

Age of Gossip in Networks with Community Structure

Baturalp Buyukates Melih Bastopcu Sennur Ulukus
 Department of Electrical and Computer Engineering
 University of Maryland, College Park, MD 20742
 baturalp@umd.edu bastopcu@umd.edu ulukus@umd.edu

Abstract—We consider a network consisting of a single source and n receiver nodes that are grouped into m equal size communities, i.e., clusters, where each cluster includes k nodes and is served by a dedicated cluster head. The source node keeps versions of an observed process and updates each cluster through the associated cluster head. Nodes within each cluster are connected to each other according to a given network topology. Based on this topology, each node relays its current update to its neighboring nodes by *local gossiping*. We use the *version age* metric to quantify information timeliness at the receiver nodes. We consider disconnected, ring, and fully connected network topologies for each cluster. For each of these network topologies, we characterize the average version age at each node and find the version age scaling as a function of the network size n . Our results indicate that per node version age scalings of $O(\sqrt{n})$, $O(n^{\frac{1}{3}})$, and $O(\log n)$ are achievable in disconnected, ring, and fully connected cluster models, respectively. Finally, through numerical evaluations, we determine the version age-optimum (m, k) pairs as a function of the source, cluster head, and node update rates.

I. INTRODUCTION

Introduced in [1] to quantify timeliness in real-time status updating systems, the age of information metric has received significant attention across information, communication, networking, and queueing theory fields [2], [3]. The classical age metric increases linearly in time in the absence of any updates and drops to a smaller value when an update is received. Thus, even if the information at the source does not change, the classical age at the receiver continues to increase as time passes, because of the underlying assumption that updates get stale with time. This may not necessarily be the case in many applications, including content delivery services. To remedy this, several variants of the classical age metric have been proposed in the literature, in which the age stays the same until the information at the source changes even if no updates are received. Among these are binary freshness metric [4], age of synchronization [5], and age of incorrect information [6].

Similar in spirit, recently, a new age metric named *version age* has appeared in the literature [7], [8]. Considering each update at the source as a version change, the version age counts how many versions out-of-date the information at a particular receiver is, compared to the version at the source. Version age increases by one when the source obtains fresher information, i.e., newer version. A predecessor of version age has appeared




Fig. 1. Tiered network model where blue node at the center represents the source, yellow nodes represent the cluster heads, and green nodes represent the end users. Here, nodes in each cluster form a bi-directional ring network. Other possible network topologies within a cluster are shown in Fig. 2.

in [9], which considers timely tracking of Poisson counting processes, that entails minimizing the count difference, i.e., version difference, between the process and its estimate.

Reference [7] characterizes the version age in memoryless gossip networks composed of n arbitrarily connected nodes. To deliver information to the receiver nodes, the source in [7] employs a Poisson updating scheme (exponential inter-update times), as has been done previously in the context of social networks [10], timely tracking [9], [11], and timely cache updating [12]–[14]. In addition to the updates arriving directly from the source, in [7], nodes relay their update versions to each other. Also referred to as *gossiping*, this network activity improves the age scaling since each node can receive updates from each other as well as from the source. In particular, [7] shows that the version age scales as $O(\sqrt{n})$ in a bi-directional ring network and as $O(\log n)$ in a fully connected network, where n is the number of nodes. Earlier works on age scaling have considered the classical age metric and achieved $O(1)$ scaling in multicast networks [15]–[19] using a centralized transmission scheme administered by the source, and $O(\log n)$ scaling in distributed peer-to-peer communication networks [20] using a hierarchical local cooperation scheme.

