

Cover Time in Edge-Uniform Stochastically-Evolving Graphs^{*}

Ioannis Lamprou^(✉), Russell Martin, and Paul Spirakis

Department of Computer Science, University of Liverpool, Liverpool, UK
{Ioannis.Lamprou,Russell.Martin,P.Spirakis}@liverpool.ac.uk

Abstract. We define a new general model of stochastically evolving graphs, namely the *Edge-Uniform Stochastically-Evolving Graphs*. In this model, each possible edge of an underlying general static graph evolves independently being either alive or dead at each discrete time step of evolution following a (Markovian) stochastic rule. The stochastic rule is identical for each possible edge and may depend on the previous $k \geq 0$ observations of the edge's state.

We examine two kinds of random walks for a single agent taking place in such a dynamic graph: (i) The *Random Walk with a Delay (RWD)*, where at each step the agent chooses (uniformly at random) an incident possible edge (i.e. an incident edge in the underlying static graph) and then it waits till the edge becomes alive to traverse it. (ii) The more natural *Random Walk on what is Available (RWA)* where the agent only looks at alive incident edges at each time step and traverses one of them uniformly at random. Our study is on bounding the *cover time*, i.e. the expected time until each node is visited at least once by the agent.

For *RWD*, we provide the first upper bounds for the cases $k = 0, 1$ by correlating *RWD* with a simple random walk on a static graph. Moreover, we present a modified electrical network theory capturing the $k = 0$ case and a mixing-time argument toward an upper bound for the case $k = 1$. For *RWA*, we derive the first upper bounds for the cases $k = 0, 1$, too, by reducing *RWA* to an *RWD*-equivalent walk with a modified delay. Finally, for the case $k = 1$, we prove that when the underlying graph is complete, then the cover time is $\mathcal{O}(n \log n)$ (i.e. it matches the cover time on the static complete graph) under only a mild condition on the edge-existence probabilities determined by the stochastic rule.

Keywords: dynamic graphs, random walk, cover time, stochastically-evolving network, edge-independent

^{*} Regular Submission to SSS 2017

Corresponding Author: Ioannis Lamprou, E-mail: Ioannis.Lamprou@liverpool.ac.uk,
Address: Room 310, Ashton Building, Ashton Street, Computer Science Department,
University of Liverpool, L69 3BX, UK, Mobile: +447521714871

1 Introduction

In the modern era of Internet, modifications in a network topology can occur extremely frequently and in a disorderly way. Communication links may fail from time to time, while connections amongst terminals may appear or disappear intermittently. Thus, classical (static) network theory fails to capture such ever-changing processes. In an attempt to fill this void, different research communities have given rise to a variety of theories on *dynamic networks*. In the context of algorithms and distributed computing, such networks are usually referred to as *temporal graphs* [16]. A temporal graph is represented by a (possibly infinite) sequence of subgraphs of the same static graph. That is, the graph is *evolving* over a set of (discrete) time steps under a certain group of deterministic or stochastic rules of evolution. Such a rule can be edge- or graph-specific and may take as input some graph instances observed in previous time steps.

In this paper, we focus on stochastically evolving temporal graphs. We define a new model of evolution where there exists a single stochastic rule which is applied *independently* to each edge. Furthermore, our model is general in the sense that the underlying static graph is allowed to be a general connected graph, i.e. with no further constraints on its topology, and the stochastic rule can include any finite number of past observations.

Assume now that a single mobile agent is placed on an arbitrary node of a temporal graph evolving under the aforementioned model. Next, the agent performs a simple random walk; at each time step, after the graph instance is fixed according to the model, the agent chooses uniformly at random a node amongst the neighbors of its current node and visits it. The *cover time* of such a walk is the expected number of time steps until the agent has visited each node at least once. Herein, we prove some first bounds on the cover time for a simple random walk as defined above, mostly via the use of Markovian theory.

Random walks constitute a very important primitive in terms of distributed computing. Examples include their use in information dissemination [1] and random network structure [3]; also, see the short survey in [7]. In this work, we consider a single random walk as a fundamental building block for other more distributed scenarios to follow.

1.1 Related Work

A paper which is very relevant with respect to ours is the one of Clementi et al. [9], where they consider the flooding time in *Edge-Markovian* dynamic graphs. In such graphs, each edge independently follows an one-step Markovian rule and their model appears as a special case of ours (matches our case $k = 1$). Further work under this Edge-Markovian paradigm includes [4, 10].

Another work related to our paper is the one of Avin et al. [2] where they define the notion of a *Markovian Evolving Graph*, i.e. a temporal graph evolving over a set of graphs G_1, G_2, \dots , where the process transits from G_i to G_j with probability p_{ij} , and consider random walk cover times. Note that their approach becomes intractable if applied to our case; each of the possible edges evolves

independently, thence causing the state space to be of size 2^m , where m is the number of possible edges in our model.

Clementi et al. [11] study the broadcast problem when at each time step the graph is selected according to the well-known $G_{n,p}$ model. Also, Yamauchi et al. [21] study the rendezvous problem for two agents on a ring when each edge of the ring independently appears at every time step with some fixed probability p . Lastly, there exist a few papers considering random walks on different models of stochastic graphs, e.g. [15, 18, 19], but without considering the cover time.

In the analysis to follow, we employ several seminal results around the theory of random walks and Markov chains. For random walks, we base our analysis on the seminal work in [1] and the electrical network theory presented in [8, 12], while for results regarding the mixing time of a Markov chain we cite textbooks [14, 17].

1.2 Our Results

We define a new general model for stochastically evolving graphs where each possible edge evolves independently, but all of them evolve according to the same stochastic rule. Furthermore, the stochastic rule may take into account the last k states of a given edge. Special cases of our model have appeared in previous literature, e.g. in [11, 21] for $k = 0$ and in the line of work starting from [9] for $k = 1$, however they only consider special graph topologies (like ring and clique). On the other hand, the model we define is general in the sense that no assumptions, aside from connectivity, are made on the topology of the underlying graph and any amount of history is allowed into the stochastic rule. Thence, we believe it can be valued as a basis for more general results to follow capturing search or communication tasks in such dynamic graphs.

We hereby provide the first known upper bounds relative to the cover time of a simple random walk taking place in such stochastically evolving graphs for $k = 0$ and $k = 1$. To do so, we make use of a simple, yet fairly useful, modified random walk, namely the *Random Walk with a Delay (RWD)*, where at each time step the agent is choosing uniformly at random from the incident edges of the static underlying graph and then waits for the chosen edge to become alive in order to traverse it. Moreover, we consider the natural random walk on such graphs, namely the *Random Walk on What's Available (RWA)*, where at each time step the agent only considers the currently alive incident edges and chooses to traverse one out of them uniformly at random.

For the case $k = 0$, that is when each edge appears at each round with a fixed probability p , we prove that the cover time for *RWD* is upper bounded by $2m(n-1)/p$, where n (respectively m) is the number of vertices (respectively edges) of the underlying graph. The result can be obtained both by a careful mapping of the *RWD* walk to its corresponding simple random walk on the static graph and by generalizing the standard electrical network theory literature in [8, 12]. Later, we proceed to prove that the cover time for *RWA* is upper bounded by $2m(n-1)/(1-(1-p)^\delta)$ where δ is the min degree of the underlying graph. The main idea here is to reduce *RWA* to an *RWD* walk where at each step the traversal delay is lower bounded by $(1-(1-p)^\delta)$.

