

Multiple Random Walks on Paths and Grids

Andrej Ivaskovic, Adrian Kosowski, Dominik Pająk, Thomas Sauerwald

► **To cite this version:**

Andrej Ivaskovic, Adrian Kosowski, Dominik Pająk, Thomas Sauerwald. Multiple Random Walks on Paths and Grids . STACS 2017 - 34th Symposium on Theoretical Aspects of Computer Science, Mar 2017, Hannover, Germany. pp.1-14, 10.4230/LIPIcs.STACS.2017.44 . hal-01669223

HAL Id: hal-01669223

<https://hal.inria.fr/hal-01669223>

Submitted on 7 Feb 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Multiple Random Walks on Paths and Grids*

Andrej Ivašković¹, Adrian Kosowski², Dominik Pajak³, and Thomas Sauerwald⁴

1 University of Cambridge, Computer Laboratory, Cambridge, UK

2 INRIA Paris, Paris, France

3 Wrocław University of Science and Technology, Wrocław, Poland

4 University of Cambridge, Computer Laboratory, Cambridge, UK

Abstract

We derive several new results on multiple random walks on “low dimensional” graphs.

First, inspired by an example of a *weighted* random walk on a path of three vertices given by Efremenko and Reingold [7], we prove the following dichotomy: as the path length n tends to infinity, we have a super-linear speed-up w.r.t. the cover time if and only if the number of walks k is equal to 2. An important ingredient of our proofs is the use of a continuous-time analogue of multiple random walks, which might be of independent interest. Finally, we also present the first tight bounds on the speed-up of the cover time for any d -dimensional grid with $d \geq 2$ being an arbitrary constant, and reveal a sharp transition between linear and logarithmic speed-up.

1998 ACM Subject Classification G.3 Probability and Statistics

Keywords and phrases random walks, randomized algorithms, parallel computing

Digital Object Identifier 10.4230/LIPIcs.STACS.2017.44


1 Introduction

Markov chains are a family of stochastic processes first considered by Markov in 1906. More than a century later, Markov chains are an ubiquitous tool in many areas, including computer science, mathematics, physics and many others (for more background on Markov chains we refer to the textbook [13]). Random walks on graphs are a fundamental instance of a Markov chain, where a particle starts at a given vertex, and at each time-step $1, 2, \dots$ moves to a random adjacent vertex. One fundamental quantity of random walks is the time it takes to visit all vertices, known as the *cover time*.

In computer science, random walks, or slightly more general, Markov chains, have been instrumental in various breakthrough results. In complexity theory for instance, random walks have played a key role in understanding the space-complexity of algorithms, amplifying the success probability of randomized algorithms and derandomizing algorithms (pseudorandomness). Many sampling tasks in areas like streaming algorithms or sub-linear algorithms can be only solved via random walks. In distributed computing applications, *multiple* random walks are used in protocols for dynamic and anonymous networks, which take advantage of the Markovian property and location-independence of the random walk process, as well as the benefits of random walk parallelization. Such applications include the fundamental tasks of graph exploration with a team of agents [1, 11, 5] and of spreading

* The first and the fourth author were supported by the ERC Starting grant “Dynamic March” funded by the European Research Council. The second and the third author were supported by the National Science Centre, Poland – grant number 2015/17/B/ST6/01897.





■ **Figure 1** Illustration of the asymptotic values of the speed-up $S_{cov}^{(k)}$ of the cover time of k independent multiple random walks with identical start vertices over the cover time of a single random walk.

information in a network using multiple tokens [17]. Concurrent random walks are also used in token circulation routines, ensuring mutual exclusion properties in self-stabilizing computation [9], sampling applications [18] and population protocols [16].

Although many of these applications involve multiple and possibly concurrent random walks, comparatively little is known about the theory of multiple random walks. In particular, one might be interested in the so-called **speed-up** that quantifies “how much are k random walks faster than one”. In one of the first studies, Broder et al. [3] derived bounds on cover time of k random walks, each starting from a stationary distribution. Their main motivation was to derive time-space tradeoffs for the s - t -connectivity problem on undirected graphs. However, being able to launch several random walks at randomly chosen vertices might not be always possible in large sampling applications such as crawling massive networks like Facebook or Twitter. Also from a theoretical perspective, it might be natural to study a “worst-case” version of the speed-up.

This might have been one reason why in 2006, Alon et al. [1] introduced the following intriguing mathematical notion of the speed-up: the speed-up $S_{cov}^{(k)}$ is the ratio of the single worst-case cover time of one random walk to the worst-case cover time of k independent and concurrent random walks (see Section 2 for a formal definition and further background on (concurrent) random walks). This definition of speed-up is very natural and captures the “parallelism” of random walks in the time they explore a network.

Unfortunately, determining speed-up $S_{cov}^{(k)}$, or equivalently determining the k -walk cover time exactly turns out to be surprisingly difficult, even for basic topologies [1, 7]. Nonetheless, Alon et al. [1] and subsequent studies [7, 19, 8] have revealed a surprisingly broad spectrum of how the speed-up behaves, sometimes even on the same graph, as illustrated in Figure 1.

Based on the two extreme cases of an expander and a cycle, Alon et al. [1] conjectured that the speed-up is always at most $O(k)$ and at least $\Omega(\log k)$. While some progress has been made towards establishing weaker version of these conjectures (cf. [8, 7]), it is still unknown whether they hold in full generality.

From a mathematical perspective, multiple random walks pose a number of intricate challenges, since unlike in case of a single random walk, we often have to look into complex



■ **Figure 2** The Markov chain given by Efremenko and Reingold [7]. The cover time for the single random walk equals $\frac{5}{1-\alpha}$, while the cover time for the two random walks starting from any endpoint is $\frac{2.25}{1-\alpha} + o(1/(1-\alpha))$, as $\alpha \rightarrow 1$.

distributions such as the hitting time distribution (as opposed to simply estimating the expectation in case of a single random walk). In addition to that, we often have to work with “short” random walks, meaning that the random walks’ distribution may be still far from the stationary distribution. This challenge has been identified in several previous works, including in sampling [14] and property-testing [6]. It was also mentioned by Efremenko and Reingold [7], who further highlighted the general complexity of multiple random walks by the following remarkable example.

