Multiple Random Walks on Graphs: Mixing Few to Cover Many

Nicolás Rivera
Department of Computer Science & Technology, University of Cambridge, UK
Instituto de Ingeniería Matemática, University of Valparaíso, Chile

Thomas Sauerwald
Department of Computer Science & Technology, University of Cambridge, UK

John Sylvester
Department of Computer Science & Technology, University of Cambridge, UK
School of Computing Science, University of Glasgow, UK

Abstract

Random walks on graphs are an essential primitive for many randomised algorithms and stochastic processes. It is natural to ask how much can be gained by running k multiple random walks independently and in parallel. Although the cover time of multiple walks has been investigated for many natural networks, the problem of finding a general characterisation of multiple cover times for *worst-case* start vertices (posed by Alon, Avin, Koucký, Kozma, Lotker, and Tuttle in 2008) remains an open problem.

First, we improve and tighten various bounds on the *stationary* cover time when k random walks start from vertices sampled from the stationary distribution. For example, we prove an unconditional lower bound of $\Omega((n/k)\log n)$ on the stationary cover time, holding for any n-vertex graph G and any $1 \leq k = o(n \log n)$. Secondly, we establish the *stationary* cover times of multiple walks on several fundamental networks up to constant factors. Thirdly, we present a framework characterising *worst-case* cover times in terms of *stationary* cover times and a novel, relaxed notion of mixing time for multiple walks called the *partial mixing time*. Roughly speaking, the partial mixing time only requires a specific portion of all random walks to be mixed. Using these new concepts, we can establish (or recover) the *worst-case* cover times for many networks including expanders, preferential attachment graphs, grids, binary trees and hypercubes.

2012 ACM Subject Classification Theory of computation → Random walks and Markov chains; Mathematics of computing → Stochastic processes

Keywords and phrases Multiple Random walks, Markov Chains, Random Walks, Cover Time

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.107

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2011.07893` [39]

Funding N.R., T.S. & J.S. were supported by ERC Starting Grant no. 679660 (DYNAMIC MARCH). N.R. was supported by FONDECYT grant number 3210805. J.S. was partially supported by ESPRC grant number EP/T004878/1.

Acknowledgements We thank Jonathan Hermon for some interesting and useful discussions, and Przemysław Gordinowicz for his feedback on an earlier version of this paper.

1 Introduction

A random walk on a graph is a stochastic process that at each time step chooses a neighbour of the current vertex as its next state. The fact that a random walk visits every vertex of a connected, undirected graph in polynomial time was first used to solve the undirected

© Nicolás Rivera, Thomas Sauerwald, and John Sylvester;
licensed under Creative Commons License CC-BY 4.0

48th International Colloquium on Automata, Languages, and Programming (ICALP 2021).
Editors: Nikhil Bansal, Emanuela Merelli, and James Worrell; Article No. 107; pp. 107:1–107:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$s-t$ connectivity problem in logarithmic space [4]. Since then random walks have become a fundamental primitive in the design of randomised algorithms which feature in approximation algorithms and sampling [32, 40], load balancing [23, 42], searching [19, 33], resource location [24], property testing [12, 26, 27], graph parameter estimation [7, 11] and biological applications [8, 20].

The fact that random walks are local and memoryless (Markov property) ensures they require very little space and relatively unaffected by changes in the environment, e.g., dynamically evolving graphs or graphs with edge failures. These properties make random walks a natural candidate for parallelisation, where running parallel walks has the potential of lower time overheads. One early instance of this idea are space-time trade-offs for the undirected $s-t$ connectivity problem [9, 18]. Other applications involving multiple random walks are sublinear algorithms [13], local clustering [6, 43] or epidemic processes on networks [29, 38].

Given the potential applications of multiple random walks in algorithms, it is important to understand fundamental properties of multiple random walks. The *speed up*, first introduced in [5], is the ratio of the worst-case cover time by a single random walk to the cover time of k parallel walks. Following [5] and subsequent works [15, 16, 22, 25, 41] our understanding of when and why a speed up is present has improved. In particular, various results in [5, 15, 16] establish that as long as the lengths of the walks are not smaller than the mixing, the speed-up is linear in k. However, there are still many challenging open problems, for example, understanding the effect of different start vertices or characterising the magnitude of speed-up in terms of graph properties, a problem already stated in [5]: "*...which leads us to wonder whether there is some other property of a graph that characterises the speed-up achieved by multiple random walks more crisply than hitting and mixing times*." Addressing the previous questions, we introduce new quantities and couplings for multiple random walks, that allow us to improve the state-of-the-art by refining, strengthening or extending results from previous works.

While there is an extensive body of research on the foundations of (single) random walks (and Markov chains), it seems surprisingly hard to transfer these results and develop a systematic theory of *multiple* random walks. One of the reasons is that processes involving multiple random walks often lead to questions about *short* random walks, e.g., shorter than the mixing time. Such short walks may arise in applications including generating random walk samples in massively parallel systems [28, 40], or in applications where random walk steps are expensive or subject to delays (e.g., when crawling social networks like Twitter [11]). The challenge of analysing *short* random walks (shorter than mixing or hitting time) has been mentioned not only in the area of multiple cover times (e.g., [15, Sec. 6]), but also in the contexts of concentration inequalities for random walks [31, p. 863] and property testing [13].

1.1 Our Contribution

Our first set of results provide several tight bounds on $t_{\mathsf{cov}}^{(k)}(\pi)$ in general (connected) graphs, where $t_{\mathsf{cov}}^{(k)}(\pi)$ is the expected time for each vertex to be visited by at least one of k independent walks each started from a vertex independently sampled from the stationary distribution π.

The main findings of Section 3 include:

- Proving general bounds of $\mathcal{O}\big(\big(\frac{|E|}{kd_{\mathsf{min}}}\big)^2\log^2 n\big)$, $\mathcal{O}\big(\frac{\max_{v\in V}\mathbb{E}_\pi[\,\tau_v\,]}{k}\log n\big)$ and $\mathcal{O}\big(\frac{|E|\log n}{kd_{\mathsf{min}}\sqrt{1-\lambda_2}}\big)$ on $t_{\mathsf{cov}}^{(k)}(\pi)$, where d_{min} is the minimum degree, $\mathbb{E}_\pi[\,\tau_v\,]$ is the single-walk hitting time of $v\in V$ from a stationary start vertex and λ_2 is the second largest eigenvalue of the transition matrix of the walk. All three bounds are tight for certain graphs. The first bound improves over [9], the second result is a Matthew's type bound for multiple random walks, and the third yields tight bounds for non-regular expanders such as preferential attachment graphs.

- We prove that for any graph G and $1 \leq k = o(n \log n)$, $t_{\mathsf{cov}}^{(k)}(\pi) = \Omega((n/k)\log n)$. Weaker versions of this bound were obtained in [16], holding only for certain values of k or under additional assumptions on the mixing time. Our result matches the fundamental $\Omega(n \log n)$ lower bound for single random walks ($k = 1$) [17], and generalises it in the sense that the total amount of work by all k stationary walks together for covering is always $\Omega(n \log n)$. We establish the $\Omega((n/k)\log n)$ bound by reducing the multiple walk process to a single, reversible Markov chain, and applying a general lower bound on stationary cover times [3].
- A technical tool that provides a bound on the lower tail of the cover time by k walks from stationary for graphs with a large and (relatively) symmetric set of hard to hit vertices (Lemma 9). When applied to 2d tori and binary trees this yields a tight lower bound.

