

ARTICLE

The power of two choices for random walks

Agelos Georgakopoulos¹, John Haslegrave¹, Thomas Sauerwald^{2*} and John Sylvester²

¹Mathematics Institute, University of Warwick, Coventry, UK and ²Department of Computer Science & Technology, University of Cambridge, Cambridge, UK

*Corresponding author. Email: thomas.sauerwald@cl.cam.ac.uk

(Received 2 September 2020; revised 31 March 2021; accepted 27 April 2021)

Abstract

We apply the power-of-two-choices paradigm to a random walk on a graph: rather than moving to a uniform random neighbour at each step, a controller is allowed to choose from two independent uniform random neighbours. We prove that this allows the controller to significantly accelerate the hitting and cover times in several natural graph classes. In particular, we show that the cover time becomes linear in the number n of vertices on discrete tori and bounded degree trees, of order $\mathcal{O}(n \log \log n)$ on bounded degree expanders, and of order $\mathcal{O}(n(\log \log n)^2)$ on the Erdős–Rényi random graph in a certain sparsely connected regime. We also consider the algorithmic question of computing an optimal strategy and prove a dichotomy in efficiency between computing strategies for hitting and cover times.

2020 MSC Codes: Primary 05C81, 60J10; Secondary 68R10, 68Q17

1. Introduction

The power of choice paradigm asserts that when a random process is offered a choice between two or more uniformly selected options, as opposed to being supplied with just one, then a series of choices can be made to improve the overall performance. This idea was first applied to the ‘balls into bins’ model [5, 9, 31], where it was proved that the power of choice decreases the maximum load from $\Theta\left(\frac{\log n}{\log \log n}\right)$ to $\Theta(\log \log n)$ when assigning n balls to n bins. The power of choice was later extensively studied for random graphs under the broader class of rule-based random graph processes, known as Achlioptas processes, see for example [1, 10, 11, 33, 34] and references therein. The power of choice has also been studied with regard to the Preferential Attachment process for growing a random connected graph; in this context, the choices may have a powerful effect on the degree distribution, see e.g. [24, 30].

In this paper, we extend the power-of-two-choices paradigm to random walk on a graph. We show that for many natural classes of graphs, this results in a significant speed-up of the cover and hitting times, which are the expected times to visit all vertices or any fixed vertex from a worst-case start vertex. We study the *choice random walk (CRW)*, which at every step is offered two uniformly random independently sampled neighbours (with repetition) of the current location and (with full

A preliminary version of this paper appeared at *The 11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of LIPIcs, pages 76:1–76:19 [22]

knowledge of the graph) must choose one as the next step; see Section 2 for more details. We prove that the cover time of CRW decreases to $\Theta(n)$ for grids (i.e. finite quotients of \mathbb{Z}^d) and bounded degree trees on n vertices, and that the cover time of expander graphs decreases to $\mathcal{O}(n \log \log n)$. We note that for simple random walk (SRW) these cover times are all $\Omega(n \log n)$ and some are $\Omega(n^2)$ [2]. We also consider computational questions relating to choosing a good strategy: we show that an optimal strategy for minimising a hitting time can be computed in polynomial time, but choosing an optimal strategy for minimising the cover time is NP-hard. See Section 1.2 for more details and other results.

Part of our motivation is to improve the efficiency of random walks used in algorithmic applications such as searching, routing, self-stabilization and query processing in wireless networks, peer-to-peer networks and other distributed systems. One practical setting where routing using the power of choice walk may be advantageous is in relatively slowly evolving dynamic networks such as the internet. For example, say a packet has a target destination v and each node stores a pointer to a neighbour which it believes leads most directly to v . If this network is perturbed, then the deterministic scheme may get stuck in ‘dead ends’, whereas a random walk would avoid this fate. The CRW which prefers edges pointed to by a node may be the best of both worlds as it would also avoid traps but may see a speed-up over the SRW when the original paths are still largely intact.

1.1 Related literature

To the best of our knowledge, Avin and Krishnamachari [3] were the first to apply the principle of the power of choice to random walks. However, their version only considers a simple choice rule where the vertex with fewer previous visits is always preferred, and ties are broken randomly. This is in the spirit of balanced allocations, the origin of the power-of-two-choices paradigm. Their results are mainly empirical and suggest a decrease in the variance of the cover time, and a significant improvement in visit load balancing. This is related to the greedy random walk of Orenshtein and Shinkar [32], which chooses uniformly from adjacent vertices that have not yet been visited (if possible). This model is well studied for expanders [8, 17]. The power of choice has also been studied in the context of deterministic random walks and the rotor-router model [7, 18].

Perhaps closest to our work, Azar, Broder, Karlin, Linial and Phillips [4] introduced the ε -biased random walk (ε -BRW) where at each step with probability $\varepsilon > 0$ a controller can choose a neighbour of the current vertex to move to, otherwise a uniformly random one is selected. The model is quite similar to ours in the sense that the controller has full knowledge of the graph when choosing a neighbour. They obtained bounds on the stationary probabilities and show that optimal strategies for maximising or minimising stationary probabilities or hitting times can be computed in polynomial time. There is some overlap with our results in Section 7, where in particular Theorem 7.4 uses a clever substitution from [4] to express an optimisation problem as a linear program. One major difference is that Azar, Broder, Karlin, Linial and Phillips restrict their study to time-independent strategies and do not investigate cover times. Three of the authors of this paper have recently extended [4] to the time-dependent setting and studied cover times for ε -BRWs [25]. The conference paper [22] collects some of our results on the CRW from here and on the ε -BRW from [25] giving a comparison between the two processes.

Azar, Broder, Karlin, Linial and Phillips [4] suggest that the most natural choice of bias for the ε -BRW is $\varepsilon = \Theta(1/d_{\max})$, where d_{\max} is the maximum degree. It is shown in [22, Prop. 1] that the CRW can emulate the ε -BRW provided $\varepsilon \leq 1/d_{\max}$. However, the reverse does not hold unless the bias ε is close to 1; the main obstacle is that avoiding a particular next step is much more difficult for the ε -BRW. Further evidence that the CRW is more powerful than the ε -BRW is in the cover time bounds we prove for the CRW in Theorem 6.1 and for the time-dependent version of the ε -BRW in [25, Thm. 3.2]. For the most natural choice $\varepsilon = \Theta(1/d_{\max})$, these bounds differ

by a factor which is almost linear in d_{\max} , suggesting that the CRW deals better with high-degree graphs than the ε -BRW.

With regard to complexity questions, we note that for the SRW, hitting times can be expressed as the solution to a set of n linear equations and can therefore be computed in polynomial time. Determining the complexity of computing the cover time, however, is far more challenging and still remains open [2, Open Problem 6.35]. Significant progress was made by Ding, Lee and Peres [20] who discovered a deterministic polynomial time $\mathcal{O}(1)$ -approximation algorithm for the cover time. In this paper, we show that computing an optimal strategy for the cover time of the CRW is NP-hard.

1.2 Our results

In this section, we shall present the main results we have obtained for CRW. The numbers of these theorems correspond to where they appear in the paper, although some theorem statements have been simplified for ease of exposition.

The CRW is not reversible in general; however, we show that it can emulate certain reversible chains. Combining this with the well-known connection between electrical networks and reversible Markov chains, we obtain the following general bound on the maximum hitting time $t_{\text{hit}}^{\text{two}}(G)$ between any two vertices of a graph G .

Theorem 1.1. *For any finite graph G , we have $t_{\text{hit}}^{\text{two}}(G) < \min\{3|E|, n^2\}$.*

This is tight up to constants at both ends of the density spectrum and improves considerably over the well-known $\mathcal{O}(n|E|)$ worst-case bound for the SRW. A witness to tightness for sparse graphs is traversing a path from end to end, and for dense graphs hitting a vertex connected by a single edge to a clique.

Most of this paper focuses on the cover time $t_{\text{cov}}^{\text{two}}(G)$ for CRW on a graph G under an optimal strategy. For the SRW t_{hit} , the maximum hitting time between any two vertices determines the cover time up to a $\log n$ factor by Matthews's bound [28, Ch. 11.2]. However, due to the effect of the choices, this does not apply to the CRW and so we develop other methods to bound $t_{\text{cov}}^{\text{two}}$.

The next result implies that $t_{\text{cov}}^{\text{two}}(T)$ is linear for a bounded degree tree T :

Theorem 1.2. *For every $d \in \mathbb{N}$ and every n -vertex tree T with maximum degree d , we have*

$$t_{\text{cov}}^{\text{two}}(T) \leq 8dn.$$

Our strategy for achieving this changes with time and covers the vertices of T in a prescribed order.

Next, we obtain a similar result for d -dimensional grids and tori. The proof technique is different: we show that there exists a CRW strategy for the infinite d -dimensional grid under which the CRW becomes strongly recurrent. In particular, the expected crossing time of any edge is finite. We use this to deduce

Theorem 1.3. *For any d , and any d -dimensional n -vertex torus or grid G , we have $t_{\text{cov}}^{\text{two}}(G) = \Theta(n)$ and $t_{\text{hit}}^{\text{two}}(G) = \Theta(\text{diam}(G)) = \Theta(n^{1/d})$.*

Avin and Krishnamachari [3] conjecture a speed-up for their aforementioned local power of two choice walk on the two-dimensional grid. Theorem 1.3 corroborates this for our global version of the process but does not yet prove their conjecture.

We develop a method for boosting the probabilities of rare events in the CRW, which gives bounds on hitting and cover times. Perhaps the most important application of these methods is to expander graphs:

Theorem 1.4. *For every sequence $(G_n)_{n \in \mathbb{N}}$ of bounded degree expanders, where G_n has n vertices, we have*

$$t_{\text{cov}}^{\text{two}}(G_n) = \mathcal{O}(n \log \log n).$$

Theorem 1.4 is in fact an immediate corollary of a more general bound (Theorem 6.1), bounding $t_{\text{cov}}^{\text{two}}(G)$ in terms of the hitting time (of the SRW), relaxation time and degree discrepancy. In particular, these bounds apply w.h.p. to the random d -regular graph for fixed d . Another application of these methods gives the following bounds for the Erdős–Rényi random graph, showing a significant improvement on cover time for the regime with subpolynomial growth of the average degree.

Theorem 1.5. *Let $\mathcal{G} \stackrel{d}{\sim} \mathcal{G}(n, p)$ where $np \geq c \ln n$ for any fixed $c > 1$ and $\log np = o(\log n)$. Then w.h.p.*

- (i) $t_{\text{cov}}^{\text{two}}(\mathcal{G}) = \mathcal{O}(n \cdot \log(np) \cdot \log \log n)$
- (ii) $t_{\text{hit}}^{\text{two}}(\mathcal{G}) = n^{1 - \Omega(1/\log(np))}$.

Finally, Section 7 deals with the computational complexity of computing optimal strategies to minimise hitting and cover times. We show the following dichotomy: an optimal strategy to hit a set of vertices can be computed efficiently, whereas choosing between two cover time strategies is NP-hard. More precisely, we have

Theorem 1.6. *For any graph G , $S \subset V$ and $x \in V \setminus S$, a strategy minimising the hitting time of S from x can be computed in time $\text{poly}(|V|)$.*

Notice that any strategy for covering a graph must specify a set of choice preferences from every vertex for every possible set of vertices covered and thus may have size exponential in n . This makes the second half of the dichotomy as phrased above sound somewhat modest. However, what we show is that even in the ‘on-line’ setting where one is given the set covered so far then just choosing the next step (outputting something of polynomial size) is NP-hard.

Theorem 1.7. *Given the covered set X and position v of the walk at some time, it is NP-hard to choose the next step from two neighbours of v so as to minimise the expected time for the CRW to visit every vertex not in X , assuming an optimal strategy is followed thereafter.*

Our proof shows that this remains NP-hard if G is constrained to have maximum degree 3. To the best of our knowledge, this is the first intractability result for processes involving random walks with choice.

Our results for fundamental graph topologies are summarised in Table 1, along with the corresponding hitting and cover times for the SRW for ease of comparison.

2. Preliminaries

The CRW is a discrete time stochastic process $(X_t)_{t \geq 0}$ on the vertices of a connected graph $G = (V, E)$, influenced by a *controller*. The starting state is a fixed vertex; at each time $t \in \mathbb{N}$ the controller is presented with two neighbours $\{c_1^t, c_2^t\}$ of the current state X_t chosen uniformly at random with replacement and must choose one of these neighbours as the next state X_{t+1} . We assume that at each time t the controller knows the graph G , its current position $X_t \in V$, and $\mathcal{H}_t := (X_i, \{c_1^i, c_2^i\})_{i=0}^t$, the *history* of the process so far. The controller has access to arbitrary computational resources and an infinite string of random bits ω in order to choose X_{t+1} from $\{c_1^t, c_2^t\}$. A *CRW strategy* is a function which given any t , \mathcal{H}_t and $\{c_1^t, c_2^t\} \subseteq \Gamma(X_t)$, outputs one

Table 1. All graphs are on n vertices. The third and fifth columns contain our results on the hitting and cover times of the choice random walk. The second and fourth columns give corresponding quantities for the simple random walk for comparison, which can be found in [2] and elsewhere. For random graphs, these bounds apply w.h.p

Graph family	Hitting time		Cover time	
	SRW	CRW	SRW	CRW
Subcubic ($d_{\max} \leq 3$)	$\mathcal{O}(n^2)$	$\Theta(\text{diam}(G))$	$\mathcal{O}(n^2)$	$\Theta(n)$
2-dim. torus/box	$\Theta(n \log n)$	$\Theta(n^{1/d})$	$\Theta(n \log^2 n)$	$\Theta(n)$
d -dim. torus/box	$\Theta(n)$	$\Theta(n^{1/d})$	$\Theta(n \log n)$	$\Theta(n)$
Bounded degree tree	$\mathcal{O}(n^2)$	$\mathcal{O}(n)$	$\mathcal{O}(n^2)$	$\Theta(n)$
d -regular expander	$\Theta(n)$	$n^{1-\Omega(1/\log d)}$	$\Theta(n \log n)$	$\mathcal{O}(n \log(d) \log \log n)$
$\mathcal{G}(n, \text{polylog}(n)/n)$	$\Theta(n)$	$n^{1-\Omega(1/\log \log n)}$	$\Theta(n \log n)$	$\mathcal{O}(n(\log \log n)^2)$
Complete graph	$n - 1$	$\frac{(n-1)^2}{2n-3}$	$\sim n \ln n$	$\sim (n \ln n)/2$

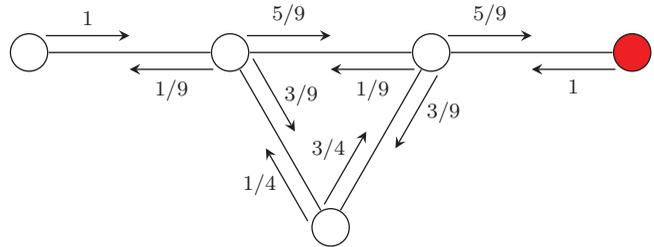


Figure 1. Optimal CRW transition probabilities on the bull graph for hitting the rightmost vertex. The corresponding Markov chain is not reversible.

of $\{c_1^t, c_2^t\}$ (where we write $\Gamma(v) := \{w:vw \in E\}$ for the neighbourhood of v). Note that any such strategy defines a Markov chain on V .