Motivated by these, in this work, our aim is to investigate *version age scaling* in more general gossip network models which exhibit a community structure; see Fig. 1. In our model, as in [7], there is a single source that has an information which is updated following a Poisson process. Each such update at the source produces a newer version of the source information. The source node sends update packets regarding




Fig. 2. Different network topologies that can be used within each cluster: (a) disconnected, (b) uni-directional ring, (c) bi-directional ring, and (d) fully connected. Fig. 1 uses the one in (c). In this figure, cluster size is $k = 6$.

the source information to multiple communities. In our work, a community refers to a set of receiver nodes that are clustered together and can only interact with each other. Each cluster is served by a cluster head, which facilitates the communication with the source node. That is, cluster heads act as gateways between the source and the end-nodes within their clusters, akin to base stations in a cellular network. Unlike the model in [7], in our model, the source cannot directly update individual nodes and updates arriving from the source go through the cluster heads. To model the various degrees of gossip within each cluster, we use disconnected, uni-directional ring, bi-directional ring, and fully connected network topologies; see Fig. 2. Based on the underlying connectivity within clusters, we characterize the version age experienced by each node.

We observe that the additional hop constituted by the cluster heads between the source and the end-nodes presents us with opportunities to optimize the version age scaling by carefully tuning the number of clusters and the cluster size. Specifically, our results indicate that even if the nodes within each community forego gossiping, i.e., disconnected networks within each cluster, we can achieve $O(\sqrt{n})$ scaling as opposed to $O(n)$. In addition, we obtain the same $O(\log n)$ scaling in the case of fully connected communities using fewer connections within clusters than [7], and further reduce the scaling result in ring networks to $O(n^{\frac{1}{3}})$ from $O(\sqrt{n})$ in [7]. Finally, through numerical evaluations, we determine the version-age optimum cluster sizes for varying update rates employed by the source, cluster heads, and the nodes within each cluster.

II. SYSTEM MODEL AND THE AGE METRIC

We consider a system where a network of n nodes is divided into m clusters, each consisting of k nodes such that $n = mk$ with $k, m \in \mathbb{Z}$; see Fig. 1. Each cluster is served by a distinct cluster head, which takes updates from the source and distributes them across that cluster. The source process is updated exogenously as a rate λ_e Poisson process. The source has a total update injection rate of λ_s , which is uniformly distributed across cluster heads such that each cluster head is updated as a rate $\frac{\lambda_s}{m}$ Poisson process. From each cluster head to its corresponding cluster, the total update injection rate is λ_c and this rate is uniformly allocated across the nodes in that

cluster. That is, each node i receives an update from its cluster head as a rate $\frac{\lambda_c}{k}$ Poisson process with $i \in \mathcal{N} = \{1, \dots, n\}$.

Nodes in each cluster are connected to each other based on a connection graph. We consider varying levels of connectivity among nodes within each cluster. These are disconnected, uni-directional ring, bi-directional ring, and fully connected networks, which are shown in Fig. 2 for a cluster of $k = 6$ nodes. Updates received from the cluster head associated with each cluster are distributed across that cluster by utilizing the connections between the nodes. A node i updates another node j as a rate λ_{ij} Poisson process. Each node in this system has a total update rate of λ , which is uniformly allocated to its neighboring nodes. That is, in the uni-directional ring, each node updates its neighbor node as a rate λ Poisson process, whereas in bi-directional ring, each node has two neighboring nodes, each of which is updated as a rate $\frac{\lambda}{2}$ Poisson process. In the fully connected cluster, each node has $k - 1$ neighbors each of which is updated as a rate $\frac{\lambda}{k-1}$ Poisson process. As a result of these local connections within a cluster, a node can receive different versions of the source update.