In Section 4 we introduce a novel quantity for multiple walks we call *partial mixing*. Intuitively, instead of mixing all (or at least half) of the k walks, we only need to mix a specified number $\tilde{k}$ of them. We put this idea on a more formal footing and prove min-max theorems which relate worst case cover times $t_{\mathsf{cov}}^{(k)}$ to partial mixing times $t_{\mathsf{mix}}^{(\tilde{k},k)}$ and stationary cover times:

- For any graph G and any $1 \leq k \leq n$, we prove that:

$$t_{\mathsf{cov}}^{(k)} \leq 16 \cdot \min_{1 \leq \tilde{k} < k} \max\left(t_{\mathsf{mix}}^{(\tilde{k},k)}, t_{\mathsf{cov}}^{(\tilde{k})}(\pi)\right).$$

For now, we omit details such as the definition of the partial mixing time $t_{\mathsf{mix}}^{(\tilde{k},k)}$ as well as some max-min characterisations that serve as lower bounds (these can be found in Section 4). Intuitively these characterisations suggest that for any number of walks k, there is an "optimal" choice of $\tilde{k}$ so that one first waits until $\tilde{k}$ out of the k walks are mixed, and then considers only these $\tilde{k}$ stationary walks when covering the remainder of the graph.

This argument involving mixing only some walks extends and generalises previous results that involve mixing all (or at least a constant portion) of the k walks [5, 15, 16]. Previous approaches only imply a linear speed-up as long as the lengths of the walks are not shorter than the mixing time of a *single* random walk. In contrast, our characterisation may still yield tight bounds on the cover time for random walks that are much shorter than the mixing time.

To demonstrate how our insights can be used, we derive worst case cover times for several well-known graph classes. Due to space limitations we could not include this in the main body of this work, however we have summarised our results in Table 1. The corresponding results with full proofs can be found in corresponding section of the full paper [39]. As a first step to calculating $t_{\mathsf{cov}}^{(k)}$, we determine the stationary cover time; this is based on our bounds from Section 3. Secondly, we derive lower and upper bounds on the partial mixing times. Finally, with the stationary cover times and partial mixing times at hand, we can apply the characterisations from Section 4 to infer lower and upper bounds on the worst case stationary times. For some of those graphs the worst case cover times were already known before, while for, e.g., binary trees and preferential attachment graphs, our bounds are new.

- For the graph families of binary trees, cycles, d-dim. tori ($d = 2$ and $d \geq 3$), hypercube, clique, and (possibly non-regular) expanders we determine the cover time up to constants, for both worst-case and stationary start vertices (see Table 1 for the quantitative results).

We believe that this new methodology constitutes some progress towards the open question of Alon et al. [5] about a characterisation of worst-case cover times.

1.2 Novelty of Our Techniques

While a lot of the proof techniques in previous work [5, 15, 16, 41] are based on direct arguments such as mixing time (or relaxation time), our work introduces a number of new methods which, to the best of our knowledge, have not been used in the analysis of cover time of multiple walks before. In particular, one important novel concept is the introduction of the *partial mixing time*. The idea is that instead of waiting for all (or a constant portion of) k walks to mix, we can just mix some $\tilde{k} \leq k$ walks to reap the benefits of coupling these $\tilde{k}$ walks to stationary walks. This then presents a delicate balancing act where one must find an optimal $\tilde{k}$ minimising the overall bound on the cover time, for example in expanders the optimal $\tilde{k}$ is linear in k whereas in binary trees it is approximately $\sqrt{k}$, and for the cycle it is roughly $\log k$. This turning point reveals something about the structure of the graph and our results relating partial mixing to hitting time of sets helps one find this. Another tool we frequently use is a reduction to random walks with geometric resets, similar to a PageRank process, which allows us to relate *multiple* walks from stationary to a *single* reversible Markov chain.

2 Notation & Preliminaries

Throughout $G = (V, E)$ will be a finite undirected, connected graph with $n := |V|$ vertices and $m := |E|$ edges. For any $k \geq 1$, let $X_t = \left(X_t^{(1)}, \ldots, X_t^{(k)}\right)$ be *multiple random walk process*, where each $X_t^{(i)}$ is an independent random walk on G. Let

$$\mathbb{E}_{u_1,\ldots,u_k}[\,\cdot\,] := \mathbb{E}[\,\cdot \mid X_0 = (u_1, \ldots, u_k)\,]$$

denote the conditional expectation where, for each $1 \leq i \leq k$, $X_0^{(i)} = u_i \in V$ is the start vertex of the i^{th} walk. Unless mentioned otherwise, walks will be *lazy*, i.e., at each step the walk stays at its current location with probability 1/2, and otherwise moves to a neighbour chosen uniformly at random. We let the random variable $\tau_{\mathsf{cov}}^{(k)}(G) = \inf\{t : \bigcup_{i=0}^{t}\{X_i^{(1)}, \ldots, X_i^{(k)}\} = V\}$ be the first time every vertex of the graph has been visited by some walk $X_t^{(i)}$. For $u_1, \ldots, u_k \in V$ let

$$t_{\mathsf{cov}}^{(k)}((u_1, \ldots, u_k), G) = \mathbb{E}_{u_1,\ldots,u_k}\left[\tau_{\mathsf{cov}}^{(k)}(G)\right], \qquad t_{\mathsf{cov}}^{(k)}(G) = \max_{u_1,\ldots,u_k \in V} t_{\mathsf{cov}}^{(k)}((u_1, \ldots, u_k), G)$$

denote the cover time of k walks from $(u_1, \ldots, u_k)$ and the cover time of k walks from worst case start positions respectively. For simplicity, we drop G from the notation if the underlying graph is clear from the context. We shall use π to denote the stationary distribution of a single random walk on on a graph G, for $v \in V$ this is given by $\pi(v) = \frac{d(v)}{2m}$ which is the degree over twice the number of edges. We use π^k, which is a distribution on V^k given by the product measure of π with itself, to denote the stationary distribution of a multiple random walk. For a probability distribution μ on V let $\mathbb{E}_{\mu^k}[\,\cdot\,]$ denote expectation with respect to k walks where each start vertex is sampled independently from μ and

$$t_{\mathsf{cov}}^{(k)}(\mu, G) = \mathbb{E}_{\mu^k}\left[\tau_{\mathsf{cov}}^{(k)}(G)\right].$$

In particular $t_{\mathsf{cov}}^{(k)}(\pi, G)$ denotes the expected cover time from independent stationary start vertices. For a set $S \subseteq V$ we define

$$\tau_S^{(k)} = \inf\{t : \text{ there exists } 1 \leq i \leq k \text{ such that } X_t^{(i)} \in S\}$$

Table 1 All results above are $\Theta(\cdot)$, that is bounded above and below by a multiplicative constant, apart from the mixing time of expanders which is only bounded from above. PA above is the preferential attachment process where each vertex has m initial links, the results hold w.h.p., see [10, 34]. Cells shaded in Yellow are new results proved in this paper with the exception that for $k = \mathcal{O}(\log n)$ upper bounds on the stationary cover time for binary trees, expanders and preferential attachment graphs can be deduced from general bounds for the worst case cover time in [5]. Cells shaded Gray in the second to last column are known results we re-prove in this paper using our partial mixing time results, for the 2-dim grid we only re-prove upper bounds. References for the second to last column are given in the corresponding section in the full version, except for the Barbell, see [15, Page 2]. The Barbell consists of two cliques on $n/2$ vertices connected by single edge; we include this in the table as an interesting example where the speed up by stationary walks is exponential in k. All other results for single walks can be found in [2], for example.