Take a path with just three vertices (see also Figure 2 on the next page), and consider the speed-up of two random walks. The worst-case start vertex for the single random walk is the middle vertex, while it is not difficult to compute that the worst-case start configuration for two random walks is when they start from the same endpoint. Taking the ratio of the single walk to the 2-walk cover time, one obtains a speed-up of 2.25! This intriguing example serves as a motivation of our study here, as it leads to the following questions:

1. Can we find other, possibly larger, graphs for which the speed-up is super-linear?
2. What happens in case of the path, if k or n grows?
3. One subtle detail in the example (Figure 2) are loop probabilities that need to be sufficiently large. Are the loop probabilities necessary, and how do the loop probabilities impact the speed-up?

In this work we will shed some light on the questions above, in particular, the second and third one. Apart from that, we also analyze the speed-up on the higher-dimensional relatives of the path and derive the first complete asymptotic characterisation of $S_{\text{cov}}^{(k)}$.

Our Contributions. In the following, we will describe our main results in more detail and sketch how to derive them. Our first result extends the example of Efremenko and Reingold in two parameters, n and k , and reveals an interesting dichotomy:

► **Theorem (Main Result 1).** *Consider a path with n vertices, where $n \rightarrow \infty$. Then the following results hold regardless of the loop-probability of the random walk:*

- For $k = 2$, the speed-up satisfies $S_{\text{cov}}^{(k)} > 2$.
- For $k \geq 3$, the speed-up satisfies $S_{\text{cov}}^{(k)} < k$.

The most challenging part of this result is the analysis for $k = 2$. We first prove that k discrete-time random walks explore the graph in nearly the same time as k continuous-time random walks (see Section 3, in particular Theorem 3.6, for a more precise statement). Then we relate the continuous-time random walks to a so-called multi-dimensional Gambler’s ruin problem, that was analyzed by Kmet and Petkovšek in [10]. This Gambler’s ruin problem provides us with a fairly accurate estimate of the cover time for the case where both random walks start from the same endpoint. However, to get a lower-bound on $S_{\text{cov}}^{(k)}$, we have to derive an upper bound on the cover time of k random walks for *any* start vertex. We eventually achieve this by a series of coupling arguments and reductions to the base case.

The sub-linear speed-up for $k \geq 3$ is also derived via the relation to the Gambler’s ruin problem; one difficulty are cases where k is a very large constant. To bootstrap from smaller

44:4 Multiple Random Walks on Paths and Grids

$d = 2 / k \in$	$t_{\text{cov}}^{(k)}$	Speed-up	$d \geq 3 / k \in$	$t_{\text{cov}}^{(k)}$	Speed-up
$[1, \log^2 n]$	$\Theta\left(\frac{n \log^2 n}{k}\right)$	linear	$[1, n^{1-2/d} \log n]$	$\Theta\left(\frac{n \log n}{k}\right)$	linear
$[\log^2 n, n]$	$\Theta\left(\frac{n}{\log \frac{k}{\ln^2 n}}\right)$	logarithmic	$[n^{1-2/d} \log n, n]$	$\Theta\left(n^{2/d} / \log\left(\frac{k}{n^{1-2/d} \log n}\right)\right)$	logarithmic

■ **Figure 3** Overview of our results on $t_{\text{cov}}^{(k)}$ for the two-dimensional grid/torus on the left hand side, and for the d -dimensional grid/torus, $d \geq 3$, on the right hand side.

to larger values of k , we derive a general result that the speed-up of the hitting time between the two endpoints is sub-linear in k (Lemma 4.5).

Finally, the fact that, as $n \rightarrow \infty$, the effect of the loop-probability becomes negligible, rests on the connection to continuous-time random walks; in fact, it holds not only for paths but also for arbitrary graphs (see Section 3.2 for our results on the impact of loop-probabilities).

For two-dimensional grids, we provide the first asymptotic characterization of $S^{(k)}$ that holds for any $1 \leq k \leq n$ (previous work [1, 8] had only determined $S^{(k)}$ up to logarithmic factors.) We find that for $1 \leq k \leq \log^2 n$, the speed-up is linear, while for $\log^2 n \leq k \leq n$, the speed-up is logarithmic. While a weaker dichotomy was established before [8], our new result demonstrates a very sharp transition behavior between these two extreme cases.

► **Theorem (Main Result 2).** *Consider the two-dimensional grid/torus with n vertices as $n \rightarrow \infty$. Then:*

- For any $k \in [1, \log^2 n]$, $t_{\text{cov}}^{(k)} = \Theta\left(\frac{n \log^2 n}{k}\right)$ and thus $S_{\text{cov}}^{(k)} = \Theta(k)$.
- For any $k \in [\log^2 n, n]$, $t_{\text{cov}}^{(k)} = \Theta\left(\frac{n}{\log(k/\log^2 n)}\right)$ and thus $S_{\text{cov}}^{(k)} = \Theta(\log^2 n \cdot \log(k/\log^2 n))$.

The upper bounds on the cover time follow from a second-moment type argument about return times (a similar technique has been used by Cooper et al. [4]). Deriving the lower bounds appears to be more challenging, in particular for small values of k , where we elaborate on our connection to continuous-time random walks and apply a lower-bound technique from Zuckerman [20] to multiple random walks.

For three or higher dimensions, the situation is very similar; the only difference is that the transition point moves further up towards n , and that all cover times are reduced by a $\log n$ -factor, which is in analogy to the behavior of single random walks.

► **Theorem (Main Result 3).** *Consider the d -dimensional grid/torus, with n vertices as $n \rightarrow \infty$, where $d \geq 3$ is any fixed constant. Then:*

- For any $k \in [1, n^{1-2/d} \log n]$, $t_{\text{cov}}^{(k)} = \Theta\left(\frac{n \log n}{k}\right)$ and thus $S_{\text{cov}}^{(k)} = \Theta(k)$.
- For any $k \in [n^{1-2/d} \log n, n]$, $t_{\text{cov}}^{(k)} = \Theta\left(n^{2/d} / \log\left(\frac{k}{n^{1-2/d} \log n}\right)\right)$ and therefore $S_{\text{cov}}^{(k)} = \Theta\left(n^{1-2/d} \cdot \log n \cdot \log\left(\frac{k}{n^{1-2/d} \log n}\right)\right)$.

