Bounds on the Age of Information for Global Channel State Dissemination in Fully-Connected Networks

Shahab Farazi[†], Andrew G. Klein[‡] and D. Richard Brown III[†]
[†]Worcester Polytechnic Institute, 100 Institute Rd., Worcester, MA 01609, Email: {sfarazi,drb}@wpi.edu
[‡]Western Washington University, 516 High St., Bellingham, WA 98225, Email: andy.klein@wwu.edu

***Abstract*—This paper studies an "age of information" problem in fully-connected wireless networks with time-varying reciprocal channels and packetized transmissions. Specifically, a scenario where each node in the network wishes to maintain a table of global channel state information (CSI) is considered. Each node updates its global CSI table in two ways: (i) direct channel measurements through standard channel estimation techniques and (ii) indirect observations of channels through CSI dissemination from other nodes in the network. Information aging, i.e., CSI staleness, occurs due to timeslotting and contention for the common channel resources. This paper derives new lower bounds for the maximum and average CSI staleness of any protocol. These bounds generalize previously developed bounds by allowing for any number of CSI estimates to be disseminated in each packet. A simple one-step greedy protocol is also proposed for any network size and any number of CSI estimates disseminated per packet. Numerical results are provided to demonstrate the achieved staleness of the greedy protocol with respect to the bounds and to also quantify CSI staleness in terms of various network parameters.**

***Index Terms*—Age of information, wireless networks, time-varying channels, global channel state information, channel estimation, data dissemination.**

I. Introduction

There has been a recent interest studying information-update systems where the "age" or "staleness" of the information is the key measure of performance [1]–[12]. Initial work in this area considered vehicular networks where vehicles exchange position and velocity information. To ensure safe operation, each vehicle must maintain current state information of all nearby vehicles while also providing timely updates of its own state to nearby vehicles. As shown in [1], careful use of the common channel resources is necessary to minimize the age of information in these types of systems. While the more recent work in this area has considered this problem in more general contexts, the common question is how to determine optimal update strategies in various settings such that age or staleness of information is minimized.

This paper considers age of information in the specific context of tracking *global* channel state information (CSI) in fully-connected wireless networks with time-varying reciprocal channels and packetized transmissions. Each node in the network maintains a table of global CSI. For the setting considered here, i.e., a fully-connected network of N nodes with reciprocal channels, the total number of channels in the network is equal to $L = (N^2 - N)/2$. Each node in the network maintains its own table of each of these L channels. Global CSI knowledge can be used to exploit available network resources and enable efficient techniques such as cooperative and distributed communication [13]–[20]. Since the channels are assumed to be time-varying, the age or staleness of the global CSI knowledge is critical.

With one node transmitting during each time slot, nodes receive updates to their global CSI tables in two ways. The first way is standard reciprocal channel estimation, i.e., when node i transmits a packet, all nodes $j \neq i$ estimate the reciprocal (i, j) channel. The staleness of these CSI table entries is then set to zero since the estimates are up to date. These direct estimates are not sufficient, however, for each node to obtain global CSI. For example, node 1 can never directly estimate the $(2, 3)$ channel. In order to obtain global CSI, each node must also *disseminate* some entries from its local CSI table when it transmits a packet. The staleness of disseminated CSI is not zero due to the latency between when the CSI was directly estimated and when it is disseminated to the other nodes in the network. The goal is to have the nodes transmit in a certain order and disseminate certain CSI, i.e., follow a particular protocol, so that the maximum and average staleness metrics are minimized across the network.

Prior work [21] developed fundamental lower bounds on the maximum and average staleness of global CSI for any valid protocol in the fully-connected reciprocal channels setting considered in this paper. The analysis in [21], however, considered two extremes for CSI dissemination: $M = 1$ and $M = N - 1$ CSI estimates disseminated per packet. Other values of M were not considered. Explicit protocols were also developed for the $M = 1$ and $M = N-1$ cases to show the achievability of these bounds to within a constant staleness gap.

The contributions of this paper are twofold. First, we generalize the lower bounds developed in [21] by deriving new lower bounds for the maximum and average staleness for any $M \in \{1, 2, 3, \ldots, N-1\}$ CSI estimates disseminated per packet. Second, we propose a simple one-step greedy protocol. In each timeslot, given M and N, the one-step greedy protocol selects the node i and the M CSI estimates from the table maintained by node i to disseminate such that the average

This work was supported by the National Science Foundation awards CCF-1302104 and CCF-1319458.

staleness improvement throughout the network is maximized after the transmission by node i. We provide numerical results that demonstrate the performance of the greedy protocol with respect to the efficient protocols developed for $M = 1$ and $M = N - 1$ in [21] and the new bounds in this paper.

