Scalability and Freshness in Time-Varying Gossip Networks*

Arunabh Srivastava Thomas Maranzatto Sennur Ulukus

Department of Electrical and Computer Engineering
University of Maryland, College Park, MD 20742
arunabh@umd.edu *tmaran@umd.edu* *ulukus@umd.edu*

Abstract

We study a source that disseminates updates across a network of n gossiping nodes. The network's topology alternates between multiple arbitrary topologies, with the switching dynamics governed by an underlying continuous-time Markov chain (CTMC). While information freshness has been extensively analyzed in static networks, this work investigates how time-varying connectivity influences freshness. To measure the timeliness of information, we employ the version age of information metric. We study the effects of fast and slow switching rates in the CTMC on the average version age of individual nodes and the entire network when all the switching rates are of the same order with respect to n. We also observe that the behavior of a vanishingly small fraction of nodes can drastically impact the average version age of the network. Motivated by this fact, we define the notion of a *typical set of nodes* in the network and find bounds on the average version ages for nodes in this set for both fast and slow switching rates. We then analyze a varying-topology network with two topologies, the ring topology and the fully-connected topology for the entire spectrum of switching rates and discuss the intricacies that arise with various combinations of switching rates. We then turn our attention to a special mobility model, considering it as a special case of the varying-topology network where nodes can exchange positions with other nodes in the network. We use the stochastic hybrid systems (SHS) framework to find recursive equations that determine the version age of a set of nodes in the network in terms of the version ages of sets of the same size and the version ages of sets containing an extra node. Using these recursive equations, we analyze gossip networks with special structures to uncover the impact of mobility on the version age. We perform numerical simulations to analyze the impact of varying topologies and mobility on the version age in gossip networks and show that the bounds obtained in our results are accurate.

*This work was presented in parts at Allerton Conference, Sept. 2024, IEEE ISIT, June 2025 and IEEE ITW, Sept. 2025.

1 Introduction

Advances in communication have propelled the human civilization to great heights and prosperity. Among these advances is the evolution of wireless connectivity, which has enabled multiple technologies such as video conferencing, remote healthcare, internet of things (IoT), and 5G mobile networks. These advancements have led to significant interest in new technologies based on the current 5G communication standards, and the emerging 6G communication standards. These include paradigms such as ultra-reliable low latency communication (URLLC), massive machine type communication (mMTC) [1] and networks with dynamic and adaptive configurations, such as environmental sensor networks, industrial automation networks and drone formations. Such networks operate under time sensitive constraints, and require efficient and adaptive information dissemination strategies.

A key challenge in time-sensitive networks lies in achieving fast and decentralized information exchange among nodes. Centralized architectures struggle with this due to scalability constraints, communication latency, and vulnerability to single points of failure. For instance, in a swarm of drones performing coordinated maneuvers, each drone must continuously share its position with nearby peers. Relying on a centralized controller for this task can introduce delays and bottlenecks as the network grows or encounters interference. This limitation can be mitigated through gossiping protocols [2, 3], which enable information to spread through local, peer-to-peer exchanges. Such protocols provide a scalable, efficient, and fault-tolerant mechanism for information dissemination in dynamic and distributed settings.

Moreover, real-time decision-making in these networks critically depends on the freshness of shared information. Stale or outdated data can lead to inefficiencies and errors, particularly in time-critical applications, such as autonomous vehicle coordination and process control in smart factories. To quantify and manage information freshness, the age of information (AoI) metric has been widely adopted [4–6]. AoI measures the time elapsed since the most recent successful update was generated, and has been extensively used to evaluate network performance [7, 8]. This has further led to the development of several related freshness metrics, including the age of incorrect information (AoII) [9], the age of synchronization (AoS) [10], the binary freshness metric (BFM) [11], and the version age of information [12–14].

In this work, we employ the version age of information (in short *version age*) to quantify information freshness. Unlike AoI, which measures continuous time since the last update, version age is a discrete measure that captures version differences. This distinction makes it particularly suitable for gossiping networks, where updates are timestamped and disseminated in discrete versions. The version age of a node in such a network is defined as the difference between the source version and the node's version. This model was first analyzed in the spatial mean-field regime [15], and later through the stochastic hybrid systems (SHS) framework [12]. Subsequent studies [16–19] have extended this framework to analyze gossiping networks with varying characteristics, including structured topologies and adversarial conditions. A comprehensive review of these developments is presented in [20]. More recently, first-passage percolation was used to characterize the distribution of version age [21].

Many prior works have analyzed mobile networks using AoI as a metric. Chaintreau et al. [15] study a gossiping network composed of multiple mobility classes, where nodes can move between classes. Using the spatial mean field regime from statistical mechanics, the authors derive a differential equation describing the joint distribution of AoI and time. In their model, each node can access a base station at a fixed rate, independent of the number of nodes in the network. In contrast, the model we consider in this paper assumes that the rate of receiving updates from the source decreases as the network size increases, capturing a more realistic congestion effect. In his masters thesis, Tripathi [22] considers AoI minimization and UAV trajectory design for both information gathering and information dissemination scenarios. The author studies networks with mobile nodes as well as stationary nodes served by mobile UAVs, and employs random walk and i.i.d. mobility models. Similarly, Wu et al. [23] addresses trajectory planning for multiple relay UAVs to minimize AoI, using reinforcement learning methods. Eldeeb et al. [24] extends this line of work by jointly optimizing energy consumption and trajectory planning alongside AoI minimization.

In this paper, we investigate the evolution of version age in gossiping networks with time-varying topologies, where the connections in the underlying network over which nodes exchange updates are time-varying and the nodes receive updates from a source which has the latest update. Our goal is to quantify the impact of network dynamics on the version age of information of nodes in the network. Specifically, we consider networks which switch between a finite set of fixed topologies. The transitions between topologies in this set occur according to the state changes of an underlying continuous time Markov chain (CTMC); see Fig. 1 for an example. We will discuss the system model in detail in Section 2.

Our main contributions are as follows: We analyze gossip networks with arbitrary topologies for a large spectrum of fast CTMC transition rates, and find bounds on the average version ages of nodes, as well as average version age of the entire network. Additionally, we establish certain fundamental properties of single-topology networks in Section 3. In Section 4, we analyze the performance of varying-topology networks with slow switching rates, where one of the topologies is the ring topology or the disconnected topology. We define a new concept, the *typical set of nodes* in Section 5, inspired by the definition of a typical set by Shannon [25]. This definition is motivated by the fact that a small subset of nodes can disproportionately increase the version age of the entire network. We prove results for varying-topology networks that hold for all but a small subset of nodes, for both fast and slow switching rates in the CTMC. In Section 6, we explore the effects of unequal switching rates in a varying-topology network with two topologies: the ring topology and the fully-connected topology. We find the average version ages of nodes in the network for the entire spectrum of switching rates and explain the nuances that arise due to this.

We then turn our attention to a special case of varying-topology network, a mobility model, in Section 7. In this model, nodes in the gossiping network can exchange positions with each other. This model corresponds to an extremely fast switching regime which is difficult to analyze using the general techniques and results developed earlier in the paper.

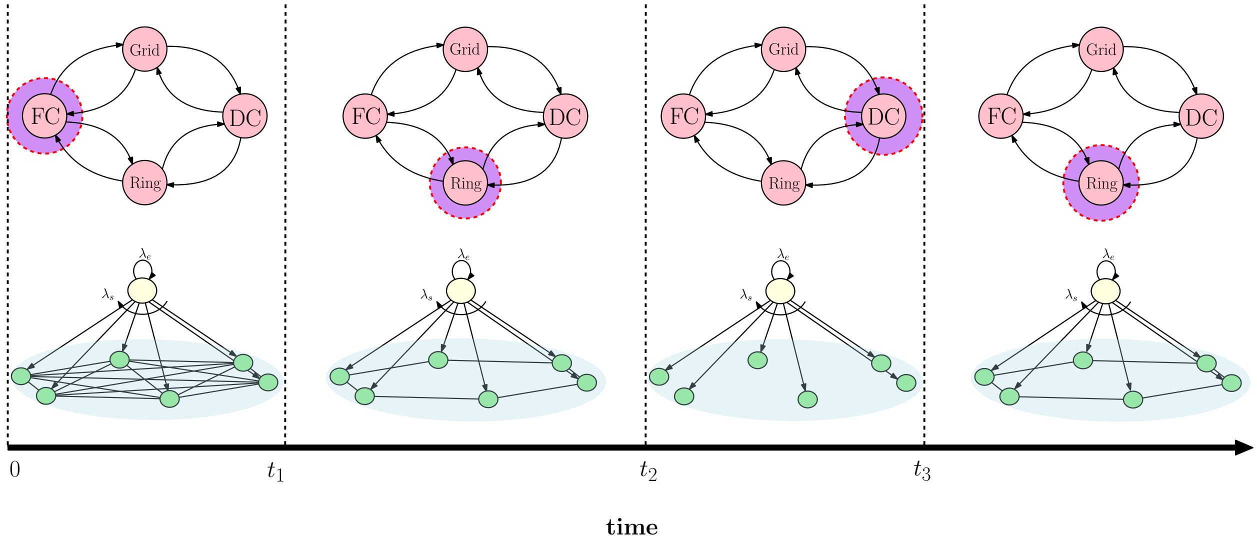

Figure 1: A description of the time evolution of the connectivity of the gossip network. The state of the CTMC represents the state of the connectivity graph (topology) of the gossip network. In this example, the CTMC has four states, representing the fully connected (FC) topology, the ring topology, the grid topology and the disconnected (DC) topology. The current topology is represented using a red dashed circle filled with purple color. In this example, the network starts in the fully connected topology, moving to the ring topology, the disconnected topology and back to the ring topology. Since the rates of the CTMC are different, the time spent in a state will be unequal and random, as shown in the figure.

Thus, we begin this section by discussing why we can use the SHS framework to analyze mobility networks, followed by using the framework to find a set of recursive equations to find the version age of a set of positions in the network. These recursive equations find the version age of a set of positions in terms of sets of positions that are one larger, and sets with the same number of positions. We then analyze how mobility affects the version age of different positions in the network using toy examples. Then, we analyze two specific gossiping networks and use the previously found recursive equations to find upper bounds for the average version age of a single node in both networks. In Section 8, we discuss many properties of varying-topology networks. We show numerical results in Section 9, and conclude the paper in Section 10.

2 System Model and the Version Age Metric

2.1 Generic Time-Varying Topologies

We consider a source node n_0 sending updates to a set of nodes $\mathcal{V} = \{1, \ldots, n\}$. The source updates itself via a Poisson process with rate λ_e. The source also sends updates to every $j \in \mathcal{V}$ via independent Poisson processes each with rate $\frac{\lambda_s}{n}$. At any time $t \geq 0$, nodes in $\mathcal{V}$ can send updates to their neighbors. However, the set of neighbors for a node may be different at different times. We model the dynamics of the edge set by considering a finite set of networks

$\mathcal{K} = \{1, 2, \ldots, K\}$ where $K < \infty$. We refer to $\mathcal{K}$ as the *set of topologies* for the system. We view $\mathcal{K}$ as the state space for a continuous time Markov chain (CTMC), with transition rates q_{kl} between states $k \in \mathcal{K}$, $l \in \mathcal{K}$, and total rate of exiting a state given by $q_k = \sum_{l \in \mathcal{K}} q_{kl}$. We also assume that the holding time in each state k of the Markov chain, defined as the random variable associated with the time it takes to leave the state, is of the same order as a function $h_k(n)$ of n. This means that the CTMC stays in a state for a time that is exponentially distributed with mean of the same order as $h_k(n)$. For example, Fig. 1 shows a case with $n = 6$ nodes, where there are $K = 4$ topologies with $\mathcal{K} = \{\mathrm{FC}, \mathrm{DC}, \mathrm{Grid}, \mathrm{Ring}\}$, with FC and DC standing for fully-connected and disconnected, respectively, and Fig. 2 shows a case with $n = 6$ nodes and $K = 2$, with the two topologies being fully-connected and ring; here, $h_1(n) = h_2(n) = h(n)$, thus the transition rates between the two topologies are of identical orders denoted by $h(n)$.

When the network is in topology k, node i sends updates to node j via another independent Poisson process with rate $\lambda_{i,j}^{(k)}$. We assume that all nodes communicate with the total rate λ, which we assume to be constant. That is, setting $\mathcal{N}^{(k)}(i)$ to be the neighbors of i in topology k, we assume $\sum_{j \in \mathcal{N}^{(k)}(i)} \lambda_{i,j}^{(k)} = \lambda$, for all $i \in \mathcal{V}$ and $k \in \mathcal{K}$. We assume $\mathcal{K}$ is constant as $n \to \infty$, but that the transition rates and probabilities can depend on n. For the rest of the paper we set $\lambda_s = \lambda$.

Thus far, we have only discussed how updates are sent, not the messages contained in each update. In our model, the source and every node in the network have an associated counting process which represents versions of updates it carries; when a node $i \in \mathcal{V} \cup n_0$ communicates to a neighbor j (because node i's Poisson process updated), i sends its current version. The version for n_0 increments by 1 if and only if the counting process N_0 for n_0 updates. A node $j \in \mathcal{V}$ updates its version at time t if and only if some other node with a larger version sends a message to j at time t. Letting $N_i(t)$ be the counting process associated with the version at node $i \in \mathcal{V}$, the version age of node i at time t is $X_i(t) = N_0(t) - N_i(t)$. Since the source has the freshest versions of updates, the version age $X_i(t)$ is how many versions node i is behind the source at time t. Further, if G is a static network, we write $X_i^G(t)$ to mean the version age of node i at time t in network G. We define the long-term average version age of a varying-topology network as $v_{av} = \lim_{t \to \infty} \frac{1}{n} \sum\limits_{i \in \mathcal{V}} X_i(t)$.

Throughout, we assume λ_e, λ are constants and let $n \to \infty$. We say that an event $\mathcal{A}$ holds with high probability (w.h.p.) if $\mathbb{P}[\mathcal{A}] \to 1$ as $n \to \infty$. We use the standard big-O definitions. We say that $f(n) = O(g(n))$ if $\limsup_{n \to \infty} \frac{f(n)}{g(n)} < \infty$, $f(n) = \Omega(g(n))$ if $\limsup_{n \to \infty} \frac{g(n)}{f(n)} < \infty$, $f(n) = \Theta(g(n))$ if $c_1 \leq \lim_{n \to \infty} \frac{f(n)}{g(n)} \leq c_2$ for some constants $c_1 \leq c_2$, $f(n) = o(g(n))$ if $\lim_{n \to \infty} \frac{f(n)}{g(n)} = 0$ and $f(n) = \omega(g(n))$ if $g(n) = o(f(n))$. We denote sets using calligraphic notation.

Our analysis is stated in the form of infinite network sequences of the form $\{\mathcal{G}(n)\}_{n=1}^{\infty}$, where network $\mathcal{G}(n)$ has n nodes. Other works in gossiping have implicitly used graph sequences in their proofs, but since the networks considered had nice properties, such as vertex transitivity (e.g., the generalized rings networks in [26]) or easy parametrization (e.g.,

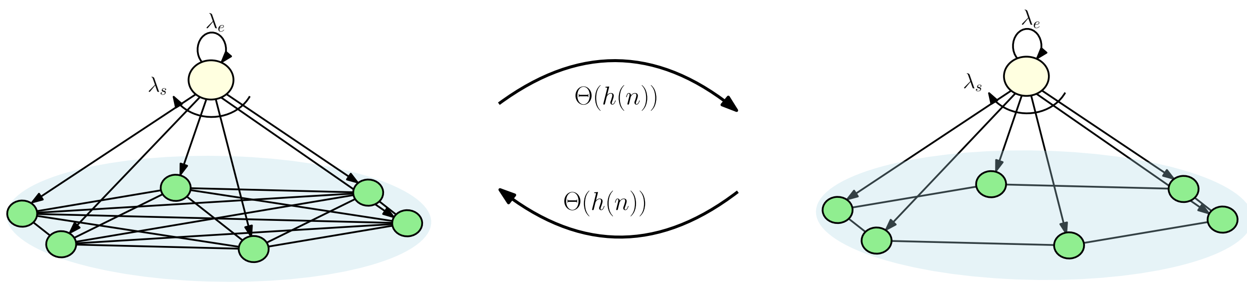

Figure 2: An example of a varying-topology network we study. Here, $n = 6$, $K = 2$, topology 1 is the fully-connected topology, and topology 2 is the ring topology. We have omitted the edge-rates, but these implicitly change as the network switches between topologies 1 and 2.

the Erdos-Renyi graph in [19]), the presentation using graph sequences could be replaced with more straightforward methods. Our primary objective is to understand what happens when large networks which have version ages scaling as $\{\Theta(f_k(n))\}_{k\in\mathcal{K}}$ interact via a CTMC with holding times $\{h_k(n)\}_{k\in\mathcal{K}}$. In this paper, we consider a running example where the topologies switch between a fully-connected and a ring topology[1] as in Fig. 2. We use this example to illustrate our results in a straightforward manner, where applicable.

2.2 A Specific Mobility Model

Here, we describe the system model for the mobility network model in Section 7, which is a special case of the varying-topology network. The nodes are placed in arbitrary fixed positions in the network at the start of the process. We name these positions as $\{1, 2, \ldots, n\}$ serially starting at an arbitrary position. We name the nodes in each position at the start of the process with the same label as the position. In terms of the source sharing updates with the nodes in the gossiping network, we adopt a broader definition when compared to the general varying-topology model. In this case, the source node shares its updates with nodes in the gossiping network as a combined rate λ Poisson process. The source sends updates to position j as a rate $\lambda_{0,j}$ Poisson process, independent of all other processes in the network. Thus, $\sum_{j\in\mathcal{V}} \lambda_{0,j} = \lambda$. We assume that the source update sharing process is fixed and independent of all other processes in the network.

Each node in the gossiping network can communicate with its neighbors as a Poisson process independent of all other processes in the network. The set of neighbors each node can communicate with depends on the position it occupies in the network. If a node in position i sends an update to a node in position j, then the rate of information flow, i.e., the rate of the Poisson process of updates sent from position i to j is $\lambda_{i,j}$.

We define the version age of a position i at time t as the version age of the node that

[1]These are two of the four well-known examples of distinct version age scalings in networks, where age order increases as connectivity decreases: Version age scales as $\log(n)$ in fully-connected networks [12]; as $n^{\frac{1}{3}}$ in grid networks [27]; as $n^{\frac{1}{2}}$ in ring networks [28]; and as n in disconnected networks.

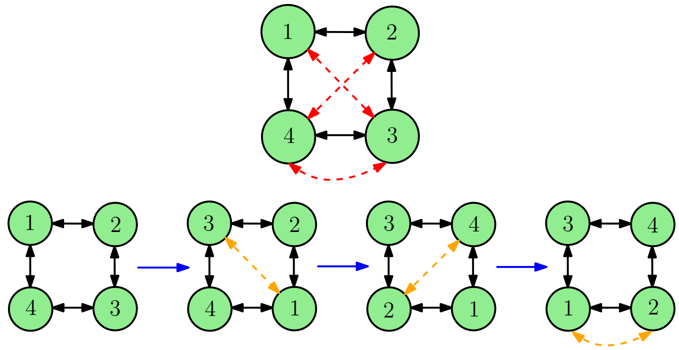

Figure 3: An illustration of the mobility model in the gossiping network. The positions where node exchanges are possible are denoted by red arrows in the network on the top. As shown by the yellow arrows in the process, nodes 1 and 3 exchange positions, followed by nodes 2 and 4, and finally nodes 1 and 2 exchange positions.

occupies position i at that time. The version age of a set $\mathcal{S}$ of fixed positions in the network is defined as $X_{\mathcal{S}}(t) = \min_{j\in\mathcal{S}} X_j(t)$. We also define the expected value and limit as $v_{\mathcal{S}}(t) = \mathbb{E}[X_{\mathcal{S}}(t)]$ and $v_{\mathcal{S}} = \lim_{t\to\infty} v_{\mathcal{S}}(t)$, respectively.

