Optimizing Age of Information with Correlated Sources

Vishrant Tripathi
vishrant@mit.edu
Massachusetts Institute of Technology
USA

Eytan Modiano
modiano@mit.edu
Massachusetts Institute of Technology
USA

ABSTRACT

We develop a simple model for the timely monitoring of correlated sources over a wireless network. Using this model, we study how to optimize weighted-sum average Age of Information (AoI) in the presence of correlation. First, we discuss how to find optimal stationary randomized policies and show that they are at-most a factor of two away from optimal policies in general. Then, we develop a Lyapunov drift-based max-weight policy that performs better than randomized policies in practice and show that it is also at-most a factor of two away from optimal. Next, we derive scaling results that show how AoI improves in large networks in the presence of correlation. We also show that for stationary randomized policies, the expression for average AoI is robust to the way in which the correlation structure is modeled. Finally, for the setting where correlation parameters are unknown and time-varying, we develop a heuristic policy that adapts its scheduling decisions by learning the correlation parameters in an online manner. We also provide numerical simulations to support our theoretical results.

CCS CONCEPTS

• **Networks** → **Network performance modeling**; **Network performance analysis**; **Mobile ad hoc networks**.

KEYWORDS

Age of Information, wireless networks, correlated sources, scheduling

ACM Reference Format:
Vishrant Tripathi and Eytan Modiano. 2022. Optimizing Age of Information with Correlated Sources . In *The Twenty-third International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc '22), October 17–20, 2022, Seoul, Republic of Korea.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3492866.3549719

1 INTRODUCTION

Monitoring, estimation, and control of systems are fundamental and well studied problems. Many emerging applications involve performing these tasks over communication networks. Examples include: sensing for IoT applications, control of robot swarms, real-time surveillance, search-and-rescue, and vehicle-to-vehicle (V2V) communication. In these settings, achieving good performance requires timely delivery of status updates from sources to destinations.

Age of Information (AoI) is a metric that captures timeliness of received information at a destination [14, 15, 26]. Unlike packet delay, AoI measures the lag in obtaining information at a destination node, and is therefore suited for applications involving time sensitive updates. Age of information, at a destination, is defined as the time that has elapsed since the last received information update was generated at the source. AoI, upon reception of a new update, drops to the time elapsed since generation of the update, and grows linearly otherwise. Over the past decade, there has been a rapidly growing body of work on analyzing AoI for queuing systems [1, 7, 8, 14, 15, 26], using AoI as a metric for scheduling policies in networks [3, 11, 12, 27, 29, 31] and for monitoring or controlling systems over networks [2, 16, 20, 23, 25]. For detailed surveys of AoI literature see [17] and [24].

Typically, AoI is used as a metric for measuring freshness of information being delivered about a source to a monitoring station. It represents a measure of distortion between the state of the system that is expected at the monitor based on past updates and the actual current state of the system. Thus, a larger age corresponds to the monitor having a higher uncertainty about the current state of the system being observed. This, in turn, means that ensuring a low average AoI can lead to higher monitoring accuracy or better control performance. While AoI is a proxy for measuring the cost of having out-of-date information, it may not properly reflect the impact of stale information on system performance. For example, a source might have a high AoI but the monitor might have a good estimate of its state because another source monitoring phenomena nearby sent updates very recently.

A typical problem formulation used to model networked monitoring and control settings is one involving multiple sources sending status updates to a central monitoring station. Many prior works have looked at optimizing information freshness metrics for such models under different assumptions on the interference constraints [12, 27], arrival processes [10, 19], costs of AoI [11, 29], and update sizes [19, 28]. However, all of these works assume that the information across different sources is *decoupled* or *uncorrelated*, i.e. update deliveries from one source only influence the AoI evolution for that source.

This is not strictly true in practice. Many monitoring and control applications involve observing information from correlated or coupled sources. Examples include: vehicular networks where vehicles communicate with their neighbors, multi-agent robotics tasks such as mapping where robots sense overlapping information, and wireless sensor networks where sensors collect spatially correlated updates or exchange information locally.

In such applications, updates from one source often contain information about the current state of other sources. The focus

MobiHoc '22, October 17–20, 2022, Seoul, Republic of Korea

ACM ISBN 978-1-4503-9165-8/22/10.
https://doi.org/10.1145/3492866.3549719

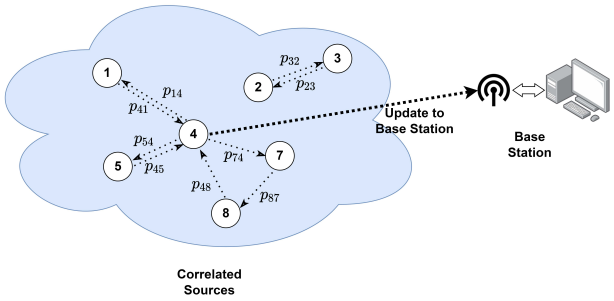

Figure 1: Sources share information locally and send updates to a base station.

of this work is to understand to role of correlation in designing scheduling policies for information freshness in wireless networks. In Sec. 2, we formulate a simple model to analyze weighted-sum average AoI in the presence of correlated sources under wireless interference constraints. In Sec. 3, we use this model to design scheduling policies that can utilize the correlation structure between sources. We formulate a convex problem that solves for the optimal stationary randomized policy and show that it is factor-2 optimal in general. We then develop a Lyapunov drift-based max-weight policy that works well in practice and show that it is also constant factor optimal. In Sec. 4, we provide scaling results that allow us to understand how the degree of correlation affects information freshness. In Sec. 5, we discuss some alternate ways to model correlation and show that the average AoI for these models remains the same as our proposed model under randomized policies. This highlights the robustness of our results to the way in which correlation is modeled. In Sec. 6, we consider the setting where correlation parameters are unknown and possibly time-varying. Here, we propose a heuristic algorithm called *EMA-max-weight* based on exponential moving averages. This algorithm attempts to both keep track of the correlation parameters and adjust the scheduling decisions in an online manner so as to keep information fresh at the base station. Finally, in Sec. 7, we show numerically that our proposed policies outperform scheduling schemes that ignore the correlation structure inherent in the problem and verify our theoretical results.

There has been some prior work in trying to understand how correlation influences information freshness. In [22], the authors consider updates from a single source that are temporally correlated. In [5, 6], the authors consider a network of cameras with overlapping fields of view and formulate a joint optimization problem that looks at processing and scheduling. In [32], the authors consider a setting with multiple sensors partially monitoring a single source, where updates from at least M sensors are required to reconstruct the state of the source. In [9], the authors consider spatially correlated updates from a random field and study the optimal density at which to place sensors. In [13], the authors consider a two hop setting where sources can send updates to multiple sensors.

To the best of our knowledge, our work is the first one to consider coupled AoI evolution. In our model, the AoI of other sources can drop whenever a source that is correlated with them transmits, since a correlated update can reduce uncertainty about the state of other sources. This couples the AoI evolution of sources and leads to fundamentally new scheduling design and scaling results.

2 SYSTEM MODEL

Consider N sources monitoring phenomena of interest and sending updates to a base station. We assume discrete slotted time and assume that only one source can transmit to the base station successfully in any given time-slot.

