

A Unified Sampling and Scheduling Approach for Status Update in Multiaccess Wireless Networks

Zhiyuan Jiang*, Sheng Zhou†, Zhisheng Niu†, Cheng Yu‡

*: Shanghai Institute for Advanced Communication and Data Science, Shanghai University, Shanghai 200444, China.

†: Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing 100084, China.

‡: Department of Electrical and Computer Engineering, Illinois Institute of Technology, Chicago 60616, USA.

Emails: zhiyjiang@foxmail.com, {sheng.zhou, niuzhs}@tsinghua.edu.cn, cheng@iit.edu

Abstract—Information source sampling and update scheduling have been treated separately in the context of real-time status update for age of information optimization. In this paper, a unified sampling and scheduling (S^2) approach is proposed, focusing on decentralized updates in multiaccess wireless networks. To gain some insights, we first analyze an example consisting of two-state Markov sources, showing that when both optimized, the unified approach outperforms the separate approach significantly in terms of status tracking error by capturing the key status variation. We then generalize to source nodes with random-walk state transitions whose scaling limit is Wiener processes, the closed-form Whittle’s index with arbitrary status tracking error functions is obtained and indexability established. Furthermore, a mean-field approach is applied to solve for the decentralized status update design explicitly. In addition to simulation results which validate the optimality of the proposed S^2 scheme and its advantage over the separate approach, a use case of dynamic channel state information (CSI) update is investigated, with CSI generated by a ray-tracing electromagnetic software.

I. INTRODUCTION

For status update (tracking) in, e.g., latency-sensitive cyber-physical systems, age of information (AoI) [1] is a performance metric that captures the end effect of status observation delay from the destination perspective, and thus enables direct and fair comparisons between systems with differences in, e.g., end-to-end delays, information source sampling rates, throughputs and lossless or lossy designs [2]. However, it can be argued that AoI is still an intermediate metric [3], especially in status update systems, whereas the ultimate goal is to optimize the status tracking accuracy at destinations with heterogeneous status information generated remotely by distributed information sources.

We focus on status update through wireless networks. Existing works in this regard implicitly adopt a *separate* approach—the sampling of information sources and the scheduling of updates/transmissions are treated separately (Fig. 1), with objectives of minimizing the real-time sampling error and AoI respectively [3]–[12]. Specifically, in terms of information sampling, Ref. [3]–[6] consider the scenario that one destination node is remotely tracking the status of one source node, with the objective to minimize, e.g., tracking error [4], [6] or AoI [5]. In this single-source scenario, optimal sampling strategies are obtained under assumptions such as random delay due to channel error [6] or limited sampling frequency [4]; these assumptions are made to model the transmissions in

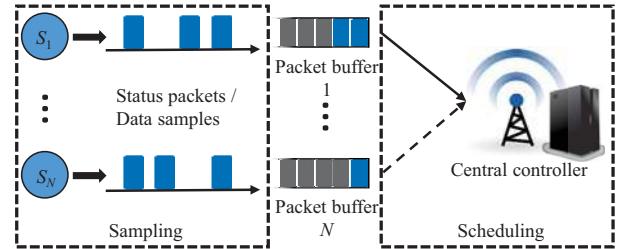


Fig. 1. Status update in a multiaccess wireless network.

wireless networks but are insufficient, due to the considerably more complex behavior of wireless networks. On the other hand, status update scheduling policies usually ignore the structures of information sources, or equivalently, assume the sampling procedure has fully exploited them, and focus on optimizing the AoI [7]–[12]. In these works, the information sources are abstracted as status update packets that arrive randomly [8], [9], [11], or as generate-at-will (active) sources [7], [10], [12]. Scheduling decisions are made to minimize the AoI, irrespective of status variation—a node may be scheduled even if its status is unchanged since the last schedule.

Despite its simplicity, this separate approach is not necessarily optimal in terms of status tracking accuracy based on a *unified* formulation, and the performance difference is unknown. There have been recent efforts [3] towards an effective AoI metric which, compared to AoI, is more directly related to status tracking accuracy in a single-link scenario. The contributions of this paper are summarized as:

- 1) The joint information sampling and transmission scheduling (S^2) problem for multiaccess wireless networks is formulated. Through an example consisting of two-state Markov nodes, it is shown that the optimal S^2 achieves smaller status tracking error than the (near) optimal separate approach. We then consider sources with random-walk state transitions—a discrete version of Wiener processes—and arbitrary error functions; a decentralized near-optimal status update scheme is proposed. Based on extensive simulations, including a use case with channel state information (CSI) sources generated by a ray-tracing software, the proposed scheme shows evident advantage over state-of-the-art separate approaches.
- 2) From a methodology perspective, to be best of our

knowledge, this is the first work that generalizes the Whittle's index to incorporate arbitrary functions which measure the status tracking accuracy, and also the first to apply a mean-field approach to derive the decentralized policy explicitly. As a special case, the Whittle's index of non-linear AoI (arbitrary reasonable non-linear functions) and the corresponding mean-field based decentralized implementation can be readily given.

II. SYSTEM MODEL AND PRELIMINARIES

A multiaccess wireless network is considered, where a central controller collects status updates from N distributed source nodes, denoted by $\{S_1, \dots, S_N\}$. The status updates are conveyed in packets which are transmitted through the wireless multiaccess channel. A time-slotted communication system is considered. In each time slot, only one packet from a source node can be transmitted; a collision happens with more than one simultaneous transmissions. The transmission error probability is $p_{e,n}$ (i.i.d. over nodes and time) if there is no collision. The source nodes only have knowledge of their own statuses in each time slot [3], [4] (expect that in Section III, a genie-aided scheduler which has knowledge of status changes of all source nodes is assumed to obtain a performance bound). The packets containing status updates are stored in the buffers at source nodes when waiting to be transmitted. For active sources, the buffers are unnecessary since status packets are generated at will.