To model the age at each node, we use the version age metric [7], [8]. We denote the version of the update at the source as $N_s(t)$, at cluster head c as $N_c(t)$, with $c \in \mathcal{C} = \{1, \dots, m\}$, and at node i as $N_i(t)$, with $i \in \mathcal{N}$, at time t . The version age at node i is given by $\Delta_i(t) = N_s(t) - N_i(t)$. Similarly, the version age at cluster head c is $\Delta_c(t) = N_s(t) - N_c(t)$. When node i has the same version as the source, its version age becomes zero, i.e., $\Delta_i(t) = 0$. When the information at the source is updated, version ages at the cluster heads and the nodes increase by 1, e.g., $\Delta_c'(t) = \Delta_c(t) + 1$. Each node i can get updates either from its cluster head or the other nodes that it is connected to within its cluster. When node i gets an update from its cluster head, its version age becomes $\Delta_i'(t) = \min\{\Delta_c(t), \Delta_i(t)\} = \Delta_c(t)$. Last equality here follows since nodes in a cluster receive the source updates through their cluster head so that they have either the same version or older versions of the information compared to their cluster head. When node i receives an update from node j , its version age becomes $\Delta_i'(t) = \min\{\Delta_i(t), \Delta_j(t)\}$. We note that if node j does not have a fresher version of the information, the version age at node i is not updated.

III. VERSION AGE WITH COMMUNITY STRUCTURE

In this section, we characterize the limiting version age of each node i , denoted by $\Delta_i = \lim_{t \rightarrow \infty} \mathbb{E}[\Delta_i(t)]$, $i \in \{1, \dots, n\}$ considering various network topologies for the clusters. Since the network model in each cluster is identical and within each cluster the network is symmetric for each of the network topologies, age processes $\Delta_i(t)$ of all users are statistically identical. Thus, in the ensuing analysis, we focus on a single cluster $c \in \mathcal{C}$ and find the average version age of a node from that cluster. For this, we follow the construction in [7] and express Δ_i in terms of Δ_S , which denotes the average version age of subset S , where $\Delta_S(t) \triangleq \min_{j \in S} \Delta_j(t)$.

We recall the following definitions from [7]: $\lambda_i(S)$ denotes the total update rate at which a node i updates the nodes in set

S . We have $\lambda_i(S) = \sum_{j \in S} \lambda_{ij}$ when $i \notin S$. Similarly, $\lambda_c(S)$ denotes the total update rate of the cluster head of a particular cluster into the set S . Finally, set of updating neighbors of a set S is $N(S) = \{i \in \{1, \dots, n\} : \lambda_i(S) > 0\}$.

With these definitions, next, in Theorem 1 below we give the resulting version age in our clustered system model as a specialization of [7, Thm. 1].

Theorem 1 *When the total network of n nodes is divided into m clusters, each of which consisting of a single cluster head and k nodes with $n = mk$, the average version age of subset S that is composed of nodes within a cluster is given by*

$$\Delta_S = \frac{\lambda_e + \lambda_c(S)\Delta_c + \sum_{i \in N(S)} \lambda_i(S)\Delta_{S \cup \{i\}}}{\lambda_c(S) + \sum_{i \in N(S)} \lambda_i(S)}, \quad (1)$$

with $\Delta_c = m \frac{\lambda_e}{\lambda_s}$.

Proof of Theorem 1 follows by applying [7, Thm. 1] to our clustered network model and noting that updates arrive at the nodes through designated cluster heads.

A. Version Age in Clustered Disconnected Networks

Nodes in a cluster are not connected to each other. Thus, the network is a two-hop multicast network, where the first hop is from the source to m cluster heads, and the second hop is from each cluster head to k nodes; combine Fig. 1 with Fig. 2(a). Multihop networks have been studied in [15]–[19] considering the classical age metric, where the source keeps sending update packets until they are received by a certain number of nodes at each hop. We do not consider such centralized management of updates, but let the source update the cluster heads as Poisson processes, and let cluster heads forward these packets to the nodes within their clusters as further Poisson processes.

Let S_1 denote an arbitrary 1-node subset of a cluster. Subset S_1 is only connected to the cluster head, i.e., $N(S_1) = \emptyset$. Using the recursion given in (1), we find

$$\Delta_{S_1} = \Delta_c + k \frac{\lambda_e}{\lambda_c} = m \frac{\lambda_e}{\lambda_s} + k \frac{\lambda_e}{\lambda_c}, \quad (2)$$

where Δ_{S_1} denotes the version age of a single node from the cluster. When the network consists of two-hops, version age is additive, in that the first term in (2) corresponds to the first hop and equals to the version age at the cluster head, whereas the second term in (2) corresponds to the version age at the second hop between the cluster head and a node.