We consider a simple model for the correlation structure between sources. At the beginning of every time-slot, each source i collects information about its own state. In addition, with probability p_{ij}, the update collected by source i also contains information about the current state of source j, for example due to overlapping fields of view between source i and j or due to spatial correlation between the processes being monitored. We assume that this information sharing or overlap happens *independently* across each *ordered pair* of sources and over time. Clearly, a value of $p_{ij} = 0$ suggests that there is never any information at i about j, while a value of $p_{ij} = 1$ suggests that i has complete information about j at all times. The overall correlation structure between the sources is described by a matrix $\mathbf{P}$ which contains the pairwise correlation probabilities. Figure 1 depicts an example of such a system, where sources have access to some local information about their neighbors and send updates to a base station. Note that $p_{ii} = 1$ for all sources, since each source is assumed to have information about itself. However, our model can also capture situations where a source fails to obtain information about itself occasionally, by setting $p_{ii} < 1$.

When a particular source transmits to the base station, it sends information about its own state to the base station. However, this update will also contain shared information about some of its neighboring sources. Thus, updates sent to the base station are *correlated*: they can contain information about multiple sources at a time. For example, when source 4 transmits to the base station in the setup depicted in Figure 1, its transmission will contain information not just about itself but also about source 1 with probability p_{41}, source 5 with probability p_{45} and source 8 with probability p_{48}.

While we focus on the broadcast interference setting with reliable channels to develop our results and insights, our work can be easily extended to general interference constraints and unreliable channels.

2.1 Age of Information

Let $u_i(t)$ be an indicator variable that denotes whether source i transmits to the base station in time-slot t. Further let $X_{ij}(t)$ be an indicator variable that denotes whether the current update at source i contains common information about source j in time-slot t. From our discussion above, we know that $X_{ij}(t) \sim Bern(p_{ij})$ independent across pairs (i, j) and also over time.

Given this correlation structure, the Age of Information for source i at the base station evolves as follows:

$$A_i(t+1) = \begin{cases} 1, & \text{if } \sum_{j=1}^{N} u_j(t) X_{ji}(t) = 1 \\ A_i(t) + 1, & \text{otherwise.} \end{cases} \tag{1}$$

The equation (1) implies that the AoI of source i at the base station drops to 1 whenever i itself transmits a new update or j

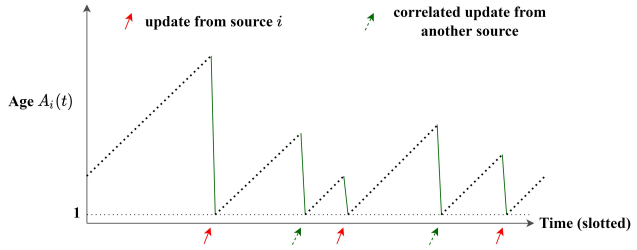

Figure 2: Correlated AoI evolution.

transmits a new update containing information about i. Figure 2 depicts this AoI evolution. Note that (1) also assumes that sources only transmit their freshest update, i.e. an older update containing more information about neighbors might not be sent since it was replaced by a newer update with lesser information about neighbors in the next time-slot.

The metric of interest in this work will be average AoI, which is simply the long-term time-average of the AoI process. Specifically,

$$\bar{A}_i \triangleq \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} A_i(t). \tag{2}$$

2.2 Goal

Given the probabilistic correlation structure and the AoI evolution described above, we want to design a wireless scheduling policy that minimizes the weighted sum of average AoI across all sources:

$$\underset{\pi}{\text{argmin}}\left(\limsup_{T\to\infty} \left[\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{N} w_i A_i(t) \right] \right). \tag{3}$$

Here, the weights $\{w_1, w_2, ..., w_N\}$ are positive real numbers that denote the relative importance of each source to the overall monitoring or control application.

3 SCHEDULING POLICIES

To solve the optimization problem (3), we first study stationary randomized policies in Sec. 3.1. These policies are amenable to analysis and provide key structural insights. Then, in Sec. 3.2, we develop a Lyapunov drift-based max-weight policy and prove performance bounds for it using the structure of the optimal stationary randomized policy.

3.1 Stationary Randomized Policies

A stationary randomized policy is described by a probability distribution $\boldsymbol{\pi}$ over the set of sources, where π_i denotes the probability of choosing source i. In every time-slot, the policy chooses which source gets to transmit by sampling from the distribution π and scheduling decisions are sampled independently across time-slots.

The following theorem relates the average AoI to the scheduling distribution $\boldsymbol{\pi}$.

Theorem 1. *Consider any stationary randomized policy with scheduling probabilities $\boldsymbol{\pi}$. If $\sum_{j=1}^{N} \pi_j p_{ji} > 0$ for a source i, then the average AoI for this source is given by:*

$$\bar{A}_i \triangleq \lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} A_i(t) = \frac{1}{\sum_{j=1}^{N} \pi_j p_{ji}}. \tag{4}$$

If $\sum_{j=1}^{N} \pi_j p_{ji} = 0$ for some source i, then the base station never receives any information regarding this source and its average AoI $\bar{A}_i$ is unbounded.

Proof. See Appendix A in the technical report [30]. □

Using the above theorem, we can formulate an optimization problem to find optimal stationary randomized policies.

Lemma 1. *Consider the space of stationary randomized policies Π^{sr}. Finding a policy $\boldsymbol{\pi}^* \in \Pi^{sr}$ that minimizes the weighted sum of average AoI is equivalent to solving the following optimization problem:*

$$\begin{aligned} \underset{\boldsymbol{\pi}}{\text{argmin}} & \sum_{i=1}^{N} \frac{w_i}{\sum_{j=1}^{N} \pi_j p_{ji}} \\ s.t. & \sum_{i=1}^{N} \pi_i \leq 1, \\ & \pi_i \geq 0, \forall i \in [N]. \end{aligned} \tag{5}$$

Proof. Using the definition of average AoI (2) and the expression derived in Theorem 1, we can simplify (3) to obtain (5). The constraints simply represent a valid stationary randomized policy given the assumption that only one source can transmit in any given time-slot. □

Next, we discuss how to solve the optimization problem (5).

Theorem 2. *The optimization problem (5) is convex in the probability distribution $\boldsymbol{\pi}$. Further, if the optimal solution $\boldsymbol{\pi}^*$ to (5) is such that $\pi_i^* > 0, \forall i \in S$ where $S \subseteq [N]$ is a subset of sources, then $\boldsymbol{\pi}^*$ can be found by solving the following system of nonlinear equations:*

$$\sum_{j=1}^{N} p_{ij} w_j \bar{A}_j^2 = \lambda, \forall i \in S, \text{ and} \tag{6}$$

$$\sum_{i\in S} \pi_i^* = 1. \tag{7}$$

Here, $\lambda > 0$ is a constant and $\bar{A}_i$ denotes the average AoI for source i under the policy $\boldsymbol{\pi}^*$ computed using (4).

Proof. See Appendix B in the technical report [30]. □

Theorem 2 establishes two key results. First, since the optimization problem (5) is convex, it can be solved efficiently by using a standard solver such as cvx [4]. Second, if the optimal policy involves scheduling some subset of sources a positive fraction of the time, then the quantity $\sum_{j=1}^{N} p_{ij} w_j \bar{A}_j^2$ is constant across all sources in this subset. This can be contrasted with the equivalent result in the uncorrelated case, where the quantity $w_j \bar{A}_j^2$ is constant across all sources [12]. We also recover the well known result from the uncorrelated case, i.e. $\pi_i^* \propto \sqrt{w_i}$ if we set $\mathbf{P} = \mathbf{I}$.

Until now, we have only discussed optimization within the space of stationary randomized policies. The following theorem shows

that this class is not too restrictive, i.e. the best stationary randomized policies are at-most a factor of two away from the best possible scheduling policy *in general*.