For $k = 1$, the stochastic rule takes into account the previous (one time step ago) state of the edge. If an edge were not present, then it becomes alive with probability p , whereas if it were alive, then it dies with probability q . Let τ_{mix} stand for the mixing time of this process. We prove that the *RWD* cover time is upper bound by $\tau_{mix} + 2m(n-1)(p^2 + q)/(p^2 + pq)$ by carefully computing the expected traversal delay at each step after mixing is attained. Moreover, we show another $2m(n-1)/\xi_{min}$ bound by considering the minimum probability guarantee of existence at each round, i.e. $\xi_{min} = \min\{p, 1-q\}$, and we discuss the trade-off between these two bounds. As far as *RWA* is concerned, we upper bound its cover time by $2m(n-1)/(1 - (1 - \xi_{min})^\delta)$ again by a reduction to an *RWD*-equivalent walk. Finally, we obtain a quite important result in the context of complete underlying graphs where we prove an upper bound of $\mathcal{O}(n \log n)$ (which matches the cover time for complete static graphs) under the soft restriction $\xi_{min} \in \Omega(\log n/n)$ via some cautious coupon-collector-type arguments.

1.3 Outline

In Section 2 we provide preliminary definitions and results regarding important concepts and tools that we use in later sections. Then, in Section 3, we define our model of stochastically evolving graphs in a more rigorous fashion. Afterwards, in Sections 4 and 5, we provide the analysis of our cover time upper bounds when for determining the current state of an edge we take into account its last 0 and 1 states, respectively. Finally, in Section 6, we cite some concluding remarks.

2 Preliminaries

Let us hereby define a few standard notions related to a simple random walk performed by a single agent on a simple connected graph $G = (V, E)$. By $d(v)$, we denote the degree (i.e. the number of neighbors) of a node $v \in V$. A simple random walk is a Markov chain where, for $v, u \in V$, we set $p_{vu} = 1/d(v)$, if $(v, u) \in E$, and $p_{vu} = 0$, otherwise. That is, an agent performing the walk chooses the next node to visit uniformly at random amongst the set of neighbors of its current node. Given two nodes v, u , the expected time for a random walk starting from v to arrive at u is called the *hitting time* from v to u and is denoted by H_{vu} . The *cover time* of a random walk is the expected time until the agent has visited each node of the graph at least once. Let P stand for the stochastic matrix describing the transition probabilities for a random walk (or, in general, a discrete-time Markov chain) where p_{ij} denotes the probability of transition from node i to node j , $p_{ij} \geq 0$ for all i, j and $\sum_j p_{ij} = 1$ for all i . Then, the matrix P^t consists of the transition probabilities to move from one node to another after t time steps and we denote the corresponding entries as $p_{ij}^{(t)}$. Asymptotically, $\lim_{t \rightarrow \infty} P^t$ is referred to as the *limiting distribution* of P . A *stationary distribution* for P is a row vector π such that $\pi P = \pi$ and $\sum_i \pi_i = 1$. That is, π is not altered after an application of P . If every state can be reached from another in a finite number of steps (i.e. P is *irreducible*) and the

transition probabilities do not exhibit periodic behavior with respect to time, i.e. $\gcd\{t : p_{ij}^{(t)} > 0\} = 1$, then the stationary distribution is *unique* and it matches the limiting distribution; this result is often referred to as the *Fundamental Theorem of Markov chains*. The *mixing time* is the expected number of time steps until a Markov chain approaches its stationary distribution. Below, let $p_i^{(t)}$ stand for the i -th row of P^t and $tvd(t) = \max_i \|p_i^{(t)} - \pi\| = \frac{1}{2} \max_i \sum_j |p_{ij}^{(t)} - \pi_j|$ stand for the *total variation distance* of the two distributions. We say that a Markov chain is ϵ -near to its stationary distribution at time t if $tvd(t) \leq \epsilon$. Then, we denote the mixing time by $\tau(\epsilon)$: the minimum value of t until a Markov chain is ϵ -near to its stationary distribution. A *coupling* (X_t, Y_t) is a joint stochastic process defined in a way such that X_t and Y_t are copies of the same Markov chain P when viewed marginally, and once $X_t = Y_t$ for some t , then $X_{t'} = Y_{t'}$ for any $t' \geq t$. Also, let T_{xy} stand for the minimum expected time until the two copies *meet*, i.e. until $X_t = Y_t$ for the first time, when starting from the initial states $X_0 = x$ and $Y_0 = y$. We can now state the following *Coupling Lemma* correlating the coupling meeting time to the mixing time:

Lemma 1 (Lemma 4.4 [14]). *Given any coupling (X_t, Y_t) , it holds $tvd(t) \leq \max_{x,y} Pr[T_{xy} \geq t]$. Consequently, if $\max_{x,y} Pr[T_{xy} \geq t] \leq \epsilon$, then $\tau(\epsilon) \leq t$.*

Furthermore, asymptotically, we need not care about the exact value of the total variation distance since, for any $\epsilon > 0$, we can force the chain to be ϵ -near to its stationary distribution after a multiplicative time of $\log \epsilon^{-1}$ steps due to the submultiplicativity of the total variation distance. Formally, it holds $tvd(kt) \leq (2 \cdot tvd(t))^k$.

Fact 1. *Suppose $\tau(\epsilon_0) \leq t$ for some Markov chain P and a constant $0 < \epsilon_0 < 1$. Then, for any $0 < \epsilon < \epsilon_0$, it holds $\tau(\epsilon) \leq t \log \epsilon^{-1}$.*

3 The Edge-Uniform Evolution Model

Let us define a novel model of a dynamically evolving graph. Let $G = (V, E)$ stand for a simple, *connected* graph, from now on referred to as the *underlying graph* of our model. The number of nodes is given by $n = |V|$, while the number of edges is denoted by $m = |E|$. For a node $v \in V$, let $N(v) = \{u : (v, u) \in E\}$ stand for the *open neighborhood* of v and $d(v) = |N(v)|$ for the (*static*) *degree* of v . Note that we make no assumptions regarding the topology of G besides connectedness. We refer to the edges of G as the *possible edges* of our model. We consider evolution over a sequence of discrete time steps (namely $0, 1, 2, \dots$) and denote by $\mathcal{G} = (G_0, G_1, G_2, \dots)$ the infinite sequence of graphs $G_t = (V_t, E_t)$ where $V_t = V$ and $E_t \subseteq E$. That is, G_t is the graph appearing at time step t and each edge $e \in E$ is either *alive* (if $e \in E_t$) or *dead* (if $e \notin E_t$) at time step t .

Let R stand for a *stochastic rule* dictating the probability that a given possible edge is alive at any time step. We apply R at each time step and at each edge *independently* to determine the set of currently alive edges, i.e. the rule is *uniform* with regard to the edges. In other words, let e_t stand for a random

variable where $e_t = 1$, if e is alive at time step t , or $e_t = 0$, otherwise. Then R determines the value of $Pr(e_t = 1|H_t)$ where H_t is also determined by R and denotes the history length (i.e. the values of e_{t-1}, e_{t-2}, \dots) considered when deciding for the existence of an edge at time step t . For instance, $H_t = \emptyset$ means no history is taken into account, while $H_t = \{e_{t-1}\}$ means the previous state of e is taken into account when deciding for its current state.

Overall, the aforementioned *Edge-Uniform Evolution* model (shortly *EUE*) is defined by the parameters G and R . In the following sections, we consider some special cases for R and provide first bounds for the cover time of G under this model. Each time step of evolution consists of two stages: in the first stage, the graph G_t is fixed for time step t following R , while in the second stage, the agent moves to a node in $N_t[v] = \{v\} \cup \{u \in V : (v, u) \in E_t\}$. Notice that, since G is connected, then the cover time under *EUE* is finite since R models edge-specific delays.

