Fundamental Bounds on the Age of Information in Multi-Hop Global Status Update Networks

Shahab Farazi, Andrew G. Klein, and D. Richard Brown III

Abstract: **This paper studies the "age of information" in a general multi-source multi-hop wireless network with explicit channel contention. Specifically, the scenario considered in this paper assumes that each node in the network is both a source and a monitor of information, that all nodes wish to receive fresh status updates from all other nodes in the network, and that only one node can transmit in each time slot. Lower bounds for peak and average age of information are derived and expressed in terms of fundamental graph properties including the connected domination number. An algorithm to generate near-optimal periodic status update schedules based on sequential optimal flooding is also developed. These schedules are analytically shown to exactly achieve the peak age bound and also achieve the average age bound within an additive gap scaling linearly with the size of the network. Moreover, the results are sufficiently general to apply to any connected network topology. Illustrative numerical examples are presented which serve to verify the analysis for several canonical network topologies of arbitrary size, as well as every connected network with nine or fewer nodes.**

Index Terms: **Age of information, explicit contention, graph theory, multi-hop, multi-source, peak and average age.**

I. INTRODUCTION

INFORMATION freshness is of critical importance in a variety of networked monitoring and control systems such as intelligent vehicular systems, channel state feedback, and environmental monitoring as well as applications such as financial trading and online learning. In these types of applications, stale information can lead to incorrect decisions, unstable control loops, and even compromises in safety and security. A recent line of research has considered information freshness from a fundamental perspective under an *Age of Information* (AoI) metric first proposed in [1] and further studied in [2]–[36]. The central idea is that there are one or more sources of information along with one or more monitors. A source generates timestamped status updates which are received by a monitor after some delay. The "age of information" is defined as the difference between the current time and the timestamp of the most recent status update at the monitor. A common theme of the AoI literature is to study and/or optimize the statistics of AoI, i.e., average age and/or peak age, as a function of the system parameters and update strategies.

The majority of the prior work on AoI has focused on the single-hop setting, where one or more sources transmit information to one or more monitors over a single hop, typically modeled as a random delay through a queue. The single-source, single-monitor, single-hop setting was studied in [2]–[8]. A more comprehensive AoI literature review of this setting can be found in [9]. Multi-source and/or multi-monitor extensions, also in the single-hop context, have been considered in [10]–[23]. These multi-source and/or multi-monitor settings often introduce additional delays in delivering updates in the single-hop setting through explicit contention for channel resources and interference constraints.

The multi-hop setting considered in this paper has received relatively little attention in the AoI literature despite the early consideration of "piggybacking" status updates over multiple hops in vehicular networks in [24]. The analysis in [24]–[29] considers age of information in specific multi-hop network structures, e.g., line, ring, and/or two-hop networks. The schedules and performance metrics derived for these specific networks are not easily extended to more general settings. Recent work in [30] considers a single source delivering updates to a monitor through a potentially multi-hop network modeled as a stochastic hybrid system (SHS). This work generalizes the results for the line network in [29], but the effects of contention and interference constraints are implicit in the sense that they are abstracted into the SHS model. A general multi-hop network setting where a single-source disseminates status updates through a gateway to the network was considered in [31]–[33]. These studies also do not consider the effects of channel contention or interference constraints as they assume that all links in the network are modeled as interference-free. A practical age control protocol to improve AoI in multi-hop IP networks was also recently proposed in [35].

The work closest to this paper is [36], which considers a setting with multiple sources, multiple monitors, and a multi-hop network with explicit contention and interference constraints. The system model in [36] assumes a setting with predefined distinct source-monitor pairs such that the information at source i is only of interest to monitor i. The analysis focuses on the development of age-optimal random transmission policies where links are activated according to a fixed probability distribution.

This paper also considers a general multi-source, multi-monitor, multi-hop setting with explicit interference constraints, but from a *global* perspective in the sense that (i) every node in the network is both a source and monitor of information and (ii)

Manuscript received November 29, 2018 approved for publication by Sastry Kompella, Guest Editor, April 15, 2019.

This research was supported by National Science Foundation awards CNS-1836695 and CNS-1836690.

S. Farazi and D. Richard Brown III are with the Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, email: sfarazi,drb@wpi.edu.

A. G. Klein is with the Department of Engineering and Design, Western Washington University, email: andy.klein@wwu.edu.

S. Farazi is the corresponding author.

Digital Object Identifier: 10.1109/JCN.2019.000038

every node wishes to receive timely status updates from all other nodes in the network. This setting is appropriate in applications where nodes both generate status updates and monitor status updates from other nodes in the network, e.g., autonomous vehicles. Our only assumption on the network is that it is connected. Also, in contrast to [36], we derive fundamental limits and develop an algorithm to generate *deterministic* schedules for general network topologies with certain age-optimal properties. The work in this paper builds on our preliminary results in [34]. The main contributions of this paper are:

1. We consider AoI in a general multi-source, multi-monitor, multi-hop setting with explicit interference constraints and global status updates. The analysis in this setting is facilitated with graph-theoretic tools rather than the queuing theoretic tools used in most of the previous AoI literature.
2. We derive uniform lower bounds on instantaneous peak age and instantaneous average age for any schedule (deterministic or random). These bounds establish a connection between the AoI metrics and the fundamental measures of a network topology such as connected domination number and average path length. These instantaneous bounds were not discussed in [34].
3. We develop an algorithm that, given a connected network topology, generates minimum-length periodic schedules which refresh all of the status updates throughout the network once per period.
4. We derive lower bounds on the peak and average age of information, as computed over any one-period interval, for any minimum-length periodic schedule. The proofs of these lower bounds were not presented in [34].
5. We show analytically that, within the class of minimum-length periodic schedules, our proposed algorithm generates schedules that exactly achieve the lower bound on peak age and achieve the lower bound on average age to within an additive gap scaling linearly with the number of nodes in the network. The proofs of these results were not presented in [34].
6. We present numerical results verifying the analysis of the peak and average age for several canonical network structures with any number of nodes and also for every connected network topology with nine or fewer nodes. The analysis of the canonical network structures was also not discussed in [34].

In particular, other than [22], the development of fundamental bounds on peak and average age has received relatively little attention in the AoI literature. The bounds derived in this paper hold for all connected network topologies in which one node transmits per time slot.

Notation: Unless otherwise noted, lowercase bold letters (e.g. $\mathbf{x}$) denote vectors, uppercase bold letters (e.g., $\mathbf{X}$) represent matrices, and calligraphic uppercase letters (e.g. $\mathcal{X}$) are used for sets. The entry in the i th row and j th column of matrix $\mathbf{X}$ is denoted by $\mathbf{X}(i,j)$, and the i th element of vector $\mathbf{x}$ is denoted by $\mathbf{x}_i$. Table 1 provides a list of parameters and notation used throughout this paper.

Table 1. Summary of notation.

Notation	Description
$\mathcal{G}$	an undirected graph representing the wireless network
$\mathcal{E}$	set of the edges in $\mathcal{G}$ representing the channels in the network
$\mathcal{V}$	set of the vertices in $\mathcal{G}$ representing the nodes in the network
$\mathcal{G}[\mathcal{U}]$	the graph induced by a vertex set $\mathcal{U} \subseteq \mathcal{V}$
N	number of the nodes in the network, i.e., $N = \lvert\mathcal{V}\rvert$
$d(i,j)$	shortest path length between two vertices i and j
$\bar{d}$	average shortest path length between all vertices in $\mathcal{G}$
δ_i	degree of vertex i
$\delta_{\max}$	maximum degree over all vertices in $\mathcal{G}$
$\mathcal{N}_k(i)$	set of vertices $j \in \mathcal{V}$ such that $1 \leq d(i,j) \leq k$
γ_c	connected domination number, i.e., cardinality of any MCDS
$\mathcal{L}$	set of pseudoleaf vertices, i.e., all vertices not in any MCDS
$H_j^{(i)}(t)$	the H_j process at node i
$\tau_j^{(i)}(t)$	timestamp of the most recent H_j process at node i
$\Delta_j^{(i)}(t)$	the age of the H_j process at node i

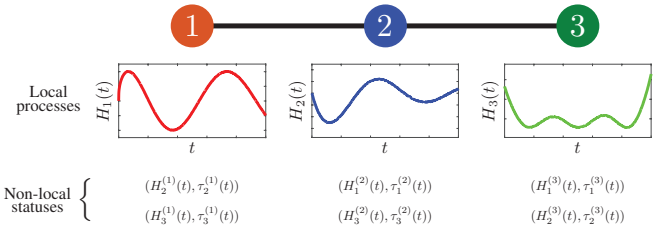

Fig. 1. Example 3-node line network. Each node is associated with a local process and maintains tables of non-local statuses for the processes of other nodes in the network.

II. SYSTEM MODEL

We consider a wireless network where connectivity of nodes is modeled by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where the vertex set $\mathcal{V}$ represents the wireless nodes and the edge set $\mathcal{E}$ represents the channels between the nodes in the network. Edge $e_{i,j}$ is in set $\mathcal{E}$ when vertices i and j are adjacent (i.e., when there exists a channel between nodes i and j such that transmissions can be successfully delivered from one directly to the other). We denote the number of nodes as $N = \lvert\mathcal{V}\rvert$ and the set of one-hop neighbors of node i as $\mathcal{N}_1(i)$, i.e., $j \in \mathcal{N}_1(i) \Leftrightarrow e_{i,j} \in \mathcal{E}$. Our only assumption with regards to the topology of the network is that it is connected, i.e., there exists a path between any two distinct vertices $i, j \in \mathcal{V}$.

We assume a setting where each node $i \in \mathcal{V}$ is associated with a local process $H_i(t)$ as illustrated in a simple three-node setting in Fig. 1. No assumptions are made about these processes other than they are time-varying and each is of timely interest to all nodes in the network. In addition to its local process, each node $i \in \mathcal{V}$ has a table of "statuses" of all of the non-local processes in the network. We denote a status as the tuple $(H_j^{(i)}(t), \tau_j^{(i)}(t))$, where $H_j^{(i)}(t)$ and $\tau_j^{(i)}(t)$ denote the most recent sample value and the corresponding timestamp of process $H_j(t)$ known to node i at time t, respectively.

