lower and upper bounds reveal a surprisingly close relationship between the cover time and broadcast time. In particular, upper bounding the cover time in terms of the broadcast time turns out to be as good as (and in some cases much better than) bounding it in terms of the spectral-gap (cf. Section 3.2). From another perspective, we derive a lower bound on the broadcast time in terms of the cover time that nicely complements the existing upper bounds on the broadcast time based on the mixing time [12, 29]. A further novel feature of this work is the use of techniques from electrical network theory to bound the broadcast time. We should note that certain difficulties in applying such methods for the study of randomized broadcast have been mentioned by Feige et al. [16].

2. Notations, Definitions and Preliminaries

Throughout this paper, let $G = (V, E)$ be an undirected, simple and connected graph of size $n = |V|$. By δ and Δ we denote the minimum and maximum degree of G , respectively. For some set $X \subseteq V$, $N(X)$ denotes the set of all neighbors of $x \in X$, and $\deg_X(u)$ is the number of edges between u and the vertices of X .

Random Walk. A *random walk* [24] on a graph G starts at a specified vertex $s \in V$ and moves in each step to a neighboring vertex chosen uniformly at random. This can be described by a *transition matrix* \mathbf{P} , where $p_{ij} = 1/\deg(i)$ if $\{i, j\} \in E(G)$, and $p_{ij} = 0$ otherwise. Then, the random walk is an infinite sequence of vertices X_0, X_1, \dots , where $X_0 := s$ is the starting point of this random walk, and X_t denotes the vertex visited by the random walk at step t . Note that X_t is a random variable with a distribution $\mathbf{p}_s(t)$ on $V(G)$. Denoting by $\mathbf{p}_s(0)$ the unit-vector (regarded as column vector) with 1 at the component corresponding to s and 0 otherwise, we obtain the iteration $\mathbf{p}_s(t+1) = \mathbf{p}_s(t) \cdot \mathbf{P}$

for every step $t \in \mathbb{N}$. It is well-known that on non-bipartite graphs, $\mathbf{p}_s(t)$ converges for $t \rightarrow \infty$ towards the *stationary distribution* vector π given by $\pi(v) = \deg(v)/(2|E|)$. For simplicity, we confine ourselves to non-bipartite graphs in the following. This causes no loss of generality as for general graphs (including bipartite ones) convergence can be ensured easily by using the transition matrix $\frac{1}{2}\mathbf{I} + \frac{1}{2}\mathbf{P}$ (with \mathbf{I} being the identity matrix) instead of \mathbf{P} . This change of the transition matrix slows down the mixing time (and the cover time) only by some constant factor [24, 30].

Mixing Time and Spectral Gap. The *mixing time* of a random walk on G is $\text{MIX}_\varepsilon(G) := \max_{s \in V} \min\{t \in \mathbb{N} : \|\mathbf{p}_s(t) - \pi\|_1 \leq \varepsilon, X_0 = s\}$. Since G is connected and non-bipartite, the eigenvalues of \mathbf{P} satisfy $\lambda_1 = 1 > \lambda_2 \geq \dots \geq \lambda_n > -1$. The following result by Sinclair shows that the spectral gap $1 - \lambda_2$ captures the mixing time up to logarithmic factors.

Theorem 2.1 ([30]). *For any graph $G = (V, E)$ and $\varepsilon > 0$,*

$$\Omega\left(\frac{\lambda_2}{1 - \lambda_2} \cdot \log\left(\frac{1}{\varepsilon}\right)\right) = \text{MIX}_\varepsilon(G) = \mathcal{O}\left(\frac{1}{1 - \lambda_2} \cdot \left(\log n + \log\left(\frac{1}{\varepsilon}\right)\right)\right).$$

Commute Time, Resistance and Cover Time. For two vertices $u, v \in V(G)$, the *hitting time* from u to v is defined as $\text{H}(u, v) := \mathbf{E}[\min\{t \in \mathbb{N} \setminus \{0\} : X_t = v, X_0 = u\}]$, i. e., the expected number of steps to reach v from u . The *commute time* $\text{C}(u, v)$ is defined as the expected number of steps to reach v when starting from u and then returning back to u , so, $\text{C}(u, v) := \text{H}(u, v) + \text{H}(v, u)$. Consider now the graph G as an electrical network where each edge represents a unit resistance. Let u and v be two vertices. Assume that one ampere were injected into vertex u and removed from vertex v . Then $\text{R}(u, v)$ is the voltage difference between u and v (for more details on electrical networks we refer the reader to [8, 24]), and is related to $\text{C}(u, v)$ as follows.

