Random Walks on Evolving Graphs with Recurring Topologies*

Oksana Denysyuk and Luís Rodrigues

INESC-ID, Instituto Superior Técnico, Universidade de Lisboa
{oksana.denysyuk,ler}@ist.utl.pt

Abstract. In this paper we consider dynamic networks that can change over time. Often, such networks have a repetitive pattern despite constant and otherwise unpredictable changes. Based on this observation, we introduce the notion of a *ρ-recurring family* of a dynamic network, which has the property that the dynamic network frequently contains a graph in the family, where frequently means at a rate $0<\rho\leq 1$. Using this concept, we reduce the analysis of max-degree random walks on dynamic networks to the case for static networks. Given a dynamic network with a ρ-recurring family $\mathcal{F}$, we prove an upper bound of $O\left(\rho^{-1}\hat{t}_{hit}(\mathcal{F})\log n\right)$ on the hitting and cover times, and an upper bound of $O\left(\rho^{-1}(1-\hat{\lambda}(\mathcal{F}))^{-1}\log n\right)$ on the mixing time of random walks, where n is the number of nodes, $\hat{t}_{hit}(\mathcal{F})$ is upper bound on the hitting time of graphs in $\mathcal{F}$, and $\hat{\lambda}(\mathcal{F})$ is upper bound on the second largest eigenvalue of the transition matrices of graphs in $\mathcal{F}$. These results have two implications. First, they yield a general bound of $O\left(\rho^{-1}n^3\log n\right)$ on the hitting time and cover time of a dynamic network (ρ is the rate at which the network is connected); this result improves on the previous bound of $O\left(\rho^{-1}n^5\log^2 n\right)$ [3]. Second, the results imply that dynamic networks with recurring families preserve the properties of random walks in their static counterparts. This result allows importing the extensive catalogue of results for static graphs (cliques, expanders, regular graphs, etc.) into the dynamic setting.

1 Introduction

In this paper we consider dynamic networks that can change over time. These networks abstract many important systems, such as mobile networks, where nodes may change neighbors as they move; and peer-to-peer networks, where nodes may connect or disconnect due to churn. A dynamic network is modeled as an *evolving graph*, which is a sequence of graphs $\mathcal{G}=\{G_i\}$ over n nodes, each graph representing a snapshot of the system at a given instant.

Much recent work has considered dynamic networks, by proposing and analyzing new algorithms [11,17,24] and by deriving new complexity bounds [18,26]. Because of their generality, dynamic networks are not only of theoretical importance, but also

* This work was partially supported by Fundação para a Ciência e Tecnologia (FCT) via the project PEPITA (PTDC/EEI-SCR/2776/2012) and via the INESC-ID multi-annual funding through the PIDDAC Program fund grant, under project PEst-OE/EEI/LA0021/2013.

F. Kuhn (Ed.): DISC 2014, LNCS 8784, pp. 333–345, 2014.

of practical relevance. At the same time, this generality makes it hard to derive strong results, which has motivated new properties that constrain the behavior of dynamic networks. Unfortunately, existing properties are either too restrictive or hard to evaluate in practice (see Section 2).

We propose a new and intuitive approach to study dynamic networks, by looking at families of graphs that recur frequently in the dynamic network. Informally, a *ρ-recurring family* of an evolving graph $\mathcal{G}$ is a family $\mathcal{F}$ of (static) graphs such that, with frequency ρ, some graph in $\mathcal{F}$ appears in the sequence $\mathcal{G}$. For example, if $\rho = 1/2$, then half of the graphs in the sequence $\mathcal{G}$ belong to $\mathcal{F}$; note that it is possible that no individual graph in $\mathcal{F}$ recurs with frequency $1/2$. Also note that the other half of the graphs in the sequence $\mathcal{G}$ may be completely arbitrary and even contain a different recurring family.

Every evolving graph has a trivial 1-recurring family, the family of all graphs. But real networks may have other more interesting recurring families because, by their own nature, these networks tend to preserve certain topological characteristics. For example, nodes in a peer-to-peer network may keep a constant number of neighbors [28]; such network has graphs with constant degree as a recurring family. Also, numerous dynamic networks build and maintain global structures, such as overlay rings [32] or routing trees [31]; in these examples, the recurring families are graphs with the required ring or tree structures. Table 1 has more examples of recurring families in various contexts.

In this paper, we focus on the study of random walks. Due to their simplicity, locality, low overhead, and correct operation under topology changes, random walks have been recently used in different types of dynamic networks for a number of applications: querying, searching, routing, topology maintenance, etc. [4,13,15,30,27].

We show that recurring families can be used to reduce the analysis of random walks in dynamic networks—which are complex—to the simpler case of static networks—which are well understood. Specifically, we give upper bounds on the behavior of random walks in dynamic networks based on similar bounds in static networks, given a recurring family.

In this study, we make two assumptions. First, we assume an *oblivious* adversary controlling the dynamic network; that is, the evolution of the graph is independent of the position of the random walk. Without this assumption, the adversary can degenerate the random walk, causing it to oscillate forever between two nodes (see Section 3). Second, we assume a *max-degree* random walk: at each node, the probability of transitioning to each neighbor is $1/d_{max}$, where d_{max} is the maximum degree of the graph. For nodes with degree $d < d_{max}$, there is a probability of $1-d/d_{max}$ of remaining at the node. Max-degree random walks are a well-behaved variant of *simple* random walks—which choose each neighbor uniformly at random—but simple random walks can have an erratic behavior in dynamic networks [3] (e.g., their cover time can be exponential).

