Susceptibility of Age of Gossip to Timestomping

Priyanka Kaswan Sennur Ulukus
Department of Electrical and Computer Engineering
University of Maryland, College Park, MD 20742
pkaswan@umd.edu *ulukus@umd.edu*

***Abstract*—We consider a fully connected network consisting of a source that maintains the current version of a file, n nodes that use asynchronous gossip mechanisms to disseminate fresh information in the network, and an adversary who infects the packets at a target node through data *timestamp manipulation,* with the intent to replace circulation of fresh packets with outdated packets in the network. We show that a single infected node increases the expected age of a fully connected network from $O(\log n)$ to $O(n)$. Further, we show that the optimal behavior for an adversary is to reset the timestamps of all outgoing packets to the current time and of all incoming packets to an outdated time. Additionally, if the adversary allows the infected node to accept a small fraction of incoming packets from the network, then a large network can manage to curb the spread of stale files coming from the infected node and pull the network age back to $O(\log n)$. Lastly, we show that if an infected node contacts only a single node instead of all nodes of the network, the system age can still be degraded to $O(n)$. These show that fully connected nature of a network can be both a benefit and a detriment for information freshness; full connectivity, while enabling fast dissemination of information, also enables fast dissipation of adversarial inputs.**

I. Introduction

Sensor networks generally have limited resources, which prevents them from implementing traditional computer security techniques, making them vulnerable to adversarial attacks. Uncertain dynamics of such networks often force them to rely on decentralized *gossip protocols* [1]–[14] for information dissemination, where information is exchanged between nodes repeatedly and asynchronously using their local status. Gossip protocols were introduced and have been widely used in the context of distributed databases. In this work, we consider the presence of an adversary in a gossip network [15]–[22], who corrupts the gossip operation by manipulating the timestamps of some data packets flowing in the network, a technique known as *timestomping* [23], with the goal of bringing about staleness and inefficiency to the network. A timestomping attack can be launched in many ways. For instance, a malicious insider node can deviate from the gossip protocol and inject old packets by rebranding them as fresh packets via timestamp manipulation, while maintaining the gossiping frequency to evade suspicion. Other methods include *meddler in the middle* (MITM) attacks, where the adversary inserts its node undetected between two nodes and manipulates communication, and *eclipse* attacks where the adversary manipulates the target node by redirecting its inbound and outbound links away from legitimate neighboring nodes to adversary controlled nodes, thereby isolating the node from the rest of the network, as encountered in gossip based blockchain networks.

Most prior works on gossip networks consider total dissemination time of a message in the network as the performance metric. For instance, [3] shows that dissemination of a single rumor to n nodes takes $O(\log n)$ rounds, [5] shows that n messages can be disseminated to n nodes in $O(n)$ time in fully connected networks using random linear coding (RLC), [6] provides an analogous result for arbitrarily connected graphs, and [7] analyzes dissemination of messages by dividing them into pieces. However, highly dynamic nature of data sources in modern applications prevents these networks from waiting for a specific message to reach all nodes of the network before fresh information can be circulated. Distributed databases [8], [9], for example, employ timestamp versioning, wherein every new information is created with a timestamp value taken from the system clock. When two nodes come in contact to exchange information, the timestamps of data at both nodes are compared and the node carrying the data with older timestamp discards its data for the fresher data of the other node.

In this regard, age of information [24]–[26] may be a more suitable indicator of network efficiency. Given $U_i(t)$ as the timestamp of the packet with node i at time t, the instantaneous age of information is given by $X_i(t) = t - U_i(t)$. The nodes wish to have access to the most up-to-date information at all times, and therefore, are prompted to decrease $X_i(t)$ by fetching packets with more recent timestamps, e.g., with higher $U_i(t)$. Gossip networks have been studied from timeliness perspective in [10]–[14]. [10], [11] derive the recursive age equations using stochastic hybrid system framework for age, [12] studies the expected version age in clustered gossip networks, [13] extends these results to the binary freshness metric, [14] considers age scaling in gossip networks using file slicing and network coding, and [18] studies the effects of jamming adversaries on gossip age in ring networks.

