Minimizing the Age of Information in Broadcast Wireless Networks

Igor Kadota, Elif Uysal-Biyikoglu, Rahul Singh and Eytan Modiano

Abstract— We consider a wireless broadcast network with a base station sending time-sensitive information to a number of clients. The Age of Information (AoI), namely the amount of time that elapsed since the most recently delivered packet was generated, captures the freshness of the information. We formulate a discrete-time decision problem to find a scheduling policy that minimizes the expected weighted sum AoI of the clients in the network. To the best of our knowledge, this is the first work to provide a scheduling policy that optimizes AoI in a wireless network with unreliable channels.

The results are twofold: first, we show that a Greedy Policy, which transmits the packet with highest current age, is optimal for the case of symmetric networks. Then, for the general network case, we establish that the problem is indexable and obtain the Whittle Index in closed-form. Numerical results are presented to demonstrate the performance of the policies.

I. INTRODUCTION

Age of Information (AoI) has been receiving increasing attention in the literature [1]–[6], particularly for applications that generate time-sensitive data such as position, command and control, or sensor data. An interesting feature of this performance metric is that it captures the freshness of the information *from the perspective of the destination*, in contrast to the long-established packet delay, that represents the freshness of the information with respect to individual packets. In particular, AoI measures the time elapsed since the generation of the packet that was most recently delivered to the destination, while packet delay measures the time elapsed from the generation of a packet to its delivery.

The two parameters that influence AoI are packet delay and packet interdelivery time. In general, controlling only one is insufficient for achieving good AoI performance. For example, consider an M/M/1 queue with a low arrival rate and a high service rate. In this setting, the queue is often empty, resulting in low packet delay. Nonetheless, the AoI may still be high, since infrequent packet arrivals may result in outdated information at the destination. Table I provides a numerical example of an M/M/1 queue with fixed service rate $\mu = 1$ and a variable arrival rate λ. The first and third rows represent a queue in which high values of expected interdelivery time and expcted packet delay, respectively, induce a high average AoI. The second row shows the system at the point of minimum average AoI [1].

A good AoI performance is achieved when packets with low delay are delivered regularly. It is important to emphasize the difference between delivering packets regularly and

providing a minimum throughput. Fig. 1 illustrates the case of two sequences of packet deliveries that have the same throughput but different delivery regularity. In general, a minimum throughput requirement can be fulfilled even if long periods with no delivery occur, as long as those are balanced by periods of consecutive deliveries.

Minimizing the AoI is particularly challenging in wireless networks with unreliable channels due to transmission errors that result in packet losses. In this work, we consider the problem of optimizing link scheduling decisions to minimize the expected weighted sum AoI of the clients in the network. To the best of our knowledge, this is the first work to provide a transmission scheduling policy that optimizes AoI in a wireless network with unreliable channels.

The problem of optimizing scheduling decisions in broadcast wireless networks with respect to throughput and delivery times has been studied extensively in the literature. Throughput maximization of traffic with strict packet delay constraints has been addressed in [7]–[10]. Interdelivery time is considered in [11], [12] as a measure of service regularity. Age of Information has been considered recently in [6].

In recent works, the problem of minimizing the AoI has been explored using different approaches. Queueing Theory is used in [1]–[3] for finding the optimal server utilization with respect to AoI. The authors in [4], [5] consider the problem of optimizing the times in which packets are generated at the source. Link scheduling optimization is considered in [6], where the complexity of the problem is established and insights into the structure of the problem are provided. In contrast, our work focuses on characterizing the scheduling policy that optimizes the AoI.

The remainder of this paper is outlined as follows. In the next section, the network model is presented. In Sec. III, we find the optimal scheduling policy for the case of symmetric networks. In Sec. IV, we consider the general network case and obtain a closed-form expression for the Whittle Index. In Sec. V, the scheduling policies are evaluated using simulations. The paper is concluded in Sec. VI.

TABLE I

Expected delay, expected interdelivery time and average AoI of a M/M/1 queue with $\mu = 1$ and variable λ.

λ (pkt/sec)	$\mathbb{E}[\text{delay}]$ (sec)	$\mathbb{E}[\text{interdelivery}]$ (sec)	Average AoI (sec)
0.01	1.01	**100.00**	**101.00**
0.53	2.13	1.89	3.48
0.99	**100.00**	1.01	**100.02**

844

II. SYSTEM MODEL

Consider a wireless single-hop network with a base station (BS) sending time-sensitive information to M clients. Let the time be slotted, with T consecutive slots forming a frame. At the beginning of every frame, the BS generates one packet per client $i \in \{1, 2, \cdots, M\}$. Those new packets replace any undelivered packets from the previous frame. Denote the frame index by the positive integer k. Each packet is associated with a single frame and client, thus, it can be unequivocally identified by the tuple (k, i).

Let $n \in \{1, \cdots, T\}$ be the index of the slot within a frame. A slot is identified by the tuple (k, n). In a slot, the BS transmits a packet to a selected client i over the wireless channel. The packet is successfully delivered to client i with probability $p_i \in (0, 1]$ and a transmission error occurs with probability $1 - p_i$. The probability of successful transmission p_i is fixed in time, but may differ across clients. The client sends a feedback signal to the BS after every transmission. The ACK/NACK reaches the BS instantaneously and without errors.

The scheduling policies considered in this paper are non-anticipative, i.e. policies that do not use future knowledge in selecting clients. Let Π be the class of non-anticipative policies and $\pi \in \Pi$ be an arbitrary admissible policy. In each slot, policy π can either idle or schedule the transmission of an undelivered packet to client i. Our goal is to characterize the scheduling policy π^* that minimizes the expected weighted sum AoI of the clients in the network. Next, we discuss this performance metric.