Our main result states that, if $\mathcal{F}$ is a ρ-recurring family of an evolving graph $\mathcal{G}$, then a max-degree random walk on $\mathcal{G}$ has hitting time and cover time of $O\left(\rho^{-1}\hat{t}_{hit}(\mathcal{F})\log n\right)$ and mixing time of $O\left(\rho^{-1}(1-\hat{\lambda}(\mathcal{F}))^{-1}\log n\right)$, where $\hat{t}_{hit}(\mathcal{F})$ is an upper bound on the hitting time of graphs in $\mathcal{F}$, and $\hat{\lambda}(\mathcal{F})$ is an upper bound on the second largest eigenvalue of the transition matrices of graphs in $\mathcal{F}$. To prove these results, we consider the *homogeneous* Markov Chains of the graphs in the recurring family, and relate these chains to the *time non-homogeneous* Markov Chain of the random walk on the dynamic network. Specifically, using arguments from matrix analysis, we analyze the transition matrices in the recurring family and obtain bounds on the algebraic properties (eigenvalues, etc.) for each matrix considered as a homogeneous Markov chain. We then relate the product of the matrices in the recurring family to the product of all matrices in the non-homogeneous chain, and map the bounds to the original dynamic network.

The obtained bounds are nearly tight and have two important implications. First, they reduce the known gap between the complexity of random walks in dynamic and static networks. In particular, in static networks, the cover time has a general upper bound of $O(n^3)$ [2][1], but in dynamic networks, the previously known bound was much higher: $O(\rho^{-1}n^5 \log^2 n)$ [3], where ρ is the frequency with which the network is connected. We reduce this gap to just a $\log n$ factor: by using the trivial ρ-recurring family of all connected graphs, we obtain a general bound of $O(\rho^{-1}n^3 \log n)$.

Second, these results imply that dynamic networks with ρ-recurring families preserve the random-walk properties of their static counterparts. It is thus possible to import the extensive catalogue of results for random walks on static graphs to the dynamic setting. For instance, it is known that random walks are especially efficient on certain families of graphs, such as expanders. For expanders, hitting time is $O(n)$, cover time is $O(n \log n)$ and mixing time is $O(\log n)$. Thus, in evolving graphs where expanders appear frequently, we can derive stronger bounds. We say that evolving graph $\mathcal{G}$ has a ρ-recurring expander if fraction ρ of the graphs in $\mathcal{G}$ are expanders. Then, it follows that, for $\mathcal{G}$, hitting time is $O\left(\rho^{-1}n \log n\right)$, cover time is $O\left(\rho^{-1}n \log n\right)$, and mixing time is $O\left(\rho^{-1} \log n\right)$, respectively.

In summary, this paper makes the following contributions:

- we introduce a novel property of evolving graphs, which we call a *ρ-recurring family*;
- using recurring families, we derive new bounds for a random walk on an evolving graph;
- we show that random walks on evolving graphs with recurring families preserve the properties of their static counterparts;
- we argue that our bounds are nearly tight and improve upon previously known bounds.

Paper Organization. The remainder of this paper is organized as follows. Section 2 discusses related work. Section 3 states the model and Section 4 defines the problem addressed in this paper. In Section 5 we state the main results and in Section 6 we discuss their implications. In Section 7 we sketch the proofs. Finally, Section 8 presents the conclusions and outlines directions for future work.

[1] More precisely, [2] proves an $O(nm)$ cover time bound of a *simple* random walk on a graph with n nodes and m edges. It can be shown that the *max-degree* random walk on a graph is equivalent to a simple random walk on the graph, augmented with sufficiently many self-loops, such that the degree of each node is d_{max}. In such augmented graph $m \leq n^2$; thus, [2] implies an $O(n^3)$ cover time bound for the max-degree random walk (we defer the complete proof to the full version of the paper).

2 Related Work

Dynamic Networks. There is a growing interest in the study of graphs that evolve over time, representing a variety of dynamic networks. Different models of dynamic networks have been proposed, each capturing specific features of some concrete scenario. Random changes of links are considered in [8]. In [9], the authors propose a model of a dynamic network where the existence of an edge in a round stochastically depends on its existence in the previous round. Adversarial networks have also received attention [8,24,16,26], representing a worst-case scenario where link changes are controlled by an adversary that tries to slow down communication. This last model covers the widest range of different network behaviors; therefore, we adopt it for our study.

Different properties have been proposed to analyze algorithms in such networks. For an extensive discussion, we refer the interested reader to [7] and [25]. In [24], the authors propose an elegant concept of T-interval connectivity and use it to study token dissemination. Evolving graph $\mathcal{G}$ is T-interval connected if, for every T consecutive rounds, there exists a connected spanning subgraph of $\mathcal{G}$ that does not change. T-interval connectivity is a strong property and may be too restrictive for some real world scenarios. Moreover, we focus on random walks, and this property is not well-suited for this problem, because the behavior of a random walk is not governed by a stable spanning subgraph. For example, a lollipop graph has cover time of $\Theta(n^3)$, despite having a line as a spanning subgraph, which has a cover time of $\Theta(n^2)$. Here, the existence of the line does not help the cover time of the lollipop. By contrast, we show that recurring families closely relate to the behavior of random walks in the evolving graph. In fact, our results imply that T-interval connectivity is not necessary for the random walk to make fast progress, as long as the evolving graph forms good topologies often enough.