Theorem 2.2 ([8]). *For any pair of vertices $u, v \in V$, $\text{C}(u, v) = 2|E| \cdot \text{R}(u, v)$.*

We will mainly be concerned with the *cover time*, which is the expected number of steps a random walk takes to visit all vertices of G . Denote by $\text{COV}(s)$ this time for a random walk which starts from s , and let $\text{COV}(G) := \max_{s \in V} \text{COV}(s)$. The cover time is related to the maximum commute time by means of $\frac{1}{2} \cdot \max_{u, v \in V} \text{C}(u, v) \leq \text{COV}(G) \leq e^3 \cdot \max_{u, v \in V} \text{C}(u, v) \ln n + n$ [8]. We restate the following bounds by Feige.

Theorem 2.3 ([13, 14]). *For any graph, $(1 - o(1)) \cdot n \ln n \leq \text{COV}(G) \leq (\frac{4}{27} + o(1)) \cdot n^3$.*

A corresponding result to Theorem 2.1 for $\text{COV}(G)$ was given by Broder and Karlin.

Theorem 2.4 ([6]). *For any regular graph $G = (V, E)$, $\text{COV}(G) = \mathcal{O}(\frac{1}{1 - \lambda_2} \cdot n \log n)$.*

Randomized Broadcast. We will consider the relationship between the cover time of random walks and the following *randomized broadcast algorithm* RBA (also known as push algorithm). Assume that at time $t = 0$ a vertex s knows of a rumor which has to be spread to all other vertices. Then, at each time-step $t = 1, 2, \dots$ every vertex that knows of the rumor chooses a neighbor uniformly at random and informs it of the rumor. Let I_t be the set of informed vertices at time t , so $I_0 = \{s\}$. The runtime of RBA is denoted by $\text{RBA}_p(G) := \max_{s \in V} \min\{t \in \mathbb{N} : \Pr[I_t = V \mid I_0 = \{s\}] \geq 1 - p\}$ for some given $0 < p < 1$. The expected runtime is $\mathbf{E}[\text{RBA}(G)] := \max_{s \in V} \{\mathbf{E}[\min\{t \in \mathbb{N} : I_t = V, I_0 = \{s\}\}]\}$. By standard arguments, we have $\mathbf{E}[\text{RBA}(G)] = \mathcal{O}(\text{RBA}_{n-1}(G)) = \mathcal{O}(\mathbf{E}[\text{RBA}(G)] \cdot \log n)$. We remark that $\text{RBA}(G)$ is at least $\max\{\log_2 n, \text{diam}(G)\}$ on any graph G , and $\text{RBA}_{n-1}(G)$ may range from $\Theta(\log n)$ (which is the case for many "nice" graphs) to $\Theta(n \log n)$ (which is the

case for the star) [16]. Sometimes we also use $\text{RBA}(s, v) := \min\{t \in \mathbb{N} : v \in I(t) \mid I(0) = \{s\}\}$ and $\text{RBA}_p(s, v) := \min\{t \in \mathbb{N} : \Pr[v \in I(t) \mid I_0 = \{s\}] \geq 1 - p\}$ for some specified $0 < p < 1$. We will frequently make use of following upper bound of Feige et al. [16].

Theorem 2.5 ([16]). *For any graph $G = (V, E)$, $\text{RBA}_{n-1}(G) = \mathcal{O}(\Delta \cdot (\log n + \text{diam}(G)))$.*

To compare the cover time with the broadcast time, we define $\mathcal{R}(G) := \frac{\text{COV}(G)}{\mathbf{E}[\text{RBA}(G)]}$.

3. Upper Bound on $\mathcal{R}(G)$ and Applications

3.1. Upper Bound on $\mathcal{R}(G)$

To prove an upper bound on $\mathcal{R}(G)$, we first prove a general inequality between *first-passage-percolation* times and broadcast times and apply then a result of Lyons et al. [25] relating first-passage-percolation to the cover time.

Definition 3.1 ([17, 25]). The *undirected first-passage-percolation* UFPP is defined as follows. All (undirected) edges $e \in E(G)$ are assigned weights $w(e)$ that are independent exponential random variable with parameter 1. Specify a vertex s . Then the *first-passage-percolation time* from s to v is defined by $\text{UFPP}(s, v) := \inf_{\mathcal{P}=(s, \dots, v)} \sum_{e \in \mathcal{P}} w(e)$, where the inf is over all possible paths from s to v in G . Note that $\text{UFPP}(s, s) = 0$.

Theorem 3.2. *For any graph $G = (V, E)$ and $s, v \in V$, $\mathbf{E}[\text{UFPP}(s, v)] \leq \frac{2}{\delta} \cdot \mathbf{E}[\text{RBA}(s, v)]$.*

Proof. In the proof we derive several (in-)equalities between different percolation and broadcast models. First we introduce a directed version of UFPP, denoted by DFPP. In this model each undirected edge $\{u, u'\} \in E(G)$ is replaced by two directed edges (u, u') and (u', u) , and all directed edges e are assigned weights $w(e)$ that are independent exponential random variable with parameter 1. Denote by $\text{DFPP}(s, v)$ the corresponding first-passage-percolation time of this directed version.

