Time Dependent Biased Random Walks

John Haslegrave[1], Thomas Sauerwald[2], and John Sylvester[2]

[1]Mathematics Institute, University of Warwick
[2]Department of Computer Science & Technology, University of Cambridge

November 25, 2019

We study the biased random walk where at each step of a random walk a "controller" can, with a certain small probability, fix the next step. This model was introduced by Azar et al. [STOC1992]; we extend their work to the time dependent setting and consider cover times of this walk. We obtain new bounds on the cover and hitting times and make progress towards resolving a conjecture of Azar et al. on maximising values of the stationary distribution. We also consider the problem of computing an optimal strategy for the controller to minimise the cover time and show this is NP-hard.

1 Introduction

Randomised algorithms have come to occupy a central place within theoretical computer science and had a profound affect on the development of algorithms and complexity theory [14, 16]. Most randomised algorithms assume access to a source of unbiased independent random bits. In practice, however, truly independent unbiased random bits are inconvenient, if not impossible, to obtain. In practice we can generate pseudo-random bits on a computer fairly effectively [12] but if computational resources are constrained the quality of these bits may suffer, in particular they may be biased or correlated. Another reason to consider the dependency of randomised algorithms on the random bits they use, other than imperfect generation, is that an adversary may seek to tamper with a source of randomness to influence the output of a randomised algorithm. This raises the natural question of whether relaxing the unbiased and independent assumptions have a notable effect on the efficacy of randomised algorithms. This is a question many researchers have studied since early in the development of randomised algorithms [3, 6, 19].

Motivated by this question Azar, Broder, Karlin, Linial and Phillips [4] introduced the ε-biased random walk (ε-BRW). This process is a walk on a graph where at each step with probability ε a controller can choose a neighbour of the current vertex to move to, otherwise a uniformly random neighbour is selected. One can see this process from two different perspectives. The first interpretation is to see the model as a simple random walk with adversarial noise,

Some results from this paper will appear in *The 11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of LIPIcs, pages 76:1–76:20 [11]

that is with probability ε the random bits used to sample the next step become corrupted by an adversary. In particular one may consider this as a toy model for any randomised algorithms which uses their random bits to find a witness to the truth of a given statement, for example [2, 18], here the objective of the adversary may be to prevent us finding a witness. For the second interpretation on may view the ε-BRW as interpolating between a purely deterministic routing algorithm and a purely random one. In particular say our task is to reach a given vertex v, if $\varepsilon = 1$ and we have complete knowledge of the graph then can traverse a shortest path to v in time at most $|V|-1$, if instead $\varepsilon = 0$ then the ε-BRW performs a simple random walk (SRW) and the expected time to reach v is $\mathcal{O}(|V|^3)$. Otherwise if $0 < \varepsilon < 1$ then we have a mix of the two algorithms and our deterministic algorithm has to solve more difficult problem akin to the single-destination shortest path problem with additional complexity added by being perturbed by the random walk steps which occur with probability $1 - \varepsilon$.

Azar et al. consider two objectives for which a controller seeks to find an optimal strategy. The first objective is to maximise weighted sums of stationary probabilities and the second is to minimise the expected hitting time of a given set of vertices. The first problem can be thought as assigning a payoff for the controller to be in a given state and a optimal strategy maximises long term payoff. They were also interested in how much a controller can boost elements of the stationary distribution and obtained bounds on the stationary probabilities achievable. They also show that optimal strategies for maximising or minimising stationary probabilities or hitting times can be computed in polynomial time. One can show, by appealing to the theory of Markov Decision Processes (MDP) [8], that for both of these tasks there is an optimal strategy which is independent of time, and indeed the stationary probabilities are only well defined for time independent strategies. Thus Azar et al. only consider strategies which are fixed at the start of the process.

We extend the work of Azar et al. [4] by studying the cover time of ε-biased random walks, that is the expected time for the walk to visit every vertex of the graph. It is clear that for our problem an optimal solution will often depend on the set of vertices already covered, and thus unlike Azar et al. we are not in the time independent setting. To deal with this we introduce the ε-time biased random walk (ε-TBRW) which is the same as the ε-BRW except that the bias matrix may depend on the history of the process. Extending the second analogy for the hitting time problem for the ε-BRW to the cover time problem for the ε-TBRW, if $\varepsilon = 1$ then we can take a shortest walk visiting all vertices which has length at most $2(|V| - 1)$ and in the $\varepsilon = 0$ case we have the cover time by a simple random walk which again takes $\mathcal{O}(|V|^3)$ steps in expectation. In contrast in the $\varepsilon = 1$ case the controller must now solve an NP-Hard problem (shortest walk finds a Hamiltonian path if one exists) whereas before the controller solved the shortest path problem, which is achievable in poly time. This motivates the difficulty of finding an optimal strategy to cover the graph when $0 < \varepsilon < 1$.

1.1 Our Results

In Section 3 we introduce a method we call the tree gadget, this is a representation of all paths of a length of a given length from a fixed start vertex in a connected graph G by embedding them into a tree. We also introduce a symmetric operator on real vectors which describes the action of the ε-TBRW. The combination of the operator and tree gadget allows to us to show that the ε-TBRW can increase the probabilities of rare events described by paths, that is:

(1) Let $u \in V$, $t > 0$, $0 \leqslant \varepsilon \leqslant 1$ and S be a set of trajectories of length t from u. Then a controller can increase the probability of being in S after t steps from u from p to $p^{1-\varepsilon}$.

This result can be applied to bound cover and hitting times in terms of the number of vertices n, the maximum, minimum and average degrees $d_{\mathsf{max}}, d_{\mathsf{min}}$ and d_{avg}, and $t_{\mathsf{rel}} := \frac{1}{1-\lambda_2}$ which is the relaxation time:

(2) For any vertex u there is a strategy so that the ε-TBRW started from u covers G in expected time at most
$$\mathcal{O}\left(\frac{n}{\varepsilon} \cdot \frac{d_{\mathsf{avg}}}{d_{\mathsf{min}}} \cdot \sqrt{t_{\mathsf{rel}}} \cdot \log\left(t_{\mathsf{rel}} \log n\right)\right).$$

(3) For any two vertices $u, v \in V$ there is a strategy so that for the ε-TBRW the expected time to reach v from u is at most
$$\mathcal{O}\left(\left(\frac{n \cdot d_{\mathsf{avg}}}{d_{\mathsf{min}}}\right)^{1-\varepsilon} \cdot t_{\mathsf{rel}} \cdot \ln n\right).$$

In Section 4.1 we study how much the controller can affect the stationary distribution of any vertex in our graph. Azar et al. [4] studied this problem and showed that for any bounded degree graph a controller increase the stationary probability of any vertex from p to $p^{1-\Omega(\varepsilon)}$. By applying the results from Section 3 we can prove a stronger bound for graphs with small relaxation time and subpolynomial degree ratio:

(4) In any graph a controller can increase the stationary probability of any vertex from p to $p^{1-\varepsilon+\delta}$, where $\delta = \ln\left(12 \cdot t_{\mathsf{rel}} \cdot \ln n\right) / |\ln p|$.

Motivated by a comment of Azar et al. stating that for regular graphs the interesting case is when ε is not substantially larger than $1/d$, we try to quantify the effect of a controller in this regime. Establishing a number of bounds and counter-examples we conclude that:

(5) A controller cannot increase any entry in the stationary distribution of a uniformly dense graph by more than a constant factor.

(6) In some regular graphs graphs of polynomial degree entries in the stationary distribution can be increased by more than a polynomial factor, in others this is not possible.

In Section 5 we consider the complexity of finding an optimal stratgies to cover a graph in minimum expected time. Azar et al. considered this problem for hitting times and showed that there is a polynomial algorithm to determine an optimal strategy, we establish a dichotomy by showing that the cover time problem is hard:

(6) Given the covered set X and position v of the walk at some time, it is NP-hard to choose the next step from all neighbours of v so as to minimise the expected time for the ε-TBRW to visit every vertex not in X, assuming an optimal strategy is followed thereafter.

Finally in Section 6 we conclude with some open problems.

We shall now formally describe the ε-Biased and ε-Time Biased random walk model and introduce some notation.

2 Preliminaries

Thought this paper we shall always consider a connected n-vertex graph $G = (V, E)$, which unless otherwise specified, will be simple and unweighted.

2.1 ε-Biased and ε-Time-Biased Random Walks.

Azar et al. [4], building on earlier work [5], introduced the ε-biased random walk (ε-BRW) on a graph G. Each step of the ε-BRW is preceded by an $(\varepsilon, 1-\varepsilon)$-coin flip. With probability $1-\varepsilon$ a step of the simple random walk is performed, but with probability ε the controller gets to select which neighbour to move to. The selection can be probabilistic, but it is time independent. Thus if $\mathbf{P}$ is the transition matrix of a random walk, then the transition matrix $\mathbf{Q}^{\varepsilon\mathrm{B}}$ of the ε-biased random walk is given by

$$\mathbf{Q}^{\varepsilon\mathrm{B}} = (1-\varepsilon)\mathbf{P} + \varepsilon\mathbf{B}, \tag{1}$$

where $\mathbf{B}$ is an arbitrary stochastic matrix chosen by the controller, with support restricted to $E(G)$. The controller of an ε-BRW has full knowledge of G.

Azar et al. focused on the problems of bias strategies which either minimise or maximise the stationary probabilities of sets of vertices or which minimise the hitting times of vertices. Azar et al. [4, Sec. 4] make the connection between Markov Decision processes and the ε-Biased walk, in particular they observe that the two tasks they study can be identified as the expected average cost and optimal first-passage problems respectively in this context [8]. The existence of time independent optimal strategies for both objectives follow from Theorems 2 and 3 respectively in [8, Ch. 3]. For this reason Azar et al. restrict to the class of unchanging strategies, where we say that an ε-Bias strategy is *unchanging* if it is independent of both time and the history of the walk.

It is clear that if we wish to consider optimal strategies to cover a graph (visit every vertex) in shortest expected time then we must include strategies which depend on the set of vertices already visited by the walk. Let $\mathcal{H}_t$ be the history of the random walk up to time t, that is the sigma algebra $\mathcal{H}_t = \sigma(X_0, \dots, X_t)$ generated by all steps of the walk up to and including time t. Thus we consider a time-dependent version, where the bias matrix $\mathbf{B}_t$ may depend on the time t and the history $\mathcal{H}_t$; we refer to this as the ε-time-biased walk (ε-TBRW).