Another popular property of the dynamic network is its *dynamic diameter* [11,24,26], which is the worst-case number of rounds required to route a piece of information from any given node to all other nodes. Intuitively, the concept of dynamic diameter is useful in the study of information spreading. Unfortunately, however, the dynamic diameter is hard to estimate in a real network, which is a practical drawback.

A number of other papers study information spreading in dynamic networks, e.g., [17,18]. Our paper differs from these works because it proposes and uses a different property to study dynamic networks (recurring families) and it focuses on a different problem, random walks.

Random Walks. Much work has considered random walks on static graphs, with the proposal of bounds for many families of graphs. For a comprehensive survey please refer to [29]. More recently, there has been growing interest in random walks in dynamic settings. In [10] the authors study random walks on a graph that evolves by adding new node with random or preferential connections to existing nodes. Since the graph grows, one never visits all nodes, and so the usual notions of hitting, cover, and mixing times (which we consider) do not apply.

In [13], the authors consider connected randomly evolving graphs where, in each round, the set of edges for a node is chosen uniformly at random. The authors show that the random walk on such evolving graph is essentially a random walk on a clique:

each transition can be seen as a random choice of a list of neighbors and then a random choice of an item in the list. Thus, the cover time of such graph is $O(n \log n)$. This result does not apply to adversarial evolving graphs (which we consider). For example, if an adversary chooses a sparse random graph and never changes it, then the cover time of such graph is $O(n^2 \log n)$ [21].

To our knowledge, [3] is the first paper to address randoms walks on adversarial evolving graphs. The authors show that the behavior of a *simple* random walk on evolving graphs can differ significantly from the static case. In particular, the cover time of a simple random walk can be exponential as demonstrated by an example of a dynamic star over nodes $0, \ldots, n-1$, where in round t the center of the star is node $t \mod (n-1)$, and the remaining nodes are leaves. In addition, all nodes have self-loops, allowing the random walk to remain at one node for several rounds. Notice that node $n-1$ is never at the center of the star. The only way the walk can reach node $n-1$ is by staying at some leaf for $n-2$ rounds until this leaf becomes a center of the star (if the walk moves to the center too soon, the process starts over, because the center will itself be a leaf again in the next round). The probability that the random walk stays at a leaf for $n-2$ consecutive rounds is $\frac{1}{2^{(n-2)}}$; hence, the cover time is $\Omega(2^n)$. Additionally, [3] gives a $O\left(\rho^{-1} d_{max}^2 n^3 \log^2 n\right)$ [3] bound on the cover time of a max-degree random walk on $\mathcal{G}$, where ρ is the fraction of connected graphs $\mathcal{G}$ and d_{max} is the maximum degree of any graph in $\mathcal{G}$. The result of [3] implies the general bound of $O\left(\rho^{-1} n^5 \log^2 n\right)$ for any evolving graph. We improve this result to the nearly tight $O(\rho^{-1} n^3 \log n)$. Our results also give stronger bounds on evolving graphs with structure, as we later explain.

Random walks in dynamic networks are also considered in [11], which studies simple random walks on connected *regular* evolving graphs. Note that the results of [11] also apply to *max-degree* random walks on any connected evolving graph. The authors formally discuss the notion of mixing time in a dynamic network and give a $O((1-\lambda)^{-1} \log n)$ bound where λ is an upper bound on the second largest (in absolute value) eigenvalue of *all* transition matrices of graphs in $\mathcal{G}$. This result is weaker than ours for two reasons. First, it considers only mixing time and not hitting nor cover times. Second, the bound in [11] is governed by the *worst* graph appearing in the entire evolving graph, whereas our bounds are governed by the good graphs that appear frequently. The authors of [11] also propose an algorithm for distributed computation of a random walk that runs in $O(\sqrt{t_{mix} \tau})$ rounds where t_{mix} is the mixing time and τ is the dynamic diameter of the evolving graph. The analysis of the running time of this algorithm can benefit from our new results on the mixing time in dynamic networks with structure (which we explain later).

3 Evolving Graph Model

We consider an undirected network with a fixed set V of n nodes, where edges between nodes may change over time. Execution proceeds in synchronous rounds, where in each round an adversary chooses the set of links connecting pairs of nodes. An execution generates an evolving graph, which is a sequence $\mathcal{G} = G_1, G_2, \ldots$ of graphs over nodes V, where G_t is a snapshot of the evolving graph in round t. We omit reference to V when it is clear from the context.

Definition 1 (**ρ-Recurring Family**)**.** *Given an evolving graph $\mathcal{G} = G_1, G_2, \ldots$ and a number $0<\rho\leq 1$, a ρ-recurring family $\mathcal{F}$ of $\mathcal{G}$ is a family of graphs such that, for every $M\geq 1$, at least $\lfloor \rho M \rfloor$ elements in $G_1, \ldots, G_M$ are in $\mathcal{F}$.*

Intuitively, this definition requires that, with frequency at least ρ, the graphs in a ρ-recurring family appear in $\mathcal{G}$. The definition can be weakened to require the frequency ρ to hold only for sufficiently large M; the results in this paper can be easily modified to work with this weaker definition.