Graph family	Cover t_{cov}	Hitting t_{hit}	Mixing t_{mix}	k-Cover Time, where $2 \le k \le n$: Worst case $t_{\mathsf{cov}}^{(k)}$	k-Cover Time, where $2 \le k \le n$: From π^k, $t_{\mathsf{cov}}^{(k)}(\pi)$
Binary tree	$n \log^2 n$	$n \log n$	n	$(n/k)\log^2 n$ **if** $k \le \log^2 n$. $(n/\sqrt{k})\log n$ **if** $k \ge \log^2 n$.	$\frac{n \log n}{k} \log\left(\frac{n \log n}{k}\right)$
Cycle	n^2	n^2	n^2	$\frac{n^2}{\log k}$	$\left(\frac{n}{k}\right)^2 \log^2 k$
2-Dim. Tori	$n \log^2 n$	$n \log n$	n	$(n/k)\log^2 n$ **if** $k \le \log^2 n$. $\frac{n}{\log(k/\log^2 n)}$ **if** $k \ge \log^2 n$.	$\frac{n \log n}{k} \log\left(\frac{n \log n}{k}\right)$
d-Dim. Tori $d \ge 3$	$n \log n$	n	$n^{2/d}$	$(n/k)\log n$ **if** $k \le n^{1-2/d}\log n$. $\frac{n^{2/d}}{\log(k/(n^{1-2/d}\log n))}$ **if** $k \ge n^{1-2/d}\log n$.	$\frac{n}{k}\log n$
Hypercube	$n \log n$	n	$\log n \log\log n$	$(n/k)\log n$ **if** $k \le n/\log\log n$. $\log n \log\log n$ **if** $k \ge n/\log\log n$.	$\frac{n}{k}\log n$
Expanders	$n \log n$	n	$\mathcal{O}(\log n)$	$\frac{n}{k}\log n$	$\frac{n}{k}\log n$
PA, $m \ge 2$	$n \log n$	n	$\mathcal{O}(\log n)$	$\frac{n}{k}\log n$	
Barbell	n^2	n^2	n^2	n^2/k	$\frac{2^{-k}n^2}{k} + \frac{n \log n}{k}$

as the first time the set S is visited by any of the k independent random walks, if $S = \{v\}$ is a singleton set we use τ_v, dropping brackets. Let

$$t_{\mathsf{hit}}^{(k)}(G) = \max_{u_1,\dots,u_k\in V} \max_{v\in V} \mathbb{E}_{u_1,\dots,u_k}\left[\tau_v^{(k)}\right]$$

be the worst case vertex to vertex hitting time. When talking about a single random walk we drop the (1) index, i.e. $t_{\mathsf{cov}}^{(1)}(G) = t_{\mathsf{cov}}(G)$; we also drop G from the notation when the graph is clear. If we wish the graph G to be clear we shall also use the notation $\mathbb{P}_{u,G}[\cdot]$ and $\mathbb{E}_{u,G}[\cdot]$. For a single random walk X_t with stationary distribution π and $x \in V$, let $d(t)$ and $s_x(t)$ be the total variation and separation distances for X_t given by

$$d(t) = \max_{x\in V} ||P_{x,\cdot}^t - \pi||_{TV} \quad \text{and} \quad s_x(t) = \max_{y\in V}\left[1 - \frac{P_{x,y}^t}{\pi(y)}\right],$$

where $P_{x,\cdot}^t$ is the t-step probability distribution of a random walk starting from x and, for probability measures μ, ν, $||\mu - \nu||_{TV} = \frac{1}{2}\sum_{x\in V} |\mu(x) - \nu(x)|$ is the total variation distance. Let $s(t) = \max_{x\in V} s_x(t)$, then for $0 \le \varepsilon \le 1$ the mixing and separation times [30, (4.32)] are

$$t_{\mathsf{mix}}(\varepsilon) = \inf\{t : d(t) \le \varepsilon\} \qquad \text{and} \qquad t_{\mathsf{sep}}(\varepsilon) = \inf\{t : s(t) \le \varepsilon\}, \tag{1}$$

and $t_{\mathsf{mix}} := t_{\mathsf{mix}}(1/4)$ and $t_{\mathsf{sep}} = t_{\mathsf{sep}}(1/e)$. A strong stationary time (SST) σ, see [30, Ch. 6] or [1], is a randomised stopping time for a Markov chain Y_t on V with stationary distribution π if

$$\mathbb{P}_u[Y_\sigma = v \mid \sigma = k] = \pi(v) \qquad \text{for any } u, v \in V \text{ and } k \ge 0. \tag{2}$$

Let $t_{\mathsf{rel}} = \frac{1}{1-\lambda_2}$ be the relaxation time of G, where λ_2 is the second largest eigenvalue of the transition matrix of the (lazy) random walk on G.

For random variables Y, Z we say that Y dominates Z ($Y \succeq Z$) if $\mathbb{P}[Y \ge x] \ge \mathbb{P}[Z \ge x]$ for all x.

3 Multiple Stationary Cover Times

We shall state our general upper and lower bound results for multiple walks from stationary in Sections 3.1 & 3.2 respectively. All proofs can be found in the full version [39].

3.1 Upper Bounds

Broder, Karlin, Raghavan, and Upfal [9] showed that for any graph G and $k \ge 1$,

$$t_{\mathsf{cov}}^{(k)}(\pi) = \mathcal{O}\left(\left(\frac{m}{k}\right)^2 \log^3 n\right).$$

We prove a general bound which improves this bound by a multiplicative factor of $d_{\mathsf{min}}^2 \log n$ which may be $\Omega(n^2 \log n)$ for some graphs.

▶ **Theorem 1.** *For any graph G and any $k \ge 1$,*

$$t_{\mathsf{cov}}^{(k)}(\pi) = \mathcal{O}\left(\left(\frac{m}{kd_{\mathsf{min}}}\right)^2 \log^2 n\right).$$

This bound is tight for the cycle if $k = n^{\Theta(1)}$, see Table 1. Theorem 1 is proved by relating the probability a vertex v is not hit up to a certain time t to the expected number of returns to v by a walk of length t from v and applying a bound by Oliveira and Peres [35].

The next bound is analogous to Matthew's bound [2, Theorem 2.26] for the cover time of single random walks from worst case, however it is proved by a different method.