A. Age of Information Formulation

Prior to introducing the expected weighted sum AoI, we characterize the AoI in the context of our system model. Let AoI_i be the positive real number that represents the Age of Information of client i. The AoI_i increases linearly in time when there is no delivery of packets to client i. At the end of the frame in which a delivery occurs, the AoI_i is updated to T. In Fig. 2, the evolution of AoI_i is illustrated for a given sample sequence of deliveries to client i.

In Fig. 3, the AoI_i is shown in detail. Let $\hat{s}_k$ denote the set of clients that successfully received packets during frame k and let the positive integer $h_{k,i}$ represent the number of frames since the last delivery to client i. At the beginning of

frame $k + 1$, the value of $h_{k,i}$ is updated as follows:

$$h_{k+1,i} = \begin{cases} h_{k,i} + 1 & , \text{ if } i \notin \hat{s}_k ; \\ 1 & , \text{ if } i \in \hat{s}_k . \end{cases} \quad (1)$$

As can be seen in Fig. 3, during frame k the area under the AoI_i curve can be divided into a triangle of area $T^2/2$ and a parallelogram of area $h_{k,i}T^2$. This area, averaged over time, captures the average Age of Information of client i. A network-wide metric for measuring the freshness of the information is the Expected Weighted Sum AoI:

$$\begin{aligned} \text{EWSAoI} &= \frac{1}{KT}\mathbb{E}\left\{\sum_{k=1}^{K}\sum_{i=1}^{M}\alpha_i\left(\frac{T^2}{2} + T^2 h_{k,i}\right)\Big|\vec{h}_1\right\} = \\ &= \frac{\sum_{i=1}^{M}\alpha_i T}{2} + \frac{T}{K}\mathbb{E}\left\{\sum_{k=1}^{K}\sum_{i=1}^{M}\alpha_i(h_{k,i})\Big|\vec{h}_1\right\}, (2) \end{aligned}$$

where α_i is the positive real value that denotes the client's weight and the vector $\vec{h}_1 = [h_{1,1}, \cdots, h_{1,M}]^T$ represents the initial values of the AoI in (1). For notational simplicity, we assume that $h_{1,i} = 1, \forall i$, and omit $\vec{h}_1$ hereafter. Manipulating the expression for EWSAoI gives us the objective function $J_K^{\pi^*}$:

$$J_K^{\pi^*} = \min_{\pi \in \Pi}\mathbb{E}\left\{\sum_{k=1}^{K}\sum_{i=1}^{M}\alpha_i h_{k,i}\right\}, \quad (3)$$

where (3) is obtained by subtracting the constant terms from (2) and multiplying the result by K/T. With the definitions of Age of Information[1] and objective function presented, in the following sections we discuss the scheduling policy that minimizes (3).

III. SYMMETRIC NETWORK

Consider a Greedy Policy that, in each slot (k, n), schedules a transmission to the client with an undelivered packet and highest value of $h_{k,i}$, with ties being broken in favor of

Sequence of deliveries to client i :

Fig. 2. On the top, five frames of a sample sequence of deliveries to client i. The upward arrows represent the times of packet deliveries. On the bottom, the evolution of the AoI_i.

[1]For ease of exposition, in this paper, the value of AoI_i is updated at the beginning of the frame that follows a successful transmission to client i, rather than immediately after the successful transmission. This update mechanism simplifies the problem while maintaining the features of interest. The solution to the more general problem will be provided in an extended version of this paper.
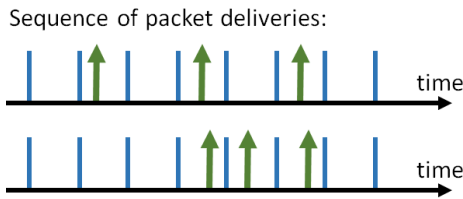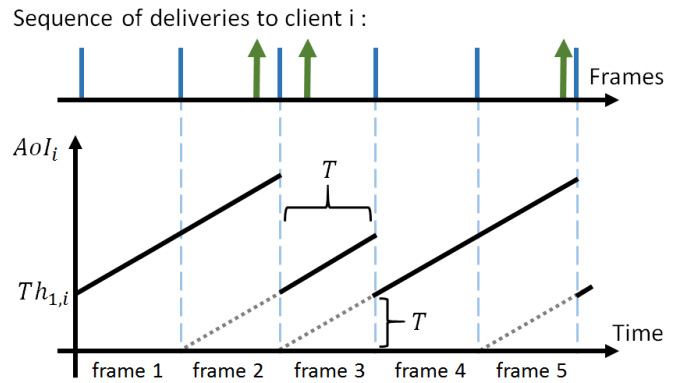
Sequence of packet deliveries:

Fig. 1. Two sample sequences of packet deliveries are represented by the green arrows. Both sequences have the same throughput, namely 3 packets over the interval, but different delivery regularity.

the client with lowest index i. Denote this policy as G. Notice that G is non-anticipative and work-conserving. Theorem 1 shows that this Greedy Policy minimizes the EWSAoI when the network is symmetric.

Theorem 1: (Optimality of Greedy for Symmetric Networks) Consider a network with all clients having equal channel reliability $p_i = p \in (0,1]$ and weight $\alpha_i = \alpha$. Among the class of admissible policies Π, the Greedy Policy attains the minimum time average sum AoI.

Proof: To show that the Greedy Policy minimizes the EWSAoI in (2), we utilize a stochastic dominance argument [13] to compare the evolution of $\vec{h}_k$ when G is employed to that when an arbitrary policy π is employed. For the sake of simplicity and without loss of optimality, in this proof we assume that π is work-conserving. Notice that for every non work-conserving policy, there is at least one work-conserving policy that is strictly dominant.