A. General Status Tracking Error Metric for \mathcal{S}^2

We assume that for source node S_n , its status is discrete and belongs to a set $\mathcal{K}_n \triangleq \{\kappa_{n,1}, \dots, \kappa_{n,K_n}\}$. A status $\kappa_{n,k}$ is generic, e.g., sensory data or images, and we assume each status update consumes one packet. We define a general measure of difference between two statuses $\kappa_{n,k}$ and $\kappa_{n,j}$ as $\delta(\kappa_{n,k}, \kappa_{n,j})$; examples include ℓ^1 , ℓ^2 norms when the statuses are described by vectors. Denote the statuses of source node S_n and the destination node under a policy π (including information source sampling and transmission scheduling) at time t as $s_n(t)$ and $\hat{s}_{n,\pi}(t)$ respectively, and $s_n(t), \hat{s}_{n,\pi}(t) \in \mathcal{K}_n$. The T -horizon time-average remote status tracking error can be expressed as $\overline{\Delta}_{n,\pi}^{(T)} \triangleq \frac{1}{T} \sum_{t=0}^{T-1} \delta(s_n(t), \hat{s}_{n,\pi}(t))$. In particular, we are interested in the infinite-horizon regime where $T \rightarrow \infty$. Denote the weighted long-time-average remote status tracking error of all sources as $\overline{\Delta}_\pi \triangleq \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{n=1}^N w_n \overline{\Delta}_{n,\pi}^{(T)}$, where $[w_1, \dots, w_N]$ is a weight vector denoting the importance of each source node. The objective is therefore to seek for a policy π that minimizes the time-average error, i.e.,

$$\min_{\pi} \overline{\Delta}_\pi, \text{ subject to } \sum_{n=1}^N u_n(t) \leq 1, \forall t, \quad (1)$$

where $u_n(t) = 1$ denotes the source node S_n transmits successfully during time slot t ; otherwise $u_n(t) = 0$.

B. Conventional Separate Approach

In a separate sampling and scheduling approach, the information sampling design concerns with the optimal sampling

time of a single source to minimize the remote tracking error, with some assumptions on the communication part, e.g., fixed delay [3], random delay [4] or unreliable channel [6]; additional constraints include limited sampling frequency [4]. Afterwards, the transmission scheduling decision is based on optimizing the AoI [7]–[9], [11] which is defined as follows. The T -horizon time-average AoI of source node S_n is $\overline{h}_{n,\pi}^{(T)} \triangleq \frac{1}{T} \sum_{t=0}^{T-1} h_{n,\pi}(t)$, where $h_{n,\pi}(t)$ denotes the AoI reported by S_n at the t -th time slot under policy π and $h_{n,\pi}(t) \triangleq t - \mu_{n,\pi}(t)$, where $\mu_{n,\pi}(t)$ denotes the sampling time of the newest status received at the destination until time t . The weighted long-time-average AoI is defined as $\overline{h}_\pi \triangleq \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{n=1}^N w_n \overline{h}_{n,\pi}^{(T)}$.

The rationality of this separate approach is that by optimizing the AoI, the tracking error can be minimized consequently by assuming the tracking error grows with the staleness of the information obtained. However, this assumption needs to be reconsidered with generic status variation. The motivating example introduced in the following section will show that this separate approach is not optimal and can be improved by a unified sampling and scheduling approach.

III. STATUS TRACKING ERROR ANALYSIS FOR A NETWORK OF TWO-STATE MARKOV SOURCES

Consider a network with N independent Markov information sources, each with two symmetrical states $\{0, 1\}$ and state change probability of $p_n \triangleq p_{n,0 \rightarrow 1} = p_{n,1 \rightarrow 0}$ ¹. We assume reliable channels in this section and later on generalize in Section IV-A. The performance metric is the long-time sample average error probability which converges to the stochastic average denoted by $\Pr(\varepsilon_n | \pi)$ given policy π , i.e., $\delta(s_n(t), \hat{s}_{n,\pi}(t)) = \mathbb{1}_{\{s_n(t) \neq \hat{s}_{n,\pi}(t)\}}$, $\Pr(\varepsilon_n | \pi) = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{1}_{\{s_n(t) \neq \hat{s}_{n,\pi}(t)\}}$. We further assume in this section that $w_i = 1, \forall i \in \{1, \dots, N\}$, namely all nodes have equal weights. The objective function in (1) is therefore $\overline{\Delta}_\pi = \frac{1}{N} \sum_{n=1}^N \Pr(\varepsilon_n | \pi)$.

A. Optimal \mathcal{S}^2 for Two-State Sources

In this section, we assume that there is a genie coordinating the status update decisions with all the information to obtain the optimal performance as a benchmark; this assumption is unrealistic and will be removed in the next section considering decentralized status update with only local information. The Markov decision process (MDP) is formulated to obtain the optimum, based on which the system state is $\mathcal{B}(t) \triangleq \{(s_1(t), \hat{s}_1(t)), \dots, (s_N(t), \hat{s}_N(t))\}$ where $s_n, \hat{s}_n \in \{0, 1\}$. The transition probability with action $u_n(t) = \mathbb{1}_{\{n=m\}}$, $\forall n \in \{1, \dots, N\}$ is

$$\Pr\{\mathcal{B}(t) \rightarrow \mathcal{B}(t+1)\} = \prod_{n \in \Psi_1} p_n \prod_{i \in \Psi_2} (1 - p_i),$$

where Ψ_1 is the set of source nodes whose statuses change at time slot t , and Ψ_2 contains those unchanged. The statuses at

¹We only consider $p_n \leq 0.5$ since a symmetrical two-state Markov source with $p_n > 0.5$ can be regarded as the modulo-2 superposition of a constant source with alternative 0 and 1 output and a source with change rate $1 - p_n$.

time $t + 1$ at the central controller remain the same as time t , i.e., $\hat{s}_n(t + 1) = \hat{s}_n(t)$, $\forall n \neq m$, except for node m who is updated to $\hat{s}_m(t + 1) = s_m(t)$. Define the T -horizon total cost function as (a constant factor of $\frac{1}{N}$ is omitted)

$$\min_{u_n(t), n \in \{1, \dots, N\}} \mathbb{E} \left[\sum_{t=0}^{T-1} \sum_{n=1}^N \mathbb{1}_{\{s_n(t) \neq \hat{s}_n(t)\}} \right]. \quad (2)$$

The optimal policy is as follows.

Theorem 1 (Optimal \mathcal{S}^2 for Two-State Markov Sources): The optimal policy for the MDP is at time t , node $S_{n_{\text{opt}}}$ updates its current status where $n_{\text{opt}} = \arg \min_{n \in \chi(t)} p_n$, and $\chi(t)$ is the set of nodes with $s_n(t) \neq \hat{s}_n(t)$, and ties are broken arbitrarily.

Proof: The optimal policy is to schedule the source node that has error, and meanwhile is the least likely to change (avoiding repetitive and hence wasted update). The detailed proof is given in Appendix A. \blacksquare

Remark 1: We have derived the optimal policy for the two-state system with a finite horizon, which is considered to be the heavy-lifting of the proof. The generalization to infinite horizon follows standard methods, cf. [13, Theorem 2]. Although significantly simplified, the considered scenario presents itself when considering a network of nodes for e.g., anomaly detection which has two states: normal and alert. \square

B. Near-Optimal Separate Approach

In this subsection, we aim to obtain the optimal separate policy and compare its performance with that of Theorem 1, in order to show whether \mathcal{S}^2 is beneficial in this scenario.

