A Whittle Index Approach to Minimizing Functions of Age of Information

Vishrant Tripathi and Eytan Modiano
Laboratory for Information & Decision Systems, MIT

Abstract—**We consider a setting where multiple active sources send real-time updates over a single-hop wireless broadcast network to a monitoring station. Our goal is to design a scheduling policy that minimizes the time-average of general non-decreasing cost functions of Age of Information. We use a Whittle index based approach to find low complexity scheduling policies that have good performance, for reliable as well as unreliable channels. We prove that for a system with two sources, having possibly different cost functions and reliable channels, the Whittle index policy is exactly optimal. For reliable channels, we also derive structural properties of an optimal policy, that suggest that the performance of the Whittle index policy may be close to optimal in general. These results might also be of independent interest in the study of restless multi-armed bandit problems with similar underlying structure. Finally, we provide simulations comparing the Whittle index policy with optimal scheduling policies found using dynamic programming, which support our results.**

I. Introduction

Many emerging applications depend on the timely delivery of status updates from a number of sources to a central monitor over a wireless network. Examples include sensor and actuator data for networked control systems, collecting information for IoT applications, mobility data in vehicular networks, and real-time surveillance and monitoring.

Age of Information (AoI) is a metric that captures timeliness of received information at a destination [1], [2]. Unlike packet delay, AoI measures the lag in obtaining information at a destination node, and is therefore suited for applications involving gathering or dissemination of time sensitive updates. Age of information, at a destination, is defined as the time that has elapsed since the last received information update was generated at the source. AoI, upon reception of a new update packet, drops to the time elapsed since generation of the packet, and grows linearly otherwise. Over the past few years, there has been a rapidly growing body of work on analyzing AoI for queuing systems [1]–[6], and using AoI as a metric for scheduling policies in networks [7]–[14].

The problem of minimizing age of information in single-hop networks was first considered in [7] and [8]. In these works, the authors considered a base station collecting time-sensitive information from a number of sources over a wireless broadcast network, where only one source can send an update at any given time. They looked at weighted linear combinations of AoI of all sources as the metric to be optimized. This prompted the design of low complexity scheduling policies that provably minimize weighted sum AoI at the base station, up to a constant multiplicative factor. These results crucially depend on the fact that for linear AoI, one can find a stationary randomized policy that is factor-2 optimal. As we will see later, this observation does not hold for general functions of AoI. In fact, stationary randomized policies can be arbitrarily worse than simple heuristic policies.

Scheduling problems with weighted linear combinations of age have also been considered with throughput constraints in [9] and with general interference constraints in [10]. AoI-based scheduling with stochastic arrivals was considered in [13], where a Whittle Index policy was shown to have good performance.

On the other hand, nonlinear cost functions of age were introduced as a natural extension to the AoI metric in [2] for characterizing how the level of dissatisfaction depends on data staleness in a more general manner. Nonlinear functions of age of information were also discussed in the context of queuing systems in [15] and [16]. These papers develop the notion of value of information and use nonlinear cost of update delays, which correspond to nonlinear age cost functions.

Nonlinear functions of age have also been discussed in the context of networked control systems in [17], [18] and [19]. In [17], the authors discuss a real time networked control system and show that the cost function is characterized as a non-decreasing, possibly nonlinear, function of AoI. In [18], the authors formulated the state estimation problem for an LTI system, where the state of a discrete-time LTI system can be observed in any time-slot by paying a fixed transmission cost. The problem of minimizing the time-average of the sum of the estimation error and transmission cost reduces to minimizing a non-decreasing age-cost function for a single source with a fixed transmission cost.

In this work, we consider a setting similar to the one in [7] and [8]. We look at a wireless broadcast network with N sources generating real-time updates that need to be sent to a monitoring station. In any time-slot, only one source can attempt a transmission to the base station. Instead of weighted sum AoI, we are interested in minimizing the time-average of *general non-decreasing cost functions* of AoI, summed over all sources. Examples of such functions include $f(x) = 2^x$, $f(x) = \log(x)$, $f(x) = \mathbb{1}_{\{x \geq 10\}}$, etc. See Fig.1 for examples. We develop a restless mutli-armed bandit formulation for the

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.

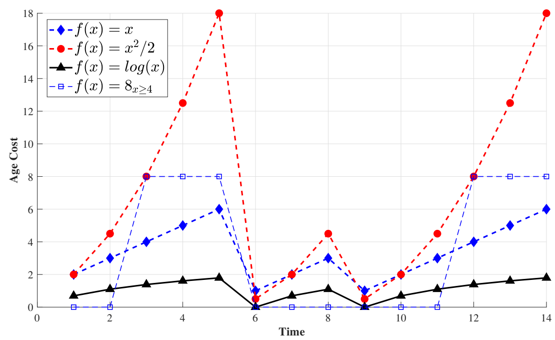

Fig. 1: Linear, quadratic, logarithmic and indicator cost functions for a sample age process. The linear process tracks the actual values of AoI.

problem and use a Whittle Index based approach to find low complexity scheduling policies that have good performance.