Lemma 3.3. *For any graph $G = (V, E)$ and $s, v \in V$, $\mathbf{E}[\text{UFPP}(s, v)] \leq 2 \cdot \mathbf{E}[\text{DFPP}(s, v)]$.*

Next consider another broadcast model denoted by $\overline{\text{SEQ}}$. At the beginning, a vertex s knows of a rumor which has to be spread to all other vertices. Once a vertex u receives the rumor at time $t \in \mathbb{R}$, it sends the rumor at each time $t + X_{1,u}, t + X_{1,u} + X_{2,u}, \dots$ to a randomly chosen neighbor, where the $X_{i,u}$ with $i \in \mathbb{N}$ are independent exponential variables with parameter $\text{deg}(u)$. Let $\overline{\text{SEQ}}(s, u)$ be the first time when u is informed.

Lemma 3.4. *For any $s, v \in V$, $\overline{\text{SEQ}}(s, v)$ and $\text{DFPP}(s, v)$ have the same distribution.*

Finally, our aim is to relate $\overline{\text{SEQ}}$ and RBA.

Observation 3.5. In any execution of RBA, there is for each $v \in V$ at least one minimal path $\mathcal{P}_{\min}(s, v) = (s = v_0 \xrightarrow{D_1} v_1 \xrightarrow{D_2} \dots \xrightarrow{D_{l-1}} v_l = v)$, such that for each i , v_i sends the rumor v_{i+1} at time $\text{RBA}(s, v_i) + D_{i+1}$, and at this time v_{i+1} becomes informed for the first time.

Using this observation and a coupling argument, we can prove the following lemma.

Lemma 3.6. *For any pair of vertices $s, v \in V$ we have $\mathbf{E}[\overline{\text{SEQ}}(s, v)] \leq \frac{\mathbf{E}[\text{RBA}(s, v)]}{\delta}$.*

We are now ready to finish the proof of Theorem 3.2. For every pair of vertices $s, v \in V$,

$$\mathbf{E}[\text{UFPP}(s, v)] \leq 2 \cdot \mathbf{E}[\text{DFPP}(s, v)] = 2 \cdot \mathbf{E}[\overline{\text{SEQ}}(s, v)] \leq \frac{2}{\delta} \cdot \mathbf{E}[\text{RBA}(s, v)].$$

■

Theorem 3.7 ([25]). *Let $s, v \in V(G)$ with $s \neq v$. Then, $R(s, v) \leq \mathbf{E}[\text{UFPP}(s, v)]$.*

Combining the two theorems above we arrive at the main result of this section.

Theorem 3.8. *For any graph $G = (V, E)$ we have for every pair of vertices $s \neq v$,*

$$C(s, v) \leq 4 \cdot \frac{|E|}{\delta} \cdot \mathbf{E}[\text{RBA}(s, v)],$$

and hence $\text{COV}(G) = \mathcal{O}\left(\frac{|E|}{\delta} \cdot \log n \cdot \mathbf{E}[\text{RBA}(G)]\right)$ or equivalently, $\mathcal{R}(G) = \mathcal{O}\left(\frac{|E|}{\delta} \cdot \log n\right)$.

3.2. Applications

We start by giving examples for which the first inequality of Theorem 3.8 is asymptotically tight. For paths and cycles with n vertices, it is well-known that $\max_{s,v} C(s, v) = \Theta(n^2)$ (e.g., [24]) and Theorem 2.5 gives $\max_{s,v} \mathbf{E}[\text{RBA}(s, v)] \leq \mathbf{E}[\text{RBA}(G)] = \mathcal{O}(n)$. Similarly, for lollipop graphs (a complete graph with $2n/3$ vertices attached by a path of length $n/3$), $\max_{s,v} C(s, v) = \Theta(n^2)$ (e.g., [24]) and $\mathbf{E}[\text{RBA}(G)] = \mathcal{O}(n)$, and therefore the first inequality of Theorem 3.8 is also asymptotically tight for this highly non-regular graph.

The following overview in Figure 3.2 is based on [2, Chapter 5, p. 11], where we have added the corresponding broadcast times. It can be seen in Figure 3.2 that the second inequality of Theorem 3.8 is matched by complete k -ary trees with $k = \mathcal{O}(1)$. For complete graphs, expanders and hypercubes, the second inequality is tight up to a factor of $\mathcal{O}(\log n)$.