Proposition 2.1. *For any connected graph G there is an optimal strategy for the ε-TBRW to cover G which is fixed over any time interval between times when a new vertex is visited.*

The result above, proven in Section 5, essentially says that there is an optimal strategy which is unchanging between times when a new vertex is discovered. A consequence of Proposition 2.1 is the existence of an optimal strategy for covering the graph can be described by a set of bias matrices $\{\mathbf{B}_U\}$ where U is a connected subset of vertices, of which there are at most $2^{|V|}$.

Let $C_v^{\varepsilon\mathsf{TB}}(G)$ denote the minimum expected time (taken over all strategies) for the ε-TBRW to visit every vertex of G starting from v, and define the *cover time* $t_{\mathsf{cov}}^{\varepsilon\mathsf{TB}}(G) := \max_{v\in V} C_v^{\varepsilon\mathsf{TB}}(G)$. Similarly let $H_x^{\varepsilon\mathsf{B}}(y)$ denote the minimum expected time for the ε-biased walk to reach y, which may be a single vertex or a set of vertices, starting from a vertex x. We do not need to provide notation for the hitting times of the ε-TBRW since, as mentioned before, there is always a time-independent optimal strategy for hitting a given vertex [4, Thm. 11], thus hitting times in the ε-TBRW and ε-BRW are the same. We also define the *hitting time* $t_{\mathsf{hit}}^{\varepsilon\mathsf{B}}(G) := \max_{x,y\in V} H_x^{\mathsf{two}}(y)$. Any unchanging strategy on a finite connected graph results in an irreducible Markov chain and thus, when appropriate, we refer to its stationary distribution as π.

We shall introduce some more notation. For a graph G let $d_{\mathsf{max}}, d_{\mathsf{min}}$ and d_{avg} denote the maximum, minimum and average degree of G respectively. Let $t_{\mathsf{rel}} := \frac{1}{1-\lambda_2}$ be the relaxation time of G, where λ_2 is the second largest eigenvalue of the transition matrix of the lazy random walk (LRW) on G with loop probability $1/2$.

3 Hitting and Cover Times in Expanders

In this section we prove that the ε-TBRW has the power to increase the probability of certain events. As a consequence of this result we obtain bounds on the cover and hitting times of the ε-TBRW on a graph G in terms of n, the extremal and average degrees, and the relaxation time.

To prove these results, we use the methods we developed in [10] to analyse the related choice random walk (CRW). These methods will apply to the ε-TBRW with relatively minor changes; the key additional contribution of the current paper is a new technical lemma which is specific to the ε-TBRW and which appears in Section 3.1.

The basic approach used in [10] is, for a given graph G, to consider events which depend only on the trajectory of the walker (that is, the sequence of vertices visited) up to some fixed time t. We use a "tree gadget" to encode all possible trajectories. This then allows us to relate the probability of a given event in the ε-TBRW to that for the SRW; the role of the technical lemma is to recursively bound the effects of an optimal strategy for the ε-TBRW at each level of the gadget.

Fix a vertex u, a non-negative integer t and a set S of trajectories of length t (here the length is the number of steps taken). Write $p_{u,S}$ for the probability that running a SRW starting from u for t steps results in a member of S. Let $q_{u,S}(\varepsilon)$ be the corresponding probability for the ε-TBRW law, which depends on the particular strategy used. We prove the following result relating $q_{u,S}(\varepsilon)$ to $p_{u,S}$.

Theorem 3.1. *Let G be a graph, $u \in V$, $t > 0$, $0 \leqslant \varepsilon \leqslant 1$ and S be a set of trajectories of length t from u. Then there exists a strategy for the ε-TBRW such that*

$$q_{u,S}(\varepsilon) \geqslant (p_{u,S})^{1-\varepsilon}.$$

Here we typically think of S encoding such events as "the walker is in a set $W \subset V$ at time t" or "the walker has visited $v \in V$ by time t"; however, the result applies to any event measurable at time t. This theorem has the following consequences for the hitting and cover times of the ε-TBRW.

Theorem 3.2. *For any graph G, and any $\varepsilon \in (0,1)$,*

$$t_{\mathsf{cov}}^{\varepsilon\mathsf{TB}}(G) = \mathcal{O}\left(\frac{n}{\varepsilon} \cdot \frac{d_{\mathsf{avg}}}{d_{\mathsf{min}}} \cdot \sqrt{t_{\mathsf{rel}}} \cdot \log\left(t_{\mathsf{rel}} \log n\right)\right).$$

Theorem 3.3. *For any graph G, any $x, y \in V$ and any $\varepsilon \in (0,1)$, we have*

$$H_x^{\varepsilon\mathsf{B}}(y) \leqslant 12 \cdot \pi(y)^{\varepsilon-1} \cdot t_{\mathsf{rel}} \cdot \ln n;$$

this bound also holds for return times. Consequently,

$$t_{\mathsf{hit}}^{\varepsilon\mathsf{B}}(G) \leqslant 12 \left(\frac{n \cdot d_{\mathsf{avg}}}{d_{\mathsf{min}}}\right)^{1-\varepsilon} \cdot t_{\mathsf{rel}} \cdot \ln n.$$

Theorems 3.2 and 3.3 are analogues for the ε-TBRW of Theorems 6.1 and 6.2 of [10], and their derivation from Theorem 3.1 follows that given in [10, Section 6.1] exactly.

The main difference between the two sets of results is that each relies on an operator which describes the random walk process being studied. The operator used here is different to those introduced in [10], and as a result so is the strength of the boosting obtainable. This highlights the versatility of the technique used to prove Theorem 3.1 in that it can be used to analyse several different "non-deterministically" random processes such as the ε-TBRW and the CRW.

Theorem 3.2 has the following consequence for expanders; a sequence of graphs (G_n) is a *sequence of expanders* if $t_{\mathsf{rel}}(G_n) = \Theta(1)$.

Corollary 3.4. *For every sequence $(G_n)_{n\in\mathbb{N}}$ of n-vertex bounded degree expanders and any fixed $\varepsilon > 0$, we have*

$$t_{\mathsf{cov}}^{\varepsilon\mathsf{TB}}(G_n) = \mathcal{O}\Big(\frac{n}{\varepsilon}\cdot \log\log n\Big).$$

3.1 The ε-Max/Average Operation

In this subsection we shall introduce an operator which models the action of the ε-TBRW. We shall then prove a bound on the output of the operator, this is used to show that the ε-TBRW can boost probabilities indexed by paths.

For $0 < \varepsilon < 1$ define the ε-max/average operator $\mathrm{MA}_\varepsilon : [0,\infty)^m \to [0,\infty)$ by

$$\mathrm{MA}_\varepsilon\left(x_1,\dots,x_m\right) = \varepsilon\cdot \max_{1\leqslant i\leqslant m} x_i + \frac{1-\varepsilon}{m}\cdot\sum_{i=1}^m x_i.$$

This can be seen as an average which is biased in favour of the largest element, indeed it is a convex combination between the largest element and the arithmetic mean.

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

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

and

$$M_\infty(x_1,\dots,x_m) = \max\{x_1,\dots,x_m\} = \lim_{p\to\infty} M_p(x_1,\dots,x_m).$$

Thus we can express the ε-max/ave operator as $\mathrm{MA}_\varepsilon(\cdot) = (1-\varepsilon)M_1(\cdot)+\varepsilon\,\mathrm{M}_\infty(\cdot)$. We use a key lemma, Lemma 3.5, which could be be described as a multivariate anti-convexity inequality.

Lemma 3.5. *Let $0<\varepsilon<1$, $m\geqslant 1$ and $\delta\leqslant \varepsilon/(1-\varepsilon)$. Then for any $x_1,\dots,x_m\in[0,\infty)$,*

$$M_{1+\delta}\left(x_1,\dots,x_m\right)\leqslant \mathrm{MA}_\varepsilon\left(x_1,\dots,x_m\right).$$

Proof. We begin by establishing the following claim.

Claim. *Let $\eta\in(0,1)$, and suppose $a,b,c\in\mathbb{R}^+$ with $c=(1-\eta)a+\eta b$. Then*

$$M_c \leqslant M_a^{(1-\eta)a/c} M_b^{\eta b/c}. \tag{2}$$

Proof of claim. Hölder's inequality states for positive reals $y_1,\dots,y_m$ and $z_1,\dots,z_m$ that

$$y_1z_1+\cdots+y_mz_m \leqslant \big(y_1^p+\cdots+y_m^p\big)^{1/p}\big(z_1^q+\cdots+z_m^q\big)^{1/q},$$

where $p,q\geqslant 1$ satisfy $1/p+1/q=1$. The desired result follows by setting $y_i = x_i^{(1-\eta)a}$, $z_i = x_i^{\eta b}$, $p=1/(1-\eta)$, $q=1/\eta$, dividing both sides by m and then taking c^{th} roots. $\diamondsuit$

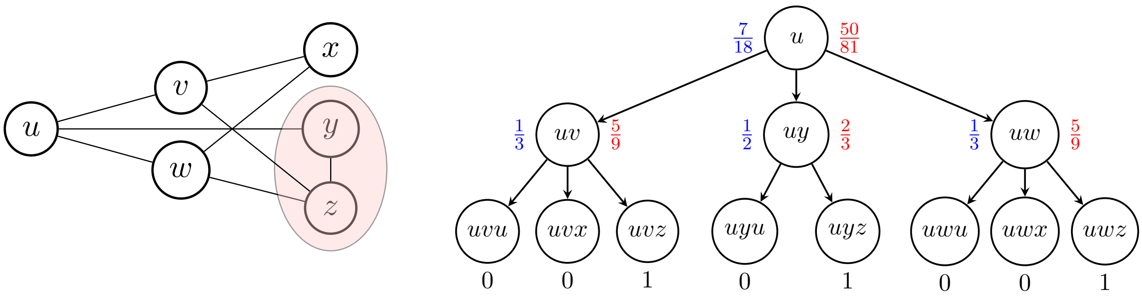

Figure 1: Illustration of a (non-lazy) walk on a non-regular graph starting from u with the objective of being at $\{y, z\}$ at step $t = 2$. The probabilities of achieving this are given in blue (left) for the SRW and in red (right) for the $\frac{1}{3}$-TBRW.