To begin with, the sampling strategy which is independent with the transmission scheduling is considered. The *sample-at-change* sampler, i.e., sampling the source whenever there is a state change, is clearly the optimal lossless sampler in terms of minimum sampling times. Moreover, since there are only two states, the optimal sampling strategy in [4], which is based on a threshold of the status change and proved optimal considering Wiener sources, reduces to sample-at-change. In general, we remark that the optimal sampling strategy is related to assumptions on communications, e.g., random delay and erasure channel. However, it is very difficult to model the overlay wireless network perfectly. Hence, we adopt the sample-at-change strategy which is also used in [3].

According to the separate approach, the sampling strategy deals with the information source structure and the scheduling policy is designed to minimize the AoI. Based on the sample-at-change sampler, the sampled packets arrive at the buffer of S_n based on a Bernoulli distribution with parameter p_n . There are rich existing works on the scheduling policy for AoI optimization in this system setting [7]–[9], [11]. It has been shown that although the optimal scheduling policy seems elusive, a closed-form Whittle's index based policy is obtained in [9, Theorem 3]. Thereby, each node only retains the most up-to-date packet, and the node with the highest index is

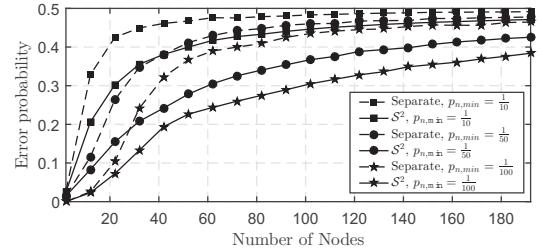


Fig. 2. Tracking error comparisons between optimal \mathcal{S}^2 and near-optimal separate approach based on the two-state Markov sources setting.

scheduled. The index is given by $m(a_n, b_n, p_n) = \frac{1}{2}x_n^2 + \left(\frac{1}{p_n} - \frac{1}{2}\right)x_n$, if $b_n > \frac{p_n}{2}(a_n^2 - a_n) + a_n$; and $\frac{b_n}{p_n}$ otherwise, where $x_n \triangleq \frac{b_n + a_n(a_n - 1)}{1 - p_n + a_n p_n} p_n$, the age of the packet at buffer of S_n is denoted by a_n , the current AoI of S_n is denoted by b_n and $b_n \triangleq h_n - a_n$. It is well-known that the Whittle's index based policy can achieve near-optimal performance of restless multi-armed bandit problems, especially with a large number of source nodes [14]. Therefore, we claim that the index policy is close to optimal to minimize the AoI of the system.

Now we have obtained the optimal \mathcal{S}^2 policy and a near-optimal separate policy in a network with independent 0-1 Markov sources, the status tracking error performance comparisons are in order.

C. Comparisons between \mathcal{S}^2 and Separate Approach

In Fig. 2, we compare the status tracking error performances of the two-state Markov source network with optimal \mathcal{S}^2 derived in Theorem 1, and the near-optimal Whittle's index policy based separate sampling and scheduling approach described in the last subsection. The transition probability p_n is uniformly distributed in $[p_{n,\min}, \frac{1}{2}]$. The simulation runs for 10^6 time slots. It is observed that the performance gap between \mathcal{S}^2 and the separate approach is evident. The reason is that the separate approach, although optimizing the status update timeliness, does not account for the real status change, i.e., instead of being event-triggered, the separate approach is timing-triggered. In this specific case, the separate approach may squander status update opportunities on nodes with large AoI but correct statuses due to the back-and-forth state transition of the two-state Markov chain. In contrast, the optimal \mathcal{S}^2 selects the node with both incorrect status and smallest p_n (least likely to change back).

Having witnessed the potential benefit by \mathcal{S}^2 , let us turn to a more general scenario and consider more practical issues.

IV. DECENTRALIZED EVENT-TRIGGERED STATUS UPDATE FOR NODES WITH RANDOM-WALK STATE TRANSITIONS

One can immediately realize that the optimal \mathcal{S}^2 proposed in Theorem 1 is not feasible, except for illustration purposes, in the sense that it requires a genie to inform the scheduler about the set of nodes where events (status changes) happen at the considered time. Moreover, even the assumption of a central scheduler to make scheduling decisions is questionable,

considering the prohibitive signalling overhead in the massive Internet-of-Things (IoT) system. In summary, there are two distinct challenges in designing event-triggered status update in a multiaccess wireless network:

- 1) How to measure the *event importance* for status update such that by scheduling nodes with important updates, the overall status tracking accuracy is improved?
- 2) How to design a *decentralized* status update scheme based on the event importance?

A natural idea to measure the event importance is by measuring the status change amplitude; however, problems arise in the case with heterogeneous nodes, e.g., different error functions due to different nodes' sensitivities towards error. To address this issue, we adopt a Whittle's index-based solution.

A high-level overview of the proposed scheme, i.e., event-triggered status update (ETSU), is as follows. The Whittle's index, which is denoted by $I_n(t)$ for node n at time t , is calculated by each source node based solely upon its local information, e.g., status change; it is therefore adopted to measure the status change importance and triggers the status update. Each node maps its index to a transmission probability $p_n(t)$ based on a pre-defined function $\Psi(I_n(t))$ which is public and identical among nodes (to be specified later), i.e., $p_n(t) = \Psi(I_n(t))$, $\forall n \in \{1, \dots, N\}$. The nodes undergo a contention period based on their distinct transmission probabilities and the winner transmits a status update packet. Intuitively, a higher index, indicating a higher importance, should be mapped to a higher transmission probability. The methodology is described here.

ETSU

Contention period:

Node S_n transmits with probability $p_n(t) = \Psi(I_n(t))$.

Transmission or collision time slot:

In case of collision, the central controller feeds back a NACK; otherwise ACK.

Go to the contention period.

ETSU can be viewed as a prioritized p -persistent carrier-sense multiple access (CSMA) scheme, and has been proposed in [9] to address AoI optimization in wireless uplinks. However, the challenges here are to design $I_n(t)$ to accommodate event importance instead of AoI in [9], and to design the mapping $\Psi(\cdot)$ explicitly for which we adopt a novel mean-field approach. In what follows, we tackle these challenges one-by-one.