We say that position i is a neighboring position of set $\mathcal{S}$ if $\lambda_{i,j} > 0$ for some $j \in \mathcal{S}$, and define the set of neighboring positions of $\mathcal{S}$ as $\mathcal{N}(\mathcal{S})$. Next, we define the rate of information flow from a node in position i into set $\mathcal{S}$ as $\lambda_i(\mathcal{S}) = \sum_{j\in\mathcal{S}} \lambda_{i,j}$ if $i \notin \mathcal{S}$ and $\lambda_i(\mathcal{S}) = 0$ if $i \in \mathcal{S}$. Similarly, we define the rate of information flow from the source to set $\mathcal{S}$ as $\lambda_0(\mathcal{S})$.

Next, we define the mobility processes in the gossiping network. A node in position i in the network can exchange positions with nodes in a subset $\mathcal{M}_i$ of the other positions in the network. The set $\mathcal{M}_i$ may be empty. We assume that all node exchange processes are independent of each other and every other process in the network. For any position i in the network and any position $j \in \mathcal{M}_i$, the exchange process is defined as a rate $\lambda_{i,j}^{\mathcal{M}}$ Poisson process, with superscript $\mathcal{M}$ signifying *mobility*. Hence, each node exchange depends on the positions, but not on the nodes present in this position. The node exchange process is illustrated in Fig. 3. We say that a gossip network has full-mobility if a node in any position can exchange positions with any other node in the network.

Finally, we provide Table 1 which summarizes the notation used throughout the paper.

3 Fast Topology Switching

In this section, we discuss the results related to the average version age of nodes in the varying-topology network when the holding times in each state of the CTMC is lower than the average version age of nodes in that state.

Suppose we have $K = |\mathcal{K}|$ distinct network topologies, topology $1, \ldots, K$, represented using $\mathcal{G}_1, \ldots, \mathcal{G}_K$, respectively. Suppose the long-term average version ages of $\mathcal{G}_k$, $k \in \mathcal{K}$ is $\Theta(f_k(n))$. It follows from the bounds in [12] that we can take the rates $\Theta(\log n) \leq f_k(n) \leq$

Notation	Definition
n	number of nodes in the gossiping network
λ_e	rate of source self-update Poisson process
λ_s	rate of updates sent from the source to the gossiping network
$\mathcal{V}$	set of all nodes in the gossiping network
$\mathcal{K}$	set of topologies in the network
K	number of topologies in the network
q_k	rate of leaving CTMC state k
$\mathcal{G}_k$	topology associated with state k in the CTMC
λ	rate of Poisson process with which each node shares information with its neighbors
$N_0(t)$	timestamp of the update with the source (node 0) at time t
$N_i(t)$	timestamp of the update with node i at time t
$X_i(t)$	version age of node i at time t
$X_{\mathcal{S}}(t)$	version age of a set $\mathcal{S}$ of nodes at time t
$v_{\mathcal{S}}$	long-term average version age of a set $\mathcal{S}$ of positions
$\lambda_{i,j}$	rate of Poisson process associated with sending information from node i to node j
$\lambda_i(\mathcal{S})$	total rate of updates going from node i to set $\mathcal{S}$ of nodes/positions
$\mathcal{N}(\mathcal{S})$	number of neighboring nodes of set $\mathcal{S}$ of nodes/positions
$\mathcal{M}_i$	set of positions in which nodes can exchange with node in position i
$\lambda_{i,j}^{\mathcal{M}}$	rate of node exchanges between positions i and j
γ	Euler-Mascheroni constant

Table 1: Notation and definitions.

$\Theta(n), \forall k \in \mathcal{K}$.[2] Furthermore, assume without loss of generality, that $f_1(n) = o(f_k(n))$, $\forall k \in \{2, \dots, K\}$, and suppose the CTMC rates satisfy $q_k = \Theta(1/h(n))$ and holding times are $\Theta(h(n))$ for all $k \in \mathcal{K}$. An example of this setup with holding times and two topologies is illustrated in Fig. 2. In this subsection, we characterize the long-term average version age of the network when $h(n)$ varies between 0 and $\Theta(f_1(n))$. For example, if the transition rates are constants then it is very unlikely an update will reach every node without a CTMC update occurring. We describe this region for $h(n)$ as the fast switching regime. The behavior of the fast switching process is studied in Theorem 1.

Theorem 1 *Let $\{\mathcal{G}_k(n)\}_{n=1}^{\infty}$, $k \in \mathcal{K}$ be K sequences of networks. Suppose $\mathcal{G}_k(n)$ has a long-term average version age $\Theta(f_k(n))$. Further, let $f_1(n) = o(f_k(n))$, $\forall k \in \mathcal{K} \backslash \{1\}$. Let the transition rates have the same order, i.e., $q_k = \Theta(1/h(n))$, $\forall k \in \mathcal{K}$. If $h(n) = O(f_1(n))$, then the long-term average version age for the time-varying system is $\Theta(f_1(n))$.*

Intuitively, Theorem 1 states that if the network switches between topologies very quickly (quicker than the version age scaling of the freshest topology), then the long-term average version age of the network is able to perform at a similar level to the freshest topology. This is because the varying-topology network never stays in staler topologies long enough for it to negatively affect the version age of any node. In our running example, Theorem 1 states that if we have two sequences of graphs, one the sequence of fully-connected topologies, and

[2]Since the version age in a fully-connected network scales as $\Theta(\log(n))$ and the version age in a disconnected network scales as $\Theta(n)$, which are the two extremes of connectivity.

the other the sequence of ring topologies, the long-term average version age of the varying-topology network with n nodes has logarithmic scaling if the holding times in both states of the CTMC have mean $O(\log n)$. This is because, the version age scaling in the fully-connected networks follows $f_1(n) = \Theta(\log n)$, and the version age scaling in the ring networks follows $f_2(n) = \Theta(\sqrt{n})$. Thus, since $f_1(n) = o(f_2(n))$, the long-term average version age scaling for the running example is $\Theta(\log n)$.

We need the following three lemmas to prove Theorem 1.

Lemma 1 *Let $\{\mathcal{G}(n)\}_{n=1}^{\infty}$ be a sequence of networks. Consider node 1 without loss of generality. Suppose $\lim_{t\to\infty} \mathbb{E}[X_1^{\mathcal{G}(n)}(t)] = \Theta(f(n))$. Then, for large enough n and t, with high probability, $X_1^{\mathcal{G}(n)}(t) = O(f(n))$.*

Lemma 2 *Suppose the holding times of each state in the CTMC are $\Theta(h(n))$. Then, if the CTMC has spent $\Theta(g(n))$ (where $g(n) = \omega(h(n))$) time in state $k \in \mathcal{K}$ at time T, with high probability, $T = \Theta(g(n))$.*

Lemma 3 *If $T = \Theta(f(n))$, then with high probability, $N_0([0,T]) = \Theta(f(n))$, where $N_0([0,T])$ is the number of times the source node updates itself between time 0 and time T.*

Lemma 1 shows that any gossiping network is able to maintain the version age of all nodes in the network at or below the node's long-term average version age (in the order sense). Lemma 2 shows that we spend the same amount of time in each state (in the order sense) as a function of n. Lemma 3 shows that the number of times source updates itself, i.e., generates a new update, is of the same order as the time that has elapsed. At a high-level, the proofs of these three lemmas follow by using the properties of CTMCs and applications of Markov and Chebyshev inequalities. Their detailed proofs are given in the Appendix. We are now ready to prove the main result of this section given in Theorem 1.

Proof: Our strategy in the proof is to show that we can design two modified networks from the original network, one fresher and one staler, such that they have order-wise the same average version age as the original network. Since these two modified networks form a lower and an upper bound, respectively, they sandwich the original network, yielding the result.

For the lower bound, we obtain a modified topology CTMC by replacing all $\mathcal{G}_k(n)$, $k \in \mathcal{K}$, with $\mathcal{G}_1(n)$, which is the freshest topology. We note that this constitutes a gossip network with time-invariant edge rates and the long-term average version age is $\Theta(f_1(n))$.

For the upper bound, we obtain a modified topology CTMC by replacing all topologies other than the freshest topology $\mathcal{G}_1(n)$ with the disconnected (DC) topology. This still gives a varying-topology network with two topologies: the freshest topology $\mathcal{G}_1(n)$ in the original setting and the absolutely stalest topology possible, DC. For the upper bound, the dissemination of information happens only through the freshest topology in this modified network and the disconnected (DC) topologies essentially act as pauses. For clarity, the modifications from the original varying-topology network to the lower bound and upper bound modified varying-topology networks are shown in Fig. 4.

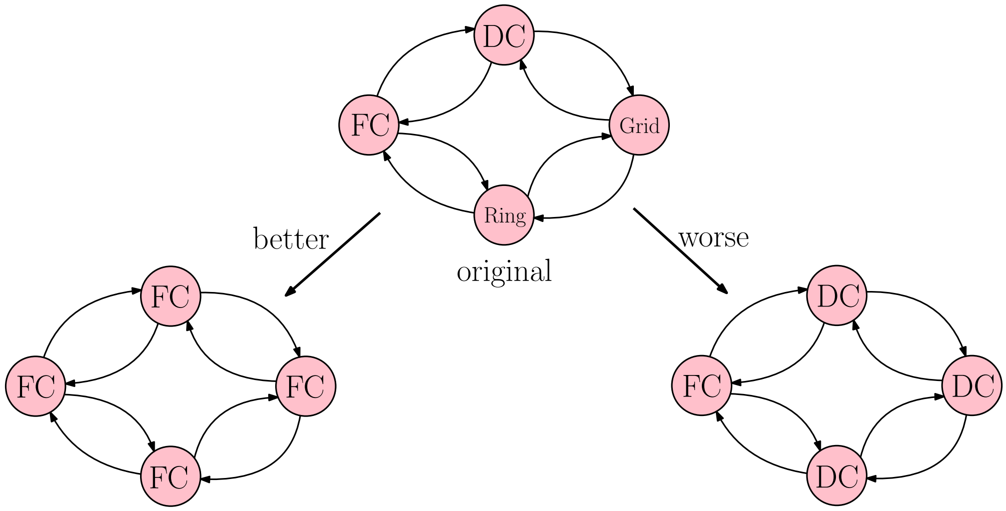

Figure 4: An example CTMC containing four topologies. The states can be any fixed topology. In this example, the fully-connected (FC), disconnected (DC), ring and grid topologies are present. In the proof of Theorem 1, we obtain a lower bound by changing all the staler topologies to the freshest topologies, which is the FC topology in this case. We also obtain an upper bound by replacing all the staler topologies other than the freshest topology with the DC topology. These are represented in the figure.

In the rest of the proof, we focus on the upper bound. Further, we focus our attention to a single node in the modified network and utilize Lemma 1, Lemma 2 and Lemma 3 to show that the average time it takes for an update to arrive from the source to this specific node is still of the same order as the time it takes for the update to travel from the source to the node in $\mathcal{G}_1(n)$. We then average over all nodes to obtain an upper bound.

It now suffices to show that in this modified varying-topology network, where $\mathcal{G}_k(n)$, $k \in \mathcal{K}\backslash\{1\}$ is replaced with the DC topology, the long-term average version age is $O(f_1(n))$. We proceed by considering the version age of a single node in this modified process. Without loss of generality, we take this node to be the node labeled 1 throughout the graph sequence. Suppose the long-term average version age of node 1 scales as $\Theta(g(n))$ in the single topology network $\mathcal{G}_1(n)$. Since $\mathcal{G}_1(n)$ is a general topology, we cannot relate node 1's version age with $f_1(n)$, because it is the long-term average version age of the topology. Then, from Lemma 1, the version age of node 1 at any time in the single topology network $\mathcal{G}_1(n)$ is $O(g(n))$ w.h.p. This means that an update sent by the source reaches node 1 in $O(g(n))$ time w.h.p. We aim to show that this results in a version age scaling of $O(g(n)+h(n))$ w.h.p. for node 1 in the modified varying-topology network. To this end, we consider two cases.

We first consider the case where $g(n) = o(h(n))$. In this case, every update that reaches node 1 reaches within one holding time in state 1 of the CTMC. This is because any update which reaches node 1, reaches in $O(g(n))$ time w.h.p. due to Lemma 1. Thus, node 1 maintains $\Theta(g(n))$ version age during the holding time of state 1 of the CTMC. Further, during the holding times of the DC topologies, node 1's version age increases by $\Theta(h(n))$ due to Lemma 3, since we spend $\Theta(h(n))$ time in all the states, and node 1 does not receive an update from the source w.h.p. since $h(n) = o(n)$. Thus, the average version age of node

1 in this case is $\Theta(h(n))$.

We next consider the case where $g(n) = \Omega(h(n))$. Let $\beta > 0$, and let $[s_1, t_1)$, $[s_2, t_2)$, $\ldots$, $[s_k, t_k)$ be the intervals of time that the CTMC is in state 1, such that $\sum_{i=1}^{k}(t_i - s_i) = \beta g'(n)$, $g'(n) = O(g(n))$. We choose $g'(n)$ in this way because an update from the source can reach node 1 in $O(g(n))$ time w.h.p. due to Lemma 1, but we do not know the exact scaling every time. If $g'(n) = o(h(n))$, then case 1 applies. Otherwise, if we consider a static gossip network on $\mathcal{G}_1(n)$ for total time $T = \sum_{i=1}^{k}(t_i - s_i)$, then the stationary property of Poisson processes gives that the distribution of arrivals of the Poisson process associated with any edge in the varying-topology network is the same in $\mathcal{G}_1(n)$ for the interval $[0, T)$ as for the varying-topology network conditioned on the intervals $[s_i, t_i)$. This means that the probability that an update from the source reaches node 1 in $\beta g'(n)$ time is the same, whether the time is spent continuously in $\mathcal{G}_1(n)$, or in multiple disjoint intervals of topology 1 in the modified varying-topology network. We can then conclude that the first passage percolation time has the same distribution in both cases.

Thus, in the modified varying-topology network, an update from the source spreads from the source through the network and reaches node 1 after $\beta g'(n)$ time has been spent in topology 1. During this time, Lemma 2 tells us that we have spent a total of $\Theta(g'(n))$ time in the modified varying-topology network, and hence the source has updated itself $\Theta(g'(n))$ times from Lemma 3. Thus, if we consider that an update was sent by the source at time s_1 and another update from the source reaches node 1 at time $t_k = \beta g'(n)$, then the source would have updated itself $\Theta(g'(n))$ times during this interval. At time s_1, the version age of node 1 would be $O(g(n))$ w.h.p. When the next update arrives, its version age in the static network would be $\Theta(g'(n))$. An additional $\Theta(g'(n))$ time has passed in the modified varying-topology network while in the disconnected topologies w.h.p. Thus, the version age of the update when it reaches node 1 in the modified varying-topology network is still $\Theta(g'(n))$ w.h.p. Thus, the version age of node 1 in the modified varying-topology network is $O(g(n))$ w.h.p. in this case, since the version age changed from $O(g(n))$ to $O(g(n) + g'(n))$ during the time interval (s_1, t_k). In this argument, we started by saying that the version age of node 1 is $O(g(n))$ at time s_1. We can make this claim because we assume that Lemma 1 is valid, and the node has already received an update directly from the source. Then, the argument we just used can be applied every time node 1 receives an update.

These two cases prove that the version age of node 1 is $O(g(n) + h(n))$ w.h.p. Averaging over all nodes and using the fact that $h(n) = O(f_1(n))$ yields that the average version age of the modified varying-topology network is $O(f_1(n))$.

Finally, combining the upper and lower bounds proves the theorem. ■

4 Slow Topology Switching

In this section, we study the impact of slow CTMC transitions on version age. We study two examples: One where the stalest network is the ring network and one where the stalest network is the DC network. As a recap, the long-term average version age of a node in the ring network is $\Theta(\sqrt{n})$ as proved in [28], and the long-term average version age of a node in the DC network is $\Theta(n)$ as can be proved using the recursive equations found in [12].

Our first result in this section shows that the average version age scales as $\Theta(\sqrt{n})$ if one of the topologies in the varying-topology network is the ring topology, the other topologies have fresher long-term average version age compared to the ring network, and the expected holding times in each state of the CTMC are $\Omega(\sqrt{n})$.

Theorem 2 *Let $\{\mathcal{G}_k(n)\}_{n=1}^{\infty}$, $k \in \mathcal{K}$ be K sequences of networks. Suppose $\mathcal{G}_k(n)$ has a long-term average version age $\Theta(f_k(n))$. Topology K is a sequence of bi-directional ring networks, which have long-term average version age $\Theta(\sqrt{n})$. Further, let $f_k(n) = o(\sqrt{n})$, $\forall k \in \mathcal{K}\backslash\{K\}$. Let the transition rates have the same order, i.e., $q_k = \Theta(1/h(n))$, $\forall k \in \mathcal{K}$. If $h(n) = \Omega(\sqrt{n})$, then the long-term average age scaling for the varying-topology network is $\Theta(\sqrt{n})$.*

This theorem tells us that if the network spends enough time, i.e., $\Omega(\sqrt{n})$, in the ring topology, then the version age of each node converges in order to the long-term average version age of the ring topology. The time spent in other topologies is not sufficient to bring down the average version age, and the ring topology dominates the average version age.

We remark that performing an analysis for a general two-topology network is more challenging since it is possible for a few nodes in the network to have a disproportionate effect on the long-term average version age of the entire network. We discuss this in more detail in the next section.

The proof of Theorem 2 requires the following lemma.

Lemma 4 *Consider a bi-directional ring network containing n nodes. Suppose a node receives an update directly from the source at time t, then the node reaches version age $\Theta(\sqrt{n})$ w.h.p. between time t and $t + \alpha\sqrt{n}$, where $\alpha > 0$ is a constant. Thus, the average version age of the node during this time interval is $\Theta(\sqrt{n})$.*

Lemma 4 shows that any node in the ring network receives updates at least once in $\Theta(\sqrt{n})$ time w.h.p. Its proof is given the Appendix.

We are now ready to prove the main result of this section given in Theorem 2.

Proof: We show that the upper bounds and lower bounds for the long-term average version age of the varying-topology network are of the same order with respect to n. In order to obtain the lower bound, we consider a modified varying-topology network which replaces all topologies fresher than the ring topology with a topology in which all nodes maintain 0 version age. We then show that the average version age of this modified varying-topology network is also $\Theta(\sqrt{n})$ by using Lemma 4. An upper bound can be obtained by replacing topologies $1, \dots, K-1$ with topology K. This leads to the edges becoming time-invariant after replacement, giving rise to the long-term average version age of $\Theta(\sqrt{n})$. Thus, the long-term average version age for each node in the varying-topology network is $O(\sqrt{n})$.