Timestomping is often used by malware authors as an anti-forensics technique to make files blend in with the rest of the system. In this work, an adversary uses timestomping with the goal of worsening the expected age in the network. Consider two nodes, A and B, that randomly come in contact to exchange information and consider the presence of adversary at node A capable of altering timestamps of all incoming and outgoing files. If node A is outdated compared to node B, the adversary would be inclined to increase the timestamp of an outgoing packet from node A to make it appear fresher so as to misguide node B into discarding its packet in favor of a staler packet, and also, decrease the timestamp of an incoming packet from node B so as to avoid its acceptance at node

A. Conversely, if node A is more up-to-date than node B, the adversary would reduce timestamps of outgoing packets and increase timestamps of incoming packets to make node B reject fresher files and node A accept staler files. More the manipulated timestamps digress away from their true value, higher are the chances of error in deciding which packet should be discarded, since this decision is based on a comparison of timestamps. At time t, the maximum error is caused when file timestamp is either changed to the current time t or the earliest time 0. Thus, we consider an adversary, who, for each packet, makes the decision of changing its timestamp to either t or 0. The adversary is *oblivious* in that it does not look into a packet and see its actual timestamp. Thus, the adversary changes the timestamp to either t or 0 probabilistically.

In this paper, we consider a gossip network where an adversary captures a node and manipulates the timestamps of packets coming into and going out of the node (see Fig. 1). We show that one infected node can single-handedly suppress the availability of fresh information in a large network of n users employing a gossip protocol, and increase the expected age in a complete graph from $O(\log n)$ found in [10] to $O(n)$. In addition, we show that the optimal action for the adversary is to always increase the timestamp of every outgoing packet to t and decrease the timestamp of every incoming packet to 0, in effect, preventing all incoming files from being accepted and actively persuading other nodes to always accept outgoing packets from the infected node. Further, we show that if the the infected node is allowed to accept even a small fraction of incoming packets from the network, then a large network can curb the spread of stale files coming from the infected node by effectively lowering the infected node's age. These observations show how the fully connected nature of a network can be both a benefit and a detriment for network staleness. Additionally, we show that if the malicious node contacts only one other node instead of all nodes of the network (see Fig. 2), the system age can still be degraded to $O(n)$, which highlights how little an effort is needed on the part of the adversary to bring down the freshness of the entire network.

II. System Model and SHS Characterization

We study a fully connected network, shown in Fig. 1, which comprises a source and n user nodes $\mathcal{N} = \{1, \ldots, n\}$. The source, alternatively referred to as node 0, is assumed to always posses the latest file packet and consequently has zero age at all times. The nodes wish to acquire the most up-to-date file to lower their average age from the source, who updates each user node as a Poisson process with rate $\frac{\lambda}{n}$. Further, a user node i randomly sends its current packet to a user node j according to a Poisson process with rate $\lambda_{ij} = \frac{\lambda}{n-1}$. Thus, all nodes send out updates after exponential inter-update times with a total rate λ. Let $U_j(t)$ denote the timestamp marked on the file stored at the node j. Then, at the receiving node j, the claimed timestamp of the incoming packet is compared with $U_j(t)$ to determine which packet should be kept. Note that a node always accepts an update from the source which generates update packets with current timestamp t.

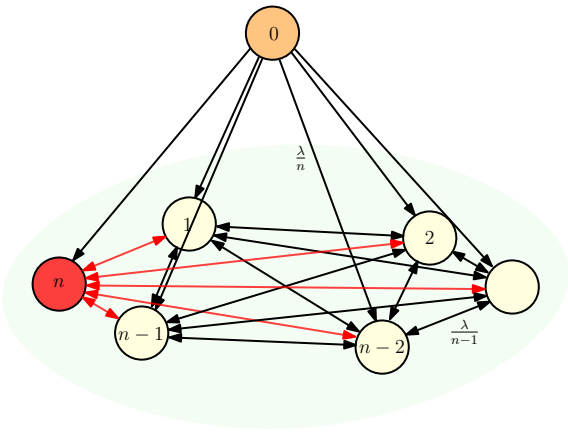

Fig. 1. Fully connected network of n nodes with an infected node.

We assume that the highest index node, node n, is under attack by an adversary that manipulates the timestamps of all incoming and outgoing packets of node n. For the outgoing packets, the adversary chooses to increase the timestamp (to current time t) with probability p and decrease the timestamp to 0 with probability $1-p$. Similarly, the adversary increases and decreases the timestamp of incoming packets with probability $1-p$ and p, respectively. We will refer to the nodes in set $\mathcal{N}_R = \{1, \ldots, n-1\}$ as *regular* nodes and node n as the infected node. We assume that the infected node always accepts packets from the source like other nodes, delivered to it with rate $\frac{\lambda}{n}$, which helps the adversary evade suspicion of malicious activity by maintaining a remote contemporary relevance of the contents of its manipulated packets.