Closer to our work, scheduling to minimize functions of age has also been considered in [11] and [12]. In [11], the authors deal with minimizing symmetric functions of age over multiple orthogonal unreliable channels and show that simple greedy schemes are asymptotically optimal. In [12], the authors formulate the general functions of age problem with reliable channels and develop a high complexity algorithm that achieves minimum age. They also derive a key structural property of the optimal policy in this setting - the optimal policy is always periodic. However, their approach does not extend to the setting with unreliable channels. In this work, we consider unreliable channels and also build upon results from [12] and [13] to derive stronger structural properties for optimal policies. These properties hint at why the performance of the heuristic Whittle index policy may be close to optimal.

The remainder of the paper is organized as follows. In Section II, we describe the general system model. In Section III, we describe the equivalent restless multi-armed bandit formulation and discuss why we use the Whittle Index approach to solve the problem. In Section IV, we discuss the functions of age problem with reliable channels, develop the Whittle Index solution for this setting, and also prove key structural properties that an optimal policy must satisfy. In Section V, we find the Whittle Index policy for the functions of age problem with unreliable channels.

II. Model

Consider a single-hop wireless network with N active sources generating real-time status updates that need to be sent to a base station. We consider a slotted system in which each source takes a single time-slot to transmit an update to the base station. Due to interference, only one of the sources can transmit in any given time-slot.

For every source i, the age of information at the base station $A_i(t)$ measures the time elapsed since it received a fresh information update from the source. We assume active sources, i.e. in any time-slot, sources can generate fresh updates at will. Let $s(t)$ be the source activated in time-slot

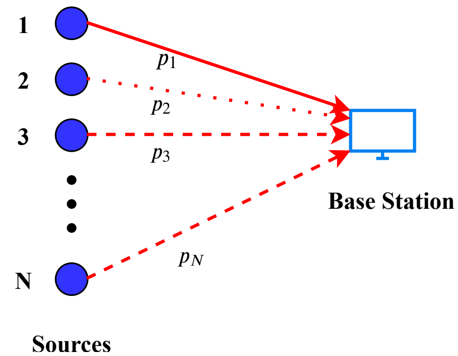

t and $u_i(t)$ be a Bernoulli random variable with parameter p_i that denotes channel reliability between the i^{th} source and the base station. Then, we have

$$A_i(t+1) = \begin{cases} A_i(t)+1, & \text{if } s(t) \neq i \text{ or } u_i(t) = 0, \\ 1, & \text{if } s(t) = i \text{ and } u_i(t) = 1. \end{cases} \quad (1)$$

In this work, we consider general cost functions of age as our metric of interest. For each source i, let $f_i(\cdot)$ denote a positive *non-decreasing* cost function.

Let π be a scheduling scheme that decides which sources to schedule in every time-slot. The age process $A_i(t)$ depends on π and the channel processes. Then, the expected average cost of age for source i is given by

$$C_i^{\text{ave}}(\pi) \triangleq \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\Big[\sum_{t=1}^{T} f_i(A_i^{\pi}(t))\Big], \quad (2)$$

where $A_i^{\pi}(t)$ is age process for the i^{th} source under policy π.

Our goal is find a schedule π that minimizes the sum of average costs of age of sources. Let Π denote the set of causal scheduling policies, then we want to solve the following optimization problem

$$C^* = \min_{\pi\in\Pi} \sum_{i=1}^{N} C_i^{\text{ave}}(\pi), \quad (3)$$

where C^* is minimum average cost and π^* is an optimal scheduling policy.

III. RMAB Formulation

In this section, we show that scheduling to minimize such a metric can be reformulated as a restless multi-armed bandit (RMAB).

Consider a restless multi-armed bandit problem with N arms. The state space for every arm i is the set of positive integers $\mathbb{Z}^+$. The state evolution of the arm depends on whether it is currently active or not. Let the state of arm i at time t be denoted by $A_i(t)$. If arm i is active in time-slot t then the state evolution is given by

$$A_i(t+1) = \begin{cases} A_i(t)+1, & \text{w.p. } 1-p_i \\ 1, & \text{w.p. } p_i. \end{cases} \quad (4)$$

If the arm is not active in time-slot t, then the state evolution is given by

$$A_i(t+1) = A_i(t) + 1. \quad (5)$$

For every arm i, there is a cost function $f_i : \mathbb{Z}^+ \to \mathbb{R}^+$ which maps the states of the arm to their associated costs. Thus, the cost of a state $\boldsymbol{x} \in \mathbb{Z}^{+^N}$ is given by $\sum_{i=1}^N f_i(x_i)$, where $\boldsymbol{x}$ is a vector of ages and x_i is the age of the i^{th} source. Given that only one arm can be activated in any time-slot, the goal of the RMAB framework is to find a scheduling policy that minimizes the total time average cost of running this system.