Applying (2), we have for any $k > \delta$ that

$$\begin{aligned} M_{1+\delta} &\leqslant M_1^{\frac{1-\delta/k}{1+\delta}} M_{k+1}^{\frac{(k+1)\delta/k}{1+\delta}} \\ &\leqslant \frac{1-\delta/k}{1+\delta} M_1 + \frac{(k+1)\delta/k}{1+\delta} M_\infty, \end{aligned}$$

using the weighted AM-GM inequality and the fact that $M_p \leqslant M_\infty$ for any p. Taking limits as $k \to \infty$, noting that $\varepsilon \geqslant \delta/(1+\delta)$, gives the required inequality. □

Remark. *The dependence of δ on ε given in Lemma 3.5 is best possible. This can be seen by setting $x_1 = 0$ and $x_i = 1$ for $2 \leqslant i \leqslant m$, and letting m tend to ∞.*

3.2 The Tree Gadget for Graphs

In this section we show how the "tree gadget" of [10] can be used to prove Theorem 3.1. This gadget encodes walks of length at most t from u in a rooted graph (G, u) by vertices of an arborescence $(\mathcal{T}_t, \mathbf{r})$, i.e. a tree with all edges oriented away from the root $\mathbf{r}$. Here we use bold characters to denote trajectories, and $\mathbf{r}$ will be the length-0 trajectory consisting of the single vertex u. The tree $\mathcal{T}_t$ consists of one node for each trajectory of length $i \leqslant t$ starting at u, and has an edge from $\mathbf{x}$ to $\mathbf{y}$ if $\mathbf{x}$ may be obtained from $\mathbf{y}$ by deleting the final vertex.

The proof of Theorem 3.1 will follow the corresponding proof in [10] closely, but we give a full proof here in order to clarify the role played by the ε-max/average operator. As usual we write $d^+(\mathbf{x})$ for the number of offspring in $\mathcal{T}_t$ of $\mathbf{x}$, and $\Gamma^+(\mathbf{x})$ for the set of offspring of $\mathbf{x}$. We denote the length of the walk $\mathbf{x}$ by $|\mathbf{x}|$. We shall extend our notation $p_{u,S}$ and $q_{u,S}(\varepsilon)$ to $p_{\mathbf{x},S}$ and $q_{\mathbf{x},S}(\varepsilon)$, defined to be the probabilities that extending $\mathbf{x}$ to a trajectory of length t, using the laws of the SRW and ε-TBRW respectively, results in an element of S. Additionally, let $W_u(k) := \bigcup_{i=0}^{k}\{X_i\}$ be the trajectory of a simple random walk X_i on G up to time k, with $X_0 = u$.

Proof of Theorem 3.1. For convenience we shall suppress the notational dependence of $q_{\mathbf{x},S}(\varepsilon)$ on ε. To each node $\mathbf{x}$ of the tree gadget $\mathcal{T}_t$ we assign the value $q_{\mathbf{x},S}$ under the the ε-TB strategy of biasing towards a neighbour in G which extends to a walk $\mathbf{y} \in \Gamma^+(\mathbf{x})$ maximising $q_{\mathbf{y},S}$. 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. Then with probability $1-\varepsilon$ we choose the next step of the walk uniformly at random in which case the probability of reaching S from $\mathbf{x}$ is just the average of $q_{\mathbf{y},S}$ over the offspring $\mathbf{y}$ of $\mathbf{x}$, otherwise we choose a maximal $q_{\mathbf{y},S}$. Thus the value of $\mathbf{x}$ is given by the ε-max/average of its offspring, that is

$$q_{\mathbf{x},S} = \mathrm{MA}_\varepsilon\left(\left(q_{\mathbf{y},S}\right)_{\mathbf{y}\in\Gamma^+(\mathbf{x})}\right). \tag{3}$$

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

$$\Phi^{(i)} = \sum_{|\mathbf{x}|=i} q_{\mathbf{x},S}^{1+\delta} \cdot \mathbb{P}\left[W_u(i) = \mathbf{x} \right]. \tag{4}$$

Notice that if $\mathbf{xy} \in E(\mathcal{T}_t)$ then

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

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

$$\Phi^{(i)} = \sum_{|\mathbf{x}|=i-1} \sum_{\mathbf{y}\in\Gamma^+(\mathbf{x})} q_{\mathbf{y},S}^{1+\delta} \cdot \frac{\mathbb{P}\left[W_u(i-1) = \mathbf{x} \right]}{d^+(\mathbf{x})}. \tag{5}$$

We now show that $\Phi^{(i)}$ is non-increasing in i. By combining (4) and (5) we can see that the difference $\Phi^{(i-1)} - \Phi^{(i)}$ is given by

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

Recalling (3), to establish $\Phi^{(i-1)} - \Phi^{(i)} \geqslant 0$ it is sufficient to show the following inequality holds whenever $\mathbf{x}$ is not a leaf:

$$\mathrm{MA}_\varepsilon\left(\left(q_{\mathbf{y},S}\right)_{\mathbf{y}\in\Gamma^+(\mathbf{x})}\right)^{1+\delta} \geqslant \frac{1}{d^+(\mathbf{x})} \sum_{\mathbf{y}\in\Gamma^+(\mathbf{x})} q_{\mathbf{y},S}^{1+\delta}.$$

By taking $(1+\delta)^{\text{th}}$ roots this inequality holds for any $\delta \leqslant \varepsilon/(1-\varepsilon)$ by Lemma 3.5, and thus for δ in this range $\Phi^{(i)}$ is non-increasing in i.

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

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

Thus since $\Phi^{(t)}$ is non-decreasing $q_{u,S}^{1+\delta} = \Phi^{(t)} \geqslant \Phi^{(0)} = p_{u,S}$. The result for the ε-TBRW follows by taking $\delta = \varepsilon/(1-\varepsilon)$. $\square$

4 Increasing and Decreasing Stationary Probabilities

In this section we shall consider a problem of how much an unchanging strategy can affect the stationary probabilities in a graph. Azar et al. studied this question and made an appealing conjecture. Our result on the hitting times of the ε-BRW will allows us to make progress towards this conjecture. We also derive some more general bounds on stationary probabilities for classes of Markov chains which include certain regimes for the ε-BRW, and tackle the question of when the stationary probability of a vertex can be changed by more than a constant factor.

4.1 A Conjecture of Azar et al. for the ε-BRW

Azar, Broder, Karlin, Linial and Phillips make the following conjecture for the ε-BRW [4, Conjecture 1].

Conjecture 4.1 (ABKLP Conjecture)**.** *In any graph, a controller can increase the stationary probability of any vertex from p to $p^{1-\varepsilon}$.*

This conjecture fails for the graph K_2, as no strategy for the ε-BRW can increase the stationary probability over that of a simple random walk. This motivates weakening the conjecture by replacing $p^{1-\varepsilon}$ by $p^{1-\varepsilon+o_n(1)}$, however this fails for the star on n vertices, and non-bipartite counterexamples may be obtained by adding a small number of extra edges to the star. In each of these counter examples there is a vertex with constant stationary probability, for large graphs this can only happen if there is a large degree discrepancy. We believe the following should hold.

Conjecture 4.2. *In any graph a controller can increase the stationary probability of any vertex from p to $p^{1-\varepsilon+\delta}$, where $\delta = \mathcal{O}(1/|\log p|)$.*

Azar et al. prove a weaker bound of $p^{1-\mathcal{O}(\varepsilon)}$ for bounded-degree regular graphs. As a corollary of Theorem 3.3 we confirm Conjecture 6.1 for any graph where t_{rel} is subpolynomial in n. Our techniques are different to those of Azar et al. and allow us to cover a larger class of graphs, including dense graphs as well as sparse ones, as well as getting closer to the conjectured bound.

Theorem 4.3. *In any graph a controller can increase the stationary probability of any vertex from p to $p^{1-\varepsilon+\delta}$, where $\delta = \ln(12 \cdot t_{\mathsf{rel}} \cdot \ln n) / |\ln p|$.*

Proof. By Theorem 3.3 for each vertex v there exists a strategy so that the return time to v is at most $12 \cdot \pi(v)^{\varepsilon-1} \cdot t_{\mathsf{rel}} \ln n$. Let q denote the stationary probability of v for this ε-B walk. Then as stationary probability is equal to the reciprocal of the return time by [15, Prop. 1.14] we have $q \geqslant \pi(v)^{1-\varepsilon}/(12 t_{\mathsf{rel}} \ln n)$, for the simple random walk $p = \pi(y)$. If we let $\delta = \ln(12 t_{\mathsf{rel}} \ln n) / |\ln \pi(y)|$ then we have

$$q/p^{1-\varepsilon+\delta} \geqslant \frac{\pi(v)^{1-\varepsilon}}{12 t_{\mathsf{rel}} \ln n} \cdot \frac{\pi(y)^{-\delta}}{\pi(v)^{1-\varepsilon}} = \frac{\exp\left(-\ln \pi(y) \cdot \frac{\ln(12 \cdot t_{\mathsf{rel}} \cdot \ln n)}{|\ln \pi(y)|}\right)}{12 t_{\mathsf{rel}} \ln n} = 1.$$

□

The dependence of δ on $|\ln p|$ in Corollary 4.3 imposes the condition that any vertex you wish to boost must have sub-polynomial degree. This condition is tight in some sense as no stationary probability bounded from below can be boosted by more than a constant factor.

4.2 Approximation by Edge-Weighted Graphs

Let $G = (V, E)$ be any connected, undirected graph with degree bound $d \leqslant C$. We will associate to every edge a positive weight given by the function $w : E \to \mathbb{R}^+$. We consider a random walk that picks an incident edge with probability proportional to its weight. Recall that the stationary distribution of this walk is given by $\pi(x) = \frac{\sum_{y \sim x} w(y,x)}{2W}$, where $W := \sum_{\{r,s\} \in E(G)} w(r,s)$ is the total sum of weights assigned.

Fix a vertex $u \in V$ and let $-1 < a < \infty$. We consider the weight function given by

$$w(r,s) = (1+a)^{\max\{d(u,r), d(u,s)\}}\,. \tag{6}$$

Note that this particular weight function satisfies the following property:

$$\forall u, v, w \colon \{u,v\}, \{u,w\} \in E(G) \colon \frac{w(u,v)}{w(u,w)} \in \{1+a, (1+a)^{-1}, 1\}. \tag{7}$$

Proposition 4.4. *Let $-1 < a < \infty$, and let G be an edge-weighted graph whose weights satisfy (7). Then, provided $\varepsilon \geqslant -a$ if $a \leqslant 0$ and $\varepsilon \geqslant a/(1+a)$ if $a > 0$, the ε-BRW can emulate the walk given by those weights.*

Proof. It suffices to prove that we may emulate a step of the walk from any given vertex x. If all edges meeting x have the same weight, we simply "bias" towards the uniform distribution on neighbours of x. Otherwise $a \neq 0$, $d = d(x) \geqslant 2$ and there are exactly two weights, w_1 and w_2, incident to x, which satisfy $w_1 = (1+a)w_2$. Suppose there are k incident edges of weight w_1 and $d-k$ of weight w_2; clearly $1 \leqslant k \leqslant d-1$. Now we need to construct a bias matrix $\mathbf{B}$ which will satisfy the walk probabilities given by (6). Note that if $w(xy) = w_1$ then $p_{x,y} = w_1/(kw_1 + (d-k)w_2) = (1+a)/(ak+d)$ and otherwise $p_{x,y} = 1/(ak+d)$.