Graph	$\text{COV}(G)$	$\mathbf{E}[\text{RBA}(G)]$	$(1 - \lambda_2)^{-1}$
path/cycle	n^2 [24]	n (Thm. 2.5)	n^2 [2, Ch. 5, p. 11]
complete $\mathcal{O}(1)$ -ary tree	$n \log^2 n$ [32, Cor. 9]	$\log n$ (Thm. 2.5)	n [2, Ch. 5, p. 11]
complete graph	$n \log n$ [24]	$\log n$ [28]	1
expander	$n \log n$ [6]	$\log n$ [29]	1
hypercube	$n \log n$ [1]	$\log n$ [16]	$\log n$ [24]
$\sqrt{n} \times \sqrt{n}$ -torus	$n \log^2 n$ [32, Thm. 4]	\sqrt{n} (Thm. 2.5)	n [24]
$K_{n/2} \times K_2$	$n \log n$	$\log n$ [29]	n
lollipop	n^3 [24]	n	n^2 [2, Ch. 5, p. 22]

Figure 2: Comparison of the asymptotic order of the cover time, broadcast time and spectral gap of various graph classes. Recall that by Theorem 2.1, $(1 - \lambda_2)^{-1}$ captures the mixing time up to logarithmic factors.

Let us consider the graph $K_{n/2} \times K_2$. One can easily verify that $\text{COV}(G) = \mathcal{O}(n \log n)$, $\mathbf{E}[\text{RBA}(G)] = \mathcal{O}(\log n)$, but $(1 - \lambda_2)^{-1} = \Omega(n)$ (and consequently $\text{MIX}_{e-1}(G) = \Omega(n)$). Comparing these values with the ones of the complete graph, we see that there are graphs with an optimal cover time and optimal broadcast time, but $(1 - \lambda_2)^{-1}$ may vary between $\Theta(1)$ and $\Omega(n)$. Hence the upper bound on the cover time based on the broadcast time can be a polynomial factor smaller than the corresponding bound (Theorem 2.4) based on the spectral gap $1 - \lambda_2$. On the other hand, the following remark shows that by using the broadcast time instead of the spectral gap, we never lose more than a $\log^2 n$ factor:

Remark 3.9. For any regular graph G , the second bound of Theorem 3.8 implies

$$\text{COV}(G) = \mathcal{O}\left(\frac{1}{1 - \lambda_2} \cdot n \log^3 n\right).$$

In addition, Theorem 3.8 implies directly the following well-known bounds.

- (1) Since $\mathbf{E}[\text{RBA}(G)] = \mathcal{O}(n)$ for regular graphs [12, Prop. 1], we obtain $\max_{u,v} \mathcal{C}(u, v) = \mathcal{O}(n^2)$ for regular graphs [2, Ch. 6, Cor. 9].
- (2) For bounded degree graphs, $\mathbf{E}[\text{RBA}(G)] = \mathcal{O}(\text{diam}(G))$ (by Theorem 2.5) implies $\max_{u,v} \mathcal{C}(u, v) = \mathcal{O}(n \text{diam}(G))$ [2, Ch. 6, Cor. 8].
- (3) Since $\max_{u,v} \mathbf{E}[\text{RBA}(u, v)] = \mathcal{O}(n)$ [16], we obtain $\max_{u,v} \mathcal{C}(u, v) = \mathcal{O}(n^3)$ [2, Ch. 6, Thm. 1].

Finally, we give an application of Theorem 3.8 to certain power law random graphs (such networks are used to model real world networks [9]).

Definition 3.10. Given an n -dimensional vector $\mathbf{d} = (d_1, d_2, \dots, d_n)$, the *generalized random graph* $G(\mathbf{d})$ is constructed as follows. Each edge $\{i, j\}, 1 \leq i, j \leq n$ exists with prob. $\frac{d_i \cdot d_j}{\sum_{k=1}^n d_k}$, independently of all other edges.

Theorem 3.11 ([11]). *Let \mathbf{d} be a vector such that for all i , $d_i > \log^c n$, where $c > 2$ is some constant, and the number of vertices with expected degree d is proportional to $(d - \log^c n)^{-1}$. Then, $G(\mathbf{d})$ satisfies $\text{RBA}_{n-1}(G(\mathbf{d})) = \mathcal{O}(\log n)$ with probability $1 - o(1)$.*

Since the number of edges satisfies $|E(G(\mathbf{d}))| = \mathcal{O}(n \log^c n)$ with probability $1 - o(1)$ [9], we obtain by combining the theorem above with Theorem 3.8:

Corollary 3.12. *For $G(\mathbf{d})$ as in Theorem 3.11 we have $\text{COV}(G(\mathbf{d})) = \mathcal{O}(n \log^2 n)$ with probability $1 - o(1)$.*

4. Lower Bounds on $\mathcal{R}(G)$

4.1. Sparse Graphs

Definition 4.1. Given a graph $G = (V, E)$, a set $\Pi \subseteq E(G)$ is called a *cutset separating* $u \in V$ from $v \in V$ if every path from u to v includes an edge of Π .