A. Whittle's Index for Nodes with Random-Walk Transitions

Two assumptions on state transitions and status tracking error metric are in order before diving into the index derivation.

- 1) We consider source nodes with state transitions modeled as random walk on a one-dimensional line. Formally, $s_n(t+1) = s_n(t) + a_n(t)$, where

$$\Pr\{a_n(t) = i | s_n(t)\} = \Pr\{a_n(t) = i\} = q_{n,i}, \quad (3)$$

and $q_{n,1} = 0.5$, $q_{n,-1} = 0.5$, $\forall n$. Hence the status evolution of each node is a random walk with i.i.d. (among nodes) and time-homogeneous increments.

- 2) The status tracking error functions $\delta_n(s_n(t), \hat{s}_n(t))$, $\forall n \in \{1, \dots, N\}$, is only related the absolute status difference, i.e., $\delta_n(s_n(t), \hat{s}_n(t)) = \delta_n(|s_n(t) - \hat{s}_n(t)|)$, and satisfy: $\forall n \in \{1, \dots, N\}$, $\delta_n(0) = 0 \leq \delta_n(1) \leq \dots \leq \delta_n(d) \leq \dots$, and there exists $D > 0$, $\forall d \geq D$, $\delta_n(d) > 0$.

The random-walk state transitions capture the essence of many real-world physical phenomenons, e.g., the stock price and the path traced by a foraging animal; the Wiener process considered in the literature [4] is the scaling limit of a random walk with very small steps. Note that the state values s_n should be viewed as normalized status values, and the end effect of heterogeneous status values, e.g., temperature status ranging in $[-30, 30]$ and an image pixel in $[0, 255]$, is captured by distinct error functions $\delta_n(\cdot)$. Several extensions to the random walk assumption are considered trivial and thus not treated specifically: Random walk with different step sizes which, as mentioned before, can be treated by formulating different error functions $\delta_n(\cdot)$; Asymmetrical random walk where, say $q_{n,1} > 0.5$, can be treated with adding a constant drift of $(2q_{n,1} - 1)\tau$, where τ denotes the last update time, to the tracking status $\hat{s}_n(t)$; For random walk with stay probability, i.e., $q_{n,0} \neq 0$ which can be different among nodes, the corresponding index is the one without stay probability multiplied by a factor of $1/(1 - q_{n,0})$ by the standard Markov chain uniformization technique [15, Chapter 6.4]. One other reason to consider random walks is that state transitions of information sources in real world are difficult to model and thus often unknown, whereas random walks offer an approximation.

The second assumption, together with the random walk assumption, essentially makes the problem *state-homogeneous*, i.e., the status tracking difference is sufficient statistics. Therefore, the system state can be simplified to be one-dimensional, i.e., $(s_n(t), \hat{s}_n(t)) \Leftrightarrow (d_n(t) \triangleq |s_n(t) - \hat{s}_n(t)|)$. The state-homogeneous error functions are still considered quite general since a wide range of metrics, e.g., ℓ^1 , ℓ^2 norms and threshold-type error functions, can be included. The requirement of non-decreasing of $\delta_n(d)$ is also reasonable, considering a larger status difference should have a larger status tracking error.

The Whittle's index is a technique to deal with the restless multi-armed bandit (RMAB) problem which usually suffers from the curse of dimensionality when considering a large number of arms. Our status update problem clearly fits into the RMAB framework since the status update decisions for nodes are binary, like pulling arms of bandit machines, and the status changes even if the node is not scheduled i.e., restless. The philosophy of the Whittle's index is to decompose the N -dimensional RMAB into N one-dimensional MDP problems; solving each one-dimensional problem is equivalent to comparing each bandit machine to a machine with constant payoff (or cost) m_n , which is also equivalent to solving a relaxed Lagrange dual problem [14]. By this methodology, the index m_n can be viewed as a measure of the event importance optimized for multi-node status update. The challenge when

applying the Whittle's index is that it is only defined for a subset of RMAB problems which are indexable.

Definition 1 (Indexability): For the decomposed problem, given auxiliary costs m_1 and m_2 , and the sets of states that the optimal action is to idle denoted by Π_{m_1} and Π_{m_2} respectively, the RMAB is indexable if $\forall 0 \leq m_1 < m_2$, $\Pi_{m_1} \subseteq \Pi_{m_2}$, and $\Pi_0 = \emptyset$; for $\Pi_{+\infty}$ is the entire state space. \square

Several existing works [7], [9]–[11], [16] have adopted the Whittle's index. Hence we omit the index formulation and problem decomposition, diving directly into solving the decomposed problems which are described as follows (the node and time indices are omitted for brevity).

$$f(d) + \hat{J}^* = \min \left\{ \begin{array}{l} \delta(d) + \frac{f(d+1)}{2} + \frac{f(d-1)}{2}, \\ m + \frac{f(1)}{2} + \frac{f(0)}{2} \end{array} \right\}, \quad (4)$$

with $d \geq 1$. Considering the long-time average cost MDP formulation, the relative cost function is denoted by $f(d)$, the optimal average cost is denoted by \hat{J}^* , and the auxiliary constant cost in the Whittle's index methodology is denoted by m . We can prescribe $f(0) = 0$ and the state-0 will transit to state-1 with probability 0.5; otherwise stays at state-0.

Theorem 2 (Whittle's Index for Random-Walk Information Sources): The Whittle's index, denoted by $I_{\text{RW},n}(d_n)$ with status difference $d_n = |s_n - \hat{s}_n|$, is

$$I_{\text{RW},n}(d_n) = w_n \sum_{i=1}^{d_n} (2i - d_n) \delta_n(i). \quad (5)$$

The problem is indexable.

Proof: A sketch of the proof is given in Appendix B. \blacksquare

Corollary 1 (Index for Unreliable Channels): Assuming the transmission error probability is $p_{e,n}$ (no collision), the corresponding Whittle's index, denoted by $\hat{I}_{\text{RW},n}(d_n, p_{e,n})$, is the index without transmission error multiplied by $1 - p_{e,n}$, i.e., $\hat{I}_{\text{RW},n}(d_n, p_{e,n}) \xrightarrow{p_{e,n} \rightarrow 0} I_{\text{RW},n}(d_n)(1 - p_{e,n}) + \mathcal{O}(1)$. \square

Proof: This result has an intuitive explanation: For a node with transmission error probability $p_{e,n}$, it takes approximately $1/(1 - p_{e,n})$ times to reach a successful transmission whose equivalent service charge (system is willing to pay) is $I_{\text{RW},n}(d_n)$, and hence the corresponding index with error is the one without transmission error divided by $1/(1 - p_{e,n})$ times of trials. However, the exact expression with general value of $p_{e,n}$ is more complicated and cumbersome without insights. Therefore, we give a sketch of the proof in Appendix C for the case with $p_{e,n} \rightarrow 0$ which has this clear structure. \blacksquare

The Whittle's index policy which, at each time, updates the node with the largest index has been shown to have a strong performance for a wide range of RMAB problems [7], [9], [11], [16], especially in scenarios with a large number of nodes. The results in Theorem 2 and Corollary 1 are, as far as we know, the first results that establish closed-form Whittle's indices and their indexability for general and heterogeneous cost functions $\delta_n(\cdot)$.