▶ **Theorem 2.** *For any graph G and any $k \geq 1$, we have*

$$t_{\mathsf{cov}}^{(k)}(\pi) = \mathcal{O}\left(\frac{\max_{v\in V} \mathbb{E}_\pi\left[\,\tau_v\,\right]\log n}{k}\right).$$

This bound is tight for many graphs, see Table 1. Since the acceptance of this paper, the stronger bound $t_{\mathsf{cov}}^{(k)}(\pi) = \mathcal{O}(t_{\mathsf{cov}}/k)$ has been proved by Hermon & Sousi [21]. This bound implies Theorem 2 by a simple application of the aforementioned Matthew's Bound for single random walks. A version of Theorem 2 for $t_{\mathsf{cov}}^{(k)}$ was established by Alon et al. [5] provided $k = \mathcal{O}(\log n)$, the restriction on k is necessary (for worst case) as witnessed by the cycle. Theorem 2 also gives the following explicit bound.

▶ **Corollary 3.** *For any graph G and any $k \geq 1$, we have*

$$t_{\mathsf{cov}}^{(k)}(\pi) = \mathcal{O}\left(\frac{m}{kd_{\mathsf{min}}}\sqrt{t_{\mathsf{rel}}}\log n\right).$$

Proof. Use $\max_{v\in V} \mathbb{E}_\pi\left[\,\tau_v\,\right] \leq 20m\sqrt{t_{\mathsf{rel}}+1}/d_{\mathsf{min}}$ from [35, Theorem 1] in Theorem 2. ◀

Notice that, for all values $k \geq 1$, this bound is tight for any expander with $d_{\mathsf{min}} = \Omega(m/n)$, such as preferential attachment graphs (see Table 1 and the full version for more details).

We also establish the following two bounds for classes of graphs with "not too large" return probabilities.

▶ **Lemma 4.** *Let G be any graph satisfying $\pi_{\mathsf{min}} = \Omega(1/n)$, $t_{\mathsf{rel}} = o(n)$ and $\sum_{i=0}^{t} P_{vv}^i = \mathcal{O}(1 + t\pi(v))$ for any $t \leq t_{\mathsf{rel}}$. Then for any $1 \leq k \leq n$,*

$$t_{\mathsf{cov}}^{(k)}(\pi) = \Theta\left(\frac{n}{k}\log n\right).$$

The bound above applied to a broad class of graphs with expander like properties but large relaxation time, this includes the hypercube and high dimensional grids. The following bound holds for graphs with sub-harmonic return times, this includes binary trees and 2d-grid/tori.

▶ **Lemma 5.** *Let G be any graph with $\sum_{i=0}^{t} P_{v,v}^i = \mathcal{O}(t/n + \log t)$ for any $t \leq n(\log n)^2$ for all $v \in V$ and $t_{\mathsf{mix}} = \mathcal{O}(n)$. Then for any $1 \leq k \leq (n\log n)/3$,*

$$t_{\mathsf{cov}}^{(k)}(\pi) = \mathcal{O}\left(\frac{n\log n}{k}\log\left(\frac{n\log n}{k}\right)\right).$$

3.2 Lower Bounds

Generally speaking, lower bounds for random walks are more challenging to derive than upper bounds. In particular, the problem of obtaining a lower bound for the cover time of a simple random walk on an undirected graph was open for many years [2]. This was finally resolved by Feige [17] who proved $t_{\mathsf{cov}} \geq (1 - o(1))n\log n$. We prove a generalisation of this bound, up to constants, that holds for k random walks which start from stationarity (thus also for worst case).

▶ **Theorem 6.** *There exists a constant $c > 0$ such that for any graph G and $1 \leq k \leq c \cdot n \log n$,*

$$t_{\mathsf{cov}}^{(k)}(\pi) \geq c \cdot \frac{n}{k} \cdot \log n.$$

We remark that in this section all results hold (and are proven) for *non-lazy* random walks, which by stochastic domination implies that the same result also holds for *lazy* random walks. Theorem 6 is tight, uniformly for all $1 \leq k \leq n$, for the hypercube, expanders and high-dimensional tori, see Table 1. We note that [16] proved this bound for any start vertices under the additional assumption that $k \geq n^{\epsilon}$, for some constant $\varepsilon > 0$. One can track the constants in the proof of Theorem 6 and show that $c > 2 \cdot 10^{-11}$, we have not optimised this but note that $c \leq 1$ must hold in either condition of Theorem 6 due to the complete graph.

To prove this result we introduce the *geometric reset graph*, which allows us to couple the multiple random walk to a single walk to which we can apply a lower bound by Aldous [3]. The random reset graph is a small modification to a graph G which gives an edge-weighted graph $\widehat{G}(x)$ such that the simple random walk on $\widehat{G}(x)$ emulates a random walk on G with $\mathrm{Geo}(x)$ resets to stationarity, where $\mathrm{Geo}(x)$ is a geometric random variable with expectation $1/x$.

▶ **Definition 7** (The Geometric Reset Graph $\widehat{G}(x)$). *For any graph G the undirected, edge-weighted graph $\widehat{G}(x)$, where $0 < x \leq 1$, consists of all vertices $V(G)$ and one extra vertex z. All edges from G are included with edge-weight 1. Further, z is connected to each vertex $u \in G$ by an edge with edge-weight $x \cdot d(u)/(1-x)$, where $d(u)$ is the degree of vertex u in G.*

Given a graph with edge weights $\{w_e\}_{e \in E}$ the probability a *non-lazy* random walk moves from u to v is given by $w_{uv}/\sum_{w \in V} w_{uw}$. Thus the walk on $\widehat{G}(x)$ behaves as a random walk in G, apart from that in any step, it may move to the extra vertex z with probability $\frac{x \cdot d(u)/(1-x)}{x \cdot d(u)/(1-x) + d(u)} = x$. Once the walk is at z it moves back to a vertex $u \in V \backslash \{z\}$ with probability proportional to $d(u)$. Hence the stationary distribution $\widehat{\pi}$ of the random walk on $\widehat{G}(x)$ is proportional to π on G, and for the extra vertex z we have

$$\widehat{\pi}(z) = \frac{\sum_{v \in V} x d(u)/(1-x)}{\sum_{v \in V} d(u) + \sum_{v \in V} x d(u)/(1-x)} = \frac{x/(1-x)}{1 + x/(1-x)} = x.$$

Using the next lemma we can then obtain bounds on the multiple stationary cover time by simply bounding the cover time in the augmented graph $\widehat{G}(x)$ for some x.

▶ **Lemma 8.** *Let G be any graph G, $k \geq 1$ and $x = Ck/T$ where $C > 30$ and $T \geq 5Ck$. Then*

$$\mathbb{P}_{\pi^k, G}\left[\tau_{\mathsf{cov}}^{(k)} > \frac{T}{10Ck}\right] > \mathbb{P}_{\widehat{\pi}, \widehat{G}(x)}[\tau_{\mathsf{cov}} > T] - \exp\left(-\frac{Ck}{50}\right).$$

The coupling above will also be used later in the paper to prove a lower bound for the stationary cover time of the binary tree and 2-dimensional grid when k is small.

The next result we present utilises the second moment method to obtain a lower bound which works very well for $k = n^{\Theta(1)}$ walks on symmetric (e.g., transitive) graphs. In particular, we apply this to get tight lower bounds for cycles, 2-dim. tori and binary trees (see the corresponding section of the full version).