Proposition 4.2 ([23, p. 59],[26]). *For $\{\Pi_i\}_{i=1}^k, k \in \mathbb{N}$, being disjoint cutsets separating u from v , $\mathcal{R}(u, v) \geq \sum_{i=1}^k |\Pi_i|^{-1}$.*

Zuckerman [32] proved that for any two vertices u, v on a tree, $\mathcal{H}(u, v) \geq \text{dist}(u, v)^2$. Using Proposition 4.2, we obtain the following generalization (a similar, but less tight bound follows from a result of [7]).

Corollary 4.3. *For any $u, v \in V$ of any graph G , $\mathcal{C}(u, v) \geq 2 \cdot \text{dist}(u, v)^2$. On the other hand, there are graphs G and $u, v \in V$ such that $\mathcal{H}(u, v) = \Theta(\text{dist}(u, v)) = o(\text{dist}(u, v)^2)$.*

We remark that Corollary 4.3 is exact for paths (cf. [24]). Combining Corollary 4.3 with the known bounds from Theorem 2.3 and Theorem 2.5 yields:

Proposition 4.4. *For any graph G with maximum degree δ , $\mathcal{R}(G) = \Omega\left(\frac{\sqrt{n}}{\Delta} \cdot \sqrt{\log n}\right)$.*

As demonstrated by the $\sqrt{n} \times \sqrt{n}$ -torus where $\text{RBA}_{n-1}(G) = \Theta(\sqrt{n})$ (by Theorem 2.5) and $\text{COV}(G) = \Theta(n \log^2 n)$ [32], this bound is tight up to a factor of $\log^{3/2} n$ for bounded degree graphs. The next result improves over Proposition 4.4 for dense graphs.

Theorem 4.5. *For any graph G with $\Delta = \mathcal{O}(\delta)$, $\mathcal{R}(G) = \Omega(\frac{\sqrt{n}}{\sqrt{\delta}} \cdot \frac{1}{\log n})$.*

4.2. Dense Graphs

In this section we present results that are tailored for dense graphs, e.g., graphs with minimum degree $\Theta(n)$. Consider a random walk $X_0 = s, X_1, \dots$ on G starting from s . Denote the number of visits to u until time t as $W_t(s, u) := |\{0 \leq t' \leq t : X_{t'} = u\}|$.

Definition 4.6 ([20, 31]). Consider a graph $G = (V, E)$ and a random walk starting from $s \in V$. Let

$$\text{BLA}(s) := \mathbf{E} \left[\min \left\{ t \in \mathbb{N} \mid \forall u \in V : \frac{1}{2} \cdot t\pi(u) \leq W_t(s, u) \leq 2 \cdot t\pi(u) \right\} \right].$$

Then, the *blanket time* of G is defined as $\text{BLA}(G) := \max_{s \in V} \text{BLA}(s)$.

Theorem 4.7 ([20]). *For any graph $G = (V, E)$, $\text{BLA}(G) = \mathcal{O}(\text{COV}(G) \cdot (\log \log n)^2)$.*

We also require the following simple graph-theoretical lemma.

Lemma 4.8. *For every graph G , there is a 2-cover X of G with $|X| \leq \lceil \frac{n}{\delta} \rceil$, i. e., there is a set $X \subseteq V$ such that for all $v \in V$ there is an $x \in X$ with $\text{dist}(x, v) \leq 2$.*

Interestingly, it is known that there are graphs with minimum degree $\frac{n}{2}$ for which every 1-cover (i. e., dominating set) is of size $\Theta(\log n)$ [4], while the lemma above shows that every such graph has a 2-cover of constant size. We now prove the main result of Section 4.

Theorem 4.9. *For any graph with $\Delta = \mathcal{O}(\delta)$, $\mathbf{E}[\text{RBA}(G)] = \mathcal{O}(\frac{1}{\delta} \cdot \text{BLA}(G) + \frac{n^2}{\delta^2} \cdot \log^2 n)$.*

The following corollary follows immediately from Theorem 4.9 and Theorem 4.7.

Corollary 4.10. *For any graph $G = (V, E)$ with $\Delta = \mathcal{O}(\delta)$ we have $\mathcal{R}(G) = \Omega(\frac{\delta^2}{n} \cdot \frac{1}{\log n})$.*

Combining Corollary 4.10 with Theorem 3.8 for graphs with minimum degree $\Theta(n)$, we see that the cover time equals the broadcast time multiplied by n up to logarithmic factors. It is worth mentioning that for graphs with $\delta \geq \lfloor \frac{n}{2} \rfloor$, Chandra et al. [8, Theorem 3.3] proved that $\text{COV}(G) = \Theta(n \log n)$. As pointed out by the same authors, $\text{COV}(G)$ may be between $n \log n$ and $\Theta(n^2)$ if $\delta < \lfloor \frac{n}{2} \rfloor$. Now, Corollary 4.10 provides a parameter (the broadcast time) that captures the cover time not only for $\delta \geq \lfloor \frac{n}{2} \rfloor$, but also for $\delta = \Omega(n)$.