We say that a CRW strategy is *unchanging* if it is independent of both time and the history of the walk. We say that an unchanging strategy is *reversible* if the Markov chain it defines is reversible. We recall that any reversible Markov chain is identically distributed with a random walk on an edge-weighted graph as explained, for example, in [2], we shall make use of this representation. For many graphs with a high degree of symmetry, we can find good reversible strategies, and we can then use tools from the theory of reversible Markov chains to analyse the CRW on these graphs. The strategies we consider may use random bits in addition to those used for choosing $\{c_1^t, c_2^t\}$; we say a strategy is *deterministic* if no additional random bits are used.

If we are trying to minimise the expected hitting time of a given vertex, it is easy to see that there is an unchanging, deterministic optimal strategy. However, it need not be reversible; an example where it is not is given in Figure 1. We shall use reversible strategies to bound the hitting time of the optimal strategy; these will also in general not be deterministic.

For a strategy α and for a vertex v and distinct neighbours i, j let $\alpha_{v,i}^j \in [0, 1]$ be the probability that when the walk is at v it chooses i when offered $\{i, j\}$ as choices, that is, $\alpha_{v,i}^j := \mathbb{P}[X_{t+1} = i \mid X_t = v, c^t = \{i, j\}]$ (this probability is also conditional on \mathcal{H}_t but we suppress this for notational convenience). These are the only parameters we may vary, but we shall find it convenient to define $\alpha_{v,i}^i := 1/2$ for each i adjacent to v . Thus,

$$\text{for each } v \in V: \alpha_{v,i}^j \in [0, 1] \text{ and } \alpha_{v,i}^i = 1 - \alpha_{v,i}^j \text{ for all } i, j \in \Gamma(v). \tag{1}$$

The transition probabilities $q_{v,i}$ for the strategy α are then given by:

$$q_{v,i} = \frac{2 \sum_{j \in \Gamma(v)} \alpha_{v,i}^j}{d(v)^2}. \tag{2}$$

For a family of parameters $\alpha_{v,i}^j$, to yield a valid set of transition probabilities $q_{i,j}$ they must satisfy

$$\sum_{i \in \Gamma(v)} \sum_{j \in \Gamma(v)} \alpha_{v,i}^j = \frac{d(v)^2}{2}, \quad (3)$$

for every $v \in V$. Notice that any weights satisfying (1) also satisfy (3).

Let $C_v^{\text{two}}(G)$ denote the minimum expected time (taken over all strategies) for the CRW to visit every vertex of G starting from v and define the *cover time* $t_{\text{COV}}^{\text{two}}(G) := \max_{v \in V} C_v^{\text{two}}(G)$. Analogously, let $H_x^{\text{two}}(y)$ denote the minimum expected time for the CRW to reach y , which may be a single vertex or a set of vertices, starting from a vertex x and define the *hitting time* $t_{\text{hit}}^{\text{two}}(G) := \max_{x,y \in V} H_x^{\text{two}}(y)$. We drop the superscript from this notation when referring to the associated quantities for the SRW.

3. Bounds from weighted graphs

In this section, we analyse CRW strategies which emulate a random walk on a weighted graph. We prove a tight general bound on hitting times and show that any vertex of a graph with maximum degree 3 can be hit in time proportional to its distance from the start vertex.

3.1 An extremal hitting time result

In this section, we prove that $t_{\text{hit}}^{\text{two}}(G) = \mathcal{O}(e(G))$ for an arbitrary graph G , where $e(G)$ is the number of edges. This bound is best possible up to the implied constants: for sparse graphs, the path has $t_{\text{hit}}^{\text{two}}$ around $2e(G)$. For dense graphs, a clique with a pendant path, where the length of the path is growing much slower than the size of the clique, gives $t_{\text{hit}}^{\text{two}}$ around $3n^2/8$.

Lemma 3.1. *Fix a vertex v , and partition its neighbours into two sets, A and B . There is an unchanging strategy for the CRW such that whenever the walker is at v , it moves to a random neighbour according to the probability distribution in which every vertex in B is twice as likely as every vertex in A .*

Proof. Fix some number $p \in [0, 1]$ and consider the following strategy for moving from v . If offered two choices from the same set, choose between them uniformly at random, but if offered one choice from A and one choice from B , choose the one from A with probability p . Clearly, all elements of A are equiprobable, as are all elements of B , so it is sufficient to show that for some p this strategy chooses an element of A with probability $q = \frac{|A|}{|A|+2|B|}$. If this is the case, each element of A will be chosen with probability $\frac{1}{|A|+2|B|}$ and each element of B w.p. $\frac{2}{|A|+2|B|}$. If $p = 1/2$ the probability of choosing an element of A is $\frac{|A|}{|A|+|B|} \geq q$, and if $p = 0$ then it is $q' := \left(\frac{|A|}{|A|+|B|}\right)^2$. Since $(|A| + |B|)^2 \geq |A|(|A| + 2|B|)$, we have $q' \leq q$, and hence for some $p \in [0, 1/2]$ we have the required probability by continuity. \square

By considering the strategy at each vertex separately, we immediately get the following consequence.

Corollary 3.2. *Let $G = (V, E)$ be a locally finite weighted graph with weight function $w: E \rightarrow \mathbb{R}_+$, having the property that for any two incident edges xy, xz either $w(xy) = w(xz)$, or $w(xy) = 2w(xz)$, or $2w(xy) = w(xz)$. Then there is an unchanging strategy for the CRW on G which simulates the random walk defined by the weights w .*

Here, by the *random walk defined by the weights* w , we mean the reversible Markov chain where the transition probability from a vertex x to a neighbour y is proportional to $w(xy)$. For a weighted graph (G, w) , write $w(G) = \sum_{e \in E(G)} w(e)$.

Lemma 3.3. *Let (G, w) be a finite weighted graph, and let x be a vertex such that every edge incident with x has weight 1. Then for any vertex y adjacent to x , we have*

$$H_y(x) \leq w(G) + w(G \setminus x).$$

Proof. Since the stationary distribution is given by $\pi_v = \frac{1}{2w(G)} \sum_{u \sim v} w(uv)$, we have expected return time to x of $\pi_x^{-1} = \frac{2w(G)}{d(x)}$, see for example [2, Sect. 3.2]. Thus,

$$\frac{2w(G)}{d(x)} = 1 + \sum_{z \sim x} \frac{1}{d(x)} H_z(x),$$

implying

$$H_y(x) \leq \sum_{z \sim x} H_z(x) = 2w(G) - d(x) = w(G) + w(G \setminus x). \quad \square$$

We now restate and prove our result for CRW hitting times.

Theorem 1.1 *For any finite graph G , we have $t_{\text{hit}}^{\text{two}}(G) < \min\{3|E|, n^2\}$.*

Proof. We have to show that the above bounds apply to $H_y^{\text{two}}(x)$ for two arbitrary vertices x, y . Define a weight function $w: E(G) \rightarrow \mathbb{R}_+$ by $w(uv) = 2^{-\min(d(u,x), d(v,x))}$. Note that w satisfies the requirements of Corollary 3.2, so we can bound $H_y^{\text{two}}(x)$ by the corresponding hitting time of the random walk on (G, w) . We will now bound the latter hitting time.

Write d for the maximum distance of a vertex from x , and V_k for the set of vertices at distance exactly k from x . Note that if $y \in V_{k+1}$ then

$$H_y(x) \leq H_y(V_k) + \max_{z \in V_k} H_z(x),$$

and consequently

$$\max_{y \in V(G)} H_y(x) \leq \sum_{k=0}^{d-1} \max_{z \in V_{k+1}} H_z(V_k).$$

For each $0 \leq k \leq d - 1$, let G_k be the simple weighted graph obtained by deleting $\bigcup_{i < k} V_i$ and identifying vertices in V_k to give a vertex v_k ; if a vertex in V_{k+1} has multiple edges to V_k , delete all but one of them to leave a simple graph. Since removing edges between V_{k+1} and V_k cannot reduce the hitting time of V_k , we have for any $z \in V_{k+1}$ that $H_z^G(V_k) \leq H_z^{G_k}(v_k)$. Note that the latter hitting time is unchanged by multiplying all weights by 2^k , and since every $z \in V_{k+1}$ is adjacent to v_k in G_k , by Lemma 3.3 we have $H_z^{G_k}(v_k) \leq 2^k(w(G_k) + w(G_k \setminus v_k))$. Thus,

$$\max_{y \in V(G)} H_y(x) \leq \sum_{k=0}^{d-1} 2^k(w(G_k) + w(G_k \setminus v_k)).$$

If e is an edge between V_j and V_{j+1} , then the contribution of e to the k th term of the above sum is 2^{k-j+1} if $k < j$, at most 1 if $k = j$ and 0 otherwise, so its total contribution is less than 3 and is less than 2 if e is one of the edges deleted to make G_j simple. If e is an edge within V_j , then its contribution to the k th term is 2^{k-j+1} if $k < j$ and 0 otherwise, so its total contribution is less than 2. The first bound follows. Note that of the edges of the first type which are not deleted,

there is exactly one from each vertex (other than x) to a vertex in a lower layer of G , and so these edges form a tree. Thus, there are $n - 1$ such edges, whose contribution is bounded by 3, and at most $\binom{n}{2} - (n - 1)$ other edges, whose contribution is bounded by 2, giving a bound of $2\binom{n}{2} + n - 1 = n^2 - 1$. \square

3.2 Cover times of subcubic graphs

In this section, we prove the CRW cover time of any subcubic graph is linear in the number of vertices, where we remind the reader that a subcubic graph is a graph with maximum degree 3.

Proposition 3.4. *Let G be any connected graph of maximum degree 3. Then, $H_u^{\text{two}}(v) \leq 9$ for any $uv \in E(G)$. If in addition G is finite with n vertices, then $t_{\text{cov}}^{\text{two}}(G) = \Theta(n)$.*

Proof. For each $w \neq v$, choose a neighbour $f(w)$ such that $d(f(w), v) < d(w, v)$. Set $f(v) = u$. Let (X_t) be a CRW starting at u using the strategy of choosing $X_{t+1} = f(X_t)$ whenever possible. Couple this with a random walk Y_t on \mathbb{Z} starting at 1, with $Y_{t+1} = Y_t - 1$ if $X_{t+1} = f(X_t)$ and $Y_{t+1} = Y_t + 1$ otherwise. Clearly $Y_t \geq \text{dist}(X_t, v)$ at every step of the walk, and so X_t reaches v on or before the first time t that $Y_t = 0$. Since $\mathbb{P}[Y_{t+1} = Y_t - 1] \geq 5/9$, we have

$$H_u^{\text{two}}(v) \leq \mathbb{E}[\min\{t: Y_t = 0\} \mid \{Y_0 = 1\}] \leq 9. \tag{4}$$

If G has n vertices, let v be any vertex and choose a spanning walk in the graph starting at v and having at most $2n - 3$ edges. Such a walk always exists, for example, a depth-first exploration of a spanning tree. Proceed in $2n - 3$ rounds, in each round using the strategy above to hit the next vertex of the walk. Each round has expected duration at most 9 by (4), and so $t_{\text{cov}}^{\text{two}}(G) \leq 18n - 27$. \square

Remark. Since $t_{\text{hit}}^{\text{two}}(G) \leq t_{\text{cov}}^{\text{two}}(G)$, this is also linear. Even for 3-regular graphs the diameter could grow linearly, so this is the best possible.

4. Trees

In this section, we show that $t_{\text{cov}}^{\text{two}}(T) = \Theta(n)$ for trees T of bounded degree. Even more, we will prove that we can even specify an arbitrary (closed) walk W traversing each edge of T once in each direction and cover the vertices of T in the order dictated by W in linear expected time. This is the gist of the following result:

Theorem 1.2 *For every $d \in \mathbb{N}$ and every tree T with maximum degree d , we have*

$$\sum_{x,y \in V(T), xy \in E(T)} H_x^{\text{two}}(y) \leq 8d|V(T)|.$$

This result will be proved by realising a strategy to cover T as a sequence of weighted walks, and then bounding the hitting times in these walks using the Essential Edge Lemma. We shall now remind the reader of the setting and statement of this lemma: we say that an edge vx of a graph is *essential* if its removal would disconnect the graph, into two components $A(v, x)$ and $A(x, v)$, say, containing v and x , respectively. Let $E(v, x)$ be the set of edges of $A(v, x)$.