▶ **Lemma 9.** *Let G be any graph. Let $\alpha \in (0, 1)$ be a fixed real constant and define $p_v(t) = \mathbb{P}_{\pi}[\tau_v \leq t]$ for $t \geq 1$. Suppose there exists a subset $S \subseteq V$, and real numbers $p > 0$ and $0 \leq \varepsilon < 1$ such that for all $v \in S$ we have $p(1-\varepsilon) \leq p_v(t) \leq p$, with $p \leq \alpha(\log n)/k$, and that $\min_{v \in S} \pi(v) = \Omega(1/|S|)$. If in addition $p^2 k = o(1)$, then*

$$\mathbb{P}_{\pi^k}\left[\tau_{\mathsf{cov}}^{(k)} < t\right] = \mathcal{O}\left(\frac{(\log n)^2 n^{\alpha(1+\varepsilon)}}{k}\right).$$

4 Mixing Few Walks to Cover Many Vertices

In this section we present several bounds on $t_{\mathsf{cov}}^{(k)}$, the multiple cover time from worst case start vertices, based on $t_{\mathsf{cov}}^{(k)}(\pi)$, the multiple cover time from stationarity, and a new notion that we call *partial* mixing time. The intuition behind this is that on many graphs such as cycles or binary trees, only a certain number, say $\tilde{k}$ out of k walks will be able to reach vertices that are "far away" from the initial distribution. That means covering the whole graph hinges on how quickly these $\tilde{k}$ "mixed" walks cover the graph G, however, we also need to take into account the number of steps needed to "mix" those. Theorem 16 (see Subsection 4.2) makes this intuition more precise and suggests that the best strategy for covering a graph might be when $\tilde{k}$ is chosen so that the time to mix $\tilde{k}$ out of k walks and the stationary cover time of $\tilde{k}$ walks are approximately equal. The first subsection (Subsection 4.1) contains details of our new notions of mixing for multiple random walks, the second contains the bounds on worst-case cover times we derive from these and the third contains some bounds on the multiple mixing times.

4.1 Two Notions of Mixing for Multiple Random Walks

Recall the definition of strong stationary time (SST) given by (2) in Section 2. Then, for any graph G, and any $1 \leq \tilde{k} < k$, we define the *partial mixing time*:

$$\begin{aligned} t_{\mathsf{mix}}^{(\tilde{k},k)}(G) &= \inf\left\{ t \geq 1 : \textit{there exists an SST } \tau \text{ such that } \min_{v\in V} \mathbb{P}_v[\tau \leq t] \geq \tilde{k}/k \right\} \\ &= \inf\left\{ t \geq 1 : s(t) \leq 1 - \tilde{k}/k \right\}. \end{aligned} \tag{3}$$

Note that the two definitions above are equivalent by the following result.

▶ **Proposition 10** ([1, Proposition 3.2]). *If σ is an SST then $\mathbb{P}[\sigma > t] \geq s(t)$ for any $t \geq 0$. Furthermore there exists an SST for which equality holds.*

This notion of mixing, based on the idea of separation distance and strong stationary times for single walks, will be useful for establishing an upper bound on the worst case cover time. For lower bounds on the cover time we will now introduce another notion of mixing for multiple random walks based on a different property of mixing times of single walks.

For single random walks, there is a fundamental connection between mixing times and hitting times of sets. In particular if we let

$$t_{\mathsf{H}}(\alpha) = \max_{u\in V, S\subseteq V:\, \pi(S)\geq\alpha} \mathbb{E}_u[\tau_S], \qquad \text{and} \qquad t_{\mathsf{H}} := t_{\mathsf{H}}(1/4),$$

then the following theorem shows this *large-set hitting time* is equivalent to the mixing time.

▶ **Theorem 11** ([36] and independently [37]). *Let $\alpha < 1/2$. Then there exist positive constants $c(\alpha)$ and $C(\alpha)$ so that for every reversible chain*

$$c(\alpha) \cdot t_{\mathsf{H}}(\alpha) \leq t_{\mathsf{mix}}(\alpha) \leq C(\alpha) \cdot t_{\mathsf{H}}(\alpha).$$

In order to prove a lower bounds on the cover time we seek to replace the partial mixing time by an analogue of hitting times of large sets, adapted to multiple walks:

$$t_{\mathsf{large-hit}}^{(\tilde{k},k)}(G) := \min\left\{ t \geq 1 : \min_{u\in V, S\subseteq V:\, \pi(S)\geq 1/4} \mathbb{P}_u[\tau_S \leq t] \geq \frac{\tilde{k}}{k} \right\}. \tag{4}$$

Note that both of our mixing times, (4) and (3), are only defined for $\tilde{k} < k$. However, by the union bound, there exists a $C < \infty$ such that if we run k walks for $Ct_{\mathsf{mix}} \log k$ steps then all k walks will be close to uniform in total variation norm.

In the following four results, we present some simple relations between $t_{\mathsf{mix}}^{(\tilde{k},k)}$ and $t_{\mathsf{large-hit}}^{(\tilde{k},k)}$, and t_{mix}, where t_{mix} is the total variation mixing time for a single random walk given by (1).

First we show a simple upper bound in terms of the single walk mixing time.

▶ **Lemma 12.** *There exists a constant $C < \infty$ such that for any graph and $1 \leq \tilde{k} < k$ we have*

(i) $t_{\mathsf{mix}}^{(\tilde{k},k)} \leq 2 \cdot t_{\mathsf{mix}} \cdot \log\left(\frac{4k}{k-\tilde{k}}\right)$,

(ii) $t_{\mathsf{large-hit}}^{(\tilde{k},k)} \leq C \cdot t_{\mathsf{mix}} \cdot \log\left(\frac{k}{k-\tilde{k}}\right)$.

The partial mixing time can be bounded from below quite simply by mixing time.

▶ **Lemma 13.** *For any graph and $1 \leq \tilde{k} < k$ we have*

$$t_{\mathsf{mix}}^{(\tilde{k},k)} \geq t_{\mathsf{mix}}\left(1 - \frac{\tilde{k}}{k}\right).$$

We would prefer a bound in terms of $t_{\mathsf{mix}} := t_{\mathsf{mix}}(1/4)$ instead of $t_{\mathsf{mix}}(1 - \tilde{k}/k)$ as the former is easier to compute for most graphs. The following Lemma establishes such a lower bound for both notions of mixing time at the cost of a $\tilde{k}/k$ factor.

▶ **Lemma 14.** *There exists some constant $c > 0$ such that for any graph and $1 \leq \tilde{k} < k$ we have*

(i) $t_{\mathsf{mix}}^{(\tilde{k},k)} \geq c \cdot \frac{\tilde{k}}{k} \cdot t_{\mathsf{mix}}$,

(ii) $t_{\mathsf{large-hit}}^{(\tilde{k},k)} \geq c \cdot \frac{\tilde{k}}{k} \cdot t_{\mathsf{mix}}$.

We leave as an open problem whether our two notions of mixing for multiple random walks are equivalent up to constants, but the next result gives partial progress in one direction.

▶ **Lemma 15.** *For any graph and $1 \leq \tilde{k} < k/4$ we have*

$$t_{\mathsf{large-hit}}^{(\tilde{k},k)} \leq t_{\mathsf{mix}}^{(4\tilde{k},k)} + 1 \leq 2t_{\mathsf{mix}}^{(4\tilde{k},k)}.$$

4.2 Upper and Lower Bounds by Partial Mixing

Armed with our new notions of mixing time for multiple random walks from Section 4.1, we can now use them to prove upper and lower bounds on the worst case cover time in terms of stationary cover times and partial mixing times. We begin with the upper bound.