4 Cover Time with Zero-Step History

We hereby analyze the cover time of G under *EUE* in the special case when no history is taken into consideration for computing the probability that a given edge is alive at the current time step. Intuitively, each edge appears with a fixed probability p at every time step independently of the others. More formally, for all $e \in E$ and time steps t , $Pr(e_t = 1) = p \in [0, 1]$.

4.1 Random Walk with a Delay

A first approach toward covering G with a single agent is the following: The agent is randomly walking G as if all edges were present and, when an edge is not present, it just waits for it to appear in a following time step. More formally, suppose the agent arrives on a node $v \in V$ with (static) degree $d(v)$ at the second stage of time step t . Then, after the graph is fixed for time step $t+1$, the agent selects a neighbor of v , say $u \in N(v)$, uniformly at random, i.e. with probability $\frac{1}{d(v)}$. If $(v, u) \in E_{t+1}$, then the agent moves to u and repeats the above procedure. Otherwise, it remains on v until the first time step $t' > t+1$ such that $(v, u) \in E_{t'}$ and then moves to u . This way, p acts as a *delay* probability, since the agent follows the same random walk it would on a static graph, but with an expected delay of $\frac{1}{p}$ time steps at each node. Notice that, in order for such a strategy to be feasible, each node must maintain knowledge about its neighbors in the underlying graph; not just the currently alive ones. From now on, we refer to this strategy for the agent as the *Random Walk with a Delay* (shortly *RWD*).

Now, let us upper bound the cover time of *RWD* by exploiting its strong correlation to a simple random walk on the underlying graph G . Below, let C_G stand for the cover time of a simple random walk on the static graph G .

Theorem 1. *For any connected underlying graph G , the cover time under *RWD* is expectedly C_G/p .*

Proof. Consider a simple random walk, shortly *SRW*, and an *RWD* (under the *EUE* model) taking place on a given connected graph G . Given that *RWD* decides on the next node to visit uniformly at random based on the underlying graph, that is in exactly the same way *SRW* does, we use a coupling argument to enforce *RWD* and *SRW* to follow the exact same trajectory (i.e. sequence of visited nodes) in G .

Then, let the trajectory end when each node in G has been visited at least once and denote by T the total number of node transitions made by the agent. Such a trajectory under *SRW* will cover all nodes in expectedly $E[T] = C_G$ time steps. On the other hand, in the *RWD* case, for each transition we have to take into account the delay experienced until the chosen edge becomes available. Let $D_i \geq 1$ be a random variable where $1 \leq i \leq T$ standing for the actual delay corresponding to node transition i in the trajectory. Then, the expected number of time steps till the trajectory is realized is given by $E[D_1 + \dots + D_T]$. Since the random variables D_i are independent and identically distributed (by the edge-uniformity of our model), T is a stopping time for them and all of them have finite expectations, then we can apply Wald's Equation [20] to get $E[D_1 + \dots + D_T] = E[T] \cdot E[D_1] = C_G \cdot 1/p$. \square

For an explicit general bound on *RWD*, it suffices to use $C_G \leq 2m(n-1)$ proved by Aleliunas et al. in [1].

A Modified Electrical Network. Another way to analyze the above procedure is to make use of a modified version of the standard literature approach of electrical networks and random walks [8, 12]. This point of view gives us in addition expressions for the hitting time between any two nodes of the underlying graph. That is, we hereby (in Lemmata 2, 3 and Theorem 2) provide a generalization of the results given in [8, 12] thus correlating the hitting and commute times of *RWD* to an electrical network analog and reaching a conclusion for the cover time similar to the one of Theorem 1.

In particular, given the underlying graph G , we design an electrical network, $N(G)$, with the same edges as G , but where each edge has a resistance of $r = \frac{1}{p}$ ohms. Let $H_{u,v}$ stand for the hitting time from node u to node v in G , i.e. the expected number of time steps until the agent reaches v after starting from u and following *RWD*. Furthermore, let $\phi_{u,v}$ declare the electrical potential difference between nodes u and v in $N(G)$ when, for each $w \in V$, we inject $d(w)$ amperes of current into w and withdraw $2m$ amperes of current from a single node v . We now upper-bound the cover time of G under *RWD* by correlating $H_{u,v}$ to $\phi_{u,v}$.

Lemma 2. *For all $u, v \in V$, $H_{u,v} = \phi_{u,v}$ holds.*

In the lemma below, let $R_{u,v}$ stand for the *effective resistance* between u and v , i.e. the electrical potential difference induced when flowing a current of one ampere from u to v .

Lemma 3. *For all $u, v \in V$, $H_{u,v} + H_{v,u} = 2mR_{u,v}$ holds.*

Theorem 2. *For any connected underlying graph G , the cover time under the *RWD* is at most $2m(n-1)/p$.*

4.2 Random Walk on what's Available

Random Walk with a Delay does provide a nice connection to electrical network theory. However, depending on p , there could be long periods of time where the agent is simply standing still on the same node. Since the walk is random anyway, waiting for an edge to appear may not sound very wise. Hence, we now analyze the strategy of a *Random Walk on what's Available* (shortly *RWA*). That is, suppose the agent has just arrived at a node v after the second stage at time step t and then E_{t+1} is fixed after the first stage at time step $t + 1$. Now, the agent picks uniformly at random only amongst the alive edges at time step $t + 1$, i.e. with probability $\frac{1}{d_{t+1}(v)}$ where $d_{t+1}(v)$ stands for the degree of node v in G_{t+1} . The agent then follows the selected edge to complete the second stage of time step $t + 1$ and repeats the strategy. In a nutshell, the agent keeps moving randomly on available edges and only remains on the same node if no edge is alive at the current time step. Below, let $\delta = \min_{v \in V} d(v)$ and $\Delta = \max_{v \in V} d(v)$.

Theorem 3. *For any connected underlying graph G with min-degree δ , the cover time for RWA is at most $2m(n - 1)/(1 - (1 - p)^\delta)$.*

Proof. Suppose the agent follows *RWA* and has reached node $u \in V$ after time step t . Then, G_{t+1} becomes fixed and the agent selects uniformly at random a neighboring edge to move to. Let M_{uv} (where $v \in \{w \in V : (u, w) \in E\}$) stand for a random variable taking value 1 if the agent moves to node v and 0 otherwise. For $k = 1, 2, \dots, d(u) = d$, let A_k stand for the event that $d_{t+1}(u) = k$. Therefore, $Pr(A_k) = \binom{d}{k} p^k (1 - p)^{d-k}$ is exactly the probability k out of the d edges exist since each edge exists independently with probability p . Now, let us consider the probability $Pr(M_{uv} = 1 | A_k)$: the probability v will be reached given that k neighbors are present. This is exactly the product of the probability that v is indeed in the chosen k -tuple (say p_1) and the probability that then v is chosen uniformly at random (say p_2) from the k -tuple. $p_1 = \binom{d-1}{k-1} / \binom{d}{k} = \frac{k}{d}$ since the model is edge-uniform and we can fix v and choose any of the $\binom{d-1}{k-1}$ k -tuples with v in them out of the $\binom{d}{k}$ total ones. On the other hand, $p_2 = \frac{1}{k}$ by uniformity. Overall, we get $Pr(M_{uv} = 1 | A_k) = p_1 \cdot p_2 = \frac{1}{d}$. We can now apply the total probability law to calculate

$$Pr(M_{uv} = 1) = \sum_{k=1}^d Pr(M_{uv} = 1 | A_k) Pr(A_k) = \frac{1}{d} \sum_{k=1}^d \binom{d}{k} p^k (1 - p)^{d-k} = \frac{1}{d} (1 - (1 - p)^d)$$

To conclude, let us reduce *RWA* to *RWD*. Indeed, in *RWD* the equivalent transition probability is $Pr(M_{uv} = 1) = \frac{1}{d} p$, accounting both for the uniform choice and the delay p . Therefore, the *RWA* probability can be viewed as $\frac{1}{d} p'$ where $p' = (1 - (1 - p)^d)$. To achieve edge-uniformity we set $p' = (1 - (1 - p)^\delta)$ which lower bounds the delay of each edge and finally we can apply the same *RWD* analysis by substituting p by p' . Applying Theorem 2 completes the proof. \square

The value of δ used to lower-bound the transition probability may be a harsh estimate for general graphs. However, it becomes quite more accurate in the special case of a d -regular underlying graph where $\delta = \Delta = d$.

