

Optimizing Age of Information in Wireless Networks with Throughput Constraints

Igor Kadota, Abhishek Sinha and Eytan Modiano
Laboratory for Information & Decision Systems, MIT

Abstract—Age of Information (AoI) is a performance metric that captures the freshness of the information from the perspective of the destination. The AoI measures the time that elapsed since the generation of the packet that was most recently delivered to the destination. In this paper, we consider a single-hop wireless network with a number of nodes transmitting time-sensitive information to a Base Station and address the problem of minimizing the Expected Weighted Sum AoI of the network while simultaneously satisfying timely-throughput constraints from the nodes.

We develop three low-complexity transmission scheduling policies that attempt to minimize AoI subject to minimum throughput requirements and evaluate their performance against the optimal policy. In particular, we develop a randomized policy, a Max-Weight policy and a Whittle's Index policy, and show that they are guaranteed to be within a factor of two, four and eight, respectively, away from the minimum AoI possible. In contrast, simulation results show that Max-Weight outperforms the other policies, both in terms of AoI and throughput, in every network configuration simulated, and achieves near optimal performance.

I. INTRODUCTION

The Age of Information (AoI) is a performance metric that measures the time that elapsed since the generation of the packet that was most recently delivered to the destination. This metric captures the freshness of the information from the *perspective of the destination*. Consider a cyber-physical system such as an automated industrial plant, a smart house or a modern car, where a number of sensors are transmitting time-sensitive information to a monitor over unreliable wireless channels. Each sensor samples information from a physical phenomena (e.g. pressure of the tire, quantity of fuel, proximity to obstacles and engine rotational speed) and transmits this data to the monitor. Ideally, the monitor receives fresh information about every physical phenomena continuously. However, due to limitations of the wireless channel, this is often impractical. In such cases, the system has to manage the use of the available channel resources in order to keep the monitor updated. In this paper, we develop three low-complexity transmission scheduling policies and analyze their performance in terms of the freshness of the information at the monitor, namely the Age of Information.

Let every packet be time-stamped with the time it was generated. Denote by $\tau_i[m]$ the time-stamp of the m th packet delivered by sensor i to the monitor. Assume that at time t , the

m th packet delivered by sensor i is the most recent. Then, the Age of Information associated with sensor i at time t is given by $h_i(t) = t - \tau_i[m]$. While the monitor does not receive new packets from sensor i , the value of $h_i(t)$ increases linearly with t , representing the information getting older. As soon as the monitor receives a new packet from sensor i , the corresponding time-stamp is instantaneously updated from $\tau_i[m]$ to $\tau_i[m+1]$, reducing the value of $h_i(t)$ by $\tau_i[m+1] - \tau_i[m]$. Notice that at the moment packet $(m+1)$ is delivered to the monitor, the value of $h_i(t)$ matches the delay of the packet. This makes sense because, at that moment, the information at the monitor is as old as the information contained in packet $(m+1)$. It follows naturally that *a good AoI performance is achieved when packets with low delay are delivered regularly*.

In order to provide good AoI performance, the scheduling policy must control how the channel resources are allocated to the different sensors in the network. Depending on the channel conditions and network configuration, this can mean that some sensors get to transmit repeatedly, while other sensors less often. The frequency at which information is delivered to the monitor is of particular importance in sensor networks. Clearly, a sensor that measures the quantity of fuel requires a lower update frequency (i.e. throughput) than a sensor that is measuring the proximity to obstacles in order to avoid collisions. For capturing this attribute, we associate a minimum timely-throughput requirement with each sensor in the network. Hence, in addition to providing good AoI performance, the scheduling policies should also fulfill timely-throughput constraints from the individual sensors.

A framework for modeling wireless networks with timely-throughput requirements was proposed in [1] together with two debt-based scheduling policies that fulfill any feasible requirements. Generalizations of this model to different network configurations were proposed in [2]–[4]. Scheduling policies that maximize throughput and also provide service regularity in wireless networks were studied in [5] and [6]. The problem of minimizing AoI was introduced in [7]. In [7]–[10], different queueing systems are analyzed and the optimal server utilization with respect to AoI is found. In [11]–[13], the authors optimize the process of generating information updates in order to minimize AoI. The design of scheduling policies based on AoI is considered in [14]–[20].

An important observation is that high throughput does not guarantee low AoI. Consider an M/M/1 queue with high arrival rate and low service rate. In this system, the queue is often filled, resulting in high throughput and high packet

This work was supported by NSF Grants AST-1547331, CNS-1713725, and CNS-1701964, and by Army Research Office (ARO) grant number W911NF-17-1-0508.

delay. This high delay means that packets being served contain outdated information. Hence, despite the high throughput, the AoI may still be high. In this paper, we develop policies that minimize AoI subject to minimum throughput requirements, where timely-throughput is modeled as in [1]. To the best of our knowledge, this is the first work to consider AoI-based policies that provably satisfy throughput constraints of multiple destinations simultaneously.

The remainder of this paper is outlined as follows. In Sec. II, the network model and performance metrics are formally presented. Then, in Sec. III, three low-complexity scheduling policies are proposed and analyzed. In Sec. IV, those policies are simulated and compared to the state-of-the-art in the literature. The paper is concluded in Sec. V.

II. SYSTEM MODEL

Consider a single-hop wireless network with a Base Station (BS) receiving time-sensitive information from M nodes. Let the time be slotted, with slot index $k \in \{1, 2, \dots, K\}$, and consider a wireless channel that allows at most one packet transmission per slot. In each slot k , the BS either idles or selects a node $i \in \{1, 2, \dots, M\}$ for transmission. Let $u_i(k)$ be the indicator function that is equal to 1 when the BS selects node i during slot k , and $u_i(k) = 0$ otherwise. When $u_i(k) = 1$, node i samples fresh information, generates a new packet and sends this packet over the wireless channel. The packet from node i is successfully received by the BS with probability $p_i \in (0, 1]$ and a transmission error occurs with probability $1 - p_i$. The probability p_i does not change with time, but may differ between nodes.

The transmission scheduling policy controls the decision of the BS in each slot k , which is represented by the set of values $\{u_i(k)\}_{i=1}^M$. The *interference constraint* associated with the wireless channel imposes that

$$\sum_{i=1}^M u_i(k) \leq 1, \forall k \in \{1, \dots, K\}, \quad (1)$$

meaning that at any given slot k , the scheduling policy can select at most one node for transmission. Let $d_i(k)$ be the random variable that indicates when a packet from node i is *delivered* to the BS. If node i transmits a packet during slot k , i.e. $u_i(k) = 1$, then $d_i(k) = 1$ with probability p_i and $d_i(k) = 0$ with probability $1 - p_i$. On the other hand, if node i does not transmit, i.e. $u_i(k) = 0$, then $d_i(k) = 0$ with probability one. It follows that $\mathbb{E}[d_i(k) | u_i(k)] = p_i u_i(k)$ and, applying the law of iterated expectations

$$\mathbb{E}[d_i(k)] = p_i \mathbb{E}[u_i(k)]. \quad (2)$$

In this paper, we consider non-anticipative scheduling policies, i.e. policies that do not use future knowledge in making decisions. Denote by Π the class of non-anticipative policies and let $\pi \in \Pi$ be an arbitrary admissible policy. *Our goal is to design low-complexity scheduling policies that belong to Π , provide close to optimal AoI performance and, at the same time, guarantee a minimum throughput level for each*

individual destination. Next, we formally introduce both performance metrics, throughput and AoI, and define a measure for “closeness to optimality”.