Since the processes $\{H_j(t)\}$ are time-varying and of timely interest to all nodes in the network, each node $i \in \mathcal{V}$ seeks to maintain a table of "fresh" statuses with recent timestamps. For simplicity, we assume each node can sample its own local process without delay such that $H_i^{(i)}(t) = H_i(t)$ and $\tau_j^{(i)}(t) = t$ for all $i \in \mathcal{V}$ and all t. The remaining statuses must be updated via broadcast transmissions containing status updates from other nodes in the system. We make the following assumptions re-

garding these broadcast transmissions:

1. Transmissions are time slotted, require one unit of time to complete, and are received at times $t = 1, 2, 3, \cdots$.
2. Only one node can transmit in each time slot[1].
3. Each transmission contains one status, i.e., one sample and its corresponding timestamp, of a process. The transmitting node may transmit the status of its own local process or its status of another node's process.
4. Transmissions from node i are received reliably by all nodes in the one-hop neighborhood of node i, denoted by $\mathcal{N}_1(i)$, while nodes outside the one-hop neighborhood receive nothing.

As each node in the network is both a source and a monitor of information, there are inherent tradeoffs in how fresh the status at each node can be in this setting. For example, in the three-node network in Fig. 1, node 2 can keep nodes 1 and 3 updated with fresh status updates of $H_2(t)$ by transmitting new samples of its local process in every time slot. While statuses $H_2^{(1)}(t)$ and $H_2^{(3)}(t)$ remain fresh (we omit the timestamps here for notational convenience), statuses $H_1^{(2)}(t)$ and $H_3^{(2)}(t)$ become stale since node 2 does not receive updates from node 1 or node 3. Moreover, $H_3^{(1)}(t)$ and $H_1^{(3)}(t)$ also become stale since nodes 1 and 3 receive no status updates from each other in this example. To formalize these notions, we define the age metric for the status of process $H_j(t)$ at node i below.

Definition 1 (Age) Given the status of process j at node i denoted as $(H_j^{(i)}(t), \tau_j^{(i)}(t))$, the *age* of this process at time $t \geq \tau_j^{(i)}(t)$ is defined as $\Delta_j^{(i)}(t) \triangleq t - \tau_j^{(i)}(t)$.

The age metric here is a multi-source multi-monitor generalization of the single-source single-monitor age metric first proposed in [2]. Note that the age $\Delta_j^{(i)}(t)$ is non-negative and is not defined for $t < \tau_j^{(i)}(t)$ or if no status update for process $H_j(t)$ has been received at node i. Under our previous assumption about the ability of each node to instantaneously sample its local process, the local age $\Delta_i^{(i)}(t) = 0$. While our model could be extended to include the effect of non-zero delay in sampling local processes, such delays would appear as a simple additive term in the various age expressions below; for simplicity, these local ages are ignored in the metrics defined in Section III and in the dynamic model describing the evolutions of the ages below.

Given Definition 1 and the assumed time slotted transmissions, we can describe the dynamics of each age in the system with a simple discrete time model similar to [13], [14]. Specifically, given a status update from node i regarding process j, the age at each node $m \in \mathcal{V}$ with $m \neq j$ is updated at integer times $t = n$ according to

$$\Delta_j^{(m)}[n+1] = \begin{cases} 1 & m \in \mathcal{N}_1(i) \text{ and } i = j \\ \Delta_j^{(i)}[n] + 1 & m \in \mathcal{N}_1(i),\ i \neq j, \\ & \text{ and } \Delta_j^{(i)}[n] < \Delta_j^{(m)}[n] \\ \Delta_j^{(m)}[n] + 1 & \text{otherwise.} \end{cases} \quad (1)$$

In order for node $m \neq j$ to update its status of process j and reduce the corresponding age $\Delta_j^{(m)}(t)$, it must (i) receive the status update transmission, i.e., be within the one-hop neighborhood of the transmitting node, and (ii) the status update must be fresher than the current status at node m. Otherwise, the age simply increases by one. The first case in (1) corresponds to the case when node i transmits a status update of its local process $H_i(t)$. In this case, since transmissions require unit time to complete, the local age at the start of the transmission is $\Delta_i^{(i)}[n] = 0$ and the age when nodes $m \in \mathcal{N}_1(i)$ receive the status update is $\Delta_i^{(m)}[n+1] = 1$. The second case in (1) corresponds to the case when node i transmits a status update of a non-local process $H_j(t)$ with $j \neq i$. In this case, nodes receiving the transmission update their statuses to match that at node i if the status from node i is fresher. When no update is received or the update is staler than the current status at node m, i.e., the third case in (1), the age simply increases by one. Note that, for $t \in [n, n+1)$, since all (non-local) ages increase linearly with time, we can write $\Delta_j^{(m)}(t) = \Delta_j^{(m)}[n] + (t - n)$.

The scalar age update model in (1) can be straightforwardly extended to a vector age update model given by

$$\boldsymbol{\Delta}[n+1] = \boldsymbol{A}[n]\boldsymbol{\Delta}[n] + \mathbf{1}, \quad (2)$$

where $\boldsymbol{\Delta}[n] \in \mathbb{Z}^{N^2-N}$, $\boldsymbol{A}[n] \in \mathbb{Z}^{(N^2-N)\times(N^2-N)}$, and $\mathbf{1} \in \mathbb{Z}^{N^2-N}$ is a vector of ones. Note that the local ages $\Delta_i^{(i)}(t)$ are not included in $\boldsymbol{\Delta}[n]$ since they are always zero. From (1), it is clear that $\boldsymbol{A}[n]$ is a matrix with elements equal to zero or one. It is also evident that the rows of $\boldsymbol{A}[n]$ each have at most one element equal to one. An enumerated list of some additional relevant properties of $\boldsymbol{A}[n]$ is given in the Appendix. This model and the properties discussed in the Appendix will be useful in the analysis of the age metrics described below.

III. AGE METRICS AND SCHEDULES

In addition to the individual ages $\Delta_j^{(i)}(t)$ defined in Section II, we are interested in characterizing certain statistics of the ages across the network. The most common age statistic studied in the literature is the average age, which we will consider here as well. In the single-source single-monitor literature, the average age is computed as an average over time. Here, we first consider the *instantaneous* peak and *instantaneous* average age, where the peak and average are calculated over the node indices at a fixed value of t. These are only defined when all of the constituent ages are defined, i.e., every node in the network has received at least one status update for every process. Hence, we denote by $\bar{t}$ a time such that all ages $\Delta_j^{(i)}(t)$ are defined. Given Definition 1 and $\bar{t}$, we now define the instantaneous peak age at any point in time $t \geq \bar{t}$.

Definition 2 (Instantaneous peak age) For any $t \geq \bar{t}$, the *instantaneous peak age* is defined as

$$\Delta_{\text{peak}}(t) \triangleq \max \boldsymbol{\Delta}(t). \quad (3)$$

Note that t is fixed here and the maximum is computed over the N^2-N elements of the vector $\boldsymbol{\Delta}(t)$. Along the same lines, we define the instantaneous average age at any point $t \geq \bar{t}$ below.

Definition 3 (Instantaneous average age) For any $t \geq \bar{t}$, the

[1] A setting where multiple nodes can transmit simultaneously, subject to general interference constraints, is studied in [37].

instantaneous average age is defined as

$$\Delta_{\text{avg}}(t) \triangleq \frac{\mathbf{1}^\top \boldsymbol{\Delta}(t)}{N^2 - N}. \tag{4}$$

Note that the instantaneous average age represents the average of the $N^2 - N$ ages of the non-local statuses, i.e., the zero-age local statuses are not included in the average.

We define a *schedule* as a sequence of transmissions indexed by integer time, transmitting node i, and process j. A schedule can be equivalently expressed as a series of state update matrices $\boldsymbol{A}[n]$ in (2). One of the main contributions of this paper is in establishing fundamental limits on peak and average age for any schedule satisfying the assumptions listed in Section II. To illustrate the concept of a schedule, we provide an example for the three-node line network shown in Fig. 1. The age vector in this example is defined as

$$\boldsymbol{\Delta}(t)=[\Delta_2^{(1)}(t),\Delta_3^{(1)}(t),\Delta_1^{(2)}(t),\Delta_3^{(2)}(t),\Delta_1^{(3)}(t),\Delta_2^{(3)}(t)]^\top. \tag{5}$$

An undefined age in $\boldsymbol{\Delta}(t)$ is denoted by "$-$". The initial state $\boldsymbol{\Delta}[0]=[-,-,-,-,-,-]^\top$. The notation TS$k$ below corresponds to time slot k, a transmission occurring over $t\in(k-1,k]$.

TS1: (i=1, j=1) Suppose node 1 transmits a sample of its local process H_1 sampled at time t=0. This update from node 1 is received by node 2 at t=1, resulting in $\boldsymbol{\Delta}[1]=[-,-,1,-,-,-]^\top$.

TS2: ($i=2, j=1$) Suppose node 2 relays the update received in TS1 to node 3. This update (of process H_1) is received by node 3 at time $t=2$, resulting in $\boldsymbol{\Delta}[2]=[-,-,2,-,2,-]^\top$. Note that node 1 ignores this transmission since it contains a status update regarding its local process.

TS3: (i=2, j=2) Suppose node 2 now transmits a sample of its local process sampled at t=2. This update from node 2 is received by nodes 1 and 3 at t=3, resulting in $\boldsymbol{\Delta}[3]=[1,-,3,-,3,1]^\top$.

TS4: ($i=3, j=3$) Suppose node 3 transmits a sample of its local process sampled at time t=3. This update from node 3 is received by node 2 at time t=4, resulting in $\boldsymbol{\Delta}[4]=[2,-,4,1,4,2]^\top$.

TS5: ($i=2, j=3$) Similar to TS2, suppose node 2 relays the update received in TS4 to node 1. This update is received by node 1 at t=5, resulting in $\boldsymbol{\Delta}[5]=[3,2,5,2,5,3]^\top$.

Note that all nodes have received updates for all processes at $t=5$, hence we can set $\bar{t}=5$ and compute instantaneous peak and average ages for all $t\geq\bar{t}$. This schedule can naturally be repeated to construct a periodic schedule for all $t\geq 0$. Table 2 summarizes this periodic schedule and Fig. 2 plots the evolution of the six relevant ages in $\boldsymbol{\Delta}(t)$ as a function of t for three periods of this schedule. The instantaneous peak and average ages for this schedule are shown in the fourth subplot of Fig. 2. Observe that all of the ages and the instantaneous age metrics are periodic in this example due to the periodicity of the schedule.