Theorem 2 *In a clustered network of disconnected users, the version age of a single user scales as $O(\sqrt{n})$.*

Theorem 2 follows by selecting $k = \sqrt{n}$ with $m = \frac{n}{k} = \sqrt{n}$ in (2) for fixed $\lambda_e, \lambda_s, \lambda_c$, which do not depend on n . Theorem 2 indicates that when nodes are grouped into \sqrt{n} clusters, an age scaling of $O(\sqrt{n})$ is achievable even though users forego gossiping. With the absence of cluster heads, i.e., when the source is uniformly connected to each of the n users, the version age scaling of each disconnected user would be $O(n)$. By utilizing clusters, we incur an additional hop, but significantly improve the scaling result from $O(n)$ to $O(\sqrt{n})$.

B. Version Age in Clustered Ring Networks

Nodes in each cluster form a ring network. We consider two types of ring clusters: uni-directional ring as shown in Fig. 2(b) and bi-directional ring as shown in Fig. 2(c).

First, we consider the uni-directional ring and observe that an arbitrary subset of j adjacent nodes S_j has a single neighbor node that sends updates with rate λ for $j \leq k-1$. Each such subset S_j receives updates from the cluster head with a total rate of $j \frac{\lambda_c}{k}$. Next, we use the recursion in (1) to write

$$\Delta_{S_j} = \frac{\lambda_e + j \frac{\lambda_c}{k} \Delta_c + \lambda \Delta_{S_{j+1}}}{j \frac{\lambda_c}{k} + \lambda}, \quad (3)$$

for $j \leq k-1$ where Δ_c is the version age at the cluster head. We note that when $j = k$ the network becomes a simple two-hop network similar to that of Section III-A and we find $\Delta_{S_k} = m \frac{\lambda_e}{\lambda_s} + \frac{\lambda_e}{\lambda_c}$.

Next, we consider the bi-directional ring and observe that an arbitrary subset S_j that consists of any adjacent j nodes has two neighbor nodes, each with an incoming update rate of $\frac{\lambda}{2}$ for $j < k-1$. When $j = k-1$, S_j has a single neighboring node that sends updates with a total rate $2 \frac{\lambda}{2} = \lambda$. For $j \leq k-1$, the cluster head sends updates to subset S_j with a total rate of $j \frac{\lambda_c}{k}$. With all these, when we apply the recursion in (1), we obtain exactly the same formula given in (3).

Lemma 1 *Both uni-directional and bi-directional ring cluster models yield the same version age for a single node when each node in a cluster has a total update rate of λ .*

Lemma 1 follows from the fact that either type of ring cluster induces the same recursion for an arbitrary subset of any adjacent j nodes within a cluster as long as the total update rate per node λ is the same. Thus, in the remainder of this paper, we only consider the bi-directional ring cluster model.

Before focusing on age scaling in a clustered network with a ring topology in each cluster, we revisit the ring network in [7], and provide a proof of the $1.25\sqrt{n}$ age scaling result observed therein as a numerical result. We show that the approximate theoretical coefficient is $\sqrt{\frac{\pi}{2}} = 1.2533$.

Lemma 2 *For the ring network model considered in [7], the version age of a user scales as $\Delta_{S_1} \approx \sqrt{\frac{\pi}{2}} \frac{\lambda_e}{\lambda} \sqrt{n}$.*