A. Minimum Throughput Requirement

Let q_i be a strictly positive real value that represents the minimum throughput requirement of node i . Using the random variable $d_i^\pi(k)$, we define the *long-term throughput* of node i when policy π is employed as

$$\hat{q}_i^\pi := \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{E}[d_i^\pi(k)]. \quad (3)$$

Then, we express the *minimum throughput constraint* of each individual node as

$$\hat{q}_i^\pi \geq q_i, \forall i \in \{1, \dots, M\}. \quad (4)$$

In this paper, we assume that $\{q_i\}_{i=1}^M$ is a *feasible* set of minimum throughput requirements, i.e. there exists a policy $\pi \in \Pi$ that satisfies all K interference constraints in (1) and all M throughput constraints in (4) simultaneously. As shown in [1, Lemma 5], the inequality

$$\sum_{i=1}^M \frac{q_i}{p_i} \leq 1, \quad (5)$$

is a necessary and sufficient condition for the feasibility of $\{q_i\}_{i=1}^M$. Throughout this paper, we assume that (5) is satisfied with strict inequality. Next, we present the AoI metric.

B. Age of Information

The Age of Information depicts how old the information is from the perspective of the BS. Let $h_i(k)$ be the positive integer that represents the AoI associated with node i at the beginning of slot k . If the BS does not receive a packet from node i during slot k , then $h_i(k+1) = h_i(k) + 1$, since the information at the BS is one slot older. In contrast, if the BS receives a packet from node i during slot k , then $h_i(k+1) = 1$, because the received packet was generated at the beginning of slot k . The evolution of $h_i(k)$ follows

$$h_i(k+1) = \begin{cases} 1 & , \text{ if } d_i(k) = 1; \\ h_i(k) + 1 & , \text{ otherwise.} \end{cases} \quad (6)$$

The average AoI of node i during the first K slots is captured by $\mathbb{E}[\sum_{k=1}^K h_i(k)]/K$, where the expectation is with respect to the randomness in the channel and the scheduling policy. For measuring the freshness of the information of the entire network when policy π is employed, we use the Expected Weighted Sum AoI

$$\mathbb{E}[J_K^\pi] = \frac{1}{KM} \mathbb{E} \left[\sum_{k=1}^K \sum_{i=1}^M \alpha_i h_i(k) \mid \vec{h}(1) \right], \quad (7)$$

where $\vec{h}(1) = [h_1(1), \dots, h_M(1)]^T$ is the vector of initial AoI in (6) and $\alpha_i > 0$ is the weight of node i . For simplicity, we assume that $h_i(1) = 1, \forall i$, and omit $\vec{h}(1)$ henceforth.

C. Optimization Problem

With the definitions of AoI and throughput, we present the optimization problem that is central to this paper.

AoI Optimization

$$\text{OPT}^* = \min_{\pi \in \Pi} \left\{ \lim_{K \rightarrow \infty} \frac{1}{KM} \mathbb{E} \left[\sum_{k=1}^K \sum_{i=1}^M \alpha_i h_i(k) \right] \right\} \quad (8a)$$

$$\text{s.t. } \hat{q}_i^\pi \geq q_i, \forall i; \quad (8b)$$

$$\sum_{i=1}^M u_i(k) \leq 1, \forall k. \quad (8c)$$

The minimum throughput constraints are depicted in (8b) and the interference constraints are in (8c). The scheduling policy that results from (8a)-(8c) is referred to as *AoI-optimal*.

For a given network setup (M, p_i, q_i, α_i) , let OPT^* be the Expected Weighted Sum AoI achieved by the AoI-optimal policy π^* . Similarly, let OPT_η be the AoI achieved by some policy $\eta \in \Pi$. The optimality ratio of η is given by

$$\psi^\eta = \frac{\text{OPT}_\eta}{\text{OPT}^*}, \quad (9)$$

and we say that policy η is ψ^η -optimal. Naturally, the closer ψ^η is to 1, the better is the AoI performance of policy η . The optimality ratio is used in the upcoming sections to compare the performance of different scheduling policies.

III. SCHEDULING POLICIES

In this section, we propose three low-complexity scheduling policies with strong AoI performances. The first two provably satisfy the throughput constraints for every feasible set $\{q_i\}_{i=1}^M$ and the third accounts for the throughput constraints, but provides no guarantee. To evaluate the AoI performance of each policy, we find their corresponding optimality ratio ψ^η . Moreover, in Sec. IV, we simulate and compare these policies to the state-of-the-art in the literature.

Prior to introducing the policies, we obtain a lower bound to the AoI optimization (8a)-(8c) which is used in the derivation of the optimality ratios ψ^η . Then, we present three scheduling policies: 1) Optimal Stationary Randomized policy; 2) Max-Weight policy; and 3) Whittle's Index policy. The first is obtained by solving the AoI optimization (8a)-(8c) over the class of Stationary Randomized Policies. The second and third policies are derived using Lyapunov Optimization [21] and the Restless Multi-Armed Bandit framework [22], respectively.

A. Lower Bound

In this section, we use a sample path argument to derive a lower bound to the AoI optimization (8a)-(8c).

Theorem 1. *The optimization problem in (10a)-(10c) provides a lower bound L_B to the AoI optimization (8a)-(8c), namely $L_B \leq \text{OPT}^*$ for every network setup (M, p_i, q_i, α_i) .*

Lower Bound

$$L_B = \min_{\pi \in \Pi} \left\{ \frac{1}{2M} \sum_{i=1}^M \alpha_i \left(\frac{1}{\hat{q}_i^\pi} + 1 \right) \right\} \quad (10a)$$

$$\text{s.t. } \hat{q}_i^\pi \geq q_i, \forall i; \quad (10b)$$

$$\sum_{i=1}^M u_i(k) \leq 1, \forall k. \quad (10c)$$

Proof. Consider a scheduling policy $\pi \in \Pi$ that satisfies all throughput and interference constraints running on a network for the time-horizon of K slots. Let Ω be the sample space associated with this network and let $\omega \in \Omega$ be a sample path. For a given sample path ω , the total number of packets delivered by node i during the K slots is denoted $D_i(K) = \sum_{k=1}^K d_i(k)$ and the inter-delivery time associated with each of those deliveries is denoted $I_i[m]$. In particular, let $I_i[m]$ be the number of slots between the $(m-1)$ th and m th packet deliveries from node i , $\forall m \in \{1, \dots, D_i(K)\}^1$. After the last packet delivery from node i , the number of remaining slots is R_i . Hence, the time-horizon can be written as

$$K = \sum_{m=1}^{D_i(K)} I_i[m] + R_i, \forall i \in \{1, 2, \dots, M\}. \quad (11)$$