II. System Model

Consider a fully-connected network with N single-antenna nodes communicating over time-varying reciprocal channels. The complex channel gain between two nodes i and j at time n is denoted by $h_{i,j}[n]$ and assuming reciprocity, we have $h_{i,j}[n] = h_{j,i}[n]$. The network's topology is described by a complete graph with $N \geq 3$ vertices and $L = (N^2 - N)/2$ edges, representing all of the N nodes and L reciprocal channels in the network, respectively. Each node in the network maintains its own local table of estimates of these L complex channel gains. During each time slot, one node transmits a packet of length P words which includes $M \in \{1, 2, 3, \ldots, N - 1\}$ *disseminated* CSI estimates.

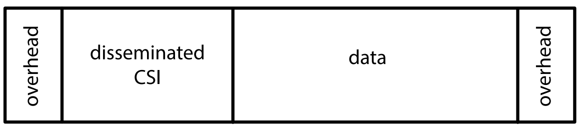

Fig. 1. Example fixed-length packet showing overhead, data, and CSI dissemination. The CSI dissemination consists of M channel estimates and each channel estimate (including its associated timestamp) has a length of one word. The data and overhead consists of D words. The total packet length is $P = D + M$ words.

Figure 1 represents the general structure of a packet exchanged among the nodes in the network. All packets are assumed to be received reliably. Each fixed-length packet contains overhead, data, and M disseminated CSI estimates. Since node k cannot estimate a channel to which it is not directly connected, i.e., the channel between nodes i and j for $i \neq j \neq k$, it uses the disseminated CSI information embedded in the transmitted packets by either nodes i or j, to obtain an estimate of the (i, j) channel. Assuming a length of D words for the data plus overhead, each packet has a length of $P = D + M$ words. Although Fig. 1 shows a particular packet structure, the position of the overhead, data, and disseminated CSI within any packet does not affect our results.

Each node that receives the transmitted packet by node i does two things:

1) It directly estimates the channel $h_{i,j}[n]$, which can be obtained via a known training sequence in the packet, e.g., a known preamble embedded in the overhead, and/or through blind channel estimation techniques.
2) It extracts the disseminated CSI and uses it to update any "staler" CSI in its local table.

Note that every disseminated CSI also includes a timestamp of when it was obtained. This allows each node to determine if the disseminated CSI is fresher than any CSI in its table.

We denote the k^{th} node's estimate of the (i, j) channel measured directly from a packet transmitted at time n as $\hat{h}_{i,j}^{(k)}[n]$. Since each node $i \in \{1, \ldots, N\}$ maintains its own table of L global CSI estimates, the total number of CSI estimates in the network is NL.

As an explicit example of CSI estimation and dissemination, suppose at time $n = 1$ node 6 directly estimates the $(5, 6)$ channel and updates its local CSI table with $\hat{h}_{5,6}^{(6)}[1]$. Then, at time $n = 2$ suppose node 6 disseminates this estimate to other nodes in the network. In the process of this dissemination, each node $j \neq 6$ updates its *direct* CSI estimate, i.e., stores $\hat{h}_{6,j}^{(j)}[2]$ in its local CSI table, and also checks its local CSI table to determine if its estimate of the $(5, 6)$ channel has a timestamp prior to the timestamp of the disseminated $(5, 6)$ estimate from node 6. Each node j with a staler $(5, 6)$ estimate stores the indirectly obtained CSI estimate $\hat{h}_{5,6}^{(j)}[1]$ in its local CSI table.

The following definitions establish the notions of a "protocol" and of "staleness", in units of "words", used throughout this paper.

Definition 1 (Protocol)**.** *A protocol is a sequence of transmitting nodes and the channel indices they disseminate.*

Protocols can be *deterministic* [21], [22] or *random* [23].

Definition 2 (Staleness)**.** *The staleness $s_{i,j}^{(k)}[n]$ of the CSI estimate $\hat{h}_{i,j}^{(k)}[n']$ with timestamp n' at time $n \geq n'$ is $(n-n')P$ words.*

Directly estimated CSI has a staleness of zero in the timeslot in which it is estimated. All indirectly estimated CSI has a minimum staleness of P words. If a given CSI estimate is not updated in a timeslot, either due to direct estimation or dissemination, then its staleness increases by P words in that timeslot.

Using any ordering of the individual staleness terms $s_{i,j}^{(k)}[n]$, we can denote a global staleness vector $\boldsymbol{s}[n] \in \mathbb{Z}^{NL}$. It is not difficult to see that, given a transmitting node and disseminated channel indices, the staleness vector obeys a simple time-varying linear update equation

$$\boldsymbol{s}[n+1] = \boldsymbol{A}[n] \left(\boldsymbol{s}[n] + \mathbf{1}P \right) \tag{1}$$

where $\boldsymbol{A}[n] \in \mathbb{Z}^{NL \times NL}$ is a time-varying update matrix with entries equal to either zero or one. As an explicit example, and omitting the time index for notational convenience, we can define the global staleness vector in an $N = 3$ node network as

$$\boldsymbol{s} = \left[s_{1,2}^{(1)}, s_{1,3}^{(1)}, s_{2,3}^{(1)}, s_{1,2}^{(2)}, s_{1,3}^{(2)}, s_{2,3}^{(2)}, s_{1,2}^{(3)}, s_{1,3}^{(3)}, s_{2,3}^{(3)} \right]^\top .$$