Let SH_k^π be the random variable that represents the sum of the elements of $\vec{h}_k$ when π is employed. Using this notation and the symmetry assumptions of Theorem 1, the objective function in (3) becomes

$$
\begin{aligned}
J_K^{\pi^*} &= \alpha \min_{\pi \in \Pi} \mathbb{E}\left\{ \sum_{k=1}^{K} \sum_{i=1}^{M} h_{k,i} \right\} = \\
&= \alpha \min_{\pi \in \Pi} \mathbb{E}\left\{ \sum_{k=1}^{K} SH_k^\pi \right\}.
\end{aligned}
\tag{4}
$$

For introducing the concept of stochastic dominance, denote the stochastic process associated with the sequence $\{SH_k^\pi\}_{k=1}^K$ as SH^π and its sample path as sh^π. Let $\mathbb{D}$ be the space of all sample paths sh^π. Define by $\mathscr{F}$ the set of measurable functions $f : \mathbb{D} \to \mathbb{R}^+$ such that $f(sh^G) \le f(sh^\pi)$ for every $sh^G, sh^\pi \in \mathbb{D}$ which satisfy $sh_k^G \le sh_k^\pi, \forall k$.

Definition 2: (Stochastic Dominance) We say that SH^G is stochastically smaller than SH^π and write $SH^G \le_{st} SH^\pi$ if $P\{f(SH^G) > z\} \le P\{f(SH^\pi) > z\}, \forall z \in \mathbb{R}, \forall f \in \mathscr{F}$.

Since $f(SH^\pi)$ is positive valued, $SH^G \le_{st} SH^\pi$ implies[2] $\mathbb{E}[f(SH^G)] \le \mathbb{E}[f(SH^\pi)], \forall f \in \mathscr{F}$. Knowing that one function that satisfies the conditions in $\mathscr{F}$ is $f(SH^\pi) = \sum_{k=1}^{K} SH_k^\pi$, it follows that if $SH^G \le_{st} SH^\pi, \forall \pi \in \Pi$, then $\mathbb{E}\{\sum_{k=1}^{K} SH_k^G\} \le \mathbb{E}\{\sum_{k=1}^{K} SH_k^\pi\}, \forall \pi \in \Pi$, which is our target expression in (4). Therefore, it follows that for establishing the optimality of G,
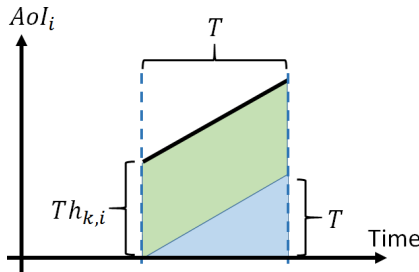

Fig. 3. Area under AoI_i during any frame k in terms of $h_{k,i}$ and T.

[2] Recall that for any positive valued X, it follows that $\mathbb{E}[X] = \int_{x=0}^{\infty}(1 - P\{X \le x\})dx = \int_{x=0}^{\infty} P\{X > x\}dx$.

it is sufficient to confirm that SH^G is stochastically smaller than $SH^\pi, \forall \pi \in \Pi$.

Stochastic dominance can be demonstrated using its definition directly. However, this is often complex for it involves comparing the probability distributions of SH^G and SH^π. Instead, we use the following result from [13], which is also used in works such as [7], [10], [14]: for verifying that $SH^G \le_{st} SH^\pi$, it is sufficient to show that there exist two stochastic processes $\widetilde{SH}^G$ and $\widetilde{SH}^\pi$ such that

(i) SH^π and $\widetilde{SH}^\pi$ have the same probability distribution;

(ii) $\widetilde{SH}^G$ and $\widetilde{SH}^\pi$ are on a common probability space;

(iii) SH^G and $\widetilde{SH}^G$ have the same probability distribution;

(iv) $\widetilde{SH}_k^G \le \widetilde{SH}_k^\pi$, with probability 1, $\forall k$.

This result allows us to establish stochastic dominance between SH^G and SH^π by properly designing the auxiliary processes $\widetilde{SH}^G$ and $\widetilde{SH}^\pi$. This design is achieved by utilizing Stochastic Coupling.

Prior to discussing stochastic coupling, we introduce the channel state. Let $E_i(k,n) \sim Ber(p)$ be the random variable that represents the channel state of client i during slot (k,n)

$$
E_i(k,n) = \begin{cases} 1, & \text{w.p. } p \quad [Channel\ ON]; \\ 0, & \text{w.p. } 1-p \quad [Channel\ OFF]. \end{cases}
\tag{5}
$$

The channel state of each client is independent of the channel state of other clients and of scheduling decisions. Notice that the BS has no knowledge of the channel state of the clients before transmissions.

Stochastic coupling is a method utilized for comparing stochastic processes by imposing a common underlying probability space. We use stochastic coupling to construct $\widetilde{SH}^\pi$ and $\widetilde{SH}^G$ based on SH^π and SH^G, respectively.

Let the process $\widetilde{SH}^\pi$ be identical to SH^π. Their (common) probability space is associated with the outcome of the transmission in each slot, i.e. it only depends on scheduling decisions and channel states. Now, let us construct $\widetilde{SH}^G$ on the same probability space as $\widetilde{SH}^\pi$. For that, we couple $\widetilde{SH}^G$ to $\widetilde{SH}^\pi$ by dynamically connecting the channel state of policy G to the channel state of π as follows. Suppose that in slot (k,n), policy π schedules client j while G schedules client i, then, for the duration of that slot, we assign $E_i(k,n) \leftarrow E_j(k,n)$. This iterative assignment imposes that, at every slot, the channel state of G is identical to the channel state of π. For example, if the channel associated with the client selected by policy π during slot (k,n) is ON, then, the channel state of G is also ON, regardless of the client selected by policy G. Notice that the assignment $E_i(k,n) \leftarrow E_j(k,n)$ is only possible because the channel state $E_i(k,n)$ is independent and identically distributed for all clients, which is the same reason for $\widetilde{SH}^G$ and SH^G having the same probability distribution.

