The Age of Gossip in Networks

Roy D. Yates
WINLAB, Department of Electrical and Computer Engineering
Rutgers University
ryates@winlab.rutgers.edu

2021 IEEE International Symposium on Information Theory (ISIT) | 978-1-5386-8209-8/21/$31.00 ©2021 IEEE | DOI: 10.1109/ISIT45174.2021.9517796

***Abstract*—A source node updates its status as a point process and also forwards its updates to a network of observer nodes. Within the network of observers, these updates are forwarded as point processes from node to node. Each node wishes its knowledge of the source to be as timely as possible. In this network, timeliness is measured by a discrete form of age of information: each status change at the source is referred to as a version and the age at a node is how many versions out of date is its most recent update from the source. This work introduces a method for evaluating the average version age at each node in the network when nodes forward updates using a memoryless gossip protocol. This method is then demonstrated by version age analysis for a collection of simple networks. For gossip on a complete graph with symmetric updating rates, it is shown that each node has average age that grows as the logarithm of the network size.**

I. Introduction

Gossip is a popular mechanism to convey status information in a distributed systems and networks. The efficacy of gossip mechanisms for distributed computation [1], [2] and message dissemination [3] is well known. While it is also known that gossip can be inefficient relative to more complex or application-specific algorithms, gossip mechanisms remain an attractive option when protocols need to be simple or the network connectivity is time-varying [4]. For example, gossip protocols could be a good choice for low latency vehicular safety messaging. And yet, while vehicular message exchange was the early motivation for age of information (AoI) research [5], [6], there has been little (if any) effort to examine AoI for gossip protocols.

In this work, we begin to re-examine gossip from an age-of-information (AoI) perspective [7], [8]. Specifically, a source wishes to share its status update messages with a network of n nodes. These nodes, which can be viewed monitors of the source, employ gossip to randomly forward these update messages amongst themselves in order that all nodes have timely knowledge of the state of the source.

This work extends AoI analysis in a class of status sampling networks, a networking paradigm that is consistent with gossip models in that short messages, representing samples of a node's status update process, are delivered as point processes to neighbor nodes. This "zero service time" model may be useful in a high speed network in which updates represent small amounts of protocol information (requiring negligible time for transmission) that are given priority over data traffic. This model has also been widely used in the age analysis of energy harvesting updaters [9]–[15] where updating rates are constrained by energy rather than bandwidth. While the transmission of a single update may be negligible, the update rates are limited so that protocol information in the aggregate does not consume an excessive fraction of network resources.

Prior work on status sampling networks [16], [17] developed tools for analyzing age in line networks in which each node i only received updates from node $i-1$. The key advance of this work is the development of an average age analysis method for monitors that receive updates via multiple network paths.

This work was supported by NSF award CCF-1717041.

II. System Model and Summary of Results

Status updates of a source node 0 are shared via a network with a set of nodes $\mathcal{N} = \{1, 2, \ldots, n\}$. Motivated by sensor networks in which accurate clocks may be unavailable, timeliness at each node is measured by *update versions*. The source node 0 maintains the current (fresh) version of its status and thus node 0 always has version age $X_0(t) = 0$. Starting at time $t = 0$, status updates at the source node 0 occur as a rate λ_{00} Poisson process $N_0(t)$. That is, at time $t > 0$, the most recent update at the source is version $N_0(t)$. If the current update at node i is version $N_i(t)$, then the age at node i, as measured in versions, is $X_i(t) = N_0(t) - N_i(t)$. An example of version age sample paths is depicted in Figure 1. In particular, if the source has an update at time t, the age at each node i becomes $X_i'(t) = X_i(t) + 1$.

In this work, we develop a method for evaluating the limiting average age $\lim_{t\to\infty} \mathrm{E}[X_i(t)]$, which we refer to as the version AoI at node i. Building on prior work [17], [18], this paper employs the methodology of the stochastic hybrid system (SHS) to analyze the convergence of the expected age.

Specifically, we assume the nodes forward updates using gossip. Node i sends its most recent update to node j as a rate λ_{ij} Poisson process. If node i sends its update to node j at time t, the age at node j becomes

$$X_j'(t) = \min[X_i(t), X_j(t)]. \tag{1}$$

Implicit in (1) is that updates are version-stamped so that node j only adopts fresher updates from node i.

The SHS approach is to develop a set of ordinary differential equations for $\mathrm{E}[X_i(t)]$ that enables the evaluation of the limiting age $\lim_{t\to\infty} \mathrm{E}[X_i(t)]$. As we see in (1), this will require the characterization of age variables such as $X_{\{i,j\}}(t) \equiv \min(X_i(t), X_j(t))$. More generally, for arbitrary subsets $S \subseteq \mathcal{N}$, the analysis will need to track the age

$$X_S(t) \equiv \min_{j\in S} X_j(t) \tag{2a}$$

and its expected value

$$v_S(t) \equiv \mathrm{E}[X_S(t)]. \tag{2b}$$

We can interpret $X_S(t)$ as the status age of an observer of updates arriving at any node in S and we may refer to $X_S(t)$ as the (version) age of subset S.

The main result of the paper is the development of a system of linear equations for the calculation of the limiting average age $\bar{v}_S = \lim_{t\to\infty} \mathrm{E}[X_S(t)]$. To describe this system of equations, define the update rate of node i into set S as

$$\lambda_i(S) \equiv \begin{cases} \sum_{j\in S} \lambda_{ij} & i \notin S, \\ 0 & i \in S, \end{cases} \tag{3}$$

and the set of updating neighbors of S as

$$N(S) \equiv \{i \in \mathcal{N} \colon \lambda_i(S) > 0\}. \tag{4}$$

With this notation, we state our main result.