According to the evolution of $h_i(k)$ in (6), the slot that follows the $(m-1)$ th packet delivery from node i has an AoI of $h_i(k) = 1$. Since the m th packet is delivered only after $I_i[m]$ slots, we know that $h_i(k)$ evolves as $\{1, 2, \dots, I_i[m]\}$. This pattern is repeated throughout the entire time-horizon, including the last R_i slots. As a result, the time-average Age of Information of node i can be expressed as

$$\begin{aligned} \frac{1}{K} \sum_{k=1}^K h_i(k) &= \frac{1}{K} \left[\sum_{m=1}^{D_i(K)} \frac{(I_i[m] + 1)I_i[m]}{2} + \frac{(R_i + 1)R_i}{2} \right] \\ &= \frac{1}{2} \left[\frac{D_i(K)}{K} \frac{1}{D_i(K)} \sum_{m=1}^{D_i(K)} I_i^2[m] + \frac{R_i^2}{K} + 1 \right], \end{aligned} \quad (12)$$

where the last equality uses (11) to replace the two linear terms by K .

Define the operator $\bar{\mathbb{M}}[\mathbf{x}]$ that computes the sample mean of any set $\mathbf{x}$. In particular, let the sample mean of $I_i[m]$ and $I_i^2[m]$ be

$$\bar{\mathbb{M}}[I_i] = \frac{1}{D_i(K)} \sum_{m=1}^{D_i(K)} I_i[m]; \quad (13)$$

$$\bar{\mathbb{M}}[I_i^2] = \frac{1}{D_i(K)} \sum_{m=1}^{D_i(K)} I_i^2[m]. \quad (14)$$

Substituting $\bar{\mathbb{M}}[I_i^2]$ into (12) and then applying Jensen's inequality, yields

$$\frac{1}{K} \sum_{k=1}^K h_i(k) \geq \frac{1}{2} \left(\frac{D_i(K)}{K} (\bar{\mathbb{M}}[I_i])^2 + \frac{R_i^2}{K} + 1 \right), \quad (15)$$

¹Naturally, $I_i[1]$ is the number of slots between the first slot and the first packet delivery from node i .

combining (11) into (13) and then substituting the result in (15), gives

$$\frac{1}{K} \sum_{k=1}^K h_i(k) \geq \frac{1}{2} \left(\frac{1}{K} \frac{(K - R_i)^2}{D_i(K)} + \frac{R_i^2}{K} + 1 \right). \quad (16)$$

By minimizing the LHS of (16) analytically with respect to the variable R_i , we have

$$\frac{1}{K} \sum_{k=1}^K h_i(k) \geq \frac{1}{2} \left(\frac{K}{D_i(K) + 1} + 1 \right). \quad (17)$$

Taking the expectation of (17) and applying Jensen's inequality, yields

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E}[h_i(k)] \geq \frac{1}{2} \left(\frac{1}{\mathbb{E}\left[\frac{D_i(K)}{K}\right] + \frac{1}{K}} + 1 \right). \quad (18)$$

Applying the limit $K \rightarrow \infty$ to (18) and using the definition of throughput in (3), gives

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{E}[h_i(k)] \geq \frac{1}{2} \left(\frac{1}{\hat{q}_i^\pi} + 1 \right). \quad (19)$$

Combining (19) and the objective function in (7), yields

$$\begin{aligned} \lim_{K \rightarrow \infty} \mathbb{E}[J_K^\pi] &= \lim_{K \rightarrow \infty} \frac{1}{M} \sum_{i=1}^M \frac{\alpha_i}{K} \sum_{k=1}^K \mathbb{E}[h_i(k)] \\ &\geq \frac{1}{2M} \sum_{i=1}^M \alpha_i \left(\frac{1}{\hat{q}_i^\pi} + 1 \right). \end{aligned} \quad (20)$$

Finally, substituting (20) into the AoI optimization (8a)-(8c) gives the Lower Bound (10a)-(10c). ■

To obtain the expression in (20), we applied Jensen's inequality twice and minimized (16) analytically with respect to R_i . Each of those steps could have led to a loose lower bound L_B . However, in the next section, we use this lower bound to obtain a tight optimality ratio, $\psi^R < 2$, for a Stationary Randomized policy. Moreover, we evaluate the tightness of L_B using numerical results in Sec. IV.

B. Optimal Stationary Randomized policy

Denote by Π_R the class of Stationary Randomized Policies and let $R \in \Pi_R$ be a scheduling policy that, in each slot k , selects node i with probability $\mu_i \in (0, 1]$ and idles with probability μ_{idle} . Each policy in Π_R is fully characterized by the set of scheduling probabilities $\{\mu_i\}_{i=1}^M$, where $\mu_i = \mathbb{E}[u_i(k)]$, $\forall i, \forall k$ and $\mu_{idle} = 1 - \sum_{i=1}^M \mu_i$. Next, we find the Optimal Stationary Randomized policy R^* that solves the AoI optimization (8a)-(8c) over the class $\Pi_R \subset \Pi$ and derive the associated optimality ratio ψ^R .

Proposition 2. Consider a policy $R \in \Pi_R$ with scheduling probabilities $\{\mu_i\}_{i=1}^M$. The long-term throughput and the expected time-average AoI of node i can be expressed as

$$\hat{q}_i^R = p_i \mu_i; \quad (21)$$

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{E}[h_i(k)] = \frac{1}{p_i \mu_i}. \quad (22)$$

Proof. In any given slot k , the BS receives a packet from node i if this node is scheduled and the corresponding packet transmission is successful. The probability of this event is $p_i \mu_i$. Moreover, the inter-delivery times $I_i[m]$ of node i are i.i.d. with $\mathbb{P}\{I_i[m] = n\} = p_i \mu_i (1 - p_i \mu_i)^{n-1}$, $\forall n \in \{1, 2, \dots\}$.

Clearly, under policy R , the sequence of packet deliveries is a renewal process. Thus, we can use renewal theory to derive (21) and (22). In particular, by the definition of long-term throughput (3) and the expression for the expected time-average AoI of node i , we have

$$\hat{q}_i^R = \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{E}[d_i(k)] \stackrel{(a)}{=} \frac{1}{\mathbb{E}[I_i[m]]} = p_i \mu_i; \quad (23)$$

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{E}[h_i(k)] \stackrel{(b)}{=} \frac{\mathbb{E}[I_i^2[m]]}{2\mathbb{E}[I_i[m]]} + \frac{1}{2} = \frac{1}{p_i \mu_i}. \quad (24)$$

where (a) follows from the elementary renewal theorem and (b) from its generalization for renewal-reward processes [23, Sec. 5.7]. ■

Substituting both expressions from Proposition 2 into the AoI optimization (8a)-(8c) gives the equivalent optimization problem over the class Π_R presented below.