This establishes the equivalence between the functions of age problem discussed earlier and a corresponding restless multi-armed bandit. Observe that the "restless" part of our construction cannot be dropped, since the states of the arms do not freeze when they are not active and there is no way to reformulate our problem as just a multi-armed bandit problem. If that were the case, we could have found an optimal policy by solving for the Gittins index [20]. However, finding optimal policies for restless bandits is much harder. The usual approach is to find the Whittle Index policy which provides good performance under certain conditions, namely *indexability* of the RMAB problem.

In [7] and [8], the authors develop three methods to solve the minimum age scheduling problem. First, they look at stationary randomized policies, where a source i is scheduled at random with a fixed probability p_i. They find a stationary randomized policy that is factor-2 optimal for weighted sum AoI. However, this result does not hold for general functions: even the best stationary randomized policies in our setting can lead to an unbounded overall cost, despite there being very simple policies that have bounded cost. We demonstrate this with a simple example.

Consider two identical sources with cost functions given by $f(x) = 3^x$ and reliable channels, i.e $p_1 = p_2 = 1$. Any stationary randomized policy schedules at least one of the sources with probability less than or equal to 0.5. For this source, the average cost is lower bounded by $\lim_{T\to\infty} \sum_{t=1}^T (3^t) \frac{0.5^T}{T}$ since with probability at least 0.5, it does not get to transmit and its age increases by 1 in every time-slot. Observe that this lower bound goes to infinity and hence the average cost also goes to infinity for all stationary randomized policies. On the other hand, a simple round-robin scheme that schedules the two sources in alternating time-slots guarantees bounded cost for both sensors. Thus, stationary randomized policies can be infinitely worse than the optimal policy for the functions of age problem.

The second method developed for age-based scheduling in [7], [8] uses a Max-Weight approach. The authors design a quadratic Lyapunov function for the weighted sum of linear functions of AoI and find the max-weight policy - the policy that maximizes the amount of negative drift in the Lyapunov function in every time-slot. Performance guarantees for the max-weight policy crucially rely on the fact that there exists a stationary randomized policy that is factor-2 optimal for linear functions of age. Since this is not the case for general functions of age, we cannot develop a similar Max-Weight policy for the general functions of age problem.

This finally leaves us with the third method - using a Whittle Index based approach. In the following two sections, we use the RMAB formulation to establish indexability for the functions of age problem and derive a Whittle Index policy. We also show that for the case with 2 sources and reliable channels, the Whittle index policy is exactly optimal. This is a novel result since the optimality of Whittle Index policies is typically shown either only asymptotically, or in symmetric settings for finite systems. On the other hand, our optimality result holds for two asymmetric sources.

IV. Reliable Channels

We first look at the problem with reliable channels between the sources and the base station. This leads to simpler analysis and a better understanding of the problem. Consider the setup described in Section I with channel reliability $u_i(t) = 1$, for all i and t. In other words, the probability of success $p_i = 1, \forall i$.

In Section III, we showed that the functions of age minimization problem is equivalent to a restless multi-armed bandit problem. Next, we use a Whittle Index based approach to try and solve the problem.

The first step in the Whittle Index approach is to formulate the *decoupled problem*, where we consider a single arm in isolation with a fixed charge required to activate the arm.

Definition *Decoupled Problem*
Consider a single arm with the state space $\mathbb{Z}^+$ and an associated non-decreasing cost function $f : \mathbb{Z}^+ \to \mathbb{R}^+$. Let the state of the arm be $A(t)$. Its evolution is given by

$$A(t+1) = \begin{cases} A(t) + 1, & \text{if not active at time t} \\ 1, & \text{otherwise.} \end{cases}$$

There is a strictly positive activation charge C to be paid in every time-slot that the arm is pulled.

Our goal is to find a scheduling policy that minimizes the time-average cost of running this system. Assuming that the cost function $f(\cdot)$ is non-negative and non-decreasing, we solve the decoupled problem using dynamic programming. The case when the activation charge is set to zero is trivial. The optimal policy is to always activate the arm. So, we consider C to be strictly positive. The single source decoupled problem has also been solved in a networked control system setting in [18].

Theorem *1:* The optimal policy for the decoupled problem is a stationary threshold policy. Let H satisfy

$$f(H) \le \frac{\sum_{j=1}^H f(j) + C}{H} \le f(H+1). \quad (6)$$

Then, the optimal policy is to activate the arm at time-slot t if $A(t) \geq H$ and to let it rest otherwise. If no such H exists, the optimal policy is to never activate the arm.

Proof: See Appendix A. ■

Theorem 1 establishes that the optimal policy for the decoupled problem has a threshold structure. We now want to show that the *indexability* property also holds for the decoupled problem. The indexability property states that as the activation charge C increases from 0 to ∞, the set of states for which it is optimal to activate the arm decreases monotonically from the entire set $\mathbb{Z}^+$ to the empty set $\{\phi\}$.