We denote the long-term average age at node i by v_i, where $v_i = \lim_{t\to\infty} \mathbb{E}[X_i(t)]$, and wish to study its extent of deterioration through timestomping. Note that the actual instantaneous age at node i is $X_i(t) = t - \bar{U}_i(t)$, where $\bar{U}_i(t)$ indicates the true packet generation time, which can be different from the claimed timestamp $U_i(t)$ if the file timestamp has been tampered with. For a set of nodes S at time t, let $X_{N(S)}(t)$ indicate the actual instantaneous age of the node claiming to possess the most recent timestamped packet in set S, i.e., $X_{N(S)}(t) = X_{\arg\max_{j\in S} U_j(t)}(t)$. We define $v_S = \lim_{t\to\infty} \mathbb{E}[X_{N(S)}(t)]$. Here we would like to point out that in a network without adversary where all files are marked with true timestamps, $X_{N(S)}(t)$ reduces to $X_S(t) = \min_{j\in S} X_j(t)$ defined in [10], since the node with highest timestamp will also have the lowest age in the set S.

Reference [10] demonstrates how stochastic hybrid system (SHS) models yield linear equations useful for deriving long-term average age at nodes in a gossip network of n users with a given topology. Due to the presence of a timestomping adversary, we choose the continuous state for our SHS model as $(\boldsymbol{X}(t), \boldsymbol{U}(t)) \in \mathbb{R}^{2n}$, where $\boldsymbol{X}(t) = [X_1(t), \ldots, X_n(t)]$ denotes the instantaneous ages at the n nodes and $\boldsymbol{U}(t) = [U_1(t), \ldots, U_n(t)]$ denotes the timestamps marked on the packets at the n nodes at time t. The convenience of the SHS based age characterization follows from the presence of a single discrete mode with trivial stochastic differential

equation $(\dot{\boldsymbol{X}}(t), \dot{\boldsymbol{U}}(t)) = (\mathbf{1}_n, \mathbf{0}_n)$, where the age at each node grows at unit rate when there is no update transfer, since the timestamps of the node packets do not change between such transitions. Consider a test function $\psi : \mathbb{R}^{2n} \times [0, \infty) \to \mathbb{R}$ that is time-invariant, i.e., its partial derivative with respect to t is $\frac{\partial \psi(\boldsymbol{X}, \boldsymbol{U}, t)}{\partial t} = 0$, such that we are interested in finding its long-term expected value $\mathbb{E}[\psi] = \lim_{t\to\infty} \mathbb{E}[\psi(\boldsymbol{X}(t), \boldsymbol{U}(t), t)]$. Since the test function only depends on the continuous state values $(\boldsymbol{X}, \boldsymbol{U})$ and is time-invariant, for simplicity, we will drop the third input t and write $\psi(\boldsymbol{X}, \boldsymbol{U}, t)$ as $\psi(\boldsymbol{X}, \boldsymbol{U})$, which we assume to flow according to the differential equation $\dot{\psi}(\boldsymbol{X}(t), \boldsymbol{U}(t)) = 1$. Let $\mathcal{L}$ correspond to the set of directed edges (i, j), such that node i sends updates to node j on this edge according to a Poisson process of rate λ_{ij}, with this transition resetting the state $(\boldsymbol{X}, \boldsymbol{U})$ at time t to $\phi_{i,j}(\boldsymbol{X}, \boldsymbol{U}, t) \in \mathbb{R}^{2n}$ post transition. Defining $\mathbb{E}[\psi(\phi_{i,j})] = \lim_{t\to\infty} \mathbb{E}[\psi(\phi_{i,j}(\boldsymbol{X}(t), \boldsymbol{U}(t), t))]$, [27, Thm. 1] yields

$$0 = 1 + \sum_{(i,j)\in\mathcal{L}} \lambda_{ij}(\mathbb{E}[\psi(\phi_{i,j})] - \mathbb{E}[\psi]) \tag{1}$$

which is similar to derivations in [10], where the left side becomes 0 as expectations stabilize. We will be using this equation repeatedly by defining a series of time-invariant test functions appropriate for our analysis. For more details, the reader is encouraged to look at references [27] and [10].

III. Age Scaling in the Presence of an Adversary

Note that packets arriving at infected node n from a node $i \in \mathcal{N}_R$ with rate $\frac{\lambda}{n-1}$ Poisson process are accepted (or discarded) with probability $1 - p$ (or p) when the adversary changes timestamp of incoming packet to t (or 0) to make it appear fresh (or stale). This is equivalent to packets arriving at node n from node i with thinned Poisson process with rate $\lambda_{in} = \frac{(1-p)\lambda}{n-1}$ such that these packets are always accepted. The remaining packets are always discarded and have no effect on age dynamics of the system. Similarly, as the outgoing packets from the infected node n are accepted at node $i \in \mathcal{N}_R$ with probability p, this is equivalent to node n sending packets with timestamp t to node i with a thinned Poisson process of rate $\lambda_{ni} = \frac{p\lambda}{n-1}$ such that these packets are always accepted.