Lemma 4.1. Essential Edge Lemma [2, Lem. 5.1]. *Let G be any graph with edge weights $\{w_{ij}\}_{ij \in E}$. If vx is an essential edge, then*

$$H_v(x) = \frac{2 \sum_{\{i,j\} \in E(v,x)} w_{ij}}{w_{v,x}} + 1,$$

where $H_v(x)$ is the hitting time of the reversible Markov chain defined by these edge weights.

We now define the CRW strategies we use in the proof of Theorem 1.2. Given a tree T , we pick an arbitrary ‘root’ vertex $r \in V(T)$. In order to obtain an upper bound on $H_x^{\text{two}}(y)$ for $x, y \in V(T)$ such that $xy \in E(T)$, we follow the (unchanging) strategy σ^{xy} making the following choices at each vertex v :

$$\begin{aligned} &\text{Reduce the distance to } y \text{ if possible. Otherwise, choose uniformly} \\ &\text{an option that increases distance to } r \text{ if at least one is available.} \end{aligned} \tag{5}$$

In other words, σ^{xy} prefers the unique neighbour w of v with $d(w, y) < d(v, y)$, avoids the unique neighbour z with $d(z, r) < d(v, r)$ and is indifferent among all other neighbours of v .

We emphasise that r was an arbitrary vertex, but it is important for our calculations below that it is fixed for all σ^{xy} , $x, y \in V(T)$.

Since the strategy σ^{xy} is unchanging, there is an assignment of weights $w_{x,y}(e)$, $e \in E(T)$ such that the corresponding random walk (as defined after Corollary 3.2) is equidistributed with the CRW under strategy σ^{xy} when both walks start at x and stop when first visiting y . These weights can be multiplied by any positive constant without changing the random walk they define, and we normalise by fixing $w_{x,y}(xy) = 1$ for concreteness. The rest of the weights can be calculated explicitly, and so we can apply the Lemma 4.1 to give the bound:

$$H_x^{\text{two}}(y) < 2 \sum_{e \in E(T)} w_{x,y}(e), \tag{6}$$

with the understanding that we set $w_{x,y}(e) = 0$ if y separates x from e as this edge does not contribute to the sum in Lemma 4.1.

The latter formula expresses $H_x^{\text{two}}(y)$ as a sum of contributions of each $e \in E(T)$. The main surprise in the proof of Theorem 1.2 is the following lemma, which says that for each $e \in E(T)$, the sum of these contributions $w_{x,y}(e)$ over all $H_x^{\text{two}}(y)$, $x, y \in V(T)$, $xy \in E(T)$, is bounded. An obvious double-counting argument involving (6) will then establish Theorem 1.2.

Lemma 4.2. *For every $d \in \mathbb{N}$, every tree T with maximum degree d , and every edge $e \in E(T)$, we have $\sum_{x,y \in V(T), xy \in E(T)} w_{x,y}(e) \leq 4d$.*

We emphasise that this sum is taken over all ordered pairs of adjacent vertices. The proof of Theorem 1.2 is based on the fact that, for a fixed e , the values $w_{x,y}(e)$ decay fast with the distance $d(xy, e)$, and even more so in the direction of r . (Here, we define the distance between two edges $xy, wz \in E$ to be $d(xy, wz) := \min\{d(x, w), d(y, z), d(x, z), d(y, w)\}$.) The following two propositions will yield quantitative bounds on the speed of this decay.

Proposition 4.3. *Let G be any graph, $x \in V(G)$, and $v \in N(x)$. Consider a CRW strategy that when at x always chooses v when that choice is available, otherwise it chooses each of the available options independently with probability $1/2$. Then for every $w \neq v \in N(x)$, the transition probabilities satisfy $q_{x,w}/q_{x,v} < 1/2$.*

Proof. Let $d := d(x)$. We have $q_{x,v} = 1 - \left(\frac{d-1}{d}\right)^2 = \frac{d^2 - (d^2 - 2d + 1)}{d^2} = \frac{2d-1}{d^2}$ and $q_{x,w} = \left(\frac{d-1}{d}\right)^2 \frac{1}{d-1}$ since each $w \neq v$ is chosen with equal probability, and only when v is not among the options. Thus, $q_{x,w}/q_{x,v} = \left(\frac{d-1}{d}\right)^2 \frac{d^2}{(d-1)(2d-1)} = \frac{d-1}{2d-1} < 1/2$ as claimed. □

Proposition 4.4. *Let G be any graph, $x \in V(G)$, and $v, z \in N(x)$ where $v \neq z$. Consider a CRW strategy that when at x always chooses v when that choice is available, never chooses z unless there is no other option, and it chooses each of the other available options independently with probability $1/2$. Then the transition probabilities satisfy $q_{x,z}/q_{x,v} = \frac{1}{2d(x)-1}$.*

Proof. As in the proof of Proposition 4.3, we have $q_{x,v} = \frac{2d(x)-1}{d(x)^2}$. Easily, $q_{x,z} = \frac{1}{d(x)^2}$. Thus, $q_{x,z}/q_{x,v} = \frac{d(x)^2}{(2d(x)-1)d(x)^2} = \frac{1}{2d(x)-1}$ as claimed. \square

Armed with these propositions, we can now prove Lemma 4.2.

Proof of Lemma 4.2. Fix $e \in E(T)$. We split $\Sigma := \sum_{x,y \in V(T), xy \in E(T)} w_{x,y}(e)$ as a sum $\Sigma = \sum_{i \in \mathbb{N}} \Sigma_i$ of ‘layers’ Σ_i , corresponding roughly to distance from e , and show that Σ_i decays exponentially in i .

Let P be the path from e to r in T (excluding e), and let L_0 be the set of all edges of P and all edges incident with P (including e). Let $\Sigma_0 := \sum_{x,y \in V(T), xy \in L_0} w_{x,y}(e)$ be the total weight assigned to e by pairs of adjacent vertices of L_0 . Define $L_i, i \geq 1$ recursively as the set of edges incident with L_{i-1} not contained in $\bigcup_{j < i} L_j$, and let $\Sigma_i := \sum_{x,y \in V(T), xy \in L_i} w_{x,y}(e)$. See Figure 2 for an illustration of the different sets L_i .

Claim *The following inequalities hold*

- (i) $\Sigma_0 \leq 2d$,
- (ii) $\Sigma_i \leq \Sigma_{i-1}/2$.

Proof of Claim (i). Let x_1, x_2, \dots, x_k , where $x_k = r$ be the vertices of P as they appear from e to r . Recall that $w_{x_{i+1},x_i}(e) = 0$ for every $i \geq 1$, as e is separated from x_{i+1} by removing x_i and thus does not contribute to the sum in the formula for $H_{x_{i+1}}(x_i)$ from Lemma 4.1. In the other direction, we claim $w_{x_i,x_{i+1}}(e) < (1/2)^i$ for every $i \geq 1$. Indeed, by Proposition 4.3, we have $\frac{w_{x_i,x_{i+1}}(x_{i-1},x_i)}{w_{x_i,x_{i+1}}(x_i,x_{i+1})} < 1/2$ because this ratio coincides with the ratio of the corresponding transition probabilities $\frac{q_{x_i,x_{i-1}}}{q_{x_i,x_{i+1}}}$ by the definitions. Moreover, at each $x_j, 1 \leq j < i$, the strategy $\sigma^{x_i x_{i+1}}$ makes the same choices as $\sigma^{x_j x_{j+1}}$, hence

$$\frac{w_{x_i,x_{i+1}}(x_{j-1}, x_j)}{w_{x_i,x_{i+1}}(x_j, x_{j+1})} = \frac{w_{x_j,x_{j+1}}(x_{j-1}, x_j)}{w_{x_j,x_{j+1}}(x_j, x_{j+1})} < 1/2,$$

by Proposition 4.3 again, with the convention that $x_0 x_1 = e$. Our claim follows by multiplying these fractions for j ranging from 1 to i .

For each of the at most $d - 1$ other edges $x_i z \neq e$ of L_0 incident with x_i , where $i \geq 1$, we use the rather generous bound $w_{x_i,z}(e) < (1/2)^{i-1}$, which is true by similar arguments because $\frac{w_{x_i,z}(x_{i-1},x_i)}{w_{x_i,z}(x_i,z)} < 1$. Again we have $w_{z,x_i}(e) = 0$.

Finally, we have $w_{x_0,x_1}(e) = w_{x_1,x_0}(e) = 1$ since $e = x_0 x_1$. Adding these contributions, we obtain $\Sigma_0 \leq 2 + \sum_{i \geq 0} (d - 1)(1/2)^i = 2 + 2(d - 1) = 2d$ as claimed. \diamond

Proof of Claim (ii). Let $yv \in L_{i-1}$. We will bound the contribution of the edges incident with yv to Σ_i in terms of the contribution of yv to Σ_{i-1} . For this, let $vw \in L_i$. Note that v separates w from r and so first, $w_{w,v}(e) = 0$. Second, this implies σ^{vw} avoids moving from v to y whenever possible by (5). Thus, Proposition 4.4 yields $\frac{w_{v,w}(yv)}{w_{v,w}(vw)} = \frac{1}{2d(v)-1}$. Moreover, when at a vertex other than v , the strategies σ^{vw} and σ^{yv} make the same choices since the directions of r as well as of the corresponding target vertex coincide. Therefore, $\frac{w_{y,v}(f)}{w_{y,v}(g)} = \frac{w_{v,w}(f)}{w_{v,w}(g)}$ for every two edges f, g incident with a common vertex on the y - e path. It follows that $w_{v,w}(e) = \frac{w_{y,v}(e)}{2d(v)-1}$, and summing over all such neighbours w of v we obtain $\sum_{vw \in L_i} w_{v,w}(e) \leq \frac{w_{y,v}(e)(d(v)-1)}{2(d(v)-1)+1} < w_{y,v}(e)/2$. Applying this to

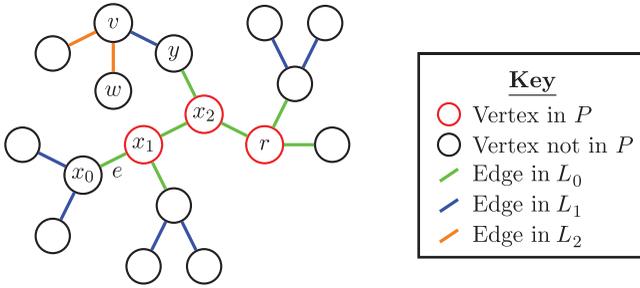


Figure 2. Illustration of the edge sets L_i and path P in the proof of Lemma 4.2. The vertices y, v and w are labelled consistently with the proof of Lemma 4.2, Claim (ii) in the case where $yv \in L_1$ and $vw \in L_2$.

each edge $yv \in L_{i-1}$, and adding together, noting that at most one end vertex v of yv is incident with edges in L_i by construction, we finally deduce

$$\Sigma_i = \sum_{yv \in L_{i-1}} \sum_{vw \in L_i} w_{v,w}(e) < \sum_{yv \in L_{i-1}} w_{y,v}(e)/2 = \Sigma_{i-1}/2,$$

as desired. ◇

Combining both parts of the Claim proves our statement, as $\Sigma = \sum_i \Sigma_i \leq 2\Sigma_0 \leq 4d$. □

It is now easy to complete the proof of Theorem 1.2.

Proof of Theorem 1.2. By (4.1), we have

$$\sum_{x,y \in V(T), xy \in E(T)} H_x^{\text{two}}(y) < \sum_{x,y \in V(T), xy \in E(T)} 2 \sum_{e \in E(T)} w_{x,y}(e).$$

Changing the order of summation, and then applying Lemma 4.2 to each summand, we bound the right-hand side by

$$2 \sum_{e \in E(T)} \left(\sum_{x,y \in V(T), xy \in E(T)} w_{x,y}(e) \right) \leq 2 \sum_{e \in E(T)} 4d = 2(|V(T)| - 1)4d \leq 8d|V(T)|. \quad \square$$

5. Infinite graphs and cover time of tori

In this section, we bound the cover time of the d -dimensional discrete torus \mathbb{Z}_k^d , which has $n = k^d$ vertices. Here, we think of the dimension d as being fixed while the side length k grows. In order to prove a linear bound on the cover time, we will instead consider the infinite limit \mathbb{Z}^d and infinite (but locally finite) graphs more generally.

For infinite graphs, it is meaningless to ask about the CRW cover time, but still interesting to ask about hitting times. The most fundamental question is whether these can be made finite, which corresponds to asking for positive recurrence.

Definition 5.1. A graph is *positive choice recurrent (PCR)* if there exists an unchanging strategy for the CRW such that the expected return time to any given vertex is finite. A graph is *strongly PCR* if for every $p \in (0, 1)$ there exists an unchanging CRW strategy such that expected return times are finite for the process which, at every time step, takes a step of the CRW with that strategy with probability p and a step of the SRW otherwise.

A natural question is whether there is a strategy under which the walk becomes a transient Markov chain. The answer is always yes: fixing a root r and giving each edge uv weight

$2^{\min(d(u,r),d(v,r))}$ produces a suitable weighting to apply Corollary 3.2. This weighted graph is transient because any infinite geodesic starting at the root has total resistance 2 (see e.g. [29, Thm. 2.3]), and taking other edges into account cannot increase the effective resistance to infinity.

While positive recurrence is the property which will be useful to us, we might also ask for the weaker property of *choice recurrence*, where we simply require return times to be almost surely finite. It is possible for a graph to be choice recurrent but not PCR; indeed, there are graphs which are recurrent under the SRW but not PCR.