We now focus on the lower bound. We first prove the result when $h(n) = \Theta(\sqrt{n})$. The modified varying-topology network is created by replacing topologies $1, \dots, K-1$ with a

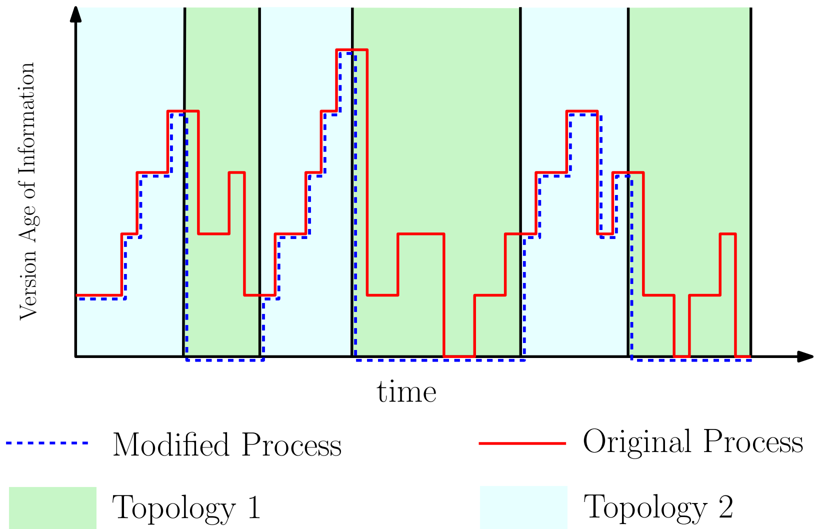

Figure 5: This plot shows the evolution of the version age process in a two-topology network when topology 2 is a ring network. The actual sample path is represented by the solid red lines. The modified sample path obtained when we replace topology 1 with the topology described in the proof of Theorem 2 is represented with the dashed blue lines. We can see that the area under the curve for the modified process is smaller than the area under the curve for the original process.

network which contains infinite rate connections between each node in the network and the source node. Hence, whenever the source updates itself, the nodes also receive the update immediately, and thus maintain 0 version age. This is illustrated in Fig. 5.

Consider a time interval which starts when we exit state K, and ends when we exit state K again. We showed in Lemma 4 that the area of the version age curve for any node in topology K is $\Theta(n)$ w.h.p. Next, the area of the version age curve during the time spent in topologies $1, \ldots, K-1$ is 0, since the version age of each node is 0 during this interval. Thus, the total area of the version age curve in the time interval is $\Theta(n)$ w.h.p. We also showed in Lemma 2 that w.h.p. we spend $\Theta(\sqrt{n})$ time in the CTMC if we spent $\Theta(\sqrt{n})$ time in topology K. Thus, the entire time interval has length $\Theta(\sqrt{n})$ w.h.p. Therefore, the modified varying-topology network is able to maintain $\Theta(\sqrt{n})$ average version age in the time interval w.h.p. This leads to the average version age being $\Omega(\sqrt{n})$ w.h.p. in the original varying-topology network. The upper bound and lower bound yield the fact that the long-term average version age is $\Theta(\sqrt{n})$.

We can extend this proof to the case where $h(n) = \Omega(\sqrt{n})$ by dividing time into time intervals of length $\Theta(\sqrt{n})$, where topology K is visited and exited within each time interval.

Then, due to the earlier case, the average version age in each time interval is $\Theta(\sqrt{n})$ w.h.p. Thus, the average version age in the entire interval is $\Omega(\sqrt{n})$ w.h.p. The rest of the proof follows similarly to the case where $h(n) = \Theta(\sqrt{n})$. $\blacksquare$

We can use a similar proof technique to show that the average version age is linear in a varying-topology network where one of the topologies is disconnected and the holding times are linear or superlinear in n.

Theorem 3 *Let $\{\mathcal{G}_k(n)\}_{n=1}^{\infty}$, $k \in \mathcal{K}$ be K sequences of networks. Suppose $\mathcal{G}_k(n)$ has a long-term average version age $\Theta(f_k(n))$. Topology K is a disconnected network that has a long-term average version age $\Theta(n)$. Let the transition rates have the same order, i.e., $q_k = \Theta(1/h(n))$, $\forall k \in \{1, \dots, K\}$. If $h(n) = \Omega(n)$, then the long-term average age scaling for the varying-topology network is $\Theta(n)$.*

5 Typical Sets and Varying Topologies

If the underlying topologies are not vertex-transitive (i.e., vertices may have different neighborhoods) then there can exist a large variation in version age of different nodes. Consider the network illustrated in Fig. 6, where the node in the fully-connected sub-network which is connected to the linear sub-network sends updates to the node in the linear sub-network with a constant rate. It can be shown that the long-term average version age for this network is $\Theta(n^{\frac{1}{3}})$. However, the nodes in the fully-connected sub-network have $\Theta(\log n)$ expected version age, while the nodes in the linear sub-network have $\Theta(n^{\frac{2}{3}})$ expected version age. This is because the linear sub-network is connected to a node of the fully-connected sub-network which has $\Theta(\log n)$ version age. Any packet sent by the node in the fully-connected sub-network reaches a node in the linear sub-network with version age proportional to how far the node is from the fully-connected sub-network. We can then average over all nodes in the linear sub-network to obtain the required version age scaling. Thus, if $h(n) = \Theta(n^{\frac{1}{3}})$ and the second topology is fully-connected, the nodes in the linear network never reach the version age they have in the single-topology network. This shows that the a small subset of nodes can have unpredictable effects on the long-term average version age of a general varying-topology network. Therefore, the average age analysis from Section 3 cannot apply to general gossip networks when $h(n) = \Omega(f_1(n))$. This discussion motivates the following definition.

Definition 1 *Let $\{\mathcal{G}(n)\}_{n=1}^{\infty}$ be a sequence of networks. Suppose $\mathcal{G}(n)$ has a long-term average version age $\Theta(f(n))$. Then, the typical set is defined as the set of nodes in $\mathcal{G}(n)$ with long-term average version age $O(f(n))$. Next, suppose $\{\mathcal{G}_1(n)\}_{n=1}^{\infty}$, $\{\mathcal{G}_2(n)\}_{n=1}^{\infty}$, ..., $\{\mathcal{G}_K(n)\}_{n=1}^{\infty}$ are sequences of networks. If $\mathcal{G}_1(n), \mathcal{G}_2(n), \dots, \mathcal{G}_K(n)$ are topologies in the varying-topology network, then the typical set of the varying-topology network composed of n nodes is the intersection of their respective typical sets.*

Analysis of the typical set of nodes describes the behavior of version age of a typical node

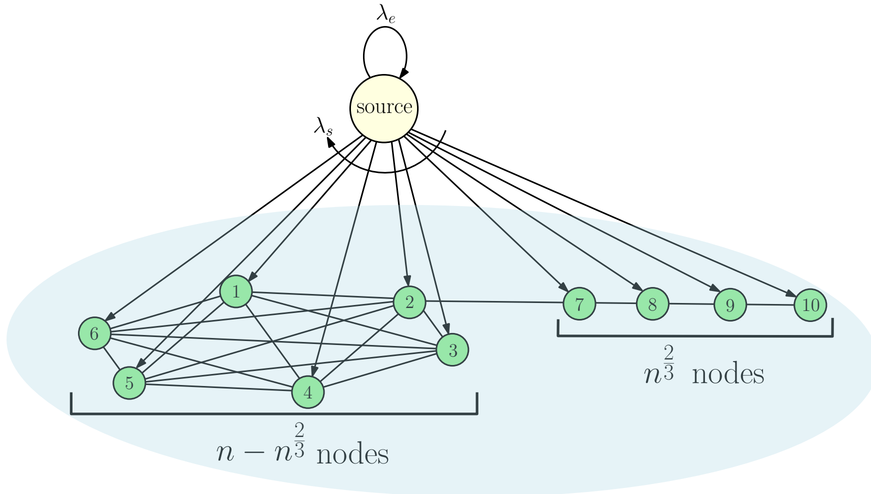

Figure 6: A network formed by combining a fully-connected network with $n - n^{\frac{2}{3}}$ nodes to a line with $n^{\frac{2}{3}}$ nodes, where the node in the fully-connected sub-network which is connected to the linear sub-network sends updates to the node in the linear sub-network with a constant rate. It can be shown that the long-term average version age for this network is $\Theta(n^{\frac{1}{3}})$. However, the nodes in the fully-connected sub-network have $\Theta(\log n)$ expected version age, while the nodes in the linear sub-network have $\Theta(n^{\frac{2}{3}})$ expected version age. This network poses challenges to the average-case analysis with slow topology switching. For example, if $h(n) = \Theta(n^{\frac{1}{3}})$ and the second topology is fully-connected, the nodes in the linear sub-network never reach the version age they have in the single-topology network.

in the varying-topology network. As we shall show in Lemma 5, the typical node is a good representative of almost all the nodes in the network, when analyzing version age behavior. In our running example, all nodes in the fully-connected topology and the ring topology are a part of the typical set of nodes. In Fig. 6, the nodes in the fully-connected part of the network and a few nodes in the linear part of the network are a part of the typical set of nodes. Now, we will show that there are at most $o(n)$ vertices outside the typical set.

Lemma 5 *Let $\mathcal{G}(n)$ be a static gossip network. The typical set of $\mathcal{G}(n)$ contains $\Theta(n)$ nodes and its compliment contains $o(n)$ nodes.*

Proof: Suppose $\mathcal{G}(n)$ has a long-term average version age $\Theta(f(n))$. Then, each node which is in the compliment set of the typical set has $\omega(f(n))$ long-term average version age. Suppose the average version age of the compliment set is $\Theta(g(n))$ where $g(n) = \omega(f(n))$, and suppose this set has kn nodes, where $0 < k < 1$ is a constant. Similarly, let the average version age of nodes in the typical set be $\Theta(h(n))$, where $h(n) = O(f(n))$. Then, we write the long-term average version age of the entire network $v_{\mathcal{G}(n)}$ in terms of $h(n)$ and $g(n)$.

$$v_{\mathcal{G}(n)} = \frac{h(n)(n - kn) + g(n)kn}{n} \tag{1}$$

$$= h(n)(1-k) + kg(n) \tag{2}$$

$$= \Theta(g(n)). \tag{3}$$

This is in contradiction to the fact that the long-term average version age of the network is $\Theta(f(n))$. Hence, the compliment set has $o(n)$ nodes, and the result follows. ■

The following result is a simple extension of Lemma 5.

Corollary 1 *Suppose we have a varying-topology network with K topologies, $\mathcal{G}_k(n)$, $\forall k \in \mathcal{K}$. Then, the typical set of nodes with respect to $\mathcal{G}_1(n), \mathcal{G}_2(n), \ldots, \mathcal{G}_K(n)$ contains $\Theta(n)$ nodes and the remaining nodes are $o(n)$ in number.*

Next, we state and prove the following theorem, which characterizes the behavior of the version age of the typical set of nodes in a varying-topology network when the holding times are $O(f_1(n))$ or $\Omega(n \log n)$. We note that for $h(n)$ between $\Omega(f_1(n))$ and $o(n \log n)$, the long-term average version age is $O(f_K(n))$, which is a universally applicable upper bound.

Theorem 4 *Suppose we have a varying-topology network with K topologies. Let $\{\mathcal{G}_1(n)\}_{n=1}^{\infty}$, $\{\mathcal{G}_2(n)\}_{n=1}^{\infty}, \ldots, \{\mathcal{G}_K(n)\}_{n=1}^{\infty}$ be the K associated sequences of networks. Suppose $\mathcal{G}_k(n)$ has a long-term average version age $\Theta(f_k(n))$, $\forall k \in \mathcal{K}$. Further, assume $f_k(n) = \Omega(f_1(n))$, $\forall k \in \mathcal{K} \backslash \{1\}$, $f_k(n) = o(f_K(n))$, $\forall k \in \mathcal{K} \backslash \{K\}$ and $q_k = \Theta(1/h(n))$, $\forall k \in \mathcal{K}$. Then, we have the following:*

1. *If $h(n) = O(f_1(n))$, then the long-term average version age scaling for the typical set in the varying-topology network is $\Theta(f_1(n))$.*
2. *If $h(n) = \Omega(n \log n)$, then the long-term average version age scaling for the typical set in the varying-topology network is $\Theta(f_K(n))$.*

Unlike the analysis in Section 3 and Section 4, which consider the entire network, Theorem 4 studies only the typical set of nodes in the varying-topology network. This theorem provides general results for the varying-topology network under both fast switching and slow switching regimes, while sacrificing the understanding of the behavior of a few stale nodes in the network. In our running example, Theorem 4 states that if the holding times have mean smaller than order $\log n$, then the long-term average version age of the typical set of nodes in the varying-topology network scales as $\Theta(\log n)$. Further, if the holding times are of order $n \log n$ or larger, then the long-term version age scaling of the typical set of nodes in the varying-topology network is $\Theta(\sqrt{n})$.

Proof: The first item is exactly the case considered in Theorem 1. For the second item, we immediately obtain the upper bound by replacing $\mathcal{G}_1(n), \ldots, \mathcal{G}_{K-1}(n)$ with $\mathcal{G}_K(n)$. A lower bound can be obtained by lower bounding the version age of nodes during their time spent in $\mathcal{G}_1(n), \ldots, \mathcal{G}_{K-1}(n)$ with 0, by replacing each of these topologies with the network we used in the proof of Theorem 2. We call this new network the modified varying-topology network.

It was shown in [21] that if a network remains in a fixed topology for a sufficiently long time, the distribution of the version age at any node converges to the distribution of the first-passage percolation time from the source to that node. This is because after a sufficient amount of time, every node in the network receives some update from the source w.h.p. We now show that $\Theta(h(n))$ time is long enough for the convergence to happen in the modified varying-topology network due to the fact that the source sends an update directly to each node during this time period w.h.p. Specifically, we obtain an $o(1)$ upper bound for the probability that the source does not send an update to a node within $\Theta(n \log n)$ time. To this end, without loss of generality, let B_i be the random variable associated with the first time that the source sends an update to node i along edge (n_0, i). Let $B = \sup_{i \in \mathcal{V}} B_i$ be the first time every node receives an update from the source. It is known that $\mathbb{E}[B] = nH^{(1)}(n)$ and $\text{Var}[B] = n^2 H^{(2)}(n)$, where $H^{(k)}(n) = \sum_{j=1}^{n} \frac{1}{j^k}$. Then, we have the following,

$$\mathbb{P}\big[|B - \mathbb{E}[B]| > n \log n\big] \leq \frac{\text{Var}[B]}{(n \log n)^2} \leq \frac{n^2 \frac{\pi^2}{6}}{(n \log n)^2} = \frac{\frac{\pi^2}{6}}{\log^2 n} = o(1). \tag{4}$$

This leads us to conclude that every node receives an update directly from the source in $\Theta(n \log n)$ time w.h.p. Thus, the distribution of the version age of each node in the typical set of $\mathcal{G}_K(n)$ is the same as the first passage time after B. Since $B = O(h(n))$, the expectation of the version age during any interval spent in $\mathcal{G}_K(n)$ is $O(f_K(n))$ for each node in the typical set, and thus the version age is also $O(f_K(n))$ at any time during this interval w.h.p., with the proof being along similar lines as the proof of Lemma 1. Thus, the average version age whenever we are in $\mathcal{G}_K(n)$ is $\Theta(f_K(n))$. From Lemma 2, we know that we spend a constant fraction of time in each topology w.h.p. Further, not adding the version ages of a node while in $\mathcal{G}_1(n), \ldots, \mathcal{G}_{K-1}(n)$ scales the long-term average version age of the node by some constant $0 < c < 1$. This gives the lower bound of $\Theta(f_K(n))$, concluding the proof. ■

6 Exploring Different Switching Rates

In this section, we consider a varying-topology network which contains only two states in the CTMC, i.e., $K = 2$. Further, state 1 is the ring topology, and state 2 is the FC topology. We have considered this setting as a running example in previous sections. In this section, by specializing to this particular setting, we aim to perform a deeper analysis of this network.

We assume that the transition rates between the two states have different orders, i.e., $h_1(n) = o(h_2(n))$ or $h_1(n) = \omega(h_2(n))$, where $h_1(n)$ is the expected holding time in the ring topology, and $h_2(n)$ is the expected holding time in the FC topology. Our goal in this section is to characterize the average version age of each node in the network for the full spectrum of transition rates between the states of the CTMC. We note that the typical sets for both graph sequences include all the nodes in the gossip network. Furthermore, since we are explicitly defining the graph sequences $\{\mathcal{G}_1(n)\}_{n=1}^{\infty}$ and $\{\mathcal{G}_2(n)\}_{n=1}^{\infty}$, we choose to use simpler notation for this section by omitting the infinite sequences as it is understood in this context. First, we state the following lemma which describes the spread of updates in the ring topology. The proof is given in the Appendix.

Lemma 6 *Suppose a node in the network starts with version age $\Theta(f(n))$ at the moment of switching to the ring topology. Suppose we stay in the ring topology for $\Theta(g(n))$ time before* *exiting. Then:*

1. *If $g(n) = \Omega(\sqrt{n})$, then the average version age of the node in the holding time of the ring topology is $\Theta(\sqrt{n})$.*

2. *If $g(n) = o(\sqrt{n})$, then the average version age of the node in the holding time of the ring topology is $\Theta(f(n) + g(n))$.*

Intuitively, Lemma 6 states that if we stay in the ring topology for sufficient, i.e., $\Omega(\sqrt{n})$ time, the node will receive an update generated by the source since the start of the current holding time of the ring topology w.h.p. If we spend less, i.e., $o(\sqrt{n})$ time, in the ring topology, then the node will not receive such an update from the source w.h.p. Using this information, we explore different regimes for $h_1(n)$ and $h_2(n)$ and find how the average version age of each node evolves depending on $h_1(n)$ and $h_2(n)$. We first note that if we spend $\Omega(\sqrt{n})$ time in the ring topology, then the average version age for each node during this holding time is $\Theta(\sqrt{n})$, as shown in Lemma 4 and restated in Lemma 6. Moreover, if we spend a total of $\log n$ time in the FC topology, even if it is in disjoint intervals, the average version age of each node in the FC topology in these intervals is $\Theta(\log n)$, as shown in Theorem 1. We primarily use these results to prove subsequent results in this section.

Theorem 5 *Suppose $h_1(n) = \Omega(\sqrt{n})$.*

1. *If $h_2(n) = O(h_1(n))$, then the average version age of each node is $\Theta(\sqrt{n})$.*

2. *If $h_2(n) = \omega(h_1(n))$, then the average version age of each node is*

$$\Theta\left(\frac{h_1(n)\sqrt{n} + h_2(n)\log n}{h_1(n) + h_2(n)}\right). \tag{5}$$

Proof: If $h_2(n) = O(h_1(n))$, then we have two cases. If $h_2(n) = \Omega(\log n)$, then the average version age of each node becomes $\Theta(\log n)$, as discussed earlier. However, if $h_2(n) = o(\log n)$, then the updates sent in the FC topology reach $o(n)$ nodes, thus making the version age $\Theta(\sqrt{n})$, since each node started with $\Theta(\sqrt{n})$ version age at the start of the holding time of the FC topology. In both cases, the average version age during the entire time horizon becomes $\Theta(\sqrt{n})$, since we have $O(\sqrt{n})$ version age during the holding times of the FC topology and $\Theta(\sqrt{n})$ version age during the holding times of the ring topology. This proves item 1.

Next, we note that the nodes have $\Theta(\sqrt{n})$ version age during each holding time in the ring topology. Moreover, the average version age of each node in the holding time of the FC topology is $\Theta(\log n)$, if $h_2(n) = \Omega(\log n)$, which is the case for item 2. The statement holds due to the fact that we spend enough time in each topology for the distribution of version age of each node to converge to the true distribution of the long-term average version age. Thus, the average version age of each node in the varying-topology network is

$$\Theta\left(\frac{h_1(n)\sqrt{n} + h_2(n)\log n}{h_1(n) + h_2(n)}\right). \tag{6}$$

This proves item 2, thus proving the lemma. ■

Next, we study the case when $h_1(n) = o(\sqrt{n})$.