Proof of Theorem 4.9. Let us briefly describe the main idea of the proof. We first show that for every vertex u there is a fixed (independent of a concrete execution of RBA) set of vertices $Y(u) \subseteq V$ of size at least $\delta/12$ such that u informs each vertex in $Y(u)$ within $\mathcal{O}((n/\delta) \cdot \log^2 n)$ steps with high probability. We then establish that if a vertex u informs v in $\mathcal{O}((n/\delta) \cdot \log^2 n)$ steps with high probability, then also v informs u in $\mathcal{O}((n/\delta) \cdot \log^2 n)$ steps with high probability. Using this fact and Lemma 4.8 we find that there is a partitioning of V into a constant number of partitions with the following property: once a vertex in such a partition becomes informed, the whole partition becomes informed within $\mathcal{O}((n/\delta) \cdot \log^2 n)$ steps. Finally, we use a coupling between the random walk and the broadcast algorithm to show that if the random walk covers the whole graph quickly, then the rumor will also be quickly propagated from one partition to the other partitions. The formal proof follows.

Lemma 4.11. *For each $u \in V$ there is a set $Y(u) \subseteq V$ (independent of the execution of RBA) of size at least $\delta/12$ such that for every $v \in Y(u)$, $\text{RBA}_{n-4}(u, v) \leq 16C_1 \frac{n}{\delta} \log^2 n$, where $C_1 > 0$ is some constant.*

Lemma 4.12. *For any two vertices u, v in a graph G with $\Delta = \mathcal{O}(\delta)$, $\text{RBA}_{n-4}(v, u) \leq C_2 \cdot (\text{RBA}_{n-4}(u, v) + \log n)$, where $C_2 > 0$ is some constant.*

Consider the undirected auxiliary graph $\widehat{G} = (\widehat{V}, \widehat{E})$ defined as follows: $\widehat{V} := V$ and $\{u, v\} \in \widehat{E}$ iff

$$\max\{\text{RBA}_{n-4}(u, v), \text{RBA}_{n-4}(v, u)\} \leq C_2 \cdot \left(16C_1 \frac{n}{\delta} \log^2 n + \log n\right).$$

By the two lemmas above, $\delta(\widehat{G}) \geq \delta/12$. Hence Lemma 4.8 implies the existence of a 2-cover $\{u_1, u_2, \dots, u_k\}$, $k \leq \lceil n/\delta \rceil$, of \widehat{G} . Therefore, the sets $U_i := \{v \in \widehat{V} \mid \text{dist}_{\widehat{G}}(v, u_i) \leq 2\}$, $1 \leq i \leq k$ form a (possibly non-disjoint) partitioning of \widehat{V} . Take a disjoint partitioning V_1, V_2, \dots, V_k such that for every $1 \leq i \leq k$, $V_i \subseteq U_i$. Consider now the directed graph $G' := (V', E')$ with $V' := \{V_1, V_2, \dots, V_k\}$ and

$$E' := \left\{ (V_i, V_j) \mid \exists u \in V_i, 1 \leq t \leq 4 \cdot \frac{\deg(u)}{2|E|} \cdot \text{BLA}(G) : N_{t,u} \in V_j \right\},$$

where $N_{t,u} \in N(u)$ is the vertex to which the random walk moves after the t -th visit of u .

Claim 4.13. *Let $s \in V_i$. With prob. $1/2$, there is a path from V_i to every V_j in G' .*

Reconsider now the partitioning V_1, V_2, \dots, V_k , $k \leq \lceil n/\delta \rceil$, of $\widehat{V} = V$. Let $\text{part}(u)$ be the function which assigns a vertex u the index of its partition. Let \mathcal{B} be the event that $\forall u \in V : V_{\text{part}(u)} \subseteq I_{\text{RBA}(s,u) + \mathcal{O}(\frac{n}{\delta} \log^2 n)}$ holds, i. e., for all $u \in V$ it holds that once u is informed, the partition $V_{\text{part}(u)}$ becomes completely informed within further $\mathcal{O}(\frac{n}{\delta} \log^2 n)$ steps. Fix some arbitrary vertex $u \in V$ and consider another vertex $w \in V_{\text{part}(u)}$. By definition of \widehat{G} and Lemma 4.12, there is path of length at most 4 from u to w in \widehat{G} . Hence once u is informed, w becomes informed within the next $\mathcal{O}(\frac{n}{\delta} \log^2 n)$ steps with probability $1 - 4n^{-4}$. Applying the union bound over $u \in V$ and $w \in V_{\text{part}(u)}$, we get $\Pr[\mathcal{B}] \geq 1 - 4n^{-2}$.