Remark. Proposition 3.4 implies that any graph of maximum degree 3 is PCR. This is not true for higher degrees, since for the infinite 4-regular tree any strategy is more likely to move away from a given target vertex than towards it.

Note that $\mathbb{Z}^d = \mathbb{Z}^{d-1} \square \mathbb{Z}$, where \square indicates the Cartesian product. We will need the following result about Cartesian products of PCR graphs.

Lemma 5.2. *If G is PCR, H is strongly PCR and both G, H are regular, then $G \square H$ is PCR.*

Proof. Define the p -product of two time-homogeneous Markov chains A, B to be the chain with state space $S(A) \times S(B)$ where at each time step with probability p a transition of B occurs, and otherwise a transition of A occurs. If both chains are irreducible and positive recurrent, then so is the p -product (this follows easily from the existence of stationary distributions). Now we define a strategy for the CRW on $G \square H$ as follows. If at least one of the choices given is a move in the H co-ordinate, we make such a move. Now the probability of exactly one of the options being a move in H is $\frac{2rs}{(r+s)^2}$, where G is r -regular and H is s -regular, and the probability of both options being moves in H is $\frac{s^2}{(r+s)^2}$. Thus, conditional on at least one option being in H , both are in H with probability $\frac{s}{2r+s}$. There is a strategy on H , for this probability of having two choices, which reaches the root in finite time; whenever we move in the H co-ordinate we use this strategy. If both choices are moves in G , then we follow the appropriate strategy for the random walk with two choices in G . The resulting Markov chain is the $\frac{2rs+s^2}{(r+s)^2}$ -product of positive recurrent Markov chains on G and H , hence positive recurrent. □

The same argument shows that if in addition G is strongly PCR, then so is $G \square H$. Lemma 5.2 allows us to conclude that \mathbb{Z}^d is PCR and consequently obtain a bound on its cover times and hitting times.

Theorem 1.3 *For any d , the cover time of the CRW in the finite d -dimensional torus \mathbb{Z}_k^d or grid $[k]^d$ is $\Theta(n)$ and the hitting time is $\Theta(k) = \mathcal{O}(n^{1/d})$, where n is the number of vertices.*

Proof. Note that \mathbb{Z} is strongly PCR, since always moving towards 0 if possible gives a random walk which moves towards 0 with probability $\frac{1}{2} + \frac{p}{4} > \frac{1}{2}$ (where p is the probability of taking a step of the CRW). Inductively applying Lemma 5.2 implies \mathbb{Z}^d is PCR for any d , and so the hitting time to a neighbour is some constant c_d . This gives an upper bound on the hitting time to a neighbour in \mathbb{Z}_k^d or $[k]^d$, and the strategy of visiting the vertices of a Hamilton path in order gives a cover time of less than $c_d k^d$. Similarly the hitting time $H_x^{two}(y)$ is bounded by constant times the length of the shortest path between x and y , and the worst-case value of this is $d \lfloor k/2 \rfloor$ for the torus and dk for the grid, so $\Theta(k)$. Both bounds are trivially best possible up to the constant. □

6. Hitting and cover times in expanders

In this section, we prove bounds on the cover and hitting times of the CRW on a graph G in terms of fundamental parameters. First, we introduce our notation. Let G be a graph with n vertices, and write d_{\max} , d_{\min} and d_{avg} for the maximum, minimum and average degree of G , respectively. Let t_{rel} be the *relaxation time* of G , defined as $\frac{1}{1-\lambda_2}$, where λ_2 is the second largest eigenvalue of the transition matrix of the lazy random walk (LRW) on G with loop probability $1/2$. Recall that t_{hit} is maximum over all pairs $uv \in V$ of the expected time it takes the SRW to reach u from v . Our first result bounds the CRW cover time.

Theorem 6.1. *For any connected n -vertex graph G the following holds*

$$t_{\text{cov}}^{\text{two}}(G) = \mathcal{O}\left(t_{\text{hit}} \cdot \log(d_{\max}) \cdot \log\left(\frac{d_{\text{avg}} \cdot t_{\text{rel}} \cdot \log n}{d_{\min}}\right)\right),$$

We also bound hitting times. First, we define the exponent $\gamma_d = \log_d \frac{d^2}{2d-1}$; note that γ_d is increasing in d , $\gamma_d < 1$ and $1 - \gamma_d \sim 1/\log_2 d$. Also recall that for a set $S \subseteq V$ let $\pi(S) = \sum_{s \in S} \pi(s)$ be the stationary probability of S .

Theorem 6.2. *For any graph G , and any $x \in V$ and $S \subset V$, we have*

$$H_x^{\text{two}}(S) \leq 12 \cdot \pi(S)^{-\gamma_{d_{\max}}} \cdot t_{\text{rel}} \cdot \ln n;$$

this bound also holds for return times. Consequently,

$$t_{\text{hit}}^{\text{two}}(G) \leq 12 \left(\frac{n \cdot d_{\text{avg}}}{d_{\min}}\right)^{\gamma_{d_{\max}}} \cdot t_{\text{rel}} \cdot \ln n.$$

We say that a sequence of graphs (G_n) is a *sequence of expanders* if $t_{\text{rel}}(G_n) = \Theta(1)$. Theorems 6.1 & 6.2 yield the following corollary:

Theorem 1.4 *For every sequence $(G_n)_{n \in \mathbb{N}}$ of bounded degree n -vertex expanders, we have*

$$t_{\text{cov}}^{\text{two}}(G_n) = \mathcal{O}(n \log \log n) \quad \text{and} \quad t_{\text{hit}}^{\text{two}}(G_n) \leq n^\alpha \quad \text{for some fixed } \alpha < 1.$$

These are significantly less than the corresponding cover and hitting times by the SRW, which are $\Theta(n \log n)$ and $\Theta(n)$, respectively [2, Thm. 10.1].

Theorems 6.1 and 6.2 will follow from Theorem 6.3 below. For a given graph G , we consider possible *trajectories* of a (non-lazy) walker, that is, finite sequences of vertices in which any two consecutive vertices are adjacent; the length of a trajectory will be the number of steps taken. In the following, we use bold characters to denote trajectories in G and if $u \in V(G)$, then u will denote the length 0 trajectory from u . Fix a non-negative integer t and a set S of trajectories of length t . Let $p_{\mathbf{x},S}$ denote the probability that extending a trajectory \mathbf{x} to length t according to the law of a SRW results in a member of S . Let $q_{\mathbf{x},S}$ denote the corresponding probability under the CRW law; this probability will depend on the particular strategy used. This function can encode probabilities of many events of interest such as ‘the graph is covered by time t ’, ‘the walk is in a set X at time t ’ or ‘the walk has hit a vertex x by time t ’ for example. However, let us emphasise that our result in fact applies to *any* possible event.

Theorem 6.3. *Let G be a graph, $u \in V$, $t > 0$ and S be a set of trajectories of length t from u . Then, there exists a strategy for the CRW such that $q_{u,S} \geq (p_{u,S})^{\gamma_{d_{\max}}}$.*

We also give an analogue of Theorem 6.3 for bad events. This analogue, unlike Theorem 6.3, gives an exponent which does not depend on the maximum degree d_{\max} of G , and so a significant reduction is possible even if d_{\max} is large.

Theorem 6.4. *Let G be a graph, $u \in V$, $t > 0$, and S be a set of trajectories of length t from u . Then there exists a strategy for the CRW such that $q_{u,S} \leq (p_{u,S})^2$.*

Remark. The exponent 2 in Theorem 6.4 is best possible, since we have equality whenever $t = 1$ and therefore also when $t > 1$ but every trajectory of the SRW of length $t - 1$ has the same probability to reach S . Similarly the exponent given in Theorem 6.3 is best possible, as evidenced by the case where this probability is $1/d_{\max}$ for every trajectory of length $t - 1$.

After stating two technical lemmas in Section 6.2, we then explain an alternative way of considering the CRW in Section 6.3, which enables the proof of Theorems 6.3 and 6.4 to be completed. To motivate the importance of Theorem 6.3, we shall begin by showing how it implies our main results on cover time and hitting times.

6.1 Deducing Theorems 6.1 and 6.2 from Theorem 6.3

In order to prove our main bounds from the key tool, Theorem 6.3, we must first overcome the obstacle that Theorem 6.3 is expressed in terms of the SRW probabilities, whereas our bounds involve the relaxation time, which is defined in terms of the LRW. The reason for using the LRW to define relaxation time is to ensure that the associated Markov chain is aperiodic. Our next lemma resolves this issue by relating the relaxation time to SRW probabilities.

Write $p_{x,\cdot}^{(t)}$ and $\tilde{p}_{x,\cdot}^{(t)}$ for the distribution of the SRW and LRW, respectively, after t steps started at x , and write $\pi(S)$ for the stationary probability of a set S (note that the two walks have the same stationary distribution).

Lemma 6.5. *For any graph G , $S \subseteq V$ and $x \in V$, there exists $t \leq 4t_{\text{rel}} \ln n$ such that*

$$p_{x,S}^{(t)} \geq \pi(S)/3.$$

Proof. If G is bipartite, then we may find a subset $\bar{S} \subseteq S$ which lies entirely within one part satisfying $\pi(\bar{S}) \geq \pi(S)/2$. Otherwise, the SRW is aperiodic and we set $\bar{S} = S$. We now consider the multigraph \bar{G} formed from G by contracting \bar{S} to a single vertex \bar{s} , retaining all edges (with edges inside \bar{S} becoming loops at \bar{s}). Retaining edges ensures that the stationary probability of \bar{s} in \bar{G} is precisely $\pi(\bar{S})$. Let $\bar{\lambda}_2$ be the second largest eigenvalue of the LRW on \bar{G} . Then for any $x \notin \bar{S}$ and $t \geq 0$, by [28, (12.11)], we have $|\tilde{p}_{x,\bar{s}}^{(t)} - \pi(\bar{S})| \leq \sqrt{\pi(\bar{S})/\pi(x)} \cdot e^{-t(1-\bar{\lambda}_2)}$. It follows that if we run the LRW on \bar{G} for $T = \log(3/\sqrt{\pi(\bar{S})\pi(x)})/(1 - \bar{\lambda}_2)$ steps then

$$\tilde{p}_{x,\bar{s}}^{(T)} \geq \pi(\bar{S}) - \sqrt{\frac{\pi(\bar{S})}{\pi(x)}} \cdot \frac{\sqrt{\pi(\bar{S})\pi(x)}}{3} \geq \frac{2\pi(\bar{S})}{3}.$$

Now, we can express the density of the LRW by $\tilde{p}_{x,S}^{(T)} = \mathbb{E} \left[p_{x,S}^{(X_T)} \right]$, where the random variable $X_T \sim \text{Bin}(T, 1/2)$ is the number of non-lazy steps taken by the LRW in time T . Thus,

$$\max_{t \leq T} p_{x,S}^{(t)} \geq \tilde{p}_{x,S}^{(T)} \geq \frac{2\pi(\bar{S})}{3} \geq \frac{\pi(S)}{3}.$$

We can assume $n \geq 2$ or else the result holds trivially, so $\log(3/\sqrt{\pi(\bar{S})\pi(x)}) \leq \log 3 + 2 \log n \leq 4 \log n$. Finally, [2, Cor. 3.27] gives that $\bar{\lambda}_2 \leq \lambda_2$, so $T \leq 4t_{\text{rel}} \ln n$. □

Our strategy to bound the cover time will be to emulate the SRW until most of the vertices are covered, only using the additional strength of the CRW when there are few uncovered vertices remaining. We will need a simple lemma to bound how long the first stage takes.

Lemma 6.6. Let $U(t)$ be the number of unvisited vertices at time t by a SRW on a graph and let $T_{n/2^x}$ be the number of SRW steps taken before $U \leq n/2^x$. Then

$$\mathbb{E} [U(2x \cdot t_{\text{hit}})] \leq \frac{n}{2^x} \quad \text{and} \quad \mathbb{E} [T_{n/2^x}] \leq 4(x + 1)t_{\text{hit}}.$$

Proof. Let $v \in V$. Then by Markov’s inequality $\mathbb{P}_w[X_t \neq v, \forall 0 \leq t \leq 2t_{\text{hit}}] \leq 1/2$, for any $w \in V$. Thus, the probability v is not visited by time $2x \cdot t_{\text{hit}}$ is at most 2^{-x} by sub-multiplicity and so the expected number of unvisited vertices at time $2x \cdot t_{\text{hit}}$ is at most $n \cdot 2^{-x}$.

By the above $\mathbb{E} [U(2(x + 1)t_{\text{hit}})] \leq n/(2 \cdot 2^x)$ and so $\mathbb{P} [U(2(x + 1)t_{\text{hit}}) \geq n/2^x] \leq 1/2$ by Markov’s inequality. Considering sections of length $2(x + 1)t_{\text{hit}}$ separately, and continuing until one section covers the required number of vertices, we use in expectation at most two such sections, thus $\mathbb{E} [T_{n/2^x}] \leq 4(x + 1)t_{\text{hit}}$. \square

We now have what we need to prove the cover and hitting time bounds.

Proof of Theorem 6.1. For convenience, we write $\gamma = \gamma_{d_{\text{max}}}$. We first emulate the SRW (i.e. set $\alpha_{x,y}^z = 1/2$ for all $x, y, z \in V(G)$ with $y, z \in \Gamma(x)$) until all but $m = \lfloor n/\log^C n \rfloor$ vertices have been visited, for some C to be specified later. Let τ_1 be the expected time to complete this phase. Then, by Lemma 6.6, we have $\tau_1 \leq 4t_{\text{hit}} \cdot C \log_2 \log n$.