Theorem 6 *Suppose $h_1(n) = o(\sqrt{n})$.*

1. *If $h_2(n) = \omega(h_1(n))$, then the average version age of each node is*

$$\Theta\left(\frac{h_1^2(n) + h_2(n)\log n}{h_1(n) + h_2(n)}\right), \tag{7}$$

   w.h.p. Further, if $h_1(n) = O(\log n)$, then the average version age of each node is $\Theta(\log n)$ w.h.p.

2. *If $h_2(n) = O(h_1(n))$, then,*

   (a) If $h_2(n) = o(\log n)$, then the average version age scaling of each node is

$$O\left(\min\left(\sqrt{n}, \frac{h_1(n)\log n}{h_2(n)}\right)\right), \tag{8}$$

   w.h.p.

   (b) If $h_2(n) = \Omega(\log n)$, then the average version age scaling of each node is

$$\Theta\left(\frac{h_1^2(n) + h_2(n)\log n}{h_1(n) + h_2(n)}\right), \tag{9}$$

   w.h.p.

Proof: We first prove item 1. Let us consider the case where $h_1(n) = O(\log n)$. The universal lower bound for the average version age of a node in our model is $\Theta(\log n)$, as discussed previously. In order to obtain a logarithmic upper bound, we proceed similar to the proof of Theorem 1 by considering a modified varying-topology network consisting of the FC topology and the DC topology, and deduce that the upper bound is $\Theta(\log n)$ w.h.p. as well, leading to the result. Note that this result also satisfies (7).

Next, if $h_1(n) = \omega(\log n)$, then from Lemma 6, the average version age of each node during the holding time in the ring topology is $\Theta(\log n + h_1(n)) = \Theta(h_1(n))$ if the node starts with $\Theta(\log n)$ version age. Since we spend $\omega(\log n)$ time in the FC topology, each node maintains $\Theta(\log n)$ version age w.h.p., as shown in Lemma 1. Thus, we can conclude that the average version age of each node during the holding time of the ring topology is $\Theta(h_1(n))$ w.h.p., and it is $\Theta(\log n)$ w.h.p. during the holding time of the FC topology. Thus the average version age of each node over all time is

$$\Theta\left(\frac{h_1^2(n) + h_2(n)\log n}{h_1(n) + h_2(n)}\right), \tag{10}$$

w.h.p., proving item 1.

We now prove item 2. In part (a), if $h_2(n) = o(\log n)$, then in a similar way to the method discussed in Theorem 1, we need to stay in the FC topology for a total of $\Theta(\log n)$ time for every node to receive an update. Following along the lines of Lemma 2, a total of $\Theta(\frac{h_1(n)\log n}{h_2(n)})$ time would pass for the varying-topology network to have spent $\Theta(\log n)$ time in the FC topology. If $\frac{h_1(n)\log n}{h_2(n)} = o(\sqrt{n})$, then all the nodes receive the updates by $O(\frac{h_1(n)\log n}{h_2(n)})$ time. However, if $\frac{h_1(n)\log n}{h_2(n)} = \Omega(\sqrt{n})$, then the time spent in the ring topology would also accumulate and reach $\Theta(\sqrt{n})$ before every node receives an update through the FC topology. In this case, Lemma 6 takes effect and we achieve $\Theta(\sqrt{n})$ average version age. This gives us the upper bound for the average version age of each node in part (a).

Part (b) of item 2 follows in a similar way to item 1, since the average version age of each node is $\Theta(\log n)$ w.h.p. during the holding times of the FC topology, and $\Theta(h_1(n))$ w.h.p. in the holding times of the ring topology, due to Lemma 6. This concludes the proof. ■

Theorem 5 and Theorem 6 summarize the interplay of different transition rates on the version age of the two-state varying-topology network. Theorem 5 analyzes the regime where the nodes in the ring topology are able to reach saturation, and the average version age of the entire network is determined by whether the network is in the FC topology for long enough or not. On the other hand, Theorem 6 analyzes the transient regime, where the nodes are not able to reach saturation in the ring topology, and the transition rates determine how version age evolves in this case.

7 Mobility: A Special Case of Varying-Topology

In any mobile network, where nodes have arbitrary movement and are gossiping with close neighbors, the connections between the nodes change frequently. Depending on the dynamics of the network, these changes in topology could happen very quickly or more slowly. Thus, we observe that any mobile network can be modeled as a varying-topology network. Even though this is the case, the speed of topology changes between the states of the CTMC can be very fast and make the network challenging to analyze. For example, in the mobility model we consider in this section, topology changes can happen at $\Theta\left(\frac{1}{n}\right)$ rate. This regime is very hard to analyze, since there are many different topologies, and a particular node may have vastly different average version ages in each topology.

In a general varying-topology network, we cannot use the stochastic hybrid systems (SHS) methodology to find a set of recursive equations due to stability issues when we choose the version ages of the nodes as the variables. This is because, the edge rates change frequently, and achieving stability is not possible without a very slow CTMC. Moreover, there is no other variable with respect to which the edge rates are constant. That is why we did not use the SHS method in the previous sections. However, in this mobility model, if we apply the SHS methodology when we consider the version ages of positions as the variables, we see that the edge rates are constant with respect to the positions, and thus, SHS can be applied. Thus, in this section, we consider the mobility model, which is a special case of the varying-topology network, using the SHS methodology. Thus, in the process, we provide an additional method to calculate the average version ages of the nodes in this special case.

7.1 SHS Characterization of the Mobility Network

In this subsection, we use the SHS characterization following [12] and [29] and construct recursive equations to find v_S for any general set of positions in the network. We recall that the rates of all Poisson processes in the network are constant with respect to time and that all processes are memoryless. Hence, we need only one discrete state to represent the network. Thus, we define the discrete state as $\mathcal{Q} = \{0\}$ and do not use notation to represent it as a function input later. We define the continuous state of the system to be the version age of the positions 1 through n, as defined in Section 2. We have $\mathbf{X}(t) = [X_1(t), X_2(t), \ldots, X_n(t)]$. We see that all changes in the continuous state are piecewise constant and only change when a transition happens. Hence, we can write the stochastic differential equation as $\dot{\mathbf{X}}(t) = 0$.

Next, we define the transition/reset maps $\mathcal{L}$ of the SHS. We can define a transition in the network uniquely using three variables (i, j, k). Here, i is the position that initiates the action, j is the second position that participates, and k tells us whether the transition is related to gossiping or a node exchange. Let (i, j, info) denote the transitions which lead to information exchanges, corresponding to the gossiping processes where node i updates node j, and (i, j, node) as transitions that lead to node exchanges. The source related transitions are different from the regular gossip transitions. We have two types of source related transitions. The first type is when the source updates itself. We denote this as $(0, 0, \text{info})$. The second type of transition happens when the source sends an update to one of the positions. The transition that occurs when the source sends an update to position j is denoted by $(0, j, \text{info})$. Hence, the entire set of transitions can be represented by the set

$$\begin{aligned}\mathcal{L} =& \{(0,0,\text{info})\} \cup \{(0,j,\text{info}) : j \in \mathcal{V}\} \cup \{(i,j,\text{info}) : i,j \in \mathcal{V}\} \\ & \cup \{(i,j,\text{node}) : i,j \in \mathcal{V}, i < j\}. \end{aligned} \tag{11}$$

Here, we note that we have chosen the transitions denoting node exchanges such that we can uniquely identify each transition. In general, node exchanges are not initiated by nodes, and hence the transition is symmetric in both nodes participating in the exchange. Let the continuous state after a transition be $\phi_{i,j,k}(\mathbf{X(t)}, t) = [X_1'(t), X_2'(t), \ldots, X_n'(t)]$. Then,

$$X_l'(t) = \begin{cases} X_l(t) + 1, & i = 0, j = 0, k = \text{info}, \\ 0, & i = 0, j = l, k = \text{info}, \\ \min(X_l(t), X_i(t)), & i \in \mathcal{V}, j = l, k = \text{info}, \\ X_i(t), & i \in \mathcal{V}, l \in \mathcal{M}_i, k = \text{node}, \\ X_l(t), & \text{otherwise}. \end{cases} \tag{12}$$

The transition rates corresponding to the transition maps are,

$$\lambda_{i,j,k}(\mathbf{X}(t),t) = \begin{cases} \lambda_e, & i=0, j=0, k=\text{info}, \\ \lambda_{0,l}, & i=0, j\in\mathcal{V}, k=\text{info}, \\ \lambda_{i,j}, & i,j\in\mathcal{V}, k=\text{info}, \\ \lambda_{i,j}^{\mathcal{M}}, & i\in\mathcal{V}, j\in\mathcal{M}_i, k=\text{node}. \end{cases} \tag{13}$$

Now, we choose a test function $\psi_{\mathcal{S}}:\mathbb{R}\times[0,\infty)\to\mathbb{R}$ as $\psi_{\mathcal{S}}(\mathbf{X}(t),t)=X_{\mathcal{S}}(t)$, where $\mathcal{S}$ is a set of positions in the network that can never include the position of the source. Now, we write the extended generator equation as,

$$(L\psi_{\mathcal{S}})(\mathbf{X}(t)) = \sum_{i,j,k}(\psi_{\mathcal{S}}(\phi_{i,j,k}(\mathbf{X}(t)))-\psi_{\mathcal{S}}(\mathbf{X}(t)))\lambda_{i,j,k}. \tag{14}$$

Since $X_{\mathcal{S}}(t)$ is piece-wise constant, we get $\frac{\partial X_{\mathcal{S}}(t)}{\partial t}=0$ and arrive at the above equation.

For our test function, the values of $[\psi_{\mathcal{S}}(\phi_{i,j,k}(\mathbf{X}))]$ are,

$$\psi_{\mathcal{S}}(\phi_{i,j,k}(\mathbf{X}(t))) = \begin{cases} 0, & i=0, j\in\mathcal{S}, k=\text{info}, \\ X_{\mathcal{S}}(t)+1, & i=0, j=0, k=\text{info}, \\ X_{\mathcal{S}\cup\{i\}}(t), & i\in\mathcal{N}(\mathcal{S}), j\in\mathcal{S}, k=\text{info}, \\ X_{\mathcal{S}\cup\{i\}\setminus\{l\}}(t), & i\in\mathcal{M}_l, j=l, k=\text{node}, \\ X_{\mathcal{S}}(t), & \text{otherwise}. \end{cases} \tag{15}$$

We substitute this in (14), take expectation on both sides, choose $t\to\infty$, and use Dynkin's formula to obtain

$$\begin{aligned} 0 =& \lambda_e(v_{\mathcal{S}}+1-v_{\mathcal{S}}) + \sum_{i\in\mathcal{S}}\lambda_{0,j}(0-v_{\mathcal{S}}) + \sum_{i\in\mathcal{N}(\mathcal{S})}\lambda_i(\mathcal{S})(v_{\mathcal{S}\cup\{i\}}-v_{\mathcal{S}}) \\ &+ \sum_{i\in\mathcal{S}}\sum_{j\in\mathcal{M}_i\setminus\mathcal{S}}\lambda_{i,j}^{\mathcal{M}}(v_{\mathcal{S}\cup\{j\}\setminus\{i\}}-v_{\mathcal{S}}). \end{aligned} \tag{16}$$

We collect the terms of $v_{\mathcal{S}}$ on one side, rewrite $v_{\mathcal{S}\cup\{j\}\setminus\{i\}}$ as $v_{\mathcal{S}(i,j)}$, and $\mathcal{M}_i\setminus\mathcal{S}$ as $\mathcal{M}_i^{\mathcal{S}}$, and rearrange the equation to obtain,

$$v_{\mathcal{S}} = \frac{\lambda_e + \sum_{i\in\mathcal{N}(\mathcal{S})}\lambda_i(\mathcal{S})v_{\mathcal{S}\cup\{i\}} + \sum_{i\in\mathcal{S}}\sum_{j\in\mathcal{M}_i^{\mathcal{S}}}\lambda_{i,j}^{\mathcal{M}}v_{\mathcal{S}(i,j)}}{\lambda_0(\mathcal{S}) + \sum_{i\in\mathcal{N}(\mathcal{S})}\lambda_i(\mathcal{S}) + \sum_{i\in\mathcal{S}}\sum_{j\in\mathcal{M}_i^{\mathcal{S}}}\lambda_{i,j}^{\mathcal{M}}}. \tag{17}$$

This representation using positions gives rise to simultaneous recursive equations of sets of positions with the same size, i.e., same number of positions in the set, in terms of sets of positions that are one larger.

Remark 1 *The limiting average version age of a position in the network is the same as*

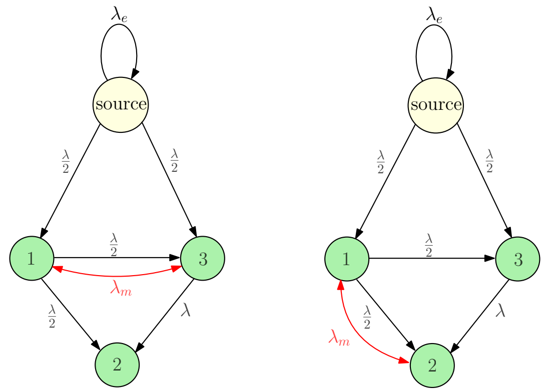

Figure 7: Toy example from [12], with mobility added in this paper. Left: nodes in positions 1 and 3 can exchange positions. Right: nodes in positions 1 and 2 can exchange positions.

the limiting average version age of a node in a symmetric network. Hence, we can use the recursive equations in (17) *to find the long-term average version age of a node in such a gossiping network.*

Remark 2 *The recursive equations in [12] are given by*

$$v_{\mathcal{S}} = \frac{\lambda_e + \sum_{i \in \mathcal{N}(\mathcal{S})} \lambda_i(\mathcal{S}) v_{\mathcal{S} \cup \{i\}}}{\lambda_0(\mathcal{S}) + \sum_{i \in \mathcal{N}(\mathcal{S})} \lambda_i(\mathcal{S})}. \tag{18}$$

Hence, when there is no mobility, the version age of a set of nodes S depends only on the version age of sets that can be constructed by adding a neighbor of S to it. On the other hand, when mobility is present, the version age of a set of positions depends not only on the version age of one larger sets but also on sets of the same size.

7.2 A Toy Example

We consider the toy example solved in [12] and apply mobility to it in two ways. This is shown in Fig. 7. We want to see how the version ages of positions are affected by mobility. In order to simplify calculations, we assume that the source sends updates as a combined rate λ, and the nodes in positions 1 and 3 are sending updates as a combined rate λ process. Moreover, the rate of node exchange between positions is λ_m in both examples.

We recall from [12] that the version age of nodes/positions in the network if there was no mobility are

$$v_1 = 2\frac{\lambda_e}{\lambda}, \qquad v_2 = 2\frac{\lambda_e}{\lambda}, \qquad v_3 = \frac{3}{2}\frac{\lambda_e}{\lambda}. \tag{19}$$

First, we solve the equations for the network where node exchanges can happen between positions 1 and 3. Using (17), we write the equations going from the largest sets to the smallest sets,

$$v_{\{1,2,3\}} = \frac{\lambda_e}{\lambda}, \tag{20}$$

$$v_{\{1,2\}} = \frac{\lambda_e + \lambda v_{\{1,2,3\}} + \lambda_m v_{\{2,3\}}}{\frac{\lambda}{2} + \lambda + \lambda_m} = \frac{2\lambda_e + \lambda_m v_{\{2,3\}}}{\frac{3}{2}\lambda + \lambda_m}, \tag{21}$$

$$v_{\{2,3\}} = \frac{\lambda_e + \lambda v_{\{1,2,3\}} + \lambda_m v_{\{1,2\}}}{\frac{\lambda}{2} + \lambda + \lambda_m} = \frac{2\lambda_e + \lambda_m v_{\{1,2\}}}{\frac{3}{2}\lambda + \lambda_m}, \tag{22}$$

$$v_{\{1,3\}} = \frac{\lambda_e}{\lambda}. \tag{23}$$

Using these, we find the version age of the three positions in the network as follows,

$$v_1 = \frac{\lambda_e}{\lambda} \frac{\lambda + \frac{5}{2}\lambda_m}{\frac{\lambda}{2} + \frac{3}{2}\lambda_m}, \tag{24}$$

$$v_3 = \frac{\lambda_e}{\lambda} \frac{\frac{3}{4}\lambda + \frac{5}{2}\lambda_m}{\frac{\lambda}{2} + \frac{3}{2}\lambda_m}, \tag{25}$$

$$v_2 = \frac{\lambda_e + \frac{\lambda}{2} v_{\{1,2\}} + \lambda_m v_{\{2,3\}}}{\frac{3\lambda}{2}} \tag{26}$$

$$= \frac{2}{3}\frac{\lambda_e}{\lambda} + \frac{1}{3}\left(\frac{2\lambda_e + \frac{2\lambda_e\lambda_m}{\frac{3}{2}\lambda + \lambda_m}}{\frac{3}{2}\lambda + \lambda_m - \frac{\lambda_m^2}{\frac{3}{2}\lambda+\lambda_m}}\right) + \frac{2}{3}\left(\frac{2\lambda_e + \frac{2\lambda_e\lambda_m}{\frac{3}{2}\lambda + \lambda_m}}{\frac{3}{2}\lambda + \lambda_m - \frac{\lambda_m^2}{\frac{3}{2}\lambda+\lambda_m}}\right) \tag{27}$$

$$= 2\frac{\lambda_e}{\lambda}. \tag{28}$$

In a similar way, we solve the equations for the network where node exchange can happen between positions 1 and 2. Using (17), we write the equations going from the largest sets to the smallest sets,

$$v_{\{1,2\}} = \frac{\lambda_e + \lambda v_{\{1,2,3\}}}{\frac{\lambda}{2} + \lambda} = \frac{4\lambda_e}{3\lambda}, \tag{29}$$

$$v_{\{2,3\}} = \frac{\lambda_e + \lambda v_{\{1,2,3\}} + \lambda_m v\{1,3\}}{\frac{\lambda}{2} + \lambda + \lambda_m} = \frac{2\lambda_e + \lambda_m v_{\{1,3\}}}{\frac{3\lambda}{2} + \lambda_m}, \tag{30}$$

$$v_{\{1,3\}} = \frac{\lambda_e + \lambda_m v_{\{2,3\}}}{\lambda + \lambda_m}. \tag{31}$$

Solving these simultaneous equations together, we get

$$v_{\{2,3\}} = \frac{\lambda_e}{\lambda} \frac{2\lambda + 3\lambda_m}{\frac{3}{2}\lambda + \frac{5}{2}\lambda_m}, \tag{32}$$

$$v_{\{1,3\}} = \frac{\lambda_e}{\lambda} \frac{\frac{3}{2}\lambda + 3\lambda_m}{\frac{3}{2}\lambda + \frac{5}{2}\lambda_m}, \tag{33}$$

and therefore,

$$v_2 = \frac{\lambda_e + \frac{\lambda}{2} v_{\{1,2\}} + \lambda v_{\{2,3\}} + \lambda_m v_1}{3\frac{\lambda}{2} + \lambda_m} \tag{34}$$

$$= \frac{\lambda_e}{\lambda} \frac{\frac{\lambda}{2} + \lambda_m}{\frac{3}{4}\lambda + 2\lambda_m} \left(\frac{5}{3} + \frac{2\lambda + 3\lambda_m}{\frac{3}{2}\lambda + \frac{5}{2}\lambda_m} + \frac{\lambda_m}{\frac{\lambda}{2} + \lambda_m} \right), \tag{35}$$

$$v_1 = \frac{\lambda_e + \lambda_m v_2}{\frac{\lambda}{2} + \lambda_m} \tag{36}$$

$$= \frac{\lambda_e}{\frac{\lambda}{2} + \lambda_m} + \frac{\lambda_e}{\lambda} \frac{\lambda_m}{\frac{3}{4}\lambda + 2\lambda_m} \left(\frac{5}{3} + \frac{2\lambda + 3\lambda_m}{\frac{3}{2}\lambda + \frac{5}{2}\lambda_m} + \frac{\lambda_m}{\frac{\lambda}{2} + \lambda_m} \right), \tag{37}$$

$$v_3 = \frac{\lambda_e + \frac{\lambda}{2} v_{\{1,3\}}}{\lambda} \tag{38}$$

$$= \frac{\lambda_e}{\lambda} \frac{\frac{9}{2}\lambda + 8\lambda_m}{3\lambda + 5\lambda_m}. \tag{39}$$

We first observe that we get the results in (19) if we substitute $\lambda_m = 0$ for both toy examples. Next, we observe that, if we choose $\lambda_m \to \infty$, we obtain, for the case where 1 and 3 can exchange positions,

$$v_1 = \frac{5}{3} \frac{\lambda_e}{\lambda}, \qquad v_2 = 2\frac{\lambda_e}{\lambda}, \qquad v_3 = \frac{5}{3} \frac{\lambda_e}{\lambda}, \tag{40}$$

and for the case where 1 and 2 can exchange positions,

$$v_1 = \frac{29}{15} \frac{\lambda_e}{\lambda}, \qquad v_2 = \frac{29}{15} \frac{\lambda_e}{\lambda}, \qquad v_3 = \frac{8}{5}\lambda. \tag{41}$$

The first thing we notice from observing the increase in version age v_3 in (40) and (41) when compared to the no mobility case in (19) is that mobility does not improve the version age of every position in the network. Further, when there is no mobility, the average version age in (19) is $\frac{11}{6}\frac{\lambda_e}{\lambda} = 1.833\frac{\lambda_e}{\lambda}$. When there is mobility where 1 and 3 can exchange positions, the average version age in (40) is $\frac{16}{9}\frac{\lambda_e}{\lambda} = 1.778\frac{\lambda_e}{\lambda}$; and when 1 and 2 can exchange positions, the average version age in (41) is $\frac{82}{45}\frac{\lambda_e}{\lambda} = 1.822\frac{\lambda_e}{\lambda}$. Hence, we observe that the average version age has decreased in both cases of mobility when compared to the network without mobility in this example. That is, mobility improved the freshness of the network in this example. We will analyze this further using simulations in Section 9.