B. Optimality of Whittle's Index with Homogeneous Nodes

The Whittle's index policy is in fact optimal for a specific scenario where nodes are homogeneous, i.e., identical error function and identical transmission error. Similar observations that the index policy is optimal in symmetrical networks are found in, e.g., [7], [16], however in different contexts where only identical transmission error is considered.

Proposition 1: Assuming random-walk state transitions and $\delta_n(\cdot) = \delta(\cdot)$, $w_n = 1$, $p_{n,e} = p_e$, $\forall n \in \{1, \dots, N\}$, the index policy based on (5) is optimal. \square

Proof: The proof technique is similar with the proof of Theorem 1. Basically, the index policy updates the node with the largest status difference in this case, which is also the myopic policy based on the MDP formulation. Again, based on backwards induction, it can be shown that the myopic policy is optimal. The details are omitted for brevity. \blacksquare

C. Design of $\Psi(\cdot)$: A Mean-Field Approach

The design of the mapping from index, representing the event importance, to transmission probability in a system with a large number of nodes can be formulated into a mean-field control problem [17]. The key idea is that when $N \rightarrow \infty$, under some conditions specified later, the control of all the nodes can be simplified to the control of the mean-field.

Definition 2 (Mean-Field): The mean field is defined as the probability distribution of the states over the set of nodes, i.e., $F(d, t) \triangleq \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{\{d_n(t) = d\}}$. \square

The general assumptions that guarantee the convergence of the optimal control to the mean-field control [17] include interchangeability of nodes, and interaction with the mean-field for each node. For ease of exposition, we consider one class of homogeneous nodes (identical error functions) in this subsection, such that the former assumption is satisfied; extensions to multiple classes of nodes are straightforward by treating nodes in each class as a mean-field. The latter assumption is also satisfied since every node only concerns with the statistical distribution of other nodes, instead of the state of each individual node.

Unlike the conventional mean-field approach that adopts the Hamilton-Jacobi-Bellman (HJB) equation [17] to account for the system dynamics, the Whittle's index solution derived in Theorem 1 has significantly simplified the control, namely the objective at each time slot is only to schedule the node with the largest index (in the homogeneous case, it is equivalent to the largest d) by considering the mean-field of the states. A schematic of the approach is specified as

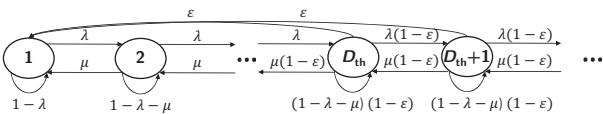
$$D_{\text{th}} \mapsto F_n(D_{\text{th}}) \xrightarrow{(a)} \overline{F}(D_{\text{th}}) \xrightarrow{\overline{F}(D_{\text{th}}^*) = 1 - \nu} D_{\text{th}}^*,$$

which can be explained as follows. First select an index threshold $I_{\text{th}}(D_{\text{th}})$, and any node with index above the threshold transmits with probability p_{tx} ; this produces a stationary distribution of each node $F_n(D_{\text{th}})$. The equivalence of (a) stems from the mean-field assumption, i.e., the state distribution over the set of nodes approaches the stationary distribution of one node in the large system limit. Finally, the index threshold is

derived which allows nodes with indices larger than the top ν -fraction of all nodes to transmit with the same probability p_{tx} . Note that ideally, the maximum index of the mean-field is scheduled; however such an approach is unstable given a continuous mean-field distribution. Therefore we adopt a mapping function that lets the top ν -fraction of nodes compete:

$$\Psi(I) = \begin{cases} p_{\text{tx}}, & \text{if } I \geq I_{\text{th}}; \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

The following theorem gives closed-form expressions of I_{th} and p_{tx} for a general Markov chains where the state transition given D_{th} can be described in the below figure. The random



walk considered in Section IV-A is a special case with $\lambda = \mu = 0.5$. Note that the transition is only among adjacent states; this is not a restriction when sufficiently small time slots are considered. The probability of successful transmission when $D \geq D_{\text{th}}$ is $\varepsilon = \frac{1}{\nu N}$, i.e., the top ν -fraction of nodes (νN nodes) get equal transmission probability.

Theorem 3 (Optimal Mean-Field Control): When $N \rightarrow \infty$, the index threshold I_{th} that allows the top ν -fraction of nodes to transmit is given by the index expression, e.g., (5), with d being the unique solution of the equation

$$\nu^{-1} = \begin{cases} \left(\frac{\lambda}{\mu}\right)^{1-d} - \frac{\lambda}{\mu} \left(\frac{1}{\lambda\beta} - \frac{\varepsilon}{\lambda-\mu}\right) + \frac{\varepsilon d}{\lambda-\mu} + 1, & \lambda \neq \mu, \\ \frac{d}{\lambda\beta} + \frac{\varepsilon(d-1)d}{2\lambda} + 1, & \lambda = \mu, \end{cases} \quad (7)$$

where $\varepsilon = \frac{1}{\nu N}$, $\beta = e_1^\top (I - \hat{P})^{-1} \mathbf{1}$ (see (15)). The transmission probability p_{tx} of the top ν -fraction of nodes is given by the unique solution of the equation

$$(1 - p_{\text{tx}})^{\nu N} = \frac{t_{\text{slot}}}{t_{\text{c}}} \left(\nu N p_{\text{tx}} + (1 - p_{\text{tx}})^{\nu N} - 1 \right), \quad (8)$$

where t_{slot} is the duration of one time slot and t_c is the duration of one contention slot.

Proof: See Appendix D.