We cover the remaining vertices in m different phases, labelled $m, m - 1, \dots, 1$, each of which reduces the number of uncovered vertices by 1. In phase i , a set of i vertices are still uncovered, and we write S_i for this set. By Lemma 6.5 for any vertex x , there is some $t \leq 4t_{\text{rel}} \log n$ such that

$$p_{x,S_i}^{(t)} \geq \frac{\pi(S_i)}{3} = \frac{1}{3} \cdot \frac{\sum_{s \in S_i} d(s)}{nd_{\text{avg}}} \geq \frac{d_{\text{min}} \cdot i}{3nd_{\text{avg}}},$$

and thus $q_{u,S_i}^{(t)} \geq (d_{\text{min}} \cdot i/(3nd_{\text{avg}}))^\gamma$ by Theorem 6.3. Since from any starting point, we can achieve this probability of hitting a vertex in S_i within the next $4t_{\text{rel}} \log n$ steps, the expected number of attempts needed to achieve this is at most $(d_{\text{min}} \cdot i/(3nd_{\text{avg}}))^{-\gamma}$, meaning that the expected time required to complete phase i is at most

$$\mathcal{O}\left(\left(\frac{n \cdot d_{\text{avg}}}{i \cdot d_{\text{min}}}\right)^\gamma \cdot t_{\text{rel}} \cdot \log n\right).$$

Hence, the expected time τ_2 to complete all m phases satisfies

$$\begin{aligned} \tau_2 &= \sum_{i=1}^{n/\log^C n} \mathcal{O}\left(\left(\frac{nd_{\text{avg}}}{id_{\text{min}}}\right)^\gamma t_{\text{rel}} \log n\right) \\ &= \mathcal{O}\left(\left(\frac{nd_{\text{avg}}}{d_{\text{min}}}\right)^\gamma t_{\text{rel}} \log n\right) \sum_{i=1}^{n/\log^C n} i^{-\gamma}. \end{aligned}$$

Then, since $\sum_{i=1}^{n/\log^C n} i^{-\gamma} \leq (n/\log^C n)^{1-\gamma} \cdot \sum_{i=1}^{n/\log^C n} i^{-1} \leq (n/\log^C n)^{1-\gamma} \cdot \log n$, we have

$$\begin{aligned} \tau_2 &= \mathcal{O}\left(\left(\frac{nd_{\text{avg}}}{d_{\text{min}}}\right)^\gamma t_{\text{rel}} \log n\right) \cdot \mathcal{O}\left(\left(\frac{n}{\log^C n}\right)^{1-\gamma} \cdot \log n\right) \\ &= \mathcal{O}\left(n \cdot \left(\frac{d_{\text{avg}}}{d_{\text{min}}}\right)^\gamma \cdot t_{\text{rel}} \cdot \frac{\log^2 n}{\log^{C(1-\gamma)} n}\right). \end{aligned} \tag{7}$$

For the first bound, we choose $C = \log\left(\left(\frac{d_{\text{avg}}}{d_{\text{min}}}\right) \cdot t_{\text{rel}} \cdot \log^2 n\right) / ((1 - \gamma) \cdot \log \log n)$ then since $\log^{C(1-\gamma)} n = (d_{\text{avg}}/d_{\text{min}})t_{\text{rel}} \cdot \log^2 n$ and $\gamma < 1$ this gives $\tau_2 = \mathcal{O}(n)$ by (7) above. Since in any graph $t_{\text{hit}} = \Omega(n)$,^a the total time is therefore $\mathcal{O}(\tau_1)$, and for this value of C we have

$$\tau_1 = \mathcal{O}\left(\frac{\log\left(\left(\frac{d_{\text{avg}}}{d_{\text{min}}}\right) \cdot t_{\text{rel}} \cdot \log^2 n\right)}{(1 - \gamma) \cdot \log \log n} t_{\text{hit}} \log \log n\right) = \mathcal{O}\left(t_{\text{hit}} \cdot \log(d_{\text{max}}) \cdot \log\left(\frac{d_{\text{avg}} \cdot t_{\text{rel}} \cdot \log n}{d_{\text{min}}}\right)\right);$$

since $1 - \gamma = \Theta(1/\log d_{\text{max}})$. □

Proof of Theorem 6.2. Write $T = 4 \cdot t_{\text{rel}} \cdot \ln n$. For any $x \in V$ and $S \subset V$, Lemma 6.5 gives a $t \leq T$ such that $p_{x,S}^{(t)} \geq \pi(S)/3$, and Theorem 6.3 consequently gives a strategy for the CRW such that $q_{x,y}^{(t)} \geq (\pi(y)/3)^\gamma$. Thus, for any target set S and start vertex x , we need in expectation at most $(3/\pi(S))^\gamma$ attempts to hit S in at most T steps, since if an attempt fails, ending at some vertex z , we have the same bound on the probability of hitting S from z . Therefore, there is a strategy for the CRW where the hitting time $H_x(S) \leq 12 \cdot \pi(S)^{-\gamma} \cdot t_{\text{rel}} \log n$. The second result follows since for any vertex $\pi(v) \geq \frac{d_{\text{min}}}{nd_{\text{avg}}}$. □

6.2 The max choice and min choice operations

In this section, we introduce two operators which represent the effect of making optimal choices for a single step of the random walk, assuming that the effects of choice on future steps are already known and prove inequalities relating them to power means.

Define the max choice operator $\text{MC}_2: [0, \infty)^m \rightarrow [0, \infty)$ as follows:

$$\text{MC}_2(x_1, \dots, x_d) = \frac{1}{m^2} \sum_{i=1}^d \sum_{j=1}^d \max\{x_i, x_j\}. \tag{8}$$

For $p \in \mathbb{R} \setminus \{0\}$, the p -power mean M_p of non-negative reals x_1, \dots, x_m is defined by:

$$M_p(x_1, \dots, x_m) = \left(\frac{x_1^p + \dots + x_m^p}{m}\right)^{1/p}.$$

We use a key lemma which could be described as a multivariate anti-convexity inequality.

Lemma 6.7. *For any $1 \leq d \leq m$ and $x_1 \dots x_d \in [0, 1]$, we have*

$$M_{\gamma_m^{-1}}(x_1, \dots, x_d) \leq \text{MC}_2(x_1, \dots, x_d).$$

Proof. By the power-mean inequality, since $\gamma_m^{-1} \leq \gamma_d^{-1}$ it is sufficient to prove the case $m = d$. We show this by induction on d ; we have equality for $d = 1$. Suppose that either $d = 2$ or $d \geq 3$ and the result holds for $d - 1$. Without loss of generality, using symmetry and homogeneity of both operators, we may assume that $\max\{x_1, \dots, x_d\} = x_d = 1$.

We first claim that we may further assume $x_1 = \dots = x_{d-1}$. If $d = 2$, this claim is trivial. If $d \geq 3$, then write $\bar{x} = M_{\gamma_d^{-1}}(x_1, \dots, x_{d-1})$. Note that

$$M_{\gamma_d^{-1}}(x_1, \dots, x_{d-1}, x_d) = M_{\gamma_d^{-1}}(\bar{x}, \dots, \bar{x}, x_d).$$

^aLet τ_v be the first time that v is visited during a random walk from u . Then $\sum_{v \neq u} \tau_v \geq \sum_{i=1}^{n-1} i = \binom{n}{2}$, since each τ_v is distinct. Thus, $t_{\text{hit}} \geq \sum_{v \neq u} H_u(v)/n \geq (n - 1)/2$.

Also we have

$$\begin{aligned} \text{MC}_2(x_1, \dots, x_{d-1}, x_d) &= \frac{2d-1}{d^2}x_d + \left(\frac{d-1}{d}\right)^2 \text{MC}_2(x_1, \dots, x_{d-1}) \\ &\geq \frac{2d-1}{d^2}x_d + \left(\frac{d-1}{d}\right)^2 M_{\gamma_{d-1}^{-1}}(x_1, \dots, x_{d-1}) \\ &\geq \frac{2d-1}{d^2}x_d + \left(\frac{d-1}{d}\right)^2 \bar{x} \\ &= \text{MC}_2(\bar{x}, \dots, \bar{x}, x_d), \end{aligned}$$

where the first inequality uses the assumption that the result holds for $n - 1$ and the second uses the power-mean inequality. Thus, replacing x_1, \dots, x_{d-1} by \bar{x}, \dots, \bar{x} does not increase the difference between the two operators, proving the claim.

Next, we claim that the function $f(x) = M_{\gamma_d^{-1}}(x, \dots, x, 1)$ is convex. Since $\text{MC}_2(x, \dots, x, 1)$ is linear, and the two functions agree at 0 (by choice of γ_d) and at 1, this will complete the proof. Indeed, we have $f(x)^{\gamma_d^{-1}} = \frac{d-1}{d}x^{\gamma_d^{-1}} + \frac{1}{d}$, giving $f'(x) = \frac{d-1}{d} \left(\frac{x}{f(x)}\right)^{\gamma_d^{-1}-1}$. Also, $\frac{x^{\gamma_d^{-1}}}{f(x)^{\gamma_d^{-1}}}$ is an increasing function of x ; since $\gamma_d^{-1} - 1 > 0$, we have $f'(x)$ is increasing, as required. \square

Lemma 6.7 will be used to prove Theorem 6.3. In order to prove Theorem 6.4, we will need a corresponding inequality for an appropriate operator. To that end, we define the *min choice* operator $\text{mC}_2 : [0, \infty)^m \rightarrow [0, \infty)$ by:

$$\text{mC}_2(x_1, \dots, x_m) = \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m \min\{x_i, x_j\}.$$

Lemma 6.8. *For any $m \geq 1$ and non-negative reals x_1, \dots, x_m , we have*

$$\text{mC}_2(x_1, \dots, x_m) \leq M_{1/2}(x_1, \dots, x_m).$$

Proof. Observe that

$$\begin{aligned} \left(\frac{1}{m} \sum_{i=1}^m \sqrt{x_i}\right)^2 &= \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m \sqrt{x_i x_j} \\ &\geq \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m \min\{x_i, x_j\}. \end{aligned} \quad \square$$

6.3 The tree gadget for graphs

In this section, we prove Theorem 6.3. To achieve this, we introduce the *tree gadget* which encodes trajectories of length at most t from u in a rooted graph (G, u) by vertices of an arborescence $(\mathcal{T}_t, \mathbf{r})$, that is, a tree with all edges oriented away from the root \mathbf{r} . Given (G, u) , we represent each trajectory of length $i \leq t$ started from u in G as a node at distance i from the root \mathbf{r} in the tree \mathcal{T}_t . The root \mathbf{r} represents the trajectory of length 0 from u . There is an edge from \mathbf{x} to \mathbf{y} in \mathcal{T}_t if \mathbf{x} is obtained from \mathbf{y} by deleting the final vertex. See Figure 3 for an illustration of the tree gadget.

Also for $\mathbf{x} \in V(\mathcal{T}_t)$ let $\Gamma^+(\mathbf{x}) = \{\mathbf{y} \in V(\mathcal{T}_t) : \mathbf{x}\mathbf{y} \in E(\mathcal{T}_t)\}$ be the offspring of \mathbf{x} in \mathcal{T}_t ; as usual we write $d^+(\mathbf{x})$ for the number of offspring. Write $|\mathbf{x}|$ for the length of the trajectory \mathbf{x} . To prove

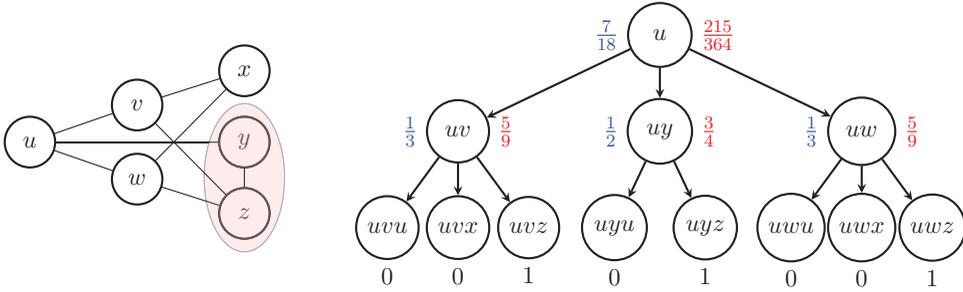


Figure 3. Illustration of a (non-lazy) walk on a non-regular graph starting from u with the objective of being in $\{y, z\}$ at step $t = 2$. The probabilities of achieving this are given in blue (left) for the SRW and in red (right) for the CRW.

Theorem 6.3, we shall need to discuss SRW trajectories; let $W_u(k) := (X_i)_{i=0}^k$ be the trajectory of a SRW X_i on G up to time k , with $X_0 = u$.

Proof of Theorem 6.3. For ease of notation, we write $\eta = 1/\gamma_{d_{\max}}$. To each node \mathbf{x} of the tree gadget \mathcal{T}_t , we assign the value $q_{\mathbf{x},S}$ under the CRW strategy of preferring the choice which extends to a trajectory $\mathbf{y} \in \Gamma^+(\mathbf{x})$ giving a higher value of $q_{\mathbf{y},S}$. This is well defined because both the strategy and the values $q_{\mathbf{x},S}$ can be computed in a ‘bottom up’ fashion starting at the leaves, where if $\mathbf{x} \in V(\mathcal{T}_t)$ is a leaf then $q_{\mathbf{x},S}$ is 1 if $\mathbf{x} \in S$ and 0 otherwise.