Suppose node 1 disseminates the $(1, 3)$ link at time n. Then

$$\boldsymbol{A}[n] = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \boldsymbol{e}_1^\top \\ \boldsymbol{e}_2^\top \\ \boldsymbol{e}_3^\top \\ \mathbf{0}^\top \\ \boldsymbol{e}_2^\top \\ \boldsymbol{e}_6^\top \\ \boldsymbol{e}_7^\top \\ \mathbf{0}^\top \\ \boldsymbol{e}_9^\top \end{bmatrix}$$

where $\boldsymbol{e}_m$ is the m^{th} standard unit vector. Rows 1-3 of $\boldsymbol{A}[n]$ reflect the fact that node 1 was transmitting and, as such, no updates are made to its local CSI table (the staleness of all CSI at node 1 increases by P). Row 4, being all zeros, reflects the direct estimation of the $(1,2)$ link at node 2. Row 5 reflects the indirect estimation of the $(1,3)$ link at node 2 from the CSI disseminated by node 1, i.e., node 2 now has the same estimate and the same staleness of the $(1,3)$ link as node 1. The remaining rows are similar, but it is worth mentioning row 8. In this case, node 3 receives the disseminated CSI of the $(1,3)$ link from node 1 but also directly estimates the $(1,3)$ link. Since the direct estimate is fresher than any disseminated estimate, node 3 ignores the disseminated CSI. This is reflected in the all-zero row 8. Some basic properties of $\boldsymbol{A}[n]$ will be listed in Section II-B.

A. Staleness Statistics

In this section, we define the maximum and average staleness statistics used in the remainder of the paper.

Definition 3 (Maximum staleness at time n)**.** *The maximum staleness at time n is defined as*

$$S_{max}[n] = \max \boldsymbol{s}[n].$$

Definition 4 (Average staleness at time n)**.** *The average staleness at time n is defined as*

$$S_{avg}[n] = \frac{1}{NL}\mathbf{1}^\top \boldsymbol{s}[n].$$

Definition 5 (Maximum staleness)**.** *The maximum staleness S_{max} is defined as*

$$S_{max} = \max_{n \geq \bar{n}} S_{max}[n]$$

for $\bar{n}$ sufficiently large such that the effect of any initial staleness state can be ignored.

Definition 6 (Average staleness)**.** *The average staleness S_{avg} is defined as*

$$S_{avg} = \mathrm{E}\left[S_{avg}[n]\right]$$

where the expectation is over $n \geq \bar{n}$ for $\bar{n}$ sufficiently large such that the effect of any initial staleness state can be ignored.

B. Basic properties of $\boldsymbol{A}[n]$

From (1) and the description of how each node updates its local CSI table through direct estimation and disseminated CSI estimates, observe that $\boldsymbol{A}[n]$ has the following properties:

- Each row of $\boldsymbol{A}[n]$ is either equal to zero or has a single non-zero entry equal to one.
- Row m of $\boldsymbol{A}[n]$ is equal to zero if the corresponding CSI estimate is directly estimated in timelot n.
- Since $N-1$ CSI estimates are directly estimated in each timeslot, exactly $N-1$ rows of $\boldsymbol{A}[n]$ are zero for all n.
- Row m of $\boldsymbol{A}[n]$ is equal to $\boldsymbol{e}_m^\top$, i.e., the transposed m^{th} standard unit vector, if the corresponding CSI estimate is not updated.
- Row m of $\boldsymbol{A}[n]$ is equal to $\boldsymbol{e}_\ell^\top$, i.e., the transposed ℓ^{th} standard unit vector, if the corresponding CSI estimate is updated to match the disseminated CSI estimate corresponding to row ℓ. Since the M disseminated CSI are received by $N-1$ nodes, it is easy to see that there are at most $(N-1)M$ such rows in $\boldsymbol{A}[n]$. In fact, due to the fact that some disseminated CSI is also directly estimated, as was illustrated in the $N=3$ example, there will be at most $(N-1)M-M$ such rows in $\boldsymbol{A}[n]$.

This last basic property allows us to characterize the dimension of the nullspace of $\boldsymbol{A}[n]$. Observe that the dimension of the nullspace of $\boldsymbol{A}[\boldsymbol{n}]$, i.e., $\text{nullity}(\boldsymbol{A}[n])$, satisfies

$$N-1 \leq \text{nullity}(\boldsymbol{A}[n]) \leq N-1+(N-1)M-M. \quad (2)$$

The $N-1$ term on both sides corresponds to the $N-1$ rows of $\boldsymbol{A}[n]$ that must be zero due to direct estimation at all nodes except the transmitting node. On the right hand side, the $(N-1)M-M=(N-2)M$ term corresponds to the "useful" disseminated CSI, i.e., each node except the transmitting node receives M disseminated CSI estimates and a total of M of these are discarded due to direct estimation. The upper bound in (2) is tight, as seen in the $N=3$ node example above, and is the key result that facilitates the development of the lower bounds on the maximum and average staleness in the following section.