Theorem 1: The expected status age $v_S(t) = \mathrm{E}[X_S(t)]$ of an observer of node set S converges to $\bar{v}_S = \lim_{t\to\infty} v_S(t)$ satisfying

$$\bar{v}_S = \frac{\lambda_{00} + \sum_{i\in N(S)} \lambda_i(S)\bar{v}_{S\cup\{i\}}}{\lambda_0(S) + \sum_{i\in N(S)} \lambda_i(S)}. \tag{5}$$

Proof of this claim is deferred to Section V-B.

In Section IV, we demonstrate the use of Theorem 1 first for the $n = 3$ node network in Figure 2 and second for the n node symmetric gossip network on a complete graph, as depicted in Figure 3 for $n = 6$ nodes. In the complete graph, $\lambda_{ij} = \lambda/(n-1)$ for all node pairs $i, j \in \mathcal{N}$. This corresponds to each node $i \in \mathcal{N}$ randomly sending its current updates to each of the other $n-1$ nodes as a rate $\lambda/(n-1)$ Poisson process. In addition, the source sends symmetrically to each node $j \in \mathcal{N}$ with Poisson rate $\lambda_{0j} = \lambda/n$. By exploiting the symmetry of the complete graph, Theorem 1 shows that the average age at a node grows as $\log n$.

Theorem 2: For the symmetric complete gossip network with the source sending updates to each node $i \in \mathcal{N}$ at rate λ/n, the average version age of each node i is

$$\frac{\lambda_{00}}{\lambda}\left[\frac{n-1}{n}\sum_{k=1}^{n-1}\frac{1}{k} + \frac{1}{n}\right] \le \lim_{t\to\infty} \mathrm{E}[X_i(t)] \le \frac{\lambda_{00}}{\lambda}\sum_{k=1}^{n}\frac{1}{k}. \tag{6}$$

Hence, as the network size n grows, the average age at each node only grows logarithmically in n. Although the communication models are different in various small ways, this average result is analogous to [3, Theorem 3.1] in which the ϵ-dissemination time, i.e. the time until the probability a source message has not reached all nodes is less than ϵ, is shown to grow as $O(\log n)$.

III. Related Work

AoI analysis started with single-source single-server first-come first-served (FCFS) queues [7], the M/M/1 last-come first-served (LCFS) queue with preemption in service [19], and the M/M/1 FCFS system with multiple sources [20]. An overview of AoI research can be found in [8]; here we focus on the most closely related contributions.

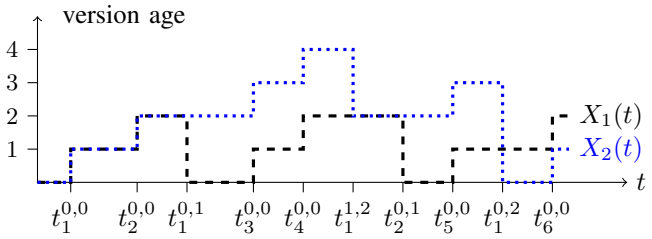

Fig. 1. Fresh updates from a source pass through the network as point processes; $t_n^{i,j}$ marks the nth update sent on link (i, j). Node 1 gets updates from the source node 0. Node 2 gets updates from both the source and also from node 1. Age is measured in versions, $X_i(t)$ records how many versions out-of-date the update at node i is relative to the source.

To evaluate AoI for a single source sending updates through a network cloud [21] or through an M/M/m server [22]–[24], out-of-order packet delivery was the key analytical challenge. The first evaluation of the average AoI over multihop network routes [25] employed a discrete-time version of the status sampling network also employed in [16], [17]. These works obtained simple AoI results because the updates followed a single path to a destination monitor. This avoided the complexity of multiple paths and the consequent accounting for repeated and out-of-order update message deliveries.

When multiple sources employ wireless networks subject to interference constraints, AoI has been analyzed under a variety of link scheduling strategies [26]–[34]. Age bounds were developed from graph connectivity properties [35] when each node needs to update every other node. For DSRC-based vehicular networks, update piggybacking strategies were developed and evaluated [6].

When update transmission times over network links are exponentially distributed, sample path arguments were used [36]–[38] to show that a preemptive Last-Generated, First-Served (LGFS) policy results in smaller age processes at all nodes of the network than any other causal policy. Note that the status sampling network model in this work can also be viewed as a network of preemptive LGFS server; see [17] for details. With that equivalence, [38] and this work can be viewed as complementary in that [38] proves the age-optimality of LGFS policies and this work provides analytic tools for the evaluation of those policies.

While these prior contributions used sawtooth age functions, version age in this work is related to various discretized age metrics. Notably, version age is similar to the age of incorrect information [39] and age of synchronization [40] metrics. For all three metrics, the age at a node is zero as long as that node has the current status of the source. Both this work and [41] employ version age to track a Poisson updating process, a problem introduced in [42]. Another interpretation of version age is that zero age is fresh, and nonzero age is stale. This is the binary freshness metric found in [43]–[45].