Theorem *2:* The *indexability* property holds for the decoupled problem.

Proof: See Appendix B in the technical report [21]. ■

The Whittle index approach states that if the decoupled problem satisfies the indexability property, we can formulate a heuristic index policy called the Whittle Index Policy that has good performance.

Definition *Whittle Index*
Consider the decoupled problem and denote by $W(h)$ the Whittle index in state h. Given indexability, $W(h)$ is the infimum charge C that makes both decisions (activate, not activate) equally desirable in state h. The expression for $W(h)$ is given by

$$W(h) = hf\big(h+1\big) - \sum_{j=1}^{h} f(j). \tag{7}$$

Observe that using (6), $C = W(h)$ is the minimum value of the activation charge that makes both actions equally desirable in state h. This gives us the expression for the Whittle index.

Let $W_i(x) := xf_i\big(x+1\big) - \sum_{j=1}^{x} f_i(j)$ represent the index function for the i^{th} decoupled problem. Using these index functions, we can define the Whittle Index Policy.

Definition *Whittle Index Policy*
Let $\pi^W(t)$ be the action taken by the Whittle Index Policy at time t. Then $\pi^W(t)$ is given by

$$\begin{aligned}\pi^W(t) &= \arg\max_{1\leq i\leq N} \Big\{ W_i\big(A_i(t)\big) \Big\} \\ &= \arg\max_{1\leq i\leq N} \Bigg\{ A_i(t) f_i\Bigg(A_i(t)+1\Bigg) - \sum_{j=1}^{A_i(t)} f_i(j) \Bigg\}.\end{aligned}$$

By the monotonicity of $f_i(\cdot)$, it is easy to see that the functions $W_i(\cdot)$ are also monotonically non-decreasing. This is because $W_i(h) - W_i(h-1) = h\big(f_i(h+1) - f_i(h)\big) \geq 0, \forall h$ since $f_i(\cdot)$ is non-decreasing.

Consider the cost functions to be weighted linear functions of AoI, i.e let $f_i(A_i(t)) = w_i A_i(t)$, with positive weights w_i. This is the setting considered in [7] and [8]. The Whittle Index for source i is then given by $W_i(A_i(t)) = w_i(A_i^2(t) + A_i(t))/2$. This is the same as the Whittle index found in [7], where the authors showed that the Whittle policy is optimal for symmetric settings when all the weights are equal. We also establish that for $N = 2$, the Whittle index policy is optimal even for asymmetric settings.

Theorem *3:* For the functions of age problem with reliable channels and two sources, the Whittle index policy is exactly optimal.

Proof: See Appendix F in the technical report [21]. ■

Next, we discuss some general properties that an optimal policy satisfies. These properties help us establish the optimality of the Whittle index policy for $N = 2$.

A. Properties of an Optimal Policy

For the functions of age problem, a policy is stationary if it depends only on the current values of age. A cyclic policy is one that repeats a finite sequence of actions in a fixed order. We define the space of policies that are stationary and periodic.

Definition *Stationary Cyclic Policies*
A stationary cyclic policy is a stationary policy that cycles through a finite subset of points in the state space, repeating a fixed sequence of actions in a particular order.

In [12], the authors show that for reliable channels there exists an optimal policy that is stationary, cyclic and can be found by solving the minimum average cost cycle problem over a large graph.

We look at this cyclic policy and analyze its properties. If there are multiple such cycles, we consider a cycle with the shortest length. We denote the length of the cycle by T and age vectors on the cycle to be $\boldsymbol{x_1}, \ldots, \boldsymbol{x_T}$. Let the corresponding scheduling decisions be $d_1, \ldots, d_T$. This implies that for state $\boldsymbol{x_k}$, taking action d_k leads to the state $\boldsymbol{x_{k+1}}$, where the subscripts cycle back to $1, 2, \ldots$ after T.

We establish an important structural property that such an optimal policy must satisfy, which we call the *strong-switch-type* property. We call the policies that satisfy this property *strong-switch-type* policies.

Definition *Strong-switch-type Policies*
Consider a stationary policy π that maps every point in

the state space ${\mathbb{Z}^+}^N$ to the set of arms $\{1,\ldots,N\}$. We say that such a policy is strong-switch-type if

$$\pi(x_1,\ldots,x_N)=i$$

implies

$$\pi(x'_1,\ldots,x'_N)=i,$$

for all $\boldsymbol{x}$ and $\boldsymbol{x'}$ such that $x'_i \geq x_i$ and $x'_j \leq x_j, \forall j \neq i$.

In words, the strong-switch-type property implies that if a policy decides to activate arm i for a state vector $\boldsymbol{x}$, then for a state vector $\boldsymbol{x'}$ with a higher age for the i^{th} source and lower ages for all the other sources, it still decides to activate source i. Note that our definition of strong-switch-type policies is a stronger version of the switch-type policies introduced in [13].

Theorem *4:* For the functions of age problem with reliable channels, no state-action pairs that are a part of the shortest length optimal cyclic policy can violate the strong-switch-type property.