Note that every evolving graph has a 1-recurring family, the family of all (including disconnected) graphs on V. Also, if $\mathcal{F}$ is a ρ-recurring family of $\mathcal{G}$ then we can add any graphs to $\mathcal{F}$ and still have a ρ-recurring family of $\mathcal{G}$. Generally, we are interested in small recurring families, because our bounds are based on the worst graph in the family.

The paper focuses on random walks on $\mathcal{G}$. We assume an *oblivious adversary* that determines the evolving graph without knowledge of the random walk. Without this assumption, an adaptive adversary could degenerate the random walk using a simple strategy: in odd rounds, the adversary provides the current position v_i of the random walk with a single edge to some fixed node v_j. In even rounds the adversary provides v_j with a single edge to v_i. Under this strategy, the random walk oscillates between v_i and v_j forever. Such a random walk never converges (it has infinite mixing time, hitting time, etc).

4 Random Walk Definition

We assume that, in round one, a random walk starts at some node of a given evolving graph and, in each round, it moves from a node to one of its neighbors with certain probability. We consider a max-degree random walk: at every node, we move to a given neighbor with fixed probability $1/d_{max}$, where d_{max} is the maximum degree of the graph or an upper bound on the maximum degree (if d_{max} is unknown, we can let $d_{max} = n$); with probability $1-d/d_{max}$ we do not move, where d is the node degree. The max-degree random walk can be seen as a simple random walk on a graph augmented with self-loops so that every node has the same degree d_{max}. We further make the standard assumption of an aperiodic random walk; this can be ensured, for example, by avoiding bipartite graphs or by assuming that all nodes have self-loops.

Max-degree aperiodic random walks are attractive for two reasons. First, in steady state, it is easy to show that every node has equal probability; this property is useful for applications that require fairness, such as fair token circulation [20]. Second, the random walk avoids the poor exponential behavior that simple random walks may exhibit [3].

We are interested in the following asymptotic properties of the random walk, which are natural extensions of the properties of random walks on static graphs. Given evolving graph $\mathcal{G}$:

- Hitting time $t_{hit}(\mathcal{G})$ is the maximal expected number of rounds before the random walk visits some node of $\mathcal{G}$;
- Cover time $t_{cov}(\mathcal{G})$ is the expected number of rounds before the random walk visits every node of $\mathcal{G}$ at least once;

- Mixing time $t_{mix}(\mathcal{G})$ is the expected number of rounds before reaching the steady state distribution of the random walk on $\mathcal{G}$ (if such distribution exists).

5 Statement of the Main Results

We now state the upper bounds on the hitting time, cover time, and mixing time of random walks on evolving graphs, based on the properties of a recurring family of that graph. The graph will generally have many recurring families; the bounds apply to each of them.

For a family $\mathcal{F}$ of graphs, let $\hat{t}_{hit}(\mathcal{F})$ be an upper bound on the hitting time of the graphs in $\mathcal{F}$ and $\hat{\lambda}(\mathcal{F})$ be an upper bound on the second largest eigenvalue of the transition matrices of the graphs in $\mathcal{F}$. Our main results are the following:

Theorem 2. *Let $\mathcal{G}$ be an evolving graph over n nodes and $\mathcal{F}$ be a ρ-recurring family of $\mathcal{G}$. The hitting time and cover time of a max-degree random walk on $\mathcal{G}$ are bounded by*

$$t_{hit}(\mathcal{G}) \le t_{cov}(\mathcal{G}) = O\left(\rho^{-1}\hat{t}_{hit}(\mathcal{F})\log n\right).$$

Theorem 3. *Let $\mathcal{G}$ be an evolving graph over n nodes and $\mathcal{F}$ be a ρ-recurring family of $\mathcal{G}$. The mixing time of a max-degree random walk on $\mathcal{G}$ is bounded by*

$$t_{mix}(\mathcal{G}) = O\left(\rho^{-1}(1-\hat{\lambda}(\mathcal{F}))^{-1}\log n\right).$$

The bounds on cover time and mixing time are tight in the sense that there is an evolving graph that matches the bounds; meanwhile, the bounds on hitting time are within $\log n$ factor from the optimal. Specifically, take an evolving graph $\mathcal{G}$ that is a static expander, that is $\mathcal{G} = G, G, \ldots$ where G is an expander. Then $\mathcal{G}$'s hitting time is $\Theta(n)$, its cover time is $\Theta(n\log n)$, and its mixing time is $O\left((1-\hat{\lambda}(\{G\}))^{-1}\log n\right)$. We see that the cover and mixing times match Theorems 2 and 3, while the hitting time is within a $\log n$ factor.

Thus, the behavior of the random walk on evolving graphs can be studied via its recurring families. Doing so allows importing the results on static graphs to the dynamic setting. We next give several applications of this idea.