IV. Applications of Theorem 1

To utilize Theorem 1, suppose we wish to calculate the average age at node n. We start with $S = \{n\}$ and generate an equation for $\bar{v}_{\{n\}}$ in terms of the variables $\bar{v}_{\{i,n\}}$ for nodes i such that $\lambda_{i,n} > 0$. For each such node i, the next step is to apply (5) recursively with $S = \{i, n\}$. This generates an

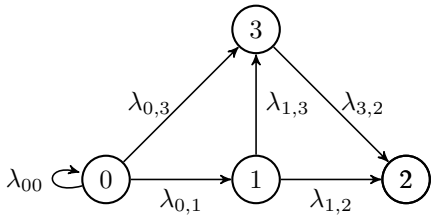

Fig. 2. Updates generated at node 0 are forwarded to nodes 1, 2, and 3.

equation for each $\bar{v}_{\{i,n\}}$ in terms of variables $\bar{v}_{i,j,n}$ for each node j that sends updates to one or both nodes in $\{i,n\}$.

In general, at stage k, we construct equations for $\bar{v}_S$ for sets S with size $|S| = k$ in terms of variables $\bar{v}_{S'}$ such that each S' has size $|S'| = k+1$. In the worst case, this procedure terminates at stage $k = n$ when $S = \mathcal{N}$. For a fully connected graph, this procedure generates equations for all $2^n - 1$ non-empty subsets of $\mathcal{N}$. On the other hand, when the network graph is sparse, substantially fewer equations may be generated.

In the next three sections, we demonstrate Theorem 1 with three examples; a three-node toy network with arbitrary rates, version age analysis of the n-node symmetric complete graph that provides the proof of Theorem 2, and an n-node symmetric ring network.

A. Toy example of Theorem 1

Here we demonstrate Theorem 1 by solving for the average version age $\bar{v}_{\{2\}}$ at node 2 for the network shown in Figure 2. The recursive application of (5) with $S = \{2\}$, $S = \{1,2\}$, $S = \{2,3\}$ and $S = \{1,2,3\}$ yields

$$\bar{v}_{\{2\}} = \frac{\lambda_{0,0} + \lambda_{1,2}\bar{v}_{\{1,2\}} + \lambda_{3,2}\bar{v}_{\{2,3\}}}{\lambda_{1,2} + \lambda_{3,2}}, \tag{7a}$$

$$\bar{v}_{\{1,2\}} = \frac{\lambda_{0,0} + \lambda_{3,2}\bar{v}_{\{1,2,3\}}}{\lambda_{0,1} + \lambda_{3,2}}, \tag{7b}$$

$$\bar{v}_{\{2,3\}} = \frac{\lambda_{0,0} + (\lambda_{1,2} + \lambda_{1,3})\bar{v}_{\{1,2,3\}}}{\lambda_{0,3} + \lambda_{1,2} + \lambda_{1,3}}, \tag{7c}$$

$$\bar{v}_{\{1,2,3\}} = \frac{\lambda_{0,0}}{\lambda_{0,1} + \lambda_{0,3}}. \tag{7d}$$

We note that (7d) is an example of the general result that $\bar{v}_{\mathcal{N}} = \lambda_{00}/\lambda_0(\mathcal{N})$. For this n, it follows from (7) that

$$\bar{v}_{\{2\}} = \frac{\lambda_{0,0}}{\lambda_{1,2} + \lambda_{3,2}} \left[1 + \frac{\lambda_{1,2}}{\lambda_{0,1} + \lambda_{3,2}} \left(1 + \frac{\lambda_{3,2}}{\lambda_{0,1} + \lambda_{0,3}}\right) + \frac{\lambda_{3,2}}{\lambda_{0,3} + \lambda_{1,2} + \lambda_{1,3}} \left(1 + \frac{\lambda_{1,2} + \lambda_{1,3}}{\lambda_{0,1} + \lambda_{0,3}}\right)\right]. \tag{8}$$

The solution (8) is complicated because it includes a variety of special cases. For example, when $\lambda_{1,2} \to \infty$, $\bar{v}_{\{1,2,3\}}$ and $\bar{v}_{\{1,2\}}$ are unchanged but $\bar{v}_{\{2\}} \to \bar{v}_{\{1,2\}}$ because nodes 1 and 2 become equivalent to a single node with update rates $\lambda_{0,1}$ from node 0 and $\lambda_{3,2}$ from node 3. On the other hand, when $\lambda_{1,2} \to 0$, $\bar{v}_{\{1,2,3\}}$ is unchanged while

$$\bar{v}_{\{2\}} \to \frac{\lambda_{0,0}}{\lambda_{3,2}} + \bar{v}_{\{2,3\}}, \quad \bar{v}_{\{2,3\}} \to \frac{\lambda_{0,0} + \lambda_{1,3}\bar{v}_{\{1,2,3\}}}{\lambda_{0,3} + \lambda_{1,3}}. \tag{9}$$

In this case, the solution for $\bar{v}_{\{2\}}$ reflects the path diversity offered by the two paths from the source to node 2.

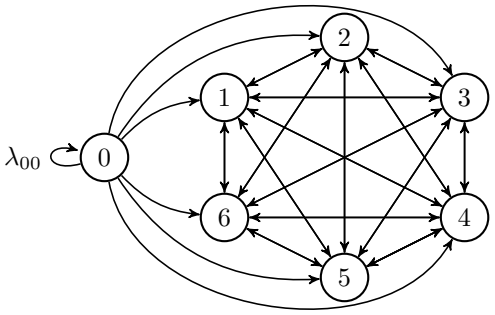

Fig. 3. Updates generated at node 0 at rate λ_{00} are forwarded to nodes in $\mathcal{N} = \{1,\ldots,6\}$ which form a complete graph. Node 0 sends updates to each node $i \in \mathcal{N}$ at rate $\lambda/6$. Each node $i \in \mathcal{N}$ send updates to every other node j at rate $\lambda/5$.