We first consider the case $a > 0$, i.e. $w_1 > w_2$. It is sufficient to assume $\varepsilon = \frac{a}{1+a}$, since if it is larger we may use the ε-BRW to emulate the $\frac{a}{1+a}$-BRW. In this case set

$$\mathbf{B}_{x,z} = \begin{cases} \frac{da+2d-k}{dak+d^2} & \text{if } w(xz) = w_1 \\ \frac{d-k}{dak+d^2} & \text{if } w(x,z) = w_2. \end{cases}$$

This gives $\sum_{z \sim x} \mathbf{B}_{x,z} = 1$, all entries are positive and

$$p_{x,z} = \frac{a}{1+a} \cdot \mathbf{B}_{x,z} + \frac{1}{1+a} \cdot \frac{1}{d} = \begin{cases} \frac{a+1}{ka+d} & \text{if } w(xz) = w_1 \\ \frac{1}{ka+d} & \text{if } w(x,z) = w_2. \end{cases}$$

The case $c < 0$ may be reduced to the previous case by replacing c with $c' = \frac{-c}{1+c}$, noting that $\varepsilon \geqslant -c$ is equivalent to $\varepsilon \geqslant \frac{c'}{1+c'}$. □

Theorem 4.5. *Let G be any graph such that $d_{\mathsf{max}} \geqslant 3$ and let $\varepsilon > 0$. Then*

(i) a controller for the ε-BRW can increase the stationary probability of any vertex from p to $p^{1-\widetilde{\varepsilon}}$, where

$$\widetilde{\varepsilon} = \frac{-\ln(1-\varepsilon)}{\ln(d_{\mathsf{max}}-1)}\left(1 - \frac{\ln(n \cdot p)}{\ln n}\right) > 0.$$

(ii) If $d_{\mathsf{max}} \leqslant n^{1/4}$ then a controller for the ε-BRW can decrease the stationary probability of

any vertex from p to $p^{1+\widetilde{\varepsilon}}$, where

$$\widetilde{\varepsilon} = -\ln(1-\varepsilon)\left(\frac{1}{\ln d_{\mathsf{max}}} - \frac{3}{\ln n}\right) - \frac{1}{\ln n} > 0.$$

Proof. For (i), consider the weighting scheme $w(r,s) = (1-\varepsilon)^{\max\{d(u,r),d(u,s)\}}$. Observe that there are at most $d_{\mathsf{max}}(d_{\mathsf{max}}-1)^{i-1}$ vertices at distance exactly i from u (and also edges from vertices at distance $i-1$ to those at i). Thus if we consider the total weight W of the graph then for any r,

$$\begin{aligned} W &\leqslant \sum_{i=1}^{r} d_{\mathsf{max}}(d_{\mathsf{max}}-1)^{i-1} \cdot (1-\varepsilon)^{i-1} + n \cdot d_{\mathsf{avg}} \cdot (1-\varepsilon)^r \\ &\leqslant (2(d_{\mathsf{max}}-1)^r + n \cdot d_{\mathsf{avg}}) \cdot (1-\varepsilon)^r. \end{aligned}$$

Thus if we let $r = \lfloor \ln(n)/\ln(d_{\mathsf{max}}-1) \rfloor$ then $W \leqslant d_{\mathsf{avg}} \cdot n^{1+\kappa}$, where $\kappa = \ln(1-\varepsilon)/\ln(d_{\mathsf{max}}-1) < 0$. For any $u \in V$ it follows that $\pi'(u) \geqslant d(u)/d_{\mathsf{avg}} \cdot n^{1+\kappa} = n \cdot \pi(u)/n^{1+\kappa}$ and so for $\delta \geqslant 0$,

$$\frac{\pi'(u)}{\pi(u)^{1+\kappa+\delta}} \geqslant \frac{n \cdot \pi(u)}{n^{1+\kappa}} \cdot \frac{n^{1+\kappa+\delta}}{(n \cdot \pi(u))^{1+\kappa+\delta}} = (n \cdot \pi(u))^{-\kappa-\delta} \cdot n^{\delta} \geqslant 1,$$

where the final inequality holds by taking $\delta = |\kappa \ln(n\pi(u))|/\ln n$.

For (ii), use the weighting scheme

$$w(r,s) = \left(1 + \frac{\varepsilon}{1-\varepsilon}\right)^{\max\{d(u,r),d(u,s)\}} = (1-\varepsilon)^{-\max\{d(u,r),d(u,s)\}}.$$

Note that the total number of edges of weight at most $(1-\varepsilon)^{-r+1}$ in any graph of max degree d_{max} is at most $\sum_{i=1}^{r} d_{\mathsf{max}}(d_{\mathsf{max}}-1)^{i-1} \leqslant 3(d_{\mathsf{max}}-1)^r$. Thus, setting $r = \lfloor \log_{d_{\mathsf{max}}-1}(n/300) \rfloor$, there are at most $n/100$ such edges. It follows that $W \geqslant (d_{\mathsf{avg}} - 1/100)n \cdot (1-\varepsilon)^{-r} \geqslant (d_{\mathsf{avg}} - 1/100) \cdot n^{1+\kappa}$, where $\kappa = -\frac{\ln(1-\varepsilon)}{\ln n}\left\lfloor \frac{\ln(n/300)}{\ln(d_{\mathsf{max}}-1)} \right\rfloor$. Thus for any $u \in V$ we have $\pi'(u) \leqslant d(u)/(d_{\mathsf{avg}} - 1/100) \cdot n^{1+\kappa})$ and so for $\delta > 0$

$$\frac{\pi(u)^{1+\kappa-\delta}}{\pi'(u)} \geqslant \frac{d_{\mathsf{avg}}(1 - \frac{1}{100 d_{\mathsf{avg}}}) \cdot n^{1+\kappa}}{d(u)} \cdot \frac{(n \cdot \pi(u))^{1+\kappa-\delta}}{n^{1+\kappa-\delta}} \geqslant \frac{99}{100} \cdot \frac{d(u)^{\kappa-\delta}}{d_{\mathsf{avg}}^{\kappa-\delta}} \cdot n^{\delta}. \tag{8}$$

If $d(u) \geqslant d_{\mathsf{avg}}$ we can take $\delta = \mathcal{O}(1/\log n)$ and (8) is greater than one, otherwise let $\delta = (-2 \cdot \kappa \ln(n \cdot \pi(u)) + 1)/\ln(n)$. Now $\kappa = -\frac{\ln(1-\varepsilon)}{\ln n}\left\lfloor \frac{\ln(n/300)}{\ln(d_{\mathsf{max}}-1)} \right\rfloor \geqslant -\ln(1-\varepsilon)\left(\frac{1}{\ln d_{\mathsf{max}}} - \frac{1}{\ln n}\right)$ thus $-\kappa \ln(n\pi(u)) \leqslant \kappa \ln d_{\mathsf{max}}$. It follows that p can be decreased to $p^{1+\widetilde{\varepsilon}}$ where

$$\widetilde{\varepsilon} = \kappa - \delta \geqslant -\ln(1-\varepsilon)\left(\frac{1}{\ln d_{\mathsf{max}}} - \frac{3}{\ln n}\right) - \frac{1}{\ln n}. \qquad \square$$

Corollary 4.6. *Let G be any graph satisfying $d_{\mathsf{max}} \leqslant n^{1/4}$ and $\varepsilon > 0$. Then a controller for the ε-BRW can increase the stationary probability of any vertex from p to $p^{1-3\varepsilon/(4\ln d_{\mathsf{max}})}$ and decrease it from p to $p^{1+\varepsilon/(4\ln d_{\mathsf{max}})-1/\ln n}$.*

Proof. The statement holds for paths & cycles and for graphs such that $d_{\mathsf{max}} \geqslant 3$ this follows from Theorem 4.5 since $-\ln(1-x) \geqslant x$ for any $x \leqslant 1$. $\square$

4.3 Bounds on Stationary Probabilities

In this section we shall focus on a class of (non-nessesarily reversible) Markov chains which resemble a simple walk on an almost regular graph.

Let Q be a transition matrix supported on G. For c, C such that $0 < c \leqslant 1 \leqslant C < \infty$ we say that Q is a (c, C)-*simple walk* on G if for every edge $uv \in E(G)$,

$$\frac{c}{d(u)} \leqslant q_{u,v} \leqslant \frac{C}{d(u)}.$$

We begin with a simple result for reversible (c, C)-simple walks to motivate what we wish to obtain in the non-reversible case.

Proposition 4.7. *Let $G = (V, E)$ be a connected, d-regular edge-weighted graph with diameter D. Then the following holds for any (c, C)-Simple walk:*

$$\pi_{\max} \leqslant C^D \cdot \pi_{\min} \qquad \textit{and} \qquad \pi_{\mathsf{min}} \geqslant c^D \pi_{\max}.$$

Proof. By reversibility $\pi(x)p_{x,y} = \pi(y)p_{y,x}$ for all $x, y \in V$. Thus let $v_0, v_d \in V$ satisfy $\pi(v_0) = \pi_{\max}$ and $\pi(v_d) = \pi_{\min}$. There exists a path $v_0, v_1, \ldots, v_d$, where $d \leqslant D$, thus

$$\pi(v_0) = \frac{p_{v_1,v_0}}{p_{v_0,v_1}}\pi(v_1) = \frac{p_{v_1,v_0}}{p_{v_0,v_1}} \cdots \frac{p_{v_d,v_{d-1}}}{p_{v_{d-1},v_d}}\pi(v_d) \leqslant C^D \pi(v_d),$$

it follows that $\pi_{\max} \leqslant C^D \cdot \pi_{\min}$. Similarly on can show that $\pi_{\mathsf{min}} \geqslant c^D \pi_{\mathsf{max}}$. ∎

Proposition 4.7 shows that the stationary distributions of reversible (c, C)-simple walks on almost regular graphs with constant diameter behave well. We shall prove an analogous result for stationary distributions of (c, C)-simple walks which are not reversible, where diameter is replaced by mixing time. Let t_{mix} be the mixing time given by

$$t_{\mathsf{mix}} = \inf\left\{t : \max_{x \in V} \|p_x^{(t)} - \pi\|_{TV} \leqslant \frac{1}{e}\right\} \qquad \text{where} \qquad \|\mu - \nu\|_{TV} = \frac{1}{2}\sum_{y \in V} |\mu(y) - \nu(y)|.$$