Returning to our four conditions, it follows from the coupling method described above that (i), (ii) and (iii) are satisfied. Thus, the only condition that remains to be shown is

$$
\text{(iv) } \widetilde{SH}_k^G \le \widetilde{SH}_k^\pi, \text{with probability 1, } \forall k.
\tag{6}
$$

This condition is established by characterizing the evolution of $\widetilde{SH}_k^G$ and $\widetilde{SH}_k^\pi$ with k. The details are omitted due to length constraints, but the intuition behind (iv) is straightforward. Fig. 4 shows a numerical example of the evolution of $\widetilde{SH}_k^G$ and $\widetilde{SH}_k^\pi$. Consider a network with $M = 3$ clients, $T = 1$ slots in a frame, $K = 5$ frames, initialization vector $\vec{h}_1 = [4,3,1]^T$ and channel reliability $p \in (0,1]$. For the sequence of *coupled channel states* given in Fig. 4, the evolution of $\vec{h}_k$ when policies π and G are employed is displayed. Recall from (1) that the quantity $h_{k+1,i}$ is updated to 1 after a successful delivery to client i during frame k and is updated to $h_{k,i}+1$, otherwise.

As can be seen, during slots in which the channel is OFF, all policies yield the same result, namely $h_{k+1,i}$ is updated to $h_{k,i}+1$ for all clients i. However, during slots in which the channel is ON, the Greedy Policy achieves the lowest sum of the elements of $\vec{h}_k$ by selecting the client with highest value of $h_{k,i}$, i.e. by reducing the highest $h_{k,i}$ to 1. Therefore, it follows that employing the Greedy Policy in every slot yields the minimum value of $\widetilde{SH}_k^G$ for all k. With the last condition established, the proof is complete. ∎

We showed the optimality of Greedy for symmetric networks. For general networks, with clients possibly having different channel reliability p_i and weights α_i, scheduling decisions based exclusively on $\vec{h}_k$ may not be optimal. In the next section, we obtain the Whittle Index Policy associated with the general case of the AoI minimization problem in (3). As expected, the index of client i is a function of $h_{k,i}$, p_i and α_i.

IV. GENERAL NETWORK

One possible approach for finding a policy that minimizes the EWSAoI is to optimize the objective function in (3) using Dynamic Programming [15]. A negative aspect of this
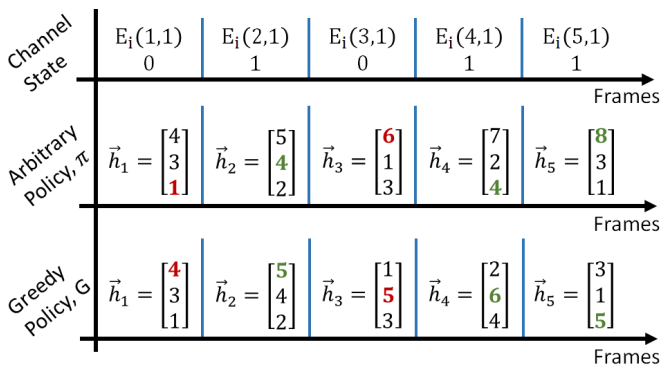

Fig. 4. Evolution of $\widetilde{SH}_k^\pi$ and $\widetilde{SH}_k^G$ for a network with $M = 3$, $T = 1$, $K = 5$, $\vec{h}_1 = [4,3,1]^T$ and unreliable channels. Recall that channels are ON when $E_i(k,n) = 1$ and OFF when $E_i(k,n) = 0$. Successful deliveries are represented in green and failed transmissions in red. On the top, channel states associated with the scheduling decisions of the arbitrary policy π. Notice that, due to coupling, policy G has the same channel states. On the middle, the evolution of $\vec{h}_k$ when policy π is employed. On the bottom, the evolution of $\vec{h}_k$ when policy G is employed. Comparing the sum of $\vec{h}_k$ over time for both policies, we have $\widetilde{SH}^\pi = \{8;11;10;13;12\}$ and $\widetilde{SH}^G = \{8;11;9;12;9\}$, and we see that $\widetilde{SH}_k^G \leq \widetilde{SH}_k^\pi, \forall k$.

approach is that evaluating the optimal policy for each state of the network can be computationally demanding, especially for networks with a large number of clients. To overcome this problem, known as the curse of dimensionality, we propose a simple Index Policy [16], also known as Whittle Index Policy.

For designing the Index Policy, the AoI minimization problem is transformed into a relaxed Restless Multi-Armed Bandit (RMAB) problem. First, we note that each client in the AoI problem can be seen as a restless bandit and thus the AoI problem can be posed as a RMAB problem. Then, we consider the relaxed version of the RMAB problem, called the Decoupled Model, in which clients are examined separately. The Decoupled Model associated with any client i adheres to the network model with $M = 1$, except for the addition of a *service charge*. The service charge is a fixed cost per transmission $C > 0$ that is incurred by the network every time the BS transmits a packet. Since the Decoupled Model considers only a single client, hereafter in this section, we omit the client index i.