Figure 3 summarizes our results for two-dimensional grid/torus and d -dimensional grid/torus for $d \geq 3$.

Organisation. In Section 2 we provide the required mathematical notation and basic definitions. Section 3 deals with the connection between continuous and discrete-time multiple random walks. Then in Section 4 we analyze the two cases $k = 2$ and $k \geq 3$ for a path with n vertices. Section 5 contains our results for 2-dimensional grids (the results for $d \geq 3$ -dimensional grids are derived in a similar fashion). We close our paper in Section 6 by pointing to some interesting open questions.

2 Notation and Definitions

Unless mentioned otherwise, we always consider simple random walks on an undirected, simple, connected and unweighted graph $G = (V, E)$, where $n = |V|$ is the number of vertices. More specifically, a random walk starts from an arbitrary prespecified vertex and then at each step $1, 2, \dots$ the random walk moves to a neighbor chosen uniformly at random. This can be encoded by a transition matrix P , where $p_{u,v} = 1/\deg(u)$ if $\{u, v\} \in E(G)$, $p_{u,v} = 0$ otherwise. Further, we denote by $p_{u,v}^t$ the probability for a random walk starting from u to be at vertex v at time-step t . Finally, we denote by π the *stationary distribution* that satisfies $\pi(u) = \frac{\deg(u)}{2|E|}$.

By $\tau^{(k)}((u_1, \dots, u_k), v)$ we denote the first passage time, i.e., the first time one of k independent and concurrent random walks hits v , where the i -th random walk starts from vertex u_i . Then the *hitting time* is the expectation of the first passage time, in symbols

$$t_{\text{hit}}^{(k)}((u_1, \dots, u_k), v) := \mathbf{E} \left[\tau^{(k)}((u_1, \dots, u_k), v) \right].$$

Similarly, $\zeta^{(k)}(u_1, \dots, u_k)$ denotes the first time k independent and concurrent random walks have visited all vertices in V , where again the i -th random walk starts from vertex u_i . Then the *cover time* is the expectation of that random variable, in symbols

$$t_{\text{cov}}^{(k)}(u_1, \dots, u_k) := \mathbf{E} \left[\zeta^{(k)}(u_1, \dots, u_k) \right].$$

For single random walks, we will often drop the ⁽¹⁾-superscript and simply write $t_{\text{cov}}(u)$ and $t_{\text{hit}}(u, v)$. Finally, we define $t_{\text{cov}} := \max_u t_{\text{cov}}(u)$ and $t_{\text{cov}}^{(k)} := \max_u t_{\text{cov}}^{(k)}(u, \dots, u)$. Since in many cases we will work with start configurations where all k random walks start from the same vertex, we may just write \vec{u} instead of (u, \dots, u) .

The *speed-up* for the cover time is the following function of k (and the underlying graph):

$$S_{\text{cov}}^{(k)} := \frac{t_{\text{cov}}}{t_{\text{cov}}^{(k)}} = \frac{\max_u t_{\text{cov}}(u)}{\max_u t_{\text{cov}}^{(k)}(\vec{u})}.$$

This has been the common definition of speed-up introduced by Alon et al. [1] and the focus of subsequent works [7, 8, 19]. However, we will also briefly discuss a slightly different version of $S_{\text{cov}}^{(k)}$ at the end of Section 4. Similarly, the speed-up for the hitting time between u and v is the following function:

$$S_{\text{hit}}^{(k)}(u, v) := \frac{t_{\text{hit}}(u, v)}{t_{\text{hit}}^{(k)}(\vec{u}, v)}.$$

If the random walk is *lazy*, i.e. it stays with some non-zero probability $\alpha > 0$ at the current vertex and otherwise moves to a randomly chosen neighbor, we adjust the notation and write $t_{\text{hit}}^{(\alpha)}(u, v)$, $t_{\text{hit}}^{(k, \alpha)}(\vec{u}, v)$ and $S_{\text{hit}}^{(k, \alpha)}(u, v)$ to reflect the dependency on α .

3 Continuous-Time Multiple Random Walks

3.1 Relating Continuous to Discrete-Time Multiple Random Walks

In the following we will relate the original, k discrete-time random walks to a continuous-time variant. In the continuous-time variant, the waiting time between any two transitions of a random walk are independent and identically distributed exponential random variables. We will denote the corresponding quantities by a superscript $\widetilde{}$, e.g., $\widetilde{t_{\text{cov}}}^{(k)}$ is the cover time of

k concurrent continuous-time random walks. In terms of all conventional and “reasonable” definitions of speed-ups, it turns out that the mean of the waiting time scales all hitting and cover times linearly, and is therefore irrelevant. Therefore the mean 1 is chosen for the sake of convenience unless mentioned otherwise. That means, the waiting time between any two transitions of a random walk are independent exponential random variables with parameter 1 (for further details we refer the reader to [13, Chapter 20.1]).

Furthermore, it is a well-known fact from the theory of Poisson processes that if we focus on one of the k walks, then the times at which a transition happens follows a Poisson point process. In fact something analogous holds if we consider all transitions among k walks, as described in the next lemma:

► **Lemma 3.1.** *Let G be any graph. Then continuous-time multiple random walks, with λ as the waiting-time parameter for each of the k walks, are a Poisson point process in which in each transition exactly one of the k walks performs a transition, and the waiting-times between two successive transitions are exponentially distributed random variables with parameter $k\lambda$.*

We continue to establish that for single random walks all hitting times and cover times are identical for discrete-time and continuous-time random walks.

► **Lemma 3.2.** *For a single random walk on an arbitrary graph G , we have the following equalities: (i) for any pair of vertices $u, v \in V$, $t_{\text{hit}}(u, v) = \widetilde{t}_{\text{hit}}(u, v)$ and (ii) for any node $u \in V$, $t_{\text{cov}}(u) = \widetilde{t}_{\text{cov}}(u)$.*

The proof of Lemma 3.2 makes use of a natural coupling between two types of random walks and Wald’s equation ([13, Exercise 6.6]).