Note that $t_{\mathsf{mix}} = \mathcal{O}\Big(\frac{\log n}{1-\lambda_2}\Big)$ by [15, Thm. 12.3]. We shall also need the separation time, defined

$$t_{\mathsf{sep}} = \inf\left\{t : \max_{x,y \in V} \left|\frac{p_{x,y}^{(t)}}{\pi_P(y)} - 1\right| \leqslant \frac{1}{e}\right\}.$$

Proposition 4.8. *Let G be an almost-regular graph, where γ is such that $d_{\mathsf{avg}}/\gamma \leqslant d(u) \leqslant \gamma d_{\mathsf{avg}}$ for all $u \in V$. Let Q be a (c, C)-simple walk on G with stationary distribution π_Q. Finally, let $\tau = 4t_{\mathsf{mix}}$, where t_{mix} is the mixing time of the lazy random walk on G. Then for any $u \in V$*

$$\frac{c^\tau}{2\gamma n} \leqslant \pi_Q(u) \leqslant \frac{2\gamma C^\tau}{n}.$$

Proof. Given a Markov chain H we call $(I + H)/2$ the lazy version of H, where I is the identity matrix. Let P be the transition matrix of a lazy random walk on G with stationary distribution π_P. Since making a Markov Chain lazy does not alter the stationary distribution, we may take the lazy version of Q. Now we have made Q lazy switching from P^τ to Q^τ can increase (resp. decrease) the probabilities of any fixed trajectory by a factor of at most C^τ (resp. c^τ).

By monotonicity of the separation distance, $\frac{e-1}{e}\pi_P(v) \leqslant p_{u,v}^{(\tau)} \leqslant \frac{e+1}{e}\pi_P(v)$ for any $u, v \in V$ and $\tau \geqslant t_{\mathsf{sep}}$. Recall also that for any Markov Chain P the stationary distribution satisfies $\pi_P \cdot P^\tau = \sum_{v \in V} \pi_P(v) \cdot p_{v,u}^{(\tau)} = \pi_P$. Thus for $\tau = t_{\mathsf{sep}}$ we have

$$\begin{aligned}
\pi_Q(u) &= \sum_{v \in V} \pi_Q(v) \cdot q_{v,u}^{(\tau)} \\
&\leqslant C^\tau \cdot \sum_{v \in V} \pi_Q(v) \cdot p_{v,u}^{(\tau)} \\
&\leqslant C^\tau \cdot \frac{e+1}{e} \cdot \frac{d_{\mathsf{max}}}{n d_{\mathsf{avg}}} \cdot \sum_{v \in V} \pi_Q(v) \\
&\leqslant 2C^\tau \gamma / n.
\end{aligned}$$

The lower bound is shown analogously but with constant $c^\tau(e-1)/(e\gamma) \geqslant c^\tau/(2\gamma)$. Noting that $t_{\mathsf{sep}} \leqslant 4t_{\mathsf{mix}}$ by [1, Thm. 4.6], the result follows. □

Although the main focus of this section is on increasing stationary probabilities we shall also consider by how much a controller can reduce stationary probabilities.

Corollary 4.9. *For G and any $\varepsilon < 1/(5t_{\mathsf{mix}})$ a controller cannot decrease the stationary probability of any almost regular graph by more than a constant factor.*

Proof. Observe that the ε-TBRW is (c, C)-simple with $c = 1 - \varepsilon$. It follows from Proposition 4.8 that for any $\pi_Q(v)/\pi_P(v) \geqslant (1-\varepsilon)^\tau/(2\gamma^2)$ for $\tau = 4t_{\mathsf{mix}}$, where $\gamma < \infty$ constant since the graph is almost regular. By the Bernoulli inequality $(1-\varepsilon)^\tau \geqslant 1 - \varepsilon\tau$, and the result follows. □

Proposition 4.10. *Let Q be a (c, C)-simple walk on G with stationary distribution π_Q. Then $\max_{u \in V} \pi_Q(u) \leqslant C/d_{\mathsf{min}}$.*

Proof. For any $u \in V$ we have

$$\pi(u) = \sum_{v \in V} \pi(v) p_{v,u} \leqslant \frac{C}{d_{\mathsf{min}}} \sum_{v \in V} \pi(v) \leqslant \frac{C}{d_{\mathsf{min}}}.$$

□

4.3.1 Everywhere Dense Graphs

We say that a graph is everywhere dense if has minimum degree $\Omega(n)$.

Lemma 4.11. *Any n vertex graph with minimum degree d_{min} has diameter at most $3n/d_{\mathsf{min}}$.*

Proof. Let ρ be a shortest path between two vertices at distance $\operatorname{diam}(G)$. Then $\sum_{x \in V(P)} d(x) = \sum_{y \in V(G)} |\Gamma(y) \cap V(\rho)|$. Since ρ is a shortest path, each vertex y is adjacent to at most three vertices of ρ. It follows that $d_{\mathsf{min}} \operatorname{diam}(G) \leqslant 3n$, which gives the required bound. □

Lemma 4.11 in combination with Proposition 4.7 shows that the ratio $d_{\mathsf{min}} : d_{\mathsf{max}}$ of the extremal stationary distribution of reversible (c, C)-Simple walks on everywhere dense graphs is bounded. We wish to prove the same statement for non-reversible (c, C)-simple walks; notice that this result will cover a class of graphs where the mixing time can be as large as $\Theta(n^2)$.

Theorem 4.12. *For any graph G with minimum degree $d_{\mathsf{min}} \geqslant \alpha \cdot n$ for some constant $\alpha > 0$ then any (c, C)-simple walk satisfies $\frac{\kappa}{n} \leqslant \pi(u) \leqslant \frac{C}{\alpha n}$ for any $u \in V$, where $\kappa(c, C, \alpha) > 0$.*

The proof of this will make use of the following two results from a recent unpublished work by Patel, Sauerwald and Sudholt [17]:

Theorem 4.13. *Let $G = (V, E)$ be a graph on n vertices with minimum degree $\delta \geqslant \alpha n$ for some constant $0 < \alpha \leqslant 1$ and let $r_0 := 2^{r_0+1}$ where $r_0 := \left\lceil \log_{(1+(\alpha/8))} \left(\frac{1}{\alpha}\right) \right\rceil$. Then there exist constants $c_s, c_d, \beta > 0$, all depending only on α, and a partition of V into $k \leqslant r_0$ disjoint sets $U_1, U_2, \dots, U_k$ with the following properties:*

(i) for every $1 \leqslant i \leqslant k$, $|U_i| \geqslant c_s n$;

(ii) for every $1 \leqslant i \leqslant k$, $\delta_{U_i} \geqslant c_d n$; and

(iii) for every $1 \leqslant i \leqslant k$, the induced subgraph $G[U_i]$ of U_i is a β-expander.

The above theorem says there is a decomposition of any uniformly dense graph into a finite collection of uniformly dense expanders. The following theorem says that any one of these expanders will have small mixing time.

Theorem 4.14. *For every graph $G = (V, E)$ with minimum degree $\alpha \cdot n$ for some constant $\alpha > 0$, the mixing time $t_{mix}(G)$ of the lazy random walk on G satisfies*

$$t_{mix}(G) = O\left(\frac{1}{\Phi(G)}\right).$$

We shall also need the following basic lemma.

Lemma 4.15. *Let $A = \{v \in V : \pi(v) > \kappa/n\}$ and $\pi_{max} \leqslant \frac{C}{n}$ for some C, κ. Then $|A| \geqslant \frac{1-\kappa}{C-\kappa} n$.*

Proof. Observe that

$$1 = \sum_{x \in A} \pi(x) + \sum_{y \in A^c} \pi(y) \leqslant \frac{C}{n}|A| + \frac{\kappa}{n}(n - |A|),$$

rearranging this gives $n(1-\kappa) \leqslant |A|(C-\kappa)$, the result follows. □

In order to prove Theorem 4.12 we need a few more definitions. For two sets $A, B \subset V$ the *ergodic flow* $Q(A, B)$ is given by

$$Q(A, B) = \sum_{a \in A, b \in B} \pi(a) p_{a,b}. \tag{9}$$

For any set A and any Markov chain P, we have $Q(A, A^c) = Q(A^c, A)$. Let

$$\rho_{U_1}(u) = \frac{\sum_{v \in U_1} \pi(v) q_{v,u}}{\pi(U_1^c)} \qquad \text{for} \qquad u \in U_1^c, \tag{10}$$

be the *ergodic exit distribution* from U_1 [1, (3.65)].

Proof of Theorem 4.12. Recall that $G = (V, E)$ is an everywhere-dense graph, where the minimum degree satisfies $\delta \geqslant \alpha \cdot n$ for some constant $\alpha > 0$ and Q is a (c, C)-simple random walk for some $0 < c \leqslant 1$ and $1 \leqslant C < \infty$. The upper bound follows from Proposition 4.10, since if $d_{\mathsf{min}} \geqslant \alpha n$ then this shows that $\max_{u \in V} \pi_Q(u) \leqslant C/\alpha n$. The proof of the lower bound is more technical, let $\kappa, c_1, c_2, \dots > 0$ and $C_1, C_2 \dots < \infty$ be constants defined and used later.

We shall now begin the proof for the lower bound. By Theorem 4.13 G has a decomposition into uniformly dense induced expanders $G[U_1], \dots, G[U_k]$, where $k(\alpha) < \infty$. We shall show by induction on the number of expanders that each has minimum stationary probability κ/n.