Suppose \mathbf{x} is not a leaf. The controller is presented with two uniformly random offspring $\mathbf{y}, \mathbf{z} \in \Gamma^+(\mathbf{x})$ and chooses \mathbf{y} if $q_{\mathbf{y},S} \geq q_{\mathbf{z},S}$ and \mathbf{z} otherwise. Thus, we have

$$q_{\mathbf{x},S} = \frac{1}{d^+(\mathbf{x})^2} \sum_{\mathbf{y}, \mathbf{z} \in \Gamma^+(\mathbf{x})} \max\{q_{\mathbf{y},S}, q_{\mathbf{z},S}\} = \text{MC}_2 \left((q_{\mathbf{y},S})_{\mathbf{y} \in \Gamma^+(\mathbf{x})} \right). \tag{9}$$

We define the following potential function $\Phi^{(i)}$ on the i^{th} generation of the tree gadget \mathcal{T} :

$$\Phi^{(i)} = \sum_{|\mathbf{x}|=i} q_{\mathbf{x},S}^\eta \cdot \mathbb{P} [W_u(i) = \mathbf{x}]; \tag{10}$$

where the sum ranges over all trajectories \mathbf{x} of length i . Notice that if $\mathbf{xy} \in E(\mathcal{T}_t)$ then

$$\mathbb{P} [W_u(|\mathbf{y}|) = \mathbf{y}] = \mathbb{P} [W_u(|\mathbf{x}|) = \mathbf{x}] / d^+(\mathbf{x}).$$

Also since each \mathbf{y} with $|\mathbf{y}| = i$ has exactly one parent \mathbf{x} with $|\mathbf{x}| = i - 1$, we can write

$$\Phi^{(i)} = \sum_{|\mathbf{x}|=i-1} \sum_{\mathbf{y} \in \Gamma^+(\mathbf{x})} q_{\mathbf{y},S}^\eta \cdot \frac{\mathbb{P} [W_u(i-1) = \mathbf{x}]}{d^+(\mathbf{x})}. \tag{11}$$

We now show that $\Phi^{(i)}$ is non-increasing in i . By combining (10) and (11), we can see that the difference $\Phi^{(i-1)} - \Phi^{(i)}$ is given by:

$$\sum_{|\mathbf{x}|=i-1} \left(q_{\mathbf{x},S}^\eta - \frac{1}{d^+(\mathbf{x})} \sum_{\mathbf{y} \in \Gamma^+(\mathbf{x})} q_{\mathbf{y},S}^\eta \right) \mathbb{P} [W_u(i-1) = \mathbf{x}].$$

Recalling (9), to establish $\Phi^{(i-1)} - \Phi^{(i)} \geq 0$, it is sufficient to show the following inequality holds whenever \mathbf{x} is not a leaf:

$$\text{MC}_2 \left((q_{\mathbf{y},S})_{\mathbf{y} \in \Gamma^+(\mathbf{x})} \right)^\eta \geq \frac{1}{d^+(\mathbf{x})} \sum_{\mathbf{y} \in \Gamma^+(\mathbf{x})} q_{\mathbf{y},S}^\eta.$$

Raising both sides to the power $1/\eta = \gamma_{d_{\max}}$, since $d^+(\mathbf{x}) \leq d_{\max}$ this inequality holds by Lemma 6.7, and thus $\Phi^{(i)}$ is non-increasing in i .

Observe $\Phi^{(0)} = q_{u,S}^\eta$. Also if $|\mathbf{x}| = t$ then $q_{\mathbf{x},S} = 1$ if $\mathbf{x} \in S$ and 0 otherwise. It follows that

$$\Phi^{(t)} = \sum_{|\mathbf{x}|=t} q_{\mathbf{x},S}^\eta \cdot \mathbb{P} [W_u(t) = \mathbf{x}] = \sum_{|\mathbf{x}|=t} \mathbf{1}_{\{\mathbf{x} \in S\}} \cdot \mathbb{P} [W_u(t) = \mathbf{x}] = p_{u,S}.$$

Thus, since $\Phi^{(t)}$ is non-increasing $q_{u,S}^\eta = \Phi^{(0)} \geq \Phi^{(t)} = p_{u,S}$, as required. □

Theorem 6.4 now follows similarly to Theorem 6.3.

Proof of Theorem 6.4. Construct the tree gadget to height t . We associate each node \mathbf{x} with the probability $q_{\mathbf{x},S}$ under a strategy which always prefers the smaller value. For a leaf this is simply the indicator function $\mathbf{1}_{\{\mathbf{x} \in S\}}$, whereas for an internal vertex it is given by $mC_2((q_{\mathbf{y},S})_{\mathbf{y} \in \Gamma^+(\mathbf{x})})$. We define a potential function Ψ by:

$$\Psi^{(i)} = \sum_{|\mathbf{x}|=i} \mathbb{P} [W_u(i) = \mathbf{x}] \sqrt{q_{\mathbf{x},S}}.$$

As before,

$$\Psi^{(t)} = \sum_{|\mathbf{x}|=t} \mathbb{P} [W_u(t) = \mathbf{x}] \cdot \mathbf{1}_{\{\mathbf{x} \in S\}} = p_{u,S}.$$

Further, for each internal vertex \mathbf{x} we have, using Lemma 6.8,

$$\begin{aligned} \mathbb{P} [W_u(|\mathbf{x}|) = \mathbf{x}] \sqrt{q_{\mathbf{x},S}} &= \mathbb{P} [W_u(|\mathbf{x}|) = \mathbf{x}] \sqrt{mC_2((q_{\mathbf{y},S})_{\mathbf{y} \in \Gamma^+(\mathbf{x})})} \\ &\leq \mathbb{P} [W_u(|\mathbf{x}|) = \mathbf{x}] \sqrt{M_{1/2}((q_{\mathbf{y},S})_{\mathbf{y} \in \Gamma^+(\mathbf{x})})} \\ &= \sum_{\mathbf{y} \in \Gamma^+(\mathbf{x})} \mathbb{P} [W_u(|\mathbf{x}|) = \mathbf{y}] \sqrt{q_{\mathbf{y},S}}. \end{aligned}$$

Summing over all \mathbf{x} at level i , we obtain $\Psi^{(i)} \leq \Psi^{(i+1)}$ for each $i < t$, and consequently $\sqrt{q_{u,S}} = \Psi^{(0)} \leq \Psi^{(t)} = p_{u,S}$, as required. □

6.4 Random graphs

We now consider CRW hitting and cover times in the Erdős–Rényi random graph $\mathcal{G}(n, p)$. This is the probability distribution over all n -vertex simple graphs generated by sampling each possible edge independently with probability p , see [12] for more details.

Theorem 1.5 *Let $\mathcal{G} \stackrel{d}{\sim} \mathcal{G}(n, p)$ where $np \geq c \ln n$ for any fixed $c > 1$ and $\log np = o(\log n)$. Then w.h.p.*

- (i) $t_{\text{cov}}^{\text{two}}(\mathcal{G}) = \mathcal{O}(n \cdot \log(np) \cdot \log \log n)$
- (ii) $t_{\text{hit}}^{\text{two}}(\mathcal{G}) = n^{1-\Omega(1/\log(np))}$.

Proof. To begin, we show that the graph is almost regular w.h.p.

Claim *For p as above, $d_{\min}, d_{\max} = \Theta(np)$ w.h.p.*

Proof of claim. In $\mathcal{G} \stackrel{d}{\sim} \mathcal{G}(n, p)$ since each edge is independent with probability p , each degree $d(u)$ is distributed as a binomial random variable $\text{Bin}(n-1, p)$. The Chernoff bound

[14, Thm. 3.2] states that for any λ , $\mathbb{P} [\text{Bin} (n, p) \geq np + \lambda] \leq \exp \left(-\frac{\lambda^2}{2(np + \lambda/3)} \right)$. Thus, by a union bound over all vertices $d_{\max} \leq 5np$ w.h.p.. For d_{\min} note that the expected number of vertices of degree k is given by $x_k = n \binom{n-1}{k} p^k (1-p)^{n-1-k}$. We shall consider $k = \kappa np$ for $\kappa \leq 1/2$, in this case $x_k/x_{k-1} = \frac{np}{k(1-p)} \geq 2$ and so the expected number of vertices with degree k is $\mathcal{O}(x_k)$. Observe that $x_k \leq \left(\frac{np}{k}\right)^k \exp(\ln n - np)$ and so if $np \geq 5 \log n$ then choosing $k = np/e$ yields $x_k \leq (e^2)^{np/e} \exp(\log n - np) = e^{-\Omega(\log n)}$. Otherwise since $np \geq c \ln n$ where $c > 1$ fixed, setting $k = \kappa np$ we have $x_k \leq \left(\frac{ce}{\kappa}\right)^{\kappa \ln n} \exp(- (c-1) \ln n)$, thus if κ satisfies $\kappa \ln \left(\frac{ce}{\kappa}\right) < (c-1)$ then $x_k = o(1)$. Choosing $\kappa = (c-1)^2/500$ suffices. Since in either case, we showed $x_k = o(1)$ for some $k = \Omega(np)$ by Markov's inequality $d_{\min} = \Omega(np)$ w.h.p. \diamond

Cooper & Frieze [16] show that for $np = c \ln n$, $c > 1$ w.h.p. the conductance of $G(n, p)$ is at least $1/6$, implying that $t_{\text{rel}} = \mathcal{O}(1)$ [28, Thm. 13.14]. For larger values of np , Coja-Oghlan [15, Thm. 1.2] showed that there exists some $c < \infty$ such that for $np \geq c \log n$ the spectral gap of the normalised Laplacian of $\mathcal{G}(n, p)$ is $1 - \mathcal{O}(1/\sqrt{np})$ w.h.p. Since the normalised Laplacian \mathcal{L} is similar to the random walk Laplacian \mathcal{L}' , and the later is given by $\mathcal{L}' = I - P$, we see that also in this range $t_{\text{rel}} = \mathcal{O}(1)$. We have shown that, in this regime, $\mathcal{G}(n, p)$ is almost regular and has constant relaxation time w.h.p., thus $t_{\text{hit}} = \mathcal{O}(n)$ w.h.p. by [13, Thm. 5.2]. Theorems 6.1 & 6.2 now yield the results. \square

Thus, the CRW gives a significant improvement in the cover and hitting times whenever degrees of $\mathcal{G}(n, p)$ are subpolynomial in n .

7. Computing optimal choice strategies

In this section, we focus on the following problem: given a graph G and an objective, how can we compute a series of choices for the walk which achieves the given objective in optimal expected time? In particular, we consider the following computational problems related to our main objectives of max/minimising hitting times, cover times and stationary probabilities π_ν .

Stat (G, w): Find a CRW strategy min/maximising $\sum_{\nu \in V} w_\nu \pi_\nu$ for vertex weights $w_\nu \geq 0$.

Hit (G, ν, S): Find a CRW strategy minimising $H_\nu^{\text{two}}(S)$ for a given $S \subseteq V(G)$ and $\nu \in V(G)$.

Cov (G, ν): Find a CRW strategy minimising $C_\nu^{\text{two}}(G)$ for a given $\nu \in V(G)$.

The analogous problems to **Stat** (G, w) and **Hit** (G, ν, S) were studied in [4] for the BRW. While **Stat** is not one of our primary objectives, we include it here both as a natural problem to consider but also because of its relationship to **Hit** in the case where w is the indicator function of a set S ; we shall abuse notation by writing **Stat**(G, S) for this case. Clearly for **Stat**, we must restrict ourselves to unchanging strategies for the stationary probabilities π_ν to be well defined; we shall show that **Hit** also has an unchanging optimal strategy.

For **Hit** and **Cov**, there are two possible interpretations of what it means to ‘find’ a CRW strategy. Perhaps the most natural is to compute a sequence of optimal choices in an online fashion, that is at each time step to compute which of the two offered choices to accept. For any particular walk, with suitable memoisation, at most a polynomial number of such computations will be required for either problem: which choice to accept depends only on the current vertex, the two choices, and in the case of **Cov** the vacant set, which can change at most n times. We might alternatively want to compute a complete optimal strategy in advance; for **Hit** this requires only a polynomial number of single-choice computations, but for **Cov** the number of possible situations

our strategy must cover will be exponential. However, we shall show that Cov is hard even for individual choices.

7.1 A polynomial-time algorithm for Stat and Hit

First, we show how the (unknown) optimal values $H_x^{two}(v)$ determine an optimal strategy for $\text{Hit}(G, \cdot, v)$. In the following two lemmas, we will need to work with a multigraph F ; in this context, the choice offered at each stage is between two random edges from the current vertex.

Lemma 7.1. *Let F be a multigraph and fix a vertex v . Let $v = v_0, v_1, \dots$ be an ordering of the vertices such that for all $i < j$, we have $H_{v_i}^{two}(v) \leq H_{v_j}^{two}(v)$. Let β be the deterministic unchanging strategy given by $\beta_{v_i, v_j}^{v_k} = 1$ whenever $j < k$. Then β is optimal (among all strategies) for $\text{Hit}(F, x, v)$ for every $x \neq v$, and also for the problem of minimising $\mathbb{E}_v [\tau_v^+]$.*

Proof. Fix an optimal strategy α for $\text{Hit}(F, x, v)$, and for each $y \in \Gamma(x)$ write q_y for the probability that the first step under this strategy is from x to y . Recall that $q_y = \sum_{z \in \Gamma(x)} \frac{2\alpha_{x,y}^z}{d(x)^2}$. Now given that the first step is at y , an optimal strategy for the remaining steps is precisely an optimal strategy for $\text{Hit}(F, y, v)$, and thus

$$H_x^{two}(v) = 1 + \sum_{y \in \Gamma(x)} q_y H_y^{two}(v).$$

Suppose there exist $y, z \in \Gamma(x)$ with $H_y^{two}(v) < H_z^{two}(v)$ but $\alpha_{x,y}^z < 1$ at the first step. By instead (at time 1 only) always choosing y in preference to z , the expected hitting time is decreased by $\frac{2}{d(x)^2}(1 - \alpha_{x,y}^z)(H_z^{two}(v) - H_y^{two}(v))$, a contradiction. Thus, we have $\alpha_{x,y}^z = 1$ if $H_y^{two}(v) < H_z^{two}(v)$ and $\alpha_{x,y}^z = 0$ if $H_y^{two}(v) > H_z^{two}(v)$. If $H_y^{two}(v) = H_z^{two}(v)$, then the expected hitting time does not depend on $\alpha_{x,y}^z$, and so any strategy satisfying these conditions at time 1, and thereafter following an optimal strategy, is itself optimal.