7.3 Fully-Connected Network with a Single Disconnected Node

In this subsection, we apply the recursive equations found in (17) to find an upper bound for the version age of the gossiping network described in Fig. 8. The gossiping nodes are divided into two parts, one part contains a FC network of $n-1$ nodes, which we call $\mathcal{V}_{FC}$, and the other part contains one node, which we call $\mathcal{V}_s$. We label the positions in $\mathcal{V}_{FC}$ as

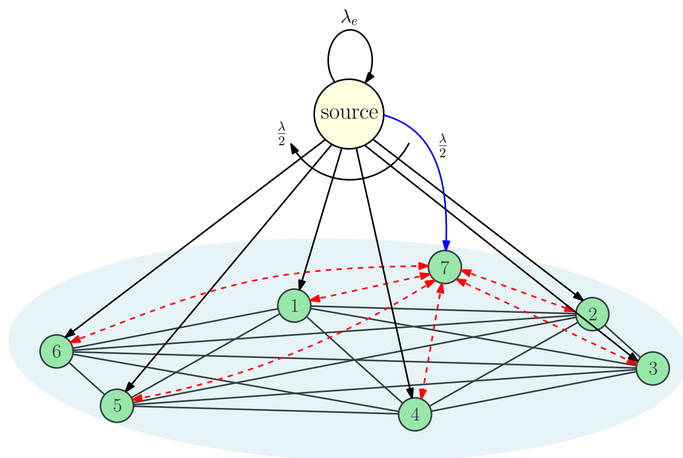

Figure 8: A gossiping network where $\mathcal{V}_{FC}$ is a fully connected network of 6 nodes $\{1,2,3,4,5,6\}$ and $\mathcal{V}_s$, represented by position 7, is a single node. Both parts receive updates as a rate $\frac{\lambda}{2}$ Poisson process from the source. Further, every node in $\mathcal{V}_{FC}$ can exchange positions with the node in $\mathcal{V}_s$, as shown by the red dashed arrows.

$\{1,\dots,n-1\}$ in order and the single node in $\mathcal{V}_s$ as n. Both parts are served with the same rate from the source, i.e., the source sends updates to each part as a rate $\frac{\lambda}{2}$ Poisson process. Further, the position of $\mathcal{V}_s$ can access and be accessed by all nodes in $\mathcal{V}_{FC}$ with rate $\lambda_m=\lambda$ for each node exchange. Hence, $\mathcal{M}_j=\{n\}$, $j\in\{1,2,\dots,n-1\}$, and $\mathcal{M}_n=\{1,2,\dots,n-1\}$. Moreover, $\lambda_{i,n}^{\mathcal{M}}=\lambda_{n,i}^{\mathcal{M}}=\lambda$, $i\in\{1,2,\dots,n-1\}$. We find an upper bound for the average version age of a single node in the network in this section. We compare it with the average version age in the same network when there is no mobility in Section 9.

Due to the symmetry of the network, each position in $\mathcal{V}_{FC}$ of the network has the same version age. Moreover, each set of the same size in $\mathcal{V}_{FC}$ has the same version age. It is easy to see that the version age of the single position in $\mathcal{V}_s$ is lower than the positions in $\mathcal{V}_{FC}$, since it receives updates directly from the source much more frequently. Next, due to the symmetry in the network, a set consisting of the single position in $\mathcal{V}_s$ and some nodes in $\mathcal{V}_{FC}$ will have the same average version age as long as the same number of positions is chosen in $\mathcal{V}_{FC}$. Hence, if we are given that the size of the set of positions is a, then there are two types of sets possible. The first type of set is the one which includes a positions from $\mathcal{V}_{FC}$. The second type of set includes $a-1$ positions in $\mathcal{V}_{FC}$ and the single position in $\mathcal{V}_s$. Let the representation of the set that has α positions in $\mathcal{V}_{FC}$ and β positions in $\mathcal{V}_s$ be (α,β). Then, we can write down the recursive equations for the two types of sets that can be created with a positions as

$$v_{(a-1,1)}=\frac{\lambda_e+\frac{(a-1)(n-a-2)\lambda}{n-2}v_{(a,1)}+(n-a-2)\lambda v_{(a,0)}}{\frac{(a-1)\lambda}{2(n-1)}+\frac{\lambda}{2}+\frac{(a-1)(n-a-2)\lambda}{n-2}+(n-a-2)\lambda}, \tag{42}$$

$$v_{(a,0)}=\frac{\lambda_e+\frac{a(n-a-1)\lambda}{n-2}v_{(a+1,0)}+a\lambda v_{(a-1,1)}}{\frac{a\lambda}{2(n-1)}+\frac{a(n-a-1)\lambda}{n-2}+a\lambda}. \tag{43}$$

We can see that $v_{(a-1,1)} < v_{(a,0)}$, due to the higher update rate for the position in $\mathcal{V}_s$. Hence, at each stage, we can write $v_{(a,0)}$ in terms of $v_{(a,1)}$ and $v_{(a+1,0)}$. Then, we can obtain an upper bound for the recursive equation by replacing $v_{(a,1)}$ with $v_{(a+1,0)}$. In order to simplify notation, we define

$$c^{(a)} = \frac{(a-1)\lambda}{2(n-1)} + \frac{\lambda}{2} + \frac{(a-1)(n-a-2)\lambda}{n-2} + (n-a-2)\lambda, \tag{44}$$

$$d^{(a)} = \frac{a\lambda}{2(n-1)} + \frac{a(n-a-1)\lambda}{n-2} + a\lambda. \tag{45}$$

Then, we substitute the equation for $v_{(a-1,1)}$ in (43) as

$$v_{(a,0)} = \left(1 - \frac{(n-a-2)a\lambda^2}{c^{(a)}d^{(a)}}\right)^{-1} \left(\frac{\lambda_e}{d^{(a)}}\left[1 + \frac{a\lambda}{c^{(a)}}\right] + \frac{\frac{a(n-a-1)\lambda}{n-2}v_{(a+1,0)}}{d^{(a)}} + \frac{\frac{a(a-1)(n-a-2)\lambda^2}{n-2}v_{(a,1)}}{c^{(a)}d^{(a)}}\right) \tag{46}$$

Now, we replace $v_{(a,1)}$ with $v_{(a+1,0)}$ and write an upper bound and simplify to get,

$$v_{(a,0)} \leq \frac{1}{z^{(a)}}(\lambda_e x^{(a)} + y^{(a)} v_{(a+1,0)}), \tag{47}$$

where $x^{(a)} = c^{(a)} + a\lambda$, $y^{(a)} = \frac{a(n-a-1)\lambda c^{(a)}}{n-2} + \frac{a(a-1)(n-a-2)\lambda^2}{n-2}$ and $z^{(a)} = c^{(a)}d^{(a)} - (n-a-2)a\lambda^2$. Then, we write out the recursive upper bound and obtain,

$$v_{(1,0)} \leq \lambda_e \sum_{i=1}^{n-2} \frac{x^{(i)}}{z^{(i)}} \prod_{j=1}^{i-1} \frac{y^{(j)}}{z^{(j)}} + \prod_{j=1}^{n-2} \frac{y^{(i)}}{z^{(i)}} v_{(n-1,0)}, \tag{48}$$

where $v_{(n-1,0)}$ is a constant, and depends on $v_{\mathcal{V}}$ and $v_{(n-2,1)}$. Using the recursive equations, we can write the value of $v_{(n-1,0)}$,

$$v_{(n-1,0)} = \left(1 - \frac{(n-1)\lambda^2}{\left(\frac{3\lambda}{2} + \frac{\lambda(n-2)}{2(n-1)} + \lambda\right)\left(\frac{\lambda}{2} + (n-1)\lambda\right)}\right)^{-1} \left(\frac{\lambda_e}{\frac{\lambda}{2} + (n-1)\lambda}\left[1 + \frac{2(n-1)\lambda}{\frac{3\lambda}{2} + \frac{\lambda}{2}\frac{n-2}{n-1} + \lambda}\right]\right). \tag{49}$$

Now, we want to simplify (48) in order to find a simple upper bound. First, we simplify $\frac{y^{(i)}}{z^{(i)}}$. In order to simplify notation, we define,

$$n_1 = \frac{i(i-1)(n-i-1)}{2(n-1)(n-2)}, \tag{50}$$

$$n_2 = \frac{i(n-i-1)(2n-2i-1)}{2(n-2)}, \tag{51}$$

$$n_3 = \frac{(i-1)i(n-i-2)(n-i-1)}{(n-2)^2}, \tag{52}$$

$$d_1 = \frac{i(4n-7)}{4(n-1)}, \tag{53}$$
$$d_2 = \frac{i(n-i-2)(2n-3)}{2(n-2)}, \tag{54}$$
$$d_3 = \frac{(i-1)i(2n-2i-3)}{2(n-1)(n-2)}, \tag{55}$$
$$d_4 = \frac{(i-1)i(n-i-2)(n-i-1)}{(n-2)^2}. \tag{56}$$

Then for $i < n-2$, we have,

$$\begin{aligned}
\frac{y^{(i)}}{z^{(i)}} &\approx \frac{n_1+n_2+n_3}{d_1+d_2+d_3+d_4} && (57)\\
&\le \frac{n_1+n_2+n_3}{\frac{i(n-i-2)(2n-3)}{2(n-2)} + \frac{(i-1)i(n-i-2)(n-i-1)}{(n-2)^2}} && (58)\\
&\approx \frac{\frac{(n-i-1)(2n-2i-1)}{2(n-2)} + \frac{(i-1)(n-i-2)(n-i-1)}{(n-2)^2}}{\frac{(n-i-2)(2n-3)}{2(n-2)} + \frac{(i-1)(n-i-2)(n-i-1)}{(n-2)^2}} && (59)\\
&= \frac{1 + \frac{(2n-2i-1)(n-2)}{2(i-1)(n-i-2)}}{1 + \frac{(2n-3)(n-2)}{2(i-1)(n-i-1)}} && (60)\\
&\approx \frac{1 + \frac{n-2}{i-1}}{1 + \frac{(n-2)^2}{(i-1)(n-i-1)}} && (61)\\
&= \frac{1 + \frac{n-2}{i-1}}{1 + \frac{n-2}{i-1} + \frac{n-2}{n-i-1}} && (62)\\
&= \frac{1}{1 + \frac{(n-2)(i-1)}{(n-i-1)(n+i-3)}} && (63)\\
&\le 1. && (64)
\end{aligned}$$

Also, clearly $\frac{y^{(n-2)}}{z^{(n-2)}} \le 1$. From this, we can conclude that $\prod_{j=1}^{i-1} \frac{y^{(j)}}{z^{(j)}} \le 1, \forall i$. (48) now becomes,

$$v_{(1,0)} \le \lambda_e \sum_{i=1}^{n-2} \frac{x^{(i)}}{z^{(i)}} + \prod_{j=1}^{n-2} \frac{y^{(i)}}{z^{(i)}} v_{(n-1,0)}. \tag{65}$$

We now simplify $\frac{x^{(i)}}{z^{(i)}}, i < n-2$.

$$\begin{aligned}
\frac{x^{(i)}}{z^{(i)}} &= \frac{n-2+\frac{i-1}{2(n-1)}+\frac{1}{2}+\frac{(i-1)(n-i-2)}{n-2}}{d_1+d_2+d_3+d_4} && (66)\\
&\le \frac{n-2+1+\frac{(i-1)(n-i-2)}{n-2}}{d_2+d_3+d_4} && (67)
\end{aligned}$$

$$= \frac{1}{\lambda} \frac{1+\frac{1}{n-2}+\frac{(i-1)(n-i-2)}{(n-2)^2}}{\frac{i(2n-3)(n-i-2)}{2(n-2)^2}+\frac{(i-1)i(n-i-1)(n-i-2)}{(n-2)^3}} \tag{68}$$

$$\leq \frac{1}{\lambda} \frac{\frac{5}{4}}{\frac{i(n-i-2)}{n-2}+\frac{(i-1)i(n-i-1)(n-i-2)}{(n-2)^3}} \tag{69}$$

$$\leq \frac{1}{\lambda} \frac{\frac{5}{4}}{\frac{i(n-i-2)}{n-2}} \tag{70}$$

$$= \frac{5}{4}\frac{1}{\lambda}\left(\frac{1}{i}+\frac{1}{n-i-2}\right). \tag{71}$$

We go from (68) to (69) due to the fact that $\frac{1}{n-2} \approx 0$ and $\frac{(i-1)(n-i-2)}{(n-2)^2} \leq \frac{1}{4}$ when $i < n-2$. Further, $\frac{x^{(n-2)}}{z^{(n-2)}} \leq \frac{1}{\lambda} \leq \frac{5}{4\lambda}$. Hence, we can conclude that,

$$v_{(1,0)} \leq \frac{5}{4}\frac{\lambda_e}{\lambda}\sum_{i=1}^{n-3}\left(\frac{1}{i}+\frac{1}{n-i-2}\right)+C \tag{72}$$

$$\leq \frac{5}{2}\frac{\lambda_e}{\lambda}\log n + C', \tag{73}$$

where $C' = v_{(n-1,0)} + 2\gamma + \frac{5\lambda_e}{4\lambda}$ where $\gamma \approx 0.577$ is the Euler-Mascheroni constant. Hence, we can conclude that $v_{(1,0)} = O(\log n)$. Since all but one node have the same version age, we can conclude that the version age of a single node in such a network scales as $O(\log n)$.

7.4 Disconnected Gossip Network with Full-Mobility

In this subsection, we consider a gossiping network where a node in any position can communicate with only one node in an adjacent position, i.e., a 1-regular graph or matching, as shown in Fig. 9. For notational ease, we assume there are an even number of nodes. A node in a particular position in the network can also exchange positions with any other node in the network. Therefore, $\mathcal{M}_i = \mathcal{V}\backslash i$, $\forall i \in \mathcal{V}$. The rate of all node exchanges is the same and is set as λ_m. Hence, $\lambda_{i,j}^{\mathcal{M}} = \lambda_{j,i}^{\mathcal{M}} = \lambda_m$, $i, j \in \mathcal{V}, i \neq j$. Any set of positions in this network can be uniquely defined by the number of adjacent node pairs and single nodes of an adjacent pair it contains, as described in Fig. 9. Hence, each set of positions with the same number of adjacent node pairs and single nodes in it has the same version age. Hence, each set of positions can be represented as (α, β) where α is the number of adjacent node pairs and β is the number of single nodes.

In order to find a recursion to find the upper bound for the version age of a single node in the network, we will obtain a recursive upper bound on the set $(a, 0)$ in terms of the set $(a+1, 0)$ using the fact that given $2a$ nodes, the set with the highest version age is the set with a node pairs, and given $2a+1$ nodes, the set with the highest version age is $(a, 1)$. These statements are true due to the fact that if we consider sets of nodes of the same size, then the set of positions with the highest version age is the one which has the least amount of gossip coming to it from nodes outside the set. The impact of mobility is the same for

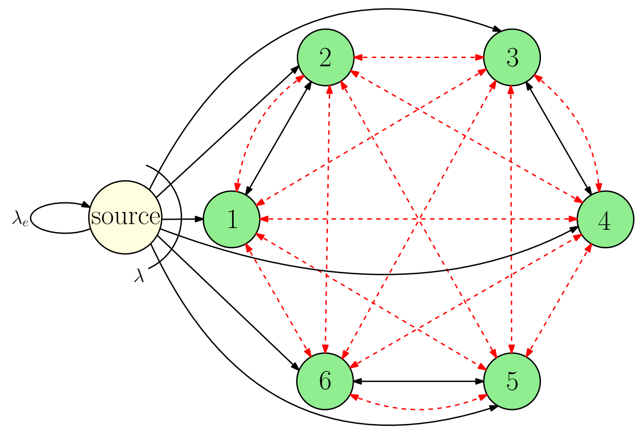

Figure 9: A disconnected gossiping network which has 6 nodes, and nodes in two adjacent positions gossip with each other, but not with nodes in any other position, i.e., a node in each position can gossip with a node only in one other position. This is represented by solid black double-sided arrows. Further, a node in any position can exchange positions with nodes in any other position in the network. This is represented by red dashed arrows. $\{1,2\}$, $\{3,4\}$ and $\{5,6\}$ are adjacent node pairs. $\{1\}$ is a single node of the node pair $\{1,2\}$.