Proof: See Appendix C in the report [21]. ∎

We can prove this result for general values of N. However, to extend the strong-switch-type property over the entire state-space, we consider systems with up to three sources.

Theorem *5:* There exists an optimal stationary policy for the functions of age problem with reliable channels and up to three sources that has the strong-switch-type property *over the entire state-space*.

Proof: We have already established that points on the minimum average cost cycle satisfy the strong-switch-type property. We extend this policy over the entire state space while maintaining the strong-switch property to obtain a well defined stationary policy in Appendix D in the technical report [21]. ∎

While we prove this result for up to three source and reliable channels, we believe that the strong-switch-type property is a natural property that some optimal policy must have in general, due to monotonicity of cost functions.

We now define the space of policies that can be found as a result of the Whittle Index based approach.

Definition *Index Policies*
Consider a stationary policy π that maps every point in the state space ${\mathbb{Z}^+}^N$ to the set of arms $\{1,\ldots,N\}$. We say that such a policy is an index policy if

$$\pi(x_1,\ldots,x_N)=\arg\max_{1\leq i\leq N}\left\{F_i(x_i)\right\}$$

for all $\boldsymbol{x}$, where $F_i:\mathbb{Z}^+\rightarrow\mathbb{R}$ are monotonically non-decreasing functions for all i.

Observe that if F_i are the same as W_i in the above definition, then we get back the Whittle Index Policy. Also, note that an index policy always satisfies the strong-switch-type property by definition. This is because the index functions $F_i(\cdot)$ are monotonically non-decreasing. We now show that index policies are in fact the same as strong-switch-type policies.

Theorem *6:* For the functions of age problem, every policy that is strong-switch-type is also an index policy.

Proof: The proof is based on induction on the number of sources. We assume that every strong-switch-type policy can be represented as an index policy for systems with N sources. Using this fact, we show that strong-switch-type policies can also be represented as index policies for systems with $N+1$ sources. We also show that the two types of policies are equivalent for the single source decoupled problem, thus completing the proof. The details are in Appendix E in the technical report [21]. ∎

An important point to notice is that while we use the reliability of channels in the proof of Theorem 5, we do not use any such condition for the proof of Theorem 6. Thus, strong-switch-type policies are equivalent to index policies regardless of channel connectivity.

Theorems 5 and 6 together imply the following corollary.

Corollary *1:* For the functions of age problem with reliable channels and up to three source, there exists a stationary optimal policy that is an index policy.

In other words, there exists an optimal policy that looks like the Whittle Index policy in that the arm to be activated has the maximum value among *monotone index functions* that take as arguments only the states of individual arms. This hints at why the performance of Whittle Index policies may be close to optimal.

Observe that the Whittle Index policy would be optimal in general if we could show that it achieves a cost that is minimum among the space of index policies and that the strong-switch-type property holds for some optimal policy. We show that this is indeed the case for $N=2$. However, we later provide an example that shows that the Whittle policy is not optimal, but only close to optimal, for $N=4$.

We leave the question of whether the Whittle index policy is at most a constant factor away from optimal in general to future work. We believe that the structural properties introduced here provide a recipe to proving constant factor optimality of the Whittle index policy, even for general bandit problems with similar underlying structure.

V. Unreliable Channels

We now consider independent Bernoulli channels between every source and the base station, with probability of success p_i for source i. We derive a Whittle index in this setting and establish indexability of the RMAB problem by enforcing a bounded cost condition on the functions $f_i(\cdot)$.

An important fact to notice is that monotonicity in itself is not sufficient to ensure that the system has finite average cost even for $N=1$, in the case of unreliable channels. Consider a single source case where $f(a)=3^a$ and the probability of success $p=0.5$. If the source attempts a transmission in every time-slot, the expected average cost satisfies

$$\limsup_{T\to\infty}\sum_{t=1}^{T}(3^t)\frac{0.5^T}{T} \leq \limsup_{T\to\infty}\frac{1}{T}\mathbb{E}\Big[\sum_{t=1}^{T}3^{A(t)}\Big], \quad (8)$$

since with probability 0.5, the transmission fails and age increases by 1 in every time-slot. However, observe that the summation on the left goes to infinity and thus the expected average cost goes to infinity. This happens despite the source attempting a transmission in every time-slot. To prevent such a situation from happening we enforce the following *bounded cost* condition on the age cost functions f_i in addition to monotonicity

$$\sum_{h=1}^{\infty} f_i(h)(1-p_i)^h < \infty. \quad (9)$$

It can be shown that this condition ensures that the single source case has bounded cost. We define the decoupled problem in this case as follows

Definition *Decoupled Problem*
Consider a single arm with the state space $\mathbb{Z}^+$, probability of success p and an associated non-decreasing cost function $f:\mathbb{Z}^+\to\mathbb{R}^+$ that satisfies the *bounded cost* condition. Let the state of the arm be $A(t)$. If the arm is active at time t, its evolution is given by

$$A(t+1)=\begin{cases}A(t)+1, & \text{w.p. } 1-p\\ 1, & \text{w.p. } p.\end{cases}$$

If the arm is not active in time-slot t, then the state evolution is given by

$$A(t+1)=A(t)+1. \quad (10)$$

There is a strictly positive activation charge C to be paid in every time-slot that the arm is pulled.