Therefore, based on transition (i, j) at time t, the reset map $\phi_{i,j}(\boldsymbol{X}, \boldsymbol{U}, t) = [X'_1, \ldots, X'_n, U'_1, \ldots, U'_n]$ can be described by

$$U'_\ell = \begin{cases} t, & i = 0, j \in \mathcal{N}, \ell = j \\ \max\{U_i, U_\ell\}, & i, j \in \mathcal{N}_R, \ell = j \\ t, & i = n, j \in \mathcal{N}_R, \ell = j \\ t, & i \in \mathcal{N}_R, j = n, \ell = j \\ U_\ell, & \text{otherwise} \end{cases} \tag{2}$$

and

$$X'_\ell = \begin{cases} 0, & i = 0, j \in \mathcal{N}, \ell = j \\ X_{N(\{i,\ell\})}, & i, j \in \mathcal{N}_R, \ell = j \\ X_n, & i = n, j \in \mathcal{N}_R, \ell = j \\ X_i, & i \in \mathcal{N}_R, j = n, \ell = j \\ X_\ell, & \text{otherwise} \end{cases} \tag{3}$$

Here, $X_{N(S)} = X_{\arg\max_{j\in S} U_j}$ for state $(\boldsymbol{X}, \boldsymbol{U})$ and a subset of nodes S. Since all regular nodes have statistically similar age processes, every arbitrary set S_k of k regular nodes will have the same expected age v_{S_k}, $S_k \subseteq \mathcal{N}_R$, with $v_{S_1} = v_1$. We pick our first test function to be $\psi(\boldsymbol{X}, \boldsymbol{U}) = X_{N(S_k)}$, which is modified upon transition (i, j) to $\psi(\phi_{i,j}(\boldsymbol{X}, \boldsymbol{U}, t)) = X'_{N(S_k)}$. This in turn is characterized using (2) and (3) as

$$X'_{N(S_k)} = \begin{cases} 0, & i = 0, j \in S_k \\ X_{N(S_k\cup\{i\})}, & i \in \mathcal{N}_R \backslash S_k, j \in S_k \\ X_n, & i = n, j \in S_k \\ X_{N(S_k)}, & \text{otherwise} \end{cases} \tag{4}$$

Noting that λ_{ij} is $\frac{\lambda}{n}$ when $i = 0$, and it is $\frac{\lambda}{n-1}$ when $i, j \in \mathcal{N}_R$ and considering the thinned Poisson processes related to node n, using (1), this test function yields,

$$\begin{aligned} 0 =& 1 + \frac{k\lambda}{n}(0 - v_{S_k}) + \frac{(n-k-1)k\lambda}{n-1}(v_{S_{k+1}} - v_{S_k}) \\ &+ \frac{kp\lambda}{n-1}(v_n - v_{S_k}) \end{aligned} \tag{5}$$

which upon rearrangment gives

$$v_{S_k} = \frac{\frac{1}{k\lambda} + \frac{n-k-1}{n-1} v_{S_{k+1}} + \frac{pv_n}{n-1}}{\frac{1}{n} + \frac{n-k-1}{n-1} + \frac{p}{n-1}} \tag{6}$$

Our second test function is simply $\psi(\boldsymbol{X}, \boldsymbol{U}) = X_n$, i.e., the age at infected node, such that its (i, j) transition map is

$$X'_n = \begin{cases} 0, & i = 0, j = n \\ X_i, & i \in \mathcal{N}_R, j = n \\ X_n, & \text{otherwise} \end{cases} \tag{7}$$

which, upon proceeding similarly to (5) and (6), gives

$$v_n = \frac{\frac{1}{\lambda} + (1-p)v_1}{\frac{1}{n} + (1-p)} \tag{8}$$

Our goal is to obtain an analytical expression for expected age of a regular node $v_{S_1} = v_1$, by making use of (6) and (8).