Proof: From recursive application of [7, Eqn. (17)], we obtain

$$\Delta_{S_1} = \frac{\lambda_e}{\lambda} \left(\sum_{i=1}^{n-1} a_i^{(n)} + a_{n-1}^{(n)} \right), \quad (4)$$

where $a_i^{(n)}$ is given for $i = 1, \dots, n-1$ as

$$a_i^{(n)} = \prod_{j=1}^i \frac{1}{1 + \frac{j}{n}}. \quad (5)$$

We note that $a_i^{(n)}$ decays fast in i , and consider $i = o(n)$,

$$-\log(a_i^{(n)}) = \sum_{j=1}^i \log \left(1 + \frac{j}{n} \right) \approx \sum_{j=1}^i \frac{j}{n} = \frac{i(i+1)}{2n} \approx \frac{i^2}{n} \quad (6)$$

where we used $\log(1+x) \approx x$ for small x , and ignored the i term relative to i^2 . Thus, for small i , we have $a_i^{(n)} \approx e^{-\frac{i^2}{2n}}$. For large i , $a_i^{(n)}$ converges quickly to zero due to multiplicative terms in $\prod_{j=1}^i \frac{1}{1+j/n}$, and this approximation still holds. Thus, we have $\sum_{i=1}^{n-1} a_i^{(n)} \approx \sum_{i=1}^{n-1} e^{-\frac{i^2}{2n}}$. For large n , by using Riemann sum approximation with steps $\frac{1}{\sqrt{n}}$, we obtain

$$\frac{1}{\sqrt{n}} \sum_{i=1}^{n-1} a_i^{(n)} \approx \frac{1}{\sqrt{n}} \sum_{i=1}^{n-1} e^{-\frac{i^2}{2n}} = \int_0^\infty e^{-\frac{t^2}{2}} dt = \sqrt{\frac{\pi}{2}}. \quad (7)$$

Thus, we get $\sum_{i=1}^{n-1} a_i^{(n)} \approx \sqrt{\frac{\pi}{2}} \sqrt{n}$. By inserting this in (4), we obtain the age scaling of a user as $\Delta_{S_1} \approx \sqrt{\frac{\pi}{2}} \frac{\lambda_e}{\lambda} \sqrt{n}$. ■

Next, we focus on age scaling in a clustered network with a ring topology in each cluster. From recursive application of (3) along with Δ_{S_k} , we obtain

$$\Delta_{S_1} = \frac{\lambda_e}{\lambda} \left(\sum_{i=1}^{k-1} b_i^{(k)} \right) + \Delta_c \left(1 - b_{k-1}^{(k)} \right) + \Delta_{S_k} b_{k-1}^{(k)}, \quad (8)$$

where similar to (5), $b_i^{(k)}$ is given for $i = 1, \dots, k-1$ as

$$b_i^{(k)} = \prod_{j=1}^i \frac{1}{1 + \frac{j \lambda_e}{k \lambda}}. \quad (9)$$

When k is large, $b_{k-1}^{(k)}$ goes to zero, and Δ_{S_1} in (8) becomes

$$\Delta_{S_1} \approx \frac{\lambda_e}{\lambda} \left(\sum_{i=1}^{k-1} b_i^{(k)} \right) + \Delta_c \approx \sqrt{\frac{\pi}{2}} \frac{\lambda_e}{\sqrt{\lambda \lambda_c}} \sqrt{k} + m \frac{\lambda_e}{\lambda_s}, \quad (10)$$

where the second approximation follows as in the proof of Lemma 2. Terms in (10) are $O(\sqrt{k})$ and $O(m)$, respectively. In [7], there is a single cluster, i.e., $m = 1$ and $k = n$, and thus, the version age scaling is $O(\sqrt{n})$. In our model, by carefully adjusting the number of clusters and the cluster sizes, we can improve this $O(\sqrt{n})$ scaling result to $O(n^{\frac{1}{3}})$.

Theorem 3 *In a clustered network with a ring topology in each cluster, the version age of a single user scales as $O(n^{\frac{1}{3}})$.*

Theorem 3 follows by selecting $m = n^{\frac{1}{3}}$ with $k = \frac{n}{m} = n^{\frac{2}{3}}$ in (10) for fixed $\lambda_e, \lambda_s, \lambda_c, \lambda$, which do not depend on n .