▶ **Theorem 16.** *For any graph G and any $1 \leq k \leq n$,*

$$t_{\mathsf{cov}}^{(k)} \leq 16 \cdot \min_{1 \leq \tilde{k} < k} \max\left(t_{\mathsf{mix}}^{(\tilde{k},k)}, t_{\mathsf{cov}}^{(\tilde{k})}(\pi)\right).$$

Proof. Fix any $1 \leq \tilde{k} < k$. It suffices to prove that with k walks starting from arbitrary positions running for

$$t := t_{\mathsf{mix}}^{(\tilde{k},k)} + 2t_{\mathsf{cov}}^{(\tilde{k})}(\pi) \leq 4 \cdot \max\left(t_{\mathsf{mix}}^{(\tilde{k},k)}, t_{\mathsf{cov}}^{(\tilde{k})}(\pi)\right)$$

steps, we cover G with probability at least $1/4$. Consider a single walk $X_1(t)$ on G. From (3), we have that at time $T = t_{\mathsf{mix}}^{(\tilde{k},k)}$ there exists a probability measure ν_v on V such that,

$$P_{v,w}^T = (1 - s_v(T))\pi(w) + s_v(T)\nu_v(w).$$

Therefore, we can generate $X_1(T)$ as follows: with probability $1 - s_v(T) \geq \tilde{k}/k$ we sample from π, otherwise we sample from ν_v. If we now consider k independent walks, the number of walks that are sampled at time T from π has a binomial distribution $\text{Bin}(k, \tilde{k}/k)$ with k trials and probability $\tilde{k}/k$, whose expectation is $\tilde{k}$. Since the expectation $\tilde{k}$ is an integer it is equal to the median, thus with probability at least $1/2$, at least $\tilde{k}$ walks are sampled from the stationary distribution. Now, consider only the $\tilde{k}$ independent walks starting from π. After $2t_{\mathsf{cov}}(\pi, \tilde{k})$ steps, these walks will cover G with probability at least $1/2$, due to Markov's inequality.

We conclude that in t time steps, from any starting configuration of the k walks, the probability we cover the graph is at least $1/4$. Hence in expectation, after (at most) 4 periods of length t we cover the graph. ◀

This theorem improves on various results in [5] and [15] which bound the worst case cover time by mixing all k walks, and it also generalises a previous result in [16, Lemma 3.1], where most walks were mixed, i.e., $\tilde{k} = k/2$.

We also prove a lower bound for cover times, however this involves the related definition of partial mixing time based on the hitting times of large sets.

▶ **Theorem 17.** *For any graph G with $\pi_{\max} = \max_u \pi(u)$ and any $1 \leq k \leq n$,*

$$t_{\mathsf{cov}}^{(k)} \geq \frac{1}{16} \cdot \max_{1 \leq \tilde{k} < k} \min\left(t_{\mathsf{large-hit}}^{(\tilde{k},k)}, \frac{1}{\tilde{k}\pi_{\max}} \right).$$

Further, for any regular graph G any $\delta > 0$ fixed, there is a constant $C = C(\delta) > 0$ such that

$$t_{\mathsf{cov}}^{(k)} \geq C \cdot \max_{n^\delta \leq \tilde{k} < k} \min\left(t_{\mathsf{large-hit}}^{(\tilde{k},k)}, \frac{n \log n}{\tilde{k}} \right).$$

As we will see later, both Theorem 16 and Theorem 17 yield asymptotically tight (or tight up to logarithmic factors) lower and upper bounds for many concrete networks. To explain why this is often the case, note that both bounds include one non-increasing function in $\tilde{k}$ and one non-decreasing in function in $\tilde{k}$. That means both bounds are optimised when the two functions are as close as possible. Then balancing the two functions in the upper bound asks for $\tilde{k}$ such that $t_{\mathsf{mix}}^{(\tilde{k},k)} \approx t_{\mathsf{cov}}^{(\tilde{k})}(\pi)$. Similarly, balancing the two functions in the first lower bound demands $t_{\mathsf{large-hit}}^{(\tilde{k},k)} \approx n/\tilde{k}$ (assuming $\pi_{\max} = O(1/n)$). Hence for any graph G where $t_{\mathsf{mix}}^{(\tilde{k},k)} \approx t_{\mathsf{large-hit}}^{(\tilde{k},k)}$, and also $t_{\mathsf{cov}}^{(\tilde{k})}(\pi) \approx n/\tilde{k}$, the upper and lower bounds will be close. This turns out to be the case for many networks (see the corresponding section in the full version).

One exception where Theorem 17 is far from tight is the cycle, we shall also prove a min-max theorem but with a different notion of partial cover time which is tight for the cycle.

For a set $S \subseteq V$ we let $\tau_{\mathsf{cov}}^{(k)}(S)$ be the first time that every vertex in S has been visited by at least one of the k walks, thus $\tau_{\mathsf{cov}}^{(k)}(V) = \tau_{\mathsf{cov}}^{(k)}$. Then we define the set cover time

$$t_{\mathsf{large-cov}}^{(k)} = \min_{S:\pi(S)\geq 1/4} \max_{\mu} \mathbb{E}_{\mu^k}\left[\tau_{\mathsf{cov}}^{(k)}(S) \right],$$

where the first minimum is over all sets $S \subseteq V$ satisfying $\pi(S) \geq 1/4$ and the second is over all probability distributions μ on the set $\partial S = \{x \in S : \text{exists } y \in S^c,\ xy \in E\}$.

▶ **Theorem 18.** *For any graph G and any $1 \leq k \leq n$,*

$$t_{\mathsf{cov}}^{(k)} \geq \frac{1}{4} \cdot \max_{1 \leq \tilde{k} < k} \min\left(t_{\mathsf{large-hit}}^{(\tilde{k},k)}, t_{\mathsf{large-cov}}^{(k)}\right).$$

4.3 Geometric Lower Bounds on the Large-Hit and Large-Cover Times

We will now derive two useful lower bounds on $t_{\mathsf{large-hit}}^{(\tilde{k},k)}$, one based on the conductance of the graph, and a second one based on the distance to a large set the random walk needs to hit.

For two sets $A, B \subseteq V$ the *ergodic flow* $Q(A,B)$ is defined by $Q(A,B) = \sum_{a\in A, b\in B} \pi(a) P_{a,b}$, where $P_{a,b}$ denotes the transition matrix of a (lazy) single random walk. We define the *conductance* $\Phi(S)$ of a set $S \subseteq V$ with $\pi(S) \in (0, 1/2]$ to be

$$\Phi(S) = \frac{Q(S, S^c)}{\pi(S)} \quad \text{and let} \quad \Phi(G) = \min_{S \subseteq V, 0 < \pi(S) \leq 1/2} \Phi(S).$$

▶ **Lemma 19.** *For any graph G with conductance $\Phi(G)$, any $1 \leq \tilde{k} \leq k$,*

$$t_{\mathsf{large-hit}}^{(\tilde{k},k)} \geq \frac{\tilde{k}}{k} \cdot \frac{2}{\Phi(G)}.$$

We remark that a similar bound was used implicitly in [41, Proof of Theorem 1.1], where $t_{\mathsf{cov}}^{(k)} \geq \sqrt{\frac{n}{k \cdot \Phi(G)}}$ was shown.