Table 2. Example schedule for the three-node line network in Fig. 1.

time slot	transmitting node index i	transmitted process index j
TS1, TS6, $\cdots$	1	1
TS2, TS7, $\cdots$	2	1
TS3, TS8, $\cdots$	2	2
TS4, TS9, $\cdots$	3	3
TS5, TS10, $\cdots$	2	3

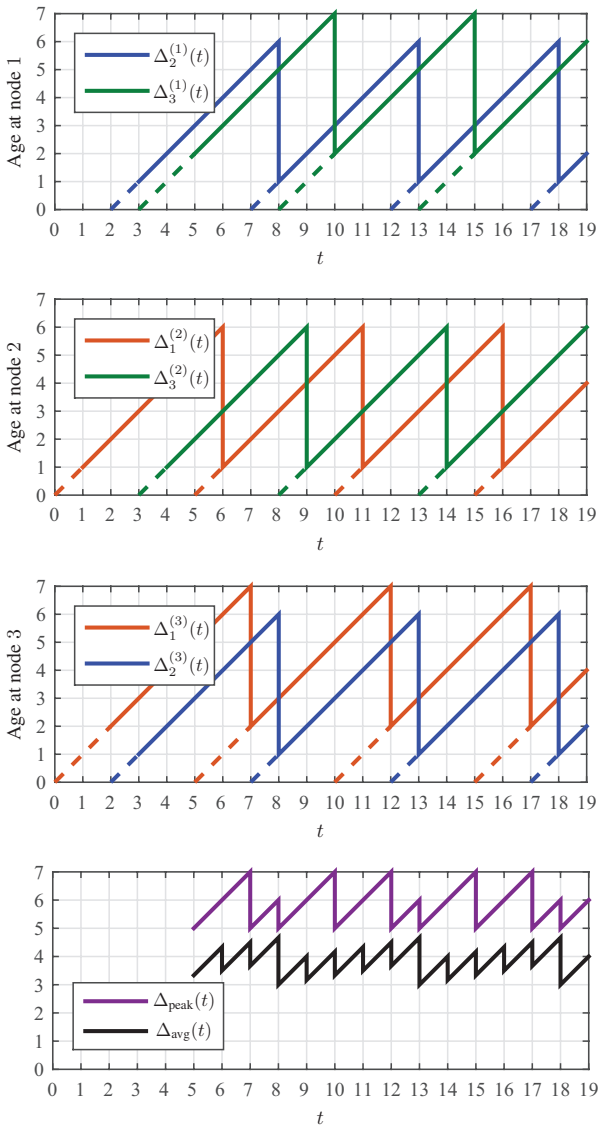

Fig. 2. Instantaneous age metrics for the example schedule for a three-node line network in Fig. 1.

We conclude this section with two additional age metrics of interest. The peak and average ages over an arbitrary time interval $t\in[t_0,t_1)$ with $\bar{t}\leq t_0<t_1$ are defined below.

Definition 4 (Peak age) The *peak age* for $t\in[t_0,t_1)$ with $\bar{t}\leq t_0<t_1$ is defined as

$$\Delta_{\text{peak}}(t_0,t_1) \triangleq \sup_{t_0\leq t<t_1} \Delta_{\text{peak}}(t). \tag{6}$$

Definition 5 (Average age) The *average age* for $t\in[t_0,t_1)$ with $\bar{t}\leq t_0<t_1$ is defined as

$$\Delta_{\text{avg}}(t_0,t_1) \triangleq \frac{1}{t_1-t_0}\int_{t_0}^{t_1} \Delta_{\text{avg}}(t)\,dt. \tag{7}$$

Note the supremum is used in (6) as $\Delta_{\text{peak}}(t)$ is a right-continuous piecewise linear function with discontinuities occurring at integer times when status updates are received. It is also worth mentioning here that Definition 4 differs from other definitions of "peak age" in the literature, e.g., [5], [36] define peak age as the average of the age peaks.

When we omit the time indices and use the notation Δ_{peak} and Δ_{avg}, this implies $t_0 = \bar{t}$ and $t_1 \to \infty$. In our running three-node example with the periodic schedule shown in Table 2 and the corresponding ages shown in Fig. 2, we can set $\bar{t} = 5$ and compute the peak and average ages as $\Delta_{\text{peak}} = 7$ and $\Delta_{\text{avg}} \approx 3.83$.

IV. FUNDAMENTAL BOUNDS ON THE INSTANTANEOUS AVERAGE AGE AND INSTANTANEOUS PEAK AGE

In this section we derive fundamental lower bounds on the *instantaneous* peak and average age metrics in Definitions 2 and 3 under the network and schedule assumptions stated in Section II. As these lower bounds rely on certain properties of the graph $\mathcal{G}$ describing the network, we first review these properties and define the concept of a "pseudo-leaf vertex".

A. Preliminary Definitions and Notation

A set $\mathcal{S} \subset \mathcal{V}$ of vertices in a graph is called a *dominating set* if every vertex not in $\mathcal{S}$ is adjacent to a vertex in $\mathcal{S}$ [38]. A minimum connected dominating set (MCDS) $\mathcal{S} \subset \mathcal{V}$ is a dominating set satisfying (i) the subgraph induced by $\mathcal{S}$ is connected and (ii) $\mathcal{S}$ has the smallest cardinality among all connected dominating sets of $\mathcal{G}$. The cardinality of any MCDS is called *connected domination number* of $\mathcal{G}$ and is denoted by γ_c. Although, in general, graphs do not have a unique MCDS, all MCDSs of a graph $\mathcal{G}$ have the same cardinality γ_c [39], [40]. Because every vertex *not* in a given MCDS is a one-hop neighbor of a vertex in the MCDS, i.e., $\mathcal{N}_1(\mathcal{S}) = \mathcal{G}$ for any MCDS $\mathcal{S}$, a collection of vertices in a given MCDS is often referred to as the "backbone" of the network in the context of broadcast routing for ad-hoc networks [40].

A *leaf* vertex of graph $\mathcal{G}$ is any vertex $i \in \mathcal{V}$ with degree of one. To the best of our knowledge, although the notion of a leaf vertex is well defined and commonly understood, there is no commonly accepted name for the following graph object. As such, we define a "pseudo-leaf vertex" below.

Definition 6 (Pseudo-leaf vertex) A vertex $i \in \mathcal{V}$ is a *pseudo-leaf vertex* if it is not a member of any MCDS. That is, $i \in \mathcal{V}$ is a pseudo-leaf vertex if $i \notin \mathcal{U}$ where

$$\mathcal{U} \triangleq \mathcal{S}_1 \cup \mathcal{S}_2 \cup \cdots \cup \mathcal{S}_M. \tag{8}$$

and $\mathcal{S}_k \subset \mathcal{V}$ for $k = \{1, 2, \cdots, M\}$ represent all possible MCDSs of $\mathcal{G}$. Further, we refer to the set of all pseudo-leaf vertices of $\mathcal{G}$ by $\mathcal{L} \triangleq \mathcal{V} - \mathcal{U}$.

Under this definition, every true leaf (i.e., every vertex with degree one) is also a pseudo-leaf. A graph may have additional non-leaf vertices that satisfy the conditions of a pseudo-leaf vertex.

Finally, recall the degree δ_i of node $i \in \mathcal{V}$ is defined as the the number of edges of vertex i or, equivalently, the cardinality of the number of one-hop neighbors of node i, i.e., $\delta_i = |\mathcal{N}_1(i)|$. We define the maximum degree of the graph as $\delta_{\max} = \max_{i \in \mathcal{V}} \delta_i$.

To illustrate the key ideas, consider the 5-node network in Fig. 3 which has two MCDSs, shown as $\mathcal{S}_1$ and $\mathcal{S}_2$, both with cardinality $\gamma_c = 2$. While only node 1 is a true leaf vertex, nodes 1 and 5 are pseudo-leaf vertices as they are not members of any MCDS. The cardinality of the pseudo-leaf set is $|\mathcal{L}| = 2$. The maximum degree of this graph is $\delta_{\max} = \delta_2 = 3$. These graph parameters play an important role in the bounds developed in the following section.

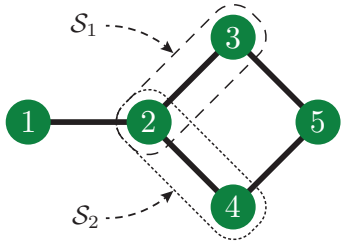

Fig. 3. A 5-node "pan" network. The two MCDS's are shown as $\mathcal{S}_1 = \{2, 3\}$ and $\mathcal{S}_2 = \{2, 4\}$ with cardinality $\gamma_c = 2$. The set of pseudo-leaf vertices is $\mathcal{L} = \mathcal{V} - (\mathcal{S}_1 \cup \mathcal{S}_2) = \{1, 5\}$.

B. Lower Bounds on Instantaneous Peak and Instantaneous Average Ages

This section presents lower bounds on the instantaneous peak and instantaneous average ages for all $t \geq \bar{t}$ such that all of the constituent ages in $\mathbf{\Delta}(t)$ are defined. To facilitate the presentation of these bounds, we first present the following Lemma to characterize the number of time slots required to update all $N^2 - N$ statuses in the network.

Lemma 1 (Number of time slots to update all statuses) Given $\mathcal{G}$ with $N = |\mathcal{V}|$, there exists a schedule that updates all $N^2 - N$ statuses in the network in $T^* \triangleq N\gamma_c + |\mathcal{L}|$ time slots. Moreover, any schedule of length $T = T^* - K$ time slots for $K \in \{0, 1, \cdots, T^*\}$ updates at most $N^2 - N - K$ statuses.

Proof: We first show the sufficiency of T^* time slots by construction. A schedule satisfying the conditions in Section II that updates all statuses in the network in T^* time slots can be constructed by flooding the status from each node throughout the network sequentially for $i = 1, \cdots, N$, i.e., the local status of node i is transmitted to its one-hop neighbors and this status is retransmitted in subsequent time slots to the remaining nodes in the network via an optimal flooding tree constructed from the graph's MCDS [41]. As shown in [41, Theorem 1], the number of time slots required to propagate an update of $H_i(t)$ throughout the network is $\gamma_c + \mathbb{1}_{i \in \mathcal{L}}$ where $\mathbb{1}_{i \in \mathcal{L}}$ is an indicator function equal to one when i is a pseudo-leaf vertex and equal to zero otherwise. It follows that a schedule that performs this optimal flooding sequentially for each node $i = 1, \cdots, N$ requires $|\mathcal{L}|(\gamma_c + 1) + (N - |\mathcal{L}|)\gamma_c = N\gamma_c + |\mathcal{L}|$ time slots to complete and updates all statuses in the network.