III. Lower bounds on maximum and average staleness for general M

In this section we derive lower bounds on the maximum and average staleness metrics for global CSI dissemination in fully-connected networks. These lower bounds hold for any protocol, whether deterministic or random, and they generalize the lower bounds derived in [21] which considered only the extreme choices of $M \in \{1, N-1\}$; in this work, the bounds apply to any fixed $M \in \{1,2,3,\dots,N-1\}$. Since the network is fully-connected, each node has at most $N-1$ useful directly estimated CSI parameters to disseminate in its packet [21]. As shown by simulations in section V, the achievable maximum and average staleness are not necessarily minimized by either of the extreme choices $M \in \{1, N-1\}$ CSI estimates disseminated in each packet, and the best choice of M depends on the number of nodes N and the amount of data plus overhead D in each packet. Intuitively, as M increases, more CSI is disseminated in each packet and each node updates more entries in its local CSI table. This improves the staleness metrics. On the other hand, the packet length P increases with M, which has a negative impact on staleness. Thus, there are two competing forces in choosing the optimal amount of CSI to disseminate per packet.

From (1) and given an initial state at time n_0, we can write

$$\boldsymbol{s}[n] = \boldsymbol{\Phi}[n,n_0]\boldsymbol{s}[n_0] + P\sum_{t=n_0}^{n-1}\boldsymbol{\Phi}[n,t]\mathbf{1} \quad (3)$$

where

$$\boldsymbol{\Phi}[n,\tau] = \begin{cases} \boldsymbol{A}[n-1]\boldsymbol{A}[n-2]\cdots\boldsymbol{A}[\tau] & n-\tau>0 \\ \boldsymbol{I}_{NL} & n-\tau=0 \\ \text{undefined} & n-\tau<0. \end{cases}$$

Observe that, like $\boldsymbol{A}[n]$, each row of $\boldsymbol{\Phi}[n,t]$ for $n \geq t$ is either all zeros or contains a single non-zero element equal to one.

In order to develop the lower bounds on the maximum and average staleness, we begin with a basic but useful Lemma.

Lemma 1. *For any matrices $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ and $\boldsymbol{B} \in \mathbb{R}^{n\times n}$,*

$$\mathrm{nullity}(\boldsymbol{AB}) \leq \mathrm{nullity}(\boldsymbol{A}) + \mathrm{nullity}(\boldsymbol{B})$$

Proof: From [24], we have $\mathrm{rank}(\boldsymbol{AB}) \geq \mathrm{rank}(\boldsymbol{A}) + \mathrm{rank}(\boldsymbol{B}) - n$. Since the rank and nullity of any $n \times n$ matrix must sum to n, we can write

$$n - \mathrm{nullity}(\boldsymbol{AB}) \geq n - \mathrm{nullity}(\boldsymbol{A}) + n - \mathrm{nullity}(\boldsymbol{B}) - n$$

which simplifies directly to the desired result. ∎

The utility of this result is that it can be used to provide a convenient upper bound on the dimension of the nullspace of the transition matrix $\boldsymbol{\Phi}[n,\tau]$ in terms of the dimension of the nullspaces of the constituent matrices $\boldsymbol{A}[n-1]$, $\boldsymbol{A}[n-2]$, $\dots$, $\boldsymbol{A}[\tau]$ for $n > \tau$.

We now derive a useful property of the transition matrices in the following Lemma.

Lemma 2. $\boldsymbol{1}^\top \boldsymbol{\Phi}[n,\tau]\boldsymbol{1} > 0$ *for all $n-\tau$ satisfying*

$$0 \leq n - \tau < \Delta^* = \lceil \Delta \rceil$$

where

$$\Delta := \frac{NL}{N-1+(N-2)M}.$$

Proof: Since $\boldsymbol{\Phi}[n,\tau] \in \mathbb{R}^{NL\times NL}$ is composed of elements equal to zero or one for all $n > \tau$, then $\boldsymbol{1}^\top \boldsymbol{\Phi}[n,\tau]\boldsymbol{1} > 0$ if and only if

$$\mathrm{nullity}(\boldsymbol{\Phi}[n,\tau]) < NL.$$

From Lemma 1 and (2), we have

$$\mathrm{nullity}(\boldsymbol{\Phi}[n,\tau]) \leq (n-\tau)(N-1+(N-2)M).$$

Hence, to satisfy $\mathrm{nullity}(\boldsymbol{\Phi}[n,\tau]) < NL$ for integer n and τ, it is sufficient for $0 \leq n-\tau < \left\lceil \frac{NL}{N-1+(N-2)M} \right\rceil$, which shows the desired result. ∎

We present one additional Lemma that will also be used in the development of the bounds.