A. Case 1: $p = 1$

In this case, the adversary blocks all incoming packets from the regular nodes, and misleads them into accepting all packets sent by infected node n through timestamp manipulation. Let $y_k = v_{S_k}\frac{n-k}{n-1}$ and using $\frac{1}{n-1} \approx \frac{1}{n}$ for large n, (6) becomes

$$y_k = \frac{n-k}{n-k+1}\left(y_{k+1} + \frac{1}{k\lambda} + \frac{v_n}{n-1}\right) \tag{9}$$

Starting from $y_1 = v_1$, and successively substituting for $y_2, y_3, \ldots, y_{n-1}$, we obtain

$$v_1 = \frac{1}{\lambda}\sum_{k=1}^{n-1}\frac{n-k}{nk} + v_n\sum_{k=1}^{n-1}\frac{n-k}{n(n-1)} \tag{10}$$

$$= \frac{1}{\lambda}\sum_{k=1}^{n-1}\frac{1}{k} - \frac{1}{\lambda}\frac{n-1}{n} + \frac{v_n}{n(n-1)}\sum_{k=1}^{n-1}k \tag{11}$$

$$= O(\log n) + \frac{v_n}{2} \tag{12}$$

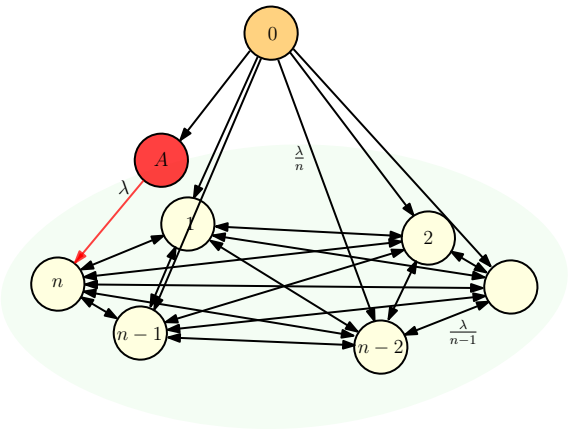

Fig. 2. MITM attack on fully connected network of n nodes.

since $\sum_{k=1}^{n-1} \frac{1}{k}$ grows asymptotically as $\log n$ and $\frac{n-1}{n} \approx 1$. The v_n in the second term can in turn be obtained by substituting $p = 1$ in (8), giving $v_n = \frac{n}{\lambda}$. Hence,

$$v_1 = O(\log n) + \frac{n}{2\lambda} = O(n) \tag{13}$$

To put this deterioration in age scaling into perspective, remember that in a fully disconnected network with no gossiping [12], expected age at each node also scales as $O(n)$, to be exact, $\frac{n}{\lambda}$, a fact that will come handy later.

B. Case 2: p=0

In this case, the infected node accepts all files from the $n-1$ regular nodes but does not transmit any files, even when it possesses the latest file, thereby limiting its contribution, positive or negative, to the system age.

Substituting $p = 0$ in (6) and assuming $\frac{1}{n-1} \approx \frac{1}{n}$ for large n gives $y_k = \frac{1}{k\lambda} + y_{k+1}$ which, upon solving iteratively for $k = \{1, \ldots, n-1\}$, gives $v_1 = \frac{1}{\lambda} \sum_{k=1}^{n-1} \frac{1}{k} = O(\log n)$. Hence, for large n, such an adversary has negligible effect on the network age. In addition, putting $p = 0$ in (8) gives

$$v_n = \frac{\frac{1}{\lambda} + v_1}{\frac{1}{n} + 1} \approx \frac{1}{\lambda} + v_1 \tag{14}$$

Hence v_n scales as $O(\log n)$ similar to v_1, which indicates that the degradation in age caused by adversary upon increasing timestamps of incoming files is negligible.

C. Case 3: $0 < p < 1$

In this case, the adversary partially allows node n to receive incoming files from the gossip network. In (6), plugging $p = 1$ in denominator gives lower bound $\frac{n-k}{n-k+1}\left(y_{k+1} + \frac{1}{k\lambda} + \frac{pv_n}{n-1}\right) < y_k$ and plugging $p = 0$ in denominator gives upper bound $y_k < y_{k+1} + \frac{1}{k\lambda} + \frac{pv_n}{n-1}$, and together with techniques employed in cases 1 and 2, we can bound v_1 as

$$O(\log n) + \frac{pv_n}{2} < v_1 < O(\log n) + pv_n \tag{15}$$

Clearly how age scales at the infected node dictates the age scaling for the regular nodes in the rest of the network. For a fixed $p < 1$, choosing $n \gg \frac{1}{1-p}$ can result in $v_n \approx \frac{\frac{1}{\lambda} + (1-p)v_1}{(1-p)} = O(1) + v_1$, which, when combined with (15) yields $O(\log n)$ age scaling for both v_1 and v_n. Hence, if the infected node is allowed to accept a small fraction of incoming packets from the network, then a large network can manage to curb the spread of stale files coming from the infected node by sustaining a low age at all nodes.[1]