To show the second part of the lemma (which also establishes the necessity of T^* time slots to update all statuses), let $K = K_1 + \cdots + K_N$ with $K_i \in \{0, \cdots, \gamma_c + \mathbb{1}_{i \in \mathcal{L}}\}$ for all $i \in \{1, \cdots, N\}$. Let $\gamma_c + \mathbb{1}_{i \in \mathcal{L}} - K_i$ be the number of time slots allocated to propagate an update of $H_i(t)$. Observe that at least K_i nodes in the network do not receive an update of process $H_i(t)$. Hence at least $K = K_1 + \cdots + K_N$ statuses of the $N^2 - N$ total statuses in the network are not updated, which shows the desired result. □

Remarks:

1. A periodic version of a sequential flooding schedule of length T^* described in the proof of Lemma 1 is presented in Section V. Lemma 1 implies that there does not exist any periodic schedule with period less than T^* time slots such that all statuses are updated.
2. For a fixed number of nodes N, it can be shown that T^* obeys the inequality $N \leq T^* \leq N^2 - 2N + 2$, where the left side is true with equality for the complete graph and the right side is true with equality for the path graph. This suggests that some network topologies require on the order of N time slots to update all statuses in the network, whereas others require on the order of N^2 time slots to update all statuses in the network. Unsurprisingly, the network topology has a significant impact on the minimum number of time slots T^* required to update all statuses in the network.

Lemma 1 also implies certain properties about the state transition matrix for the discrete time model in (2). Given an initial age state of $\boldsymbol{\Delta}[n_0]$, the age vector at time $n \geq n_0,\ n \in \mathbb{Z}$ can be written as

$$\boldsymbol{\Delta}[n] = \boldsymbol{\Phi}[n, n_0]\boldsymbol{\Delta}[n_0] + \sum_{k=n_0}^{n-1} \boldsymbol{\Phi}[n, k+1]\mathbf{1} \tag{9}$$

where $\boldsymbol{\Phi}[n, m]$ is the discrete time time-varying state transition matrix defined in (17) in the Appendix. Note that a sequence of transmissions $\boldsymbol{A}[n_0], \cdots, \boldsymbol{A}[n-1]$ updates all statuses if and only if $\boldsymbol{\Phi}[n, n_0] = \boldsymbol{A}[n-1]\cdots\boldsymbol{A}[n_0] = 0$. From Lemma 1, since T^* transmissions are necessary to update all statuses in the network, $\boldsymbol{\Phi}[n, m]$ must have one or more non-zero rows for all $m \in \{n - T^* + 1, \cdots, n\}$. This implies there are always at least T^* non-zero terms in the summation of (9), which will be useful in the following results.

We now present the lower bounds on the instantaneous peak and instantaneous average age of information.

Theorem 1 (Lower bound on instantaneous peak age) The instantaneous peak age of information for any schedule at time $t \geq \bar{t}$ is lower bounded by

$$\Delta_{\text{peak}}(t) \geq \Delta^*_{\text{peak,inst}} \triangleq T^* = N\gamma_c + |\mathcal{L}|. \tag{10}$$

Proof: For $t \geq \bar{t}$ and $t \in [n, n+1)$, we can write

$$\begin{aligned}\Delta_{\text{peak}}(t) &= \max \boldsymbol{\Delta}[n] + (t-n),\\ &\geq \max \boldsymbol{\Delta}[n],\\ &\geq \boldsymbol{e}_i^\top \boldsymbol{\Delta}[n]\end{aligned}$$

for all $i \in \{1, \cdots, N^2 - N\}$, where $\boldsymbol{e}_i$ is the i th standard basis vector. From (9), we can set $n_0 = 0$ and $n \geq \bar{t} \geq T^*$ to write

$$\begin{aligned}\boldsymbol{\Delta}[n] &= \boldsymbol{\Phi}[n, 0]\boldsymbol{\Delta}[0] + \sum_{k=0}^{n-1} \boldsymbol{\Phi}[n, k+1]\mathbf{1},\\ &= \sum_{k=0}^{n-1} \boldsymbol{\Phi}[n, k+1]\mathbf{1},\\ &\geq \sum_{k=n-T^*}^{n-1} \boldsymbol{\Phi}[n, k+1]\mathbf{1}, \end{aligned}\tag{11}$$

where the equality follows from the fact that $n \geq \bar{t}$ implies all statuses have been updated by time n, i.e., $\boldsymbol{\Phi}[n, 0] = 0$, and the inequality follows from the fact that each term in the sum is non-negative. Observe that there are T^* terms in the sum and that that all $\boldsymbol{\Phi}[n, k+1]$ in the sum are non-zero. Hence, there must exist at least one i such that $\boldsymbol{e}_i^\top \boldsymbol{\Phi}[n, n-T^*+1]\mathbf{1} = 1$. Moreover, from Property VI in the Appendix, $\boldsymbol{e}_i^\top \boldsymbol{\Phi}[n, n-T^*+1]\mathbf{1} = 1$ implies $\boldsymbol{e}_i^\top \boldsymbol{\Phi}[n, k+1]\mathbf{1} = 1$ for all $k \in \{n - T^*, \cdots, n-1\}$. Hence, given i such that $\boldsymbol{e}_i^\top \boldsymbol{\Phi}[n, n-T^*+1]\mathbf{1} = 1$, we can write

$$\Delta_{\text{peak}}(t) \geq \boldsymbol{e}_i^\top \sum_{k=n-T^*}^{n-1} \boldsymbol{\Phi}[n, k+1]\mathbf{1} = T^*,$$

which shows the desired result. □

Theorem 2 (Lower bound on instantaneous average age) The instantaneous average age of information for any schedule is lower bounded by

$$\Delta_{\text{avg}}(t) \geq \Delta^*_{\text{avg,inst}} \triangleq \frac{\mathbf{1}^\top \boldsymbol{s}}{N^2 - N}, \tag{12}$$

where

$$\boldsymbol{s} \triangleq [\max(\boldsymbol{x}_1, \boldsymbol{y}_1), \max(\boldsymbol{x}_2, \boldsymbol{y}_2), \cdots, \max(\boldsymbol{x}_{T^*}, \boldsymbol{y}_{T^*})], \tag{13a}$$

$$\boldsymbol{x} \triangleq [N^2 - N, N^2 - N - \delta_{\max}, \cdots, N^2 - N - (T^* - 1)\delta_{\max}], \tag{13b}$$

$$\boldsymbol{y} \triangleq [T^*, T^* - 1, \cdots, 1]. \tag{13c}$$

Proof: Along the same lines as Theorem 1, for $t \geq \bar{t}$ and $t \in [n, n+1)$, we can write

$$\begin{aligned}\Delta_{\text{avg}}(t) &= \frac{\mathbf{1}^\top \boldsymbol{\Delta}[n]}{N^2 - N} + (t - n),\\ &\geq \frac{\mathbf{1}^\top \sum_{k=n-T^*}^{n-1} \boldsymbol{\Phi}[n, k+1]\mathbf{1}}{N^2 - N},\\ &= \frac{\sum_{k=n-T^*}^{n-1} s[n, k+1]}{N^2 - N},\end{aligned}$$

where the inequality follows from the fact that $t - n \geq 0$ and from following the same steps that led to (11). The term $s[n, k+1]$ denotes the number of non-zero elements, i.e., the number of statuses not updated, in $\boldsymbol{\Phi}[n, k+1]$. Observe that

(P1) $s[n, n] = N^2 - N$ from the fact that $\boldsymbol{\Phi}[n, n] = \boldsymbol{I}_{N^2-N}$.

(P2) $s[n, k+1] - s[n, k] \leq \delta_{\max}$ since at most $\delta_{\max}$ statuses can be updated in a time slot.

(P3) $s[n, n-T^*+K] \geq K$ for all $K \in \{1, \cdots, T^*\}$ from Lemma 1.

The minimal sequence $\boldsymbol{s} = [s[n, n], \cdots, s[n, n-T^*+1]]$ satisfying these properties can be constructed by first defining

$$\begin{aligned}\boldsymbol{x} &= [N^2 - N, N^2 - N - \delta_{\max}, \cdots, N^2 - N - (T^* - 1)\delta_{\max}],\\ \boldsymbol{y} &= [T^*, T^* - 1, \cdots, 1]\,.\end{aligned}$$

Note that $\boldsymbol{x}$ captures the constraints imposed by (P1) and (P2). Similarly $\boldsymbol{y}$ captures the constraints imposed by (P3). Then

$$\boldsymbol{s} = [\max(\boldsymbol{x}_1, \boldsymbol{y}_1), \max(\boldsymbol{x}_2, \boldsymbol{y}_2), \cdots, \max(\boldsymbol{x}_{T^*}, \boldsymbol{y}_{T^*})]$$

is the minimal sequence satisfying all of the properties, which establishes the lower bound and shows the desired result. □

Again referring to the three-node example in Section II, we have $N = 3$, $\gamma_c = 1$ and $|\mathcal{L}| = 2$. These graph parameters imply $\Delta^*_{\text{peak,inst}} = T^* = 5$ and it is evident from the fourth subplot in Fig. 2 that the example periodic schedule achieves this lower bound at times $t = \{5, 7, 8, 10, 12, 13, 15, 17, 18, \cdots\}$. For the lower bound on instantaneous average age in the three-node example, we have $\delta_{\max} = 2$ and we can calculate $\boldsymbol{x} = [6, 4, 2, 0, -2]$ and $\boldsymbol{y} = [5, 4, 3, 2, 1]$, resulting in $\boldsymbol{s} = [6, 4, 3, 2, 1]$. It follows then that $\Delta^*_{\text{avg,inst}}$=16/6 $\approx$ 2.67. Fig. 2 shows that the example periodic schedule reaches a minimum value of $\Delta_{\text{avg}}(t) = 3$ at times $t = \{8, 13, 18, \cdots\}$. Hence, unlike the instantaneous peak age, there is a gap between the bound and the minimum instantaneous age achieved in this example.

V. T^*-PERIODIC SCHEDULES

In this section, we consider the class of periodic schedules with period T^* as defined in Lemma 1. Such schedules have the property that $\boldsymbol{A}[n + kT^*] = \boldsymbol{A}[n]$ for all $n = 0, 1, \cdots, T^* - 1$ and all $k = 0, 1, \cdots$. The study of periodic schedules is motivated by the central goal of maintaining fresh statuses at each node in the network and by the recent results in [19], [20] where round robin schedules, i.e., schedules in which a series of transmissions is repeated, were shown to be optimal in terms of minimizing peak age. As noted earlier, a periodic schedule with period $T < T^*$ cannot update all statuses and, consequently, this class of schedules is of little interest since some ages are never defined. Periodic schedules with period $T > T^*$ are also of limited interest since the time between updates of at least some statuses (and, hence, the corresponding peak ages of these statuses) will be larger than necessary.