As before, our goal is to find a scheduling policy that minimizes the time-average cost of running this system.

Theorem *7:* The optimal policy for the decoupled problem is a stationary threshold policy. Let H satisfy

$$\begin{aligned}p^2(H-1)\Big(\sum_{k=H}^{\infty}f(k)(1-p)^{k-H}\Big)-p\Big(\sum_{j=1}^{H-1}f(j)\Big)\\ \leq C\\ \leq p^2H\Big(\sum_{k=H+1}^{\infty}f(k)(1-p)^{k-H-1}\Big)-p\Big(\sum_{j=1}^{H}f(j)\Big)\end{aligned} \quad (11)$$

Then, the optimal policy is to activate the arm at time-slot t if $A(t)\geq H$ and to let it rest otherwise. If no such H exists, the optimal policy is to never activate the arm.

Proof: See Appendix G in the technical report [21]. ∎

Observe that taking the limit as $p\to 1$ in Theorem 7, we get back the threshold policy for reliable channels derived in Theorem 1. We now establish indexability and derive the functional form of the Whittle Index.

Theorem *8:* The *indexability* property holds for the decoupled problem. Denote by $W(h)$ the Whittle index in state h. Given indexability, $W(h)$ is the infimum charge C that that makes makes both decisions (activate, not activate) equally desirable in state h. The expression for $W(h)$ is given by

$$W(h)=p^2h\big(\sum_{k=1}^{\infty}f(k+h)(1-p)^{k-1}\big)-p\big(\sum_{j=1}^{h}f(j)\big). \quad (12)$$

Proof: See Appendix H in the report [21]. ∎

Again, observe that taking the limit as $p\to 1$, we get back the Whittle Index derived in Section IV. Further, if we assume that the cost functions are weighted linear functions of AoI, i.e. $f_i(A_i(t))=w_iA_i(t)$ where all the weights are positive, then the index functions for the Whittle policy are given by $W_i(A_i(t))=w_ip_iA_i(t)(A_i(t)+\frac{1+(1-p_i)}{1-(1-p_i)})/2$. This corresponds to the index policy developed in [7], where the authors showed that for symmetric settings when all the weights and channels probabilities are equal, the Whittle index policy is optimal.

VI. Simulations

First, we compare the optimal policy, found using dynamic programming, with the Whittle index policy for two sources. We consider six different settings in total - 3 sets of functions, each with reliable and unreliable channels.

For settings A_1 and A_2, the cost functions are chosen to be $f_1(x)=13x$ and $f_2(x)=x^2$. In A_1, we consider reliable channels, i.e. $p_1=p_2=1$. In A_2, we consider unreliable channels, specifically $p_1=0.9$ and $p_2=0.5$. For settings B_1 and B_2, the cost functions are chosen to be $f_1(x)=x^2$ and $f_2(x)=3^x$. In B_1, we consider reliable channels, i.e. $p_1 = p_2 = 1$. In B_2, we consider unreliable channels, specifically $p_1 = 0.65$ and $p_2 = 0.8$. For settings C_1 and C_2, the cost functions are chosen to be $f_1(x) = x^3/2$ and $f_2(x) = 10\log(x)$. In C_1, we consider reliable channels, i.e. $p_1 = p_2 = 1$. In C_2, we consider unreliable channels, specifically $p_1 = 0.55$ and $p_2 = 0.75$. Simulation results are presented in Table I.

Setting	Optimal Cost	Whittle Index Cost
A_1 (reliable)	21.95	21.95
A_2 (unreliable)	36.12	36.28
B_1 (reliable)	8.48	8.48
B_2 (unreliable)	23.16	23.37
C_1 (reliable)	5.69	5.69
C_2 (unreliable)	21.54	21.54

TABLE I: Cost of the Whittle index policy and the optimal dynamic programming policy for 2 sources.

We find the optimal cost for each setting using finite horizon dynamic programming over a horizon of 500 time-slots. For reliable channels, we find the cost of the Whittle index policy by simply implementing it once over 500 time-slots. For unreliable channels, we estimate the expected Whittle index cost by averaging the performance of the Whittle policy over 500 independent runs.

Observe that the Whittle index policy is exactly optimal when the channels are reliable, as expected from our theoretical results. The expected cost for the Whittle index policy is very close to the optimal cost for unreliable channels as well. Also, for the same set of functions, having unreliable channels increases the cost compared to reliable channels, as is expected.

Next, we compare the optimal policy with the Whittle index policy for more than two sources. Simulation results are presented in Table II.