C. Version Age in Clustered Fully Connected Networks

Nodes in each cluster form a fully connected network where each node is connected to all the other nodes within its cluster with rate $\frac{\lambda}{k-1}$. We find the version age for a subset of j nodes S_j in a cluster. Each such subset j has $k-j$ neighbor nodes in addition to the cluster head associated with their cluster. Using the recursion given in (1), we find

$$\Delta_{S_j} = \frac{\lambda_e + \frac{j \lambda_e}{k} \Delta_c + \frac{j(k-j)\lambda}{k-1} \Delta_{S_{j+1}}}{\frac{j \lambda_c}{k} + \frac{j(k-j)\lambda}{k-1}}, \quad (11)$$

for $j \leq k-1$, where Δ_c is equal to $m \frac{\lambda_e}{\lambda_s}$. The average version age of the whole cluster is $\Delta_{S_k} = \Delta_c + \frac{\lambda_e}{\lambda_c} = m \frac{\lambda_e}{\lambda_s} + \frac{\lambda_e}{\lambda_c}$.

Next, we present bounds for Δ_{S_1} .

Lemma 3 *When $\lambda_c = \lambda$, in a clustered network with fully connected topology in each cluster, the version age of a single node satisfies*

$$\begin{aligned} & \frac{(k-1)^2 + k}{k^2} \Delta_c + \frac{\lambda_e}{\lambda} \left(\frac{k-1}{k} \sum_{\ell=1}^{k-1} \frac{1}{\ell} + \frac{1}{k} \right) \\ & \leq \Delta_{S_1} \leq \Delta_c + \frac{\lambda_e}{\lambda} \left(\sum_{\ell=1}^k \frac{1}{\ell} \right). \end{aligned} \quad (12)$$

Proof: We use steps similar to those in the proof of [7, Thm. 2] and also consider the additional hop from the source to the cluster heads. For $\lambda_c = \lambda$, we take $j = k-\ell$ and (11) becomes

$$\Delta_{S_{k-\ell}} = \frac{\frac{1}{k-\ell} \frac{\lambda_e}{\lambda} + \frac{1}{k} \Delta_c + \frac{\ell}{k-1} \Delta_{S_{k-\ell+1}}}{\frac{1}{k} + \frac{\ell}{k-1}}, \quad (13)$$

for $\ell \leq k-1$ and $\Delta_{S_k} = \Delta_c + \frac{\lambda_e}{\lambda}$, where Δ_c is the age at the cluster head. Defining $\hat{\Delta}_{S_\ell} \triangleq \Delta_{S_{k-\ell+1}}$, we get

$$\hat{\Delta}_{S_{\ell+1}} = \frac{\frac{1}{k-\ell} \frac{\lambda_e}{\lambda} + \frac{1}{k} \Delta_c + \frac{\ell}{k-1} \hat{\Delta}_{S_\ell}}{\frac{1}{k} + \frac{\ell}{k-1}}. \quad (14)$$

Next, one can show that $\hat{\Delta}_{S_{\ell+1}}$ satisfies the following

$$\hat{\Delta}_{S_{\ell+1}} \leq \frac{\frac{1}{k-\ell} \frac{\lambda_e}{\lambda} + \frac{1}{k} \Delta_c + \frac{\ell}{k} \hat{\Delta}_{S_\ell}}{\frac{1}{k} + \frac{\ell}{k}}. \quad (15)$$

Defining $\tilde{\Delta}_{S_\ell} \triangleq \frac{\ell}{k} \hat{\Delta}_{S_\ell}$ and plugging it in (15), we get

$$\tilde{\Delta}_{S_{\ell+1}} = \frac{\ell+1}{k} \hat{\Delta}_{S_\ell} \leq \frac{1}{k-\ell} \frac{\lambda_e}{\lambda} + \frac{1}{k} \Delta_c + \tilde{\Delta}_{S_\ell}. \quad (16)$$