6 Implications

General Bound. In some cases, little is known about the topology of the dynamic network $\mathcal{G}$; its changes over time can be arbitrary and unpredictable. However, if we only know that there exists some $\rho > 0$ such that, for every $M \ge 1$, $\mathcal{G}$ is connected in at least ρM rounds, we can apply our results to obtain non-trivial bounds[2]. For such $\mathcal{G}$, we can take the ρ-recurring family of all connected graphs and obtain the following result:

[2] If $\mathcal{G}$ does not remain connected for any fraction ρ of rounds, its mixing, hitting, and cover times can be infinite.

Theorem 4. *Let $\mathcal{G}$ be an evolving graph such that, for every $M \geq 1$, $\mathcal{G}$ is connected in at least ρM rounds. Then, the hitting time and cover time of a max-degree random walk on $\mathcal{G}$ are bounded by*

$$t_{cov}(\mathcal{G}) = O\left(\rho^{-1} n^3 \log n\right).$$

This theorem improves on the cover time bound of $O\left(\rho^{-1} n^5 \log^2 n\right)$ in [3]. The proof is by direct application of Theorem 2.

Relevant ρ-Recurring Families. We can model many dynamic networks by evolving graphs with structure. For instance, many mobile ad hoc networks have cliques as recurring families. Cliques have excellent mixing and hitting times of only $\Theta(1)$ and $\Theta(n)$. However, unfavorable topologies can emerge frequently, such as lollipop and barbell graphs, which have poor mixing and hitting times of $\Theta(n^3)$. If $\mathcal{G}'$ forms a clique at least a fraction $\rho > 0$ of the time—we say that $\mathcal{G}'$ has a ρ-recurring clique—then even if $\mathcal{G}'$ has frequent lollipops and barbell graphs, our results show that its behavior is governed by the good topologies. Here, the ρ-recurring clique provides an intuitive example: when the network forms a clique, the random walk can jump to any node, irrespective of the remaining rounds. It is thus easy to see that the random walk quickly covers the network. Theorems 2 and 3 yield a bound of $O(\rho^{-1} n \log n)$ on the cover and hitting times of $\mathcal{G}'$ and a bound of $O(\rho^{-1} \log n)$ on the mixing time. By contrast, the result in [11], which is governed by the worst graphs in $\mathcal{G}'$, yields a much looser bound of $O(n^3 \log n)$ on the mixing time (and no results on hitting and cover time). With a little more work, we can further improve the bounds of Theorems 2 and 3 using the same proof techniques, to obtain tight bounds for all metrics, as stated in the following theorem:

Theorem 5. *If evolving graph $\mathcal{G}$ has a ρ-recurring clique, then the mixing time of a max-degree random walk on $\mathcal{G}$ is $O(\rho^{-1})$, the hitting time is $O(\rho^{-1} n)$, and the cover time is $O(\rho^{-1} n \log n)$.*

Expander graphs are another important recurring family in many dynamic networks. For instance, some unstructured peer-to-peer overlays seek to maintain good expansion properties [28]. Our results imply that an evolving graph with ρ-recurring bounded-degree expander has $O(\rho^{-1} n \log n)$ cover time and $O(\rho^{-1} \log n)$ mixing time. Thus, regardless of arbitrary topologies generated during transition periods, a random walk on such evolving graph preserves the properties of its good static topologies.

In Table 1 we illustrate more implications of our results. All the graphs in the table have $(d_{max} - d_{min}) < c$, for some constant c. This property minimizes the difference between simple and max-degree random walks, allowing us to use the bounds for simple random walks in static graphs (the intuition is that adding bounded holding probabilities does not change the asymptotic behavior of the random walk).

Unions as ρ-Recurring Families. We further note that a recurring family $\mathcal{F}$ can be defined as a union of multiple well-known families of graphs. As an example, consider a network arranged in a ring in which one fixed link is intermittent (e.g., it may be an unstable link in a radio network). We can model such network as an evolving graph. When the link is present, the graph is a ring; when the link is absent, the graph is a chain. We may have no information about what fraction ρ of the time the graph is a ring or chain, making it impossible to apply our results to either ring or chain. However,

Table 1. New bounds obtained from Theorems 2, 3, and 5, for a max-degree random walk on evolving graphs, with the corresponding recurring families

ρ-recurring family	**occurrence in dynamic networks**	**cover time**	**mixing time**	**static ref**
cliques	mobile ad-hoc networks mesh networks	$O(\rho^{-1}n\log n)$	$O(\rho^{-1})$	[14]
regular and nearly regular	structured overlays(rings) unstructured overlays	$O(\rho^{-1}n^2\log n)$	$O(\rho^{-1}n^2\log n)$	[22]
2-dim grids	sensor networks	$O(\rho^{-1}n\log^2 n)$	$O(\rho^{-1}n\log n)$	[12]
bound.degree trees	routing overlays	$O(\rho^{-1}n^2\log n)$	$O(\rho^{-1}n^2\log n)$	[5]
d-regular expanders	unstructured overlays	$O(\rho^{-1}n\log n)$	$O(\rho^{-1}\log n)$	[6]

we can take the ρ-recurring family to contain *both* the ring and chain, and in this case $\rho = 1$. Then, Theorems 2 and 3 give strong bounds of $O\left(n^2\log n\right)$ on the hitting and cover times, and of $O\left(n^2\log n\right)$ on the mixing time. In this example, the intermittent link was fixed, but the example carries through identically even if the intermittent link varies over time.