Theorem 3. *Consider an optimal stationary randomized policy $\boldsymbol{\pi}^*$ that is a solution to* (5) *and an optimal policy $\boldsymbol{\pi}^{opt}$ that solves the general problem* (3). *Let $\bar{\mathbf{A}}^*$ denote the average AoIs under $\boldsymbol{\pi}^*$ and $\bar{\mathbf{A}}^{opt}$ denote the average AoIs under $\boldsymbol{\pi}^{opt}$. Then,*

$$\frac{\sum_{i=1}^N w_i \bar{A}_i^*}{\sum_{i=1}^N w_i \bar{A}_i^{opt}} \leq 2. \tag{8}$$

Proof. See Appendix C in the technical report [30]. □

3.2 Max-Weight Policy

Motivated by the Lyapunov drift-based policies proposed in [12, 24, 27], we next look at an alternative way to design scheduling policies that take correlation into account. Consider the quadratic Lyapunov function given by:

$$L(t) \triangleq \sum_{i=1}^N w_i A_i^2(t). \tag{9}$$

Then, the *quadratic max-weight policy* chooses a scheduling decision that minimizes the one-slot Lyapuonv drift in every time-slot.

$$\pi^{qmw}(t) = \underset{i\in[N]}{\operatorname{argmin}} \mathbb{E}\Big[L(t+1) - L(t)\Big|\mathbf{A}(t)\Big]. \tag{10}$$

This simplifies to:

$$\pi^{qmw}(t) = \underset{i\in[N]}{\operatorname{argmin}} \sum_{j=1}^N w_j p_{ij} A_j(t)\Big(A_j(t)+2\Big). \tag{11}$$

The following theorem provides an upper-bound on the performance of the quadratic max-weight policy.

Theorem 4. *Consider the quadratic max-weight policy $\boldsymbol{\pi}^{qmw}$ and an optimal policy $\boldsymbol{\pi}^{opt}$ that solves the general problem* (3). *Let $\bar{\mathbf{A}}^{qmw}$ denote the average AoIs under $\boldsymbol{\pi}^{qmw}$ and $\bar{\mathbf{A}}^{opt}$ denote the average AoIs under $\boldsymbol{\pi}^{opt}$. Then,*

$$\frac{\sum_{i=1}^N w_i \bar{A}_i^{qmw}}{\sum_{i=1}^N w_i \bar{A}_i^{opt}} \leq 4. \tag{12}$$

Proof. See Appendix D in the technical report [30]. □

Observe that computing the decisions in the equation above does not require us to explicitly solve for the optimal stationary randomized policy $\boldsymbol{\pi}^*$. Next, we will develop a max-weight policy that utilizes the optimal stationary randomized policy $\boldsymbol{\pi}^*$ to get even better performance guarantees.

Consider an optimal stationary randomized policy $\boldsymbol{\pi}^*$ that solves the optimization problem (5). Using $\boldsymbol{\pi}^*$, we define the quantities:

$$\alpha_i \triangleq \frac{w_i}{\sum_{j=1}^N \pi_j^* p_{ji}}, \forall i \in [N]. \tag{13}$$

We use the quantities α_i to construct the following linear Lyapunov function:

$$L(t) \triangleq \sum_{i=1}^N \alpha_i A_i(t). \tag{14}$$

Then, our new max-weight policy chooses a scheduling decision that minimizes the one-slot Lyapuonv drift in every time-slot.

$$\pi^{mw}(t) = \underset{i\in[N]}{\operatorname{argmin}} \mathbb{E}\Big[L(t+1) - L(t)\Big|\mathbf{A}(t)\Big]. \tag{15}$$

The structure of the max-weight policy can be obtained by simplifying the expression in (15) above:

$$\begin{aligned}\pi^{mw}(t) &= \underset{i\in[N]}{\operatorname{argmin}}\Big(\sum_{j=1}^N p_{ij}\alpha_j A_j(t)\Big)\\ &= \underset{i\in[N]}{\operatorname{argmin}}\Big(\alpha_i A_i(t) + \sum_{j\neq i} p_{ij}\alpha_j A_j(t)\Big).\end{aligned} \tag{16}$$

Note that the max-weight policy proposed in [24, Sec. 3.2.4] for the uncorrelated setting schedules the source i with the highest value of $\sqrt{w_i}A_i(t)$. On the other hand, our correlated max-weight policy (16) adds up the "value" of each possible AoI reduction including correlated updates, weighted by the of probability of correlation.

The following theorem shows that the new max-weight policy defined in (16) enjoys a similar factor-2 optimality guarantee as the optimal stationary randomized policy which is better than the factor-4 guarantee for the quadratic max-weight policy described in (11).

Theorem 5. *Consider the max-weight policy $\boldsymbol{\pi}^{mw}$ and an optimal policy $\boldsymbol{\pi}^{opt}$ that solves the general problem* (3). *Let $\bar{\mathbf{A}}^{mw}$ denote the average AoIs under $\boldsymbol{\pi}^{mw}$ and $\bar{\mathbf{A}}^{opt}$ denote the average AoIs under $\boldsymbol{\pi}^{opt}$. Then,*

$$\frac{\sum_{i=1}^N w_i \bar{A}_i^{mw}}{\sum_{i=1}^N w_i \bar{A}_i^{opt}} \leq 2. \tag{17}$$

Proof. See Appendix E in the technical report [30]. □

While we show a factor of two optimality result above, we will show via simulations in Sec. 7 that the max-weight policy performs almost as well as the theoretical lower bound derived in Appendix C in the technical report [30], and also outperforms the optimal stationary randomized policy in practice.

4 SCALING

In this section, we consider how correlation improves information freshness as network sizes scale. For deriving our scaling results, we will focus on the equal weights setting, i.e. $w_i = \frac{1}{N}, \forall i \in [N]$. However, similar scaling results will hold qualitatively for general weight configurations.

The lemma below characterizes the minimum average AoI of a network with N uncorrelated sources with equal weights. We will use this as a comparison baseline to see how the degree of correlation improves AoI.

Lemma 2. *Consider N uncorrelated sources sending updates to a base station. Let the scheduling weights for all sources be equal, i.e. $w_i = \frac{1}{N}, \forall i$. Then, the optimal weighted sum AoI satisfies the following:*

$$\sum_{i=1}^N w_i \bar{A}_i^{opt} = \frac{N+1}{2} \sim \Theta(N). \tag{18}$$

Proof. In the setting with symmetric weights, the greedy or the round-robin policy is known to be optimal [12]. It is easy to see that the average AoI under the round-robin policy is simply the average of the sequence 1, 2, ..., N which is $\frac{N+1}{2}$. □

4.1 An Upper Bound

We derive a general upper bound for the weighted-sum average AoI by representing information about the correlation matrix $\mathbf{P}$ for N sources with a directed graph.

Consider a correlation threshold $p \in (0, 1)$. We are interested in the entries of the correlation matrix above this threshold, i.e the pairs of sources that are significantly correlated. To do so, we construct a directed graph $\mathcal{G}(V, E)$ on the set of sources. For every ordered pair of sources (i, j) such that $p_{ij} > p$, we add the edge (i, j) to the graph $\mathcal{G}$. Trivially, every node in $\mathcal{G}$ must have a self-loop, since $p_{ii} = 1, \forall i$. We will show that the average AoI of the network can be upper-bounded by analyzing the properties of these constructed graphs. To derive this upper bound, we first need to define the notion of a vertex cover for a directed graph.

Vertex Cover: Given a directed graph $\mathcal{G}(V, E)$, a vertex cover is defined to be a set of vertices $S \subseteq V$ if for every vertex $i \in V$, there exists a vertex $j \in S$ such that the edge (j, i) is in the set E.