Claim 4.14. *Conditioned on the events \mathcal{A} and \mathcal{B} , all vertices of G become informed after $\mathcal{O}\left(\frac{1}{\delta} \cdot \text{BLA}(G) + \frac{n^2}{\delta^2} \cdot \log^2 n\right)$ steps.*

To finish the proof of the Theorem, we apply the union bound to get $\Pr[\mathcal{A} \wedge \mathcal{B}] \geq 1 - \frac{1}{2} - 4n^{-2}$. So, with probability larger than $1/3$, all vertices of G become informed after at most $\mathcal{O}\left(\frac{1}{\delta} \cdot \text{BLA}(G) + \frac{n^2}{\delta^2} \cdot \log^2 n\right)$ steps. Thus for every $k \in \mathbb{N}$, we succeed after $\mathcal{O}\left(k \cdot \left(\frac{1}{\delta} \cdot \text{BLA}(G) + \frac{n^2}{\delta^2} \cdot \log^2 n\right)\right)$ steps with probability $1 - (2/3)^k$ and hence the expected broadcast time is $\mathcal{O}\left(\frac{1}{\delta} \cdot \text{BLA}(G) + \frac{n^2}{\delta^2} \cdot \log^2 n\right)$. ■

4.3. Discussion

We first complement the lower bounds on $\mathcal{R}(G)$ by some concrete graphs. By a construction based on Harary graphs [18] and the two-dim. torus we obtain the following.

Theorem 4.15. *For any $\sqrt{n} \leq d \leq n - 1$, there is a d -regular graph G with $\mathcal{R}(G) = \mathcal{O}(d \cdot \log n)$. Moreover, for any $1 \leq d \leq \sqrt{n}$ there is a graph with minimum degree d and maximum degree $d + 1$ such that $\mathcal{R}(G) = \mathcal{O}(\sqrt{n} \cdot \log^2 n)$.*

While for certain degrees, a small polynomial gap remains between the examples of Theorem 4.15 and the bounds of Theorem 4.5 and Theorem 4.9 (cf. Figure 1), the quotient between cover time and diameter is minimized up to logarithmic factors by these examples.

Proposition 4.16. *For any graph G with $\Delta = \mathcal{O}(\delta)$, $\frac{\text{COV}(G)}{\text{diam}(G)} = \Omega(\max\{\sqrt{n}, \delta\} \cdot \sqrt{\log n})$.*

So far, in all considered graphs with a (nearly) optimal cover times and high broadcast time, the latter was caused by a large diameter. Therefore, one could try to throw in the lower bounds on $\text{diam}(G)$ and ask the following question: Does $\text{COV}(G) = \mathcal{O}(\text{polylog}(n) \cdot \max\{n \log n, \text{diam}(G)^2\}) \Leftrightarrow \mathbf{E}[\text{RBA}(G)] = \mathcal{O}(\text{polylog}(n) \cdot \max\{\text{diam}(G), \log n\})$ hold? The answer is that both directions can be refuted by counter-examples, even for graphs where minimum and maximum degree coincide (up to constant factors).

5. Conclusion

Inspired by the intuition of Chandra et al. [8] about the relationship between cover time of random walks and the runtime of randomized broadcast, we devised the first formal results relating both times. As our main result in Section 3, we proved that the cover time of any graph G is upper bounded by $\mathcal{O}(\frac{E}{\delta} \log n)$ times the broadcast time. This result is tight for many graphs (at least up to a factor of $\log n$) and gives an upper bound on the cover time that is at least as good (and in certain cases much tighter than) the previous bound based on the spectral gap [6]. Moreover, this result implies several classic bounds on the cover time and an almost optimal upper bound on the cover time of certain random graphs that are used to model real world networks. In Section 4 we derived lower bounds on the ratio between the cover time and broadcast time. Together with our upper bound of Section 3, we established a surprisingly strong correspondence between the cover time and broadcast time on dense graphs. This positive result was complemented by the construction of certain graphs to demonstrate that this strong correspondence cannot be extended to sparser graphs. Nevertheless, our lower and upper bounds show that the relationship between cover time and broadcast time is substantially stronger than the relationship between any of these parameters and the mixing time (or the closely related spectral gap). In particular, our findings provide evidence for the following hierarchy for regular graphs:

$$\text{low mixing time} \Rightarrow \text{low broadcast time} \Rightarrow \text{low cover time},$$

which extends the following known relations: low mixing time \Rightarrow low cover time ([1, 6, 10]) and low mixing time \Rightarrow low broadcast time ([5, 12, 29]).

References

- [1] D. J. Aldous. On the Time Taken by Random Walks on Finite Groups to Visit Every State. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, pages 361–374, 1983.
- [2] D. J. Aldous and J. A. Fill. *Reversible Markov Chains and Random Walks on Graphs*. (draft at <http://www.stat.berkeley.edu/aldous/RWG/book.html>), 2002.
- [3] R. Aleliunas, R.M. Karp, R.J. Lipton, L. Lovász, and C. Rackoff. Random Walks, Universal Traversal Sequences, and the Complexity of Maze Problems. In *20th IEEE Symp. on Found. of Computer Science (FOCS'79)*, pages 218–223, 1979.