The solution to the Decoupled Model lays the foundation for the design of the Index Policy. The Index Policy is a low-complexity heuristic that has been extensively used in the literature [10], [12], [17] and is known to have a strong performance in a range of applications [18], [19]. The challenge associated with this approach is that the Index Policy is only defined for problems that are *indexable*, a condition that is often difficult to establish. The Decoupled Model is formulated and solved next. Indexability and the Index Policy are discussed in Sec. IV-C. A detailed introduction to the Whittle Index Policy can be found in [16], [20].

A. Decoupled Model

In this section, the Decoupled Model is formulated as a Dynamic Program (DP). For presenting the cost-to-go function, which is central to the DP, we first introduce the state, control, transition and objective of the model. Then, using the expression of the cost-to-go, we establish in Proposition 3 a key property of the Decoupled Model which is used in the characterization of its optimal scheduling policy.

Consider the network model from Sec. II with a single client. Recall that at the beginning of every frame, the BS replaces any undelivered packet from the previous frame with a new packet. Let the set $s_{k,n}$ represent the delivery status of this new packet at the beginning of slot (k,n). If the packet has been successfully delivered to the client by the beginning of slot (k,n), then $s_{k,n} = 1$, and if the packet is still undelivered, $s_{k,n} = \emptyset$. The tuple $(s_{k,n}, h_k)$ depicts the system state, for it provides a complete characterization of the network at slot (k,n).

Denote by $u_{k,n}$ the scheduling decision (or control) in time slot (k,n). This quantity is equal to 1 if the BS transmits the packet in slot (k,n), and $u_{k,n} = 0$ otherwise. The BS can only transmit undelivered packets, i.e. if $s_{k,n} = 1$, then $u_{k,n} = 0$.

State transitions are different at frame boundaries and within frames. At the boundary between frames $k-1$ and

k, namely, in the transition from slot $(k-1,T)$ to slot $(k,1)$, each component of the system state $(s_{k,n},h_k)$ evolves in a distinct way. Since the BS generates a new packet at the beginning of slot $(k,1)$, we have that $s_{k,1} = \emptyset$. The evolution of h_k is divided into two cases: i) when the BS transmits the packet during slot $(k-1,T)$, i.e. $u_{k-1,T} = 1$, the value of h_k depends on the feedback signal, as follows

$$P(h_k = h_{k-1} + 1 | h_{k-1}) = 1 - p ; \qquad \text{[NACK received]} \quad (7)$$

$$P(h_k = 1 | h_{k-1}) = p ; \qquad \text{[ACK received]} \quad (8)$$

and ii) when the BS idles, i.e. $u_{k-1,T} = 0$, the transition is deterministic

$$P(h_k = 1 | h_{k-1}) = 1 , \qquad \text{if } s_{k-1,T} = 1 ; \quad (9)$$

$$P(h_k = h_{k-1} + 1 | h_{k-1}) = 1 , \qquad \text{if } s_{k-1,T} = \emptyset . \quad (10)$$

For state transitions that do not occur at frame boundaries, the quantity h_k remains fixed and the set $s_{k,n}$ evolves according to the scheduling decisions and feedback signals. If the BS idles during slot $(k,n-1)$, the delivery status of the packet does not change, thus

$$P(s_{k,n} = s_{k,n-1} | s_{k,n-1}) = 1 . \quad (11)$$

If the BS transmits during slot $(k,n-1)$, the set $s_{k,n}$ depends upon the outcome of the transmission, as given by

$$P(s_{k,n} = \emptyset | s_{k,n-1}) = 1 - p ; \qquad \text{[NACK received]} \quad (12)$$

$$P(s_{k,n} = 1 | s_{k,n-1}) = p . \qquad \text{[ACK received]} \quad (13)$$

The last concept to be discussed prior to the cost-to-go function is the objective. The objective function of the Decoupled Model, $\tilde{J}_K^{\pi^*}$, is analogous to (3), except that it introduces the service charge C and evolves in slot increments (instead of frame increments). The expression for $\tilde{J}_K^{\pi^*}$ is given by

$$\tilde{J}_K^{\pi^*} = \min_{\pi \in \Pi} \mathbb{E} \left\{ \sum_{k=1}^{K} \sum_{n=1}^{T} \left(\alpha \, h_k + C \, \mathbb{I}_{\{u_{k,n}=1\}} \right) \right\} , \quad (14)$$

where $\mathbb{I}_{\{u_{k,n}=1\}}$ is an indicator function that takes the value 1 if $\{u_{k,n} = 1\}$ is true and takes the value 0, otherwise.

The cost-to-go function $J_{k,n}(s_{k,n},h_k)$ associated with the optimization problem in (14) has two forms. For the *last* slot of any frame k, namely slot (k,T), the cost-to-go is expressed as

$$J_{k,T}(s_{k,T},h_k) = \alpha h_k + \min_{u_{k,T}} \left\{ C \, \mathbb{I}_{\{u_{k,T}=1\}} + \mathbb{E}[J_{k+1,1}(\emptyset, h_{k+1})] \right\}$$
$$(15)$$

and for slots *other than the last*

$$J_{k,n}(s_{k,n},h_k) = \alpha h_k + \min_{u_{k,n}} \left\{ C \, \mathbb{I}_{\{u_{k,n}=1\}} + \mathbb{E}[J_{k,n+1}(s_{k,n+1}, h_k)] \right\}$$
$$(16)$$

Given a network setup (K,T,p,α,h_1,C), it is possible to use backward induction on (15) and (16) to compute the optimal scheduling policy π^* for the Decoupled Model. However, for the purpose of designing the Index Policy, it is not sufficient to provide an algorithm that computes π^*. The Index Policy is based on a complete characterization

of π^*. Proposition 3 provides a key feature of the optimal scheduling policy which is used in its characterization.