7 Proofs

In this section we sketch the proofs of the main results. Due to space constraints, the complete proofs have been deferred to the full version of the paper.

Preliminaries and Main Technique. Let $\mathcal{G} = G_1, G_2, \ldots$ be an evolving graph; in each round t, A_{G_t} denotes the transition probability matrix of the random walk on G_t. If $\mathbf{p}_t = (p_1, p_2, \ldots, p_n)$ is the probability distribution on the nodes in round t, then, the probability distribution on the nodes in round $t+1$ is calculated by $\mathbf{p}_{t+1} = \mathbf{p}_t A_{G_t}$. Hence, the random walk on $\mathcal{G}$ can be modeled as a stochastic process that holds the Markov property, i.e., each transition of the random walk depends only on its current position and the transition probabilities in a given round. This kind of stochastic processes is known in the literature as *time non-homogeneous* Markov chains [23].

We model the random walk on evolving graphs as a time non-homogeneous Markov chain and work with products of stochastic matrices. For conciseness, we denote by $\mathcal{G}_{\rho,\mathcal{F}}$ an evolving graph $\mathcal{G}$ with a ρ-recurring family $\mathcal{F}$. In the analysis of $\mathcal{G}_{\rho,\mathcal{F}}$, we rely on the common algebraic properties of the stochastic matrices of graphs in $\mathcal{F}$. We then use the fact that, for any $M > 0$, in a set of M matrices, there are at least ρM matrices with those properties, to obtain the overall bounds.

For the mixing time we use the well known relation to the second largest eigenvalue. For the hitting and cover times, we bound the spectral radii of principal submatrices (i.e., matrices resulting from deleting an i-the row and an i-th column). The bound on the spectral radii of the principal submatrices is related to the hitting time of the *homogeneous* Markov chain.

Below we summarize the notation used in our proofs.

Notation

- $||\mathbf{v}||_p = \sqrt[p]{\sum_{i=1}^{n} |v_i|^p}$ for some vector $\mathbf{v} = (v_1, \ldots, v_n)$.
- $\lambda_1(A) \geq \lambda_2(A) \geq \ldots \geq \lambda_n(A)$ are eigenvalues of square matrix A.
- $\lambda(A) = \max\{|\lambda_2(A)|, |\lambda_n(A)|\}$.
- $\hat{\lambda}(\mathcal{F}) = \max_{G \in \mathcal{F}} \lambda(A_G)$, where A_G is the transition matrix of a max-degree random walk on G.
- $\delta(A) = \max_i |\lambda_i(A)|$ is spectral radius of matrix A .
- $\left\|\mathbf{v} - \mathbf{w}\right\|_{\text{TV}} = \max_{X \in \Omega} |\mathbf{v}(X) - \mathbf{w}(X)|$ denotes the total variation of two probability measures $\mathbf{v}$ and $\mathbf{w}$ over Ω.
- $t_{hit}(G)$ is the hitting time of graph G.
- $\hat{t}_{hit}(\mathcal{F})$ denotes upper bound on the hitting times of all graphs in family $\mathcal{F}$.
- $\mathcal{G}_{\rho,\mathcal{F}}$ denotes an evolving graph $\mathcal{G}$ with a ρ-recurring family $\mathcal{F}$.

Recall that we make the standard aperiodicity assumption. Moreover, as a result of using the max-degree strategy, the transition probability matrices A_{G_t}, in each round t, are symmetric and doubly stochastic (i.e., every row sums to one and every column sums to one). Therefore, each A_{G_t} has eigenvector $\frac{\mathbf{i}}{n} = \left(\frac{1}{n}, \frac{1}{n}, \ldots\right)$ with a corresponding eigenvalue $\lambda_1\left(A_{G_t}\right) = 1$.

Also, since the matrices A_{G_t} are real symmetric with all entries $0 \leq a_{i,j} \leq 1$, for any $i, j \leq n$, all eigenvalues of A_{G_t} are real (see e.g. [19]). In particular, $-1 < \lambda_n\left(A_{G_t}\right) \leq \ldots \leq \lambda_1\left(A_{G_t}\right) \leq 1$ (the strict inequality follows from aperiodicity). Also, when G_t is connected, $\lambda(A_{G_t}) = \max\{|\lambda_2(A_{G_t})|, |\lambda_n(A_{G_t})|\} < 1$. Hence, if the evolving graph is connected in sufficiently many rounds, the resulting time non-homogeneous Markov chain is ergodic and has unique stationary distribution $\pi = \frac{\mathbf{i}}{n}$.

Mixing Time. We start by bounding the mixing time. The *convergence rate* of the Markov chain is the rate at which the chain approaches stationary distribution. For *homogeneous* chains, the spectral gap of the transition matrix, i.e. the difference between the largest and the second largest eigenvalues in absolute value, defines the convergence rate to the stationary distribution [1].

The following Lemma 6 bounds the convergence rate of a max-degree random walk on an evolving graph in any given round t.

Lemma 6. *If $\mathbf{p}_t = (p_1, \ldots, p_n)$ is a probability distribution on nodes of G_t, then*

$$\left\|\mathbf{p}_{t+1} - \tfrac{\mathbf{i}}{n}\right\|_2^2 \leq \lambda^2(A_{G_t})\left\|\mathbf{p}_t - \tfrac{\mathbf{i}}{n}\right\|_2^2.$$

The following Lemma 7 establishes the monotonicity property of distribution p_t: whenever G_t belongs to $\mathcal{F}$, the random walk on $\mathcal{G}_{\rho,\mathcal{F}}$ gets closer to the stationary distribution at a known rate, while never moving away from the stationary distribution in the remaining rounds.