The following lemmas will be useful to lower bound worst case cover times on cycles/tori.

▶ **Lemma 20.** *Let G be a d-dimensional torus with constant $d \geq 2$ (or cycle, $d = 1$), $u \in V$ and S be a set with $|S| \geq n/2$. Then for any $\tilde{k} \leq k/2$,*

$$t_{\mathsf{large-hit}}^{(\tilde{k},k)} = \Omega\left((\mathrm{dist}(u,S))^2 / \log(k/\tilde{k}) \right).$$

▶ **Lemma 21.** *Let $S \subseteq V$ be a subset of vertices with $\pi(S) \geq 1/4$, $t \geq 2$ be an integer and $k \geq 100$ such that for every $u \in S$,*

$$\sum_{s=0}^{t} P_{u,u}^{s} \geq 32 \cdot t \cdot \pi(u) \cdot k.$$

Then for any starting distribution μ of $k/8$ walks,

$$\mathbb{E}_{\mu^k}\left[\tau_{\mathsf{cov}}^{(k)}(S)\right] \geq t/5.$$

5 Conclusion & Open Problems

In this work, we derived several new bounds on multiple stationary and worst-case cover times. We also introduced a new quantity called *partial mixing time*, which extends the definition of mixing time from single random walks to multiple random walks. By means of a min-max characterisation, we proved that the partial mixing time connects the stationary and worst-case cover times, leading to tight lower and upper bounds for many graph classes.

In terms of worst-case bounds, Theorem 1 implies that for any regular graph G and any $k \geq 1$, $t_{\mathsf{cov}}^{(k)}(\pi) = \mathcal{O}\left(\left(\frac{n}{k}\right)^2 \log^2 n\right)$. This bound is tight for the cycle when k is polynomial in n but not for smaller k. We suspect that for any $k \geq 1$ the cycle is (asymptotically) the worst case for $t_{\mathsf{cov}}^{(k)}(\pi)$ amongst regular graphs, which suggests $t_{\mathsf{cov}}^{(k)}(\pi) = \mathcal{O}\left(\left(\frac{n}{k}\right)^2 \log^2 k\right)$.

Some of our results have been only proven for the independent stationary case, but it seems plausible they extend to the case where the k random walks start from the *same* vertex. For example, extending the bound $t^{(k)}_{\mathsf{cov}}(\pi) = \Omega((n/k)\log n)$ to this case would be very interesting.

In Theorem 2 we prove $t^{(k)}_{\mathsf{cov}}(\pi) = \mathcal{O}((\max_{v\in V} \mathbb{E}_\pi[\tau_v] \log n)/k)$, can we prove the stronger bound $t^{(k)}_{\mathsf{cov}}(\pi) = \mathcal{O}(1/k \cdot t_{\mathsf{cov}}(\pi, G))$ without assuming anything on the mixing time of G?

Although our min-max characterisations involving partial mixing time yields tight bounds for many natural graph classes, it would be interesting to establish a general approximation guarantee (or find graph classes that serve as counter-examples). For the former, we believe techniques such as Gaussian Processes and Majorising Measures used in the seminal work of Ding, Lee and Peres [14] could be very useful.

References

1 David Aldous and Persi Diaconis. Strong uniform times and finite random walks. *Adv. in Appl. Math.*, 8(1):69–97, 1987. doi:10.1016/0196-8858(87)90006-6.

2 David Aldous and James Allen Fill. Reversible Markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014. URL: https://www.stat.berkeley.edu/~aldous/RWG/book.html.

3 David J. Aldous. Lower bounds for covering times for reversible Markov chains and random walks on graphs. *J. Theoret. Probab.*, 2(1):91–100, 1989. doi:10.1007/BF01048272.

4 Romas Aleliunas, Richard M. Karp, Richard J. Lipton, László Lovász, and Charles Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *20th Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 29-31 October 1979*, pages 218–223. IEEE Computer Society, 1979. doi:10.1109/SFCS.1979.34.

5 Noga Alon, Chen Avin, Michal Koucký, Gady Kozma, Zvi Lotker, and Mark R. Tuttle. Many random walks are faster than one. *Combin. Probab. Comput.*, 20(4):481–502, 2011. doi:10.1017/S0963548311000125.

6 Reid Andersen, Fan R. K. Chung, and Kevin J. Lang. Local graph partitioning using pagerank vectors. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings*, pages 475–486. IEEE Computer Society, 2006. doi:10.1109/FOCS.2006.44.

7 Anna Ben-Hamou, Roberto I. Oliveira, and Yuval Peres. Estimating graph parameters via random walks with restarts. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1702–1714. SIAM, 2018. doi:10.1137/1.9781611975031.111.

8 Lucas Boczkowski, Brieuc Guinard, Amos Korman, Zvi Lotker, and Marc P. Renault. Random walks with multiple step lengths. In Michael A. Bender, Martin Farach-Colton, and Miguel A. Mosteiro, editors, *LATIN 2018: Theoretical Informatics - 13th Latin American Symposium, Buenos Aires, Argentina, April 16-19, 2018, Proceedings*, volume 10807 of *Lecture Notes in Computer Science*, pages 174–186. Springer, 2018. doi:10.1007/978-3-319-77404-6_14.

9 Andrei Z. Broder, Anna R. Karlin, Prabhakar Raghavan, and Eli Upfal. Trading space for time in undirected s-t connectivity. *SIAM J. Comput.*, 23(2):324–334, 1994. doi:10.1137/S0097539790190144.

10 Colin Cooper and Alan Frieze. The cover time of the preferential attachment graph. *J. Combin. Theory Ser. B*, 97(2):269–290, 2007. doi:10.1016/j.jctb.2006.05.007.

11 Colin Cooper, Tomasz Radzik, and Yiannis Siantos. Estimating network parameters using random walks. *Social Netw. Analys. Mining*, 4(1):168, 2014. doi:10.1007/s13278-014-0168-6.

12 Artur Czumaj, Morteza Monemizadeh, Krzysztof Onak, and Christian Sohler. Planar graphs: Random walks and bipartiteness testing. *Random Struct. Algorithms*, 55(1):104–124, 2019. doi:10.1002/rsa.20826.

13 Artur Czumaj and Christian Sohler. Testing expansion in bounded-degree graphs. *Comb. Probab. Comput.*, 19(5-6):693–709, 2010. doi:10.1017/S096354831000012X.

14 Jian Ding, James R. Lee, and Yuval Peres. Cover times, blanket times, and majorizing measures. *Annals of Mathematics*, 175(3):1409–1471, 2012. doi:10.4007/annals.2012.175.3.8.

15 Klim Efremenko and Omer Reingold. How well do random walks parallelize? In Irit Dinur, Klaus Jansen, Joseph Naor, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21-23, 2009. Proceedings*, volume 5687 of *Lecture Notes in Computer Science*, pages 476–489. Springer, 2009. doi:10.1007/978-3-642-03685-9_36.