IV. MITM Attack on Fully Connected Network

In previous sections, the adversarial node was in direct contact with all other nodes due to fully connected nature of the network, and the adversary could raise the system age to $O(n)$ with $p = 1$. Here, an interesting question to ask is if the network could do better if the adversary instead had access to only one node. To this end, we consider the network model of Fig. 2, where the adversary, which we will refer to as node A, intercepts the updates to node n coming from the source. In turn the adversary sends updates with rate λ, after changing the timestamps of every outgoing packet to current time, only to node n.

Clearly the expected age at the adversary, denoted by v_A, scales as $O(n)$ since it is isolated from the gossip network and only receives updates from the source with rate $\frac{\lambda}{n}$. The two reset maps useful for our analysis are

$$X'_{N(S_k \cup \{n\})} = \begin{cases} 0, & i = 0, j \in S_k \\ X_{N(S_{k+1} \cup \{n\})}, & i \in \mathcal{N}_R \backslash S_k, j \in S_k \cup \{n\} \\ X_A, & i = A, j = n \\ X_{N(S_k \cup \{n\})}, & \text{otherwise} \end{cases} \tag{16}$$

and

$$X'_{N(S_k)} = \begin{cases} 0, & i = 0, j \in S_k \\ X_{N(S_{k+1})}, & i \in \mathcal{N}_R \backslash S_k, j \in S_k \\ X_{N(S_k \cup \{n\})}, & i = n, j \in S_k \\ X_{N(S_k)}, & \text{otherwise} \end{cases} \tag{17}$$

We claim $v_{S_k \cup \{n\}} \geq \frac{v_A}{2}$, a loose lowerbound that is trivially verified with induction as follows. Invoking (1) regarding (16) for $k = n - 1$ results in

$$v_{S_{n-1} \cup \{n\}} = \frac{\frac{1}{\lambda} + v_A}{\frac{n-1}{n} + 1} \geq O(1) + \frac{v_A}{2} \geq \frac{v_A}{2} \tag{18}$$

which verifies the claim for $k = n - 1$. Next, we assume the claim holds for $k + 1$, i.e., $v_{S_{k+1} \cup \{n\}} \geq \frac{v_A}{2}$, and verify for k. Invoking (1) regarding (16) for $k \leq n - 2$ and using $\frac{1}{\lambda} > 0$ in the numerator and $\frac{k}{n} \leq 1$ in the denominator gives

$$v_{S_k \cup \{n\}} = \frac{\frac{1}{\lambda} + \frac{(k+1)(n-1-k)}{n-1} v_{S_{k+1} \cup \{n\}} + v_A}{\frac{k}{n} + \frac{(k+1)(n-1-k)}{n-1} + 1} \tag{19}$$

[1]More generally, the adversary can increase timestamps of outgoing and incoming packets with probability p and $1 - q$, respectively, which changes (8) to $v_n = \frac{\frac{1}{\lambda} + (1-q)v_1}{\frac{1}{n} + (1-q)}$. Nevertheless, if $q = 1$, similar to Case 1, $v_n = \frac{n}{\lambda}$, and since (15) remains unchanged, we get again $v_1 = O(n)$. Likewise, when $q < 1$, similar to Case 3, choosing $n \gg \frac{1}{1-q}$ again brings down age at all nodes to $O(n)$. Hence, the adversary can best worsen the age to $O(n)$ by not allowing incoming packets to infected node.

$$\geq \frac{\frac{(k+1)(n-1-k)}{n-1} v_{S_{k+1}\cup\{n\}}}{\frac{(k+1)(n-1-k)}{n-1}+2} + \frac{v_A}{\frac{(k+1)(n-1-k)}{n-1}+2} \quad (20)$$

$$\geq \frac{\frac{(k+1)(n-1-k)}{n-1}\frac{v_A}{2}}{\frac{(k+1)(n-1-k)}{n-1}+2} + \frac{\frac{2v_A}{2}}{\frac{(k+1)(n-1-k)}{n-1}+2} \quad (21)$$

$$= \frac{v_A}{2} \quad (22)$$

Finally, we re-invoke (1) for (17) which results in

$$v_{S_k} = \frac{\frac{1}{k\lambda} + \frac{n-k-1}{n-1} v_{S_{k+1}} + \frac{v_{S_k\cup\{n\}}}{n-1}}{\frac{1}{n} + \frac{n-k-1}{n-1} + \frac{1}{n-1}} \quad (23)$$