B. Proof of Theorem 2

We now use Theorem 1 to find the average version age of a node for the n-node complete graph, as depicted for $n = 6$ in Figure 3. Here the symmetry of the complete graph is essential to derive Theorem 2. In the absence of symmetry, the recursion of Theorem 1 would generate equations for all $2^n - 1$ nontrivial subsets of $\mathcal{N}$.

Let S_j denote an arbitrary j-node subset of the complete graph. By symmetry, the age processes $X_{S_j}(t)$ for all subsets S_j are statistically identical. Hence we define $\tilde{v}_j = \bar{v}_{S_j}$. Moreover, each subset S_j has $|N(S_j)| = n - j$ neighbor nodes i that send updates to S_j at rate $\lambda_i(S_j) = j\lambda/(n-1)$. For each such neighbor i, $S_j \cup \{i\}$ is a $j+1$ node subset S_{j+1}. Also, because the source symmetrically updates all nodes in $\mathcal{N}$, each subset S_j receives updates from the source node at rate $\lambda_0(S_j) = j\lambda/n$. Thus Theorem 1 yields

$$\tilde{v}_j = \frac{\lambda_{00} + |N(S_j)|\lambda_i(S_j)\tilde{v}_{j+1}}{\lambda_0(S_j) + N(S_j)\lambda_i(S_j)} = \frac{\lambda_{00} + \frac{j(n-j)\lambda}{n-1}\tilde{v}_{j+1}}{\frac{j\lambda}{n} + \frac{j(n-j)\lambda}{n-1}}. \tag{10}$$

For $j = n$, $S_j = S_n$ is the set of all nodes. With all nodes in the observer set, the neighbor set $N(\mathcal{N})$ is empty, $\lambda_0(\mathcal{N}) = \lambda$, and Theorem 1 yields $\tilde{v}_n = \bar{v}_{\mathcal{N}} = \lambda_{00}/\lambda$. With this initial condition, (10) enables iterative computation of $\tilde{v}_{n-1}, \tilde{v}_{n-2}, \ldots$ until we reach $\tilde{v}_1$, the average age of a single node. However to complete the proof, let $j = n - k$, implying

$$\tilde{v}_{n-k} = \frac{\frac{\lambda_{00}}{(n-k)\lambda} + \frac{k}{n-1}\tilde{v}_{n-k+1}}{\frac{1}{n} + \frac{k}{n-1}}. \tag{11}$$

With the definition $\hat{v}_k \equiv \tilde{v}_{n-k+1}$, (11) becomes

$$\hat{v}_{k+1} = \frac{\frac{\lambda_{00}}{(n-k)\lambda} + \frac{k}{n-1}\hat{v}_k}{\frac{1}{n} + \frac{k}{n-1}} \le \frac{\frac{\lambda_{00}}{(n-k)\lambda} + \frac{k}{n}\hat{v}_k}{\frac{1}{n} + \frac{k}{n}}. \tag{12}$$

The upper bound in (12) holds iff $\hat{v}_k \le n\lambda_{00}/(n-k)\lambda$. Since $\hat{v}_1 = \tilde{v}_n = \lambda_{00}/\lambda$ this requirement holds at $k = 1$ and can be shown by induction to hold for all k. Defining $y_k = k\hat{v}_k/n$, it follows from (12) that

$$y_{k+1} \le \frac{\lambda_{00}}{(n-k)\lambda} + y_k. \tag{13}$$

Since $y_1 = \lambda_{00}/(n\lambda)$, it follows from (13) that

$$y_n \le \frac{\lambda_{00}}{\lambda} \sum_{k=0}^{n-1} \frac{1}{n-k} = \frac{\lambda_{00}}{\lambda} \sum_{k=1}^{n} \frac{1}{k}. \tag{14}$$

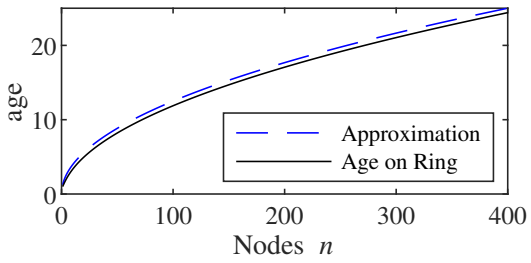

Fig. 4. Average version age of a node on the symmetric n node ring with $\lambda_{00}/\lambda = 1$. The dashed blue line $1.25\sqrt{n}$ is an empirical approximation, but not an upper bound; the age will exceed the approximation for $n > 40401$.

Since $y_n = \hat{v}_n = \tilde{v}_1$, this completes the proof of the Theorem 2 upper bound. For the lower bound, the equality in (12) implies

$$\hat{v}_{k+1} \geq \frac{n-1}{k-1}\left[\frac{\lambda_{00}}{(n-k)\lambda} + \frac{k}{n-1}\hat{v}_k\right]. \tag{15}$$

Defining $\hat{y}_k \equiv k\hat{v}_k/(n-1)$, (15) implies

$$\hat{y}_{k+1} \geq \frac{\lambda_{00}}{(n-k)\lambda} + \hat{y}_k. \tag{16}$$

It follows from (16) that $\hat{v}_n = (n-1)\hat{y}_n/n$ satisfies the lower bound of Theorem 2.

C. Age on a Symmetric Ring

In the ring network, the source sends updates to each node at rate λ/n while each node i sends updates to each of its neighbor nodes[1] $i+1$ and $i-1$ at rate $\lambda/2$. Thus the network resembles the complete graph of Figure 3, except the interior transitions are deleted.