16 Robert Elsässer and Thomas Sauerwald. Tight bounds for the cover time of multiple random walks. *Theoret. Comput. Sci.*, 412(24):2623–2641, 2011. doi:10.1016/j.tcs.2010.08.010.

17 Uriel Feige. A tight lower bound on the cover time for random walks on graphs. *Random Structures Algorithms*, 6(4):433–438, 1995. doi:10.1002/rsa.3240060406.

18 Uriel Feige. A spectrum of time-space trade-offs for undirected s-t connectivity. *J. Comput. Syst. Sci.*, 54(2):305–316, 1997. doi:10.1006/jcss.1997.1471.

19 Christos Gkantsidis, Milena Mihail, and Amin Saberi. Hybrid search schemes for unstructured peer-to-peer networks. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies, 13-17 March 2005, Miami, FL, USA*, pages 1526–1537. IEEE, 2005. doi:10.1109/INFCOM.2005.1498436.

20 Brieuc Guinard and Amos Korman. Tight bounds for the cover times of random walks with heterogeneous step lengths. In Christophe Paul and Markus Bläser, editors, *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France*, volume 154 of *LIPIcs*, pages 28:1–28:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.STACS.2020.28.

21 Jonathan Hermon and Perla Sousi. Covering a graph with independent walks, 2021. arXiv:2104.00665.

22 Andrej Ivaskovic, Adrian Kosowski, Dominik Pajak, and Thomas Sauerwald. Multiple random walks on paths and grids. In Heribert Vollmer and Brigitte Vallée, editors, *34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, 2017, Hannover, Germany*, volume 66 of *LIPIcs*, pages 44:1–44:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.STACS.2017.44.

23 David R. Karger and Matthias Ruhl. Simple efficient load balancing algorithms for peer-to-peer systems. In Phillip B. Gibbons and Micah Adler, editors, *SPAA 2004: Proceedings of the Sixteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures, June 27-30, 2004, Barcelona, Spain*, pages 36–43. ACM, 2004. doi:10.1145/1007912.1007919.

24 David Kempe, Jon M. Kleinberg, and Alan J. Demers. Spatial gossip and resource location protocols. In Jeffrey Scott Vitter, Paul G. Spirakis, and Mihalis Yannakakis, editors, *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*, pages 163–172. ACM, 2001. doi:10.1145/380752.380796.

25 Ralf Klasing, Adrian Kosowski, Dominik Pajak, and Thomas Sauerwald. The multi-agent rotor-router on the ring: a deterministic alternative to parallel random walks. In Panagiota Fatourou and Gadi Taubenfeld, editors, *ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013*, pages 365–374. ACM, 2013. doi:10.1145/2484239.2484260.

26 Akash Kumar, C. Seshadhri, and Andrew Stolman. Finding forbidden minors in sublinear time: A nˆ1/2+o(1)-query one-sided tester for minor closed properties on bounded degree graphs. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 509–520. IEEE Computer Society, 2018. doi:10.1109/FOCS.2018.00055.

27 Akash Kumar, C. Seshadhri, and Andrew Stolman. Random walks and forbidden minors II: a poly$(d\, \epsilon^{-1})$-query tester for minor-closed properties of bounded degree graphs. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 559–567. ACM, 2019. `doi:10.1145/3313276.3316330`.

28 Jakub Lacki, Slobodan Mitrovic, Krzysztof Onak, and Piotr Sankowski. Walking randomly, massively, and efficiently. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 364–377. ACM, 2020. `doi:10.1145/3357713.3384303`.

29 Henry Lam, Zhenming Liu, Michael Mitzenmacher, Xiaorui Sun, and Yajun Wang. Information dissemination via random walks in d-dimensional space. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1612–1622. SIAM, 2012. `doi:10.1137/1.9781611973099.128`.

30 David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009. With a chapter by James G. Propp and David B. Wilson.

31 Pascal Lezaud. Chernoff-type bound for finite Markov chains. *The Annals of Applied Probability*, 8(3):849–867, 1998. `doi:10.1214/aoap/1028903453`.

32 László Lovász. Random walks on graphs: a survey. In *Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993)*, volume 2 of *Bolyai Soc. Math. Stud.*, pages 353–397. János Bolyai Math. Soc., Budapest, 1996.

33 Qin Lv, Pei Cao, Edith Cohen, Kai Li, and Scott Shenker. Search and replication in unstructured peer-to-peer networks. In Kemal Ebcioglu, Keshav Pingali, and Alex Nicolau, editors, *Proceedings of the 16th international conference on Supercomputing, ICS 2002, New York City, NY, USA, June 22-26, 2002*, pages 84–95. ACM, 2002. `doi:10.1145/514191.514206`.

34 Milena Mihail, Christos H. Papadimitriou, and Amin Saberi. On certain connectivity properties of the internet topology. *J. Comput. Syst. Sci.*, 72(2):239–251, 2006. `doi:10.1016/j.jcss.2005.06.009`.

35 Roberto I. Oliveira and Yuval Peres. Random walks on graphs: new bounds on hitting, meeting, coalescing and returning. In Marni Mishna and J. Ian Munro, editors, *Proceedings of the Sixteenth Workshop on Analytic Algorithmics and Combinatorics, ANALCO 2019, San Diego, CA, USA, January 6, 2019*, pages 119–126. SIAM, 2019. `doi:10.1137/1.9781611975505.13`.

36 Roberto Imbuzeiro Oliveira. Mixing and hitting times for finite Markov chains. *Electron. J. Probab.*, 17:no. 70, 12, 2012. `doi:10.1214/EJP.v17-2274`.

37 Yuval Peres and Perla Sousi. Mixing times are hitting times of large sets. *J. Theoret. Probab.*, 28(2):488–519, 2015. `doi:10.1007/s10959-013-0497-9`.

38 Alberto Pettarin, Andrea Pietracaprina, Geppino Pucci, and Eli Upfal. Tight bounds on information dissemination in sparse mobile networks. In Cyril Gavoille and Pierre Fraigniaud, editors, *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011*, pages 355–362. ACM, 2011. `doi:10.1145/1993806.1993882`.

39 Nicolás Rivera, Thomas Sauerwald, and John Sylvester. Multiple random walks on graphs: Mixing few to cover many, 2020. `arXiv:2011.07893`.

40 Atish Das Sarma, Danupon Nanongkai, Gopal Pandurangan, and Prasad Tetali. Distributed random walks. *J. ACM*, 60(1):2:1–2:31, 2013. `doi:10.1145/2432622.2432624`.

41 Thomas Sauerwald. Expansion and the cover time of parallel random walks. In Andréa W. Richa and Rachid Guerraoui, editors, *Proceedings of the 29th Annual ACM Symposium on Principles of Distributed Computing, PODC 2010, Zurich, Switzerland, July 25-28, 2010*, pages 315–324. ACM, 2010. `doi:10.1145/1835698.1835776`.

42 Thomas Sauerwald and He Sun. Tight bounds for randomized load balancing on arbitrary network topologies. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 341–350. IEEE Computer Society, 2012. doi:10.1109/FOCS.2012.86.

43 Daniel A. Spielman and Shang-Hua Teng. A local clustering algorithm for massive graphs and its application to nearly linear time graph partitioning. *SIAM J. Comput.*, 42(1):1–26, 2013. doi:10.1137/080744888.