5 Cover Time with One-Step History

We now turn our attention to the case where the current state of an edge affects its next state. That is, we take into account a history of length one when computing the probability of existence for each edge independently. A Markovian model for this case was introduced in [9]; see Table 1. The left side of the table accounts for the current state of an edge, while the top for the next one. The respective table box provides us with the probability of transition from one state to the other. Intuitively, another way to refer to this model is as the *Birth-Death* model: a dead edge becomes alive with probability p , while an alive edge dies with probability q .

Let us now consider an underlying graph G evolving under the *EUE* model where each possible edge independently follows the aforementioned stochastic rule of evolution. In order to bound the *RWD* cover time, we apply a two-step analysis. First, we bound the mixing time of the Markov chain defined by Table 1 for a single edge and then for the whole graph by considering all m independent edge processes evolving together. Lastly, we estimate the cover time for a single agent after each edge has reached the stationary state of Birth-Death.

	<i>dead</i>	<i>alive</i>
<i>dead</i>	$1 - p$	p
<i>alive</i>	q	$1 - q$

Table 1. Birth-Death chain for a single edge [9]

On the other hand, for *RWA*, we make use of the "being alive" probabilities $\xi_{min} = \min\{p, 1 - q\}$ and $\xi_{max} = \max\{p, 1 - q\}$ in order to bound the cover time by following a similar argument to the one of Theorem 3 (starting again from an *RWD* analysis). In the special case of a complete underlying graph, we employ a coupon-collector-like argument to achieve an improved upper bound.

5.1 RWD for General (p, q) -Graphs via Mixing

As a first step, let us prove the following upper-bound inequality, which helps us break our analysis to follow into two separate phases.

Lemma 4. *Let $\tau(\epsilon)$ stand for the mixing time for the whole-graph chain up to some total variation distance $\epsilon > 0$, $C_{\tau(\epsilon)}$ for the expected time to cover all nodes after time step $\tau(\epsilon)$ and C for the cover time of G under *RWD*. Then, $C \leq \tau(\epsilon) + C_{\tau(\epsilon)}$ holds.*

The above upper bound discards some walk progress, however, intuitively, this may be negligible in some cases: if the mixing is rapid, then the cover time $C_{\tau(\epsilon)}$ dominates the sum, whereas, if the mixing is slow, this may mean that edges appear rarely and thence little progress can be made anyway.

Phase I: Mixing Time. Let P stand for the Birth-Death Markov chain given in Table 1. It is easy to see that P is irreducible and aperiodic and therefore its limiting distribution matches its stationary distribution and is unique. We hereby provide a coupling argument to upper-bound the mixing time of the Birth-Death

chain for a single edge. Let X_t, Y_t stand for two copies of the Birth-Death chain given in Table 1 where $X_t = 1$ if the edge is alive at time step t and $X_t = 0$ otherwise. We need only consider the initial case $X_0 \neq Y_0$. For any $t \geq 1$, we compute the meeting probability $Pr(X_t = Y_t | X_{t-1} \neq Y_{t-1}) = Pr(X_t = Y_t = 1 | X_{t-1} \neq Y_{t-1}) + Pr(X_t = Y_t = 0 | X_{t-1} \neq Y_{t-1}) = p(1 - q) + q(1 - p)$.

Definition 1. Let $p_0 = p(1 - q) + q(1 - p)$ denote the meeting probability under the above Birth-Death coupling for a single time step.

We now bound the mixing time of Birth-Death for a single edge.

Lemma 5. *The mixing time of Birth-Death for a single edge is $\mathcal{O}(p_0^{-1})$.*

Proof. Let T_{xy} denote the meeting time of X_t and Y_t , i.e. the first occurrence of a time step t such that $X_t = Y_t$. We now compute the probability the two chains meet at a specific time step $t \geq 1$:

$$\begin{aligned} Pr[T_{xy} = t] &= Pr(X_t = Y_t | X_{t-1} \neq Y_{t-1}, X_{t-2} \neq Y_{t-2}, \dots, X_0 \neq Y_0) &= \\ &= Pr(X_t = Y_t | X_{t-1} \neq Y_{t-1}) \cdot Pr(X_{t-1} \neq Y_{t-1} | X_{t-2} \neq Y_{t-2}) \cdot \dots \cdot Pr(X_1 \neq Y_1 | X_0 \neq Y_0) \cdot Pr(X_0 \neq Y_0) &= \\ &= p_0 \cdot (1 - p_0)^{t-1} \end{aligned}$$

where we make use of the total probability law and the one-step Markovian evolution. Finally, we accumulate and then bound the probability the meeting time is greater to some time-value t :

$$Pr[T_{xy} \leq t] = \sum_{i=1}^t Pr[T_{xy} = i] = \sum_{i=1}^t p_0(1 - p_0)^{i-1} = p_0 \frac{1 - (1 - p_0)^t}{p_0} = 1 - (1 - p_0)^t$$

Then, $Pr[T_{xy} > t] = (1 - p_0)^t \leq e^{-p_0 t}$, by applying the inequality $1 - x \leq e^{-x}$ for all $x \in \mathbb{R}$. By setting $t = c \cdot p_0^{-1}$ for some constant $c \geq 1$, we get $Pr[T_{xy} > c \cdot p_0^{-1}] \leq e^{-c}$ and apply Lemma 1 to bound $\tau(e^{-c}) \leq c \cdot p_0^{-1}$. \square

The above result analyzes the mixing time for a single edge of the underlying graph G . In order to be mathematically accurate, let us extend this to the Markovian process accounting for the whole graph G . Let G_t, H_t stand for two copies of the Markov chain consisting of m independent Birth-Death chains; one per edge. Initially, we define a graph $G^* = (V^*, E^*)$ such that $V^* = V$ and $E^* \subseteq E$; any graph with these properties is fine. We set $G_0 = G^*$ and $H_0 = \overline{G^*}$ which is a worst-case starting point since each pair of respective G, H edges has exactly one alive and one dead edge. To complete the description of our coupling, we enforce that when a pair of respective edges meets, i.e. when the coupling for a single edge as described in the proof of Lemma 5 becomes successful, then both edges stop applying the Birth-Death rule and remain at their current state. Similarly to before, let $T_{G,H}$ stand for the meeting time of the two above defined copies, that is, the time until all pairs of respective edges have met. Furthermore, let $T_{x,y}^e$ stand for the meeting time associated with edge $e \in E$.