Corollary 2 (Closed-Form Solution for Random-Walk Nodes): When $\lambda = \mu$, the solution of (7) is

$$d = \frac{1}{2} - \frac{\nu N}{\beta} + 2\sqrt{\frac{\nu^2 N^2}{\beta^2} - \frac{\nu N}{\beta} + \frac{1}{4} - 2\lambda\nu N(1 - \nu^{-1})},$$

$$\beta = \frac{1 + a/2b + \sqrt{a^2/4b^2 - 1}}{a + 2b},$$

where $a \triangleq 2\lambda(1 - \frac{1}{\nu N}) + \frac{1}{\nu N}$, and $b \triangleq -\lambda(1 - \frac{1}{\nu N})$. \square

Proof: The proof is based on an inversion lemma of Toeplitz tridiagonal matrices given in [18, Theorem 3.3]. The details are omitted due to lack of space. \blacksquare

The above methodology can be readily applied to derive the decentralized mean-field control for non-linear AoI cost (arbitrary function of AoI satisfying the non-decreasing assumption) minimization problem [19]. Define the AoI cost function as $g(h)$ and consider the generate-at-will sources whereby each update drops the AoI to one. The following corollary gives the mean-field control based on Whittle's index and scheduling design of ETSU.

Corollary 3 (Mean-Field Control for Non-Linear AoI):

The index threshold with arbitrary AoI cost function $g(h)$, $\forall h \geq 1$ is given by

$$I_{\text{th}} = \sum_{h=1}^{H_{\text{th}}} (g(H_{\text{th}}) - g(h)) (1 - p_e),$$

where $H_{\text{th}} = \frac{(1-\nu)N}{1-p_e}$, the transmission error probability is p_e , and by setting ν we let the top ν -fraction of nodes to transmit with probability given by (8).

We have thus far addressed the two issues proposed at the beginning of the section, by applying Theorem 2 to decide the event importance and Theorem 3 based on mean-field approach to determine the importance-to-transmission-probability mapping of $\Psi(\cdot)$ in (6).

D. Performance Evaluations

Compared with the centralized optimal status update scheme, the sub-optimality of ETSU comes from three possible aspects, which will be investigated in order in this, and following sections. First, the Whittle's index policy (Theorem 2), in contrast to the MDP optimal solution, is strictly speaking sub-optimal. Additionally, although the index in Theorem 2 is precise, the index approximation in Corollary 1 with unreliable channels needs to be validated. This aspect is evaluated in Fig. 3(left), where it is observed that the approximate index with channel error, being derived assuming transmission failure probability approaches zero, is close to optimum with a wide range of p_e . The x -axis is $p_{e,1}$ and we set $p_{e,2} = 0.9$; we also let the error functions reflect the error sensitivities, i.e., $\delta_1(d) = d$ (error-tolerant) and $\delta_2(d) = e^d - 1$ (error-sensitive).

The second aspect of sub-optimality comes from the mean-field-based transmission probability mapping, which assumes the number of nodes is large. Fig. 3 (right) simulates scenarios with various numbers of nodes. Since it is impossible to obtain the MDP-based optimum with many nodes due to curse of dimensionality, the performance of ETSU is compared with centralized index policy by Theorem 2 which has been shown near-optimal based on Fig. 3(left). We adopt $\delta_n(d) = d$, $\forall n$, $\frac{t_{\text{slot}}}{t_c} = 10$, and $p_{e,n}$ is uniformly generated from $[0, 0.3]$. The transmission probability mapping function $\Psi(\cdot)$ is given by Corollary 2 and we let $\nu = 5/N$ such that there are about 5 nodes with indices above the threshold at each time slot. In general, ETSU achieves near-optimal performance with a moderate number of nodes. Furthermore, the separate approach is also simulated with AoI-based scheduling (with

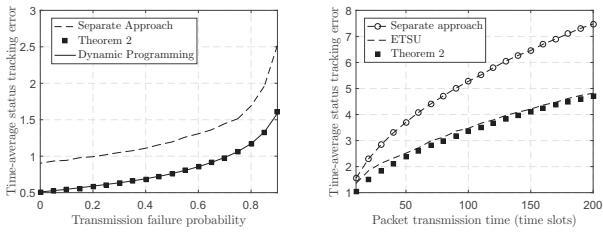


Fig. 3. Performance evaluations. (a) Theorem 2 is compared with the optimum by MDP; (b) ETSU is compared with the centralized index policy (no random access) given by Theorem 2.

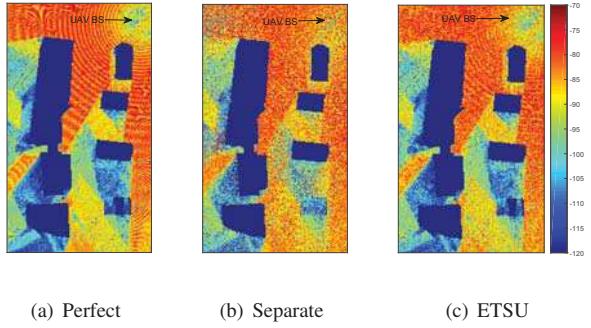


Fig. 4. CSI (in decibel-milliwatts) update based on ETSU and a separate approach [9] with a UAV BS traversing over the area.

channel error) according to index policy $I_n(h) = (1 - p_{e,n})h^2$ derived in [7]; its performance is outperformed by ETSU in all cases.

The third part is the modeling error of the information source nodes. We model them as random walk transitions, with several trivial extensions to, e.g., stay probability and asymmetrical walk discussed in Section IV-A. Nevertheless, the current model cannot encompass state-inhomogeneous transitions which render the index policy untractable. In this regard, we test the algorithms based on real-world information sources in Section V: CSI variations with an unmanned aerial vehicle (UAV) base station (BS) traversing an area.

V. CASE STUDY: ETSU FOR DYNAMIC CSI UPDATE

We study a distributed CSI update use case where there are a large number of terminals (1 m apart in a 140 m \times 250 m area and a total of $251 \times 141 \approx 3.5 \times 10^4$) reporting measured CSI to a sink node. The CSI variation is generated by the Wireless InSite® ray-tracing simulator and letting a UAV BS (transmitter with a height of 50 m) traverse the intended area from the bottom left corner to the top right with a velocity of 10 km/h. We assume each CSI update occupies one time slot of length 0.5 ms. The error function in ETSU is chosen as the squared error function, i.e., $\delta_n(s_n(t), \hat{s}_n(t)) = |s_n(t) - \hat{s}_n(t)|^2$. To be fair, we assume perfect CSI at the sink node at the beginning and evaluate the status tracking performance when the UAV flies past the area (at the top right corner). The CSI tracked at the sink node by a

separate approach and ETSU is shown in Fig. 4, together with the perfect CSI. We assume reliable channel in this case. The separate approach based on [9] implements sample-at-change sampling and scheduling that minimizes the average AoI of nodes, while neglecting the fact that the key status is not the one that is the stalest, but the one with, roughly speaking, the largest error. In contrast, ETSU, although assuming random-walk state transitions which does not match the CSI transitions, achieve evidently better CSI tracking performance.