The following theorem relates the average AoI to the size of the *minimum vertex cover* of the graph $\mathcal{G}$ constructed using the correlation threshold p.

Theorem 6. *Consider N sources with the correlation matrix $\mathbf{P}$. Given a correlation threshold $p > 0$, construct a directed graph $\mathcal{G}$ that represents pairs of source with correlation higher than the threshold. Assuming equal weights $w_i = \frac{1}{N}, \forall i$, the optimal weighted sum average AoI satisfies:*

$$\sum_{i=1}^{N} w_i \bar{A}_i^{opt} \leq \frac{N_{cov}}{p}, \tag{19}$$

where N_{cov} is the size of a minimum vertex cover for the graph $\mathcal{G}$ and p is the correlation threshold.

Proof. See Appendix A. □

This upper-bound allows us to relate the degree of correlation or information sharing between the sources to the average AoI. If a small subset of sources, of say size $O(\log N)$, are highly correlated with all other sources in the network, then average AoI with correlation is only $O(\log N)$ as well. In this case, correlation leads to a significant reduction in the average AoI compared to the uncorrelated case, which is $O(N)$ as shown in Lemma 2. If no such small vertex coverings exist or if correlation probabilities are very small, then a scheduler likely needs to communicate with a large fraction of all the sources and the average AoI would grow as $O(N)$, similar to the uncorrelated case.

Further, we show that it is possible to construct a correlation matrix such that the average AoI of the correlated max-weight given by (16) is $O(1)$ while the average AoI of the uncorrelated max-weight policy from [12] is $\Theta(N)$. This means that the performance gap between policies that consider and ignore correlation can grow linearly with the size of the network.

Theorem 7. *Given any number of sources N with equal weights $\frac{1}{N}$, we consider two policies - the correlated max-weight policy π^{mw} proposed in Sec. 3.2 and the max-weight policy that does not take correlation into account π^u which was proposed in [12]. Let the average AoI of a source i under π^{mw} be $\bar{A}_i^{mw}$ and under π^u be $\bar{A}_i^u$. Then, there exists a correlation matrix $\mathbf{P} \in [0, 1]^{N \times N}$ such that the following holds:*

$$\frac{\sum_{i=1}^{N} w_i \bar{A}_i^u}{\sum_{i=1}^{N} w_i \bar{A}_i^{mw}} \sim \Omega(N). \tag{20}$$

Proof. Consider N sources with the following correlation matrix:

$$\mathbf{P} = \begin{bmatrix} 1 & p & p & \cdots & p \\ p & 1 & 0 & \cdots & 0 \\ p & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \cdots & 1 & \vdots \\ p & 0 & 0 & \cdots & 1 \end{bmatrix}. \tag{21}$$

The stationary randomized policy $\boldsymbol{\pi}^{(1)}$ that puts all scheduling weight on source 1 achieves a weighted-sum AoI of:

$$\sum_{i=1}^{N} w_i \bar{A}_i = \frac{1}{N} + \sum_{i=2}^{N} \frac{1}{N} \frac{1}{p} \leq \frac{1}{p}. \tag{22}$$

Thus, the optimal stationary randomized policy $\boldsymbol{\pi}^*$ also achieves weighted-sum average AoI that is at-least as good as that of $\boldsymbol{\pi}^{(1)}$. This implies:

$$\sum_{i=1}^{N} w_i \bar{A}_i^* \leq \frac{1}{p}, \tag{23}$$

where $\bar{A}_i^*$ is the average AoI of source i under the optimal stationary randomized policy $\boldsymbol{\pi}^*$. In Appendix E in the technical report [30], we show that the performance of the max-weight policy is upper-bounded by the policy of the optimal stationary randomized policy. Thus, we get:

$$\sum_{i=1}^{N} w_i \bar{A}_i^{mw} \leq \sum_{i=1}^{N} w_i \bar{A}_i^* \leq \frac{1}{p}, \tag{24}$$

where $\bar{A}_i^{mw}$ is the average AoI of source i under the max-weight policy $\boldsymbol{\pi}^{mw}$ described by (16).

Now consider the performance of the uncorrelated max-weight policy $\boldsymbol{\pi}^u$. Since all the weights are symmetric, the optimal policy is greedy or max-AoI-first. In Appendix G in the technical report [30], we show that the max-AoI-first policy behaves similar to a round-robin policy on sources 2, ..., N while occasionally scheduling source 1. Specifically, we derive a lower bound on the weighted-sum AoI under policy $\boldsymbol{\pi}^u$:

$$\sum_{i=1}^{N} w_i \bar{A}_i^u \geq \frac{(N-1)^2}{2N - (N+1)(1-p)^{N-1}}. \tag{25}$$

This bound holds under the assumtion that $p \geq \frac{1}{N-1}$. Combining (24) and (25), we get:

$$\frac{\sum_{i=1}^{N} w_i \bar{A}_i^u}{\sum_{i=1}^{N} w_i \bar{A}_i^{mw}} \geq \frac{p(N-1)^2}{2N - (N+1)(1-p)^{N-1}}. \tag{26}$$

Assuming p to be a fixed constant that does not depend on N completes the proof. Note that a similar result also holds for p

scaling with N, for example if $p \sim \frac{1}{N^\beta}, \beta < 1$ the performance gap is $\Omega(N^{1-\beta})$. □

The intuition behind this result is straightforward: if there is one source that is correlated even by a small amount with all other sources then scheduling just that source all the time should be sufficient. The gap in performance is achieved by assuming that other sources are not correlated with each other.

Theorem 7 provides a key insight: *the gain in performance cannot be obtained by using policies that ignore correlation.* It is necessary to design scheduling policies that take the correlation structure into account, especially for correlation graphs where the degree distribution is highly skewed. Next, we consider scaling in the special case where the correlation matrix can be represented by random geometric graphs. These graphs are commonly used to model wireless sensors networks monitoring spatial phenomenon.

4.2 Random Geometric Graphs

Consider N sources distributed uniformly at random on the unit square $[0,1]^2$. Each source has information about itself, by definition. Thus $p_{ii} = 1, \forall i$. Further, if the distance between sources i and j is less than a threshold r, then we set $p_{ij} = p$ and $p_{ji} = p$, otherwise we set $p_{ij} = 0$ and $p_{ji} = 0$. This leads to a symmetric correlation matrix $\mathbf{P}$. Constructing a graph that connects sources with correlation leads to the random geometric graph $\mathcal{G}(N, r)$.

The following theorem looks at the scaling of AoI under this geometric correlation structure.

Theorem 8. *Consider a symmetric correlation matrix generated by creating a random geometric graph $\mathcal{G}(N, r)$ on the two dimensional unit square and setting correlation probabilities for neighbors to be p. Assuming equal weights $w_i = \frac{1}{N}, \forall i$, the weighted sum average AoI satisfies:*

$$\sum_{i=1}^{N} w_i \bar{A}_i^{opt} \leq \frac{2}{pr^2}. \tag{27}$$

Proof. See Appendix H in the technical report [30]. □

Note that the connectivity threshold of a random geometric graphs occurs at $r \sim \Theta\left(\sqrt{\frac{\log N}{N}}\right)$. For this choice of r, we can see that the overall AoI of the network is $O\left(\frac{N}{p \log N}\right)$. This leads to a factor of $\log N$ reduction over the uncorrelated case analyzed in Lemma 2, assuming p is a constant. Further, if $r \sim \Theta(1)$, then the AoI is also $\Theta(1)$ and there is factor of N reduction in AoI compared to the uncorrelated case.