Noting that $\tilde{\Delta}_{S_1} = \frac{\hat{\Delta}_{S_1}}{k} = \frac{\Delta_{S_k}}{k} = \frac{1}{k} (\Delta_c + \frac{\lambda_e}{\lambda})$, we write

$$\tilde{\Delta}_{S_k} \leq \Delta_c + \frac{\lambda_e}{\lambda} \left(\sum_{\ell=1}^k \frac{1}{\ell} \right). \quad (17)$$

Since $\tilde{\Delta}_{S_k} = \hat{\Delta}_{S_k} = \Delta_{S_1}$, (17) presents an upper bound to the version age of a single node. For the lower bound, we follow similar steps starting from (14). Detailed steps are omitted here due to space limitations. ■

From (12), we see that for large n with $\lambda_c = \lambda$, the version age of a single node Δ_{S_1} satisfies

$$\Delta_{S_1} \approx m \frac{\lambda_e}{\lambda_s} + \frac{\lambda_e}{\lambda} \log k. \quad (18)$$

Theorem 4 *In a clustered network with a fully connected topology in each cluster, the version age of a single user scales as $O(\log n)$.*

Theorem 4 follows in multiple different ways. For instance, it follows by selecting $m = 1$ and $k = \frac{n}{m} = n$. That is, we have a single fully connected network of n users as in [7]. Theorem 4 also follows by selecting $m = \log n$ and $k = \frac{n}{m} = \frac{n}{\log n}$. That is, we have $\log(n)$ fully connected clusters with $\frac{n}{\log n}$ users in each cluster. Thus, version age obtained under a smaller cluster size with less connections is the same




Fig. 3. Version age of a node with fully connected, ring, and disconnected cluster models with $n = 120$, (a) $\lambda_e = 1, \lambda_s = 1, \lambda_c = 1$, and $\lambda = 1$, (b) $\lambda_e = 1, \lambda_s = 10, \lambda_c = 1$, and $\lambda = 1$, (c) $\lambda_e = 1, \lambda_s = 10, \lambda_c = 10$, and $\lambda = 1$, (d) $\lambda_e = 1, \lambda_s = 10, \lambda_c = 1$, and $\lambda = 2$.

as that obtained when all nodes are connected to each other. In particular, in our model with $m = \log n$, each node has $O(\frac{n}{\log n})$ connections in comparison to $O(n)$ in [7].

Finally, we note that, a recurring theme in the analysis of clustered networks is the fact that the version age at an end-node Δ_{S_1} is *almost* additive in the version age at the cluster head Δ_c as seen in (2), (8), and (12). It is *exactly* additive in the case of disconnected clusters in (2).

IV. NUMERICAL RESULTS

We have seen in Section III that the version age depends on update rates $\lambda_e, \lambda_s, \lambda_c$, and λ . In this section, we explore the effects of these rates on the age via numerical results.

First, we take $\lambda_e = 1, \lambda_s = 1, \lambda_c = 1, \lambda = 1$, and $n = 120$. We plot the version age of a node for the considered cluster models with respect to k . We see in Fig. 3(a) that for the fully connected cluster model, the version age decreases with k and thus, the version age-optimal cluster size is $k^* = 120$, i.e., all n nodes are grouped in a single cluster. In the ring cluster model, the version age is minimized when $k^* = 30$. In the disconnected cluster model, the version age is minimized when we have $k^* = 10$ or $k^* = 12$. From these, we deduce that when the topology has less connectivity in a cluster, the optimal cluster size is smaller. Further, a topology with larger connectivity within a cluster achieves a lower version age.

Second, we consider the same setting as in Fig. 3(a) but take $\lambda_s = 10$ in Fig. 3(b). Here, the version age decreases with increasing k at first due to increasing number of connections within a cluster and the increase in the update rate between the source and each cluster head (as the number of clusters decreases with increasing k). However, as k continues to increase, the decrease in the update rate from the cluster head to the nodes starts to dominate and the version age increases

for all cluster models. In Fig. 3(b), we see that the optimal cluster size is $k^* = 12$ in fully connected clusters, $k^* = 8$ in ring clusters, $k^* = 3$ and $k^* = 4$ in disconnected clusters.