Lemma 6. *The mixing time for any underlying graph G where each edge independently applies the Birth-Death rule is at most $\mathcal{O}(p_0^{-1} \log m)$.*

Phase II: Cover Time After Mixing. We can now proceed to apply Lemma 4 by computing the expected time for *RWD* to cover G after mixing is attained. As before, we use the notation $C_{\tau(\epsilon)}$ to denote the cover time after the whole-graph process has mixed to some distance $\epsilon > 0$ from its stationary state in time $\tau(\epsilon)$. The following remark is key in our motivation toward the use of stationarity.

Fact 2. *Let D be a random variable capturing the number of time steps until a possible edge becomes alive under *RWD* once the agent selects it for traversal. For any time step $t \geq \tau(\epsilon)$, the expected delay for any single edge traversal e under *RWD* is the same and equals $E[D|e_t = 1]Pr(e_t = 1) + E[D|e_t = 0]Pr(e_t = 0)$.*

That is, due to the uniformity of our model, all edges behave similarly. Furthermore, after convergence to stationarity has been achieved, when an agent picks a possible edge for traversal under *RWD*, the probability $Pr(e_t = 1)$ that the edge is alive for any time step $t \geq \tau(\epsilon)$ is actually given by the stationary distribution in a simpler formula and can be regarded independently of the edge's previous state(s).

Lemma 7. *For any constant $0 < \epsilon < 1$ and $\epsilon' = \epsilon \cdot \frac{\min\{p,q\}}{p+q}$, it holds that $C_{\tau(\epsilon')} \leq 2m(n-1) \cdot (1+2\epsilon) \frac{p^2+q}{p^2+pq}$.*

Proof. We compute the stationary distribution π for the Birth-Death chain P by solving the system $\pi P = \pi$. Thus, we get $\pi = [\frac{q}{p+q}, \frac{p}{p+q}]$.

From now on, we only consider time steps $t \geq \tau(\epsilon')$, i.e. after the chain has mixed, for some $\epsilon' = \epsilon \cdot \frac{\min\{p,q\}}{p+q} \in (0, 1)$. We have $tvd(t) = \frac{1}{2} \max_i \sum_j |p_{ij}^{(t)} - \pi_j| \leq \epsilon'$ implying that for any edge e , we get $Pr(e_t = 1) \leq (1+2\epsilon) \frac{p}{p+q}$. Similarly, $Pr(e_t = 0) \leq (1+2\epsilon) \frac{q}{p+q}$. Let us now estimate the expected delay until the *RWD*-chosen possible edge at some time step t becomes alive. If the selected possible edge exists, then the agent moves along it with no delay (i.e. we count 1 step). Otherwise, if the selected possible edge is currently dead, then the agent waits till the edge becomes alive. This will expectedly take $1/p$ time steps due to the Birth-Death chain rule. Overall, the expected delay is at most $1 \cdot (1+2\epsilon) \frac{p}{p+q} + \frac{1}{p} \cdot (1+2\epsilon) \frac{q}{p+q} = (1+2\epsilon) \frac{p^2+q}{p^2+pq}$, where we condition on the above cases.

Since for any time $t \geq \tau(\epsilon)$ and any edge e , we have the same expected delay to traverse an edge, we can extract a bound for the cover time by considering an electrical network with each resistance equal to $(1+2\epsilon) \frac{p^2+q}{p^2+pq}$. Applying Theorem 2 completes the proof. \square

The following theorem is directly proven by plugging into the inequality of Lemma 4 the bounds computed in Lemmata 6 and 7.

Theorem 4. *For any connected underlying graph G and the Birth-Death rule, the cover time of *RWD* is $\mathcal{O}(p_0^{-1} \log m + mn \cdot (p^2 + q)/(p^2 + pq))$.*

5.2 RWD and RWA for General (p, q) -Graphs via Min-Max

In the previous subsection, we employed a mixing-time argument in order to reduce the final part of the proof to the zero-step history case. Let us hereby derive another upper bound for the cover time of *RWD* (and then extend it for *RWA*) via a min-max approach. The idea here is to make use of the "being alive" probabilities to prove lower and upper bounds for the cover time parameterized by $\xi_{min} = \min\{p, 1 - q\}$ and $\xi_{max} = \max\{p, 1 - q\}$.

Let us consider an *RWD* walk on a general connected graph G evolving under *EUE* with a zero-step history rule dictating $Pr(e_t = 1) = \xi_{min}$ for any edge e and time step t . We refer to this walk as the *Upper Walk with a Delay*, shortly *UWD*. Below, we make use of *UWD* in order to bound the cover time of *RWD* and *RWA* in general (p, q) -graphs.

Lemma 8. *For any connected underlying graph G and the Birth-Death rule, the cover time of *RWD* is at most $2m(n - 1)/\xi_{min}$.*

Notice that the above upper bound improves over the one in Theorem 4 for a wide range of cases, especially if q is really small. For example, when $q = \Theta(m^{-k})$ for some $k \geq 2$ and $p = \Theta(1)$, then Lemma 8 gives $O(mn)$ whereas Theorem 4 gives $O(m^k)$ since the mixing time dominates the whole sum. On the other hand, for relatively big values of p and q , e.g. in $\Omega(1/m)$, then mixing is rapid and the upper bound in Theorem 4 proves better.

Let us now turn our attention to the *RWA* case with the subsequent theorem.

Theorem 5. *For any connected underlying graph G evolving under the Birth-Death rule, the cover time of *RWA* is at most $2m(n - 1)/(1 - (1 - \xi_{min})^\delta)$.*

Proof. Suppose the agent follows *RWA* with some stochastic rule R of the form $Pr(e_t = 1|H_t)$ which incorporates some history H_t when making a decision about an edge at time step t . Let us now proceed in fashion similar to the proof of Theorem 3. Assume the agent follows *RWA* and has reached node $u \in V$ after time step t . Then G_{t+1} becomes fixed and the agent selects uniformly at random an alive neighboring node to move to. Let M_{uv} (where v is a neighbor to u) stand for a random variable taking value 1 if the agent moves to v at time step $t + 1$ and 0 otherwise. For $k = 0, 1, 2, \dots, d(u) = d$, let $A_k(H_t)$ stand for the event that $d_{t+1} = k$ given some history H_t about all incident possible edges of u . We compute $Pr(M_{uv} = 1) = \sum_{k=1}^d Pr(M_{uv} = 1|A_k(H_t))Pr(A_k(H_t))$. Similarly to the proof of Theorem 3, $Pr(M_{uv} = 1|A_k(H_t)) = p_1 \cdot p_2 = 1/d$ where p_1 is the probability v is indeed in the chosen k -tuple (which is k/d) and p_2 is the probability it is chosen uniformly at random from the k -tuple (which is $1/k$). Thus, we get $Pr(M_{uv} = 1) = \frac{1}{d} \sum_{k=1}^d Pr(A_k(H_t)) = \frac{1}{d}(1 - Pr(A_0(H_t)))$ where A_0 is the event no edge becomes alive at this time step.

Moving forward, by definition, *UWD* depicts a zero-step history *RWD* walk. Let us denote by *UWA* its *RWA* corresponding walk. Furthermore, let P_U be equal to the probability $Pr(M_{uv} = 1)$ under the *UWA* walk. Then, we can substitute p by ξ_{min} to apply Theorem 3 and get $P_U = \frac{1}{d}(1 - (1 - \xi_{min})^d)$. In the

Birth-Death model, we know $Pr(A_0(H_1)) \leq (1 - \xi_{min})^d$ since each possible edge becomes alive with probability at least ξ_{min} . Thus, it follows $P_U \leq Pr(M_{uv} = 1)$.