5 ROBUSTNESS

In this section, we consider a different way to model correlation between a set of sources. We will show that for stationary randomized policies, this model is equivalent to our proposed model. Later, in Sec. 7, we will show via numerical results that the scheduling policies designed in Sec. 3 also perform well for this new correlation structure. This suggests that if the "amount of correlation" is the same on average, then the AoI performance tends to be similar, regardless of how correlation is modeled.

As before, consider a set of N sources and a correlation matrix $\mathbf{P}$ with entries $p_{ij} \in [0,1]$ that represent the amount of information shared at source i about source j. However, instead of a Bernoulli random variable indicating whether i either does or does not currently have information about j, we now assume that i always has a constant p_{ij} fraction of information regarding the state of j. In other words, an update from i can reduce the current uncertainty regarding j by a fraction of p_{ij}. This correlation structure is convenient for modeling settings where different sources are sensing information and have *partially* overlapping ranges of sensing, for example cameras with overlapping fields of view as studied in [5, 6]. Thus, p_{ij} can be viewed as the fraction of source j's range that source i also covers.

Consequently, the evolution of AoI for this model differs from the one studied in earlier sections. Let $u_i(t)$ be indicator variables that denote whether source i transmits in time-slot t or not. Then, AoI for source i evolves as follows:

$$A_i(t+1) = (1 - p_{ji})A_i(t) + 1, \text{ if } u_j(t) = 1. \tag{28}$$

The equation above formalizes the notion that uncertainty regarding source i drops by a fraction p_{ji}, when j sends an update.

The following theorem establishes the equivalence between this new model and the one proposed in Sec. 2, under stationary randomized policies. In fact, we prove the equivalence result for a much more general class of correlation structures, where correlation is defined by i.i.d. random variables $X_{ij}(t) \in [0,1]$ such that $\mathbb{E}[X_{ij}(t)] = p_{ij}$. Our original model assumes $X_{ij}(t) \sim Bern(p_{ij})$, while the model proposed in this section assumes $X_{ij}(t) = p_{ij}$ and both of them belong to this class of correlation structures.

Theorem 9. *Consider N sources and a correlation matrix $\mathbf{P}$ with the AoIs for each source evolving according to:*

$$A_i(t+1) = A_i(t) + 1 - \left(\sum_{j=1}^{N} u_j(t) X_{ji}(t)\right) A_i(t), \forall i, t. \tag{29}$$

Here $X_{ji}(t)$ are i.i.d random variables such that $X_{ji}(t) \in [0,1]$ and $\mathbb{E}[X_{ji}(t)] = p_{ji}$. Given any stationary randomized policy with scheduling probabilities $\boldsymbol{\pi}$, if $\sum_{j=1}^{N} \pi_j p_{ji} > 0$ for a source i, then the average AoI for this source is given by:

$$\bar{A}_i \triangleq \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} A_i(t)\right] = \frac{1}{\sum_{j=1}^{N} \pi_j p_{ji}}. \tag{30}$$

If $\sum_{j=1}^{N} \pi_j p_{ji} = 0$ for some source i, then the base station never receives any information regarding this source and its average AoI $\bar{A}_i$ is unbounded.

Proof. See Appendix I in the technical report [30]. □

Observe that the expression for the average AoI is the same as the one derived in Theorem 1. Thus the procedure to find the optimal stationary randomized policy and the policy itself also remain the same. We will show later via simulations that the performance of the max-weight policy is also similar for different distributions of $X_{ji}(t)$, which suggests that our analysis is fairly robust to the way in which correlation is modeled.

6 LEARNING THE CORRELATION MATRIX

Until now, we have focused on cases where the correlation structure is known in advance and fixed over time, and we use this information to analyze and optimize AoI. In this section, we consider the setting when the correlation matrix is unknown and possibly varying with time.

6.1 Online Setting

Learning the correlation matrix in an online setting where the $\mathbf{P}$ changes over time, even slowly, is a challenging problem. This setting is of interest because correlation between source tends to be time-varying in practice, especially in settings involving mobility.

We want to implement a max-weight style policy that gradually updates its policy parameters to be able to track changes in the environment. However, note that the max-weight policy proposed in Sec. 3.2 in (16) requires us to solve for the optimal stationary randomized policy. Thus, whenever the correlation matrix changes, we would have to recalculate $\boldsymbol{\pi}^*$ and the Lyapunov function weights α_i. To avoid this added computation, we will use the quadratic max-weight policy, given by (11). Recall that this policy is based on a quadratic Lyapunov function and does not require us to calculate the the optimal stationary randomized policy $\boldsymbol{\pi}^*$ repeatedly as $\boldsymbol{P}$ varies.

Algorithm 1 uses an exponential moving average to keep track of the correlation probabilities and then runs the *quadratic* max-weight scheduler from (11) using the estimated correlation matrix. We call this the *EMA-max-weight* policy.

Intuitively, if the probabilities change slowly over time, the exponential moving average estimate should be able to closely track the actual correlation matrix and the EMA-max-weight policy would perform similar to a max-weight policy that knows the entire sequence of correlation matrices in advance. We confirm that this is indeed the case via simulations in Sec. 7.

Algorithm 1: Exponential Moving Average Max-Weight

Input : parameter $\alpha > 0$

Start by assuming no correlation, i.e. set

$$\hat{\mathbf{P}}(1) = \mathbf{I}$$

while $t \in 1, ..., T$ **do**

Run quadratic max-weight using $\hat{\mathbf{P}}(t)$:

$$s = \underset{i\in[N]}{\operatorname{argmin}} \sum_{j=1}^{N} w_j \hat{p}_{ij}(t) A_j(t) \Big(A_j(t) + 2 \Big).$$

Schedule source s and receive correlated updates

Update AoIs for every source j:

$$A_j(t+1) = \begin{cases} 1, & \text{if } s \text{ sent an update about } j, \\ A_j(t)+1, & \text{otherwise.} \end{cases}$$

Update the correlation matrix, for source s

$$\hat{p}_{sj}(t+1) = \begin{cases} (1-\alpha)\hat{p}_{sj}(t) + \alpha, & \text{if } s \text{ sent an update about } j, \\ (1-\alpha)\hat{p}_{sj}(t), & \text{otherwise.} \end{cases}$$

For sources other than s:

$$\hat{p}_{ij}(t+1) = \hat{p}_{ij}(t), \forall i \neq s.$$

end

7 NUMERICAL RESULTS

First, we consider random geometric graphs and see how different scheduling policies perform as the number of sources N increases. In particular, we simulate graphs $\mathcal{G}(N, r)$ on the unit square $[0, 1]^2$ where $r = 1.1\sqrt{\frac{\log N}{N}}$ is slightly above the connection threshold. For each pair of nodes (i, j) in this graph, we set $P_{ij} = P_{ji} = 0.7$ if the nodes are closer than the distance r and set $P_{ij} = P_{ji} = 0$ otherwise. We fix all weights to be equal, i.e. $w_i = \frac{1}{N}$.

For each value of N, we compare the performance of four different policies - 1) the uniform stationary randomized policy which would have been the optimal stationary randomized policy if we ignore correlation, 2) our optimal stationary randomized policy which solves (5), 3) the max-Age-first or greedy policy which would have been the optimal max-weight policy if we ignore correlation and 4) our proposed max-weight policy from Sec. 3.2.