For the ring graph, let $R_j = \{i, i+1, \ldots, i+j-1\}$ denote an arbitrary contiguous j-node subset of the ring. By symmetry, the age processes $X_{R_j}(t)$ for all subsets R_j are statistically identical. Hence we define $\tilde{v}_j = \bar{v}_{R_j}$. Moreover, for $j < n-1$, each subset R_j has $|N(R_j)| = 2$ neighbor nodes $k \in \{i-1, j\}$ that send updates to R_j at rate $\lambda_k(R_j) = \lambda/2$. For each neighbor k, $R_j \cup \{k\}$ is a $j+1$ node subset R_{j+1}. In addition, for a subset R_{n-1}, there is a single neighbor k who sends updates at rate $\lambda_k(R_{n-1}) = \lambda$ (at rate $\lambda/2$ to neighbor nodes $k+1$ and $k-1$ that are the head and tail of R_{n-1}. Thus $|N(R_j)|\lambda_k(R_j) = \lambda$ for $j \in \{1, \ldots, n-1\}$. Also, because the source symmetrically updates all nodes in $\mathcal{N}$, each subset R_j receives updates from the source node at rate $\lambda_0(R_j) = j\lambda/n$. Thus Theorem 1 yields

$$\tilde{v}_j = \frac{\lambda_{00} + |N(R_j)|\lambda_k(R_j)\tilde{v}_{j+1}}{\lambda_0(R_j) + N(R_j)\lambda_k(R_j)} = \frac{\lambda_{00} + \lambda\tilde{v}_{j+1}}{j\lambda/n + \lambda}. \tag{17}$$

For $j = n$, $R_j = R_n$ is the set of all nodes. With all nodes in the observer set, the neighbor set $N(\mathcal{N})$ is empty, $\lambda_0(\mathcal{N}) = \lambda$, and Theorem 1 yields $\tilde{v}_n = \bar{v}_{\mathcal{N}} = \lambda/\lambda_{00}$, as it does for all graphs. With this initial condition, (17) enables iterative computation of $\tilde{v}_{n-1}, \tilde{v}_{n-2}, \ldots$ until we reach $\tilde{v}_1$, the average age of an individual node. As of this writing, the downward iteration (17) has not yet yielded a simple bound for $\tilde{v}_1$. However, as the numerical evaluation is nearly trivial, an age plot is presented in Figure 4 for $\lambda_{00}/\lambda = 1$. From the figure, it is empirically observed that $\Delta_{\text{ring}}(n) \approx 1.25\sqrt{n}$.

This numerical evidence may seem surprising since $O(n \log n)$ dissemination time has been reported for the ring graph [3]. However, to enable age comparisons with the complete graph, the ring model in this work sends its fresh updates randomly to the ring. By contrast, the ring graph model in [3] assumes the source is a node on the ring and thus the dissemination time to all nodes must be $\Omega(n)$. If source updates were passed only to a single node on the ring, the average age would indeed grow as $O(n)$.

[1]We assume node indexing modulo the n node ring, i.e., node $n+1$ refers to node 1 and node $1-1$ refers to node n.

V. Stochastic Hybrid Systems for AoI Analysis

In this section we use a stochastic hybrid system (SHS) model to derive Theorem 1. While there are many SHS variations [46], this work follows [17], [18], which employ the model and notation in [47]. In general, the SHS is described by a discrete state $q(t) \in \mathcal{Q} = \{0, 1, \ldots, q_{\max}\}$ that evolves as a point process, a continuous component $\mathbf{X}(t) \in \mathbb{R}^n$ described by a stochastic differential equation in each state $q \in \mathcal{Q}$, and a set $\mathcal{L}$ of transition/reset maps that correspond to both changes in the discrete state and jumps in the continuous state.

A. Version AoI for gossip networks as an SHS

The operation of the gossip network is memoryless; each node i sends its current update to node j as a Poisson process of rate λ_{ij}. Hence, the SHS discrete state space is the trivial set $\mathcal{Q} = \{0\}$. Furthermore, because age is measured in versions, the normally continuous age state $\mathbf{X}(t)$ in fact becomes discrete. That is, $\mathbf{X}(t)$ changes only when there is a transition that corresponds to an update being forwarded. In the absence of such a transition, the stochastic differential equation of the SHS is trivially $\dot{\mathbf{X}}(t) = \mathbf{0}$.

The remaining component of the SHS model is the set $\mathcal{L}$ of discrete transition/reset maps. In the gossip network, $\mathcal{L}$ corresponds to the set of directed edges (i, j) over which node i updates node j. However, because of the special role of node 0 as the source, there are three kinds of transitions. First, $(i, j) = (0, 0)$ corresponds to the source node generating a new version so that the version age at all other nodes k increases by one. The second type of transition is given by $(0, j)$, corresponding to the source node 0 sending the current version to node j, reducing the age at node j to zero. In the third type, a gossiping node i forwards its current update to node j; node j accepts the update if it is a fresher than its existing version. To summarize, the set of transitions is

$$\mathcal{L} = \{(0,0)\} \cup \{(0,j) \colon j \in \mathcal{N}\} \cup \{(i,j) \colon i, j \in \mathcal{N}\}, \tag{18}$$

transition (i, j) occurs at rate $\lambda_{i,j}$, and in that transition the age vector becomes $\phi_{i,j}(\mathbf{X}) = \begin{bmatrix} X_1' & \cdots & X_n' \end{bmatrix}$ such that

$$X_k' = \begin{cases} X_k + 1 & i = 0, j = 0, k \in \mathcal{N}, \\ 0 & i = 0, k = j \in \mathcal{N}, \\ \min(X_i, X_j) & i \in \mathcal{N}, k = j \in \mathcal{N}, \\ X_k & \text{otherwise.} \end{cases} \tag{19a}$$

Because of the generality and power of the SHS model, complete characterization of the $\mathbf{X}(t)$ process is often impossible. The approach in [47] is to define test functions $\psi(q, \mathbf{X}, t)$ whose expected values $\mathrm{E}[\psi(q(t), \mathbf{X}(t), t)]$ are performance measures of interest that can be evaluated as functions of time; see [47], [48], and the survey [46] for additional background.