Optimization over Randomized policies

$$\text{OPT}_{R^*} = \min_{R \in \Pi_R} \left\{ \frac{1}{M} \sum_{i=1}^M \frac{\alpha_i}{p_i \mu_i} \right\} \quad (25a)$$

$$\text{s.t. } p_i \mu_i \geq q_i, \forall i; \quad (25b)$$

$$\sum_{i=1}^M \mu_i \leq 1, \forall k. \quad (25c)$$

Notice that under the class Π_R , conditions (25c) and (8c) are equivalent. The Optimal Stationary Randomized policy R^* is characterized by the set $\{\mu_i^*\}_{i=1}^M$ that solves (25a)-(25c).

Theorem 3 (Optimality Ratio for R^*). The optimality ratio of R^* is such that $\psi^R < 2$, namely the Optimal Stationary Randomized policy is 2-optimal for every network setup.

Proof. Let $\hat{q}_i^L$ be the throughput associated with the policy that solves the Lower Bound (10a)-(10c). Consider the policy $R \in \Pi_R$ with long-term throughput $\hat{q}_i^R = p_i \mu_i = \hat{q}_i^L$ for each node i . Since $\hat{q}_i^R = \hat{q}_i^L$, it follows that R satisfies all throughput constraints. Comparing L_B in (10a) with the objective function associated with R , namely OPT_R , yields

$$\frac{\text{OPT}_R}{2} < L_B \rightarrow \psi^R = \frac{\text{OPT}_{R^*}}{\text{OPT}_R} \leq \frac{\text{OPT}_R}{L_B} < 2, \quad (26)$$

where OPT^* comes from (8a) and OPT_{R^*} from (25a). Recall that $L_B \leq \text{OPT}^* \leq \text{OPT}_{R^*} \leq \text{OPT}_R$. ■

Corollary 4. *The Optimal Stationary Randomized policy R^* is also the solution for the Lower Bound problem (10a)-(10c).*

Proof. Using the same argument as in the proof of Theorem 3, in particular $\hat{q}_i^R = p_i \mu_i = \hat{q}_i^L$, it follows that the scheduling policy that solves the Optimization over Randomized policies (25a)-(25c) also solves the Lower Bound (10a)-(10c). ■

Theorem 5 (Optimal Stationary Randomized policy). *The scheduling probabilities $\{\mu_i^*\}_{i=1}^M$ that result from Algorithm 1 are the unique solution to (25a)-(25c) and, thus, characterize the Optimal Stationary Randomized policy R^* .*

Algorithm 1 Unique solution to KKT Conditions

```

1:  $\gamma_i \leftarrow \alpha_i p_i / M q_i^2, \forall i \in \{1, 2, \dots, M\}$ 
2:  $\gamma \leftarrow \max_i \{\gamma_i\}$ 
3:  $\mu_i \leftarrow (q_i / p_i) \max\{1; \sqrt{\gamma_i / \gamma}\}, \forall i$ 
4:  $S \leftarrow \mu_1 + \mu_2 + \dots + \mu_M$ 
5: while  $S < 1$  do
6:   decrease  $\gamma$  slightly
7:   repeat steps 3 and 4 to update  $\mu_i$  and  $S$ 
8: end while
9:  $\mu_i^* = \mu_i, \forall i$ , and  $\gamma^* = \gamma$ 
10: return  $(\mu_1^*, \mu_2^*, \dots, \mu_M^*, \gamma^*)$ 
    
```

Proof. To find the set of scheduling probabilities $\{\mu_i^*\}_{i=1}^M$ that solve the optimization problem (25a)-(25c), we analyze the KKT Conditions. Let $\{\lambda_i\}_{i=1}^M$ be the KKT multipliers associated with the relaxation of (25b) and γ be the multiplier associated with the relaxation of (25c). Then, for $\lambda_i \geq 0, \forall i, \gamma \geq 0$ and $\mu_i \in (0, 1], \forall i$, we define

$$\mathcal{L}(\mu_i, \lambda_i, \gamma) = \frac{1}{M} \sum_{i=1}^M \frac{\alpha_i}{p_i \mu_i} + \sum_{i=1}^M \lambda_i (q_i - p_i \mu_i) + \gamma \left(\sum_{i=1}^M \mu_i - 1 \right), \quad (27)$$

and, otherwise, we define $\mathcal{L}(\mu_i, \lambda_i, \gamma) = +\infty$. Then, the KKT Conditions are

- (i) Stationarity: $\nabla_{\mu_i} \mathcal{L}(\mu_i, \lambda_i, \gamma) = 0$;
- (ii) Complementary Slackness: $\gamma(\sum_{i=1}^M \mu_i - 1) = 0$;
- (iii) Complementary Slackness: $\lambda_i(q_i - p_i \mu_i) = 0, \forall i$;
- (iv) Primal Feasibility: $p_i \mu_i \geq q_i, \forall i$, and $\sum_{i=1}^M \mu_i \leq 1$;
- (v) Dual Feasibility: $\lambda_i \geq 0, \forall i$, and $\gamma \geq 0$.

Since $\mathcal{L}(\mu_i, \lambda_i, \gamma)$ is a *convex function*, if there exists a vector $(\{\mu_i^*\}_{i=1}^M, \{\lambda_i^*\}_{i=1}^M, \gamma^*)$ that satisfies all KKT Conditions, then this vector is unique. Hence, the scheduling policy $R^* \in \Pi_R$ that optimizes (25a)-(25c) is also unique and is characterized by $\{\mu_i^*\}_{i=1}^M$. Next, we find the vector $(\{\mu_i^*\}_{i=1}^M, \{\lambda_i^*\}_{i=1}^M, \gamma^*)$.

To assess stationarity, $\nabla_{\mu_i} \mathcal{L}(\mu_i, \lambda_i, \gamma) = 0$, we calculate the partial derivative of $\mathcal{L}(\mu_i, \lambda_i, \gamma)$ with respect to μ_i . It follows from the derivative that

$$\frac{\alpha_i}{M p_i \mu_i^2} + \lambda_i p_i = \gamma, \forall i. \quad (28)$$

From complementary slackness, $\gamma(\sum_{i=1}^M \mu_i - 1) = 0$, we know that either $\gamma = 0$ or $\sum_{i=1}^M \mu_i = 1$. Equation (28) shows that the value of γ can only be zero if $\lambda_i = 0$ and $\mu_i \rightarrow \infty$, which violates $\mu_i \in (0, 1]$. Hence, we obtain

$$\gamma > 0 \quad \text{and} \quad \sum_{i=1}^M \mu_i = 1. \quad (29)$$

Notice that $\sum_{i=1}^M \mu_i = 1$ implies in $\mu_{idle} = 0$.