Lemma 3. $\boldsymbol{0} \preceq \boldsymbol{\Phi}[n,t-1]\boldsymbol{1} \preceq \boldsymbol{\Phi}[n,t]\boldsymbol{1} \preceq \boldsymbol{1}$ *for all $n \geq t$ where $\preceq$ corresponds to element-wise inequality.*

Proof: The inequality $\boldsymbol{\Phi}[n,t]\boldsymbol{1} \preceq \boldsymbol{1}$ follows directly from the fact that each row of $\boldsymbol{\Phi}[n,\tau]$ for $n \geq \tau$ is either all zeros or contains a single non-zero element equal to one. Given any $\boldsymbol{0} \preceq \boldsymbol{u} \preceq \boldsymbol{v}$, this fact also implies $\boldsymbol{0} \preceq \boldsymbol{\Phi}[n,\tau]\boldsymbol{u} \preceq \boldsymbol{\Phi}[n,\tau]\boldsymbol{v}$ for all $n \geq \tau$. Then, setting $\boldsymbol{u} = \boldsymbol{A}[t-1]\boldsymbol{1}$ and $\boldsymbol{v} = \boldsymbol{1}$ and noting that $\boldsymbol{0} \preceq \boldsymbol{u} \preceq \boldsymbol{v}$, we can write

$$\boldsymbol{0} \preceq \boldsymbol{\Phi}[n,t-1]\boldsymbol{1} = \boldsymbol{\Phi}[n,t]\boldsymbol{u} \preceq \boldsymbol{\Phi}[n,t]\boldsymbol{v} = \boldsymbol{\Phi}[n,t]\boldsymbol{1} \preceq \boldsymbol{1}$$

which is the desired result. ∎

An implication of this result is that if, for some fixed $\tau < n$, a particular element of $\boldsymbol{\Phi}[n,\tau]\boldsymbol{1}$ is equal to one, then that same element must also be equal to one for all $\boldsymbol{\Phi}[n,\tau+\ell]\boldsymbol{1}$ for all $\ell = 1,\dots,n-\tau$.

The following two Theorems present the lower bounds on maximum and average staleness.

Theorem 1 (Lower bound on maximum staleness)**.** *The maximum staleness of any protocol is lower bounded by*

$$S_{max} \geq S^*_{max} = (\Delta^* - 1)\,P.$$

Proof: From Definition 3 and (3), we can write

$$\begin{aligned}
S_{max}[n] &= \max \boldsymbol{s}[n] \\
&= \max\left(\boldsymbol{\Phi}[n,n_0]\boldsymbol{s}[n_0] + P\sum_{t=n_0}^{n-1}\boldsymbol{\Phi}[n,t]\boldsymbol{1}\right) \\
&\geq \max\left(P\sum_{t=n_0}^{n-1}\boldsymbol{\Phi}[n,t]\boldsymbol{1}\right)
\end{aligned}$$

where the inequality results from the fact that $\boldsymbol{\Phi}[n,n_0]\boldsymbol{s}[n_0]$ only contains non-negative elements. Note that the sum contains $n-n_0$ terms. Lemma 2 implies that at least $\boldsymbol{\Phi}[n,n-1]\boldsymbol{1},\dots,\boldsymbol{\Phi}[n,n-\Delta^*+1]\boldsymbol{1}$ must be non-zero. Hence

$$S_{max}[n] \geq \max\left(P\sum_{t=n-\Delta^*+1}^{n-1}\boldsymbol{\Phi}[n,t]\boldsymbol{1}\right) = (\Delta^*-1)\,P.$$

The final equality follows from the fact that there are Δ^*-1 terms in the summation and, according to Lemma 3, at least one consistent element of each term is equal to one. The desired result then follows directly from Definition 5. ∎

Theorem 2 (Lower bound on average staleness)**.** *The average staleness of any protocol is lower bounded by*

$$S_{avg} \geq S^*_{avg} = \lambda(\Delta^*-1)P$$

where $\lambda := 1 - \frac{\Delta^*}{2\Delta}$.

Proof: From Definition 4 and following a similar approach as in the proof of Theorem 1, we can write

$$\begin{aligned}
S_{avg}[n] &= \frac{1}{NL}\boldsymbol{1}^\top\boldsymbol{s}[n] \\
&\geq \frac{P}{NL}\sum_{t=n-\Delta^*+1}^{n-1}\boldsymbol{1}^\top\boldsymbol{\Phi}[n,t]\boldsymbol{1} \\
&\geq \frac{P}{NL}\sum_{t=n-\Delta^*+1}^{n-1}\mathrm{rank}\left(\boldsymbol{\Phi}[n,t]\right) \\
&\geq \frac{P}{NL}\sum_{t=n-\Delta^*+1}^{n-1} NL-(n-t)(N-1+(N-2)M) \\
&= P\sum_{t=n-\Delta^*+1}^{n-1} 1-\frac{n-t}{\Delta}
\end{aligned}$$