To wrap up, UWA can be viewed as an RWD walk with delay probability $(1 - (1 - \xi_{min})^d)$ which lower bounds the $(1 - Pr(A_0(H_t)))$ probability associated with RWA . Inverting the inequality to account for the delays, we have $C \leq C_U$ for the cover times. Finally, Theorem 3 gives $C_U \leq 2m(n-1)/(1 - (1 - \xi_{min})^d)$. \square

5.3 RWA for Complete (p, q) -Graphs

We now proceed towards providing an upper bound for the cover time in the special case when the underlying graph G is complete, i.e. between any two nodes there exists a possible edge for our model. We utilize the special topology of G to come up with a different analytical approach and derive a better upper bound than the one given in Theorem 5. In this case, let $|V| = n + 1$ to make the calculations to follow more presentable. In other words, each node has n possible neighbors. Below, again, let $\xi_{min} = \min\{p, 1 - q\}$ and $\xi_{max} = \max\{p, 1 - q\}$. Also, let $d_t(v)$ stand for a random variable depending on the Birth-Death process and denoting the actual degree of $v \in V$ at time step t . Since all nodes have the same static degree, we simplify the notation to d_t .

Lemma 9. *For some constants $\beta \in (0, 1)$ and $\alpha \geq 3/\beta^2$, if $\xi_{min} \geq \alpha \frac{\log n}{n}$, then it holds with high probability that $d_t \in [(1 - \beta)\xi_{min}n, (1 + \beta)\xi_{max}n]$.*

Theorem 6. *For any complete underlying graph G and the Birth-Death rule with $\xi_{min} \geq \alpha \frac{\log n}{n}$, for some constant $\alpha \geq 3$, the cover time of RWA is $\mathcal{O}(n \log n)$.*

Proof. At some time step t , $i + 1$ out of the $n + 1$ nodes of G have already been visited at least once, while $n + 1 - (i + 1) = n - i$ nodes remain unvisited. The agent now lies on some arbitrary node $v \in V$. Let us consider all n possible edges with v as their one endpoint: $n - i$ of them lead to an unvisited node. That is, each possible edge leads to an unvisited node with probability $\frac{n-i}{n}$. This observation holds for all edges, therefore also for alive edges at node v at time step t . We denote the alive edges by e_1, e_2, \dots, e_{d_t} . Then, let U_1, U_2, \dots, U_{d_t} stand for random variables where $U_j = 1$ if e_j leads to an unvisited node (that is with probability $\frac{n-i}{n}$) and $U_j = 0$ otherwise. We calculate

$$Pr[\bigcup_{j=1}^{d_t} U_j = 1] = 1 - Pr[\bigcap_{j=1}^{d_t} U_j = 0] = 1 - Pr[U_j = 0]^{d_t} = 1 - \left(1 - \frac{n-i}{n}\right)^{d_t}$$

In order for an unvisited node to be visited at this step, it is required that at least one such node can be reached via an alive edge *and* that such an edge will be selected by RWA . Below, let M_i stand for a random variable where $M_i = 1$ if one of the i unvisited nodes is chosen to be visited and $M_i = 0$ otherwise. Furthermore, let R stand for a random variable where $R = 1$ if RWA selects an edge leading to an unvisited node and $R = 0$ otherwise. We compute

$$Pr[M_i = 1] = Pr[R = 1 | \exists j : U_j = 1] \cdot Pr[\bigcup_{j=1}^{d_t} U_j = 1] \geq \frac{1}{d_t} \cdot \left(1 - \left(1 - \frac{n-i}{n}\right)^{d_t}\right)$$

since if at least one unvisited node can be reached, then it will be reached with probability at least $\frac{1}{d_t}$ due to the uniform choice of *RWA*. To lower-bound the above probability, we make use of the auxiliary inequalities $1 - x \leq e^{-x}$ for any $x \in \mathbb{R}$ and $e^x \leq 1 + x + \frac{1}{2}x^2$ for any $x \leq 0$.

$$\begin{aligned} Pr[M_i = 1] &\geq \frac{1}{d_t} \cdot (1 - (1 - \frac{n-i}{n})^{d_t}) && \geq \frac{1}{d_t} \cdot (1 - e^{-\frac{n-i}{n}d_t}) && \geq \\ &\geq \frac{1}{d_t} \cdot (1 - (1 - \frac{n-i}{n}d_t + \frac{1}{2}(-\frac{n-i}{n}d_t)^2)) && \geq \frac{1}{d_t} \cdot (\frac{n-i}{n}d_t - \frac{1}{2}(\frac{n-i}{n}d_t)^2) = \\ &= \frac{n-i}{n} - \frac{1}{2}(\frac{n-i}{n})^2d_t && \geq \frac{n-i}{n} - \frac{1}{2}\frac{(n-i)^2}{n}\xi \end{aligned}$$

where in the last inequality $\xi = (1 + \beta)\xi_{max}$ follows by Lemma 9. Then, let t_i stand for the time until one of the i unvisited nodes is visited and thus $\mathbb{E}[t_i] = 1/Pr[M_i = 1]$ for any $i = 1, 2, \dots, n-1$. Overall, the cover time is given by $\sum_{i=1}^{n-1} \mathbb{E}[t_i] \leq \sum_{i=1}^{n-1} (\frac{n-i}{n} - \frac{1}{2n}(n-i)^2\xi)^{-1} \leq \int_1^{n-1} (\frac{n-x}{n} - \frac{1}{2n}(n-x)^2\xi)^{-1} dx$. We compute $\int_1^{n-1} (\frac{n-x}{n} - \frac{1}{2n}(n-x)^2\xi)^{-1} dx = n \log(|\frac{2}{x-n} + \xi|) \Big|_1^{n-1} = n(\log(|-2 + \xi|) - \log(|\frac{2}{1-n} + \xi|))$. Then, $\log(|-2 + \xi|) = \log(2 - \xi) \leq \log 2$ since $\xi \in [0, 1]$ and $\log(|\frac{2}{1-n} + \xi|) = \log(|\frac{2-\xi(n-1)}{1-n}|) = \log(|2 - \xi(n-1)|) - \log(|1-n|) = \log(\xi(n-1) - 2) - \log(n-1) \geq \log(2) - \log(n-1)$ since $2 - \xi(n-1) \leq 0$ and $\log(\xi(n-1) - 2) \geq \log(2)$ for a sufficiently large choice of α at Lemma 9. \square

Notice that the latter bound matches exactly the cover time upper bound for a simple random walk on a complete static graph. Intuitively, the condition $\xi_{min} \in \Omega(\log n/n)$ indicates the graph instance G_t is almost surely connected at each time step t given that each graph instance can be viewed as "lower-bounded" by a $G(n, \xi_{min})$ Erdős-Rényi graph. In other words, an expected degree of $\Omega(\log n)$ alive edges at each time step suffices to explore the complete graph at asymptotically the same time as in the case when all n of them are available.

6 Further Work

Our results can directly be extended for any history length considered by the stochastic rule. Of course, if we wish to take into account the last k states of a possible edge, then we need to consider 2^k possible states, thus making some tasks computationally intractable for large k . On the other hand, the min-max guarantee is easier to deal with for any value of k . Finally, it remains open whether the $\mathcal{O}(n \log n)$ bound can be extended for a wider family of underlying graphs, thus making progress over the general bound stated in Theorem 5.

Our model seems to be on the opposite end of the Markovian evolving graph model introduced in [2]. There, the evolution of possible edges directly depends on the family of graphs selected as possible instances. Thus, a new research direction we suggest is to devise another model of *partial* edge-dependency.