Based on dual feasibility, $\lambda_i \geq 0$, we can separate nodes $i \in \{1, \dots, M\}$ into two categories: nodes with $\lambda_i > 0$ and nodes with $\lambda_i = 0$.

Category 1) node i with $\lambda_i > 0$. It follows from complementary slackness, $\lambda_i(q_i - p_i \mu_i) = 0$, that

$$\mu_i = \frac{q_i}{p_i}. \quad (30)$$

Plugging this value of μ_i into (28) gives the inequality $\lambda_i p_i = \gamma - \gamma_i > 0$, where we define the constant

$$\gamma_i := \frac{\alpha_i p_i}{M q_i^2}. \quad (31)$$

Category 2) node i with $\lambda_i = 0$. It follows from (28) that

$$\gamma = \gamma_i \left(\frac{q_i}{p_i \mu_i} \right)^2 \rightarrow \mu_i = \frac{q_i}{p_i} \sqrt{\frac{\gamma_i}{\gamma}}. \quad (32)$$

In summary, for any fixed value of $\gamma > 0$, the scheduling probability of node i is

$$\mu_i = \frac{q_i}{p_i} \max \left\{ 1; \sqrt{\frac{\gamma_i}{\gamma}} \right\}. \quad (33)$$

Notice that for a decreasing value of γ , the probability μ_i remains fixed or increases. Our goal is to find the value of γ^* that gives $\{\mu_i^*\}_{i=1}^M$ satisfying the condition $\sum_{i=1}^M \mu_i^* = 1$.

Proposed algorithm to find γ^ :* start with $\gamma = \max\{\gamma_i\}$. Then, according to (33), all nodes have $\mu_i = q_i/p_i$ and, by the feasibility condition in (5), it follows that

$$\sum_{i=1}^M \mu_i = \sum_{i=1}^M \frac{q_i}{p_i} \leq 1. \quad (34)$$

Now, by gradually decreasing γ and adjusting $\{\mu_i\}_{i=1}^M$ according to (33), we can find the unique γ^* that fulfills $\sum_{i=1}^M \mu_i^* = 1$. The solution γ^* exists since $\gamma \rightarrow 0$ implies in $\sum_{i=1}^M \mu_i \rightarrow \infty$. The uniqueness of γ^* follows from the monotonicity of μ_i with respect to γ . This process is described in Algorithm 1 and illustrated in Fig. 1.

Algorithm 1 outputs the set of scheduling probabilities $\{\mu_i^*\}_{i=1}^M$ and the parameter γ^* . The set $\{\lambda_i^*\}_{i=1}^M$ is obtained using (28). Hence, the unique vector $(\{\mu_i^*\}_{i=1}^M, \{\lambda_i^*\}_{i=1}^M, \gamma^*)$ that solves the KKT Conditions is found. ■

In order to fulfill the throughput constraints (25b), every scheduling policy in Π_R must allocate at least $\mu_i \geq q_i/p_i$ to each node i . What differentiates policies in Π_R is how they distribute the remaining resources, $1 - \sum_{i=1}^M q_i/p_i$, between nodes. According to Algorithm 1, the Optimal Stationary Randomized policy R^* supplies additional resources, $\mu_i^* > q_i/p_i$,

to nodes with high value of γ_i , namely nodes with a high priority α_i or a low value of q_i/p_i . Notice that if a node with low q_i/p_i was given the minimum required amount of resources, it would rarely transmit and its AoI would be high. In contrast, policy R^* allocates the minimum required, $\mu_i^* = q_i/p_i$, to nodes with low priority α_i or high q_i/p_i .

The policies $R \in \Pi_R$ discussed in this section are as simple as possible. They select nodes randomly, according to fixed scheduling probabilities $\{\mu_i\}_{i=1}^M$ calculated offline by Algorithm 1. Despite their simplicity, it was shown that R^* is 2-optimal regardless of the network setup (M, p_i, q_i, α_i) . In the following sections, we develop scheduling policies that take advantage of additional information, such as the current AoI of each node, for selecting nodes in an adaptive manner.

C. Max-Weight policy

Using techniques from Lyapunov Optimization, we derive the Max-Weight policy associated with the AoI optimization (8a)-(8c). Max-Weight is a scheduling policy designed to reduce the expected increase in the Lyapunov Function. The Lyapunov Function outputs a positive scalar that is large when the network is in *undesirable states*, namely when nodes have high AoI or less throughput than the minimum required q_i . Intuitively, the Max-Weight policy keeps the network in desirable states by controlling the growth of the Lyapunov Function. Prior to presenting the Max-Weight policy, we introduce the notions of throughput debt, network state, Lyapunov Function and Lyapunov Drift.

Let $x_i(k)$ be the throughput debt associated with node i at the beginning of slot k . The throughput debt evolves as

$$x_i(k+1) = kq_i - \sum_{t=1}^k d_i(t). \quad (35)$$

The value of kq_i can be interpreted as the minimum number of packets that node i should have delivered by slot $k+1$ and $\sum_{t=1}^k d_i(t)$ is the total number of packets actually delivered in the same interval. Define the operator $(\cdot)^+ = \max\{(\cdot), 0\}$ that computes the positive part of a scalar. Then, the positive part of the throughput debt is given by $x_i^+(k) = \max\{x_i(k); 0\}$. A large debt $x_i^+(k)$ indicates to the scheduling policy $\pi \in \Pi$

that node i is lagging behind in terms of throughput. In fact, strong stability of the process $x_i^+(k)$, namely

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{E}[x_i^+(k)] < \infty, \quad (36)$$

is sufficient to establish that the minimum throughput constraint, $\hat{q}_i^\pi \geq q_i$, is satisfied [21, Theorem 2.8].

Denote by $S_k = (h_i(k), x_i^+(k))_{i=1}^M$ the network state at the beginning of slot k and define the Lyapunov Function by

$$L(S_k) := \frac{1}{2} \sum_{i=1}^M \left(\alpha_i h_i^2(k) + V [x_i^+(k)]^2 \right), \quad (37)$$

where V is a positive real value that depicts the importance of the throughput constraints. Observe that $L(S_k)$ is large when nodes have high AoI or high throughput debt. To measure the expected change in the Lyapunov Function from one slot to the next, we define the Lyapunov Drift

$$\Delta(S_k) := \mathbb{E} \{ L(S_{k+1}) - L(S_k) | S_k \}. \quad (38)$$

The Max-Weight policy is designed to keep $L(S_k)$ small by reducing $\Delta(S_k)$ in every slot k . Next, we present an upper bound on $\Delta(S_k)$ that can be readily used to design the Max-Weight policy. The derivation of this upper bound is centered around the evolution of $h_i(k)$ in (6) and the evolution of $x_i^+(k)$ in (35). The complete proof can be found in [24] and the upper bound follows

$$\Delta(S_k) \leq - \sum_{i=1}^M \mathbb{E} \{ u_i(k) | S_k \} W_i(k) + B(k), \quad (39)$$

where $W_i(k)$ and $B(k)$ are given by

$$W_i(k) = \frac{\alpha_i p_i}{2} h_i(k) [h_i(k) + 2] + V p_i x_i^+(k); \quad (40)$$

$$B(k) = \sum_{i=1}^M \left\{ \alpha_i \left[h_i(k) + \frac{1}{2} \right] + V \left[x_i^+(k) q_i + \frac{1}{2} \right] \right\}. \quad (41)$$

Both $W_i(k)$ and $B(k)$ are fully characterized by the network state S_k and network setup (M, p_i, q_i, α_i) . Hence, both can be used by admissible policies for making scheduling decisions. However, notice that the term $B(k)$ in (39) is not affected by the choice of $u_i(k)$. Thus, for minimizing the upper bound in (39), the Max-Weight policy selects, in each slot k , the node with highest value of $W_i(k)$, with ties being broken arbitrarily. Denote the Max-Weight policy as MW .