Third, we increase the update rate of the cluster heads and take $\lambda_c = 10$. We see in Fig. 3(c) that the optimum value of k increases compared to the second case when cluster heads have a larger update rate in all the cluster models. We find $k^* = 20$ in fully connected clusters, $k^* = 15$ in ring clusters, and $k^* = 10$ or $k^* = 12$ in disconnected clusters.

Fourth, we study the effect of update rates among the nodes. For this, we take $\lambda_c = 10, \lambda_e = 1, \lambda_s = 1, \lambda = 2$. We see in Fig. 3(d) that as the communication rate between the nodes increases, the optimal cluster size increases, and it is equal to $k^* = 24$ in fully connected clusters, and $k^* = 10$ in ring clusters. As there is no connection between nodes in the case of disconnected clusters, the optimum cluster size remains the same, i.e., $k^* = 3$ or $k^* = 4$, compared to Fig. 3(b).

REFERENCES

- [1] S. K. Kaul, R. D. Yates, and M. Gruteser. Real-time status: How often should one update? In *IEEE Infocom*, March 2012.
- [2] 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.
- [3] 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.
- [4] J. Cho and H. Garcia-Molina. Effective page refresh policies for web crawlers. *ACM Trans. on Database Systems*, 28(4):390–426, December 2003.
- [5] J. Zhong, R. D. Yates, and E. Soljanin. Two freshness metrics for local cache refresh. In *IEEE ISIT*, June 2018.
- [6] 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.
- [7] R. D. Yates. The age of gossip in networks. In *IEEE ISIT*, July 2021.
- [8] B. Abolhassani, J. Tadrous, A. Eryilmaz, and E. Yeh. Fresh caching for dynamic content. In *IEEE Infocom*, May 2021.
- [9] M. Bastopcu and S. Ulukus. Who should Google Scholar update more often? In *IEEE Infocom*, July 2020.
- [10] S. Ioannidis, A. Chaintreau, and L. Massoulié. Optimal and scalable distribution of content updates over a mobile social network. In *IEEE Infocom*, April 2009.
- [11] M. Bastopcu and S. Ulukus. Timely tracking of infection status of individuals in a population. In *IEEE Infocom*, May 2021.
- [12] M. Bastopcu and S. Ulukus. Information freshness in cache updating systems. *IEEE Trans. on Wireless Communications*, 20(3):1861–1874, March 2021.
- [13] M. Bastopcu and S. Ulukus. Maximizing information freshness in caching systems with limited cache storage capacity. In *Asilomar Conference*, November 2020.
- [14] P. Kaswan, M. Bastopcu, and S. Ulukus. Freshness based cache updating in parallel relay networks. In *IEEE ISIT*, July 2021.
- [15] J. Zhong, E. Soljanin, and R. D. Yates. Status updates through multicast networks. In *Allerton Conference*, October 2017.
- [16] J. Zhong, R. D. Yates, and E. Soljanin. Multicast with prioritized delivery: How fresh is your data? In *IEEE SPAWC*, June 2018.
- [17] B. Buyukates, A. Soysal, and S. Ulukus. Age of information in two-hop multicast networks. In *Asilomar Conference*, October 2018.
- [18] B. Buyukates, A. Soysal, and S. Ulukus. Age of information in multihop multicast networks. *Journal of Communications and Networks*, 21(3):256–267, July 2019.
- [19] B. Buyukates, A. Soysal, and S. Ulukus. Age of information in multicast networks with multiple update streams. In *Asilomar Conference*, November 2019.
- [20] B. Buyukates, A. Soysal, and S. Ulukus. Scaling laws for age of information in wireless networks. *IEEE Trans. on Wireless Communications*, 20(4):2413–2427, April 2021.