Figure 3 plots the performance of these four policies as N increases. Each data point is computed by averaging over 10 random graph instances and running the policy for 10000 time-slots for each such instance. We also plot a lower bound for average AoI computed using analysis from Appendix C in the technical report [30]. We observe that our proposed methods clearly outperform

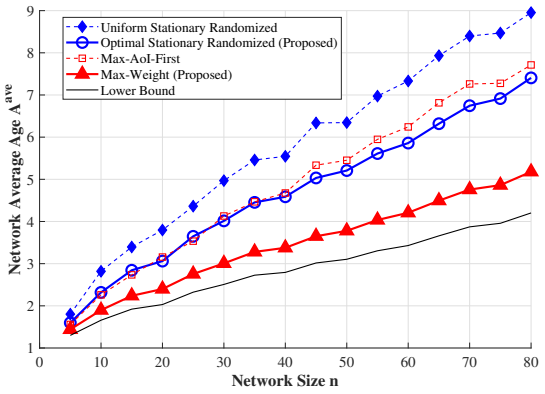

Figure 3: Average AoI vs network size N for random geometric graphs.

policies that ignore the correlation between sources for scheduling. In particular, our proposed max-weight policy outperforms max-AoI-first by almost 33%.

Next, we perform the same exercise for hyperbolic geometric graphs. Hyperbolic geometric graphs are generated by choosing points on a two-dimensional hyperbolic space and connecting vertices that are closer than the distance R where distance is measured along hyperbolic geodesics. The most important feature of hyperbolic random graphs is that when their parameters are chosen

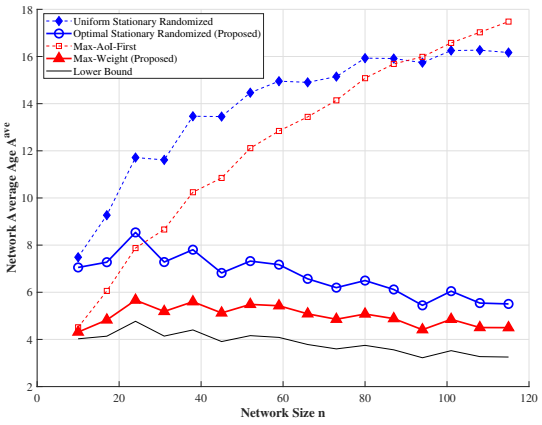

Figure 4: Average AoI vs network size N for hyperbolic geometric graphs.

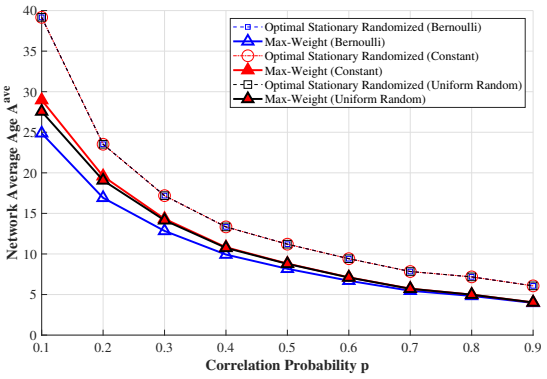

Figure 5: Average AoI vs correlation probability p.

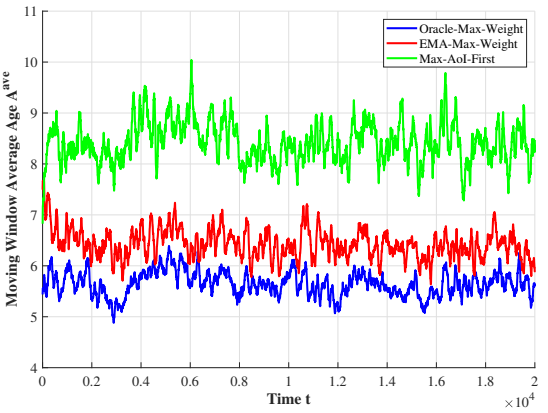

Figure 6: Moving window average AoI vs time.

appropriately, the degree distribution of nodes becomes *scale-free* [18, 21]. Importantly, for our simulation setup, we expect that scale-free degree distributions lead to correlation matrices with lower vertex covering numbers. According to Theorem 6, this should lead to larger performance gaps between policies that consider correlation and policies that ignore it. Figure 4 confirms this: as the number of sources increase, the gap in performance between our proposed policies and policies that ignore correlation also increases.

Next, we compare the performance of the optimal stationary randomized policy and our proposed max-weight policy for different correlation models. Specifically, we consider three correlation models: 1) Bernoulli correlation: $X_{ij}(t) \sim Bern(p_{ij})$, which is the main focus of this work, 2) Constant correlation: $X_{ij}(t) = p_{ij}$ as introduced in Sec. 5, and 3) Uniform correlation: $X_{ij}(t) = p_{ij} + Unif([-0.1, 0.1])$, i.e. correlation is chosen uniformly at random from the interval $[p_{ij} - 0.1, p_{ij} + 0.1]$.

Figure 5 compares the performance of the two policies (optimal stationary randomized and max-weight) on random geometric graphs $\mathcal{G}(N = 90, r = 0.25)$ while varying the correlation parameter p from 0.1 to 0.9 for the three different correlation models discussed above.

We observe that for the stationary randomized policy, the average AoI values are the same across correlation models for each value of p. This is consistent with the result we derived in Theorem 9. For the max-weight policy, the average AoI values are close to one another for each value of p. Further, as the correlation increases, the gap between the average AoI for different correlation models decreases. Combined with our observation from Sec. 5, where we showed that average AoI under stationary randomized policies is the same for a large class of correlation models, this supports the claim that our results are fairly robust to the way in which correlation is modeled. Another important observation from Figure 5 is the inverse dependence of the average AoI on the correlation probability p, consistent with our upper-bound for random-geometric graphs from Theorem 8.

Finally, we consider a setting with time-varying correlation probabilities. Specifically, we consider a random geometric graph with $N = 90$ sources and a connectivity radius of 0.25 on the unit square $[0, 1]^2$. However, unlike the previous simulations, we assume that these sources are *mobile* and move according to Brownian motion on $[0, 1]^2$ with a maximum velocity of 0.01. As the sources move around randomly, the distances between them change and so does the corresponding correlation matrix $\mathbf{P}$. Figure 6 plots the performance of three different scheduling policies in this time-varying setting. To measure the performance, we consider a windowed time-average of the network AoI with a window size of 100. As the sources connectivity changes, we see the windowed time-average AoI change in response for each scheduling policy.

The three policies we consider are as follows. First, we consider the max-AoI-first policy which is completely oblivious of the correlation while making scheduling decisions. Second, we consider the EMA-max-weight policy from Sec. 6.1 which has no information about correlation in the beginning but gradually learns the matrix $\mathbf{P}$ and adapts to changes in it using an exponential moving average. Specifically, we set the learning rate $\alpha = 0.4$. Third, we consider the hypothetical *oracle-max-weight* policy. This is an omniscient policy that knows the current correlation matrix exactly in each

time-slot and uses this information to run the quadratic max-weight scheduler (11).

We observe that EMA-max-weight is able to track the performance of the oracle-max-weight policy with only a small gap, indicating that it is able to learn the correlation structure and adapt in response to it. Further, max-AoI-first is not able to do so and has a larger gap in performance compared to the oracle-max-weight policy.

8 CONCLUSION

In this work, we formulated a simple model to study the timely monitoring of correlated sources over a wireless network. Using this model, we proposed new scheduling policies that optimize weighted-sum average Age of Information (AoI) in the presence of correlation. These policies have constant-factor optimality guarantees. We derived scaling results that illustrate how AoI improves in large networks in the presence of correlation and discussed how our model is relatively robust to correlation modeling assumptions. Lastly, We also developed a novel approach based on exponential moving averages that schedules correlated sources in a time-varying setting.