where the final inequality follows from Lemma 1 and the properties of $\mathbf{\Phi}[n, t]$. This result can be simplified by reindexing the sum to write

$$\begin{aligned}
S_{avg}[n] &\geq P \sum_{m=1}^{\Delta^*-1} 1 - \frac{m}{\Delta} \\
&= P\left(\Delta^* - 1 - \frac{1}{\Delta}\sum_{m=1}^{\Delta^*-1} m\right) \\
&= P\left(\Delta^* - 1 - \frac{1}{\Delta}\cdot\frac{(\Delta^*-1)\Delta^*}{2}\right) \\
&= P(\Delta^*-1)\left(1 - \frac{\Delta^*}{2\Delta}\right) \\
&= \lambda(\Delta^*-1)P.
\end{aligned}$$

Since this result does not depend on n, we have

$$S_{avg} \geq S^*_{avg} = \lambda(\Delta^* - 1)P$$

which is the desired result. ■

IV. Greedy CSI Dissemination

A deterministic protocol is considered to be *efficient* if it achieves within a constant gap of the lower bounds on the maximum and average staleness, independent of the number of nodes in the network [21]. For the $M = 1$ and $M = N-1$ cases, it was shown in [21] that an efficient deterministic protocol can be constructed by considering an Eulerian tour and a round-robin transmission order, respectively, where the transmitting node during each time slot disseminates its M freshest direct estimates. However, for $M \in \{2, 3, 4, \ldots, N-2\}$, the resulting combinatorics are challenging, and it is not obvious how to construct an efficient protocol in general. That is, it is not clear how to make the best choice of transmitting node in each time slot, and the choice of which CSI should be disseminated among the $\binom{N-1}{M} = \frac{(N-1)!}{(N-1-M)!M!}$ different sets of direct estimates from its table to disseminate. Hence, here we present a "greedy" protocol that minimizes the instantaneous average staleness throughout the network, and generates protocols for any choice of N and M.

One-Step Greedy Protocol:

■ I: initialization.
∗ each of the N nodes shares its staleness table with the rest of the network.
∗ set time slot, $n \leftarrow 0$.
■ II-a: compute the average staleness improvement over all combinations.
for $i_{\text{node}} = 1 : N$ **do**
  for $j_{\text{set}} = 1 : \binom{N-1}{M}$ **do**
    compute $S_{avg,improve}(i_{\text{node}}, j_{\text{set}})$, which represents the average staleness improvement throughout the network, given that node i_{node} disseminates its $j^{\text{th}}_{\text{set}}$ set of M direct estimates.
  end
end
■ II-b: choose the node and its set of M estimates that maximize the average staleness improvement, and resolve ties.
if there is a tie between two different nodes **then**
  ∗ select the node that has least recently transmitted.
else
  if there is a tie between two different sets of M direct estimates at the same node **then**
    ∗ select the set that updates greater number of staler estimates throughout the network.
  end
end
■ III: the selected node disseminates it selected set of M direct estimates.
■ IV: all nodes update their local tables with any disseminated CSI that is fresher than the CSI currently in their table.
■ V: $n \leftarrow n + 1$, and repeat from II-a.

The one-step greedy protocol can be run in parallel at all nodes in a distributed fashion. We note that the protocol could omit the initialization step to minimize startup overhead, if desired, by assuming that all nodes have no CSI knowledge. Although sharing the tables at the beginning takes some time, it permits the protocol to make use of the existing CSI knowledge throughout the network, and therefore may lead to an overall staleness improvement at startup. In the long run, however, this initial CSI is insignificant since the CSI tables are continually updated. Finally, we note that the resulting greedy protocol for a given choice of N, M, and D could be precomputed offline and saved in a lookup table at all nodes.

Since in each time slot, the greedy protocol determines the transmitting node and its set of M direct estimates based on maximizing the average staleness during only the current time slot, it is called one-step greedy protocol. Using a protocol that minimizes the instantaneous staleness in each time slot is somewhat myopic, and may not lead to the best protocol in terms of minimizing steady-state staleness. Thus, this one-step greedy protocol could be extended to minimize the average staleness improvement over a window of two or more consecutive time slots, leading to a possible improvement in steady-state staleness. However, the amount of combinations to search over (i.e., node order and CSI to disseminate) increases exponentially with the window size.

Note that the protocols generated by the one-step greedy protocol eventually result in a periodic protocol since, for any fixed N, there are a finite number of parameters throughout the network, i.e., NL CSI parameters, and they can be modeled as a finite state-space system. The resulting periodic protocol, however, is dependent on the initial conditions (i.e., the staleness values throughout the network at startup). Interestingly, for the $M \geq (N-1)/2$ region, always the greedy protocol generates the same N-periodic round-robin protocol in [21] regardless of initialization, and in the resulting protocol each transmitting node always disseminates its M freshest directly estimated CSI parameters. Intuitively, for the $M \geq (N-1)/2$ region, all CSI parameters throughout the network can be updated at least once when during every N consecutive time slots, each of the N nodes transmits exactly once and disseminates its M freshest estimates.