Claim. *If there are $c_1 \cdot n$ vertices with $\pi(u) \geqslant c_2/n$ in a graph with a decomposition into uniformly dense induced expanders then at least one has minimum stationary probability at least c_3/n.*

Proof of claim. By the pigeonhole principal, there must be at least one expander $G[U_i]$ containing a subset $U \subseteq U_i$ of vertices with stationary distribution at least c_2/n where $|U| \geqslant c_4 n$. Now since $G[U_i]$ is an expander by Theorem 4.14 the mixing time t_{mix} of $G[U_i]$ is bounded and so is the separation time, say $t_{\mathsf{sep}} \leqslant C_1$. Thus for any $u, v \in U_i$ the LRW satisfies $p_{u,v}^{(C_1)} \geqslant (1-1/e)\pi(v) \geqslant \alpha/(2n)$, recall also that since Q is (c, C)-simple we have $q_{u,v}^{(C_1)} \geqslant c^{C_1} \cdot p_{u,v}^{(C_1)}$. Thus for any $u \in U_i$,

$$\pi(u) = \sum_{v \in V} \pi(v) \cdot P_{v,u}^{(C_1)} \geqslant \sum_{x \in U} \frac{c_2}{n} \cdot \frac{c^{C_1}\alpha}{2n} \geqslant c_4 n \cdot \frac{c_2}{n} \cdot \frac{c^{C_1}\alpha}{2n} \geqslant \frac{c_3}{n},$$

for some $c_3 > 0$ as claimed. ♢

Now, for the base case we know that $\pi_{max} \leqslant C/(\alpha n)$ and so by Lemma 4.15 there are at least $c_5 n$ vertices $v \in V$ such that $\pi(v) \geqslant c_6/n$. By the claim there is at least one induced expander with minimum stationary probability at least c_7/n. If there was only one expander in the decomposition of G we are done, so assume there was more than. Let $A \subset V$ be a set with $\min_{x \in A} \pi(x) \geqslant c_7 \cdot n$ and $A^c = U_1 \cup \dots \cup U_\ell$ be the remainder of G where $G[U_1], \dots, G[U_\ell]$ are uniformly dense expanders.

Thus if we start the random walk Q at any vertex $u \in A^c$ then, since every vertex in $G[A^c]$ has degree at least $c_8 \cdot n$, with probability at least c_8 the walk will not go back to A immediately. Furthermore, since $q_{x,y} \leqslant C/\alpha n$ for all $x, y \in V$, for any $t \geqslant 1$ the random walk's distribution $q_{u,\cdot}^{(t)}$ is always bounded by C_2/n pointwise. Hence at any time $t \geqslant 1$, the probability for the walk to return to A is upper bounded by

$$\sum_{x \in A^c} q_{u,x}^{(t)} \cdot q_{x,A} \leqslant \sum_{x \in A^c} \frac{C_2}{n} \cdot \frac{C \deg_A(x)}{\deg(x)} \leqslant \frac{C_2}{c_8} \cdot \frac{C|E(A, A^c)|}{n^2} = C_3 \cdot \Phi(A).$$

It follows that $\mathbb{E}_u[\tau_A] \geqslant c_8/(C_3\Phi(A)) = c_9/\Phi(A)$ for any $u \in A^c$, and thus for any probability distribution μ on A^c we have $\mathbb{E}_\mu[\tau_A] \geqslant c_9/\Phi(A)$. In particular let ρ_A be the ergodic exit distribution from A (10) then $\mathbb{E}_{\rho_A}[\tau_A] \geqslant c_9/\Phi(A)$, however $\mathbb{E}_{\rho_A}[\tau_A] = \pi(A^c)/Q(A^c, A)$ by [1, (3.69)], where $Q(A, A^c)$ is the ergodic flow (9). Thus

$$\pi(A^c)/Q(A^c, A) \geqslant c_9/\Phi(A). \tag{11}$$

Observe that by the definition of ergodic flow

$$Q(A, A^c) = \sum_{u \in A} \sum_{v \in A^c} \pi(u) \cdot q_{u,v} \geqslant \sum_{u \in A} \sum_{v \in A^c} \frac{c_7}{n} \cdot \frac{c \cdot \mathbf{1}(uv \in E)}{n} \geqslant c_{10}\Phi(A). \tag{12}$$

By combining (11) and (12) we have $\pi(A^c) \geqslant c_9 \cdot c_{10}$. Thus since $\pi_{\mathsf{max}} \leqslant C/(\alpha n)$ there at least $c_{11}n$ vertices with stationary distribution at least c_{12}/n in A^c and thus by the claim at least one of the induced expanders $G[U_1], \dots, G[U_\ell]$ has minimum stationary probability c_{13}/n. Now we can just redefine the set A as the union of the (former) set A and the new expander(s) and continue this argument inductively until we have considered all expanders. $\square$

The Markov chain exhibited in the proof of Theorem 4.12 satisfies an "approximate" notion of reversibility, i.e., for any $u, v \in V$ we have both $\pi_u \approx \pi_v$ as well as $P_{u,v} \approx P_{v,u}$. Note that if we drop the assumption that $P_{u,v} \approx P_{v,u}$ and only assume that every pair u, v satisfies $P_{u,v} = O(1/n)$, then it is possible to construct chains with $\pi_{\min} = O(1/n^2)$. For example, take two cliques of size $n/2$, and add $n/2$ directed edges from one clique to the other, but only add one directed edge in the opposite direction.

4.4 Biasing in d-Regular Graphs with $\varepsilon = \Theta(1/d)$

Referring to the ε-BRW on d-regular graphs, Azar et al. [4] state,

> *The interesting situation is when ε is not substantially larger than $1/d$; otherwise, the process is dominated by the controller's strategy.*

Recall the AKBLP Conjecture states that a controller can boost the stationary probability of any vertex from p to $p^{1-\varepsilon}$ and notice that for d-regular graphs with $d = \omega(\log n)$ this boost from p to $p^{1-\varepsilon}$ does not change the stationary probabilities by more than a constant factor. For this reason we shall focus on the following question for d-regular graphs with $\varepsilon = \Theta(1/d)$ of

When can we change the stationary distribution by more than a constant factor? (13)

As noted when $d = \omega(\log n)$ this question is stronger than the AKBLP conjecture and we think it is quite natural. We begin with a corollary of Proposition 4.12 which shows the answer to (13) is negative for uniformly dense graphs.

Corollary 4.16 (Corollary of Theorem 4.12)**.** *For any graph G with minimum degree $d_{\mathsf{min}} \geqslant \alpha \cdot n$ for some constant $\alpha > 0$ then controller of an ε-biased walk with $\epsilon = O(1/n)$ cannot change the stationary distribution by more than a constant factor.*

Proof. The ε-BRW is a (c, C)-simple walk in this regime, for some $0 < c, C < \infty$. $\square$

We now consider graphs of lower degree. Our first example shows that for $d = \mathsf{poly}(n)$ being arbitrarily close to linear, there are graphs for which we can answer (13) in the affirmative. These graphs do not only have the largest possible diameter $\approx n/d$, they also feature several bottlenecks.

Proposition 4.17. *Fix any $0 < \alpha < 1$ and let $d = n^\alpha$, $\varepsilon = \Theta(1/d)$. Then there exists a d-regular graph for which the stationary distribution p of any given vertex can be boosted by the ε-TB random walk to $\Omega(p^\alpha)$.*

Proof. Let $d = n^\alpha$ and $\ell = n^{1-\alpha}$ and consider the $(\ell, K_{d,d})$-ring pictured in Figure 2. The $(\ell, K_{d,d})$-ring has $N = 2\ell(d+1)$ many vertices and is $d+1$-regular graph, thus in our case $N \sim 2n$.

Let x, u be the end points of one of the edges which connects two $K_{d,d}$'s, and $u_1, \dots, u_d$ be the vertices in the $K_{d,d}$ attached to u (see picture). Assuming that x is closer to the target vertex we wish to boost the ε-TB strategy is clear: we should prefer the walk at u to visit x

and thus set $B_{u,x} = 1$ and $B_{u,u_i} = 0$, for all $1 \leqslant i \leqslant d$ where B is the bias matrix. Now we see that

$$\frac{w(u,x)}{w(u,u_i)} = \frac{\varepsilon + (1-\varepsilon)/(d+1)}{(1-\varepsilon)/(d+1)} = 1 + \frac{\varepsilon(d+1)}{(1-\varepsilon)} = 1 + \Omega(1).$$

We seek to bound the total weight W. If we sum from the target v, where we set the adjacent weights to one, then we see that the weights in i^{th} $K_{d,d}$ away from v must have weights that are at most $(1+\Omega(1))^{-i}$, thus

$$W \leqslant 2\sum_{i=0}^{\ell}(1+\Omega(1))^{-i}(d^2+2d+1) = \mathcal{O}(d^2).$$

Now we see a boosting since the original SRW stationary distribution of v was $p = 1/N \sim 1/2n$ however under this ε-TB boosting strategy this is now p' where

$$p' \geqslant d/\mathcal{O}(d^2) = \Omega(1/d) = \Omega(n^{-\alpha}).$$

□

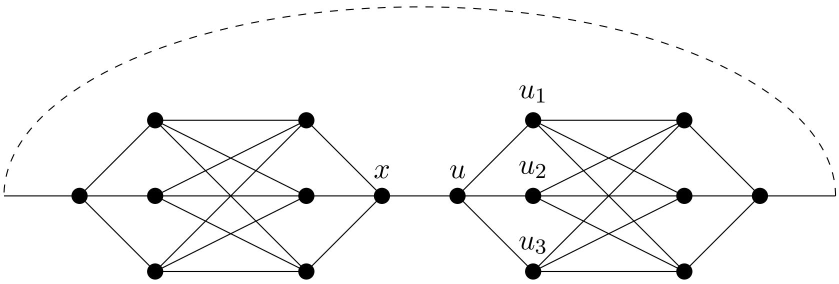

Figure 2: The $(\ell, K_{d,d})$-ring consists of ℓ complete bipartite graphs on d vertices arranged in a cycle. The $(\ell, K_{3,3})$-ring, for some $\ell \geqslant 2$, is shown above.

Our next example show that for the Erdős-Rényi random graph the answer to (13) is negative.

Proposition 4.18. *Let $\beta > 0$ and $\mathcal{G} \overset{d}{\sim} \mathcal{G}(n,p)$ where $np \sim n^a$ for some $a > 0$. Then w.h.p.*

(i) $t_{\mathsf{mix}}(\mathcal{G}) \leqslant \lceil 1/a \rceil$.

(ii) The (β/np)-BRW can only increase the stationary probability of any vertex from p to at most $p^{1-\delta}$ where $\delta \leqslant (4\lceil 1/a\rceil + o(1))\ln(1+\beta)/\ln(n)$.

(iii) The (β/np)-BRW can only decrease the stationary probability of any vertex from p to at least $p^{1+\delta}$ where $\delta \leqslant (8\lceil 1/a\rceil + o(1))\ln(n)/(np)$.

Proof. Hildebrand [13] showed that for $np = (\log n)^{\alpha}$, where $\alpha > 2$ and any $\eta > 0$ the mixing time of $\mathcal{G}(n,p)$ is bounded from above by $\left\lceil \frac{\log n}{\log(np)}(1+\eta)\right\rceil$ w.h.p.. Hildebrand's proof can be adapted so that it holds for any $np = n^a$, $a > 0$. Similarly this gives $t_{\mathsf{mix}} \leqslant \lceil (1+\eta)/a\rceil$ w.h.p., since all degrees are in the range $np \pm 3\sqrt{np\log n}$ w.p. $1 - ne^{-3\log n}$ by Chernoff bounds. If follows that for any fixed $a > 0$ such that $a \neq \lceil 1/a \rceil$ there is an $\eta > 0$ such that if we take n large enough $\lceil (1+\eta)/a\rceil = \lceil 1/a\rceil$. This result is tight as it agrees with the diameter [7].