Let $y_k = v_{S_k}\frac{n-k}{n-1}$, using $\frac{1}{n-1} \approx \frac{1}{n}$ for large n, (23) becomes

$$y_k = \frac{n-k}{n-k+1}\left(y_{k+1} + \frac{1}{k\lambda} + \frac{v_{S_k\cup n}}{n-1}\right) \quad (24)$$

$$\geq \frac{n-k}{n-k+1} y_{k+1} + \frac{(n-k)v_{S_k\cup\{n\}}}{(n-k+1)(n-1)} \quad (25)$$

Starting from $y_1 = v_1$, we successively substitute for $y_2, y_3, \ldots, y_{n-1}$ and use $v_{S_k\cup\{n\}} \geq \frac{v_A}{2}$ to obtain

$$v_1 \geq \frac{1}{n(n-1)} \sum_{k=1}^{n-1} (n-k) v_{S_k\cup n} \quad (26)$$

$$\geq \frac{v_A}{2n(n-1)} \sum_{k=1}^{n-1} (n-k) = \frac{v_A}{4} \quad (27)$$

Hence, v_1 scales at least as $O(n)$ for all regular nodes. This result is far from intuitive, for it brings home the point how an adversary, with so little an effort as sending tampered packets to just one node, can bring down the freshness of an entire large gossip network.

V. Numerical Results

We simulate a fully connected network of size n and allow it to gossip for a total time of $1000n$, choosing $\lambda = 1$.

Fig. 3 shows the expected age at a regular node and the infected node for all three cases of p of Section III. Focusing on the case $p = 1$ shown in red color, the age at the infected node v_n grows as $\frac{n}{\lambda}$ and the age at a regular node v_1 grows as $\frac{v_n}{2} = \frac{n}{2\lambda}$, as was analytically suggested in (12) and (13). On the other extreme, $p = 0$ gives logarithmic age scaling at all nodes, with the infected node age v_n just slightly above the regular node age v_1, in accordance with (14). In the third case of $p = 0.99$, which allows the infected node to accept 1% of incoming gossip, we observe that the infected node age v_n initially begins to grow linearly but later starts to scale logarithmically for larger values of n as n becomes $n \gg \frac{1}{1-p}$. These results imply that the best course of action for the adversary should be to block all incoming traffic and actively send out outdated timestomped packets.

Fig. 4 shows the expected age at different types of nodes when the adversary is positioned between the source and a node. The red line shows the lower bound $\frac{v_A}{4}$ of (27), where the age at a regular node v_1 lies above this lower bound. Adversary age v_A grows as $O(n)$ by virtue of being an isolated

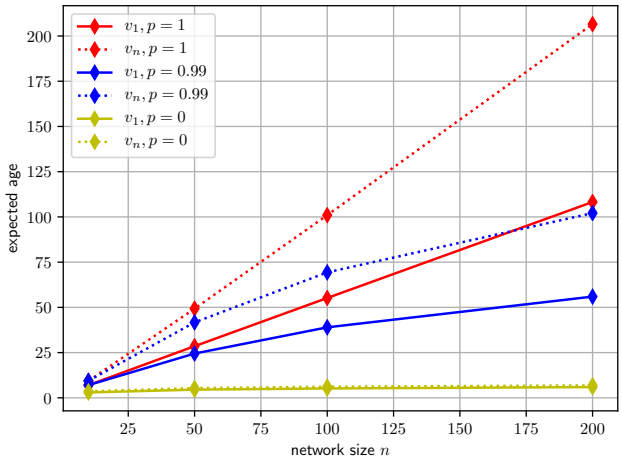

Fig. 3. Node capture attack on fully connected network of n nodes.

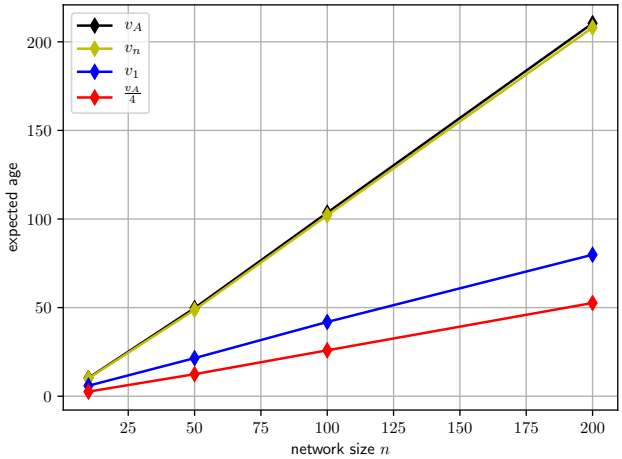

Fig. 4. MITM attack on fully connected network of n nodes.

node. Finally, though (19) yields a loose lower bound of $\frac{v_A}{2}$, the graph shows that the age at the node that is in contact with adversary, v_n, closely follows adversary age v_A.