For settings D_1 and D_2, we consider 3 sources. The cost functions are chosen to be $f_1(x) = x^2$, $f_2(x) = 3^x$ and $f_3(x) = x^4$. In D_1, we consider reliable channels, i.e. $p_1 = p_2 = p_3 = 1$. In D_2, we consider unreliable channels, specifically $p_1 = 0.66$, $p_2 = 0.8$ and $p_3 = 0.75$.

For settings E_1 and E_2, we consider 4 sources. The cost functions are chosen to be $f_1(x) = x^3$, $f_2(x) = 2^x$, $f_3(x) = 15x$ and $f_4(x) = x^2$. In E_1, we consider reliable channels, i.e. $p_1 = p_2 = p_3 = 1$. In E_2, we consider unreliable channels, specifically $p_1 = 0.7$, $p_2 = 0.9$, $p_3 = 0.67$ and $p_4 = 0.8$.

No. of Sources	Setting	Optimal Cost	Whittle Index Cost
3	D_1 (reliable)	44.23	44.23
	D_2 (unreliable)	161.19	161.39
4	E_1 (reliable)	73.36	73.36
	E_2 (unreliable)	129.02	130.94
4	F_1 (reliable)	87.66	88.27
	F_2 (unreliable)	158.35	159.81

TABLE II: Cost of the Whittle index policy and the optimal policy for more than 2 sources.

For settings F_1 and F_2, we consider 4 sources. The cost functions are chosen to be $f_1(x) = x^3$, $f_2(x) = e^x$, $f_3(x) = 15x$ and $f_4(x) = x^2$. In F_1, we consider reliable channels, i.e. $p_1 = p_2 = p_3 = 1$. In F_2, we consider unreliable channels, specifically $p_1 = 0.8$, $p_2 = 0.85$, $p_3 = 0.75$ and $p_4 = 0.66$.

We observe that the cost of the Whittle index policy is the same as that obtained using dynamic programming for settings D_1 and E_1. However, for setting F_1, we observe a small gap in performance between the two policies, thus giving us an example that shows that the *Whittle index policy need not be optimal*, in general. We also verify that the optimal policy found using dynamic programming follows a cyclic pattern that satisfies the strong-switch-type property and is distinct from the Whittle index policy. This is also in line with our discussion on structural properties.

We also note that computing the optimal policy using dynamic programming becomes progressively harder in terms of space and time complexity for larger values of N, as the state-space to be considered grows exponentially with N. The Whittle index policy, on the other hand, is very easy to compute and implement with only a linear increase in space and time complexity with the number of sources. Also, as is evident from simulations, the performance of the Whittle policy is close to optimal in every setting considered, thus making it a very good low complexity heuristic.

VII. Conclusion

In this work, we presented the problem of minimizing functions of age of information over a wireless broadcast network. We used a restless multi-armed bandit approach to establish indexability of the problem and found the Whittle index policy. For the case with two sources and reliable channels, we were able to show that the Whittle index policy is exactly optimal. We also established structural properties of an optimal policy, for the case with reliable channels. These properties hint at why the performance of the Whittle index policy is close to optimal in general.

A possible direction of future work is to try and prove constant factor optimality of the Whittle index policy in general, using the structural properties developed in this work. Another interesting extension would be to consider sources with stochastic arrivals instead of active sources.

References

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

[2] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," *IEEE Trans. Inf. Theory*, vol. 63, pp. 7492–7508, Nov. 2017.

[3] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Minimizing the age of information through queues," *IEEE Transactions on Information Theory*, 2019.

[4] L. Huang and E. Modiano, "Optimizing age-of-information in a multi-class queueing system," in *2015 IEEE International Symposium on Information Theory (ISIT)*, pp. 1681–1685, IEEE, 2015.

[5] Y. Inoue, H. Masuyama, T. Takine, and T. Tanaka, "A general formula for the stationary distribution of the age of information and its application to single-server queues," *arXiv preprint arXiv:1804.06139*, 2018.

[6] C. Kam, S. Kompella, G. D. Nguyen, J. E. Wieselthier, and A. Ephremides, "On the age of information with packet deadlines," *IEEE Transactions on Information Theory*, vol. 64, no. 9, pp. 6419–6428, 2018.