Let $\varepsilon = \beta/np$ and observe that under any strategy $P_{x,y} \leqslant \frac{\beta}{np} + \frac{1-\beta/np}{d(x)} = \frac{1+\beta+o(1)}{d(x)}$ and similarly $P_{x,y} \geqslant \frac{1-\beta/np}{d(x)} = \frac{1}{d(x)}\left(1 - \frac{\beta+o(1)}{np}\right)$. Thus we can apply Proposition 4.8 with $\gamma = 1 + o(1)$,

$c = 1 + \beta + o(1)$ and $\tau = 4\lceil 1/a \rceil$. Thus for any u and any (β/np)-Bias strategy Q,

$$(1 - o(1))\frac{(1 - (\beta + o(1))/np)^{4\lceil 1/\alpha \rceil}}{2n} \leqslant \pi_Q(u) \leqslant \frac{2(1 + \beta + o(1))^{4\lceil 1/\alpha \rceil}}{n}.$$

Assume that there is a δ so that the stationary probability of any vertex can be raised from p to $p^{1-\delta}$, thus there is a strategy giving a transition matrix Q such that $\pi_Q(u) \cdot n^{1-\delta} \geqslant 1 - o(1)$. Our bound on $\pi_Q(u)$ implies that $\delta \ln n \leqslant (4\lceil 1/a \rceil + o(1)) \ln(1 + \beta + o(1))$.

Similarly if the controller seeks to decrease the stationary probability to $p^{1+\delta}$ then we have $\delta \ln n \leqslant -(4\lceil 1/a \rceil + o(1)) \ln(1 - (\beta + o(1))/np)$. Recall that $\ln(1 - x) \leqslant -x/\sqrt{1 - x}$ for $x \in (-1, 0]$ and thus $\delta \leqslant (4\beta\lceil 1/a \rceil + o(1))\frac{\ln n}{np}$. □

5 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 strategy for the ε-TBRW which achieves the given objective in optimal expected time? A strategy consists of a family of controller bias matrices $\{\mathbf{B}(\mathcal{H}_t)\}$, where $t \geqslant 0$ is the time and $\mathcal{H}_t$ is the history of the walk up to time t. Azar et al. [4] considered the following computational problems

`Stat` (G, w): Find an ε-bias strategy min/maximising $\sum_{v \in V} w_v \cdot \pi_v$ for vertex weights $w_v \geqslant 0$.

`Hit` (G, v, S): Find an ε-bias strategy minimising $\sum_{v \in V} \ell_v \cdot H_v^{\varepsilon\mathsf{B}}(S)$ for a given $S \subseteq V(G)$, $v \in V(G)$ and vertex weights $\ell_v \geqslant 0$.

Notice that for `Stat` to make sense we must fix an unchanging strategy and there exists an unchanging optimal strategy for `Hit`, see Section 1. Azar et al. showed `Stat`, `Hit` $\in \mathsf{P}$.

Theorem 5.1 (Theorems 6 & 12 in [4]). *Let G be any connected directed graph, $v \in V(G)$ and $S \subseteq V(G)$. Then* `Stat` (G, w) *and* `Hit` (G, v, S) *can be solved in polynomial time.*

We introduce the following computational problem not considered by Azar et al.

`Cov` (G, v): Find an ε-TB strategy minimising $C_v^{\varepsilon\mathsf{TB}}(G)$ for a given $v \in V(G)$.

Unlike for `Stat` and `Hit`, an optimal strategy for `Cov` on essentially any graph cannot to be unchanging as it will need to adapt to the knowledge of which vertices remain uncovered. Proposition 2.1 shows that there is an optimal strategy for `Cov` which is conditionally independent of time, in that no more information from $\mathcal{H}_t$ than the set of uncovered vertices is used. This fact means that an optimal strategy for `Cov` can be described using only finitely many bias matrices.

We show that the ε-TBRW exhibits the same dichotomy as the CRW studied in [10], in that while optimising `Hit` admits a polynomial-time algorithm, even computing an individual bias matrix $\mathbf{B}(\mathcal{H}_t)$ from an optimal strategy for `Cov` is NP-hard. We may view this as an on-line approach to solving `Cov`, where we compute only the specific bias matrices needed as the random walk progresses; clearly this is an easier problem than precomputing an entire optimal strategy. Note that at most n bias matrices will need to be computed in the course of any given walk, since an optimal bias matrix only depends on the uncovered set, which changes at most n times; however, a full optimal strategy may require exponentially many such matrices.

Our proof will be similar to that of [10] for the CRW, but the dependence of the construction for the ε-TBRW on ε creates some additional difficulties.

5.1 A Hardness Result for $\texttt{Cov}\,(G, v)$

We show that in general even computing a single optimal bias matrix for $\texttt{Cov}\,(G, v)$ is NP-hard. To that end we introduce the following problem, which in fact represents computing a single row of a matrix. The input is a graph G, a current vertex u, and a visited set X which must be connected and contain u.

$\texttt{NextStep}\,(G, u, X)$: Output a probability distribution over the neighbours of u so the next step minimises the expected time for the ε-TBRW to visit every vertex not in X, assuming an optimal strategy is followed thereafter.

Any such problem may arise during the ε-TBRW on G starting from any vertex in X, no matter what strategy was followed up to that point, since with positive probability the bias coin did not allow the controlled to influence any previous walk steps.

Before giving the proof of hardness, we need some lemmas on the performance of the ε-BRW on paths.

Lemma 5.2. *For any $\varepsilon > 0$, the ε-BRW started from the end point 0 of path hits a vertex at distance k in expected time*

$$\frac{k}{\varepsilon} - \frac{1+\varepsilon}{2\varepsilon^2} + \mathcal{O}\left(\frac{1}{\varepsilon^2}\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^{k+1}\right),$$

with $\min\{1/\varepsilon, k\}$ many expected returns to the origin 0 before first hitting k.

Proof. This is a birth-death chain on $\{0, \dots, k\}$, we shall follow the notation from [15, Sec. 2.5]. This chain has probability $p_k = (1+\varepsilon)/2$ of going from $k-1$ to k and probability $p_k = (1-\varepsilon)/2$ of going from k to $k-1$ and thus has weights $w_k = \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^k$. Thus the expected hitting time of k from 0 is given by

$$\begin{aligned}\sum_{i=1}^{k}\mathbb{E}_{i-1}\left[\,\tau_i\,\right] &= \sum_{i=1}^{k}\frac{2}{1-\varepsilon}\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^i\sum_{j=0}^{i-1}\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^j\\ &= \sum_{i=1}^{k}\frac{1}{\varepsilon}\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^i\left(\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^i - 1\right)\\ &= \frac{k}{\varepsilon} - \frac{1+\varepsilon}{2\varepsilon^2} + \mathcal{O}\left(\frac{1}{\varepsilon^2}\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^{k+1}\right),\end{aligned}$$

as claimed. For returns, since there are no cycles in a path the ε-BRW is reversible. Thus if we fix the edge connecting the endpoint to the first vertex of the path to have resistance 1 then the effective resistance between 0 and the target is $\sum_{i=0}^{k-1}\left(\frac{1-\varepsilon}{1+\varepsilon}\right) \leqslant \frac{1+\varepsilon}{2\varepsilon} \leqslant 1/\varepsilon$. Since there is only one edge from the end point of the path this also bounds the expected number of returns before hitting k by [15, Prop. 9.5]. This bound is not good when ε is small, in any case however the bias only decreases resistance the expected number of returns is at most k. □

We define the k-subdivision of a graph G to be the graph G' where every edge is replaced by a path of length k. We call $V(G)$ the branch vertices of G'.

Lemma 5.3. *Let G' be the k-subdivision of any n-vertex graph and let $k \geqslant 3\ln(n/\varepsilon)/\ln\left(\frac{1+\varepsilon}{1-\varepsilon}\right)$. Then there exists an ε-bias strategy for any $u \in V(G')$ such that the expected return time to u is bounded by $2\min\{1/\varepsilon, k\}$. For the ε bias strategy going from one branch vertex x to another y where $xy \in E(G)$ the probability the walk hits some $z \in V(G)$ before y is at most $d(x)\frac{1+\varepsilon}{1-\varepsilon}\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^k$.*

Proof. For a given vertex u we take the following reversible unchanging strategy: for all vertices on interior of the $d(u)$ paths of length k adjacent to u we bias towards u. On the rest of the graph we do not bias, simulating a simple random walk. This strategy is reversible since the paths are all of the same length k and weights are uniform over the rest of the graph. If we set the weight of edges adjacent to u to be 1 then the weight of edges at the ends of the paths adjacent to u is $w = \left(\frac{1-\varepsilon}{1+\varepsilon}\right)^k$, all edges not on a path with end point u also have this weight. Since the return time is the reciprocal of the stationary distribution and this is given by $W/d(u)$ where W is the total weight of the edges. Note the contribution from paths adjacent to u is

$$W_1 = d(u)\sum_{i=0}^{k-1}\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^i = \frac{d(u)\,(1+\varepsilon)}{2\varepsilon}\left(1-\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^k\right) \geqslant 1, \tag{14}$$

and the contribution W_2 from the rest of the graph is at most $n^2k\cdot\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^k$. It follows by (14) that if we take $k \geqslant 2\ln(nk)/\ln\left(\frac{1+\varepsilon}{1-\varepsilon}\right)$ then $W_1 \geqslant W_2$ and the return time is at most $2W_1/d(u)$. One can check that any $k \geqslant 3\ln(n/\varepsilon)/\ln\left(\frac{1+\varepsilon}{1-\varepsilon}\right)$ satisfies this. For such a k we have $W_1 \leqslant d(u)\min\{1/\varepsilon, k\}$ by (14), thus the return time is at most $2\min\{1/\varepsilon, k\}$.