Proposition 3: Consider the Decoupled Model and its optimal scheduling policy π^*. During any frame k, the optimal policy either: (i) idles in every slot; or (ii) transmits until the packet is delivered or the frame ends.

Proof: The proof follows from the analysis of the backward induction algorithm on (15) and (16). For this proof, we assume that the algorithm has been running and that the values of $J_{k+1,1}(s_{k+1,1},h_{k+1})$ for all possible system states are known. The proof is centered around the backward induction during frame k and for a fixed value of h_k.

First, we analyze the (trivial) case in which the packet has already been delivered. Consider any slot (k,n) with $s_{k,n} = 1$. In this case, the optimal scheduling policy always idles.

For the case of an undelivered packet, we start by analyzing the last slot of the frame, namely slot (k,T). It follows from the cost-to-go in (15) that the optimal scheduling decision $u_{k,T}^*$ depends only on the expression

$$C - p \left[J_{k+1,1}(\emptyset, h_k + 1) - J_{k+1,1}(\emptyset, 1) \right] . \quad (17)$$

The optimal policy idles in slot (k,T) if (17) is non-negative and transmits if (17) is negative.

By analyzing the cost-to-go function in (16), which is associated with the optimal scheduling decisions in the remaining slots of frame k, it is possible use mathematical induction to establish that:

- if it is optimal to transmit in slot $(k,n+1)$, then it is also optimal to transmit in slot (k,n); and
- if it is optimal to idle in slot $(k,n+1)$, then it is also optimal to idle in slot (k,n).

We conclude that if (17) is non-negative, the optimal policy idles in every slot of frame k, and if (17) is negative, the optimal policy transmits until the packet is delivered or until frame k ends. ∎

Let $\Gamma \subset \Pi$ be the subclass of all scheduling policies that satisfy the conditions in Proposition 3. Since $\pi^* \in \Gamma$, we can reduce the scope of the Decoupled Model to policies in Γ without loss of generality. In the following section, we redefine the Decoupled Model so that scheduling decisions are made only once per frame, rather than once per slot. This new model is denoted Frame-Based Decoupled Model.

B. Frame-Based Decoupled Model

Denote by u_k the scheduling decision at the beginning of frame k. We let $u_k = 0$ if the BS idles in every slot of frame k and $u_k = 1$ if the BS transmits repeatedly until the packet is delivered or the frame ends.

Since this discrete-time decision problem evolves in frames, we can fully represent its state by h_k. State transitions follow the evolution of h_k in (1) and can be divided into two cases: i) when the BS idles during frame $k-1$, i.e. $u_{k-1} = 0$, then

$$P(h_k = h_{k-1} + 1 | h_{k-1}) = 1 , \quad (18)$$

and ii) when the BS transmits, i.e. $u_{k-1} = 1$, the transition depends on whether the packet was delivered or discarded

during frame $k-1$, as follows

$$P(h_k = h_{k-1}+1 | h_{k-1}) = (1-p)^T ; \quad \text{[discarded]} \quad (19)$$
$$P(h_k = 1 | h_{k-1}) = 1-(1-p)^T ; \quad \text{[delivered]} \quad (20)$$

The objective function of the Frame-Based Decoupled Model, $\hat{J}_K^{\pi^*}$, is given by

$$\hat{J}_K^{\pi^*} = \min_{\pi \in \Gamma} \mathbb{E} \left\{ \sum_{k=1}^{K} \left(T\alpha h_k + \hat{C} \mathbb{I}_{\{u_k=1\}} \right) \right\}, \quad (21)$$

where $\hat{C} = C(1-(1-p)^T)/p$ is the expected value of the service charge incurred during a frame in which $u_k = 1$.

Notice that the Frame-Based Decoupled Model is equivalent to the Decoupled Model when the optimization is carried over the policies in Γ. Both models have the same optimal scheduling policy π^*. Next, we characterize π^* for the network in steady-state.

For analyzing the system in steady-state, we consider the problem over an infinite horizon, namely $K \to \infty$. The state and control of the system in steady-state are denoted h and u, respectively. Then, Bellman equations are given by $S(1) = 0$ and

$$S(h) + \lambda = \min \left\{ \begin{array}{c} T\alpha h + S(h+1); \\ \hat{C} + T\alpha h + (1-p)^T S(h+1) \end{array} \right\}, \quad (22)$$

for all $h \in \{1, 2, \cdots\}$, where λ is the optimal average cost and $S(h)$ is the differential cost-to-go. Notice that the upper part of the minimization is associated with idling in every slot, i.e. $u = 0$, and the lower part with transmitting until the packet is delivered or the frame ends, i.e. $u = 1$, with ties being broken in favor to idling. The stationary scheduling policy that solves Bellman equations[3] is given in Proposition 4.

Proposition 4: (Threshold Policy) Consider the Frame-Based Decoupled Model over an infinite horizon. The stationary scheduling policy π^* that solves Bellman equations (22) is a threshold policy in which the BS transmits when the frame has a state $h \geq H$ and idles when $1 \leq h < H$, where the threshold H is given by

$$H = \left\lfloor 1 - A + \sqrt{A^2 + \frac{2C}{pT\alpha}} \right\rfloor, \quad (23)$$

and the value of A is

$$A = \frac{1}{2} + \frac{(1-p)^T}{(1-(1-p)^T)}.$$

Proof: The expression of the objective function in (21) outlines a trade-off between the cost of transmitting in a frame (expected service charge $\hat{C}$) and the cost of not transmitting (increasing value of h). Naturally, we expect that the optimal scheduling decision is to idle during frames in which h is low (avoiding the service charge) and to

[3]In general, Expected Average Cost problems over an infinite horizon and with *infinite state space* are challenging to address. For the Frame-Based Decoupled Model, it can be shown that [15, Proposition 5.6.1] is satisfied under some additional conditions on Γ. The results in [15, Proposition 5.6.1] and Proposition 4 are sufficient to establish the optimality of the stationary scheduling policy π^*.

transmit when h is high (attempting to reduce the value of h). The proof consists of three parts: i) assume that π^* is a threshold policy on h, ii) solve Bellman equations under this assumption, and iii) show that the solution is consistent. The details of the proof are omitted due to length constraints. ∎

The complete characterization of the policy π^* that optimizes the Decoupled Model and the Frame-Based Decoupled Model gives us the background to establish indexability and to obtain the Index Policy for the AoI minimization problem.