It follows by induction that following β for k turns and thereafter following α is optimal; since this gives arbitrarily good approximations of the expected hitting time under β , β is itself optimal for $\text{Hit}(F, x, v)$, and, since the definition of β does not depend on x , for $\text{Hit}(F, y, v)$ for any $y \neq v$.

Next, we show that β is also an optimal strategy for minimising $\mathbb{E}_v [\tau_v^+]$. Suppose not, and let γ be an optimal strategy. Write q_x^γ for the probability of moving from v to x at time 1 under γ , and $H_v^\gamma(v^+)$ for $\mathbb{E}_v [\tau_v^+]$ under γ . Now

$$\begin{aligned} H_v^\gamma(v^+) &= 1 + \sum_{x \in \Gamma(v)} q_x^\gamma H_x^\gamma(v) \\ &\geq 1 + \sum_{x \in \Gamma(v)} q_x^\gamma H_x^\beta(v), \end{aligned}$$

by optimality of β for $\text{Hit}(F, x, v)$. Suppose $\gamma_{v,x}^\gamma \neq \beta_{v,x}^\gamma$ for some $x, y \in \Gamma(v)$. Replacing $\gamma_{v,x}^\gamma$ and $\gamma_{v,y}^\gamma$ by $\beta_{v,x}^\gamma$ and $\beta_{v,y}^\gamma$, respectively, changes $\sum_{x \in \Gamma(v)} q_x^\gamma H_x^\beta(v)$ by $\frac{2}{d(v)^2}(\beta_{v,x}^\gamma - \gamma_{v,x}^\gamma)(H_x^{two}(v) - H_y^{two}(v))$, which is non-positive by choice of β . Thus, after a sequence of such changes, we obtain

$$\begin{aligned} H_v^\gamma(v^+) &\geq 1 + \sum_{x \in \Gamma(v)} q_x^\gamma H_x^\beta(v) \\ &= H_v^\beta(v^+). \end{aligned} \quad \square$$

Remark. In particular, recalling that for an unchanging strategy $\pi_v = 1/\mathbb{E}_v[\tau_v^+]$, it follows that β is an optimal strategy for $\text{Stat}(F, \{v\})$. However, this is true in a somewhat stronger sense, since optimality for Stat only requires minimising $\mathbb{E}_v[\tau_v^+]$ among unchanging strategies, whereas Lemma 7.1 shows that β minimises this quantity among all strategies; we shall need this extra strength.

Note that there may be other deterministic unchanging optimal strategies for $\text{Hit}(F, x, v)$. For example, if there are multiple vertices with the same optimal hitting time, we may choose between them arbitrarily, and in particular may have a cyclic order of preference which is not consistent with any single ordering. The following lemmas will enable us to show that a good enough approximation to an optimal strategy must itself be optimal.

Lemma 7.2. *Let F be a multigraph with at most n vertices and at most $\binom{n}{2}$ edges and fix a vertex v . Let α be any unchanging strategy for $\text{Stat}(F, \{v\})$. Suppose there exist vertices x, y, z with $y, z \in \Gamma(x)$, $H_y^{\text{two}}(v) < H_z^{\text{two}}(v)$ and $\alpha_{x,y}^z \leq 1/2$. Then π_v^α differs from the optimal value by at least $n^{-4(n+1)}(H_z^{\text{two}}(v) - H_y^{\text{two}}(v))$.*

Proof. First, we bound $H_v^\alpha(v^+) - H_v^\beta(v^+)$, where β is as described in Lemma 7.1. Consider the strategy of following α until the first time the walk either reaches v or is at x and offered a choice between y and z , and in the latter case following β until v is reached. The difference between this strategy and following α is $p(\alpha_{x,y}^z H_y^\alpha(v) + \alpha_{x,z}^y H_z^\alpha(v) - H_y^\beta(v))$, where p is the probability of the latter event occurring before the walk returns to v . Note that

$$\begin{aligned} \alpha_{x,y}^z H_y^\alpha(v) + \alpha_{x,z}^y H_z^\alpha(v) - H_y^\beta(v) &\geq (\alpha_{x,y}^z - 1)H_y^\beta(v) + \alpha_{x,z}^y H_z^\beta(v) \\ &= (1 - \alpha_{x,y}^z)(H_z^{\text{two}}(v) - H_y^{\text{two}}(v)) \\ &\geq (H_z^{\text{two}}(v) - H_y^{\text{two}}(v))/2 \end{aligned}$$

by Lemma 7.1 and the assumptions. Further,

$$p \geq 2 \left(\frac{1}{d_{\max}(F)^2} \right)^{d(v,x)+1} \geq \binom{n}{2}^{-2n},$$

since with at least this probability, the walk is forced along a specific shortest path to x , then offered a choice of y or z .

Thus, the difference in $\mathbb{E}_v[\tau_v^+]$ between α and this hybrid strategy is at least

$$\zeta := \frac{1}{2} \binom{n}{2}^{-2n} (H_z^{\text{two}}(v) - H_y^{\text{two}}(v)),$$

and since β minimises this quantity among all strategies by Lemma 7.1, the same bound applies to the difference between α and β , giving

$$\pi_\alpha(v)^{-1} \geq \pi_\beta(v)^{-1} + \zeta,$$

and consequently

$$\pi_\alpha(v) \leq \pi_\beta(v) - \zeta \frac{\pi_\beta(v)^2}{1 + \pi_\beta(v)\zeta}. \tag{12}$$

We have $1 \geq \pi_\beta(v) \geq \binom{n}{2}^{-1}$ by comparison with a SRW. Also we may crudely bound $t_{\text{hit}}^{\text{two}}$ by noting that a SRW has probability at least $\binom{n}{2}^{1-n}$ of reaching any given vertex in at most $n - 1$ steps, giving $\zeta < 1$. Combining these bounds with (12) gives the required result. \square

Lemma 7.3. For any simple graph G of order n and every pair of vertices x, y with $H_x^{\text{two}}(S) < H_y^{\text{two}}(S)$ we have $H_y^{\text{two}}(S) - H_x^{\text{two}}(S) > n^{-2n^2}$.

Proof. Note that the hitting times $(h_x)_{x \in V}$ of S from x for any given unchanging strategy are uniquely determined by the equations:

$$h_x = \begin{cases} 1 + \sum_y \mathbf{P}_{xy} \cdot h_y & \text{if } x \notin S \\ 0 & \text{if } x \in S, \end{cases}$$

where \mathbf{P} is the transition matrix for the strategy. This set of equations can be written as $\mathbf{A}\mathbf{h} = \mathbf{b}$, where $\mathbf{A} := (\mathbf{I} - \mathbf{Q})$, $\mathbf{Q}_{ij} = \mathbf{P}_{ij}$ if $i \notin S$ and 0 otherwise, and \mathbf{b} is a 0-1 vector. Notice that \mathbf{A} is diagonally dominant, and from any row where equality occurs there is a path of non-zero entries to a strictly dominant row. It is straightforward to check that such a matrix is invertible: see for example [6, Lem. 3.2]. For any non-random strategy, and in particular for the optimal strategy described above, every transition probability from x is a multiple of $d(x)^{-2}$. Thus, all the elements of \mathbf{A} can be put over a common denominator D , where $D := \text{LCM}(d(x)^2)_{x \in V} < (n!)^2 < n^{2n}/2$.

We have $\mathbf{h} = \mathbf{A}^{-1}\mathbf{b} = |\mathbf{A}|^{-1}\mathbf{C}^T\mathbf{b}$, where \mathbf{C} is the matrix of cofactors. Each entry in \mathbf{C} can be put over a common denominator which is at most D^n , and so the same applies to each entry of $\mathbf{C}^T\mathbf{b}$. Also, $|\mathbf{A}| < 2^n$ by Hadamard’s inequality [26, Thm. 7.8.1]. It follows that if two hitting times differ, they differ by at least $(2D)^{-n}$. □

For any graph G and weighting $w:V \rightarrow [0, \infty)$ on the vertices of G we can phrase $\text{Stat}(G, w)$ as an optimisation problem as follows, where we shall encode our actions using the probabilities $\alpha_{x,y}^z = \mathbb{P}[X_{t+1} = y \mid X_t = x, c = \{y, z\}]$ from Section 2.

$$\begin{aligned} \text{maximise: } & \sum_{v \in V} w_v \pi(v) \\ \text{subject to: } & \pi(x) = \sum_{y \in \Gamma(x)} \pi(y) \cdot \frac{2 \sum_{z \in \Gamma(y)} \alpha_{y,x}^z}{d(x)^2}, \quad \forall x \in V \\ & \sum_{x \in V} \pi(x) = 1, \\ & \alpha_{x,z}^y \in [0, 1], \quad \forall xz, xy \in E \\ & \alpha_{x,z}^y = 1 - \alpha_{x,y}^z, \quad \forall xz, xy \in E \end{aligned} \tag{13}$$

For minimising the stationary probabilities, we maximise -1 times the objective function.

To prove Theorem 7.4, the quadratic terms in (13) can be eliminated using the same substitution as [4, Thm. 6]; we can then solve (13) as a linear program.

Theorem 7.4. For any multigraph G and weight function $w:V \rightarrow [0, \infty)$ a policy solving the problem $\text{Stat}(G, w)$ to within an additive ε factor can be computed in time $\text{poly}(|E|, \log(1/\varepsilon))$.

Proof. We prove the simple graph case; this proof may be easily extended for multigraphs with suitably adapted notation. The optimisation problem (13) above can be rephrased as a linear program by making the substitution $r_{x,y,z} = \pi(x) \cdot \alpha_{x,y}^z$. Either the Ellipsoid method or Karmarkar’s algorithm will approximate the solution to within an additive $\varepsilon > 0$ factor in time which is polynomial in the dimension of the problem and $\log(1/\varepsilon)$, see for example [23, 27]. □

We now show how one can now use this linear program to determine the hitting times.

Theorem 1.6 *For any graph G and any $S \subset V$, a solution to $\text{Hit}(G, x, S)$ for every $x \in V \setminus S$ can be computed in time $\text{poly}(n)$.*

Proof. Contract S to a single vertex v to obtain a multigraph F ; where a vertex x has more than one edge to S in G , retain multiple edges between x and v in F . Note that F has at most n vertices and at most $\binom{n}{2}$ edges. Provided that the CRW on G has not yet reached S , there is a natural correspondence between strategies on G and F with the same transition probabilities, and it follows that $H_x^{\text{two}}(S)$ for G and $H_x^{\text{two}}(v)$ for F are equal for any $x \in V(G) \setminus S$. We compute an optimal strategy to $\text{Stat}(F, \{v\})$ to within an additive error of $\varepsilon := n^{-10n^2}$; note that $\log(1/\varepsilon) = o(n^3)$ and so this may be done in time $\text{poly}(n)$ by Theorem 7.4. Applying Lemma 7.2 to F and Lemma 7.3 to G , using the equality of corresponding hitting times, implies that this strategy has $\alpha_{x,y}^z > 1/2$ whenever $H_y^{\text{two}}(v) < H_z^{\text{two}}(v)$, and so rounding each of the probabilities $\alpha_{x,y}^z$ to the nearest integer gives an optimal strategy (on F) for every x , which may easily be converted to an optimal strategy for G . \square

7.2 A hardness result for Cov

We show that in general even the online version of $\text{Cov}(G, v)$ is NP-hard. To that end, we introduce the following problem, which represents a single decision in the online version. The input is a graph G , a current vertex u , two vertices v and w which are adjacent to u , and a visited set X , which must be connected and contain u .

NextStep (G, u, v, w, X): Choose whether to move from u to either v or w so as to minimise the expected time for the CRW to visit every vertex not in X , assuming an optimal strategy is followed thereafter.

Any such problem may arise during a random walk with choice on G starting from any vertex in X , no matter what strategy was followed up to that point, since with positive probability no real choice was offered in the walk up to that point.

Theorem 1.7 *NextStep is NP-hard, even if G is constrained to have maximum degree 3.*

Proof. We give a (Cook) reduction from the NP-hard problem of either finding a Hamilton path in a given graph H or determining that none exists. This is known to be NP-hard even if H is restricted to have maximum degree 3 [21].

We shall find it more convenient to work with the following problem, which takes as input a graph G , a current vertex u and a connected visited set X containing u .

BestStep (G, u, X): Choose a neighbour of u to move to so as to minimise the expected time for the CRW to visit every vertex not in X , assuming an optimal strategy is followed thereafter.

We may solve $\text{BestStep}(G, u, X)$ by computing $\text{NextStep}(G, u, v, w, X)$ for every pair v, w of neighbours of u ; since all optimal neighbours must be preferred to all others, this will identify a set of one or more optimal choices for $\text{BestStep}(G, u, X)$. Consequently, it is sufficient to reduce the Hamilton path search problem to BestStep .

Given an n -vertex graph H , construct the graph G as follows. First replace each edge of H by a path of length $2cn^2$ through new vertices. Next add a new pendant path of length n^3 starting at the midpoint of each path corresponding to an edge of H . Finally, add edges to form a cycle consisting

of the end vertices of these pendant paths (in any order). Note that if H has maximum degree 3, so does G .