VI. CONCLUSIONS

A unified sampling and scheduling approach is proposed for status update in multiaccess wireless networks, capturing the key status variation in contrast to the conventional separate approach which samples the information sources and then schedules the nodes based on minimizing the AoI. The Whittle's index methodology is generalized to characterize the status packet importance based on arbitrary status tracking error functions. A mean-field approach is applied to derive the decentralized implementation in closed-form. As a special case of our results, we describe the closed-form decentralized non-linear AoI (arbitrary cost function of AoI) minimization scheme. The proposed ETSU evidently outperforms the separate approach as shown by extensive simulation results, including a realistic case study where we adopt ray-tracing generated CSIs as information sources with a mobile transmitter.

ACKNOWLEDGMENT

This work is sponsored in part by the Nature Science Foundation of China (No. 61701275, 61628107, 91638204, 61571265, 61621091), the US NSF CNS-1816908, and ECCS-1610874.

APPENDIX A PROOF OF THEOREM 1

Define the cost-to-go function as $V_t(\mathcal{B}(t)) = \min_{\mathbf{u}(t), \mathbf{u}(t)^\top \mathbf{1} = 1} q_t(\mathcal{B}(t), \mathbf{u}(t))$, $q_t(\mathcal{B}(t), \mathbf{u}(t)) \triangleq \mathbb{E}[V_{t+1}(\mathcal{B}(t+1)) + g_t(\mathcal{B}(t), \mathbf{u}(t))]$, where $V_T(\mathcal{B}(T)) = \mathbb{E}[g_T(\mathcal{B}(T))]$ is the terminal cost, the expected immediate cost at time t is denoted by $\mathbb{E}[g_t(\mathcal{B}(t), \mathbf{u}(t))]$, denote $\mathbf{u}(t) \triangleq [u_1(t), \dots, u_N(t)]^\top \in \{0, 1\}^N$, and the time index is omitted for brevity. According to the MDP methodology [20], the optimal total cost for the T -horizon problem is $V^*(\mathcal{B}(0)) = V_0(\mathcal{B}(0))$, and the optimal dynamic policy is given by $\mathbf{U}^* = \{\mathbf{u}(0)^*, \dots, \mathbf{u}(T-1)^*\}$, where $\mathbf{u}(t)^*$ minimizes the cost-to-go for $t \in \{0, \dots, T-1\}$. Because symmetrical two-state Markov sources are considered, it is clear that the system state can be simplified as $\mathcal{B}(t) \triangleq \{\mathbb{1}_{\{s_1(t) \neq \hat{s}_1(t)\}}, \dots, \mathbb{1}_{\{s_N(t) \neq \hat{s}_N(t)\}}\}$, indicating whether the current statuses at source and destination are identical; we further denote $d_n(t) \triangleq \mathbb{1}_{\{s_n(t) \neq \hat{s}_n(t)\}}$.

The proof includes two steps: first, we prove that the policy in Theorem 1 is the *myopic* policy, i.e., it minimizes the expected immediate cost (this is straightforward and hence omitted); secondly, a *backwards induction* (on time t) based proof is given showing that the myopic policy is indeed the

optimal policy. Define the total cost from time t on given the current state and following the myopic policy in Theorem 1 as $W_t(\hat{\mathcal{B}}(t))$. Concretely, we will show the following statements are valid. Denote $\bar{W}_{n,m}^t(s, r) \triangleq W_t(c_1, \dots, d_n(t) = s, \dots, d_m(t) = r, \dots, c_N)$ and $d_l(t) = c_l, \forall l \neq n, m$.

I. $W_t(\hat{\mathcal{B}}(t))$ is the optimal total cost from t on.

II. For any $t \geq 0$ and $p_i \leq p_j$, $\bar{W}_{i,j}(0, 1) \leq \bar{W}_{i,j}(1, 0)$.

Let us first establish the induction basis. For $t = T$, hypothesis I is obviously true since the policy of Theorem 1 is the myopic policy and the last step only concerns with the immediate cost. Hypothesis II also holds with equality.

Suppose hypothesis I and II are valid from time $t+1$ to T , at time t denote the myopic solution as $u_{n^*}(t) = 1$ and zero otherwise, which indicates that $n^* = \arg \min_{n \in \chi(t)} p_n$ where $\chi(t)$ is the set of nodes with $d_n(t) = 1$. Then $\forall l \in \chi(t)$ with $p_l \geq p_{n^*}$, it follows from the cost-to-go function, after some manipulations, that

$$\begin{aligned} q_t(\hat{\mathcal{B}}(t), n^*) - q_t(\hat{\mathcal{B}}(t), l) &= \underbrace{\mathbb{E}[g_t(\mathcal{B}(t), n^*)] - \mathbb{E}[g_t(\mathcal{B}(t), l)]}_{\mathcal{M}_1} \\ &+ \underbrace{\sum_{c_n, n \neq l, n^*} \bar{p}_{n^*, l} \left(p_{n^*} - p_l \right) (\bar{W}_{n^*, l}^{t+1}(1, 1) - \bar{W}_{n^*, l}^{t+1}(0, 0))}_{\mathcal{M}_2} \\ &+ \underbrace{(1 - p_{n^*} - p_l) (\bar{W}_{n^*, l}^{t+1}(0, 1) - \bar{W}_{n^*, l}^{t+1}(1, 0))}_{\mathcal{M}_3}, \end{aligned}$$

where $\bar{p}_{n^*, l}$ denotes the transition probability from $\hat{\mathcal{B}}(t) \setminus \{d_{n^*}, d_l\}$ to $\{c_1, \dots, c_N\} \setminus \{c_{n^*}, c_l\}$, i.e., the transition probability to a given state excluding node n^* and c_l (without loss of generality, assume $n^* < l$). Note that it is unnecessary to consider nodes with index $m \notin \chi(t)$ since their states are correct. Based on the definition of myopic policy, $\mathcal{M}_1 \leq 0$. It is straightforward that $\bar{W}_{n^*, l}^{t+1}(1, 1) - \bar{W}_{n^*, l}^{t+1}(0, 0) \geq 0$ since $\bar{W}_{n^*, l}^{t+1}(1, 1)$ includes a state with two tracking errors and $\bar{W}_{n^*, l}^{t+1}(0, 0)$ a state that corrects both; hence, combining with the $p_l \geq p_{n^*}$, we obtain $\mathcal{M}_2 \leq 0$. Since $\forall n, p_n \leq 0.5$, and based on hypothesis II, we have $\mathcal{M}_3 \leq 0$. Therefore, we arrive at the conclusion that the myopic solution at time t , i.e., n^* , is also the optimal action. With this, we have proved hypothesis I. Next, we can prove hypothesis II based on induction, by analyzing three cases on the myopic solution. The details are omitted due to lack of space. With this we conclude the induction proof.