- [4] N. Alon and J. Spencer. *The Probabilistic Method*. John Wiley & Sons, 2nd edition, 2000.
- [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(6):2508–2530, 2006.
- [6] A. Broder and A. Karlin. Bounds on the cover time. *Journal of Theoretical Prob.*, 2(1):101–120, 1989.
- [7] T. K. Carne. A transmutation formula for markov chains. *Bulletin des Sciences Mathematiques*, 2(4):399–405, 1985.
- [8] A.K. Chandra, P. Raghavan, W.L. Ruzzo, R. Smolensky, and P. Tiwari. The Electrical Resistance of a Graph Captures its Commute and Cover Times. *Computational Complexity*, 6(4):312–340, 1997.
- [9] F. Chung, L. Lu, and V. Vu. Eigenvalues of random power law graphs. *Annals of Combinatorics*, 7:21–33, 2003.
- [10] C. Cooper and A.M. Frieze. The Cover Time of Random Regular Graphs. *SIAM Journal of Discrete Mathematics*, 18(4):728–740, 2005.
- [11] R. Elsässer. On Randomized Broadcasting in Power Law Networks. In *20th International Symposium on Distributed Computing (DISC'06)*, pages 370–384, 2006.
- [12] 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'07)*, pages 163–174, 2007.
- [13] U. Feige. A Tight Lower Bound for the Cover Time of Random Walks on Graphs. *Random Structures & Algorithms*, 6(4):433–438, 1995.
- [14] U. Feige. A Tight Upper Bound for the Cover Time of Random Walks on Graphs. *Random Structures & Algorithms*, 6(1):51–54, 1995.
- [15] U. Feige. Collecting Coupons on Trees, and the Cover Time of Random Walks. *Comp. Complexity*, 6(4):341–356, 1997.
- [16] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized Broadcast in Networks. *Random Structures & Algorithms*, 1(4):447–460, 1990.
- [17] J.A. Fill and R. Pemantle. Percolation, first-passage percolation and covering times for richardson's model on the n -cube. *The Annals of Applied Probability*, 3:593–629, 1993.
- [18] F. Harary. The Maximum Connectivity of a Graph. *Proceedings of the National Academy of Sciences of the United States of America*, 48(7):1142–1146, 1962.
- [19] H.W. Hethcote. Mathematics of infectious diseases. *SIAM Review* 42, pages 599–653, 2000.
- [20] J.D. Kahn, J.H. Kim, L. Lovász, and V.H. Vu. The cover time, the blanket time and the Matthews bound. In *41st IEEE Symposium on Foundations of Computer Science (FOCS'00)*, pages 467–475, 2000.
- [21] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized Rumor Spreading. In *41st IEEE Symposium on Foundations of Computer Science (FOCS'00)*, pages 565–574, 2000.
- [22] W.O. Kermack and A.G. McKendrick. Contributions to the mathematical theory of epidemics. *Proceedings of the Royal Society*, 115A:700–721, 1927.
- [23] D.A. Levin, Y. Peres, and E.L. Wilmer. *Markov Chains and Mixing Times*. (draft at <http://www.oberlin.edu/markov/>), 2006.
- [24] L. Lovász. Random walks on graphs: A survey. *Combinatorics, Paul Erdős is Eighty*, 2:1–46, 1993.
- [25] R. Lyons, R. Pemantle, and Y. Peres. Resistance Bounds for First-Passage-Percolation and Maximum Flow. *Journal of Combinatorial Theory (Series A)*, 86(1):158–168, 1999.
- [26] C.St.J.A. Nash-Williams. Random walk and electric currents in networks. *Proceedings of the Cambridge Philosophical Society*, 55(1):181–194, 1959.
- [27] M. E. J. Newman. The spread of epidemic disease on networks. *Physical Review E* 66, 016128, 2002.
- [28] B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.
- [29] T. Sauerwald. On Mixing and Edge Expansion Properties in Randomized Broadcasting. In *18th International Symposium on Algorithms and Computation (ISAAC'07)*, pages 196–207, 2007.
- [30] A. Sinclair. Improved Bounds for Mixing Rates of Markov Chains and Multicommodity Flow. *Combinatorics, Probability & Computing*, 1:351–370, 1992.
- [31] P. Winkler and D. Zuckerman. Multiple Cover Time. *Random Structures & Alg.*, 9(4):403–411, 1996.
- [32] D. Zuckerman. A Technique for Lower Bounding the Cover Time. *SIAM Journal on Discrete Math.*, 5(1):81–87, 1992.