Before we connect the cover times and hitting times of continuous to their discrete-time counterparts, we recall the following useful lower bound on the cover time of k random walks (recall our notation $\vec{u} = (u, \dots, u)$):

► **Lemma 3.3** ([8, Theorem 4.2]¹). *Let G be any graph with n vertices and u be an arbitrary start vertex of k independent discrete-time random walks, where $n^\varepsilon \leq k \leq n$ for some arbitrary $\varepsilon > 0$. Then*

$$\Pr \left[t_{\text{cov}}^{(k)}(\vec{u}) \geq \frac{\varepsilon}{8} \cdot \frac{n}{k} \cdot \log n \right] \geq 1 - \exp \left(-n^{-\varepsilon/8} \right).$$

In the following lemma we show that the cover times of the continuous-time random walk are not much larger than those of a discrete-time random walk.

► **Lemma 3.4.** *For any graph G and $1 \leq k \leq n$,*

$$\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \leq 42 \cdot t_{\text{cov}}^{(k)}(\vec{u}).$$

Furthermore, if the graph G and k satisfy $\Pr \left[t_{\text{cov}}^{(k)}(\vec{u}) \geq f(n) \right] \geq 1 - o(1)$ for some function $f(n) = \omega(\log n)$, then the above inequality can be strengthened to:

$$\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \leq (1 + o(1)) \cdot t_{\text{cov}}^{(k)}(\vec{u}).$$

¹ While the result stated in [8] involves the worst-case start vertices, the proof does not make use of this fact which is why we immediately obtain this slightly more general statement.

Proof. Fix any integer $t \geq \frac{1}{8} \log n$. For any of the k continuous-time random walks, let \mathcal{E}_t^i be the event that occurs if the time for the i -th random walk to perform the first t transition is at most $t + 20t^{2/3}(\log n)^{1/3}$. To this end, note that the number of transitions within time $t + 20t^{2/3}(\log n)^{1/3}$ has a Poisson distribution with mean $\mu = t + 20t^{2/3}(\log n)^{1/3}$. Hence by a Chernoff bound for Poisson random variables ([2, Theorem A.1.15])

$$\begin{aligned} \Pr[\overline{\mathcal{E}}_t^i] &= \Pr[P < t] \leq \Pr\left[P < \left(1 - \frac{20t^{2/3}(\log n)^{1/3}}{\mu}\right) \cdot \mu\right] \\ &\leq \exp\left(-\frac{400t^{4/3}(\log n)^{2/3}}{2\mu}\right) \leq \exp\left(-4t^{1/3}(\log n)^{2/3}\right), \end{aligned}$$

where in the last inequality we used the fact that the assumption $t \geq \frac{1}{8} \log n$ implies $\mu \leq 41t$. Let $\mathcal{E}_t := \bigcup_{i=1}^k \mathcal{E}_t^i$. By a Union bound over all $k \leq n$ random walks,

$$\Pr[\mathcal{E}_t] \leq n \cdot \exp\left(-4t^{1/3}(\log n)^{2/3}\right) = n \cdot n^{-4 \cdot (t/\log n)^{1/3}} \leq n^{-1}.$$

To relate $\widetilde{t}_{\text{cov}}^{(k)}(\vec{u})$ (cover time for k continuous-time walks) to $t_{\text{cov}}^{(k)}(\vec{u})$ (cover time for k discrete-time walks), we consider again the straightforward coupling, where the random walk i (in discrete and continuous case) performs the same sequence of transitions. In other words, the two trajectories in discrete-time and continuous-time are identical, and only the times at which the transitions happen will be different. With this, we can now begin to upper bound $\widetilde{t}_{\text{cov}}^{(k)}(\vec{u})$ by decomposing the expectation:

$$\begin{aligned} \widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) &= \sum_{t=0}^{\infty} \Pr[\zeta^{(k)}(\vec{u}) = t] \cdot \mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t\right] \\ &\leq \Pr\left[\zeta^{(k)}(\vec{u}) \leq \frac{1}{8} \cdot \log n\right] \cdot \mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) \leq \frac{1}{8} \cdot \log n\right] \\ &\quad + \sum_{t=\frac{1}{8} \cdot \log n}^{\infty} \Pr[\zeta^{(k)}(\vec{u}) = t] \cdot \mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t\right]. \end{aligned}$$

To bound $\Pr[\zeta^{(k)}(\vec{u}) \leq \frac{1}{8} \cdot \log n]$, we can apply the bound from Theorem 3.3 for $k = n$ (since $\zeta^{(k)}(\vec{u})$ is stochastically dominated by $\zeta^{(n)}(\vec{u})$), and combine it with the simple bound $\mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) \leq \frac{1}{8} \cdot \log n\right] \leq \mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u})\right]$ to obtain

$$\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \leq e^{-n^{1/8}} \mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u})\right] + \sum_{t=\frac{1}{8} \cdot \log n}^{\infty} \Pr[\zeta^{(k)}(\vec{u}) = t] \mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t\right].$$

In the following, we shall derive an upper bound on $\mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t\right]$. First, by conditioning on the event \mathcal{E}_t we obtain

$$\begin{aligned} &\mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t\right] \\ &= \mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t, \mathcal{E}_t\right] \Pr[\mathcal{E}_t] + \mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t, \overline{\mathcal{E}}_t\right] \Pr[\overline{\mathcal{E}}_t] \quad (1) \end{aligned}$$

We shall now proceed by bounding the two conditional expectations in eq. (1). First, we have

$$\mathbf{E}\left[\widetilde{t}_{\text{cov}}^{(k)}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t, \mathcal{E}_t\right] \leq (t + 20t^{2/3}(\log n)^{1/3}),$$

since conditional on \mathcal{E}_t , all k random walks perform at least t transitions within the time-interval $[0, t + 20t^{2/3}(\log n)^{1/3}]$.