Since the simplified SHS for the gossip network is time invariant and has a trivial discrete state, it is sufficient to employ the time invariant test functions $\psi_S(\mathbf{X}) = X_S$. These test functions yield the processes

$$\psi_S(\mathbf{X}(t)) = X_S(t), \tag{20}$$

which have expected values

$$\mathrm{E}\big[\psi_S(\mathbf{X}(t))\big] = \mathrm{E}\big[X_S(t)\big] \equiv v_S(t). \tag{21}$$

The objective here is to use the SHS framework to derive a system of differential equations for the $v_S(t)$. To do so, the SHS mapping $\psi \to L\psi$ known as the extended generator is applied to every test function $\psi(\mathbf{X})$. The extended generator $L\psi$ is simply the function whose expected value is the expected rate of change of the test function ψ. Specifically, a test function $\psi(\mathbf{X}(t))$ has an extended generator $(L\psi)(\mathbf{X}(t))$ that satisfies Dynkin's formula

$$\frac{d\,\mathrm{E}[\psi(\mathbf{X}(t))]}{dt} = \mathrm{E}[(L\psi)(\mathbf{X}(t))]. \tag{22}$$

For each test function $\psi(\mathbf{X})$, (22) yields a differential equation for $\mathrm{E}[\psi(\mathbf{X}(t))]$.

From [47, Theorem 1], it follows from the trivial discrete state, the trivial stochastic differential equation $\dot{\mathbf{X}}(t) = \mathbf{0}$, and the time invariance of $\psi_S(\mathbf{X})$ in (20) that the extended generator of a piecewise linear SHS is given by

$$(L\psi_S)(\mathbf{X}) = \sum_{(i,j)\in\mathcal{L}} \lambda_{ij}[\psi_S(\phi_{i,j}(\mathbf{X})) - \psi_S(\mathbf{X})]. \tag{23}$$

With these preliminaries, we now prove Theorem 1.

B. Proof of Theorem 1

In (23), it follows from (2a), (19), and (20) that the effect on the test function of transition (i, j) is

$$\psi_S(\phi_{i,j}(\mathbf{X})) = X'_S = \min_{k\in S} X'_k. \tag{24}$$

Evaluation of (24) depends on the transition type (i, j), as given in (19). In transition $(0, 0)$, the source node has a version update and each node $k \in S$ ages by one more version. This implies $X'_k = X_k + 1$ for all $k \in \mathcal{N}$ and thus

$$X'_S = \min_{k\in S} X'_k = X_S + 1. \tag{25}$$

For other transitions (i, j), only the age X_j at node j is changed. Thus if $j \notin S$, then $X_S = \min_{k\in S} X_k$ is unchanged. However, if $j \in S$, then

$$\begin{aligned} X'_S = \min_{k\in S} X'_k &= \min(\min(X_i, X_j), \min_{k\in S\setminus\{j\}} X_k) \\ &= \min_{k\in S\cup\{i\}} X_k = X_{S\cup\{i\}}. \end{aligned} \tag{26}$$

In addition to the common (i, j) transition in which $i \in \mathcal{N}$ is a gossiping neighbor of $j \in S$, we note that (26) incorporates some special cases. If $i = 0$, then $X'_S = X_{S\cup\{0\}} = 0$ since $X_0 = 0$. On the other hand, if $i \in S$, then $S \cup \{i\} = S$ and $X'_S = X_S$. That is, an update sent by a node in S cannot reduce the age X_S.

Based on the three types of transitions, namely $(0, 0)$, $(0, j)$, and (i, j), we conclude that

$$\begin{aligned} (L\psi_S)(\mathbf{X}) = \lambda_{00}(X_S + 1 - X_S) + \sum_{j\in S} \lambda_{0j}[0 - X_S] \\ + \sum_{\substack{i>0\\ i\notin S}} \sum_{j\in S} \lambda_{ij}\big[X_{S\cup\{i\}} - X_S\big]. \end{aligned} \tag{27}$$

We note that $\mathbf{X}$, X_S, and $X_{S\cup\{i\}}$ in (27) refer to the age processes $\mathbf{X}(t)$, $X_S(t)$ and $X_{S\cup\{i\}}(t)$. With this in mind, we take the expectation of (27). On the left side of (27), $\mathrm{E}[(L\psi_S)(\mathbf{X}(t))] = \dot{v}_S(t)$ by Dynkin's formula (22). On the right side, $\mathrm{E}[X_S(t)] = v_S(t)$ and $\mathrm{E}\big[X_{S\cup\{i\}}(t)\big] = v_{S\cup\{i\}}(t)$ for all i. These substitutions yield

$$\dot{v}_S(t) = \lambda_{00} - v_S(t)\sum_{j\in S} \lambda_{0j} + \sum_{\substack{i>0\\ i\notin S}} \sum_{j\in S} \lambda_{ij}[v_{S\cup\{i\}}(t) - v_S(t)].$$