In the following, we first derive lower bounds on the peak and average age of information for all T^*-periodic schedules that update all statuses in each period. While the instantaneous age bounds developed in Section IV.B also serve as lower bounds over any period of a T^*-periodic schedule, our focus here is on the development of bounds on the peak and average age over a period of the schedule according to Definitions 4 and 5. We then present an algorithm for constructing a specific T^*-periodic schedule that updates all statuses in each period given any connected network topology. To analytically characterize the performance of this schedule, we upper bound its achieved age of information. This upper bound is then compared with the previously developed lower bounds to show the achieved peak age is tight with respect to the lower bound on peak age for any network topology and and size N. We also show that the achieved average age is asymptotically tight to the lower bound on average age as $N \to \infty$.

A. Lower Bounds on Peak and Average Age for T^*-Periodic Schedules

To establish fundamental limits for the peak and average age of information for T^*-periodic schedules, we first present the following useful Lemmas.

Lemma 2: For any T^*-periodic schedule that updates all $N^2 - N$ statuses, every status throughout the network is updated exactly once every T^* time slots.

Proof: Consider any T^*-periodic schedule that updates all of the $N^2 - N$ statuses. From Lemma 1 recall that $\gamma_c + \mathbb{1}_{i\in\mathcal{L}}$ time slots are required to propagate an update of $H_i(t)$ throughout the network. The first of these time slots corresponds to dissemination of a fresh update of $H_i(t)$ by node i and the remaining $\gamma_c + \mathbb{1}_{i\in\mathcal{L}} - 1$ time slots correspond to retransmissions of the status by nodes other than node i. While nodes in the network may receive multiple transmissions containing the status of the $H_i(t)$ process, the status at each node $j \neq i$ is only updated once per period since subsequent transmissions contain the same status and are redundant. □

As an example, consider the network shown in Fig. 3. Assume a schedule of TS1:(5,5), TS2:(4,5), and then TS3:(2,5), which corresponds to node 5 disseminating its status to all nodes in the network through the MCDS $\mathcal{S}_2$. Observe that node 3 receives transmissions regarding process $H_5(t)$ in both TS1 and TS3. These transmissions contain identical information, however. Hence node 3's status with regards to $H_5(t)$ is *updated only once* in TS1.

The main implication of Lemma 2 is that, for the class of T^*-periodic schedules that update all N^2–N statuses in the network, all status updates at each node occur with period T^*. No statuses are updated more frequently. So, over the interval $t\in[\nu,\nu+T^*)$, where ν is the time at which the status of process i is updated at node j, the age trajectory $\Delta_i^{(j)}(t)$ is simply $\Delta_i^{(j)}(t)$=t–ν+$\Delta_i^{(j)}$ where $\Delta_i^{(j)}$ is the age of the process $H_i(t)$ at the time node j is updated. In other words, for T^*-periodic schedules that update all N^2–N statuses in the network, each age trajectory $\Delta_i^{(j)}(t)$ is identical except for time shifts and the "age offsets" $\Delta_i^{(j)}$.

These time shifts and age offsets are illustrated for the T^*-periodic schedule for the three-node path network with $T^* = 5$ in Fig. 2 where $\Delta_2^{(1)} = 1$, $\Delta_3^{(1)} = 2$, $\Delta_3^{(2)} = 1$, $\Delta_1^{(2)} = 1$, $\Delta_2^{(3)} = 1$, and $\Delta_1^{(3)} = 2$. In general, note that the age offsets must satisfy $\Delta_i^{(j)} \geq d(i, j)$, where $d(i, j)$ is the distance in hops of the shortest path between nodes i and j. The following Lemma establishes an additional useful property of the age offsets $\Delta_i^{(j)}$ for networks with T^*-periodic schedules.

Lemma 3: Given $i \in \mathcal{V}$ and a T^*-periodic schedule that updates all of the $N^2 - N$ statuses,

$$\max_{j\in\mathcal{V}} \Delta_i^{(j)} \geq \gamma_c + \mathbb{1}_{i\in\mathcal{L}}.$$

Proof: This result follows directly from Lemma 1, which establishes that $\gamma_c + \mathbb{1}_{i\in\mathcal{L}}$ time slots are required to propagate an update of the status of process i throughout the network. Hence, given $i \in \mathcal{V}$, there always exists a node $j \in \mathcal{V}$ such that status updates regarding process i are received with an age offset of at least $\gamma_c + \mathbb{1}_{i\in\mathcal{L}}$. □

Lemmas 2 and 3 imply that, over any interval $[t_0, t_0 + T^*)$ and fixing $i \in \mathcal{V}$, there exists at least one age trajectory $\Delta_i^{(j)}(t)$ satisfying $\sup_{t\in[t_0,t_0+T^*)} \Delta_i^{(j)}(t) \geq \gamma_c + \mathbb{1}_{i\in\mathcal{L}} + T^*$. This forms the basis for Theorem 3 below, which establishes a lower bound on the peak age of information for T^*-periodic schedules.

Theorem 3 (Lower bound on Δ_{peak} of T^*-periodic schedules) The peak age of information for any T^*-periodic schedule over any interval $[t_0, t_0 + T^*)$ with $\bar{t} \leq t_0$ is lower bounded by

$$\Delta_{\text{peak}}(t_0, t_0 + T^*) \geq \Delta^*_{\text{peak},T^*} \triangleq T^* + \gamma_c + \mathbb{1}_{|\mathcal{L}|\geq 1}. \quad (14)$$

Proof: From Definition 4 with $\bar{t} \leq t_0$ we can write

$$\begin{aligned}\Delta_{\text{peak}}(t_0, t_0+T^*) &= \sup_{t_0 \leq t < t_0+T^*} \Delta_{\text{peak}}(t),\\ &= \sup_{t_0 \leq t < t_0+T^*} \max_{i,j \in \mathcal{V}} \Delta_i^{(j)}(t),\\ &= \max_{i,j \in \mathcal{V}} \sup_{t_0 \leq t < t_0+T^*} \Delta_i^{(j)}(t),\\ &\overset{(a)}{=} \max_{i,j \in \mathcal{V}} \left(T^* + \Delta_i^{(j)}\right),\\ &= T^* + \max_{i \in \mathcal{V}} \max_{j \in \mathcal{V}} \Delta_i^{(j)},\\ &\overset{(b)}{\geq} T^* + \gamma_c + \mathbb{1}_{|\mathcal{L}| \geq 1},\end{aligned}$$

where (a) follows from Lemma 2. Inequality (b) follows from Lemma 3 and the fact that $\max_{i \in \mathcal{V}} \mathbb{1}_{i \in \mathcal{L}} = \mathbb{1}_{|\mathcal{L}| \geq 1}$. □

The lower bound on peak age established in Theorem 3 has an intuitive interpretation. The T^* component is a consequence of the common period between updates for all statuses in the network. The $\gamma_c + \mathbb{1}_{|\mathcal{L}| \geq 1}$ component is a consequence of the maximum amount of time required to propagate a status throughout the network. Note that the only inequality in the derivation of the lower bound is from Lemma 3. This suggests a strategy for constructing T^*-periodic schedules to achieve the peak age bound with equality. This is discussed in detail in Section V.B. In fact, for the three node example in Section II, we can compute $\Delta^*_{\text{peak},T^*}=7$, which is achieved with equality as seen in Fig. 2.

Since $T^*=N\gamma_c+|\mathcal{L}|$, we can express the lower bound on peak age as $\Delta^*_{\text{peak},T^*} = (N+1)\gamma_c + |\mathcal{L}| + \mathbb{1}_{|\mathcal{L}| \geq 1}$. The role of γ_c, i.e., the connected domination number of the graph, is more evident in this expression. In graphs with $\gamma_c=1$, e.g., a star graph or a complete graph, $\Delta^*_{\text{peak},T^*} \sim \mathcal{O}(N)$. In graphs where $\gamma_c \sim \mathcal{O}(N)$, e.g., a path graph or a ring graph, $\Delta^*_{\text{peak},T^*} \sim \mathcal{O}(N^2)$.

The following theorem establishes a lower bound on the average age for networks with T^*-periodic schedules.

Theorem 4 (Lower bound on Δ_{avg} of T^*-periodic schedules) The average age of information for any T^*-periodic schedule over any interval $[t_0, t_0+T^*)$ with $\bar{t} \leq t_0$ is lower bounded by

$$\Delta_{\text{avg},T^*}(t_0, t_0+T^*) \geq \Delta^*_{\text{avg},T^*} \triangleq \frac{T^*}{2} + \bar{d}, \tag{15}$$

where $\bar{d} \triangleq \frac{1}{N^2-N} \sum d(i,j)$ is the average distance of the network.

Proof: From Definition 3 and Definition 5, we can write

$$\begin{aligned}\Delta_{\text{avg},T^*}(t_0,t_0+T^*) &= \frac{1}{T^*(N^2-N)} \sum_{\substack{i,j \in \mathcal{V}\\ j \neq i}} \int_{t_0}^{t_0+T^*} \Delta_i^{(j)}(t)dt,\\ &\overset{(a)}{=} \frac{1}{T^*(N^2-N)} \sum_{\substack{i,j \in \mathcal{V}\\ j \neq i}} \int_{0}^{T^*} (\Delta_i^{(j)}+t)dt,\\ &\overset{(b)}{\geq} \frac{1}{N^2-N} \sum_{\substack{i,j \in \mathcal{V}\\ j \neq i}} \left(d(i,j)+\frac{T^*}{2}\right),\\ &= \frac{T^*}{2} + \bar{d},\end{aligned}$$

where (a) follows from Lemma 2 and (b) follows from the fact that $\Delta_i^{(j)} \geq d(i,j)$. □

Like the lower bound on peak age, the lower bound on average age has an intuitive interpretation. The $T^*/2$ component is a consequence of the common period between updates for all statuses in the network. The $\bar{d}$ component is a consequence of the average amount of time required to propagate statuses throughout the network. For the three node example in Section II, the lower bound on average age can be calculated as $\Delta^*_{\text{avg},T^*} = \frac{5}{2}+\frac{8}{6} \approx 3.83$, which is achieved with equality since $\Delta_i^{(j)}=d(i,j)$ for all i,j in this case. This bound is not achievable in general, however, since it is not always possible to achieve $\Delta_i^{(j)}=d(i,j)$ for all i,j (see, e.g., an N=5 node path network).