To upper bound $\mathbf{E} \left[\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \mid \zeta^{(k)}(u) = t, \overline{\mathcal{E}_t} \right]$ note that $\overline{\mathcal{E}_t}$ only affects the random walks within the time interval $[0, t + 20t^{2/3}(\log n)^{1/3}]$. Therefore we have,

$$\mathbf{E} \left[\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \mid \zeta^{(k)}(u) = t, \overline{\mathcal{E}_t} \right] \leq \left(t + 20t^{2/3}(\log n)^{1/3} + \mathbf{E} \left[\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \mid \zeta^{(k)}(u) = t \right] \right).$$

Combining these two bounds and plugging these into equation (1) yields

$$\begin{aligned} \mathbf{E} \left[\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \mid \zeta^{(k)}(u) = t \right] &\leq (t + 20t^{2/3}(\log n)^{1/3}) \cdot (1 - \mathbf{Pr}[\overline{\mathcal{E}_t}]) + \\ &\quad \left(t + 20t^{2/3}(\log n)^{1/3} + \mathbf{E} \left[\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t \right] \right) \cdot \mathbf{Pr}[\overline{\mathcal{E}_t}] \\ &= t + 20t^{2/3}(\log n)^{1/3} + \mathbf{E} \left[\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t \right] \cdot \mathbf{Pr}[\overline{\mathcal{E}_t}]. \end{aligned}$$

Using $\mathbf{Pr}[\overline{\mathcal{E}_t}] \leq n^{-1}$ and rearranging gives

$$\mathbf{E} \left[\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \mid \zeta^{(k)}(\vec{u}) = t \right] \leq \frac{1}{1 - n^{-1}} \cdot \left(t + 20t^{2/3}(\log n)^{1/3} \right).$$

Using this bound and returning to the upper bound on $\widetilde{t_{\text{cov}}^{(k)}}(\vec{u})$ yields

$$\begin{aligned} \widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \cdot (1 - e^{-n^{1/8}}) &\leq \sum_{t=\frac{1}{8} \cdot \log n}^{\infty} \mathbf{Pr} \left[\zeta^{(k)}(\vec{u}) = t \right] \cdot \frac{1}{1 - n^{-1}} \cdot \left(t + 20t^{2/3}(\log n)^{1/3} \right) \quad (2) \\ &\leq \sum_{t=1}^{\infty} \mathbf{Pr} \left[\zeta^{(k)}(\vec{u}) = t \right] \cdot \frac{41t}{1 - n^{-1}}, \end{aligned}$$

where the final step holds because $t + 20t^{2/3}(\log n)^{1/3} \leq 41t$. Hence, $\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \leq 42 \cdot t_{\text{cov}}^{(k)}(\vec{u})$.

To prove the tighter inequality, we proceed similarly. Details can be found in the full version. \blacktriangleleft

The next lemma is the corresponding lower bound variant of Lemma 3.4:

► **Lemma 3.5.** *For any graph G and $1 \leq k \leq n$, there is a constant $c > 0$ (independent of k and n) so that*

$$\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \geq c \cdot t_{\text{cov}}^{(k)}(\vec{u}).$$

Furthermore, if the graph G and k satisfy $\mathbf{Pr} \left[t_{\text{cov}}^{(k)}(\vec{u}) \geq f(n) \right] \geq 1 - o(1)$ for some function $f(n) = \omega(\log n)$, then the above inequality can be strengthened to:

$$\widetilde{t_{\text{cov}}^{(k)}}(\vec{u}) \geq (1 - o(1)) \cdot t_{\text{cov}}^{(k)}(\vec{u}).$$

The proof of this Lemma is similar to that of Lemma 3.4, and is therefore given in the appendix only. It might be worth mentioning that it is not possible to drop the upper bound on k in the precondition of Lemma 3.5; clearly, as $k \rightarrow \infty$, $t_{\text{cov}}^{(k)}(u) \geq \text{diam}$, where diam is the diameter of the underlying graph, however, as $k \rightarrow \infty$, $\widetilde{t_{\text{cov}}^{(k)}}(u) \rightarrow 0$.

Combining the three previous lemmas above, we conclude the following.

► **Theorem 3.6.** *Let G be any graph. For any speed-up of the cover time with $1 \leq k \leq o(n)$, $S_{\text{cov}}^{(k)} = (1 \pm o(1)) \cdot \widetilde{S_{\text{cov}}^{(k)}}$. Moreover, for any $1 \leq k \leq n$, $S_{\text{cov}}^{(k)} = \Theta(\widetilde{S_{\text{cov}}^{(k)}})$.*

Analogous results also hold for hitting times, but in this case one would need a lower bound like $\mathbf{Pr} \left[\tau(u, v) \geq \frac{1}{8} \log n \right] \geq 1 - o(1)$, which may not be possible for all pairs of vertices.

3.2 The Effect of Loop Probabilities

Our first result demonstrates that the influence of the loop probabilities on the speed-up diminishes, as $n \rightarrow \infty$.

► **Theorem 3.7.** *Let G be any graph and $0 \leq \alpha_1 < \alpha_2 \leq 1$ be any pair of loop probabilities. Then for any pair of vertices u, v , $S_{\text{cov}}^{(k, \alpha_1)}(u) = \Theta(S_{\text{cov}}^{(k, \alpha_2)}(u))$. Further, if $\Pr \left[t_{\text{cov}}^{(k, 0)}(\vec{u}) \geq f(n) \right] \geq 1 - o(1)$ for some $f(n) = \omega(\log n)$, then $S_{\text{cov}}^{(k, \alpha_1)}(u) = (1 + o(1)) \cdot S_{\text{cov}}^{(k, \alpha_2)}(u)$.*

The proof of Theorem 3.7 is based on the following two steps. Firstly, we use the connection between continuous-time and (non-lazy) discrete-time random walks derived above. Secondly, we make use of the fact that in continuous-time, increasing the loop-probabilities from 0 to some value is just a scaling operation, which leaves the speed-up unchanged.

For hitting times, we can prove an even stronger result stating that the speed-up of hitting times between any pair of vertices is non-decreasing in the loop-probability.

► **Lemma 3.8.** *Let G be any graph. Then for any combination of n and k , the speed-up $S_{\text{hit}}^{(k, \alpha)}(u, v)$ of any hitting time pair for random walks with loop probability $\alpha > 0$ is at least as large as the speed-up $S_{\text{hit}}^{(k, 0)}(u, v)$ in the non-lazy setting.*

Similar to the proof of Lemma 3.2, the proof of Lemma 3.8 involves a natural coupling between lazy and non-lazy random walks and an application of Wald’s equation.

We now establish a useful duality between ‘very lazy’ random walks and continuous-time random walks.

► **Lemma 3.9.** *Let G be any graph and u, v be any two vertices. Then $\lim_{\alpha \rightarrow 1} S_{\text{hit}}^{(k, \alpha)}(u, v) = \widetilde{S}_{\text{hit}}^{(k)}(u, v)$ and $\lim_{\alpha \rightarrow 1} S_{\text{cov}}^{(k, \alpha)} = \widetilde{S}_{\text{cov}}^{(k)}$.*

The proof of this lemma revolves around the idea that if the loop-probability is sufficiently large, then with high probability in each step among the k discrete-time multiple random walks at most one will move to a neighboring vertex. Hence there are no “concurrency effects”, and the discrete-time and continuous-time random walks behave correspondingly.