Lemma 7. *Consider a max-degree random walk on $\mathcal{G}_{\rho,\mathcal{F}}$. Let $\hat{\lambda}(\mathcal{F})$ denote an upper bound on the second largest (in absolute value) eigenvalues of the stochastic matrices of all graphs in $\mathcal{F}$ (i.e., $\forall_{G \in \mathcal{F}} \lambda(A_G) \leq \hat{\lambda}(\mathcal{F})$). It holds that*

$$\left\|\mathbf{p}_{t+1} - \tfrac{\mathbf{i}}{n}\right\|_2^2 \leq \left(\hat{\lambda}(\mathcal{F})\right)^{2\rho t}\left\|\mathbf{p}_1 - \tfrac{\mathbf{i}}{n}\right\|_2^2.$$

We are now ready to prove the bound on the mixing time in Theorem 3. We use the standard definition of the mixing time via the total variation distance to the steady state distribution (e.g., [1]):

$$t_{mix}=min\left\{t:\left\|\mathbf{p}_t-\pi\right\|_{\mathrm{TV}}<\frac{1}{4}\right\}.$$

This definition gives the expected value of minimal random time at which the random walk has the stationary distribution.

We use the standard method of bounding the total variation distance via the 2-norm distance to the steady state distribution.

By taking $t=O\left(\rho^{-1}(1-\hat{\lambda}(\mathcal{F}))^{-1}\log n\right)$ and applying Lemma 7, we show that after t rounds, the total variation distance is less than $\frac{1}{4}$.

Hitting Time and Cover Time. We take an arbitrary node j and remove the corresponding row and column from matrix A_{G_t}. Let A' be the resulting matrix. Lemma 8 connects the largest eigenvalue of A' to the largest eigenvalue of the fundamental matrix $(\mathbb{I}-A')^{-1}$.

Lemma 8. *Let A_{G_t} be the transition probability matrix of a max-degree random walk on graph G_t. Let A' be an $(n-1)\times(n-1)$ matrix resulting from deleting the j-th row and j-th column from A_{G_t}, for some $1\le j\le n$. And let $\mathbb{I}$ be an $(n-1)\times(n-1)$ identity matrix. Then,*

$$\lambda_1\left(A'\right)=1-\frac{1}{\lambda_1\left((\mathbb{I}-A')^{-1}\right)}.$$

The following lemma uses Lemma 8 to connect the spectral radius of A' to the hitting time of the deleted node j.

Lemma 9. *Let A_{G_t} be the transition probability matrix of a max-degree random walk on graph G_t. Let A' be an $(n-1)\times(n-1)$ matrix resulting from deleting the j-th row and j-th column from A_{G_t}, for some $1\le j\le n$. Then,*

$$\delta(A')\le\begin{cases}\left(1-\frac{1}{t_{hit}(G_t)}\right) & \text{if } G_t \text{ is connected}\\ 1 & \text{otherwise.}\end{cases}$$

We now sketch the proof of Theorem 2. We take an arbitrary node i and remove the corresponding rows and columns from the matrices $A_{G_1}, A_{G_2}, \ldots$. We use the bounds on the spectral radii of these submatrices, given in Lemma 9, to obtain the bound on the spectral radius of the product of those submatrices. Then, we relate the spectral radius of the product to the hitting time of the deleted node i. The cover time is obtained by the union bound over all n nodes.

8 Conclusions

We have introduced the notion of a ρ-recurring family of evolving graphs, which has the property that the evolving graph frequently contains a graph in the family. We believe that recurring families is a natural and powerful concept to understand many real dynamic networks.

We have studied max-degree random walks and, using the concept of recurring families, derived bounds on hitting, cover, and mixing times of an evolving graph with a ρ-recurring family $\mathcal{F}$. These results imply that dynamic networks with recurring families preserve the properties of random walks in their static counterparts. This allows importing the extensive catalogue of results for static graphs into the dynamic setting.

We believe that ρ-recurring families may be useful to study other problems in dynamic networks, such as rumour spreading, information dissemination, and token circulation. We leave this as future work.

Acknowledgements. The authors are grateful to Fabian Kuhn and the anonymous referees for their valuable feedback on the previous versions of the paper.