[7] I. Kadota, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Minimizing the age of information in broadcast wireless networks," in *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 844–851, IEEE, 2016.
[8] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 26, no. 6, pp. 2637–2650, 2018.
[9] I. Kadota, A. Sinha, and E. Modiano, "Scheduling algorithms for optimizing age of information in wireless networks with throughput constraints," *IEEE/ACM Transactions on Networking*, 2018.
[10] R. Talak, S. Karaman, and E. Modiano, "Optimizing information freshness in wireless networks under general interference constraints," in *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pp. 61–70, ACM, 2018.
[11] V. Tripathi and S. Moharir, "Age of information in multi-source systems," in *GLOBECOM 2017-2017 IEEE Global Communications Conference*, pp. 1–6, IEEE, 2017.
[12] P. R. Jhunjhunwala and S. Moharir, "Age-of-information aware scheduling," in *SPCOM*, 2018.
[13] Y.-P. Hsu, E. Modiano, and L. Duan, "Scheduling algorithms for minimizing age of information in wireless broadcast networks with random arrivals: The no-buffer case," *arXiv preprint arXiv:1712.07419*, 2017.
[14] S. Farazi, A. G. Klein, J. A. McNeill, and D. R. Brown, "On the age of information in multi-source multi-hop wireless status update networks," in *2018 IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 1–5, IEEE, 2018.
[15] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, "Age and value of information: Non-linear age case," in *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 326–330, IEEE, 2017.
[16] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, "The cost of delay in status updates and their value: Non-linear ageing," *arXiv preprint arXiv:1812.09320*, 2018.
[17] J. P. Champati, M. H. Mamduhi, K. H. Johansson, and J. Gross, "Performance characterization using aoi in a single-loop networked control system," *arXiv preprint arXiv:1901.06694*, 2019.
[18] M. Klügel, M. H. Mamduhi, S. Hirche, and W. Kellerer, "Aoi-penalty minimization for networked control systems with packet loss," *IEEE Infocom AoI Workshop*, 2019.
[19] O. Ayan, M. Vilgelm, M. Klügel, S. Hirche, and W. Kellerer, "Age-of-information vs. value-of-information scheduling for cellular networked control systems," *arXiv preprint arXiv:1903.05356*, 2019.
[20] J. C. Gittins, K. D. Glazebrook, R. Weber, and R. Weber, *Multi-armed bandit allocation indices*, vol. 25. Wiley Online Library, 1989.
[21] V. Tripathi and E. Modiano, "A whittle index approach to minimizing functions of age of information," *arXiv preprint arXiv:1908.10438*, 2019.

APPENDIX

A. Proof of Theorem 1

Consider the decoupled problem described in Section IV. Let $u(t)$ be an indicator variable that denotes whether the arm is pulled or not at time t. Under a scheduling policy π that specifies the value of $u(t)$ for all instants of time, the average cost is given by

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T}\Big[f\big(A^\pi(t)\big)+Cu^\pi(t)\Big]. \tag{13}$$

We want to find a policy that minimizes this cost over the space of all policies. Let $S:\mathbb{Z}^+\to\mathbb{R}$ denote the differential cost-to-go function for this problem, let $u:\mathbb{Z}^+\to\{1,0\}$ be the stationary optimal policy and let λ denote the optimal cost. Then, the Bellman equations are given by

$$S(h)=f(h)+\min_{u(h)\in\{1,0\}}\{C,S(h+1)\}-\lambda, \forall h\in\mathbb{Z}^+. \tag{14}$$

Without loss of generality we set $S(1)=0$. Assume that the optimal policy has a threshold structure, i.e. there exists H such that it is optimal to pull the arm $(u(h)=1)$ for all states $h\geq H$ and let it rest otherwise $(u(h)=0)$. If this the case, then the Bellman equations reduce to

$$S(h)=f(h)+C-\lambda, \forall h\geq H. \tag{15}$$

Using the monotonicity of $f(\cdot)$, we conclude that $S(h+1)\geq S(h), \forall h\geq H$. We will use this fact later. For the state $H-1$, we get

$$\begin{aligned} S(H-1)&=f\big(H-1\big)-\lambda+S(H)\\ &=f\big(H-1\big)-\lambda+f(H)-\lambda+C. \end{aligned} \tag{16}$$

Repeating this k times, we get

$$S(H-k)=\sum_{j=0}^{k} f\big(H-j\big)-(k+1)\lambda+C, \tag{17}$$

for all k in $\{1,\dots,H-1\}$. Observe that since we set $S(1)=0$, we get

$$\lambda=\frac{\sum_{j=1}^{H} f(j)+C}{H}, \tag{18}$$

by putting $k=H-1$ in (17). Now assume that H further satisfies the relation given in Theorem 1, i.e.

$$f(H)\leq\frac{\sum_{j=1}^{H} f(j)+C}{H}\leq f(H+1). \tag{19}$$

Using (18), we can simplify (19) as

$$f(H)\leq\lambda\leq f(H+1). \tag{20}$$

Adding $C-\lambda$ to every term above, we get

$$\begin{aligned} &f(H)+C-\lambda\leq C\leq f(H+1)+C-\lambda\\ \implies &S(H)\leq C\leq S(H+1). \end{aligned} \tag{21}$$

Observe that we assumed $f(\cdot)$ to be non-decreasing. This combined with (20) and the Bellman equations (15) and (17) ensures that $S(\cdot)$ is also non-decreasing. Thus, if there exists a state H that satisfies (19), then the threshold policy with threshold H satisfies the Bellman equations and is hence optimal.

The one thing that remains to be shown is the case in which we cannot find some H that satisfies (19). We leave the details of this case to the technical report [21].