C. Indexability and Index Policy

Consider the Decoupled Model and its optimal scheduling policy π^*. Let $\mathscr{P}(C)$ be the set of system states for which it is optimal to idle when the service charge is C, i.e. $\mathscr{P}(C) = \{h \in \mathbb{N} | h < H\}$. Recall from (23) that the threshold H is a function of C. The definition of indexability is given next.

Definition 5: (Indexability) The Decoupled Model associated with client i is *indexable* if $\mathscr{P}(C)$ increases monotonically from $\emptyset$ to the entire state space, namely $\mathbb{N}$, as C increases from 0 to $+\infty$. The AoI minimization problem is indexable if the Decoupled Model is indexable for all clients i.

The indexability of the Decoupled Model follows directly from the expression of H in (23). Substituting $C = 0$ yields $H = 1$, what implies $\mathscr{P}(C) = \emptyset$, and the limit $C \to +\infty$ gives $H \to +\infty$ and, consequently, $\mathscr{P}(C) = \mathbb{N}$. Since this is true for the Decoupled Model associated with every client i, we conclude that the AoI minimization problem is indexable. Prior to introducing the Index Policy, we define the Index.

Definition 6: (Index) Consider the Decoupled Model and denote by $C(h)$ the Index in state h. Given indexability, $C(h)$ is the infimum service charge C that makes both scheduling decisions (idle, transmit) equally desirable in state h.

The closed-form expression for the Index $C(h)$ yields from the observation that if both scheduling decisions are equally desirable in state h, it follows from the definition of threshold that $H = h+1$ and H is integer-valued. Thus, by substituting $H = h+1$ into (23) and isolating C, we have

$$C(h) = \frac{T\alpha}{2} ph \left[h + \frac{1 + (1-p)^T}{1 - (1-p)^T} \right]. \quad (24)$$

Returning to our original problem, consider the AoI minimization problem with M clients and no service charge. Let $C_i(h_{k,i})$ be the Index associated with client i during frame k. The Index Policy is as follows: in each slot (k,n), the BS transmits to the client with an undelivered packet and highest value of $C_i(h_{k,i})$, with ties being broken in favor of the client with lowest index i. An important feature of the Index Policy is that it requires low computational resources even for networks with a large number of clients.

The performance of this heuristic policy is evaluated in the next section. However, it is possible to anticipate that the Index Policy is optimal for symmetric networks. Notice that when $\alpha_i = \alpha$ and $p_i = p$, prioritizing according to $C_i(h_{k,i})$ is identical to prioritizing according to $h_{k,i}$, i.e. the Index Policy is equivalent to the Greedy Policy. Thus, from the result in Theorem 1 (Optimality of Greedy), we conclude

that the Index Policy is the optimal scheduling policy for symmetric networks.

V. SIMULATION RESULTS

The metric of interest for evaluating the performance of the scheduling policies is the normalized EWSAoI, namely EWSAoI$/M$. Using this metric, we compare three scheduling policies: i) Greedy Policy; ii) Index Policy; and iii) Dynamic Program. The numerical results associated with the Greedy and Index Policies are *simulations* of the policies presented in Secs. III and IV, respectively. On the other hand, the results from the Dynamic Program are *computations* of the metric. The quantity EWSAoI$/M$ is computed by applying Value Iteration to (3). By definition, the Dynamic Program yields the optimal performance.

Figs. 5 and 6 evaluate the scheduling policies in a variety of network settings. In Fig. 5, we consider a two-user *symmetric* network with $T = 5$, $K = 150$ and both clients having the same $\alpha_i = 1$ and $p_i = p \in \{1/15, \cdots, 14/15\}$. In Fig. 6, we consider a two-user *general* network with $K = 200$, $p_1 = 2/3$, $p_2 = 1/10$, $T \in \{1, \cdots, 10\}$ and both clients having $\alpha_i = 1$. Recall that the service charge C is used only for the purpose of deriving the Index Policy. The service charge is not part of the AoI minimization problem and is not considered in the numerical results presented in this section.

Our results show that the performance of the Index Policy is comparable to the performance of the optimal policy (DP) in every network setting considered in Figs. 5 and 6. Moreover, the results in Fig. 5 support the optimality of the Index Policy for symmetric networks established in Sec. IV-C. Fig. 6 suggests that, in the asymmetric case, the Index Policy outperforms the Greedy Policy for all values of T.

VI. CONCLUDING REMARKS

This paper considered a wireless broadcast network with a BS sending time-sensitive information to multiple clients over unreliable channels. We studied the problem of optimizing scheduling decisions with respect to the expected weighted sum AoI of the clients in the network. Our main contributions include showing that the Greedy Policy is optimal for the case of symmetric networks; establishing indexability of the general network case; and obtaining the Whittle Index in closed-form. Numerical results demonstrate the performance of the Index Policy. Interesting extensions include consideration of stochastic arrivals, time-varying channels and multi-hop networks.

REFERENCES

[1] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proceedings of IEEE INFOCOM mini-conference*, 2012.