Theorem 6. *The Max-Weight policy satisfies any feasible set of minimum throughput requirements $\{q_i\}_{i=1}^M$.*

Theorem 7 (Optimality Ratio for MW). *For any given network setup (M, p_i, q_i, α_i) , the optimality ratio of MW is such that*

$$\psi^{MW} \leq 4 + \frac{1}{L_B} \left[V - \frac{2}{M} \sum_{i=1}^M \alpha_i \right]. \quad (42)$$

In particular, for every network with $V \leq 2 \sum_{i=1}^M \alpha_i / M$, the Max-Weight policy is 4-optimal.




Fig. 1. Illustration of Algorithm 1 in a network with 3 nodes. On the left, the initial configuration with $\gamma = \max\{\gamma_i\}$. On the right, the outcome γ^* implies that under policy R^* node 2 will operate with minimum required scheduling probability $\mu_2 = q_2/p_2$, while the other two nodes will operate with a scheduling probability that is larger than the minimum.

The proofs of Theorems 6 and 7 follow from the analysis of the expression in (39) and are provided in [24], where an alternative Max-Weight policy is shown to be 2-optimal.

Recall that the Optimal Stationary Randomized policy R^* selects nodes randomly, according to fixed scheduling probabilities $\{\mu_i^*\}_{i=1}^M$. In contrast, the Max-Weight policy MW uses feedback from the network, namely $h_i(k)$ and $x_i^+(k)$, to guide scheduling decisions. Despite the added complexity, we expect the feedback loop to improve the performance of MW . In fact, numerical results in Sec. IV demonstrate that MW outperforms R^* in every network configuration simulated. However, by comparing Theorems 3 and 7, it might seem that R^* yields a better performance than MW . This is because the analysis associated with MW is more challenging, leading to an optimality ratio ψ^{MW} that is less tight than ψ^R . Next, we develop an index policy based on Whittle's Index [22] that is surprisingly similar to MW and has a similar performance.

D. Whittle's Index policy

To find Whittle's Index, we transform the AoI optimization (8a)-(8c) into a relaxed Restless Multi-Armed Bandit (RMAB) problem. This is possible because every node in the network evolves as a restless bandit. To obtain the relaxed RMAB problem, we first substitute the K interference constraints in (8c) by the single time-averaged constraint

$$\frac{1}{K} \sum_{k=1}^K \sum_{i=1}^M \mathbb{E}[u_i(k)] \leq 1. \quad (43)$$

Next, we relax this time-averaged constraint, by placing (43) into the objective function (8a) together with the associated Lagrange Multiplier $C \geq 0$. The resulting optimization problem is called relaxed RMAB and its solution lays the foundation for the design of Whittle's Index. A detailed description of this method can be found in [22], [25], [26].

One of the challenges associated with this method is that Whittle's Index is only defined for problems that are *indexable*. Unfortunately, it can be shown that due to the throughput constraints, $\hat{q}_i^\pi \geq q_i$, the relaxed RMAB resulting from the transformation of the AoI optimization is not indexable. To overcome this, we relax the throughput constraints (8b), placing them into the objective function of (8a)-(8c) as follows

Relaxed AoI Optimization

$$\widetilde{\text{OPT}}^* = \min_{\pi \in \Pi} \left\{ \lim_{K \rightarrow \infty} \frac{1}{KM} \sum_{k=1}^K \sum_{i=1}^M \left[\alpha_i \mathbb{E}[h_i(k)] + \theta_i \left(\frac{q_i}{p_i} - \mathbb{E}[u_i(k)] \right) \right] \right\} \quad (44a)$$

$$\text{s.t. } \theta_i \geq 0, \forall i; \quad (44b)$$

$$\sum_{i=1}^M u_i(k) \leq 1, \forall k. \quad (44c)$$

Each Lagrange Multiplier θ_i is associated with a relaxation of $\hat{q}_i^\pi \geq q_i$. These multipliers are called *throughput incentives* for they represent the penalty incurred by scheduling policies that deviate from the corresponding throughput constraint.

Applying the transformation described at the beginning of this section to the relaxed AoI optimization (44a)-(44c) yields

Doubly relaxed RMAB

$$\widetilde{\text{OPT}}_D = \min_{\pi \in \Pi} \left\{ \lim_{K \rightarrow \infty} \frac{1}{KM} \sum_{k=1}^K \sum_{i=1}^M \left[\alpha_i \mathbb{E}[h_i(k)] + (C - \theta_i) \mathbb{E}[u_i(k)] - \frac{C}{M} + \frac{\theta_i q_i}{p_i} \right] \right\} \quad (45a)$$

$$\text{s.t. } \theta_i \geq 0, \forall i; \quad (45b)$$

$$C \geq 0. \quad (45c)$$

Next, we solve the doubly relaxed RMAB, establish that the relaxed AoI optimization is indexable and obtain a closed-form expression for the Whittle's Index.

The doubly relaxed RMAB is separable and thus can be solved for each individual node. Observe that a scheduling policy running on a network with a single node i can only choose between selecting node i for transmission or idling during slot k . The scheduling policy that optimizes (45a)-(45c) for a given node i is characterized next.

Proposition 8 (Threshold policy). *Consider the doubly relaxed RMAB problem (45a)-(45c) associated with a single node i . The optimal scheduling policy is a Threshold policy that, in each slot k , selects node i when $h_i(k) \geq H_i$ and idles when $1 \leq h_i(k) < H_i$. For positive fixed values of C and θ_i , if $C > \theta_i$, the expression for the threshold is*

$$H_i = \left\lfloor \frac{3}{2} - \frac{1}{p_i} + \sqrt{\left(\frac{1}{p_i} - \frac{1}{2} \right)^2 + \frac{2(C - \theta_i)}{p_i \alpha_i}} \right\rfloor. \quad (46)$$

Otherwise, if $C \leq \theta_i$, the threshold is $H_i = 1$.