any two sets of the same size due to the full-mobility of the network and the fact that due to the symmetry of the network, each node in the network has the same version age. We then start the recursion by writing the version age of a single node in the network in terms of a node pair,

$$v_{(0,1)} = \frac{\lambda_e + \lambda v_{(1,0)}}{\frac{\lambda}{n} + \lambda}, \tag{74}$$

and then write $(a,0)$ in terms of $(a+1,0)$ to find a recursive upper bound. We write $(a,0)$ in terms of $(a+1,0)$ for $a=1$ to $a=\frac{n}{2}-2$ as follows,

$$\begin{aligned}
v_{(a,0)} &= \frac{\lambda_e + 2a(n-2a)\lambda_m v_{(a-1,2)}}{\frac{2a\lambda}{n} + 2a(n-2a)\lambda_m}, && (75)\\
v_{(a-1,2)} &= \frac{\lambda_e + 2\lambda v_{(a,1)} + 2\lambda_m v_{(a,0)} + h^{(a)} v_{(a-2,4)}}{\frac{2a\lambda}{n} + 2\lambda + 2\lambda_m + h^{(a)}} && (76)\\
&\leq \frac{\lambda_e + 2\lambda v_{(a,1)} + (2\lambda_m + h^{(a)})\lambda_m) v_{(a-1,2)}}{\frac{2a\lambda}{n} + 2\lambda + 2\lambda_m + h^{(a)}}, && (77)\\
v_{(a,1)} &= \frac{\lambda_e + \lambda v_{(a+1,0)} + g^{(a)} v_{(a-1,3)}}{\frac{(2a+1)\lambda}{n} + \lambda + g^{(a)}} && (78)\\
&\leq \frac{\lambda_e + \lambda v_{(a+1,0)} + g^{(a)} v_{(a,1)}}{\frac{(2a+1)\lambda}{n} + \lambda + g^{(a)}}, && (79)
\end{aligned}$$

where $h^{(a)} = 2(a-1)(n-2(a+1))\lambda_m$ and $g^{(a)} = 2a(n-2(a+1))\lambda_m$. Now, to simplify notation, we define,

$$b^{(a)} = \frac{2a\lambda}{n} + 2a(n-2a)\lambda_m, \tag{80}$$

$$c^{(a)} = \frac{(2a+1)\lambda}{n} + \lambda + 2a(n-2(a+1))\lambda_m, \tag{81}$$

$$d^{(a)} = \frac{2a\lambda}{n} + 2\lambda + 2\lambda_m + 2(a-1)(n-2(a+1))\lambda_m, \tag{82}$$

$$f^{(a)} = 1 - \frac{2a(n-2(a+1))\lambda_m}{c^{(a)}}, \tag{83}$$

$$g^{(a)} = 1 - \frac{2(a-1)(n-2(a+1))\lambda_m}{d^{(a)}}. \tag{84}$$

Then, we have,

$$v_{(a,0)} \leq \lambda_e x^{(a)} + y^{(a)} v_{(a+1,0)}, \tag{85}$$

where

$$x^{(a)} = \frac{1}{b^{(a)}} + \frac{2a(n-2a)\lambda_m}{b^{(a)}d^{(a)}g^{(a)}} + \frac{4a(n-2a)\lambda_m\lambda}{b^{(a)}c^{(a)}d^{(a)}f^{(a)}g^{(a)}}, \tag{86}$$

$$y^{(a)} = \frac{4a(n-2a)\lambda_m\lambda^2}{b^{(a)}c^{(a)}d^{(a)}f^{(a)}g^{(a)}}. \tag{87}$$

Suppose, $\lambda_m = \frac{\lambda}{\alpha(n)}$, then, we use the fact that,

$$\frac{\frac{a(n-2a)}{\alpha(n)}}{\frac{a}{n} + 1 + \frac{1}{\alpha(n)}} + \frac{\frac{2a(n-2a)}{\alpha(n)}}{(1+\frac{2a+1}{n})(\frac{a}{n} + 1 + \frac{1}{\alpha(n)})} \leq \frac{\frac{3a(n-2a)}{\alpha(n)}}{1 + \frac{a}{n} + \frac{1}{\alpha(n)}}, \tag{88}$$

to get,

$$x^{(a)} \leq \frac{1}{\frac{2a\lambda}{n} + 2a(n-2a)\frac{\lambda}{\alpha(n)}} \left(1 + \frac{\frac{3a(n-2a)}{\alpha(n)}}{1 + \frac{a}{n} + \frac{1}{\alpha(n)}}\right), \tag{89}$$

$$y^{(a)} = \frac{1}{(1 + \frac{2\alpha(n)}{n(n-2a)})(1 + \frac{2a+1}{n})(1 + \frac{1}{\alpha(n)} + \frac{a}{n})}. \tag{90}$$

Hence, we can say that,

$$v_{(1,0)} \leq \lambda_e \sum_{i=1}^{\frac{n}{2}-2} x^{(i)} \prod_{j=1}^{i-1} y^{(j)} + v_{(\frac{n}{2}-1,0)} \prod_{j=1}^{\frac{n}{2}-1} y^{(j)}. \tag{91}$$

Define $\beta_i = \prod_{j=1}^{i-1} y^{(j)}$. Now, we write the value for $v_{(\frac{n}{2}-1,0)}$,

$$v_{(\frac{n}{2}-1,0)} = \left(1 - \frac{4\lambda_m^2}{(\frac{n-2}{n}\lambda + 2\lambda_m)(\frac{n-2}{n}\lambda + 2\lambda_m + 2\lambda)}\right)^{-1}$$

$$\times\left(\frac{\lambda_e}{\frac{n-2}{n}\lambda+2\lambda_m}+\frac{6\lambda_e\lambda_m}{(\frac{n-2}{n}\lambda+2\lambda_m)(\frac{n-2}{n}\lambda+2\lambda_m+2\lambda)}\right). \tag{92}$$

Then, the version age of a single node can be upper bounded as,

$$v_{(0,1)} \leq \frac{\lambda_e}{\lambda}+\lambda_e\sum_{i=1}^{\frac{n}{2}-2}x^{(i)}\prod_{j=1}^{i-1}y^{(j)}+v_{(\frac{n}{2}-1,0)}\prod_{j=1}^{\frac{n}{2}-1}y^{(j)} \tag{93}$$

$$\leq \frac{\lambda_e}{\lambda}+\frac{\lambda_e}{\lambda}\sum_{i=1}^{\frac{n}{2}-2}x^{(i)}\beta_i+C, \tag{94}$$

where C $= v_{(\frac{n}{2}-1,0)}$, since $y^{(j)} \leq 1$, $\forall j$.

We now find the upper bound scaling for two regions for $f(n)$. For each region, we bound $x^{(a)}$ and $y^{(a)}$ appropriately and then find the sum in (94). The first region is when $f(n)=k\in\mathbb{R}$. Starting with (89), we can upper bound $x^{(a)}$ as follows,

$$x^{(a)} \leq \frac{\frac{3a(n-2a)}{\alpha(n)}}{\frac{2a\lambda}{n}+2a(n-2a)\frac{\lambda}{\alpha(n)}} \tag{95}$$

$$\leq \frac{3}{2\lambda}, \tag{96}$$

and starting from (90), $y^{(a)}$ can be upper bounded as,

$$y^{(a)} \leq \frac{1}{1+\frac{1}{k}+\frac{a}{n}} \tag{97}$$

$$\leq \frac{1}{1+\frac{1}{k}} \tag{98}$$

$$= \frac{k}{k+1}. \tag{99}$$

We substitute this back in (94),

$$v_{(0,1)} \leq \frac{\lambda_e}{\lambda}+\frac{3\lambda_e}{2\lambda}\sum_{i=1}^{\frac{n}{2}-2}\prod_{j=1}^{i-1}\frac{k}{k+1}+C \tag{100}$$

$$= \frac{\lambda_e}{\lambda}+\frac{3\lambda_e}{2\lambda}\sum_{i=1}^{\frac{n}{2}-2}\left(\frac{k}{k+1}\right)^{i-1}+C \tag{101}$$

$$\leq \frac{\lambda_e}{\lambda}+\frac{3\lambda_e}{2\lambda}k+C', \tag{102}$$

where $C' = C+\frac{3\lambda_e}{2\lambda}$. Hence, the version age is upper bounded by a constant in this case.

In the second region, we choose $\alpha(n)$ such that $f(n)=O(n)$ and $1=o(\alpha(n))$. First, we note that we use the same bound for $x^{(a)}$ as in (96). Starting at (90), we bound $y^{(a)}$ as,

$$y^{(a)} \leq \frac{1}{1+\frac{1}{\alpha(n)}} \frac{1}{1+\frac{2a}{n}}. \tag{103}$$

Then, we can find β_i by taking the negative logarithm on both sides,

$$\begin{aligned} -\log \beta_i =& \sum_{j=1}^{i-1} \log\left(1+\frac{1}{\alpha(n)}\right) + \log\left(1+\frac{2a}{n}\right) && (104)\\ =& \sum_{j=1}^{i-1} \frac{1}{\alpha(n)} + \frac{2a}{n} && (105)\\ =& \frac{i-1}{\alpha(n)} + \frac{i(i-1)}{n} && (106)\\ \approx& \frac{i}{\alpha(n)} + \frac{i^2}{n}. && (107) \end{aligned}$$

Hence, we get,

$$\beta_i \approx e^{-\left(\frac{i}{\alpha(n)}+\frac{i^2}{n}\right)}. \tag{108}$$

We now further upper bound β_i in two different ways. First, we have,

$$\beta_i \leq e^{-\frac{i}{\alpha(n)}}. \tag{109}$$

Substituting this back in (94), we get,

$$\begin{aligned} v_{(0,1)} \leq& \frac{\lambda_e}{\lambda} + \frac{3\lambda_e}{2\lambda} \sum_{i=1}^{\frac{n}{2}-2} e^{-\frac{i}{\alpha(n)}} + C && (110)\\ =& \frac{\lambda_e}{\lambda} + \frac{3\lambda_e}{2\lambda} \alpha(n) \sum_{i=1}^{\frac{n}{2}-2} \frac{1}{f(n)} e^{-\frac{i}{\alpha(n)}} + C && (111)\\ =& \frac{\lambda_e}{\lambda} + \frac{3\lambda_e}{2\lambda} \alpha(n) \int_0^\infty e^{-x} dx + C && (112)\\ =& \frac{\lambda_e}{\lambda} + \frac{3\lambda_e}{2\lambda} \alpha(n) + C. && (113) \end{aligned}$$

Next, we bound (108) in the second way,

$$\beta_i \leq e^{-\frac{i^2}{n}}. \tag{114}$$

Substituting this back in (94), we get,

$$v_{(0,1)} \leq \frac{\lambda_e}{\lambda} + \frac{3\lambda_e}{2\lambda} \sum_{i=1}^{\frac{n}{2}-2} e^{-\frac{i^2}{n}} + C \tag{115}$$

$$= \frac{\lambda_e}{\lambda} + \frac{3\lambda_e}{2\lambda} \sqrt{n} \sum_{i=1}^{\frac{n}{2}-2} \frac{1}{\sqrt{n}} e^{-\frac{i^2}{n}} + C \tag{116}$$

$$= \frac{\lambda_e}{\lambda} + \frac{3\lambda_e}{2\lambda} \sqrt{n} \int_0^{\infty} e^{-x^2} dx + C \tag{117}$$

$$= \frac{\lambda_e}{\lambda} + \frac{3\lambda_e}{2\lambda} \frac{\sqrt{\pi}}{2} \sqrt{n} + C. \tag{118}$$

Hence, we combine (113) and (118) to get the following upper bound for the version age of a single node for the second region of $f(n)$,

$$v_{(0,1)} = O\Big(\min(\alpha(n), \sqrt{n})\Big). \tag{119}$$

We can now combine the results in (102) and (119), and have the following result.

$$v_{(0,1)} = O(\min(\alpha(n), \sqrt{n})). \tag{120}$$

Hence, if $\alpha(n) = \Omega(\sqrt{n})$, then we see that the version age of a single node scales as $O(\sqrt{n})$. If $\alpha(n) = o(\sqrt{n})$, then the version age scaling is $O(\alpha(n))$.

8 Discussion and Remarks

8.1 Remarks on General Varying-Topology Networks

Remark 3 *Lemma 1 tells us that the version age of each node in the network does not exceed the long-term average version age w.h.p. in any gossiping network with a single topology. This means that the distribution of the version age of information for any node in a gossiping network has a light tail. We conjecture that if the long-term average version age of a node is $\Theta(f(n))$, then the version age at any large time is also $\Theta(f(n))$ w.h.p.*

Remark 4 *In part 2 of Theorem 4, we require $h(n) = \omega(f(n))$ for the proof. However, we conjecture that the theorem is true for the condition $h(n) = \Omega(f(n))$.*

Remark 5 *Section 3 shows that if the switching rate is faster than the average version age of the freshest topology, then the varying-topology network is also able to maintain the same version age as the freshest topology. For example, if the freshest network is the fully-connected network, the varying-topology network will also be able to maintain $\Theta(\log n)$ version age if the holding times in each state are $O(\log n)$. This special case was explored in [30] where* *a different result than Lemma 1 was used to show that the packets reach every node in the varying-topology network in $\Theta(\log n)$ time.*

Remark 6 *In this paper, we have addressed varying-topology networks where the CTMC switching rates are at least constant. However, if the switching rates are $o(1)$ and of the same order with respect to n, then the results in Section 3 carry over. This is because if all the switching rates are of the same order with respect to n, then the results in all the lemmas and theorems in the section are unaffected.*

Remark 7 *In Section 4, we analyzed the impact of $\Omega(\sqrt{n})$ switching rates on varying-topology networks where one of the topologies is the ring topology and the other topologies have lower average version age than the ring topology. However, if we restrict ourselves to having the switching rates to be exactly $\Theta(\sqrt{n})$, then we can also include topologies which have worse average version age than the ring topology while maintaining $\Theta(\sqrt{n})$ version age. This is because we can write the proof along the lines of the proof of Theorem 1. We can obtain the lower bound easily by replacing topologies with worse version age than the ring topology with the ring topology. We can further obtain an upper bound by replacing all such topologies by the disconnected topology and observing that we do not increase the version age of any node by $\omega(\sqrt{n})$ w.h.p. This is because we do not spend more than $\Theta(\sqrt{n})$ time in the disconnected topology due to Lemma 2.*

8.2 Remarks on the Networks with Mobility

Remark 8 *We have found a recursion to find the long-term average version age of positions. This will be exactly the same as the average version age of a single node in the network. However, for sets of larger size, the average version age depends on the mobility model. In order to carry out exact calculations successfully, high amount of symmetry is needed in order to make sets of the same size have the same version age. Otherwise, we can find tight upper bounds as we did in in Section 7.3 and Section 7.4.*

Remark 9 *In a fully-connected network, adding mobility does not improve the version age of a subset of nodes in the network. To see this, we can show that even if every node can exchange its position with every other node in the network, the version age of a single node scales as $\log n$. In this case, due to the symmetry in the network, any two sets formed in the recursion will have equal version age if they are of the same size. Hence, if $|S| = k$, then in (16), we can replace v_S and $v_{S\cup\{j\}\setminus\{i\}}$ with v_k and $v_{S\cup\{j\}}$ with v_{k+1}. Then, it is easy to see that the extra mobility term gets cancelled, and we are left with just the recursion from [12]. From this, we can conclude that if all nodes can exchange positions with any other node in the network, then the version age scaling is still $\log n$. Hence, the version age of a fully-connected network does not improve with mobility.*

Remark 10 *Nodes being able to exchange positions does not improve the version age of every position in the network, and therefore, does not improve the version age of every node* *in the network. This is because nodes which are in advantageous positions in networks that have no mobility, may not be in the same position for the entire time horizon in networks with mobility, but may travel to positions which typically see less updates from the rest of the network or from the source.*

Remark 11 *The result in* (120) *tells us that if* $\lambda_m = \Omega\left(\frac{\lambda}{\log n}\right)$, *then the version age scaling is the same as a fully-connected network. Moreover, full mobility improves the version age of a single node in the disconnected network described in Section 7.4 to be at least as good as the version age in a ring network.*

Remark 12 *We observe in Section 7.2, Section 7.3 and Section 7.4 that the average version age of a single node in the network does not increase due to the addition of mobility, and in many cases decreases by orders of magnitude. That is, mobility generally improves freshness.*

Remark 13 *Consider mobility networks where positions in a subset of size* $\Theta(n)$ *are connected using one topology, and rest of the positions are connected using another topology. Node exchanges happen between these two types of positions. More specifically, a node in any position in topology* 1 *will be able to exchange with at least one node in a position in topology* 2, *and vice versa. Then, if the node exchange rates are constant independent of* n, *then the node can be seen as switching between the two topologies with very fast rate. Thus, from the node's perspective, it keeps switching between a fresher topology and a staler topology at a fast rate. Thus, the node will be able to maintain average version age equal to the average version age in the fresher topology.*

As an example, consider half the positions connected in a fully-connected fashion, and half the positions being disconnected from any other positions. Suppose a node in each position in the fully-connected topology is able to exchange positions with exactly one node in a position in the disconnected topology with constant rate. Then, we can conclude using the above argument and Theorem 1 that the average version age of each node in the mobility network is $\Theta(\log n)$. *This is a significant improvement over the case where there is no mobility and the average version age of the network without mobility is* $\Theta(n)$.

Remark 14 *In Section 7.4, we found bounds for the long-term average version age of each node in the disconnected network with full-mobility. However, a tighter bound can be obtained using the fact that an update available with a single node at the start reaches every node in* $\Theta(\log n)$ *time. This result is similar to [30, Lemma 1], where it is shown that an update at a single node in the fully-connected network reaches every other node in* $\Theta(\log n)$ *time.*

Here, the argument needs to be modified slightly. Suppose there are k *nodes which have the update at a particular time. We want to find the distribution of the stopping time when one of the* k *nodes sees an arrival for the first time so that they can share their update with a neighbor, like in [30]. Seeing an arrival would be equivalent to adding a node to the set of nodes that have the update. We note that irrespective of the positions of these* k *nodes, due to the Poisson property of memorylessness, the distribution of this time remains the same.*

In other words, when one of the nodes changes positions, the distribution of the time until it sees the next arrival is still distributed as an exponential with rate λ, and this is true for every node. Thus, the stopping time has a distribution of an exponential distribution with rate $k\lambda$, which is the same as in the proof in [30, Lemma 1]. This is because of the high mobility rates for each position, since this ensures that nodes which have updates spend a lot of time with nodes which do not have an update. Thus, the rest of the proof follows along similar lines. Thus, in the mobility network the version age of each node is $O(\log n)$. Note that this upper bound does not depend on the mobility exchange rates, as long as the total mobility rate is at least a constant. Since this network is disconnected, any network which has fully-mobility and better average version age than this disconnected network will have $O(\log n)$ average version age. Thus, we can conclude that full-mobility brings down the average version age of almost all networks to $O(\log n)$, including the ring network, grid network and generalized rings networks.

9 Numerical Results

9.1 Simulations for the Varying-Topology Model

In this section, we carry out simulations to examine how the topologies and CTMC rates affect the average version age of nodes in the gossiping network. For each simulation, we choose $\lambda_e = \lambda_s = \lambda = 1$.

First, we simulate a gossiping network where the CTMC has two states corresponding to the ring topology and the square grid topology in Fig. 10a. Note that having just the ring topology will give average version age $\Theta(\sqrt{n})$ [16,28] and the grid topology will give average version age $\Theta(n^{\frac{1}{3}})$ [27,31]. The number of nodes varies between 100 and 1024 and includes all even square integers between these two. We choose two different types of rates. In the first case, we choose the rates as $\sqrt{n}$ and in the second case, we choose the rates as $n^{\frac{1}{3}}$. We observe in Fig. 10a that for $\sqrt{n}$ holding times, the version age scales as $\Theta(\sqrt{n})$ as well, and for the latter case when the holding times are of the order of $n^{\frac{1}{3}}$, the version age scales as $\Theta(n^{\frac{1}{3}})$.

Next, we simulate the gossiping network which has two states corresponding to the ring topology and the fully-connected topology in Fig. 10b. The number of nodes varies between 200 and 1000 with steps of 200. We choose three different types of expected holding times in this case. We choose linear expected holding times n, square root expected holding times $\sqrt{n}$, and logarithmic expected holding times $\log n$. The expected holding times are the same for both states. We observe that the average version age grows slowly when the expected holding times are logarithmic, but grows faster when the expected holding times are square root and linear. We observe that the version age scaling is $\Theta(\sqrt{n})$ when the expected holding times are linear and square root, and $\Theta(\log n)$ when the expected holding times are logarithmic.