[2] L. Huang and E. Modiano, "Optimizing age-of-information in a multi-class queueing system," in *IEEE International Symposium on Information Theory (ISIT)*, 2015.

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

[4] B. T. Bacinoglu, E. T. Ceran, and E. Uysal-Biyikoglu, "Age of information under energy replenishment constraints," in *Information Theory and Applications Workshop at UCSD*, 2015.

[5] Y. Sun, E. Uysal-Biyikoglu, R. Yates, C. E. Koksal, and N. Shroff, "Update or wait: How to keep your data fresh," in *Proceedings of IEEE INFOCOM*, 2016.

[6] Q. He, D. Yuan, and A. Ephremides, "Optimizing freshness of information: On minimum age link scheduling in wireless systems," in *International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, May 2016, pp. 1–8.

[7] P. P. Bhattacharya and A. Ephremides, "Optimal scheduling with strict deadlines," *IEEE Transactions on Automatic Control*, vol. 34, pp. 721–728, July 1989.

[8] I.-H. Hou, V. Borkar, and P. R. Kumar, "A theory of qos for wireless," in *Proceedings of IEEE INFOCOM*, Apr. 2009, pp. 486–494.

[9] K. S. Kim, C.-P. Li, I. Kadota, and E. Modiano, "Optimal scheduling of real-time traffic in wireless networks with delayed feedback," in *Proceedings of IEEE Allerton Conference on Communication, Control and Computing*, Oct. 2015, pp. 1143–1149.

[10] V. Raghunathan, V. Borkar, M. Cao, and P. R. Kumar, "Index policies for real-time multicast scheduling for wireless broadcast systems," in *Proceedings of IEEE INFOCOM*, Apr. 2008.

[11] B. Li, R. Li, and A. Eryilmaz, "Throughput-optimal scheduling design with regular service guarantees in wireless networks," *IEEE/ACM Transactions on Networking*, vol. 23, no. 5, pp. 1542–1552, Oct. 2015.

[12] 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 *Proceedings of IEEE INFOCOM*, Jan. 2015.

[13] D. Stoyan, *Comparison Methods for Queues and Other Stochastic Models*, ser. Wiley Series in Probability and Statistics. Wiley, 1983.
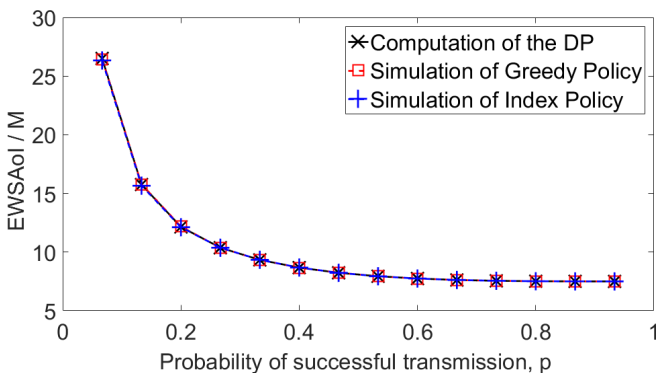
Fig. 5. Numerical results of the normalized EWSAoI for a two-user symmetric network with $T = 5, K = 150, \alpha_i = 1, p_i = p, \forall i$. The simulation result for each policy and for each value of p is an average over 1.000 runs.
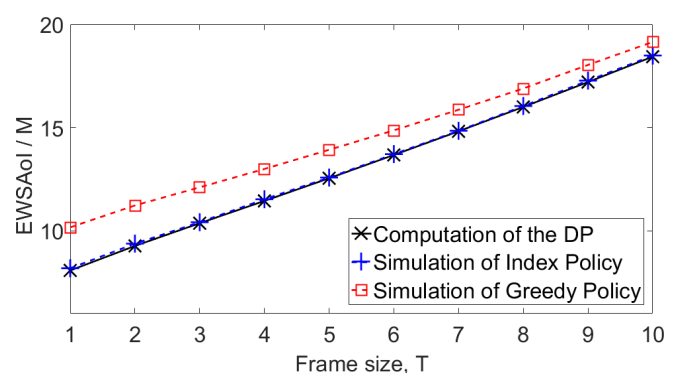

Fig. 6. Numerical results of the normalized EWSAoI for a two-user general network with $K = 200, \alpha_i = 1, p_1 = 2/3, p_2 = 1/10, \forall i$. The simulation result for each policy and for each value of T is an average over 1.000 runs.

[14] A. Ganti, E. Modiano, and J. N. Tsitsiklis, "Optimal transmission scheduling in symmetric communication models with intermittent connectivity," *IEEE Transactions on Information Theory*, vol. 53, no. 3, Mar. 2007.

[15] D. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Athena Scientific, 2005, vol. 1.

[16] P. Whittle, "Restless bandits: Activity allocation in a changing world," *Journal of Applied Probability*, vol. 25, pp. 287–298, 1988.

[17] P. Mansourifard, T. Javidi, and B. Krishnamachari, "Optimality of myopic policy for a class of monotone affine restless multi-armed bandits," in *Proceedings of IEEE Conference on Decision and Control*, Dec. 2012, pp. 877–882.

[18] K. Liu and Q. Zhao, "Indexability of restless bandit problems and optimality of whittle index for dynamic multichannel access," *IEEE Transactions on Information Theory*, vol. 56, pp. 5547–5567, Nov. 2010.

[19] R. R. Weber and G. Weiss, "On an index policy for restless bandits," *Journal of Applied Probability*, vol. 27, no. 3, pp. 637–648, 1990.

[20] J. Gittins, K. Glazebrook, and R. Weber, *Multi-armed Bandit Allocation Indices*, 2nd ed. Wiley, Mar. 2011.