Proposition 8 follows from [15, Proposition 4]. Next, we define the condition for indexability and establish that the relaxed AoI optimization is indexable. For a given value of C , let $\mathcal{I}_i(C) = \{h_i(k) \in \mathbb{N} | h_i(k) < H_i\}$ be the set of states $h_i(k)$ in which the Threshold policy idles. The doubly relaxed RMAB associated with node i is indexable if the set $\mathcal{I}_i(C)$ increases monotonically from $\emptyset$ to $\mathbb{N}$, as the value of C increases from 0 to $+\infty$. Furthermore, the relaxed AoI optimization is indexable if this condition holds for all nodes. The condition on $\mathcal{I}_i(C)$ follows directly from Proposition 8 and is true for all nodes i . Thus, we establish that the relaxed AoI optimization (44a)-(44c) is indexable.

Given indexability, we define Whittle's Index. Let $C_i(h_i(k))$ be the Whittle's Index associated with node i in state $h_i(k)$. By definition, $C_i(h_i(k))$ is the infimum value of C that makes both scheduling decisions (transmit or idle) equally desirable to the Threshold policy while in state $h_i(k)$. The scheduling decisions are equally desirable when the multiplier C is such that $H_i = h_i(k) + 1$. Using (46) to solve this equation for the value of C gives the following expression for the Index

$$C_i(h_i(k)) = \frac{\alpha_i p_i}{2} h_i(k) \left[h_i(k) + \frac{2}{p_i} - 1 \right] + \theta_i. \quad (47)$$

After establishing indexability and obtaining the expression for $C_i(h_i(k))$, we define Whittle's Index policy. The Whittle's Index policy selects, in each slot k , the node with highest value of $C_i(h_i(k))$, with ties being broken arbitrarily. Denote the Whittle's Index policy as WI .

Theorem 9 (Optimality Ratio for WI). *For any given network setup (M, p_i, q_i, α_i) , the optimality ratio of WI is such that*

$$\psi^{WI} \leq 8 + \frac{1}{L_B} \left[\frac{1}{M} \sum_{i=1}^M \theta_i - \frac{7}{2M} \sum_{i=1}^M \alpha_i \right]. \quad (48)$$

In particular, for every network with $\sum_{i=1}^M \theta_i \leq 7 \sum_{i=1}^M \alpha_i/2$, the Whittle's Index policy is 8-optimal.

The proof of Theorem 9 is provided in [24]. The arguments used for deriving ψ^{WI} are analogous to the ones for deriving ψ^{MW} in Theorem 7. Those similarities come from the fact that policies MW and WI are almost identical. Comparing the expressions for $W_i(k)$ and $C_i(h_i(k))$, in (40) and (47), respectively, we can see that both have the term $\alpha_i p_i h_i^2(k)/2$ and both have an isolated throughput term: $W_i(k)$ has $V p_i x_i^+(k)$ and $C_i(h_i(k))$ has θ_i . Naturally, we expect the performance of both policies to be similar in terms of AoI. The key difference between MW and WI lies in the throughput term. While the term $V p_i x_i^+(k)$ guarantees that MW satisfies the throughput constraint, $\hat{q}_i^\pi \geq q_i$, the positive scalar θ_i represents an incentive for WI to comply with the constraint, but provides no guarantee. The benefit of using a fixed θ_i is that there is no need to keep track of $x_i^+(k)$ for each node and at every slot k .

The results in this section hold for any given set of positive throughput incentives $\{\theta_i\}_{i=1}^M$. Next, we propose an algorithm that finds the values of θ_i which maximize a lower bound on the Lagrange Dual problem associated with the relaxed AoI optimization (44a)-(44c). Observe that $\widetilde{\text{OPT}}_D$ in (45a) is the Lagrange Dual function associated with (44a)-(44c). Thus, we can define the Lagrange Dual problem as $\max_{C, \theta_i} \{\widetilde{\text{OPT}}_D\}$ subject to $C \geq 0$ and $\theta_i \geq 0, \forall i$. Since this dual problem is challenging to address, we consider a lower bound:

$$\max_{C, \chi_i} \{\tilde{\mathcal{L}}(C, \chi_i)\} \leq \max_{C, \theta_i} \{\widetilde{\text{OPT}}_D\} \leq \text{OPT}^*. \quad (49)$$

subject to $\chi_i = C - \theta_i$, $C \geq 0$ and $\theta_i \geq 0$ for all nodes i , where

$$\begin{aligned} \tilde{\mathcal{L}}(C, \chi_i) = & \frac{1}{M} \sum_{i=1}^M \frac{\alpha_i}{p_i} - \frac{C}{M} \left[1 - \sum_{i=1}^M \frac{q_i}{p_i} \right] + \\ & + \sum_{i=1}^M \frac{\alpha_i}{M} \left[\sqrt{\frac{2\chi_i}{\alpha_i p_i} + \left[\frac{1}{p_i} - \frac{1}{2} \right]^2} - \frac{\chi_i q_i}{\alpha_i p_i} - \frac{1}{p_i} - \frac{1}{2} \right]. \end{aligned} \quad (50)$$

The throughput incentives θ_i that result from the maximization of $\tilde{\mathcal{L}}(C, \chi_i)$ are given by Algorithm 2. They are used in the next section to simulate the Whittle's Index policy. Simulation results show that the values of $\{\theta_i^*\}_{i=1}^M$ from Algorithm 2 reduce the throughput debt when compared to $\theta_i = 0$.

Algorithm 2 Throughput Incentives

```

1:  $\chi_i \leftarrow \alpha_i p_i [(1/q_i)^2 - (1/p_i - 1/2)^2]/2, \forall i$ 
2:  $C \leftarrow \max_i \{\chi_i\}$ 
3:  $\phi_i^{-1} \leftarrow p_i \sqrt{2 \min\{C; \chi_i\} / (\alpha_i p_i) + (1/p_i - 1/2)^2}, \forall i$ 
4:  $S \leftarrow \phi_1 + \phi_2 + \dots + \phi_M$ 
5: while  $S < 1$  do
6:   decrease  $C$  slightly
7:   repeat steps 3 and 4 to update  $\phi_i$  and  $S$ 
8: end while
9:  $C^* = C$  and  $\chi_i^* = \min\{C^*; \chi_i\}$  and  $\theta_i^* = C^* - \chi_i^*, \forall i$ 
10: return  $(\theta_1^*, \theta_2^*, \dots, \theta_M^*)$ 
    
```

IV. SIMULATION RESULTS

In this section, we simulate five transmission scheduling policies: 1) Optimal Randomized, R^* ; 2) Max-Weight², MW ; 3) Whittle's Index, WI ; 4) Largest Weighted-Debt First, LD ; and 5) Whittle's Index without throughput constraints, WP . The first three policies are developed in Sec. III and the last two are proposed in [1] and [15], respectively. Policy LD selects, in each slot k , the node with highest value of $x_i(k)/p_i$, where $x_i(k)$ is the throughput debt (35). It was shown in [1] that LD satisfies any set of feasible throughput requirements $\{q_i\}_{i=1}^M$. Notice that LD does not account for AoI. Policy WP was proposed in [15] for minimizing the AoI in broadcast wireless networks. It is analogous to WI but with $\theta_i = 0, \forall i$ and it does not account for minimum throughput requirements.