Employing the definitions (3) and (4) of the update rate $\lambda_i(S)$ of node i into S, and the neighbor set $N(S)$, we obtain

$$\dot{v}_S(t) = \lambda_{00} - v_S(t)\Big[\lambda_0(S) + \sum_{i\in N(S)} \lambda_i(S)\Big] + \sum_{i\in N(S)} \lambda_i(S) v_{S\cup\{i\}}(t).$$

By setting the derivatives $\dot{v}_S(t) = 0$, we obtain a linear equation for the time average age $\bar{v}_S = \lim_{t\to\infty} v_S(t)$ in terms of the necessary $\bar{v}_{S\cup\{i\}}$. This yields (5).

VI. Conclusion

This work has introduced AoI analysis tools for gossip algorithms on network graphs. In Theorem 1 we developed a set of linear equations for the computation of average version age at any node in a gossip network described by an arbitrary graph. While the general solution has exponential complexity in the number of nodes, we believe this unavoidably reflects the multiplicity of paths from the source to a node. When this method is applied to the n node complete graph, it was shown using graph symmetry that the average version age at each node grows as $\log n$. This promising result suggests that gossip networks may indeed be suitable for low latency measurement dissemination, particularly in sensor network settings.

As age analysis for gossip networks is new, considerable work remains. Since this work has examined only the simplest network graphs, age analysis over more complex graphs is needed. Age analysis of gossip for energy harvesting sensors would also be another obvious area of interest. While this work employs the version age metric, we expect to see analogous results for the traditional sawtooth age metric that tracks the evolution of time. We also believe it may be possible to derive distributional properties of the age in a gossip network by extending the moment generating function (MGF) approach to age analysis in [17].

REFERENCES

[1] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Analysis and optimization of randomized gossip algorithms," in *2004 43rd IEEE Conference on Decision and Control (CDC)(IEEE Cat. No. 04CH37601)*, vol. 5. IEEE, 2004, pp. 5310–5315.

[2] ——, "Randomized gossip algorithms," *IEEE transactions on information theory*, vol. 52, no. 6, pp. 2508–2530, 2006.

[3] D. Shah, *Gossip algorithms*. Now Publishers Inc, 2009.

[4] K. Birman, "The promise, and limitations, of gossip protocols," *ACM SIGOPS Operating Systems Review*, vol. 41, no. 5, pp. 8–13, 2007.

[5] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing age of information in vehicular networks," in *IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, 2011.

[6] S. K. Kaul, R. D. Yates, and M. Gruteser, "On piggybacking in vehicular networks," in *IEEE Global Telecommunications Conference, GLOBECOM 2011*, Dec. 2011.

[7] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. IEEE INFOCOM*, March 2012, pp. 2731–2735.

[8] R. D. Yates, Y. Sun, D. R. Brown III, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," *arXiv preprint arXiv:2007.08564*, 2020.

[9] X. Wu, J. Yang, and J. Wu, "Optimal status update for age of information minimization with an energy harvesting source," *IEEE Trans. Green Commun. Netw.*, vol. 2, no. 1, pp. 193–204, March 2018.

[10] S. Feng and J. Yang, "Age of information minimization for an energy harvesting source with updating erasures: Without and with feedback," available Online: arXiv:1808.05141.

[11] A. Arafa, J. Yang, S. Ulukus, and H. V. Poor, "Age-minimal transmission for energy harvesting sensors with finite batteries: Online policies," *IEEE Trans. Inf. Theory*, vol. 66, no. 1, pp. 534–556, January 2020.

[12] ——, "Online timely status updates with erasures for energy harvesting sensors," in *Proc. Allerton*, October 2018.

[13] ——, "Using erasure feedback for online timely updating with an energy harvesting sensor," in *Proc. IEEE ISIT*, July 2019.

[14] B. T. Bacinoglu, E. T. Ceran, and E. Uysal-Biyikoglu, "Age of information under energy replenishment constraints," in *Proc. ITA*, February 2015.

[15] W. Liu, X. Zhou, S. Durrani, H. Mehrpouyan, and S. D. Blostein, "Energy harvesting wireless sensor networks: Delay analysis considering energy costs of sensing and transmission," *IEEE Trans. Wireless Commun.*, vol. 15, no. 7, pp. 4635–4650, July 2016.

[16] R. D. Yates, "Age of information in a network of preemptive servers," in *IEEE Conference on Computer Communications (INFOCOM) Workshops*, Apr. 2018, pp. 118–123, arXiv preprint arXiv:1803.07993.

[17] ——, "The age of information in networks: Moments, distributions, and sampling," *IEEE Transactions on Information Theory*, vol. 66, no. 9, pp. 5712–5728, 2020.

[18] R. D. Yates and S. K. Kaul, "The age of information: Real-time status updating by multiple sources," *IEEE Trans. Info. Theory*, vol. 65, no. 3, pp. 1807–1827, March 2019.

[19] S. Kaul, R. Yates, and M. Gruteser, "Status updates through queues," in *Conf. on Information Sciences and Systems (CISS)*, Mar. 2012.

[20] R. Yates and S. Kaul, "Real-time status updating: Multiple sources," in *Proc. IEEE Int'l. Symp. Info. Theory (ISIT)*, Jul. 2012.