Fix a starting vertex u and a non-empty unvisited set $Y \subseteq V(H) \setminus \{u\}$ and set $X = V(G) \setminus Y$. (The purpose of the second and third stages of the construction is to make X connected without affecting the optimal strategy.) Suppose that H contains at least one path of length $|Y|$ starting at u which visits every vertex of Y ; in particular if $Y = V(H) \setminus \{u\}$ this is a Hamilton path of H . We claim that any optimal next step is to move towards the next vertex on some such path. Assuming the truth of this claim, an algorithm to find a Hamilton path starting at x , if one exists, is to set $u = x$ and $Y = V(H) \setminus \{x\}$, then find the vertex y such that moving towards y is optimal, set $u = y$ and remove y from Y , then continue. If this fails to find a Hamilton path, repeat for other possible choices of x .

To prove the claim, first we argue by induction that there is a strategy to visit every vertex in $|Y|$ in expected time $(4cn^2 + \mathcal{O}(n))|Y|$, where the implied constant does not depend on c . This is clearly true for $|Y| = 0$. Let y be the next vertex on a suitable path in H , and let z be the middle vertex of the path corresponding to the edge uy . Attempting to reach z by a straightforward strategy, the distance to z evolves as a random walk with probability $3/4$ of decreasing unless the current location is a branch vertex. We thus reach z in expected time $2cn^2$ plus an additional constant time for each visit to u , of which we expect $\mathcal{O}(d(u)) = \mathcal{O}(n)$, giving a total expected time of $2cn^2 + \mathcal{O}(n)$ (if the walker is forced to a different branch vertex first, the expected time to return from this point is polynomial in n , but this event occurs with exponentially small probability). Similarly, the time taken to reach y from z is $2cn^2 + \mathcal{O}(1)$. Once y is reached, there is (by choice of y) a path of length $|Y| - 1$ in H starting from y and visiting all of $Y \setminus \{y\}$. Thus, by induction, the required bound holds. Secondly, suppose that an optimal first step in a strategy from u moves towards a vertex y' of H which is not the first step in a suitable path. Since the expected remaining time decreases whenever an optimal step is taken, two successive optimal steps cannot be in opposite directions unless the walker visits a vertex of Y in between. Thus, the optimal strategy is to continue in the direction of y' if possible, and such a strategy reaches y' before returning to u with at least constant probability p , and this takes at least $2cn^2$ steps. Note that the expected time taken to reach another vertex of H from a vertex in H , even if the walker is purely trying to minimise this quantity, is at least $4cn^2$, and from either u or y' at least $|Y|$ such transitions are necessary to cover Y . Thus, such a strategy, conditioned on the first step being in the direction of y' , has expected time at least $4cn^2 + 2pcn^2$, which, for suitable choice of c , proves the claim. \square

7.3 Computing Cov via Markov decision processes

To compute a solution for $\text{Cov}(G, \nu)$, we can encode the cover time problem as a hitting time problem on a (significantly) larger graph.

Lemma 7.5. *For any graph $G = (V, E)$, let the (directed) auxiliary graph $\tilde{G} = (\tilde{V}, \tilde{E})$ be given by $\tilde{V} = \{(v, S) : S \subseteq V, v \in S\}$ and $\tilde{E} = \{((i, S), (j, S \cup j)) \mid ij \in E\}$. Then solutions to $\text{Cov}(G, \nu)$ correspond to solutions to $\text{Hit}(\tilde{G}, \tilde{\nu}, W)$ and vice versa, where $W = \{(u, V) \mid u \in V\}$.*

Proof. There is a natural bijection between the out-edges in G from u and those in \tilde{G} from (u, T) for any $u \in V, T \subseteq V$. This extends to a natural bijection from finite walks (which we may think of as a vertex together with a history) in G starting from ν to walks in \tilde{G} starting from $\tilde{\nu}$, and also to a measure-preserving bijection between the choices which may be offered from u and (u, T) . Thus, there is a natural bijection between strategies for the two walks, and both the choices offered and any random bits used may be coupled so that corresponding strategies produce corresponding walks. Since the walk in G has covered V if and only if the walk in \tilde{G} has hit some vertex in W , the times that these events first occur are identically distributed for corresponding strategies, and in particular the sets of optimal strategies correspond. \square

In light of Lemma 7.5, it may appear that we can solve $\text{Cov}(G, v)$ by converting it to an instance of $\text{Hit}(\tilde{G}, \tilde{v}, W)$ and appealing to Theorem 1.6. This is unfortunately not the case as \tilde{G} is a directed graph and Theorem 1.6 cannot handle directed graphs. Lemma 7.5 is still of use as we can phrase Hit in terms of Markov decision processes (MDPs) and then standard results tell us that an optimal strategy for the problem can be computed in finite time.

A MDP is a discrete time finite-state stochastic process controlled by a sequence of decisions [19]. At each step, a controller specifies a probability distribution over a set of actions which may be taken and this has a direct affect on the next step of the process. Costs are associated with each step/action and the aim of the controller is to minimise the total cost of performing a given task, for example, hitting a given state. In our setting, the actions are orderings of the vertices in each neighbourhood and the cost of each step/action is one unit of time. The problem $\text{Hit}(G, u, v)$ is then an instance of the optimal first passage problem which is known to be computable in finite time [19].

Corollary 7.6. *For any graph G and $v \in V$, an optimal policy for the problem $\text{Cov}(G, v)$ can be computed in exponential time.*

Proof. We first encode the problem $\text{Cov}(G, v)$ as the problem $\text{Hit}(\tilde{G}, \tilde{v}, W)$ as described in Lemma 7.5. Now as mentioned $\text{Hit}(\tilde{G}, \tilde{v}, W)$ is an instance of the optimal first passage problem which for a given graph \tilde{G} , start vertex \tilde{v} and target vertex W can be computed in finite time using either policy iteration or linear programming, see for example [19, Ch. 5, Cor. 1]. Examination of the linear program on [19, page. 58] reveals that there is a constraint for every ordering of the neighbours of each vertex. Since G has at most 2^n vertices and each of these has at most n neighbours, we see that there are at most $2^n \cdot n! \leq e^{n^3}$ constraints. It follows that this linear program can be solved in time $\text{poly}(e^{n^3})$, thus $\text{Cov}(G, v) \in \text{EXP}$. \square

Remark. Since in our setting actions are orderings of neighbourhoods, the space of actions may be factorial in the size of the graph. The algorithms for computing $\text{Hit}(G, u, v)$ from [19] used to establish Corollary 7.6 are polynomial in the number of actions and thus will not yield a polynomial-time algorithm for the problem. This is why we resisted appealing to MDP theory when finding a polynomial-time algorithm for $\text{Hit}(G, u, v)$ on undirected graphs in Section 7.1.

8. Summary

In this paper, we proposed a new random walk process inspired by the power of choice paradigm. We derived several quantitative bounds on the hitting and cover times and also presented a dichotomy with regard to computing optimal strategies.

While we were able to show that on an expander graph, the CRW significantly outperforms the SRW in terms of its cover time, we do not yet know the exact order of magnitude of $t_{\text{cov}}^{\text{two}}$. In fact, we do not have any lower bound on $t_{\text{cov}}^{\text{two}}$ improving the trivial $\Omega(n)$ for any sequence of bounded degree graphs. Constructing a sequence of graphs (G_n) , especially expanders, with $t_{\text{cov}}^{\text{two}}(G_n) = \omega(n)$ would be very interesting.

We have shown that $\text{Cov} \in \text{EXP}$ and that the problem is NP-hard. It would be interesting to find a complexity class for which the problem is complete, and we suspect it is PSPACE-complete.

Acknowledgements

We thank Sam Olesker-Taylor for many insightful comments on an earlier version of this work. A.G. and J.H. were supported by ERC Starting Grant no. 639046 (RGGC). J.H. was also partially

supported by UK Research and Innovation Future Leaders Fellowship MR/S016325/1. T.S. and J.S. were supported by ERC Starting Grant no. 679660 (DYNAMIC MARCH).

References

- [1] Achlioptas, D., D'Souza, R. M. and Spencer, J. (2009) Explosive percolation in random networks. *Science* **323** 1453–1455.
- [2] Aldous, D. and Fill, J. A. (2002) *Reversible Markov Chains and Random Walks on Graphs*. Unfinished monograph, recompiled 2014.
- [3] Avin, C. and Krishnamachari, B. (2008) The power of choice in random walks: An empirical study. *Comput. Netw.* **52** 44–60. (1) Performance of Wireless Networks (2) Synergy of Telecommunication and Broadcasting Networks.
- [4] Azar, Y., Broder, A. Z., Karlin, A. R., Linial, N. and Phillips, S. (1996) Biased random walks. *Combinatorica* **16** 1–18.
- [5] Azar, Y., Broder, A. Z., Karlin, A. R. and Upfal, E. (1999) Balanced allocations. *SIAM J. Comput.* **29** 180–200.
- [6] Azimzadeh, P. and Forsyth, P. A. (2016) Weakly chained matrices, policy iteration, and impulse control. *SIAM J. Numer. Anal.* **54** 1341–1364.
- [7] Beeler, K. E., Berenhaut, K. S., Cooper, J. N., Hunter, M. N. and Barr, P. S. (2014) Deterministic walks with choice. *Discrete Appl. Math.* **162** 100–107.
- [8] Berenbrink, P., Cooper, C. and Friedetzky, T. (2015) Random walks which prefer unvisited edges: Exploring high girth even degree expanders in linear time. *Random Struct. Algor.* **46** 36–54.
- [9] Berenbrink, P., Czumaj, A., Steger, A. and Vöcking, B. (2000) Balanced allocations: The heavily loaded case. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, ACM, pp. 745–754.
- [10] Bohman, T. and Frieze, A. (2002) Addendum to 'Avoiding a giant component' [*Random Struct. Algor.* **19** 75–85, 2001]. *Random Struct. Algor.* **20** 126–130.
- [11] Bohman, T. and Kravitz, D. (2006) Creating a giant component. *Combin. Probab. Comput.* **15** 489–511.
- [12] Bollobás, B. (2001) *Random Graphs*, 2nd edition, vol. 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press.
- [13] Chandra, A. K., Raghavan, P., Ruzzo, W. L., Smolensky, R. and Tiwari, P. (1996/1997) The electrical resistance of a graph captures its commute and cover times. *Comput. Complexity* **6** 312–340.
- [14] Chung, F. R. K. and Lu, L. (2006) Survey: Concentration inequalities and martingale inequalities: A survey. *Internet Math.* **3** 79–127.
- [15] Coja-Oghlan, A. (2007) On the Laplacian eigenvalues of $G_{n,p}$. *Combin. Probab. Comput.*, **16** 923–946.
- [16] Cooper, C. and Frieze, A. (2007) The cover time of sparse random graphs. *Random Struct. Algor.* **30**(1–2) 1–16.
- [17] Cooper, C., Frieze, A. M. and Johansson, T. (2018) The cover time of a biased random walk on a random cubic graph. In *29th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms, AofA 2018, June 25–29, 2018, Uppsala, Sweden*, Vol. 110 of *LIPICs* (J. A. Fill and M. D. Ward, eds), Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 16:1–16:12.
- [18] Cooper, C., Ilcinkas, D., Klasing, R. and Kosowski, A. (2009) Derandomizing random walks in undirected graphs using locally fair exploration strategies. In *Automata, Languages and Programming* (S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. Nikolettseas and W. Thomas, eds), Springer Berlin Heidelberg, pp. 411–422.
- [19] Derman, C. (1970) *Finite State Markovian Decision Processes*. Academic Press, Inc.
- [20] Ding, J., Lee, J. R. and Peres, Y. (2012) Cover times, blanket times, and majorizing measures. *Ann. Math. (2)* **175** 1409–1471.
- [21] Garey, M. R., Johnson, D. S. and Stockmeyer, L. (1976) Some simplified NP-complete graph problems. *Theoret. Comput. Sci.* **1** 237–267.
- [22] Georgakopoulos, A., Haslegrave, J., Sauerwald, T. and Sylvester, J. (2020) Choice and bias in random walks. In *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12–14, 2020, Seattle, Washington, USA*, Vol. 151 of *LIPICs* (T. Vidick, ed), Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 76:1–76:19.
- [23] Grötschel, M., Lovász, L. and Schrijver, A. (1993) *Geometric Algorithms and Combinatorial Optimization*, 2nd edition, Vol. 2 of *Algorithms and Combinatorics*. Springer-Verlag.
- [24] Haslegrave, J. and Jordan, J. (2016) Preferential attachment with choice. *Random Struct. Algor.* **48** 751–766.
- [25] Haslegrave, J., Sauerwald, T. and Sylvester, J. (2020) Time Dependent Biased Random Walks. Preprint, arXiv:2006.02475.
- [26] Horn, R. A. and Johnson, C. R. (2013) *Matrix Analysis*. 2nd edition, Cambridge University Press.
- [27] Karmarkar, N. (1984) A new polynomial-time algorithm for linear programming. *Combinatorica* **4** 373–395.
- [28] Levin, D. A., Peres, Y. and Wilmer, E. L. (2009) *Markov Chains and Mixing Times*. American Mathematical Society. With a chapter by James G. Propp and David B. Wilson.
- [29] Lyons, R. and Peres, Y. (2016) *Probability on Trees and Networks*, Vol. 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*, Cambridge University Press.

- [30] Malyshkin, Y. and Paquette, E. (2015) The power of choice over preferential attachment. *ALEA Lat. Am. J. Probab. Math. Stat.* **12** 903–915.
- [31] Mitzenmacher, M. (2001) The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distrib. Syst.* **12** 1094–1104.
- [32] Orenshtein, T. and Shinkar, I. (2014) Greedy random walk. *Combin. Probab. Comput.*, **23** 269–289.
- [33] Riordan, O. and Warnke, L. (2012) Achlioptas process phase transitions are continuous. *Ann. Appl. Probab.* **22** 1450–1464.
- [34] Spencer, J. and Wormald, N. (2007) Birth control for giants. *Combinatorica* **27** 587–628.