We simulate a network with M nodes, each having different parameters. Node i has weight $\alpha_i = (M+1-i)/M$, channel reliability $p_i = i/M$ and minimum throughput requirement $q_i = \varepsilon p_i/M$, where $\varepsilon \in [0, 1)$ represents the hardness of satisfying the throughput constraints $\hat{q}_i^\pi \geq q_i$. The larger the value of ε , the more challenging are the constraints. Notice that $\varepsilon < 1$ is necessary for the feasibility of $\{q_i\}_{i=1}^M$. Each simulation runs for a total of $K = M \times 10^6$ slots.

Figs. 2 and 3 show simulation results of networks with different sizes, namely $M \in \{5, 10, \dots, 25, 30\}$, while Fig. 4 shows networks with varying throughput constraints, in particular $\varepsilon \in \{0.7, 0.75, \dots, 0.9, 0.95, 0.999\}$. Two performance metrics are used to evaluate scheduling policies. Figs. 2 and 4 measure the Expected Weighted Sum AoI, $\mathbb{E}[J_K^\pi]$, defined in (7) and compare it with the lower bound L_B in (10a). Fig. 3 measures the maximum normalized throughput debt, defined as $\max_i \{x_i^+(K+1)/Kq_i\}$. Each data point in Figs. 2, 3 and 4 is an average over the results of 10 simulations.

Our results clearly demonstrate the superior performance of the Max-Weight policy. Fig. 3 shows that, as expected, only WI and WP violate the throughput requirements. Nevertheless, by comparing WI and WP , it is evident that the incentives θ_i^* from Algorithm 2 reduced the throughput debt. Figs. 2 and 4 show that the performance of MW , WI and WP are comparable to the lower bound. Since the lower bound is only associated with policies that fulfill the throughput

²The Max-Weight policy is simulated with $V = M^2$.

requirements, we conclude that the performance of MW is close to optimal.

V. CONCLUDING REMARKS

In this paper, we considered a single-hop wireless network with a number of nodes transmitting time-sensitive information to a Base Station over unreliable channels. We addressed the problem of minimizing the Expected Weighted Sum AoI of the network while satisfying minimum throughput requirements from the individual nodes. Three low-complexity scheduling policies were developed: Optimal Stationary Randomized

policy, Max-Weight policy, and Whittle's Index policy. The performance of each policy was evaluated both analytically and through simulation. The Max-Weight policy demonstrated the best performance in terms of both AoI and throughput. Interesting extensions include consideration of unknown channel probabilities p_i and periodic generation of packets.

REFERENCES

- [1] I.-H. Hou, V. Borkar, and P. R. Kumar, "A theory of qos for wireless," in *IEEE INFOCOM*, 2009.
- [2] I.-H. Hou and P. R. Kumar, "Scheduling heterogeneous real-time traffic over fading wireless channels," in *IEEE INFOCOM*, 2010.
- [3] —, "Admission control and scheduling for qos guarantees for variable-bit-rate applications on wireless channels," in *ACM MobiHoc*, 2009.
- [4] K. S. Kim, C.-P. Li, I. Kadota, and E. Modiano, "Optimal scheduling of real-time traffic in wireless networks with delayed feedback," in *IEEE Allerton*, 2015.
- [5] R. Li, A. Eryilmaz, and B. Li, "Throughput-optimal wireless scheduling with regulated inter-service times," in *IEEE INFOCOM*, 2013.
- [6] R. Singh, X. Guo, and P. Kumar, "Index policies for optimal mean-variance trade-off of inter-delivery times in real-time sensor networks," in *IEEE INFOCOM*, 2015.
- [7] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?," in *IEEE INFOCOM*, 2012.
- [8] C. Kam, S. Kompella, G. D. Nguyen, and A. Ephremides, "Effect of message transmission path diversity on status age," *IEEE Trans. on Info. Theory*, 2015.
- [9] M. Costa, M. Codreanu, and A. Ephremides, "On the age of information in status update systems with packet management," *IEEE Trans. on Info. Theory*, 2016.
- [10] L. Huang and E. Modiano, "Optimizing age-of-information in a multi-class queueing system," in *IEEE ISIT*, 2015.
- [11] B. T. Bacinoglu, E. T. Ceran, and E. Uysal-Biyikoglu, "Age of information under energy replenishment constraints," in *ITA*, 2015.
- [12] B. T. Bacinoglu and E. Uysal-Biyikoglu, "Scheduling status updates to minimize age of information with an energy harvesting sensor," in *IEEE ISIT*, 2017.
- [13] Y. Sun, E. Uysal-Biyikoglu, R. Yates, C. E. Koksal, and N. Shroff, "Update or wait: How to keep your data fresh," in *IEEE INFOCOM*, 2016.
- [14] Q. He, D. Yuan, and A. Ephremides, "Optimizing freshness of information: On minimum age link scheduling in wireless systems," in *IEEE WiOpt*, 2016.
- [15] I. Kadota, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Minimizing the age of information in broadcast wireless networks," in *IEEE Allerton*, 2016.
- [16] C. Joo and A. Eryilmaz, "Wireless scheduling for information freshness and synchrony: Drift-based design and heavy-traffic analysis," in *IEEE WiOpt*, 2017.
- [17] Y.-P. Hsu, E. Modiano, and L. Duan, "Age of information: Design and analysis of optimal scheduling algorithms," in *IEEE ISIT*, 2017.
- [18] R. Yates, P. Ciblat, A. Yener, and M. Wigger, "Age-optimal constrained cache updating," in *IEEE ISIT*, 2017.
- [19] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Optimizing data freshness, throughput, and delay in multi-server information-update systems," in *IEEE ISIT*, 2016.
- [20] S. Kaul and R. Yates, "Status updates over unreliable multiaccess channels," in *IEEE ISIT*, 2017.
- [21] M. J. Neely, *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan and Claypool Publishers, 2010.
- [22] P. Whittle, "Restless bandits: Activity allocation in a changing world," *Journal of Applied Probability*, 1988.
- [23] R. Gallager, *Stochastic Processes: Theory for Applications*. Cambridge University Press, 2013.
- [24] "Optimizing age of information in wireless networks with throughput constraints," Online: <http://www.igorkadota.com/publications.html>.
- [25] J. Gittins, K. Glazebrook, and R. Weber, *Multi-armed Bandit Allocation Indices*, 2nd ed. Wiley, Mar. 2011.
- [26] R. R. Weber and G. Weiss, "On an index policy for restless bandits," *Journal of Applied Probability*, 1990.




Fig. 2. Simulation of a network with fixed $\varepsilon = 0.9$ and varying size M .




Fig. 3. Simulation of a network with fixed $\varepsilon = 0.9$ and varying size M .




Fig. 4. Simulation of a network with size $M = 30$ and varying hardness ε .