[21] C. Kam, S. Kompella, and A. Ephremides, "Age of information under random updates," in *Proc. IEEE Int'l. Symp. Info. Theory (ISIT)*, 2013, pp. 66–70.

[22] ——, "Effect of message transmission diversity on status age," in *Proc. IEEE Int'l. Symp. Info. Theory (ISIT)*, June 2014, pp. 2411–2415.

[23] C. Kam, S. Kompella, G. D. Nguyen, and A. Ephremides, "Effect of message transmission path diversity on status age," *IEEE Trans. Info. Theory*, vol. 62, no. 3, pp. 1360–1374, Mar. 2016.

[24] R. D. Yates, "Status updates through networks of parallel servers," in *Proc. IEEE Int'l. Symp. Info. Theory (ISIT)*, Jun. 2018, pp. 2281–2285.

[25] R. Talak, S. Karaman, and E. Modiano, "Minimizing age-of-information in multi-hop wireless networks," in *55th Annual Allerton Conference on Communication, Control, and Computing*, Oct 2017, pp. 486–493.

[26] Q. He, D. Yuan, and A. Ephremides, "Optimal link scheduling for age minimization in wireless systems," *IEEE Trans. Info. Theory*, vol. 64, no. 7, pp. 5381–5394, July 2018.

[27] N. Lu, B. Ji, and B. Li, "Age-based scheduling: Improving data freshness for wireless real-time traffic," in *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. Mobihoc '18. New York, NY, USA: ACM, 2018, pp. 191–200. [Online]. Available: http://doi.acm.org/10.1145/3209582.3209602

[28] R. Talak, S. Karaman, and E. Modiano, "Distributed scheduling algorithms for optimizing information freshness in wireless networks," in *IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2018, pp. 1–5.

[29] ——, "Optimizing age of information in wireless networks with perfect channel state information," in *2018 16th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, 2018, pp. 1–8.

[30] R. Talak, I. Kadota, S. Karaman, and E. Modiano, "Scheduling policies for age minimization in wireless networks with unknown channel state," in *IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 2564–2568.

[31] A. Maatouk, M. Assaad, and A. Ephremides, "The age of updates in a simple relay network," in *2018 IEEE Information Theory Workshop (ITW)*, 2018, pp. 1–5.

[32] H. H. Yang, A. Arafa, T. Q. S. Quek, and H. V. Poor, "Locally adaptive scheduling policy for optimizing information freshness in wireless networks," in *2019 IEEE Global Communications Conference (GLOBECOM)*, 2019, pp. 1–6.

[33] B. Buyukates, A. Soysal, and S. Ulukus, "Age of information in multihop multicast networks," *Journal of Communications and Networks*, vol. 21, no. 3, pp. 256–267, 2019.

[34] S. Leng and A. Yener, "Age of information minimization for an energy harvesting cognitive radio," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 2, pp. 427–439, 2019.

[35] S. Farazi, A. G. Klein, and D. R. Brown, "Fundamental bounds on the age of information in multi-hop global status update networks," *Journal of Communications and Networks*, vol. 21, no. 3, pp. 268–279, 2019.

[36] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Optimizing data freshness, throughput, and delay in multi-server information-update systems," in *Proc. IEEE Int'l. Symp. Info. Theory (ISIT)*, 2016, pp. 2569–2574.

[37] ——, "Age-optimal information updates in multihop networks," in *Proc. IEEE Int'l. Symp. Info. Theory (ISIT)*, June 2017, pp. 576–580.

[38] A. M. Bedewy, Y. Sun, and N. B. Shroff, "The age of information in multihop networks," *IEEE/ACM Transactions on Networking*, vol. 27, no. 3, pp. 1248–1257, 2019.

[39] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides, "The age of incorrect information: A new performance metric for status updates," *IEEE/ACM Transactions on Networking*, vol. 28, no. 5, pp. 2215–2228, 2020.

[40] J. Zhong, R. Yates, and E. Soljanin, "Two freshness metrics for local cache refresh," in *Proc. IEEE Int'l. Symp. Info. Theory (ISIT)*, Jun. 2018, pp. 1924–1928.

[41] M. Bastopcu and S. Ulukus, "Who should google scholar update more often?" in *IEEE INFOCOM Age of Information Workshop*, may 2020, pp. 696–701.

[42] M. Wang, W. Chen, and A. Ephremides, "Reconstruction of counting process in real-time: The freshness of information through queues," in *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*, 2019, pp. 1–6.

[43] M. Bastopcu and S. Ulukus, "Information freshness in cache updating systems," *IEEE Transactions on Wireless Communications*, vol. 20, no. 3, pp. 1861–1874, 2021.

[44] ——, "Timely tracking of infection status of individuals in a population," in *IEEE INFOCOM Age of Information Workshop*, May 2021. [Online]. Available: https://arxiv.org/abs/2012.13393

[45] ——, "Cache freshness in information updating systems," in *2021 55th Annual Conference on Information Sciences and Systems (CISS)*, 2021, pp. 01–06.

[46] A. R. Teel, A. Subbaraman, and A. Sferlazza, "Stability analysis for stochastic hybrid systems: A survey," *Automatica*, vol. 50, no. 10, pp. 2435–2456, 2014.

[47] J. Hespanha, "Modelling and analysis of stochastic hybrid systems," *IEE Proceedings-Control Theory and Applications*, vol. 153, no. 5, pp. 520–535, 2006.

[48] J. P. Hespanha, "Hybrid and Switched Systems: ECE 229 – Fall 2005," https://www.ece.ucsb.edu/~hespanha/ece229/.