Acknowledgements. We would like to acknowledge an anonymous reviewer who identified an important technical error in a previous version of this extended abstract and another anonymous reviewer who suggested the use of Theorem 1 as an alternative to electrical network theory and several other useful modifications.

References

1. R. Aleliunas, R. Karp, R. Lipton, L. Lovasz, and C. Rackoff, Random walks, universal traversal sequences and the complexity of maze problems, *In 20th IEEE Annual Symposium on Foundations of Computer Science*, pp. 218-223, 1979.
2. C. Avin, M. Koucký, and Z. Lotker, How to Explore a Fast-Changing World (Cover Time of a Simple Random Walk on Evolving Graphs), *In Proceedings of the 35th international colloquium on Automata, Languages and Programming (ICALP '08)*, Springer, pp. 121-132, 2008.
3. J. Bar-Ilan and D. Zernik, Random leaders and random spanning trees, *WDAG89*, pp. 1-12, 1989.
4. H. Baumann, P. Crescenzi, and P. Fraigniaud, Parsimonious flooding in dynamic graphs, *In Proceedings of the 28th ACM symposium on Principles of distributed computing (PODC '09)*, ACM, pp. 260-269, 2009.
5. J. Bird, *Electrical Circuit Theory and Technology*, 5th Ed., Routledge, 2013.
6. G. Brightwell and P. Winkler, Maximum hitting time for random walks on graphs, *Random Structures and Algorithms*, vol. 1, pp. 263-276, 1990.
7. M. Bui, T. Bernard, D. Sohler and A. Bui, Random Walks in Distributed Computing: A Survey, *IICS 2004*, LNCS 3473, pp. 1-14, Springer, 2006.
8. A. K. Chandra, P. Raghavan, W. L. Ruzzo, and R. Smolensky, The electrical resistance of a graph captures its commute and cover times, *In Proceedings of the twenty-first annual ACM symposium on Theory of computing (STOC '89)*, ACM, pp. 574-586, 1989.
9. A. E. F. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri, Flooding time in edge-Markovian dynamic graphs, *PODC '08*, ACM, pp. 213-222, 2008.
10. A. Clementi, A. Monti, F. Pasquale, and R. Silvestri, Information Spreading in Stationary Markovian Evolving Graphs, *IEEE Trans. Parallel Distrib. Syst.*, 22, 9, pp. 1425-1432, 2011.
11. A. Clementi, A. Monti, F. Pasquale and R. Silvestri, Communication in dynamic radio networks, *PODC '08*, ACM, pp. 205-214, 2007.
12. P. G. Doyle and J. L. Snell, *Random Walks and Electric Networks*, 2006.
13. U. Feige, A tight upper bound on the cover time for random walks on graphs, *Random Structures and Algorithms*, vol. 6, pp. 51-54, 1995.
14. M. Habib, C. McDiarmid, J. Ramirez-Alfonsín, B. Reed, *Probabilistic Methods for Algorithmic Discrete Mathematics*, Springer, 1998.
15. T. Hoffmann, M. A. Porter, R. Lambiotte, Random Walks on Stochastic Temporal Networks, *Temporal Networks*, Springer-Verlag, pp. 295-313, 2013.
16. O. Michail, An Introduction to Temporal Graphs: An Algorithmic Perspective, *Internet Mathematics*, 12 (4), pp. 239-280, 2016.
17. J.R. Norris, *Markov Chains*, Cambridge University Press, 1998.
18. V. Ramiro, E. Lochin, P. Snac and T. Rakotoarivelo, Temporal random walk as a lightweight communication infrastructure for opportunistic networks, *Proceeding of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, pp. 1-6, 2014.
19. M. Starnini, A. Baronchelli, A. Barrat, and R. Pastor-Satorras, Random walks on temporal networks, *Phys. Rev. E* 85, 056115, 2012.
20. A. Wald, *Sequential Analysis*. John Wiley & Sons, New York, 1947.
21. Y. Yamauchi, T. Izumi, and S. Kamei, Mobile Agent Rendezvous on a Probabilistic Edge Evolving Ring, *In Proceedings of the Third International Conference on Networking and Computing (ICNC '12)*, pp. 103-112, 2012.

A Wald's Equation

Theorem 7 (e.g. see [20]). *Let X_1, X_2, \dots, X_N be a sequence of real-valued, independent and identically distributed random variables where N is a nonnegative integer random variable independent of the sequence (in other words, a stopping time for the sequence). If each X_i and N have finite expectations, then*

$$E[X_1 + X_2 + \dots + X_N] = E[N]E[X_1]$$

B RWA Lower Bounds

In order to derive lower bounds for *RWA*, we use as underlying graphs the following graph family commonly known as *lollipop graphs* capturing the maximum cover time for a simple random walk, e.g. see [6, 13].

Definition 2. *A lollipop graph L_n^k consists of a clique on k nodes and a path on $n - k$ nodes connected with a cut-edge, i.e. an edge whose deletion makes the graph disconnected.*

Lemma 10. *There exists an underlying graph G evolving under the zero-step history EUE such that the RWA cover time is at least $\Omega(mn/(1 - (1 - p)^\Delta))$.*

Proof. We consider the $L_n^{2n/3}$ lollipop graph which is known to attain a cover time of $\Omega(mn)$ for a simple random walk [6, 13]. Then, we can proceed similarly to the proof of Theorem 3 to calculate the *RWA* delay probability for a given vertex with degree d as $p' = (1 - (1 - p)^d)$. To achieve edge-uniformity, we upper bound p' by $(1 - (1 - p)^\Delta)$ therefore inducing an expected delay of at least $1/(1 - (1 - p)^\Delta)$. Applying Theorem 1 completes the proof. \square

Lemma 11. *There exists an underlying graph G evolving under the one-step history EUE such that the RWA cover time is at least $\Omega(mn/(1 - (1 - \xi_{max})^\Delta))$.*

Proof. Again, we consider an *RWA* walk on $L_n^{2n/3}$ to get the $\Omega(mn)$ term of the cover time. Then, we proceed in similar fashion to the proof of Theorem 5 and lower bound the probability that no edges are present at the current step by $(1 - \xi_{max})^d \leq Pr(A_0(H_1))$ given that each possible edge becomes alive with probability at most ξ_{max} . Therefore, the delay probability becomes $(1 - (1 - \xi_{max})^d) \leq (1 - (1 - \xi_{max})^\Delta)$ and we consider the *RWD*-equivalent walk to apply Theorem 1 and complete the proof. \square

C Missing Proofs from Section 4

Proof of Lemma 2. Let us denote by C_{uw} the current flowing between two neighboring nodes u and w . Then, $d(u) = \sum_{(u,w) \in E} C_{uw}$ since at each node the total inward current must match the total outward current (Kirchhoff's first

law). Moving forward, $C_{uw} = \phi_{uw}/r = \phi_{uw}/(1/p) = p \cdot \phi_{uw}$ by Ohm's law. Finally, $\phi_{uw} = \phi_{uv} - \phi_{wv}$ since the sum of electrical potential differences forming a path is equal to the total electrical potential difference of the path (Kirchhoff's second law). Overall, we can rewrite $d(u) = \sum_{(u,w) \in E} p(\phi_{u,v} - \phi_{w,v}) = d(u) \cdot p \cdot \phi_{u,v} - p \sum_{(u,w) \in E} \phi_{w,v}$. Rearranging gives

$$\phi_{u,v} = \frac{1}{p} + \frac{1}{d(u)} \sum_{(u,w) \in E} \phi_{w,v}.$$

Regarding the hitting time from u to v , we rewrite it based on the first step:

$$H_{u,v} = \frac{1}{p} + \frac{1}{d(u)} \sum_{(u,w) \in E} H_{w,v}$$

since the first addend represents the expected number of steps for the selected edge to appear due to *RWD*, and the second addend stands for the expected time for the rest of the walk.