Next, we simulate the network where the CTMC is the "original" CTMC in Fig. 4 and plot the average version age in Fig. 11. We note that the average version age is linear in a static gossip network with the disconnected topology. We choose four different types of expected holding times, n, $\sqrt{n}$, $n^{\frac{1}{3}}$ and $\log n$, for each state. We assume the probability of moving to both neighboring states is equal to half. The number of nodes varies between 100 and 1024 and includes all even square integers between these two. We observe that the version age scales linearly if the expected holding times are linear, and the version age as a function of n gradually becomes smaller as the expected holding times become smaller. We observe that the version age scaling is the same as the expected holding times (among the four rates given) in this network.

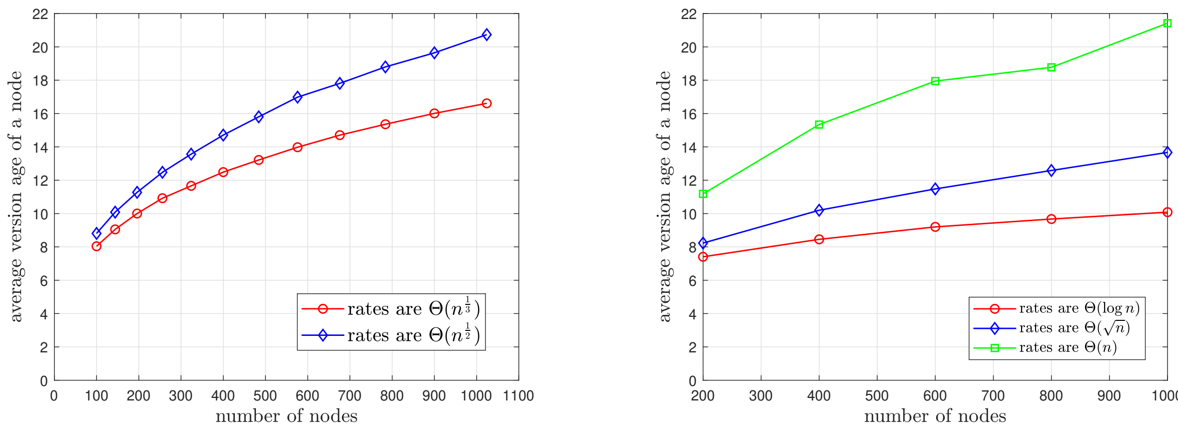

(a) (b)

Figure 10: (a) Variation of average version age of a node as the number of nodes varies when the CTMC has two states, the ring topology and the grid topology. (b) Variation of average version age of a node as the number of nodes varies when the CTMC has two states, the ring topology and the fully-connected topology.

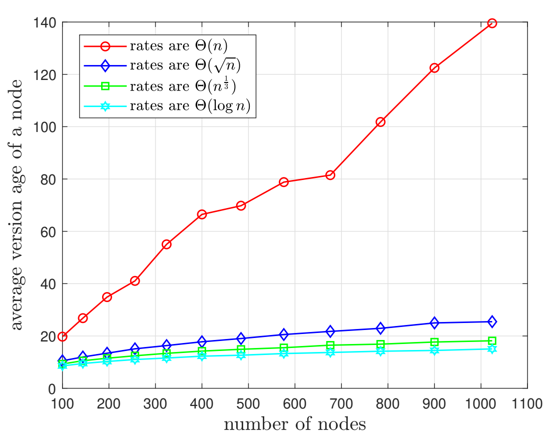

Figure 11: Variation of average version age of a node as the number of nodes varies when the CTMC has four states, the ring topology, the grid topology, the disconnected topology and the fully-connected topology.

Finally, we simulate the varying-topology network considered in Section 6. The number of nodes varies between 200 and 1000 with steps of size 200. In Fig. 12, we verify the results

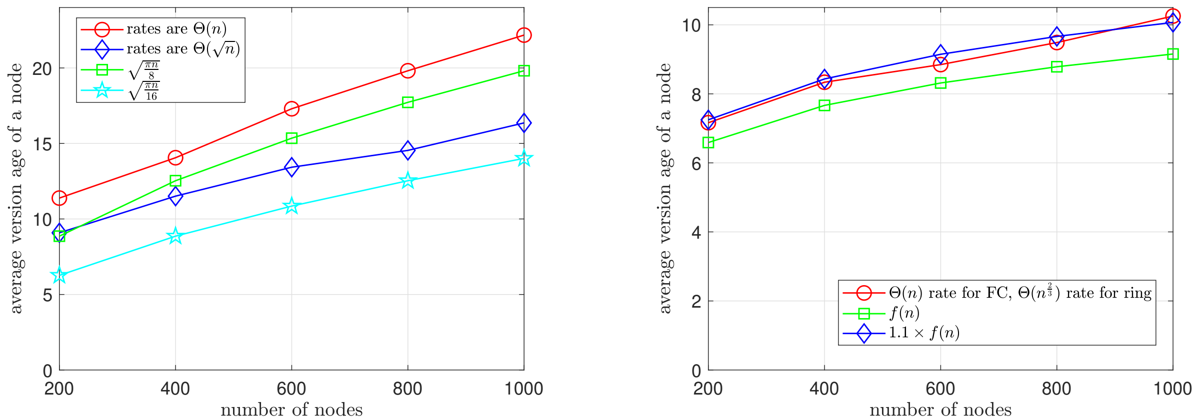

(a) (b)

Figure 12: (a) A comparison of the average version age for linear rates and square root rates of the CTMC in the varying-topology network considered in Section 6. (b) A simulation of a two-topology network described in Section 6, when the CTMC holding times are $\Theta(n)$ for the fully-connected topology, and $\Theta(n^{\frac{2}{3}})$ for the ring topology. Here, $f(n) = \frac{n^{\frac{7}{6}}+n\log n}{n+n^{\frac{2}{3}}}$.

obtained in Theorem 5. In Fig. 12a, we simulate the varying-topology network, when both CTMC holding times have expected value $\sqrt{n}$ and when both CTMC holding times have expected value n. We observe that the average version age scaling in both cases is the same as $\Theta(\sqrt{n})$, as seen due to the two $\Theta(\sqrt{n})$ reference curves. In Fig. 12b, we simulate the varying-topology network when the holding time in the ring topology has expected value $n^{\frac{2}{3}}$ and the holding time in the fully-connected topology has expected value n. We also have two curves for the function from Theorem 5 which describes the order the average version age takes. We observe that the curve is followed closely in this case. Next, in Fig. 13, we verify the results obtained in Theorem 6. In Fig. 13a, Fig. 13b and Fig. 13d, we observe that the bounds obtained are followed closely. In Fig. 13c, we observe that the upper bound obtained is not tight, which shows that the bound can be tightened further.

9.2 Simulations for the Mobility Model

We first plot the results of the toy example for different values of λ_m while choosing $\lambda_e = \lambda = 1$. We choose λ_m among $[0.001, 0.01, 0.1, 1, 10, 100, 1000]$. In Fig. 14a, we plot the average age for each of the three networks. Since the version age of positions in the case of no mobility will not depend on λ_m, the average age is constant. In both cases of node exchanges, the average version age decreases due to mobility. We can clearly see that as λ_m increases, the average version age decreases. Moreover, the average version ages are quite similar for small values of λ_m. However, when nodes in positions 1 and 3 exchange positions with each other, the version age decreases significantly more than when nodes in positions 1 and 2 exchange positions. In Fig. 14b, we plot the individual version age of each position in all three networks. We see that in the case where nodes 1 and 3 exchange positions,

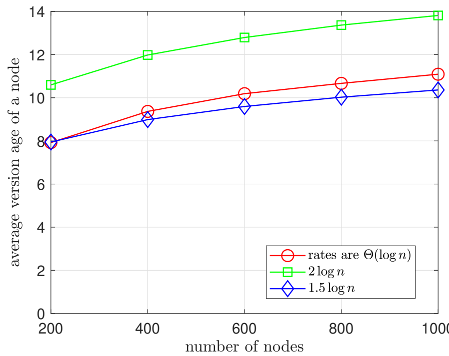

(a)

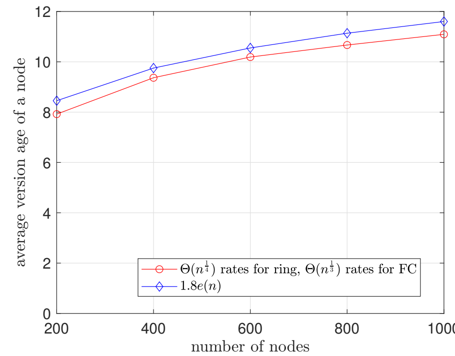

(b)

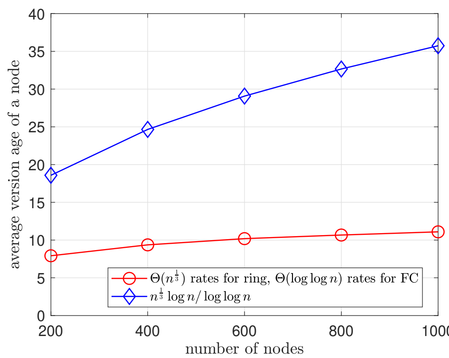

(c)

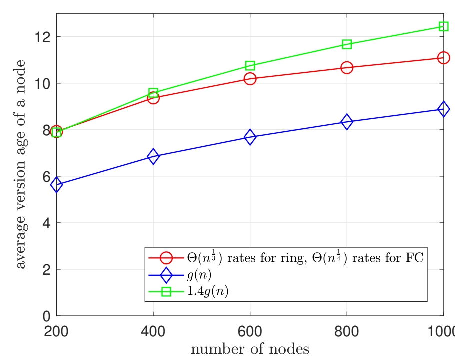

(d)

Figure 13: (a) A simulation of a two-topology network described in Section 6, when the CTMC holding times are $\Theta(\log n)$ for both topologies. (b) A simulation of a two-topology network described in Section 6, when the CTMC holding times are $\Theta(n^{\frac{1}{3}})$ for the fully-connected topology, and $\Theta(n^{\frac{1}{4}})$ for the ring topology. Here, $e(n) = \frac{n^{\frac{1}{2}}+n^{\frac{1}{3}}\log n}{n^{\frac{1}{3}}+n^{\frac{1}{4}}}$. (c) A simulation of a two-topology network described in Section 6, when the CTMC holding times are $\Theta(\log\log n)$ for the fully-connected topology, and $\Theta(n^{\frac{1}{3}})$ for the ring topology. (d) A simulation of a two-topology network described in Section 6, when the CTMC holding times are $\Theta(n^{\frac{1}{4}})$ for the fully-connected topology, and $\Theta(n^{\frac{1}{3}})$ for the ring topology. Here, $g(n) = \frac{n^{\frac{2}{3}}+n^{\frac{1}{4}}\log n}{n^{\frac{1}{3}}+n^{\frac{1}{4}}}$.

the version age of position 1 decreases significantly whereas the version age of position 3 increases significantly and the version age of position 2 remains the same. Moreover, as λ_m becomes large, the version age of position 1 and 3 becomes equal. This is because, on average, both nodes in the positions are in each position half the time. When nodes in positions 1 and 2 exchange, then the version ages of positions 1 and 2 decrease less in magnitude when compared to the previous case. Also, the average version age of position 3 increases less than the previous case. This is because the node in position 2 gets less updates, and it moves to position 1 having high version age for some time. Hence, the version updates of position 3 become less frequent.

Next, we simulate the network considered in Section 7.3. We choose $\lambda_e = \lambda = \lambda_m = 1$. We see that having mobility in the network improves the average version age of the nodes.

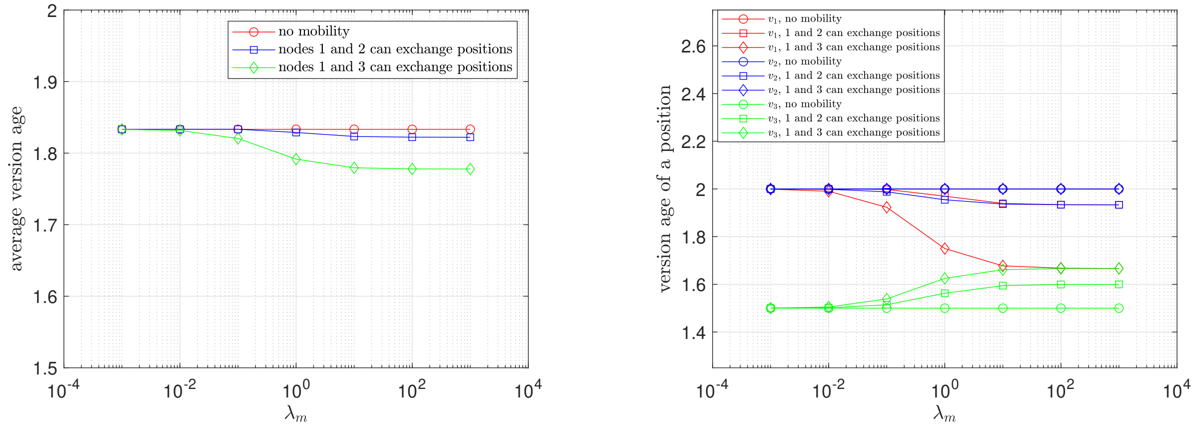

(a) (b)

Figure 14: (a) A comparison of average version age in the toy example. We compare the cases when there is no mobility, represented by red, where nodes 1 and 2 can exchange positions, represented by blue, and where nodes 1 and 3 can exchange positions, represented by green. (b) A comparison of version age of individual nodes in the toy example. The first position is given in red, the second position is given in blue and the third position is given in green.

Moreover, the upper bound we obtain is close to the simulated version age.

Finally, we simulate the network considered in Section 7.4 and plot the results in Fig. 15b. We choose $\lambda_e = \lambda = 1$ and $\lambda_m = \frac{1}{n}$. We see that the simulated version age of the network follows logarithmic scaling. Our upper bound follows square root scaling and the version age scales linearly when there is no mobility. From these simulations, we see that the version age scaling of a single node in the network considered in Section 7.4 scales as $O(\log n)$ if $\alpha(n) = O(n)$.

10 Conclusion

In this work, we examined the effect of time varying topologies and the corresponding transition rates of the underlying CTMC, which determines the active topology, on the average version age of individual nodes as well as that of the overall network. We established analytical results characterizing the temporal behavior of nodes in the gossiping network. Our findings show that when the expected holding times are smaller than the long-term average version age of the fresher topology, the network with varying topologies can sustain the same level of freshness as the fresher topology.

We further analyzed networks where one of the switching topologies is either a ring or a disconnected topology, both representing staler topologies in the varying-topology network. In these cases, we showed that the network maintains the same version age as the topology with the staler long-term average version age if the expected holding times are large enough.

Additionally, we observed that a small subset of nodes can disproportionately influence

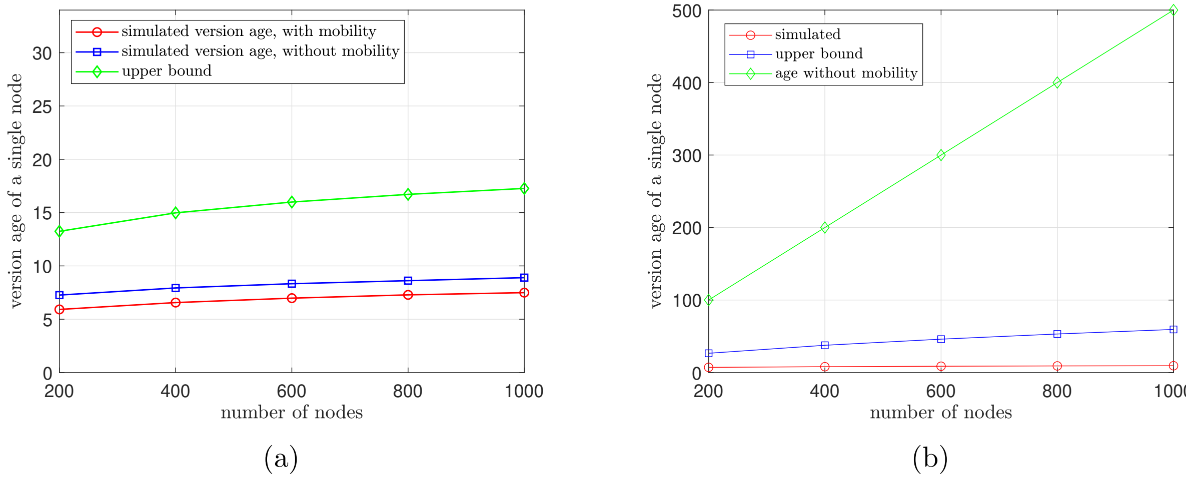

Figure 15: (a) A comparison of the mobility gossip network considered in Section 7.3. We compare the average version age of a single node when there is mobility, when there is no mobility and the upper bound we found in (73). (b) A comparison of the mobility gossip network considered in Section 7.4. We compare the average version age of a single node when there is mobility, when there is no mobility and the upper bound we found in (120).

the network's average version age, even though they do not reflect the experience of most nodes. To better represent the majority behavior, we introduced the notion of a typical set of nodes. We showed that the version age of this typical set coincides with that of the fresher (or staler) topology when the expected holding times are shorter (or longer) than the long-term average version age of the corresponding topology. Finally, we demonstrated that when the expected holding times lie between the long-term average version ages of the fresher and staler topologies, the network's version age aligns with the timescale of the CTMC transitions.

In all the above cases, we assumed that the CTMC transition rates are of the same order. Next, we analyzed a two-topology network consisting of the ring and fully-connected topology, with unequal transition rates. We analyzed the entire spectrum of connectivity from expected holding times with constant expectations to expected holding times with superlinear expectations, and found the average version age of each node.

We then performed the analysis of the mobility model, which can be considered to be a special case of the varying-topology network. We discussed why in this special case, it is possible to perform an SHS analysis. We then used the SHS framework to find recursive equations where the version age of a set of nodes can be found using the version ages of sets that are of the same size and sets that are one larger. We then found bounds for the version ages of nodes in the fully-connected network with a disconnected node, and a disconnected network with full-mobility using the recursive equations. We showed that in the disconnected network, the version age of a node is upper bounded by $O(\sqrt{n})$ in the worst case, and by a constant in the best case, which is a large improvement over the case where there is no mobility.

Finally, we performed numerical simulations to further interpret the properties of varying-topology networks and mobility networks. We also verified the results we obtained earlier in the paper.

Appendix

A Proof of Lemma 1

Proof: We first recall the main result from [21], which states that once a node has received some update generated by the source, the version age of the node follows the same distribution as the first passage percolation time from the source to the node. In order to use this result in our proof, we need to choose a large enough time so that we can be sure that node 1 has received an update w.h.p. Since the source sends updates to all nodes with an equal rate of $\frac{\lambda}{n}$, it can be shown that the source sends a direct update to all the nodes in $\Theta(n \log n)$ time w.h.p. (we also prove this in the proof of Theorem 4, item 2). Thus, we choose a time t_1 which is much larger than $n \log n$ in order. At this time, we have that the distribution of the version age of node 1 has converged to the distribution of the first passage percolation time of node 1 from the source. Thus, for all times $t > t_1$, the distribution of the version age of node 1 is the same as the distribution of the first passage percolation time w.h.p. This also implies that the limiting distribution of the version age of node 1 is the same as the distribution of the first passage percolation time as well. Thus, we have,

$$\lim_{t \to \infty} \mathbb{E}[X_1(t)] = \mathbb{E}[X_1(t)], \quad \forall t \geq t_1. \tag{121}$$

Now, we know that version age is a discrete non-negative random variable. Thus, the version age of node 1 has a probability mass function (PMF). Now, we choose a fixed time $t_2 > t_1$, and analyze the distribution and expectation of version age of node 1 at this fixed time. Thus, we have $\mathbb{E}[X_1(t_2)] = \Theta(f(n))$ from (121) and the statement of the lemma. Therefore, let $\mathbb{E}[X_1(t_2)] = Cf(n)$, $C > 0$. In order to complete the proof for Lemma 1, we want to show that the version age of node 1 is $O(f(n))$ w.h.p. at time t_2. To this end, choose a function $g(n)$ such that $g(n) \to \infty$. Then, by Markov's inequality, we have,

$$\begin{aligned}
\mathbb{P}[X_1(t_2) > f(n)g(n)] \leq & \frac{\mathbb{E}[X_1(t_2)]}{f(n)g(n)} && (122)\\
= & \frac{Cf(n)}{f(n)g(n)} && (123)\\
= & \frac{C}{g(n)}. && (124)
\end{aligned}$$

Thus, the probability goes to 0 as $n \to \infty$. Hence, the version age of node 1 is less than $f(n)g(n)$ w.h.p. This is true for every growing function $g(n)$, no matter how slowly it grows. Thus, we can conclude that $X_1(t_2) = O(f(n))$ w.h.p., proving the lemma. ■

B Proof of Lemma 2

Proof: We start by discussing the convergence result with respect to the stationary regime. Intuitively, it seems that the result follows from the stationarity and ergodicity of the finite-state CTMC. However, this gives a result only in expectation, and does not ensure a convergence result with high probability (w.h.p.). To obtain a w.h.p. result, we employ Chebyshev's inequality after calculating the expectations and variances of appropriate random variables.