4 Speed-up on the Path

For convenience, we shall consider a path with $n + 1$ nodes labelled from 0 to n , and n is even. We defer the straightforward but tedious adjustments for odd n to the full version.

Before delving into further specifics of the analyses for $k = 2$ and $k = 3$, let us briefly mention the following fact for the single random walks ($k = 1$). To determine the cover time from the vertex, say, $n/2$, we can use the Gambler’s ruin problem to infer that the time until one of the vertices 0 or n is reached equals $(n/2 - 0) \cdot (n - n/2) = n^2/4$. Then, the time until the opposite endpoint is reached equals $t_{\text{hit}}(0, n) = t_{\text{hit}}(n, 0) = n^2$ (cf. [15]). Combining this, we arrive at the following simple yet crucial observation:

► **Observation 4.1.** *For a single random walk ($k = 1$) on a path with endpoints 0 and n , $t_{\text{cov}}(n/2) = \frac{5}{4} \cdot n^2$. Furthermore, $t_{\text{cov}}(u)$ is maximized for $u = n/2$.*

Recalling the definition of $S_{\text{cov}}^{(k)}$, the strategy to establish a sub-linear speed-up for $k \geq 3$ is to find a vertex u so that $t_{\text{cov}}^{(k)}(\vec{u}) > \frac{1}{k} \cdot \frac{5}{4} \cdot n^2$. A both natural and convenient candidate is to pick the endpoint, and we shall indeed establish in the next subsection that $t_{\text{cov}}^{(k)}(\vec{0}) = t_{\text{hit}}^{(k)}(\vec{0}, n) > \frac{1}{k} \cdot \frac{5}{4} \cdot n^2$.

k	1	2	3	4	5	6	7	8	9
c_k	1	1.178	1.349	1.512	1.670	1.823	1.973	2.118	2.261

■ **Figure 4** The numerical values for c_k as given in [10] rounded down to the third significant digit.

In contrast to this, establishing a super-linear speed-up for $k = 2$ turns out to be more demanding because in this case we have to upper bound $t_{\text{cov}}^{(k)}(\vec{u})$ for *any* vertex u .

4.1 $k \geq 3$: Sub-linear Speed-up

In this section we shall prove that the speed-up for cover times on paths is sub-linear as soon as $k \geq 3$. We should mention that for sufficiently large values of k , i.e., k being super-constant, [1] proved that on the cycle with n vertices that $S^k = \Theta(\log k)$, as $n \rightarrow \infty$. Unfortunately, it seems difficult to apply the result (or proof) of [1] to the case of a path with a small constant k . Here we will show the following theorem, that captures all values of $k \geq 3$:

► **Theorem 4.2.** *Consider a discrete-time random walk on a path with endpoints 0 and n . Then for any $k \geq 3$, as $n \rightarrow \infty$, $t_{\text{cov}}^{(k)}(\vec{0}) = t_{\text{hit}}^{(k)}(\vec{0}, n) > \frac{5/4}{k} \cdot n^2 - o(n^2)$.*

Since $t_{\text{cov}}(n/2) = \frac{5}{4} \cdot n^2$ by Observation 4.1, the theorem above immediately implies the following result:

► **Theorem 4.3.** *For any path with endpoints 0 and n and $k \geq 3$, as $n \rightarrow \infty$, $S_{\text{cov}}^{(k)} < k$.*

Before proving Theorem 4.2, we list three auxiliary results. The theorem below follows from a work by Kmet and Petkovsek [10] on the k -dimensional Gambler’s ruin problem.

► **Theorem 4.4** (cf. [10]). *Consider a continuous-time random walk on a path with endpoints 0, n . Then for any $1 \leq k \leq 10$,*

$$\widetilde{t}_{\text{cov}}^{(k)}(\vec{0}) = \widetilde{t}_{\text{hit}}^{(k)}(\vec{0}, n) > \frac{c_k}{k} \cdot n^2 - o(n^2/k),$$

where the c_k ’s are the numerical values given in the table below.

The following lemma establishes that the speed-up is at most sub-linear, however, this is for the hitting time between two endpoints.

► **Lemma 4.5.** *For the path with endpoints 0 and n , we have for any pair of integers $\ell \geq 1$ and $k \geq 1$, $t_{\text{hit}}^{(\ell \cdot k)}(\vec{0}, n) \geq \frac{1}{\ell} \cdot t_{\text{hit}}^{(k)}(\vec{0}, n)$.*

Proof. Let $X = \tau^{(k)}(\vec{0}, n)$ be the random variable describing the first visit to n of a random walk starting from 1; so $\mathbf{E}[X] = t_{\text{hit}}^{(k)}(\vec{0}, n)$. For the $\ell \cdot k$ random walks, let $X_1, X_2, \dots, X_{\ell \cdot k}$ be $\ell \cdot k$ independent copies of X . We have:

$$\begin{aligned} t_{\text{hit}}^{(\ell \cdot k)}(\vec{0}, n) &= \sum_{i=1}^{\infty} \Pr[\min\{X_1, \dots, X_k\} = i] + \sum_{i=1}^{\infty} \Pr[\min\{X_1, \dots, X_k\} > i] \\ &= 1 + \sum_{i=1}^{\infty} \Pr[X > i]^\ell \end{aligned}$$

Note that $\Pr[X > \ell \cdot i]$ can be upper bounded by $\Pr[X > i]^\ell$. This follows by dividing the time into consecutive blocks of length i each, and noting that the probability that a fixed random walk does not hit n in any of the blocks is upper bounded by $\Pr[X > i]$. Therefore,

$$\begin{aligned} t_{\text{hit}}^{(\ell,k)}(\vec{0}, n) &\geq 1 + \sum_{i=1}^{\infty} \Pr[X > \ell \cdot i] = 1 - \sum_{i=1}^{\infty} \Pr[X = \ell \cdot i] + \sum_{i=1}^{\infty} \Pr[X \geq \ell \cdot i] \\ &\geq \sum_{i=1}^{\infty} \Pr[X \geq \ell \cdot i] = \mathbf{E}[X/\ell] = \frac{1}{\ell} \cdot t_{\text{hit}}^{(k)}(\vec{0}, n). \quad \blacktriangleleft \end{aligned}$$