V. Numerical Results

This section provides numerical examples to verify the analysis in the previous section and to quantify the maximum and average staleness as a function of the network parameters N, D and M. Figures 2, 3 and 4 compare the achievable maximum and average staleness of the greedy protocol with the lower bounds in Theorems 1 and 2 and the efficient deterministic protocols for $M \in \{1, N-1\}$ in [21] for $N = 8$, versus the number of CSI estimates per packet M for $D \in \{0, 2, 10\}$. The $D = 0$ case can be considered a protocol with no data or overhead where each packet is dedicated solely to CSI dissemination. Since the greedy protocol is sensitive to the initial staleness values throughout the network, for each M the protocol is run for 1000 random initializations and the minimum and maximum staleness values are chosen as the best and worst achievable staleness of the greedy protocol, respectively. The area between the best and worst achievable staleness of the greedy protocol is shaded, which represents achievable maximum and average staleness of the greedy protocol for different initializations.

Figures 2 and 4 confirm that for small and large values of D, to minimize the achievable maximum staleness it is optimal to disseminate $M = 1$ and $M = N - 1$ CSI estimates in each packet, respectively, but as Fig. 3 shows for an intermediate value of D, i.e., $D = 2$, the achievable maximum staleness is minimized when $M = 4$ CSI estimate are disseminated in each packet. Also, depending on the initial staleness tables, the greedy protocol can achieve the maximum and average staleness of the efficient deterministic protocols in [21] for $M = 1$, but for $M = N - 1$ the maximum and average staleness of the efficient deterministic protocols in [21] are always achievable by the greedy protocol.

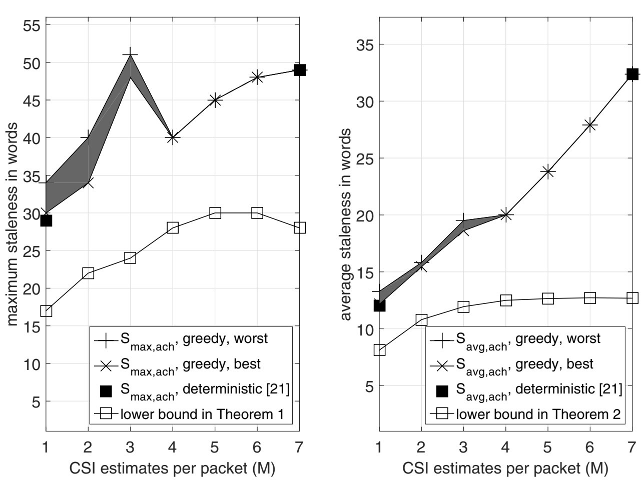

Fig. 2. Staleness versus M for $N = 8$ and $D = 0$.

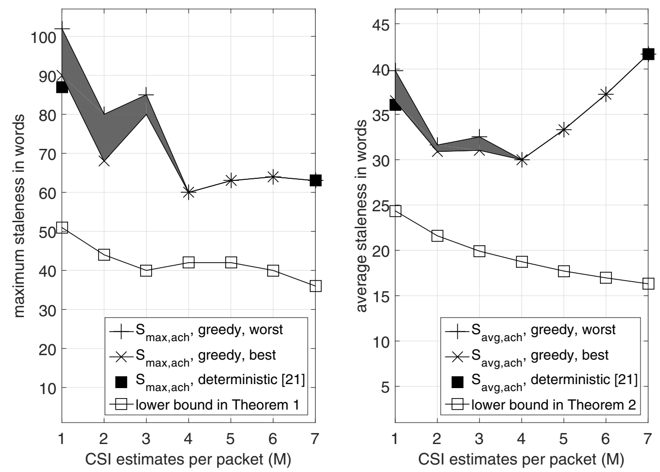

Fig. 3. Staleness versus M for $N = 8$ and $D = 2$.

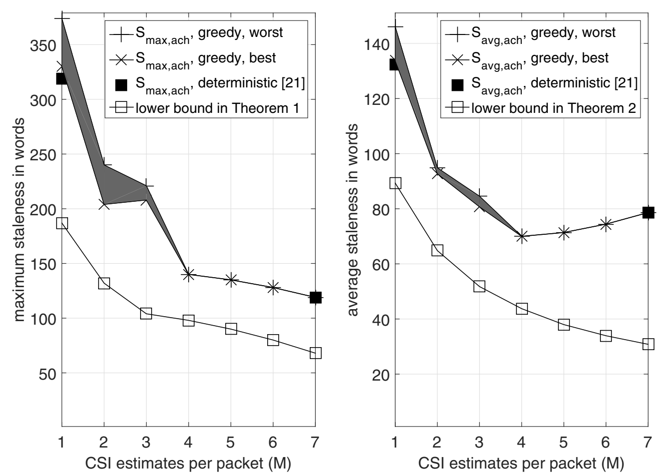

Fig. 4. Staleness versus M for $N = 8$ and $D = 10$.