For some $\beta \geq 0$, let T be the first time the CTMC has spent $\beta g(n)$ time in state k. Then, T is a stopping time, and furthermore at time T, we are in state k of the CTMC. We also start with the stationary distribution, thus the CTMC may be starting in state k or some other state of the CTMC. If the CTMC starts in state k, it will leave the state $C_k(\beta g(n))$ times, where C_k is the counting process associated with the Poisson process when in state k of the CTMC. If it starts in another state of the CTMC, it will reach state k and then return to the state $C_k(\beta g(n))$ times. We note that the time to return to state k is an inter-arrival time for a renewal process whose arrival times are the instances of return to the state k. These inter-arrival times, say $\{I_i\}_{i=1}^{C_k(\beta g(n))}$, follow a phase-type distribution; see [32].

We start by finding $\mathrm{Var}[T]$. Suppose we spend a total of T_{out} time in other states up to time T, then $\mathrm{Var}[T] = \mathrm{Var}[T_{out}]$. Further, let the time of reaching the state k for the first time be T_1, with $T_1 = 0$ if we start in the state k. Then, $T_{out} = T_1 + \sum_{i=1}^{C_k(\beta g(n))} I_i$. Further, we note that $C_k(\beta g(n)) + 1$ is a stopping time for the Poisson process associated with state k. We note that $\mathbb{E}[T_1] \leq \eta \mathbb{E}[I_1]$ for some $\eta \in \mathbb{R}$. Let $\mathrm{Var}[I_i] = \sigma^2 = \Theta((h(n))^2)$ and $\mathbb{E}[I_i] = \mu = \Theta(h(n))$. These can be derived from the definitions of the expectation and second moment of phase-type distributions which are defined in [32]. Then, we have

$$\begin{aligned}
\mathbb{E}[T_{out}] =& \mathbb{E}[T_1] + \mathbb{E}\left[\sum_{i=1}^{C_k(\beta g(n))} I_i\right] && (125)\\
\leq& (\eta - 1)\mu + \mathbb{E}\left[\sum_{i=1}^{C_k(\beta g(n))+1} I_i\right] && (126)\\
=& (\eta - 1)\mu + \mathbb{E}[C_k(\beta g(n)) + 1]\mu && (127)\\
=& (\eta - 1)\mu + \left(q_k \beta \frac{g(n)}{h(n)} + 1\right)\mu && (128)\\
\leq& \rho g(n), && (129)
\end{aligned}$$

for some $\rho \in \mathbb{R}$. The expectation is also lower bounded by $\Theta(g(n))$ if we remove T_1 from the equation and write the expectation. Thus, we let $\mathbb{E}[T_{out}] = \nu g(n)$.

Due to the strong Markov property and independence of these inter-arrival times, we have, $\mathrm{Var}[T_{out}] = \mathrm{Var}[T_1] + \sum_{i=1}^{C_k(\beta g(n))} \mathrm{Var}[I_i]$. We further see that $\mathrm{Var}[T_1] \leq \chi \mathrm{Var}[I_0]$, where $\chi \in \mathbb{R}$ and I_0 has the same phase-type distribution as I_1. This is due to the positive recurrence of the CTMC. Thus, we can write $\mathrm{Var}[T_{out}] \leq (\chi^2 - 1)(h(n))^2 + \sum_{i=1}^{C_k(\beta g(n))+1} \mathrm{Var}[I_i]$.

Define the event $\mathcal{E} = \{C_k(\beta g(n)) + 1 = \Theta(g(n)/h(n))\}$. This event happens w.h.p. since

$C_k(\beta g(n))$ follows a Poisson distribution with mean $\Theta\left(\frac{g(n)}{h(n)}\right)$. We define this event because we want to find the variance of a sum of random variables where the number of terms in the summation is a random variable which is not independent of the terms in the sum. Thus, using Wald's identity is not possible, and conditioning on the event is necessary.

Now, we can find the upper bound for $\mathrm{Var}[T|\mathcal{E}]$,

$$\mathrm{Var}[T|\mathcal{E}] = \mathrm{Var}[T_{out}|\mathcal{E}] \tag{130}$$

$$\leq (\chi^2 - 1)(h(n))^2 + \mathrm{Var}\left[\sum_{i=1}^{C_k(\beta g(n))+1} I_i \Big| \mathcal{E}\right] \tag{131}$$

$$\leq (\chi^2 - 1)(h(n))^2 + \zeta \frac{g(n)}{h(n)}(h(n))^2 \tag{132}$$

$$\leq \xi h(n) g(n), \tag{133}$$

for some $\xi \in \mathbb{R}$. Then, using Chebyshev's inequality, we write

$$\mathbb{P}\big[|T - \mathbb{E}[T]| > g(n)\big] = \mathbb{P}\big[|T - \mathbb{E}[T]| > g(n)\big|\mathcal{E}]\mathbb{P}[\mathcal{E}] + \mathbb{P}\big[|T - \mathbb{E}[T]| > g(n)\big|\mathcal{E}^{\mathcal{C}}]\mathbb{P}[\mathcal{E}^{\mathcal{C}}] \tag{134}$$

$$\leq \frac{\mathrm{Var}[T|\mathcal{E}]}{(g(n))^2} + \mathbb{P}[\mathcal{E}^{\mathcal{C}}] \tag{135}$$

$$\leq \frac{\xi h(n)}{g(n)} + o(1) \tag{136}$$

$$= o(1). \tag{137}$$

Hence, $\beta g(n) + \rho g(n) - g(n) \leq T \leq \beta g(n) + \rho g(n) + g(n)$ w.h.p. ■

C Proof of Lemma 3

Proof: Let $T = \tau f(n)$. We know that $N_0(t)$ is Poisson distributed with $\lambda_e t$. Then, we have that $\mathbb{E}[N_0([0, \tau f(n)])] = \lambda_e \tau f(n)$ and $\mathrm{Var}[N_0([0, \tau f(n)])] = \lambda_e \tau f(n)$. Hence, using Chebyshev's inequality, and using N as a shorthand for $N_0([0, \tau f(n)])$, we have

$$\mathbb{P}\big[|N - \mathbb{E}[N]| > f(n)\big] \leq \frac{\mathrm{Var}[N]}{(f(n))^2} \tag{138}$$

$$= \frac{\lambda_e \tau f(n)}{(f(n))^2} \tag{139}$$

$$= \frac{\lambda_e \tau}{f(n)} \tag{140}$$

$$= o(1). \tag{141}$$

Hence, $(\lambda_e \tau - 1) f(n) \leq N_0([0, \tau f(n)]) \leq (\lambda_e \tau + 1) f(n)$ w.h.p. ■

D Proof of Lemma 4

Proof: Let the node be labeled 1 without loss of generality. We know from Lemma 3 that in the given time interval $[t, t+\alpha\sqrt{n}]$, the source updates itself $\Theta(\sqrt{n})$ times. Thus, node 1 cannot exceed $\Theta(\sqrt{n})$ version age. Next, suppose node 1 does not reach $\Theta(\sqrt{n})$ version age in the desired time interval. Then, the node reaches some $a(n) = o(\sqrt{n})$, but does not exceed it (in order). This means that node 1 receives a fresher update than the one it has at least once every $O(a(n))$ time.

The packets that update node 1 come from nodes in a small neighborhood of size $\Theta(a(n))$. To see this, assume that a packet reaches node 1 from a node b that is $b(n) = \omega(a(n))$ distance away, i.e., there are $b(n)$ nodes between node 1 and node b. Let the time taken by the packet to reach node 1 from node b be S. Then, $S = \sum_{i=1}^{b(n)} X_i$, where X_i are i.i.d. exponential random variables with rate $\frac{\lambda}{2}$. $\mathbb{E}[S] = \alpha_1 b(n)$. Further, $\text{Var}[S] = \sum_{i=1}^{b(n)} \text{Var}[X_i] = \alpha_2 b(n)$. From Chebyshev's inequality, we have,

$$\mathbb{P}(|S - \mathbb{E}[S]| > \frac{1}{2}b(n)) \leq \frac{\text{Var}[S]}{\frac{1}{4}(b(n))^2} \tag{142}$$

$$\leq \frac{c}{b(n)} \tag{143}$$

$$\to 0, \tag{144}$$

as $n \to \infty$. Thus, we conclude that $\frac{1}{2}b(n) \leq S \leq \frac{3}{2}b(n)$ w.h.p. Thus, by the time a packet reaches node 1 from a node $b(n)$ distance away, the packet's version age would already have reached $\Theta(b(n))$ by Lemma 3. This value is strictly greater than $\Theta(a(n))$. Thus, these packets cannot maintain the version age of node 1 at $O(a(n))$. Hence, the packet can only reach node 1 in $O(a(n))$ time if the nodes are at most $\Theta(a(n))$ distance away.

Now, let the set of nodes that are $O(a(n))$ distance away from node 1 be $\mathcal{A}$. Then, the source needs to send an update to $\mathcal{A}$ at least once every $\Theta(a(n))$ time, so that the packet can reach node 1 fast enough to maintain $O(a(n))$ version age. Thus, the source needs to send $\Omega\left(\frac{\sqrt{n}}{a(n)}\right)$ updates to $\mathcal{A}$ to maintain the version age at node 1. However, the source sends updates to every node at equal rate. Since $|\mathcal{A}| = \Theta(a(n))$, the number of updates received by $\mathcal{A}$ in the given time interval is $\Theta(\frac{a(n)}{n}\sqrt{n}) = \Theta(\frac{a(n)}{\sqrt{n}})$ w.h.p. This number is not enough for node 1 to maintain $a(n)$ version age. Thus, the node cannot maintain $a(n)$ version age and its version age reaches $\Theta(\sqrt{n})$ in the desired interval.

Next, to show that the average version age during this interval is $\Theta(\sqrt{n})$, we will show that the area under the version age curve of node 1 is $\Theta(n)$. An upper bound for the area is if the node never receives a fresher update than the one it has at time t. In this case, the area is given by $Ar = \sum_{i=1}^{\beta\sqrt{n}} iY_i$ where $\beta\sqrt{n}, \beta > 0$ is the number of times the source updates itself, and Y_i's represent the inter-arrival times of the source self-update process.

Then, $\mathbb{E}[Ar] = \sum_{i=1}^{\beta\sqrt{n}} i\mathbb{E}[Y_i] = \frac{1}{2\lambda_e}\beta\sqrt{n}(\beta\sqrt{n}+1)$. Moreover,

$$\text{Var}[Ar] = \mathbb{E}[Ar^2] - (\mathbb{E}[Ar])^2 \tag{145}$$

$$=\frac{1}{\lambda_e^2}\sum_{i=1}^{\beta\sqrt{n}} i^2 + \frac{1}{\lambda_e^2}\Big(\sum_{i=1}^{\beta\sqrt{n}} i\Big)\Big(\sum_{j=1}^{\beta\sqrt{n}} j\Big) - (\mathbb{E}[Ar])^2 \tag{146}$$

$$=\frac{1}{6\lambda_e^2}\beta\sqrt{n}(\beta\sqrt{n}+1)(2\beta\sqrt{n}+1) \tag{147}$$

Thus, from Chebyshev's inequality, we have,

$$\mathbb{P}[|Ar - \mathbb{E}[Ar]| > \frac{1}{2}\beta n] \leq \frac{\text{Var}[Ar]}{\frac{1}{4}\beta^2 n^2} \tag{148}$$

$$= \frac{c}{\sqrt{n}} \tag{149}$$

$$\to 0, \tag{150}$$

as $n \to \infty$. Hence, the area upper bound is $\Theta(n)$ w.h.p. Similarly, we can obtain a lower bound for the area if we assume that there is only one continuous rise of the version age graph for node 1 to $\Theta(\sqrt{n})$, and the version age is 0 everywhere else in the time interval. We are able to do this because of the result we proved earlier which states that we will reach $\Theta(\sqrt{n})$ version age for each node at some point of time during the interval. Then, using a similar calculation as the upper bound, we can show that the lower bound of the area is also $\Theta(n)$ w.h.p. Thus, the area under the version age curve for node 1 is $\Theta(n)$ w.h.p. Thus, the average version age for node 1 in the given time interval is $\Theta(\sqrt{n})$ w.h.p. ■

E Proof of Lemma 6

Proof: Item 1 follows from Lemma 4. To prove item 2, we want to show that the node does not receive an update during the holding time in the ring network w.h.p. We only give a sketch of the proof here, since it is similar to the proof of the first part of Lemma 4. First we show that the node does not receive an update during the time spent in the ring network. For a proof by contradiction, suppose the node does receive an update during this time. Then, the update must come from a set of nodes $\Theta(f(n) + g(n))$ distance away. Thus, this set must have received an update directly from the source during the holding time. However, this set of nodes does not receive an update directly from the source during this time w.h.p. Hence, we have a contradiction. Next, the version age of the node increases by $\Theta(g(n))$ by the end of the time in the ring network. Thus, the version age of the node at the end is $\Theta(f(n) + g(n))$. Then, using an area calculation as in Lemma 4, we see that the average version age of the node during the time spent in the ring topology is $\Theta(f(n) + g(n))$. ■

References

[1] S. R. Pokhrel, J. Ding, J. Park, O. Park, and J. Choi. Towards enabling critical mMTC: A review of RLLC within mMTC. *IEEE Access*, 8:131796–131813, 2020.

[2] L. Chettri and R. Bera. A comprehensive survey on internet of things (IoT) toward 5G wireless systems. *IEEE Internet of Things Journal*, 7(1):16–32, October 2019.

[3] S. N. Swamy and S. R. Kota. An empirical study on system level aspects of internet of things (IoT). *IEEE Access*, 8:188082–188134, October 2020.

[4] S. K. Kaul, R. D. Yates, and M. Gruteser. Real-time status: How often should one update? In *IEEE Infocom*, March 2012.

[5] Y. Sun, I. Kadota, R. Talak, and E. Modiano. Age of information: A new metric for information freshness. *Synthesis Lectures on Communication Networks*, 12(2):1–224, December 2019.

[6] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus. Age of information: An introduction and survey. *IEEE Jour. on Selected Areas in Communications*, 39(5):1183–1210, May 2021.

[7] R. D. Yates and S. K. Kaul. The age of information: Real-time status updating by multiple sources. *IEEE Transactions on Information Theory*, 65(3):1807–1827, March 2018.

[8] S. Banerjee, S. Ulukus, and A. Ephremides. To re-transmit or not to re-transmit for freshness. In *IEEE WiOpt*, August 2023.

[9] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides. The age of incorrect information: A new performance metric for status updates. *IEEE/ACM Trans. on Networking*, 28(5):2215–2228, October 2020.

[10] J. Zhong, R. D. Yates, and E. Soljanin. Two freshness metrics for local cache refresh. In *IEEE ISIT*, June 2018.

[11] J. Cho and H. Garcia-Molina. Effective page refresh policies for web crawlers. *ACM Trans. on Database Systems*, 28(4):390–426, December 2003.

[12] R. D. Yates. The age of gossip in networks. In *IEEE ISIT*, July 2021.

[13] B. Abolhassani, J. Tadrous, A. Eryilmaz, and E. Yeh. Fresh caching for dynamic content. In *IEEE Infocom*, May 2021.

[14] M. Bastopcu and S. Ulukus. Who should Google Scholar update more often? In *IEEE Infocom*, July 2020.

[15] A. Chaintreau, J. Y. Le Boudec, and N. Ristanovic. The age of gossip: spatial mean field regime. *ACM SIGMETRICS Performance Evaluation Review*, 37(1):109–120, June 2009.

[16] B. Buyukates, M. Bastopcu, and S. Ulukus. Version age of information in clustered gossip networks. *IEEE Jour. on Selected Areas in Information Theory*, 3(1):85–97, March 2022.

[17] P. Kaswan and S. Ulukus. Age of gossip in ring networks in the presence of jamming attacks. In *Asilomar Conference*, October 2022.

[18] P. Mitra and S. Ulukus. Timely opportunistic gossiping in dense networks. In *IEEE Infocom*, May 2023.

[19] T. J. Maranzatto. Age of gossip in random and bipartite networks. In *IEEE ISIT*, July 2024.

[20] P. Kaswan, P. Mitra, A. Srivastava, and S. Ulukus. Age of information in gossip networks: A friendly introduction and literature survey. *IEEE Transactions on Communications*, 73(8):6200–6220, February 2025.

[21] T. J. Maranzatto and M. Michelen. Age of gossip from connective properties via first passage percolation. *IEEE Transactions on Information Theory*, 71(11):9019–9027, September 2025.

[22] V. Tripathi. *Age of information and mobility.* PhD thesis, Massachusetts Institute of Technology, 2019.

[23] F. Wu, H. Zhang, J. Wu, Z. Han, H. V. Poor, and L. Song. UAV-to-device underlay communications: Age of information minimization by multi-agent deep reinforcement learning. *IEEE Transactions on Communications*, 69(7):4461–4475, 2021.

[24] E. Eldeeb, J. M. Sant'Ana, D. E. Pérez, M. Shehab, N. H. Mahmood, and H. Alves. Multi-UAV path learning for age and power optimization in IoT with UAV battery recharge. *IEEE Transactions on Vehicular Technology*, 72(4):5356–5360, April 2023.

[25] C. E. Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, July 1948.

[26] A. Srivastava and S. Ulukus. Age of gossip on generalized rings. In *IEEE MILCOM*, October 2023.

[27] A. Srivastava and S. Ulukus. Age of gossip on a grid. In *Allerton Conference*, September 2023.

[28] B. Buyukates, M. Bastopcu, and S. Ulukus. Age of gossip in networks with community structure. In *IEEE SPAWC*, September 2021.

[29] J. Hespanha. Modeling and analysis of stochastic hybrid systems. *IEEE Proceedings – Control Theory and Applications*, 153(5):520–535, January 2006.

[30] A. Srivastava, T. J. Maranzatto, and S. Ulukus. Age of gossip with time-varying topologies. In *IEEE ISIT*, June 2025.

[31] A. Srivastava and S. Ulukus. Network connectivity–information freshness tradeoff in information dissemination over networks. *IEEE Transactions on Information Theory, to appear*, September 2025. Available at arXiv:2405.19310.

[32] M. Bladt. A review on phase-type distributions and their use in risk theory. *ASTIN Bulletin: The Journal of the IAA*, 35(1):145–161, April 2005.