VI. Conclusion

We studied the effects of timestomping attacks on the age of gossip in a large fully connected network. We showed that one infected node in such a network can increase the age at all other nodes from $O(\log n)$ to $O(n)$ through timestamp manipulation. Further, we showed that the optimal behavior for the adversary is to reset the timestamps of all outgoing packets to current time thereby disguising them as current packets and of all incoming packets to an outdated time to prevent their acceptance at the infected node. Additionally, we showed that if the adversary allows the infected node to accept even a very small fraction of the incoming packets from the network, then a large network can manage to curb the spread of stale files coming from the infected node and pull the network age back to $O(\log n)$. Lastly, we showed that if an infected node contacts only a single node instead of all nodes of the network, the system age can still be degraded to $O(n)$.

REFERENCES

[1] A. J. Demers, D. H. Greene, C. H. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart, and D. B. Terry. Epidemic algorithms for replicated database maintenance. In *ACM PODC*, August 1987.

[2] Y. Minsky. *Spreading Rumors Cheaply, Quickly, and Reliably*. PhD thesis, Cornell University, March 2002.

[3] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking. Randomized rumor spreading. In *FOCS*, November 2000.

[4] B. G. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, February 1987.

[5] S. Deb, M. Medard, and C. Choute. Algebraic gossip: a network coding approach to optimal multiple rumor mongering. *IEEE Transactions on Information Theory*, 52(6):2486–2507, June 2006.

[6] D. Mosk-Aoyama and D. Shah. Information dissemination via network coding. In *IEEE ISIT*, July 2006.

[7] S. Sanghavi, B. Hajek, and L. Massoulie. Gossiping with multiple messages. *IEEE Transactions on Information Theory*, 53(12):4640–4654, December 2007.

[8] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: Amazon's highly available key-value store. *ACM SIGOPS Oper. Syst. Rev.*, 41(6):205–220, 2007.

[9] A. Lakshman and P. Malik. Cassandra: A decentralized structured storage system. *ACM SIGOPS Oper. Syst. Rev.*, 44(2):35–40, 2010.

[10] R. D. Yates. Timely gossip. In *IEEE SPAWC*, September 2021.

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

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

[13] M. Bastopcu, B. Buyukates, and S. Ulukus. Gossiping with binary freshness metric. In *IEEE Globecom*, December 2021.

[14] P. Kaswan and S. Ulukus. Timely gossiping with file slicing and network coding. In *IEEE ISIT*, June 2022.

[15] G. D. Nguyen, S. Kompella, C. Kam, J. E. Wieselthier, and A. Ephremides. Impact of hostile interference on information freshness: A game approach. In *IEEE WiOpt*, May 2017.

[16] A. Garnaev, W. Zhang, J. Zhong, and R. D. Yates. Maintaining information freshness under jamming. In *IEEE Infocom*, May 2019.

[17] Y. Xiao and Y. Sun. A dynamic jamming game for real-time status update. In *IEEE Infocom*, April 2018.

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

[19] S. Banerjee and S. Ulukus. Age of information in the presence of an adversary. In *IEEE Infocom*, May 2022.

[20] S. Banerjee and S. Ulukus. Game theoretic analysis of an adversarial status updating system. In *IEEE ISIT*, June 2022.

[21] J. Augustine, C. Avin, M. Liaee, G. Pandurangan, and R. Rajaraman. Information spreading in dynamic networks under oblivious adversaries. In *DISC*, 2016.

[22] C. Georgiou, S. Gilbert, R. Guerraoui, and D. R. Kowalski. On the complexity of asynchronous gossip. In *PODC*, 2008.

[23] W. Minnaard. Timestomping NFTS. Master's thesis, University of Amsterdam, July 2014.

[24] A. Kosta, N. Pappas, and V. Angelakis. Age of information: A new concept, metric, and tool. *Foundations and Trends in Networking*, 12(3):162–259, November 2017.

[25] 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.

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

[27] J. Hespanha. Modeling and analysis of stochastic hybrid systems. *IEE Proc. Control Theory & Applications, Special Issue on Hybrid Systems*, 153:520–535, January 2007.