Wrapping it up, since both formulas above hold for each $u \in V \setminus \{v\}$, therefore inducing two identical linear systems of n equations and n variables, it follows that there exists a unique solution to both of them and $H_{u,v} = \phi_{u,v}$. \square

Proof of Lemma 3. Similarly to the definition of $\phi_{u,v}$ above, one can define $\phi_{v,u}$ as the electrical potential difference between v and u when $d(w)$ amperes of current are injected into each node w and $2m$ of them are withdrawn from node u . Next, note that changing all currents' signs leads to a new network where for the electrical potential difference, namely ϕ' , it holds $\phi'_{u,v} = \phi_{v,u}$. We can now apply the Superposition Theorem (see Section 13.3 in [5]) and linearly superpose the two networks implied from $\phi_{u,v}$ and $\phi'_{u,v}$ creating a new one where $2m$ amperes are injected into u , $2m$ amperes are withdrawn from v and no current is injected or withdrawn at any other node. Let $\phi''_{u,v}$ stand for the electrical potential difference between u and v in this last network. By the superposition argument, we get $\phi''_{u,v} = \phi_{u,v} + \phi'_{u,v} = \phi_{u,v} + \phi_{v,u}$, while from Ohm's law we get $\phi''_{u,v} = 2m \cdot R_{u,v}$. The proof concludes by merging these two observations and applying Lemma 2. \square

Proof of Theorem 2. Consider a spanning tree T of G . An agent, starting from any node, can visit all nodes by performing an Eulerian tour on the edges of T (crossing each edge twice). This is a feasible way to cover G and thus the expected time for an agent to finish the above task provides an upper bound on the cover time. The expected time to cover each edge twice is given by $\sum_{(u,v) \in E_T} (H_{u,v} + H_{v,u})$ where E_T is the edge-set of T with $|E_T| = n - 1$. By Lemma 3, this is equal to $2m \sum_{(u,v) \in E_T} R_{u,v} = 2m \sum_{(u,v) \in E_T} \frac{1}{p} = 2m(n - 1)/p$. \square

D Missing Proofs from Section 5

Proof of Lemma 4. The upper bound is easy to see since *RWD* covers a subset $V_0 \subseteq V$ until mixing occurs and then, after the mixing time $\tau(\epsilon)$, we require

RWD to cover the whole node-set V ; including the already visited V_0 nodes. That is, we discard any progress made by the walk during the first $\tau(\epsilon)$ time steps and require a full cover to occur afterwards. \square

Proof of Lemma 6. To start with, we calculate the probability the meeting time is bounded by some value t :

$$\begin{aligned} Pr[T_{G,H} \leq t] &= Pr[\max_{e \in E} T_{x,y} \leq t] = Pr[(T_{x,y}^{e_1} \leq t) \wedge (T_{x,y}^{e_2} \leq t) \wedge \dots \wedge (T_{x,y}^{e_m} \leq t)] = \\ &= Pr[T_{x,y} \leq t]^m = (1 - (1 - p_0)^t)^m \geq \\ &\geq 1 - m(1 - p_0)^t \geq 1 - me^{-p_0 t} \end{aligned}$$

where we successively applied the fact that the edges are independent, Bernoulli's inequality stating $(1+x)^r \geq 1+rx$ for every r and any $x \geq -1$, and the already seen inequality $1-x \leq e^{-x}$.

Moving forward, $Pr[T_{G,H} > t] \leq me^{-p_0 t}$ and after setting $t = \alpha p_0^{-1} \log m$, for some $\alpha \geq 2$ we derive that $Pr[T_{G,H} > \alpha p_0^{-1} \log m] \leq m^{1-\alpha}$. Applying Lemma 1 gives $\tau(m^{1-\alpha}) \leq \alpha p_0^{-1} \log m$. \square

Proof of Lemma 8. Regarding *UWD*, one can design a corresponding electrical network where each edge has a resistance of $1/\xi_{min}$ capturing the expected delay till any possible edge becomes alive. Applying Theorem 2, gives an upper bound of $2m(n-1)/\xi_{min}$ for the *UWD* cover time.

Let C' stand for the *UWD* cover time and C stand for the cover time of *RWD* under the Birth-Death rule. It now suffices to show $C \leq C'$ to conclude.

In Birth-Death, the expected delay before each edge traversal is either $1/p$ (in case the possible edge is dead) or $1/(1-q)$ (in case the possible edge is alive). In both cases, the expected delay is upper-bounded by the $1/\xi_{min}$ delay of *UWD* and therefore $C \leq C'$ follows since any trajectory under *RWD* will take at most as much time as the same trajectory under *UWD*. \square

Proof of Lemma 9. We provide a lower and upper bound for the expected value of d_t and determine the necessary condition under which d_t remains near its expected value. Given d_{t-1} , we get the expression $\mathbb{E}[d_t | d_{t-1}] = p(n - d_{t-1}) + (1-q)d_{t-1}$. Then, it follows $\xi_{min}n \leq \mathbb{E}[d_t | d_{t-1}] \leq \xi_{max}n$ and, by applying the expectation again, we get $\mathbb{E}[\xi_{min}n] \leq \mathbb{E}[\mathbb{E}[d_t | d_{t-1}]] \leq \mathbb{E}[\xi_{max}n]$ which is the same as $\xi_{min}n \leq \mathbb{E}[d_t] \leq \xi_{max}n$. We now bound the probability that d_t deviates from its expected value by using the Chernoff bounds $Pr[X \geq (1+\beta)\mu] \leq e^{-\frac{\beta^2\mu}{3}}$ and $Pr[X \leq (1-\beta)\mu] \leq e^{-\frac{\beta^2\mu}{2}}$ where X is a random variable with expected value μ and $\beta \in (0, 1)$. In our case, $X = d_t$ and $\mu = \mathbb{E}[d_t]$.

$$\begin{aligned} Pr[d_t \geq (1+\beta)\xi_{max}n] &\leq Pr[d_t \geq (1+\beta)\mu] \leq e^{-\frac{\beta^2\mu}{3}} \leq e^{-\frac{\beta^2\xi_{min}n}{3}} \\ Pr[d_t \leq (1-\beta)\xi_{min}n] &\leq Pr[d_t \leq (1-\beta)\mu] \leq e^{-\frac{\beta^2\mu}{2}} \leq e^{-\frac{\beta^2\xi_{min}n}{2}} \leq e^{-\frac{\beta^2\xi_{min}n}{3}} \end{aligned}$$

In order to make the above probabilities negligible with respect to n , we constrain $\xi_{min} \geq \alpha \frac{\log n}{n}$ for some constant $\alpha \geq \frac{3}{\beta^2}$. Thus, we derive $Pr[d_t \geq (1+\beta)\xi_{max}n] \leq n^{-\frac{\alpha\beta^2}{3}} = n^{-\gamma}$ and similarly $Pr[d_t \leq (1-\beta)\xi_{min}n] \leq n^{-\frac{\alpha\beta^2}{3}} = n^{-\gamma}$ for some $\gamma = \frac{\alpha\beta^2}{3} \geq 1$ in the case $\xi_{min} \geq \alpha \frac{\log n}{n}$. \square