VI. Conclusion

This paper derived lower bounds on the maximum and average staleness of any protocol in fully-connected networks with packetized transmissions and time-varying reciprocal channels, when any number of CSI estimates get disseminated per packet. A one-step greedy protocol that generates protocols for any choice of number of nodes and number of CSI estimates per packet was proposed based on maximizing the instantaneous average staleness improvement throughout the network. Simulation results quantified the achievable maximum and average staleness of the one-step greedy protocol in terms of network parameters compared to the derived bounds.

References

[1] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing age of information in vehicular networks," in *Sensor, Mesh and Ad Hoc Communications and Networks (SECON), 2011 8th Annual IEEE Communications Society Conference on*. IEEE, 2011, pp. 350–358.

[2] S. Kaul, R. Yates, and M. Gruteser, "On piggybacking in vehicular networks," in *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE*. IEEE, 2011, pp. 1–5.

[3] ——, "Real-time status: How often should one update?" in *INFOCOM, 2012 Proceedings IEEE*. IEEE, 2012, pp. 2731–2735.

[4] S. K. Kaul, R. D. Yates, and M. Gruteser, "Status updates through queues," in *Information Sciences and Systems (CISS), 2012 46th Annual Conference on*. IEEE, 2012, pp. 1–6.

[5] R. D. Yates and S. Kaul, "Real-time status updating: Multiple sources," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 2666–2670.

[6] C. Kam, S. Kompella, and A. Ephremides, "Age of information under random updates," in *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*. IEEE, 2013, pp. 66–70.

[7] ——, "Effect of message transmission diversity on status age," in *Information Theory (ISIT), 2014 IEEE International Symposium on*. IEEE, 2014, pp. 2411–2415.

[8] M. Costa, M. Codreanu, and A. Ephremides, "Age of information with packet management," in *Information Theory (ISIT), 2014 IEEE International Symposium on*. IEEE, 2014, pp. 1583–1587.

[9] M. Costa, S. Valentin, and A. Ephremides, "On the age of channel information for a finite-state Markov model," in *IEEE Intl. Conf. on Communications (ICC)*, 2015, pp. 4101–4106.

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

[11] Y. Sun, E. Uysal-Biyikoglu, R. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," in *Computer Communications, IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on*. IEEE, 2016, pp. 1–9.

[12] C. Kam, S. Kompella, G. D. Nguyen, J. E. Wieselthier, and A. Ephremides, "Age of information with a packet deadline," in *Information Theory (ISIT), 2016 IEEE International Symposium on*. IEEE, 2016, pp. 2564–2568.

[13] R. Madan, N. Mehta, A. Molisch, and J. Zhang, "Energy-efficient cooperative relaying over fading channels with simple relay selection," *IEEE Trans. Wireless Commun.*, vol. 7, no. 8, pp. 3013–3025, Aug. 2008.

[14] R. Mudumbai, D.R. Brown III, U. Madhow, and H.V. Poor, "Distributed transmit beamforming: Challenges and recent progress," *IEEE Commun. Mag.*, vol. 47, no. 2, pp. 102–110, Feb. 2009.

[15] D. R. Brown III, P. Bidigare, and U. Madhow, "Receiver-coordinated distributed transmit beamforming with kinematic tracking," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2012, pp. 5209–5212.

[16] D. R. Brown III, P. Bidigare, S. Dasgupta, and U. Madhow, "Receiver-coordinated zero-forcing distributed transmit nullforming," in *2012 IEEE Statistical Signal Processing Workshop (SSP)*, Aug. 2012, pp. 269–272.

[17] D.R. Brown III, U. Madhow, S. Dasgupta, and P. Bidigare, "Receiver-coordinated distributed transmit nullforming with channel state uncertainty," in *Conf. Inf. Sciences and Systems (CISS2012)*, Mar. 2012.

[18] V. Cadambe and S. Jafar, "Interference alignment and degrees of freedom of the k-user interference channel," *IEEE Trans. Inf. Theory*, vol. 54, no. 8, pp. 3425–3441, Aug. 2008.

[19] D.R. Brown III, U. Madhow, M. Ni, M. Rebholz, and P. Bidigare, "Distributed reception with hard decision exchanges," *IEEE Trans. Wireless Commun.*, vol. 13, no. 6, pp. 3406–3418, Jun. 2014.

[20] J. Choi, D. Love, D.R. Brown III, and M. Boutin, "Quantized distributed reception for mimo wireless systems using spatial multiplexing," *IEEE Trans. Wireless Commun.*, vol. 14, no. 13, pp. 3537–3548, Jul. 2015.

[21] 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.*, submitted for publication.

[22] 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.

[23] 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 *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2016.

[24] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge university press, 2012.