B. Algorithm for T^*-Periodic Schedule Design

This section formalizes the main idea suggested by Lemma 1 to develop a sequential flooding algorithm that generates a periodic minimum-length schedule for a given network topology. Recall from the proof of Lemma 1 that propagation of $H_i(t)$ throughout the network can be accomplished with an initial transmission by node i of its zero-delay status update, and then subsequent transmissions that propagate that status update to remaining nodes via multiple hops from nodes in a MCDS. By repeating this approach and disseminating status updates from each of the N processes in turn, a length T^* schedule emerges which can then be repeated in a periodic fashion to continuously propagate fresh status updates throughout the network. An algorithm that details this approach is summarized in Algorithm 1, where Depth-First Search$(\mathcal{G}[\mathcal{S}], i)$ denotes an ordered list of vertices generated by performing an in-order depth-first search of the graph induced by $\mathcal{S}$ where the search starts at root vertex i.

Observe that the schedule generated by Algorithm 1 obeys the following properties: (i) It uses exactly T^* transmissions to update all tables throughout the network, (ii) it is periodic, and (iii) all statuses throughout the network get updated exactly once during each period. While our algorithm makes use of the depth-first search to traverse the graph induced by the MCDS, we note that an alternate graph traversal approach (e.g., breadth-first search) could be used here. Moreover, the specific choice of graph search used does not impact the bounds presented below.

C. Guaranteed Peak and Average Age for Schedules Generated by Algorithm 1

Since Algorithm 1 implements optimal sequential flooding, it achieves $\max_{j \in \mathcal{V}} \Delta_i^{(j)} = \gamma_c + \mathbb{1}_{i \in \mathcal{L}}$ for all $i \in \mathcal{V}$. Hence, the inequality in Lemma 3 used in the derivation of the lower bound on peak age in Theorem 3 can be achieved with equality. Unlike peak age, the average age achieved by schedules generated by Algorithm 1 will not coincide with the lower bound in Theorem 4 since the inequality $\Delta_i^{(j)} \geq d(i,j)$ is not tight. In this section, we derive an *upper* bound, i.e., a performance guarantee, on the average age achieved by schedules generated by Algorithm 1. This is followed by a characterization of the gap between the upper and lower bounds developed previously.

Theorem 5 (Upper bound on achievable average age) For any interval $[t_0, t_0+T^*)$ with $t_0 \geq \bar{t}$, the average age achieved by

Algorithm 1: Schedule design to disseminate status updates throughout the network.

Input: $\mathcal{G}(\mathcal{V},\mathcal{E})$.
Output: Schedule $\boldsymbol{\pi} \in \mathbb{R}^{T^*\times 2}$.
Initialization:
$t \leftarrow 0$,
$j \leftarrow 1$.
Part I: Build one period of the schedule.
for node $i = 1 : N$ **do**
  if $\exists$ MCDS $\bar{\mathcal{S}}$ s.t. $i \in \bar{\mathcal{S}}$ **then**
    $\mathcal{S} \leftarrow \bar{\mathcal{S}}$.
  else
    $\mathcal{S} \leftarrow \bar{\mathcal{S}} \cup \{i\}$, for any MCDS $\bar{\mathcal{S}} \subset \mathcal{V}$.
  end
  $\mathcal{S}_{\text{sorted}} =$ Depth-First Search$(\mathcal{G}[\mathcal{S}], i)$.
  for $k = 1 : |\mathcal{S}_{\text{sorted}}|$ **do**
    $\boldsymbol{\pi}(j,1) = \mathcal{S}_{\text{sorted}}(k)$,
    $\boldsymbol{\pi}(j,2) = H_i^{(\boldsymbol{\pi}(j,1))}(t)$,
    $j \leftarrow j+1$.
  end
end
Part II: Status update dissemination.
for $m = 1 : T^*$ **do**
  $\boldsymbol{\pi}(m,1)$ disseminates $\boldsymbol{\pi}(m,2)$,
  $t \leftarrow t+1$.
end
Part III: Repeat from Part II.

the schedule generated by Algorithm 1 satisfies

$$\Delta_{\text{avg,Alg.1}}(t_0, t_0+T^*) \le \Delta_{\text{avg,ub}} \triangleq \frac{T^*}{2} + \gamma_c + \frac{|\mathcal{L}|}{N}. \quad (16)$$

Proof: The proof follows from the fact that Algorithm 1 implements optimal sequential flooding with age offsets satisfying $d(i,j) \le \Delta_i^{(j)} \le \gamma_c + \mathbb{1}_{i\in\mathcal{L}}$. The lower bound on $\Delta_i^{(j)}$ holds for any schedule and was used in Theorem 4. Using the upper bound here, we can write

$$\begin{aligned}\Delta_{\text{avg,Alg.1}}(t_0, t_0+T^*) &\le \frac{1}{N^2-N}\sum_{\substack{i,j\in\mathcal{V}\\ j\neq i}} \frac{T^*}{2} + \gamma_c + \mathbb{1}_{i\in\mathcal{L}},\\ &= \frac{T^*}{2} + \gamma_c + \frac{1}{N^2-N}\sum_{\substack{i,j\in\mathcal{V}\\ j\neq i}} \mathbb{1}_{i\in\mathcal{L}},\\ &= \frac{T^*}{2} + \gamma_c + \frac{|\mathcal{L}|}{N}.\end{aligned}$$

□

Corollary 1 below characterizes the gap between the upper bound on the achieved average age in Theorem 5 and the average age lower bound in Theorem 4.

Corollary 1: The gap between upper bound on the average age achieved by the schedule generated by Algorithm 1 and the lower bound in Theorem 4 satisfies $\Delta_{\text{avg,ub}} - \Delta^*_{\text{avg},T^*} < N-2$.

Proof: From (15) and (16) we can write

$$\begin{aligned}\Delta_{\text{avg,ub}} - \Delta^*_{\text{avg},T^*} &= \gamma_c + \frac{|\mathcal{L}|}{N} - \bar{d},\\ &\overset{(a)}{\le} N - 2 + \frac{|\mathcal{L}|}{N} - \bar{d},\\ &\overset{(b)}{<} N - 1 - \bar{d},\\ &\overset{(c)}{<} N - 2,\end{aligned}$$

where (a) is from $\gamma_c \le N-2$, (b) is from $\frac{|\mathcal{L}|}{N} < 1$, and (c) is from $\bar{d} \ge 1$. All of these inequalities are general properties of graphs. □

The result presented by Corollary 1 shows that the gap grows at most linearly with respect to the number of nodes, N. Next, we show that the schedule generated by Algorithm 1 is asymptotically optimal in terms of the average age when N is large.

Corollary 2: The ratio of the average age achieved by the schedule generated by Algorithm 1 to the average age lower bound in Theorem 4 goes to one as $N \to \infty$.

Proof: We can use the average age upper bound of Theorem 5 and a sandwiching argument to prove the desired result. From (15) and (16) we can write

$$\begin{aligned}\lim_{N\to\infty} \frac{\Delta_{\text{avg,ub}}}{\Delta^*_{\text{avg},T^*}} &= \lim_{N\to\infty} \frac{\frac{T^*}{2} + \gamma_c + \frac{|\mathcal{L}|}{N}}{\frac{T^*}{2} + \bar{d}},\\ &= 1 + \lim_{N\to\infty} \frac{\gamma_c + \frac{|\mathcal{L}|}{N} - \bar{d}}{\frac{T^*}{2} + \bar{d}},\\ &= 1 + \lim_{N\to\infty} \frac{\frac{T^*}{N} - \bar{d}}{\frac{T^*}{2} + \bar{d}},\\ &\le 1 + \lim_{N\to\infty} \frac{2}{N} = 1,\end{aligned}$$

where the inequality in the last step results from the fact that $\frac{\frac{T^*}{N}-\bar{d}}{\frac{T^*}{2}+\bar{d}} \le (\frac{T^*}{N})/(\frac{T^*}{2})$ since $\bar{d} \ge 0$. Moreover, since $\frac{\Delta_{\text{avg,Alg.1}}(t_0,t_0+T^*)}{\Delta^*_{\text{avg},T^*}} \ge 1$ for all N, we have $\lim_{N\to\infty} \frac{\Delta_{\text{avg,Alg.1}}(t_0,t_0+T^*)}{\Delta^*_{\text{avg},T^*}} = 1$. □

While the upper bound in Theorem 5 makes the pessimistic assumption that all nodes have a worst-case age offset of $\Delta_i^{(j)} = \gamma_c + \mathbb{1}_{i\in\mathcal{L}}$, the result in Corollary 2 shows that this pessimistic assumption is inconsequential asymptotically.

VI. NUMERICAL RESULTS

This section presents numerical examples that serve to illustrate and verify the bounds on peak and average age in Sections IV and V, and that permit comparison with the peak and average age achieved by the schedule generated by Algorithm 1.

A. Application of Bounds to Canonical Graph Topologies

To illustrate computation of the various bounds, and to observe the impact of topology on the AoI, we list in Table 3 the bounds for some canonical graph topologies [42], [43] as a function of the number of nodes N. The table shows that for topologies like complete and star where $\gamma_c \ll N$ and $\delta_{\max}$ is of order

of N, the age of information is of order of N, too. On the other hand for topologies like cycle, path, and pan where γ_c is of order of N and $\delta_{\max} \ll N$, the age of information is of order of N^2.

B. All Connected Topologies with $3 \leq N \leq 9$ Nodes

For every connected network topology with $3 \leq N \leq 9$ nodes, we make use of a database [44] that exhaustively enumerates all connected network topologies with isomorphs removed.

B.1 Instantaneous Peak and Instantaneous Average Age

Fig. 4 shows the lower bounds on the instantaneous peak and instantaneous average age of information in Theorems 1 and 2, as well as the minimum instantaneous peak and instantaneous average age achieved by the periodic schedule generated by Algorithm 1. For every graph topology in the database, the minimum instantaneous peak age achieved by Algorithm 1 is equal to the lower bound on instantaneous peak age, i.e., $\min_{t \geq \bar{t}} \Delta_{\text{peak,Alg.1}}(t) = \Delta^*_{\text{peak,inst}}$, so Theorem 1 serves as a tight lower bound on schedules generated by Algorithm 1. Meanwhile, for the instantaneous average age, there is generally a gap between the lower bound from Theorem 2 and the minimum instantaneous average age achieved by Algorithm 1. To investigate this gap, we consider the ratio computed by dividing the achieved minimum instantaneous average age by the lower bound, and have found that for all networks with a connected topology having $3 \leq N \leq 9$ nodes, schedules generated by Algorithm 1 obey

$$1 \leq \frac{\min_{t \geq \bar{t}} \Delta_{\text{avg,Alg.1},k}(t)}{\Delta^*_{\text{avg,inst}}(k)} \leq 1.783,$$

$$\frac{1}{K} \sum_{k=1}^{K} \frac{\min_{t \geq \bar{t}} \Delta_{\text{avg,Alg.1},k}(t)}{\Delta^*_{\text{avg,inst}}(k)} \approx 1.563,$$

where we index by network topology $k = \{1, \cdots, K\}$ with $K = 273,191$ representing the total number of such networks. This leads to the following three observations: (i) only for the fully-connected network topologies, where $\gamma_c = 1$, Algorithm 1 achieves the lower bound on instantaneous average age, (ii) in the worst case, Algorithm 1 yields a minimum instantaneous average age that is 78.3% above the lower bound, and (iii) averaging over all K topologies, Algorithm 1 is 56.3% above the lower bound.