Important directions of future work involve proving performance bounds on EMA-max-weight under assumptions on how quickly the correlation matrices change, and incorporating more general AoI cost functions. A drawback that is worth mentioning is that we consider correlation to be a pairwise notion - information sharing happens only between pairs of sources. However, modeling more general notions of correlation and coupling between sources while still keeping analysis tractable is a challenging open problem.

9 ACKNOWLEDGEMENTS

This work was supported in part by NSF Grant CNS-1713725.

REFERENCES

[1] Ahmed M Bedewy, Yin Sun, and Ness B Shroff. 2019. Minimizing the age of information through queues. *IEEE Trans. Information Theory* 65, 8 (2019), 5215–5232.
[2] Jaya Prakash Champati, Mohammad H Mamduhi, Karl H Johansson, and James Gross. 2019. Performance characterization using aoi in a single-loop networked control system. In *Proc. IEEE INFOCOM AoI Workshop*. 197–203.
[3] Shahab Farazi, Andrew G Klein, John A McNeill, and D Richard Brown. 2018. On the age of information in multi-source multi-hop wireless status update networks. In *Proc. IEEE Int. Workshop Signal Process. Adv. Wireless Commun. (SPAWC)*. 1–5.
[4] Michael Grant and Stephen Boyd. 2014. CVX: Matlab software for disciplined convex programming, version 2.1.
[5] Qing He, Gyorgy Dan, and Viktoria Fodor. 2018. Minimizing age of correlated information for wireless camera networks. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 547–552.
[6] Qing He, György Dán, and Viktoria Fodor. 2019. Joint assignment and scheduling for minimizing age of correlated information. *IEEE/ACM Transactions on Networking* 27, 5 (2019), 1887–1900.
[7] Longbo Huang and Eytan Modiano. 2015. Optimizing age-of-information in a multi-class queueing system. In *Proc. IEEE Int. Symp. Information Theory (ISIT)*. 1681–1685.
[8] Yoshiaki Inoue, Hiroyuki Masuyama, Tetsuya Takine, and Toshiyuki Tanaka. 2018. A general formula for the stationary distribution of the age of information and its application to single-server queues. *arXiv preprint arXiv:1804.06139* (2018).
[9] Zhiyuan Jiang and Sheng Zhou. 2019. Status from a random field: How densely should one update?. In *2019 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 1037–1041.
[10] Igor Kadota and Eytan Modiano. 2019. Minimizing the age of information in wireless networks with stochastic arrivals. *IEEE Transactions on Mobile Computing* 20, 3 (2019), 1173–1185.
[11] Igor Kadota, Abhishek Sinha, and Eytan Modiano. 2019. Scheduling algorithms for optimizing age of information in wireless networks with throughput constraints. *IEEE/ACM Trans. Netw.* 27, 4 (2019), 1359–1372.
[12] Igor Kadota, Abhishek Sinha, Elif Uysal-Biyikoglu, Rahul Singh, and Eytan Modiano. 2018. Scheduling policies for minimizing age of information in broadcast wireless networks. *IEEE/ACM Trans. Netw.* 26, 6 (2018), 2637–2650.
[13] Anders E Kalør and Petar Popovski. 2019. Minimizing the age of information from sensors with common observations. *IEEE Wireless Communications Letters* 8, 5 (2019), 1390–1393.
[14] Clement Kam, Sastry Kompella, and Anthony Ephremides. 2013. Age of information under random updates. In *Proc. IEEE Int. Symp. Information Theory (ISIT)*. 66–70.
[15] Sanjit Kaul, Roy Yates, and Marco Gruteser. 2012. Real-time status: How often should one update?. In *Proc. IEEE INFOCOM*. 2731–2735.
[16] Markus Klügel, Mohammad H Mamduhi, Sandra Hirche, and Wolfgang Kellerer. 2019. Aoi-penalty minimization for networked control systems with packet loss. In *Proc. IEEE INFOCOM AoI Workshop*. 189–196.
[17] Antzela Kosta, Nikolaos Pappas, Vangelis Angelakis, et al. 2017. Age of information: A new concept, metric, and tool. *Foundations and Trends in Networking* 12, 3 (2017), 162–259.
[18] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. 2010. Hyperbolic geometry of complex networks. *Physical Review E* 82, 3 (2010), 036106.
[19] Chengzhang Li, Shaoran Li, and Y Thomas Hou. 2019. A general model for minimizing age of information at network edge. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 118–126.
[20] Tasmeen Zaman Ornee and Yin Sun. 2019. Sampling for remote estimation through queues: Age of information and beyond. *IEEE Int. Symp. Model. Optim. Mobile, Ad Hoc Wireless Netw. (WiOpt)* (2019).
[21] Fragkiskos Papadopoulos, Dmitri Krioukov, Marián Boguná, and Amin Vahdat. 2010. Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces. In *2010 Proceedings IEEE Infocom*. IEEE, 1–9.
[22] Sudheer Poojary, Sanidhay Bhambay, and Parimal Parag. 2017. Real-time status updates for correlated source. In *2017 IEEE Information Theory Workshop (ITW)*. IEEE, 274–278.
[23] Yin Sun and Benjamin Cyr. 2019. Sampling for data freshness optimization: Non-linear age functions. *IEEE Journal Commun. Netw.* 21, 3 (2019), 204–219.
[24] Yin Sun, Igor Kadota, Rajat Talak, and Eytan Modiano. 2019. Age of information: A new metric for information freshness. *Synthesis Lectures on Communication Networks* 12, 2 (2019), 1–224.
[25] Yin Sun, Yury Polyanskiy, and Elif Uysal-Biyikoglu. 2017. Remote estimation of the Wiener process over a channel with random delay. In *Proc. IEEE Int. Symp. Information Theory (ISIT)*. 321–325.
[26] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff. 2017. Update or Wait: How to Keep Your Data Fresh. *IEEE Trans. Information Theory* 63, 11 (Nov. 2017), 7492–7508.
[27] Rajat Talak, Sertac Karaman, and Eytan Modiano. 2018. Optimizing information freshness in wireless networks under general interference constraints. In *Proc. ACM Int. Symp. Mobile Ad Hoc Netw. Comput. (MobiHoc)*. 61–70.
[28] Vishrant Tripathi, Luca Ballotta, Luca Carlone, and Eytan Modiano. 2021. Computation and communication co-design for real-time monitoring and control in multi-agent systems. In *2021 19th International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks (WiOpt)*. IEEE, 1–8.
[29] Vishrant Tripathi and Eytan Modiano. 2019. A whittle index approach to minimizing functions of age of information. In *Proc. 57th Allerton Conf. Commun. Control Comput.* IEEE, 1160–1167.
[30] Vishrant Tripathi and Eytan Modiano. 2022. *Optimizing Age of Information with Correlated Sources*. Technical Report. https://arxiv.org/abs/2208.01738.
[31] Vishrant Tripathi and Sharayu Moharir. 2017. Age of information in multi-source systems. In *Proc. IEEE Global Commun. Conf. (GLOBECOM)*. 1–6.
[32] Bo Zhou and Walid Saad. 2020. On the age of information in Internet of Things systems with correlated devices. In *GLOBECOM 2020-2020 IEEE Global Communications Conference*. IEEE, 1–6.