References

1. Aldous, D., Fill, J.A.: Reversible Markov Chains and Random Walks on Graphs. Unpublished (1995)
2. Aleliunas, R., Karp, R.M., Lipton, R.J., Lovász, L., Rackoff, C.: Random walks, universal traversal sequences, and the complexity of maze problems. In: 20th Annual Symposium on Foundations of Computer Science, FOCS 1979, pp. 218–223 (1979)
3. Avin, C., Koucký, M., Lotker, Z.: How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part I. LNCS, vol. 5125, pp. 121–132. Springer, Heidelberg (2008)
4. Bar-Yossef, Z., Friedman, R., Kliot, G.: Rawms - random walk based lightweight membership service for wireless ad hoc networks. ACM Trans. Comput. Syst. 26(2), 1–66 (2008)
5. Brightwell, G., Winkler, P.: Extremal cover times for random walks on trees. Journal of Graph Theory 14(5), 547–554 (1990)
6. Broder, A.Z., Karlin, A.R.: Bounds on the cover time. J. Theoretical Probab. 2, 101–120 (1988)
7. Casteigts, A., Flocchini, P., Quattrociocchi, W., Santoro, N.: Time-varying graphs and dynamic networks. International Journal of Parallel, Emergent and Distributed Systems 27(5), 387–408 (2012)
8. Clementi, A.E.F., Pasquale, F., Monti, A., Silvestri, R.: Communication in dynamic radio networks. In: Proceedings of the Twenty-sixth Annual ACM Symposium on Principles of Distributed Computing, PODC 2007, pp. 205–214. ACM, New York (2007)
9. Clementi, A.E.F., Macci, C., Monti, A., Pasquale, F., Silvestri, R.: Flooding time in edge-markovian dynamic graphs. In: PODC 2008, pp. 213–222. ACM, New York (2008)
10. Cooper, C., Frieze, A.: Crawling on simple models of web graphs. Internet Mathematics 1, 57–90 (2003)
11. Das Sarma, A., Molla, A.R., Pandurangan, G.: Fast distributed computation in dynamic networks via random walks. In: Aguilera, M.K. (ed.) DISC 2012. LNCS, vol. 7611, pp. 136–150. Springer, Heidelberg (2012)
12. Dembo, A., Peres, Y., Rosen, J., Zeitouni, O.: Cover times for brownian motion and random walks in two dimensions. Annals of Mathematics 160(2), 433–464 (2004)
13. Dolev, S., Schiller, E., Welch, J.L.: Random walk for self-stabilizing group communication in ad hoc networks. IEEE Transactions on Mobile Computing 5, 893–905 (2006)
14. Feige, U.: A tight lower bound on the cover time for random walks on graphs. Random Struct. Algorithms 6, 433–438 (1995)

15. Gkantsidis, C., Mihail, M., Saberi, A.: Random walks in peer-to-peer networks: algorithms and evaluation. Perform. Eval. 63, 241–263 (2006)
16. Haeupler, B.: Analyzing network coding gossip made easy. In: Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, STOC 2011, pp. 293–302. ACM, New York (2011)
17. Haeupler, B., Karger, D.: Faster information dissemination in dynamic networks via network coding. In: Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, PODC 2011, pp. 381–390. ACM, New York (2011)
18. Haeupler, B., Kuhn, F.: Lower bounds on information dissemination in dynamic networks. In: Aguilera, M.K. (ed.) DISC 2012. LNCS, vol. 7611, pp. 166–180. Springer, Heidelberg (2012)
19. Horn, R.A., Johnson, C.R.: Matrix analysis. Cambridge University Press, New York (1986)
20. Ikeda, S., Kubo, I., Okumoto, N., Yamashita, M.: Fair circulation of a token. IEEE Transactions on Parallel and Distributed Systems 13(4), 367–372 (2002)
21. Jonasson, J.: On the cover time for random walks on random graphs. Comb. Probab. Comput. 7(3), 265–279 (1998)
22. Kahn, J.D., Linial, N., Nisan, N., Saks, M.E.: On the cover time of random walks on graphs. Journal of Theoretical Probability 2(1), 121–128 (1989)
23. Kirkland, S.: Nonhomogeneous matrix products. World Scientific, River Edge (2002)
24. Kuhn, F., Lynch, N., Oshman, R.: Distributed computation in dynamic networks. In: Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, pp. 513–522. ACM, New York (2010)
25. Kuhn, F., Oshman, R.: Dynamic networks: Models and algorithms. SIGACT News 42(1), 82–96 (2011)
26. Kuhn, F., Oshman, R., Moses, Y.: Coordinated consensus in dynamic networks. In: Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, PODC 2011, pp. 1–10. ACM, New York (2011)
27. Law, C., Siu, K.-Y.: Distributed construction of random expander networks. In: Twenty-Second Annual Joint Conference of the IEEE Computer and Communications, INFOCOM 2003, vol. 3, pp. 2133–2143. IEEE Societies (2003)
28. Leitao, J., Pereira, J., Rodrigues, L.: Hyparview: A membership protocol for reliable gossip-based broadcast. In: In IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2007, pp. 419–428. IEEE Computer Society (2007)
29. Lovász, L.: Random walks on graphs: A survey (1993)
30. Massoulié, L., Le Merrer, E., Kermarrec, A.-M., Ganesh, A.: Peer counting and sampling in overlay networks: random walk methods. In: Proceedings of the Twenty-fifth Annual ACM Symposium on Principles of Distributed Computing, PODC 2006, pp. 123–132. ACM, New York (2006)
31. Perkins, C.E., Royer, E.M.: Ad-hoc on-demand distance vector routing. In: Proceedings of the Second IEEE Workshop on Mobile Computer Systems and Applications, WMCSA 1999, pp. 90–100. IEEE Computer Society, Washington, DC (1999)
32. Stoica, I., Morris, R., Karger, D., Frans Kaashoek, M., Balakrishnan, H.: Chord: A scalable peer-to-peer lookup service for internet applications. In: Proceedings of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM 2001, pp. 149–160. ACM, New York (2001)