Finally, we shall need the following simple lower bound on $t_{\text{hit}}^{(k)}(\vec{0}, n)$, that will enable us to apply Lemma 3.4:

► **Lemma 4.6.** *For any $1 \leq k \leq n$, $\Pr\left[t_{\text{hit}}^{(k)}(\vec{0}, n) \geq n^2/\log^2 n\right] \geq 1 - o(1)$.*

The basic outline of the proof of Theorem 4.2 is now as follows. We use the lower bound from Theorem 4.4 for smaller values of k as a base case. Then we use Lemma 3.4 to derive corresponding lower bounds for discrete-time random walks. Finally, we use Lemma 4.5 to “bootstrap” the lower bounds from smaller values of k to larger values.

4.2 $k = 2$: Super-linear Speed-up

First let us recall that c_2 is the constant from the two-dimensional Gambler’s ruin problem and was shown to be at most 1.179 [10] (see also Theorem 4.4 and Figure 4 in this manuscript).

We first derive several upper bounds on the cover time for four special cases. The proofs of these three estimates are non-trivial but similar. In essence, they use couplings in order to reduce the cover time to the “base case” where the two random walks start from the same endpoint, for which we have a fairly accurate estimate (Theorem 4.4).

► **Lemma 4.7.** *Let G be a path with endpoints 0 and n . Then, as $n \rightarrow \infty$:*

1. $t_{\text{cov}}^{(2)}(n/2, n/2) \leq \left(\frac{3}{8} + \frac{3c_2}{16}\right) n^2 + o(n^2) \leq 0.59602 \cdot n^2 + o(n^2)$.
2. $t_{\text{cov}}^{(2)}(n/4, n/4) \leq 0.6128 \cdot n^2 + o(n^2)$.
3. For any $n/4 \leq w \leq n/2$,

$$t_{\text{cov}}^{(2)}(w, n/2) \leq \frac{1}{2} \cdot \left(\left(2 - \frac{c_2}{2}\right) \cdot \frac{n^2}{4} \right) + \frac{1}{2} \left(\frac{1}{2} \cdot w \cdot (n - w) + \frac{1}{2} \cdot \frac{n}{2} \cdot \frac{n}{2} + \frac{c_2}{2} \cdot n^2 \right) + o(n^2),$$

where the cover time is under the implicit assumption that all vertices between w and $n/2$ are already visited.

4. For any $0 \leq w < n/4$, $t_{\text{cov}}^{(2)}(0, w) \leq \frac{c_2}{2} \cdot n^2 - \frac{c_2}{4} \cdot w^2 + o(n^2)$, where the cover time is under the implicit assumption that all vertices between 0 and w are already visited.

Equipped with the three auxiliary estimates above, we now analyze $t_{\text{cov}}^{(2)}(w, w)$.

► **Proposition 4.8.** *Let G be a path with endpoints 0 and n . Then, as $n \rightarrow \infty$:*

1. For any $n/4 \leq w \leq n/2$, $t_{\text{cov}}^{(2)}(w, w) \leq 0.6244 \cdot n^2 + o(n^2) < \frac{5}{8} \cdot n^2$.
2. For any $0 \leq w < n/4$, $t_{\text{cov}}^{(2)}(w, w) \leq 0.6128 \cdot n^2 + o(n^2) < \frac{5}{8} \cdot n^2$.

With the estimates from Lemma 4.7 at hand, the proof of the first statement is essentially a repeated application of the corresponding bounds in Lemma 4.7, depending which of the two endpoints 0 or $n/4$ the first and second random walk reach first (cf. appendix).

Finally, by using the two results in the above propositions and noting the simple symmetry that $t_{\text{cov}}^{(2)}(w, w) = t_{\text{cov}}^{(2)}(n - w, n - w)$ for any $0 \leq w \leq n/2$, we obtain:

► **Theorem 4.9.** *For the path with vertices $0, \dots, n$, as $n \rightarrow \infty$, then for any $1 \leq w \leq n$, $t_{\text{cov}}^{(2)}(w, w) < \frac{5}{8} \cdot n^2$. Consequently, as $n \rightarrow \infty$, $S_{\text{cov}}^{(k)} > 2$.*

Let us briefly mention that even if we were to redefine the speed-up as the quantity $S_{\text{cov}}^{(k)} := \max_{u \in V} \frac{t_{\text{cov}}(u)}{t_{\text{cov}}^{(k)}(\bar{u})}$, our results above for $u = n/2$ would imply a super-linear speed-up.

5 Speed-up on d-Dimensional Grid/Torus

A d -dimensional grid is the graph with vertex set $V = [-n^{1/d}/2, +n^{1/d}/2]^d$, and two vertices are connected iff they differ in one coordinate by 1 and all other coordinates are identical. A d -dimensional torus is almost same as the d -dimensional grid, but additionally we have “wrap-around” edges, e.g., for $d = 2$, we have additional edges between $(x, n/2)$ and $(x, -n/2)$, and $(n/2, x)$ and $(-n/2, x)$ for any $-n/2 \leq x \leq n/2$.

Due to space limitations, we only sketch the analysis for $d = 2$ here. The analysis for $d \geq 3$ is quite similar and even slightly easier than the one for $d = 2$. To avoid any periodicity issues, we shall always consider lazy random walks with loop probability $1/2$.

5.1 Preliminaries

► **Lemma 5.1.** *For any graph G and any integer t , $\Pr[\tau(u, v) \leq t] \geq \frac{\sum_{s=0}^t p_{uv}^s}{1 + \mathbf{E}[R_u(t)]}$, where $\mathbf{E}[R_u(t)] = \sum_{i=1}^t p_{u,u}^i$ is the expected number of returns to u until step t .*

We next note an elementary fact about the limiting behaviour of the t -step transition probabilities. It can be easily derived by relating the random walk to the infinite grid \mathbb{Z}^d and then applying the central limit theorem (e.g. [12]).