A PROOF OF THEOREM 6

Consider N sources with the correlation matrix $\mathbf{P}$. Given a correlation threshold $p > 0$, construct a directed graph $\mathcal{G}$ that represents pairs of source with correlation higher than the threshold.

Further, assume that the set $S \subseteq [N]$ is a minimum size vertex covering of the graph $\mathcal{G}$. Recall that we defined the notion of vertex covering for directed graphs in Sec. 4. Let the size of this minimum vertex covering be denoted by N_{cov}.

We will show that under a specific scheduling policy π, the average AoI of every source is upper-bounded by $\frac{N_{cov}}{p}$. Consequently, the weighted sum of average AoIs under an optimal policy will also be upper-bounded by $\frac{N_{cov}}{p}$. The scheduling policy $\boldsymbol{\pi}$ we analyze is round-robin scheduling of sources in the covering set S. This is a cyclic policy of length N_{cov} time-slots where each source in the covering set S gets scheduled once, after which the scheduling pattern repeats every N_{cov} time-slots.

Before we analyze the performance of this policy, we present a lemma that discusses the monotonicity of AoI with correlation.

Lemma 3. *Consider N sources with two different correlation matrices $\mathbf{P}$ and $\mathbf{P}'$. If $p_{ij} \geq p'_{ij}, \forall (i,j) \in [N] \times [N]$ then under a fixed sequence of scheduling decisions $\boldsymbol{\pi}$, the following holds:*

$$\bar{A}_i \leq \bar{A}'_i, \forall i \in [N]. \tag{31}$$

Here, $\bar{A}_i$ is the average AoI of source i under policy π with the correlation matrix $\mathbf{P}$, while $\bar{A}'_i$ is the average AoI of source i under policy π with the correlation matrix $\mathbf{P}'$.

Proof. The proof is easy to see via a stochastic dominance argument. Let $X_{ji}(t)$ be an indicator variable denoting whether j had information about i at time-slot t given the correlation matrix $\boldsymbol{P}$, and likewise $X'_{ji}(t)$ for the matrix $\boldsymbol{P'}$. Then, for all pairs (i,j) and for all time-slots t, $X_{ij}(t) \sim Bern(p_{ij})$ and $X'_{ij}(t) \sim Bern(p'_{ij})$, where $p_{ij} \geq p'_{ij}$. Thus, $X_{ij}(t) \geq_{st.} X'_{ij}(t)$.

Now, the AoI evolution (1) can be rewritten as:

$$\begin{aligned} A_i(t+1) &= A_i(t) + 1 - A_i(t) \sum_{j=1}^{N} u_j(t) X_{ji}(t), \text{ and} \\ A'_i(t+1) &= A'_i(t) + 1 - A'_i(t) \sum_{j=1}^{N} u'_j(t) X'_{ji}(t), \forall i, t. \end{aligned} \tag{32}$$

Since we have fixed the sequence of scheduling decisions $\boldsymbol{\pi}$ beforehand irrespective of how AoI evolves, so the scheduling decisions remain the same for both correlation matrices, i.e. $u_j(t) = u'_j(t), \forall t$. Setting $A_i(1) = A'_i(1)$, (32) implies that:

$$A_i(t) \leq_{st.} A'_i(t). \tag{33}$$

This, in turn, further implies that:

$$\frac{1}{T} \sum_{t=1}^{T} A_i(t) \leq_{st.} \frac{1}{T} \sum_{t=1}^{T} A'_i(t), \forall i, t. \tag{34}$$

Using the equation above and the definition of average AoI (2), we can conclude that:

$$\bar{A}_i \leq_{st.} \bar{A}'_i, \forall i \in [N]. \tag{35}$$

If the average AoI limits exist, then $\bar{A}_i$ and $\bar{A}'_i$ are just constants and the stochastic dominance becomes a simple inequality. This completes the proof. □

The lemma above shows a rather simple result: AoI improves with correlation. To upper-bound the average AoI of sources under the round-robin policy $\boldsymbol{\pi}$, it is sufficient to analyze AoI for a new correlation matrix $\mathbf{P}'$, which is defined as:

$$p'_{ij} = \begin{cases} p, & \text{if } p_{ij} \geq p. \\ 0, & \text{otherwise.} \end{cases} \tag{36}$$

Clearly, $\mathbf{P}'$ is element-wise smaller than $\mathbf{P}$ and $\boldsymbol{\pi}$ is a fixed sequence of scheduling decisions, so we can apply Lemma 3. Now, to analyze AoI of a source i under the policy $\boldsymbol{\pi}$ with the correlation matrix $\mathbf{P}'$, we consider two cases.

Case 1: Source i is in the vertex covering set S. In this case, note that this source is guaranteed to be scheduled at least once in every N_{cov} time-slots due to the way round-robin scheduling works.

$$\begin{aligned} \bar{A}_i &\leq \frac{\sum_{k=1}^{N_{cov}} k}{N_{cov}} \leq \frac{N_{cov}+1}{2} \\ &\leq \frac{N_{cov}}{p}. \end{aligned} \tag{37}$$

The last inequality follows since $p \leq 1$ and $N_{cov} \geq 1$.

Case 2: Source i is not in the vertex covering set S. By the definition of a vertex covering, there must be at least one source $j \in S$ such that the edge (j,i) is in the graph $\mathcal{G}$ or alternatively, that $p'_{ji} = p$. This means that whenever the scheduling policy schedules source j, the base station also receives an update regarding source i with probability p. Note that the scheduling policy π schedules source j once every N_{cov} time-slots. Then, the time intervals between two successful correlated update deliveries regarding source i by source j are $N_{cov}I_1, N_{cov}I_2, ..., N_{cov}I_K$ where $I_1, I_2, ..., I_K$ are i.i.d. geometrically distributed random variables with parameter p. It is possible that there are other sources also correlated with source i that deliver correlated updates to the base-station between two consecutive correlated updates from source j. For the sake of upper-bounding the AoI we can safely ignore any such updates. Then, the average AoI of source i can be upper-bounded by:

$$\begin{aligned} \bar{A}_i &\leq \lim_{K\to\infty} \frac{\sum_{k=1}^{K} (1 + 2 + ... + N_{cov}I_k)}{\sum_{k=1}^{K} N_{cov}I_k} \\ &\leq \lim_{K\to\infty} \frac{\sum_{k=1}^{K} N_{cov}^2 I_k^2 + N_{cov}I_k}{2\sum_{k=1}^{K} N_{cov}I_k} \\ &\leq \frac{1}{2} + N_{cov} \lim_{K\to\infty} \frac{\sum_{k=1}^{K} I_k^2}{2\sum_{k=1}^{K} I_k}. \end{aligned} \tag{38}$$

Applying the law of large numbers and using the moments of a geometric random variable, we get:

$$\bar{A}_i \leq \frac{1}{2} + N_{cov} \frac{(2-p)/p^2}{2/p} \leq \frac{N_{cov}}{p}. \tag{39}$$

The last inequality follows since a vertex covering must have at least one vertex, i.e. $N_{cov} \geq 1$.

Together, (37) and (39) imply that the average AoI $\bar{A}_i$ for every source i under the vertex cover round-robin policy $\boldsymbol{\pi}$ is upper-bounded by $\frac{N_{cov}}{p}$. Clearly, the performance of the policy that achieves minimum weighted-sum average AoI must be better than that of $\boldsymbol{\pi}$. This implies:

$$\sum_{i=1}^{N} w_i \bar{A}_i^{opt} \leq \sum_{i=1}^{N} w_i \bar{A}_i \leq \frac{N_{cov}}{p}. \tag{40}$$

This completes the proof.