B.2 Peak and Average Age for T^*-periodic Schedules

Fig. 5 compares the lower bounds on peak and average age of information for T^*-periodic schedules in Theorems 3 and 4 with the peak and average age of information achieved by the schedule generated by Algorithm 1, as well as the upper bound on achievable average age for schedules generated by Algorithm 1 given by Theorem 5. The results verify that, indeed, $\Delta_{\text{peak,Alg.1}} = \Delta^*_{\text{peak},T^*}$ for all of the considered topologies. In addition, for all $k = \{1, \cdots, K\}$ we have $\Delta^*_{\text{avg},T^*}(k) \leq \Delta_{\text{avg,Alg.1}}(k) \leq \Delta_{\text{avg,ub}}(k)$, thus verifying Theorems 4 and 5. Moreover, the numerical results on average age obey

$$1 \leq \frac{\Delta_{\text{avg,Alg.1}}(k)}{\Delta^*_{\text{avg},T^*}(k)} \leq 1.035, \quad \frac{1}{K} \sum_{k=1}^{K} \frac{\Delta_{\text{avg,Alg.1}}(k)}{\Delta^*_{\text{avg},T^*}(k)} \approx 1.008,$$

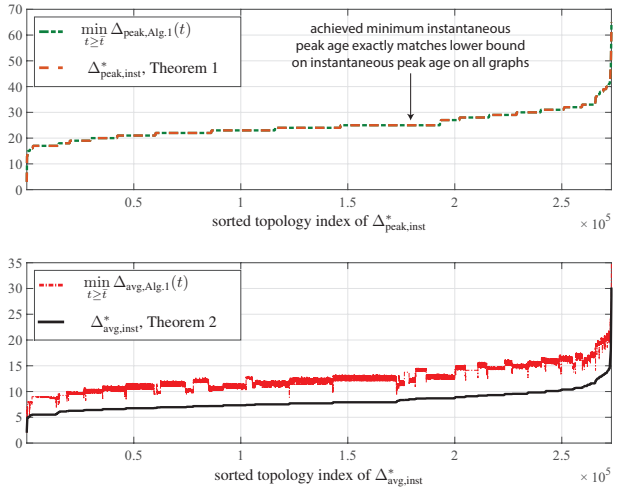

Fig. 4. Lower bounds on the instantaneous peak and instantaneous average age of information in Theorems 1 and 2, compared to the minimum instantaneous peak and instantaneous average age of information achieved by the schedule in Algorithm 1.

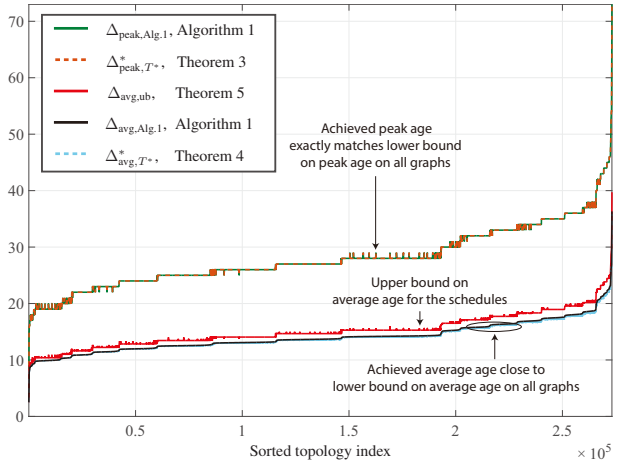

Fig. 5. Achievable peak and average age for the schedule generated by Algorithm 1 for all networks with a connected topology and $3 \leq N \leq 9$. The age values are arranged in an increasing order for the achievable average age.

where as above we index by network topology $k = \{1, \cdots, K\}$. Again, this leads to three observations: (i) For network topologies with small γ_c, the average age achieved by Algorithm 1 is very close or matches the lower bound on the average age for T^*-periodic schedules, (ii) in the worst case, Algorithm 1 yields a minimum average age that is 3.5% above the lower bound, and (iii) averaging over all K topologies, Algorithm 1 is 0.8% above the lower bound. Moreover, we note that Corollary 2 implies that for schedules generated by Algorithm 1, these ratios approach 1 as N grows large. Finally, we note that the achievable peak age is roughly twice the achievable average age, i.e., $\Delta_{\text{peak,Alg.1}}(k) \approx 2\Delta_{\text{avg,Alg.1}}(k)$ for all $k = \{1, \cdots, K\}$.

Fig. 6 represents the gap between the average age achieved by the schedule generated by Algorithm 1 and the lower bound on the average age in Theorem 4 compared to the upper bound of $N-2$ in Corollary 1 for all of the connected network topologies with $3 \leq N \leq 9$. The results show that the upper bound in Corollary 1 is conservative and for most of the topologies the

Table 3. AoI bounds for canonical graph topologies.

topology	$\lvert\mathcal{L}\rvert$	γ_c	$\delta_{\max}$	$\bar{d}$	$\Delta^*_{\text{peak,inst}}$	$\Delta^*_{\text{peak},T^*}$	$\Delta^*_{\text{avg,inst}}$	$\Delta^*_{\text{avg},T^*}$
complete	0	1	N–1	1	N	N+1	$\frac{N+1}{2}$	$\frac{N+2}{2}$
cycle	0	N–2	2	$\frac{N^2}{4(N-1)}$, N even $\frac{N+1}{4}$, N odd	N^2–2N	N^2–N–2	$\frac{N^3-4N^2+6N-1}{2(N-1)}$	$\frac{2N^3-5N^2+4N}{4(N-1)}$, N even $\frac{2N^2-3N+1}{4}$, N odd
path	2	N–2	2	$\frac{N+1}{3}$	N^2–2N+2	N^2–N+1	$\frac{N^4-4N^3+10N^2-13N+8}{2N(N-1)}$	$\frac{3N^2-4N+8}{6}$
star	N–1	1	N–1	$\frac{2(N-1)}{N}$	2N–1	2N+1	$\frac{N^3+N^2-2N+2}{2N(N-1)}$	$\frac{2N^2+3N-4}{2N}$
pan ($N>5$)	1	N–3	3	$\frac{N^3-N^2+6N-8}{4N(N-1)}$, N even $\frac{N^2+7}{4N}$, N odd	N^2–3N+1	N^2–2N–1	$\frac{N^4-6N^3+14N^2-9N+2}{2N(N-1)}$	$\frac{2N^4-7N^3+7N^2+4N-8}{4N(N-1)}$, N even $\frac{2N^3-5N^2+2N+7}{4N}$, N odd

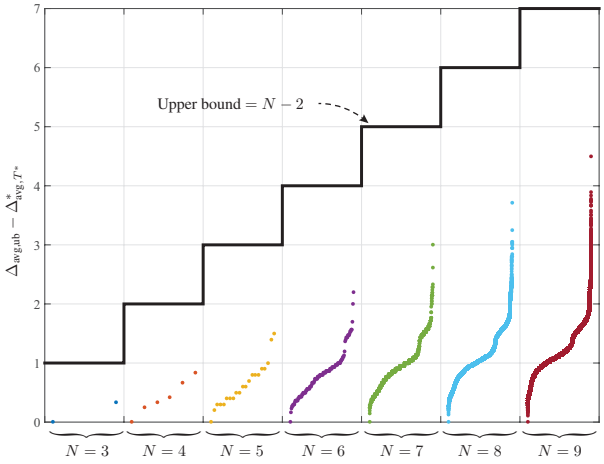

Fig. 6. The gap between the upper bound on the average age achieved by the schedule generated by Algorithm 1 and the lower bound on the average age for all networks with a connected topology and $3 \leq N \leq 9$ compared to the upper bound in Corollary 1. Observe that the dots represent the network topologies for each N.

schedule generated by Algorithm 1 achieves close to the lower bound on the average age.

VII. CONCLUSION

This paper studied the age of information problem in a general multi-source multi-hop partially-connected wireless network with nodes communicating over time slotted transmissions. We derived fundamental results that lower bound the performance of any status update dissemination schedule in terms of the peak and average age of information metrics. To the best of our knowledge, this is the first work to consider the impact of network topology on the age of information, and we found that the AoI depends on fundamental graph parameters such as the connected domination number and average shortest path length. We presented an algorithm that generates minimum-length periodic schedules for dissemination of the status updates among the nodes in the network with a connected topology. We derived upper bounds on the peak and average age achieved by the schedules designed by the proposed algorithm. A potentially interesting future direction of this work is to study the impact of network coding on the age of information.

APPENDIX A
PROPERTIES OF $\boldsymbol{A}[n]$ AND $\boldsymbol{\Phi}[n,m]$

From (1), we can derive several useful properties of the state update matrices $\boldsymbol{A}[n]$ in (2). Since the ordering of the ages in $\boldsymbol{\Delta}(t)$ is not specified, let $r(i,j)$ correspond to the row position of $\Delta_i^{(j)}(t)$ in $\boldsymbol{\Delta}(t)$. The following properties of $\boldsymbol{A}[n]$ are straightforward to verify:

I. Each row of $\boldsymbol{A}[n]$ is either equal to zero or has a single non-zero entry equal to one.

II. There are exactly δ_i all-zero rows in $\boldsymbol{A}[n]$ when node i transmits a status update of its local process $H_i(t)$. These occur at row indices $r(i,m)$ for all $m \in \mathcal{N}_1(i)$. This corresponds to case 1 in (1).

III. There are no all-zero rows in $\boldsymbol{A}[n]$ when node i transmits a status update of a non-local process $H_j(t)$. There are, however, at most δ_i rows set to match row $r(j,i)$ of $\boldsymbol{A}[n]$, i.e., the transposed $r(j,i)$'th standard unit vector, when node i transmits a status update of a non-local process $H_j(t)$. This corresponds to case 2 in (1).