For the second result it suffices to consider the neighbourhood of x in G in the statement using the same weights as above and interpreting them as electrical resistances the path to y from x has resistance $R(x,y) = \sum_{i=0}^{k-1}\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^i = \frac{1+\varepsilon}{2\varepsilon}\left(1-\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^k\right)$ and each of the $d(x)-1$ as having resistance $R(x,z) = \sum_{i=0}^{k-1}\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^i = \frac{1-\varepsilon}{2\varepsilon}\left(\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^k - 1\right)$. Thus if we identify all $z \in V(G)\backslash\{y\}$ such that $xz \in E(G)$ as a single vertex $\bar{z}$ then the effective resistance $R(x,\bar{z}) = R(x,z)/(d(x)-1)$. Then if we consider the path $\bar{z}, x, y$ and impose a voltage of 1 at $\bar{z}$ and 0 at y then the voltage at x is $\frac{R(x,y)}{R(x,\bar{z})+R(x,y)}$, this is the probability the walk from x hits $\bar{z}$ before y. After some cancellations we we have

$$\frac{R(x,y)}{R(x,\bar{z})+R(x,y)} = \frac{\frac{1+\varepsilon}{2\varepsilon}}{\frac{1-\varepsilon}{2\varepsilon(d(x)-1)}\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^k + \frac{1+\varepsilon}{2\varepsilon}} \leqslant d(x)\frac{1+\varepsilon}{1-\varepsilon}\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^k.$$

The result follows. □

We are now ready to prove the main result of this section.

Theorem 5.4. *Provided $\varepsilon \geqslant 1/|V(G)|^k$, for some fixed $k < \infty$,* `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 [9].

Given an n-vertex graph H, construct the graph G as follows. First take the k-subdivision of H, for some $k \geqslant C\Big(\frac{\ln(n/\varepsilon)}{\ln\left(\frac{1+\varepsilon}{1-\varepsilon}\right)} + 1\Big)$ where $C \geqslant 3$ determined later. Next add a new pendant path of length k^2 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. This construction is adapted from that of [10], but moving to the ε-TBRW necessitates the dependence on ε.

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. Recalling $k \geqslant C\Big(\frac{\ln(n/\varepsilon)}{\ln\left(\frac{1+\varepsilon}{1-\varepsilon}\right)} + 1\Big)$, for some $C \geqslant 3$, we make the following claim.

Claim. *Any optimal next step is to move towards the next vertex on some such path.*

Proof of claim. To prove the claim, first we argue by induction that there is a strategy to visit every vertex in $|Y|$ in expected time at most $\big(k/\varepsilon + 5/\varepsilon^2\big)\,|Y|$. 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 the strategy of biasing towards it gives an expected time at most $k/(2\varepsilon)$ to reach z by Lemma 5.2 plus an additional expected time at most $2/\varepsilon$ for each visit to u by Lemma 5.3, of which we expect $1/\varepsilon$ by Lemma 5.2, giving a total expected time of $k/(2\varepsilon) + 2/\varepsilon^2$. Note that if the walker is forced to a different branch vertex first, the expected time to return from this point is $\mathcal{O}(k/\varepsilon)$ by Lemmas 5.2 & 5.3, but this event occurs with probability at most $d(u)\frac{1+\varepsilon}{1-\varepsilon}\big(\frac{1-\varepsilon}{1+\varepsilon}$
$bigr)^k$ by Lemma 5.3 so the contribution to the expectation from this is occurrence is $\mathcal{O}(\varepsilon/n)$, negligible. Similarly, the time taken to reach y from z is $k/(2\varepsilon)+2/\varepsilon^2$. 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 aims to reach y' before another element of H. The expected time taken to vertex of H from a another vertex in H is at least $\frac{k}{\varepsilon}\,(1 - 4/C)$ by Lemma 5.2. To see this notice that for our choice of k, $k\varepsilon > C/4$ since $\ln(\frac{1+\varepsilon}{1-\varepsilon}) \geqslant \frac{\varepsilon}{1-\varepsilon}$ if $\varepsilon \leqslant 1/2$. To conclude since this strategy did not follow a Hamiltonian path it must make at least $|Y|+1$ such crossings. The result follows by taking $C \geqslant 3$ sufficiently large. $\Diamond$

To conclude, by the 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. Finally we observe that by the restrictions placed on ε, there is some satisfying k such that the graph G has at most $poly(n)$ many vertices. $\square$

Remark. *The restriction $\varepsilon \geqslant 1/|V|^k$ is not really significant as if ε is less than this for some large k then w.h.p. $1-\mathcal{O}\big(n^{-k}\big)$ no bias step will even be taken. It can then be shown that such a small ε would only affect the expected cover time by at most an additive $\mathcal{O}\big(n^{-k+3}\big)$ term which is negligible.*

5.2 Computing $\texttt{Cov}(G,v)$ via Markov Decision Processes

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

Lemma 5.5. *[Lemma 7.7 of [10]] For any graph $G = (V, E)$ let the (directed) auxiliary graph $\widetilde{G} = (\widetilde{V}, \widetilde{E})$ be given by $\widetilde{V} = V \times \mathcal{P}(V)$ and $\widetilde{E} = \{((i, S), (j, S \cup j)) \mid ij \in E\}$. Then solutions to $\texttt{Cov}(G,v)$ correspond to solutions to $\texttt{Hit}\big(\widetilde{G}, \widetilde{v}, W\big)$ and vice versa, where $W = \{(u, V) \mid u \in V\}$.*

We can now prove Proposition 2.1 which states there is an optimal strategy to cover G which is fixed over any time interval between times when a new vertex is visited.

Proof of Proposition 2.1. We shall appeal to Lemma 5.5 and consider the problem of covering G as hitting the set W in the auxiliary graph $\widetilde{G}$. This is now an instance of the optimal first-passage problem in the context of Markov Decision Processes [8], and the existence of a time independence optimal strategy follows from [8, Thm. 3, Ch. 3]. Notice that although the strategy for hitting the vertex W in $\widetilde{G}$ is independent of time this is not strictly true of the original problem. Recall $\widetilde{G}$ is a directed graph which consists of a series of undirected graphs linked by directed edges, the undirected graphs represent the subgraphs of G induced by possible visited sets and the directed edges correspond to the walk in G visiting a new vertex. Since the strategy for $\widetilde{G}$ is independent of time during the times when a new vertex is added to the covered set the strategy is fixed. $\square$

In light of Lemma 5.5 we can solve $\texttt{Cov}(G,v)$.

Corollary 5.6. *For any graph G and $v \in V$ an optimal policy for the problem $\texttt{Cov}(G,v)$ can be computed, in particular $\texttt{Cov} \in \mathsf{EXP}$.*

Proof. We first encode the problem $\texttt{Cov}(G,v)$ as the problem $\texttt{Hit}\big(\widetilde{G}, \widetilde{v}, W\big)$ as described in Lemma 5.5. The problem $\texttt{Hit}\big(\widetilde{G}, \widetilde{v}, W\big)$ for any directed graph $\widetilde{G}$ can be computed in time which is polynomial in $|V(\widetilde{G})| \leqslant 2^n$ by Theorem 5.1. $\square$

6 Concluding Remarks and Open Problems

In this paper we have extended the previous work on the ε-biased random walk to include strategies which may depend on the history of the walk. Our motivation for this is the cover time problem for which we obtain bounds using a new technique we refer to as the tree gadget. This tree gadget also allowed us to make progress on a conjecture of Azar et al. [4]. We note that this conjecture as originally stated is actually wrong, see Section 4.3. However, their conjecture only appears to fail for graphs with large entries in the stationary vector and we believe that, as noted in Section 4.3, the following slightly refined version of their conjecture should hold.

Conjecture 6.1. *In any graph a controller can increase the stationary probability of any vertex from p to $p^{1-\varepsilon+\delta}$, where $\delta = \mathcal{O}(1/|\log p|)$.*

We also showed that computing an optimal next step for the ε-TBRW to take in the online version of the covering problem is NP-Hard. We believe that there should be no short certificate certifying that a given strategy for $\texttt{Cov}$ is optimal and thus the problem should not belong to the class NP.

Conjecture 6.2. *There exists an $\varepsilon > 0$ such that $\texttt{Cov}(G,v)$ is PSPACE-Hard.*

In addition we showed that $\texttt{Cov}(G,v) \in \mathsf{EXP}$, it would also be interesting to see if this can be improved.

Acknowledgements

J.H. was supported by ERC Starting Grant no. 639046 (RGGC). T.S. and J.S. were supported by ERC Starting Grant no. 679660 (DYNAMIC MARCH).

References

[1] D. Aldous and J. A. Fill. Reversible Markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014.

[2] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *20th Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 29-31 October 1979*, pages 218–223. IEEE Computer Society, 1979.

[3] N. Alon and M. O. Rabin. Biased coins and randomized algorithms. *Advances in Computing Research*, 5:499–507, 1989.

[4] Y. Azar, A. Z. Broder, A. R. Karlin, N. Linial, and S. Phillips. Biased random walks. *Combinatorica*, 16(1):1–18, 1996.

[5] M. Ben-Or and N. Linial. Collective coin flipping. In S. Micali, editor, *Randomness and Computation*, pages 91–115. Academic Press, New York, 1989.

[6] R. B. Boppana and B. O. Narayanan. The biased coin problem. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Theory of Computing*, STOC '93, pages 252–257, New York, NY, USA, 1993. ACM.

[7] F. Chung and L. Lu. The diameter of sparse random graphs. *Adv. in Appl. Math.*, 26(4):257–279, 2001.

[8] C. Derman. *Finite State Markovian Decision Processes.* Academic Press, Inc., Orlando, FL, USA, 1970.

[9] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoret. Comput. Sci.*, 1(3):237–267, 1976.

[10] A. Georgakopoulos, J. Haslegrave, T. Sauerwald, and J. Sylvester. The power of two choices for random walks. *arXiv:1911.05170*, 2019.

[11] A. Georgakopoulos, J. Haslegrave, T. Sauerwald, and J. Sylvester. Choice and bias in random walks. In T. Vidick, editor, *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPIcs*, pages 76:1–76:19. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2020.

[12] O. Goldreich. *A primer on pseudorandom generators*, volume 55 of *University Lecture Series.* American Mathematical Society, Providence, RI, 2010.

[13] M. Hildebrand. Random walks on random simple graphs. *Random Structures Algorithms*, 8(4):301–318, 1996.

[14] R. M. Karp. An introduction to randomized algorithms. *Discrete Applied Mathematics*, 34(1-3):165–201, 1991.

[15] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009. With a chapter by James G. Propp and David B. Wilson.

[16] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[17] V. Patel, T. Sauerwald, and D. Sudholt. Finding small cuts in dense graphs.

[18] R. Solovay and V. Strassen. A fast monte-carlo test for primality. *SIAM J. Comput.*, 6(1):84–85, 1977.

[19] U. V. Vazirani and V. V. Vazirani. Random polynomial time is equal to slightly-random polynomial time. In *26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, USA, 21-23 October 1985*, pages 417–428. IEEE Computer Society, 1985.