► **Lemma 5.2.** *Consider a d -dimensional grid or torus, where d is constant, and let $n \rightarrow \infty$ be sufficiently large. Then for any pair of vertices $u, v \in V$, there is a constant $c_1 > 0$ so that $p_{u,v}^t \leq \pi(v) + c_1 \cdot t^{-d/2}$. Moreover, for any pair of vertices $u, v \in V$, there is a constant $c > 0$ so that*

$$p_{u,v}^t \geq c \cdot t^{-d/2} \cdot \exp\left(-\frac{d \cdot \|u - v\|_2^2}{t}\right) - c \cdot t^{-(d+2)/2}.$$

5.2 $d = 2$ and $k \in [1, \log^2 n]$: Linear Speed-Up

We begin this part by deriving the upper bound on the cover time. First, we point out that for $1 \leq k \leq \log n$, the desired bound $t_{\text{cov}}^{(k)} = O(n \log^2 n/k)$ follows from [1, Theorem 4]. To cover also the missing regime $\log n \leq k \leq \log^2 n$, we show the following lemma:

► **Lemma 5.3.** *For $k \in [1, \log^2 n]$ the cover time of k independent lazy random walks on a 2-dimensional torus is $O(n \log^2 n/k)$.*

This result is a fairly straightforward application of Lemma 5.1. The derivation of the lower bound is more involved, and is based on the following result:

► **Lemma 5.4** ([20, Lemma 2]). *Let $V' \subseteq V$ such that $|V'| \geq n^\delta, \delta > 0$ and let t such that for $i \in V'$, at most $1/n^\beta$ fraction of the $j \in V'$ satisfy $t_{\text{hit}}(u, v) < t$. Then for any start vertex $v \in V$, $t_{\text{cov}}(v) \geq t \cdot (\gamma \ln n - 2)$, where $\gamma = \min(\delta, \beta)$.*

5.3 $d = 2$ and $k \in [\log^2 n, n]$: Logarithmic Speed-Up

► **Lemma 5.5.** For $k \in [\log^2 n, n]$ the cover time of k independent lazy random walks on a 2-dimensional torus is $\Theta(n/\log(k/\log^2 n))$.

6 Summary and Open Problems

While a focus of this work has been the speed-up of multiple random walks on paths and grids, an important general insight is the benefit of using continuous-time random walks. In that regard, Lemma 3.4, Lemma 3.5 and Theorem 3.6 (all holding for *arbitrary* graphs) might be useful for future work, since one can conveniently switch between continuous-time and discrete-time walks.

One interesting open problem is to analyze the speed-up on the cycle, where for $k = 2$, simulations indicate that the speed-up may be very close to 2. Secondly, in light of our result that already for the cover time on the path the speed-up can be super-linear, a logical next step would be to explore whether $S_{\text{hit}}^{(k)}(u, v) \leq k$ holds (for the special case of the path and u and v being the endpoints, Lemma 4.5 provides a positive answer).

References

- 1 Noga Alon, Chen Avin, Michal Koucký, Gady Kozma, Zvi Lotker, and Mark R. Tuttle. Many random walks are faster than one. *Combinatorics, Probability & Computing*, 20(4):481–502, 2011.
- 2 Noga Alon and Joel Spencer. *The Probabilistic Method*. John Wiley & Sons, 2nd edition, 2000.
- 3 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.
- 4 Colin Cooper, Alan M. Frieze, and Tomasz Radzik. Multiple random walks and interacting particle systems. In *Proc. 36th Intl. Colloquium on Automata, Languages and Programming (ICALP'09)*, pages 399–410, 2009.
- 5 Colin Cooper, David Ilcinkas, Ralf Klasing, and Adrian Kosowski. Derandomizing random walks in undirected graphs using locally fair exploration strategies. *Distributed Computing*, 24(2):91–99, 2011.
- 6 Artur Czumaj and Christian Sohler. Testing expansion in bounded-degree graphs. *Combinatorics, Probability & Computing*, 19(5-6):693–709, 2010.
- 7 Klim Efremenko and Omer Reingold. How well do random walks parallelize? In *13th International Workshop on on Randomization and Computation (RANDOM'09)*, pages 476–489, 2009.
- 8 Robert Elsässer and Thomas Sauerwald. Tight bounds for the cover time of multiple random walks. *Theor. Comput. Sci.*, 412(24):2623–2641, 2011.
- 9 Amos Israeli and Marc Jalfon. Token management schemes and random walks yield self-stabilizing mutual exclusion. In *Proceedings of the 9th Annual ACM Symposium on Principles of Distributed Computing (PODC'90)*, pages 119–131, 1990.
- 10 Andrej Kmet and Marko Petkovšek. Gambler's ruin problem in several dimensions. *Advances in Applied Mathematics*, 28:107–118, 2002.
- 11 Adrian Kosowski. *Time and Space-Efficient Algorithms for Mobile Agents in an Anonymous Network*. Research habilitation, Université Bordeaux 1, September 2013.
- 12 Gregory F. Lawler and Vlada Limic. *Random Walk: A Modern Introduction*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2010.

- 13 David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. AMS, 2008.
- 14 Pascal Lezaud. Chernoff-type bound for finite markov chains. *Ann. Appl. Probab.*, 8(3):849–867, 1989.
- 15 László Lovász. Random walks on graphs: A survey. *Combinatorics, Paul Erdős is Eighty*, 2:1–46, 1993.
- 16 George B. Mertzios, Sotiris E. Nikolettseas, Christoforos Raptopoulos, and Paul G. Spirakis. Determining majority in networks with local interactions and very small local memory. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP'14)*, pages 871–882, 2014.
- 17 Atish Das Sarma, Anisur Rahaman Molla, and Gopal Pandurangan. Fast distributed computation in dynamic networks via random walks. In *Proceedings of the 26th International Symposium on Distributed Computing (DISC'12)*, pages 136–150, 2012.
- 18 Atish Das Sarma, Danupon Nanongkai, Gopal Pandurangan, and Prasad Tetali. Distributed random walks. *J. ACM*, 60(1):2, 2013.
- 19 Thomas Sauerwald. Expansion and the cover time of parallel random walks. In *Proceedings of the 29th Annual ACM Symposium on Principles of Distributed Computing (PODC'10)*, pages 315–324, 2010.
- 20 David Zuckerman. A technique for lower bounding the cover time. *SIAM J. Discrete Math.*, 5(1):81–87, 1992.