IV. If node m does not receive an update on process j at time $n+1$, then row $r(j,m)$ of $\boldsymbol{A}[n]$ is equal to the transposed $r(j,m)$'th standard unit vector. This corresponds to case 3 in (1).

V. For any integer $m \in \mathbb{N}$, $(\boldsymbol{A}[n])^m = \boldsymbol{A}[n]$.

VI. From Lemma 1 we have $\mathbf{1}^\top \boldsymbol{\Phi}[n,m]\mathbf{1} \geq 1$ for $n-T^*+1 \leq m \leq n$, where

$$\boldsymbol{\Phi}[n,m] = \begin{cases} \boldsymbol{A}[n-1]\boldsymbol{A}[n-2]\cdots\boldsymbol{A}[m] & n-m>0 \\ \boldsymbol{I}_{N^2-N} & n-m=0 \\ \text{undefined} & n-m<0 \end{cases}. \quad (17)$$

REFERENCES

[1] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing age of information in vehicular networks," in *Proc. IEEE SECON*, June 2011, pp. 350–358.

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

[3] S. K. Kaul, R. D. Yates, and M. Gruteser, "Status updates through queues," in *Proc. CISS*, Mar. 2012, pp. 1–6.

[4] L. Huang and E. Modiano, "Optimizing age-of-information in a multi-class queueing system," in *Proc. IEEE ISIT*, June 2015, pp. 1681–1685.

[5] M. Costa, M. Codreanu, and A. Ephremides, "On the age of information in status update systems with packet management," *IEEE Trans. Inf. Theory*, vol. 62, no. 4, pp. 1897–1910, Feb. 2016.

[6] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," *IEEE Trans. Inf. Theory*, vol. 63, no. 11, pp. 7492–7508, Aug. 2017.

[7] S. Farazi, A. G. Klein, and D. R. Brown III, "Average age of information for status update systems with an energy harvesting server," in *Proc. IEEE INFOCOM WORKKSHPS*, Apr. 2018, pp. 112–117.
[8] S. Farazi, A. G. Klein, and D. R. Brown III, "Age of information in energy harvesting status update systems: When to preempt in service?" in *Proc. IEEE ISIT*, June 2018, pp. 2436–2440.
[9] A. Kosta, N. Pappas, V. Angelakis, *et al.*, "Age of information: A new concept, metric, and tool," *Foundations and Trends® in Networking*, vol. 12, no. 3, pp. 162–259, 2017.
[10] S. Farazi, A. G. Klein, and D. R. Brown III, "On the average staleness of global channel state information in wireless networks with random transmit node selection," in *Proc. IEEE ICASSP*, Mar. 2016, pp. 3621–3625.
[11] I. Kadota, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Minimizing the age of information in broadcast wireless networks," in *Proc. Allerton Conference*, Sept. 2016, pp. 844–851.
[12] Y.-P. Hsu, E. Modiano, and L. Duan, "Age of information: Design and analysis of optimal scheduling algorithms," in *Proc. IEEE ISIT*, Aug. 2017, pp. 561–565.
[13] A. G. Klein, S. Farazi, W. He, and D. R. Brown III, "Staleness bounds and efficient protocols for dissemination of global channel state information," *IEEE Trans. Wireless Commun.*, vol. 16, no. 9, pp. 5732–5746, Sept. 2017.
[14] S. Farazi, A. G. Klein, and D. R. Brown III, "Bounds on the age of information for global channel state dissemination in fully-connected networks," in *Proc. ICCCN*, July 2017, pp. 1–7.
[15] V. Tripathi and S. Moharir, "Age of information in multi-source systems," in *Proc. IEEE Global Telecomm. Conf. (GLOBECOM)*, Dec. 2017, pp. 1–6.
[16] Q. He, D. Yuan, and A. Ephremides, "Optimal link scheduling for age minimization in wireless systems," vol. 64, no. 7, pp. 5381–5394, July 2018.
[17] R. D. Yates and S. K. Kaul, "Status updates over unreliable multiaccess channels," in *Proc. IEEE ISIT*, Aug. 2017, pp. 331–335.
[18] J. Zhong, R. D. Yates, and E. Soljanin, "Multicast with prioritized delivery: How fresh is your data?" in *in Proc. IEEE Workshop SPAWC*, June 2018.
[19] Z. Jiang, B. Krishnamachari, X. Zheng, S. Zhou, and Z. Niu, "Decentralized status update for age of information optimization in wireless multiaccess channels," in *Proc. IEEE ISIT*, June 2018, pp. 2276–2280.
[20] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," *IEEE/ACM Trans. Netw.*, vol. 26, no. 6, pp. 2637–2650, 2018.
[21] R. Talak, S. Karaman, and E. Modiano, "Distributed scheduling algorithms for optimizing information freshness in wireless networks," in *Proc. IEEE Workshop SPAWC*, June 2018.
[22] I. Kadota, A. Sinha, and E. Modiano, "Scheduling algorithms for optimizing age of information in wireless networks with throughput constraints," 2018, submitted to *IEEE/ACM Trans. Netw.*, In review.
[23] M. Costa and Y. E. Sagduyu, "Age of information with network coding," *Ad Hoc Networks*, vol. 86, pp. 15–22, 2019.
[24] S. Kaul, R. Yates, and M. Gruteser, "On piggybacking in vehicular networks," in *Proc. GLOBECOM*, Dec. 2011, pp. 1–5.
[25] J. Selen, Y. Nazarathy, L. L. Andrew, and H. L. Vu, "The age of information in gossip networks," in *Proc. Springer ASMTA*. 2013, pp. 364–379.
[26] S. Farazi, D. R. Brown III, and A. G. Klein, "On global channel state estimation and dissemination in ring networks," in *Proc. Asilomar Conf. on Signals, Systems, and Computers*, Nov. 2016, pp. 1122–1127.
[27] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Optimizing data freshness, throughput, and delay in multi-server information-update systems," in *Proc. IEEE ISIT*, July 2016, pp. 2569–2573.
[28] A. Arafa and S. Ulukus, "Age-minimal transmission in energy harvesting two-hop networks," in *Proc. IEEE GLOBECOM*, Dec. 2017, pp. 1–6.
[29] R. D. Yates, "Age of information in a network of preemptive servers," in *Proc. IEEE INFOCOM WORKSHOPS*, Apr. 2018, pp. 118–123.
[30] R. D. Yates, "The age of information in networks: Moments, distributions, and sampling," *arXiv preprint arXiv:1806.03487*, 2018.
[31] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Age-optimal information updates in multihop networks," in *Proc. IEEE ISIT*, Aug. 2017, pp. 576–580.
[32] A. M. Bedewy, Y. Sun, and N. B. Shroff, , "Age-optimal information updates in multihop networks," 2017.
[33] A. M. Bedewy, Y. Sun, and N. B. Shroff, "The age of information in multihop networks," *CoRR, http://arxiv.org/abs/1712.10061*, 2017.
[34] S. Farazi, A. G. Klein, J. A. McNeill, and D. R. Brown III, "On the age of information in multi-source multi-hop wireless status update networks," in *Proc. IEEE Workshop SPAWC*, June 2018.
[35] T. Shreedhar, S. K. Kaul, and R. D. Yates, "Acp: An end-to-end transport protocol for delivering fresh updates in the internet-of-things," *arXiv preprint arXiv:1811.03353*, 2018.
[36] R. Talak, S. Karaman, and E. Modiano, "Minimizing age-of-information in multi-hop wireless networks," in *Proc. Allerton Conference*, Oct. 2017, pp. 486–493.
[37] S. Farazi, A. G. Klein, and D. R. Brown III, "Fundamental bounds on the age of information in general multi-hop interference networks," in *Proc. INFOCOM Workshop*, Apr. 2019.
[38] D. B. West, *Introduction to Graph Theory*, 2nd ed. Prentice Hall, 2000.
[39] E. Sampathkumar and H. Walikar, "The connected domination number of a graph," *J. Math. Phys*, 1979.
[40] B. Das and V. Bharghavan, "Routing in ad-hoc networks using minimum connected dominating sets," in *Proc. IEEE ICC*. 1997, pp. 376–380.
[41] H. Lim and C. Kim, "Flooding in wireless ad hoc networks," *Comput. Commun.*, vol. 24, no. 3-4, pp. 353–363, 2001.
[42] A. Brandstadt, J. P. Spinrad, *et al.*, *Graph classes: A survey*. Siam, 1999, vol. 3.
[43] J. A. Gallian, "A dynamic survey of graph labeling," *Electron. J. Combinatorics*, vol. 16, no. 6, pp. 1–219, 2009.
[44] B. McKay, "Database of connected simple graphs." [Online]. Available: http://users.cecs.anu.edu.au/~bdm/data/graphs.html

Shahab Farazi received the B.S. and M.S. degrees in Electrical Engineering from K.N. Toosi University of Technology and Iran University of Science and Technology, Tehran, Iran, in 2010 and 2013, respectively. He is currently pursuing the Ph.D. degree with The Signal Processing and Information Networking Laboratory (spinlab), Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, Worcester, Massachusetts, USA. His research interests include optimization theory and wireless communication networks.

Andrew G. Klein received the B.S. degree from Cornell University, Ithaca, NY, the M.S. degree from the University of California, Berkeley, and the Ph.D. degree from Cornell University, all in Electrical Engineering. He joined the Department of Engineering and Design, Western Washington University, Bellingham, WA, in 2014 where he is currently an Associate Professor. Previously, he was an Assistant Professor at Worcester Polytechnic Institute, Worcester, MA, from 2007–2014, and he was a Post-Doctoral Researcher with Supélec/LSS, Paris, France, from 2006 to 2007.

D. Richard Brown III is currently a Professor and the Associate Department Head in the Department of Electrical and Computer Engineering at Worcester Polytechnic Institute, where he has been a Faculty Member since 2000. He received a PhD in Electrical Engineering from Cornell University in 2000 and MS and BS degrees in Electrical Engineering from The University of Connecticut in 1996 and 1992, respectively. From 1992-1997, he was a design engineer at General Electric Electrical Distribution and Control in Plainville, Connecticut. From August 2007 to June 2008, he held an appointment as a Visiting Associate Professor at Princeton University. From 2016-2018, he served as a Program Director at the National Science Foundation in the Computing and Communications Foundations (CCF) division of the Directorate for Computer & Information Science & Engineering (CISE). He is also currently serving as an Associate Editor for IEEE Transactions on Wireless Communications.