APPENDIX B PROOF OF THEOREM 2

Following the technique used in, e.g., [9], [10], [16], we first assume the optimal policy has a threshold-type structure and solve the Bellman equations based upon it; after obtaining the optimal solution, consistency with this assumption is checked to conclude the proof.

Assume the optimal policy to solve (4) is update when $d \geq D$, and idle when $0 \leq d < D$. Denote the upper term in the minimization in (4) as $\gamma_0(d)$ and the lower term $\gamma_1(d)$. Then for $d \geq D$, the optimal action is to update, and hence

$$f(d) = m + \frac{f(1)}{2} - \hat{J}^*, \quad d \geq D. \quad (9)$$

Likewise, the optimal action is to idle otherwise, and

$$f(d) + \hat{J}^* = \delta(d) + \frac{1}{2}f(d+1) + \frac{1}{2}f(d-1), \quad 0 < d < D,$$

and the case with $d = 0$ is specially treated with outcomes $f(1) = 2\hat{J}^*$. Together with $f(0) = 0$ we can obtain the formula for $f(d)$ when $0 \leq d \leq D$ based on the differential equation of (10).

$$f(d) = d(d+1)\hat{J}^* - 2 \sum_{i=1}^{d-1} \sum_{j=1}^i \delta(j), \quad 0 \leq d \leq D. \quad (10)$$

With the following three equations we can solve for the Whittle's index in (5).

$$f(D) = m = D(D+1)\hat{J}^* - 2 \sum_{i=1}^{D-1} \sum_{j=1}^i \delta(j), \quad \gamma_0(D) = \gamma_1(D). \quad (11)$$

The first equation follows from (9) and (10) with $d = D$. The second one is based on the fact that the auxiliary cost m should be the minimum cost that make the update decision equally beneficial give the current state. Therefore, when $d = D$, two options should be equally valuable.

Consistency with the threshold-type assumption can be checked to be satisfied with the derived optimal policy, whose details are omitted for brevity. The indexability can be verified by checking that when $m = 0$, the threshold is zero but update and idle are equally beneficial and hence one should idle; when $m \rightarrow \infty$, the threshold also goes to infinity. Additionally, the monotonicity can be proved by checking

$$I_{\text{RW},n}(d) - I_{\text{RW},n}(d-1) = d\delta_n(d) - \sum_{i=1}^{d-1} \delta_n(i) \stackrel{(a)}{\geq} \sum_{i=1}^d \delta_n(i) \geq 0.$$

The inequality (a) follows from the non-decreasing property of the error functions, i.e., $\delta_n(d) \geq \delta_n(i), \forall i \leq d$. With this, we conclude the proof.

APPENDIX C PROOF OF COROLLARY 1

The cost-to-go function is changed to

$$f(d) + \hat{J}^* = \min \left\{ \begin{array}{l} \gamma_0(d), \\ m + (1 - p_e)f(1)/2 + p_e\gamma_0(d) \end{array} \right\} \quad (12)$$

Following from (11), we can obtain $f(D) = \frac{m}{1-p_e}$. For $d \geq D$, based on the lower term in (12),

$$\frac{p_e}{2}f(d+1) - f(d) + \frac{p_e}{2}f(d-1) = p_e \left(\hat{J}^* - \delta(d) \right) - m.$$

Solving this equation recursively gives us

$$p_e f(D+1) - \left(1 - \sqrt{1 - p_e^2} \right) f(D) = \frac{p_e(m - p_e \hat{J}^*)}{1 - p_e + \sqrt{1 - p_e^2}}.$$

Analyzing this in the regime $p_e \rightarrow 0$, we can obtain

$$f(D+1) \xrightarrow{p_e \rightarrow 0} \frac{m}{1 - p_e} + \mathcal{O}(1). \quad (13)$$

Plugging (13) into (12), we obtain a similar equation as in (11), except that m is replaced with $\frac{m}{1-p_e}$, and hence the corresponding index follows immediately. The indexability also follows because $\frac{m}{1-p_e}$, compared with m , does not affect the monotonicity or index values at zero and infinity.

APPENDIX D PROOF OF THEOREM 3

The state transition matrix with parameter D_{th} (denote $\gamma \triangleq 1 - \lambda - \mu$) is $\mathbf{P}_{D_{\text{th}}} \triangleq \left[\begin{array}{c|c} \bar{\mathbf{P}} & \tilde{\mathbf{P}} \\ \hline \check{\mathbf{P}} & \hat{\mathbf{P}} \end{array} \right] =$

$$\left[\begin{array}{cc|cc} 1 - \lambda & \lambda & & \\ \mu & \gamma & \lambda & \\ \ddots & \ddots & \ddots & \\ & \mu & \gamma & \lambda \\ \hline \varepsilon & & \mu(1 - \varepsilon) & \gamma(1 - \varepsilon) & \lambda(1 - \varepsilon) \\ \vdots & & & \ddots & \ddots & \ddots \end{array} \right],$$

where $\bar{\mathbf{P}}$ is a $(D_{\text{th}} - 1)$ -dimensional square matrix. Denote the stationary distribution as $\pi \triangleq [\bar{\pi}, \hat{\pi}]$ correspondingly. Solving for the first $D_{\text{th}} - 1$ states yields (first assuming $\lambda \neq \mu$ and then generalized)

$$\pi_d = \left(\frac{\lambda}{\mu} \right)^d \left(\pi_0 - \frac{\varepsilon\sigma}{\lambda - \mu} \right) + \frac{\varepsilon\sigma}{\lambda - \mu}, \quad 0 \leq d < D_{\text{th}}, \quad (14)$$

where $\sigma \triangleq \hat{\pi} \mathbf{1}$. Solving for the remaining states yields

$$\hat{\pi} = \lambda \pi_{D_{\text{th}}-1} \mathbf{e}_1^\top (\mathbf{I} - \hat{\mathbf{P}})^{-1}, \quad (15)$$

where $\mathbf{e}_1 \triangleq [1, 0, 0, \dots]^\top$, and that $\mathbf{I} - \hat{\mathbf{P}}$ is invertable, since $\hat{\mathbf{P}}$ has eigenvalues smaller than one ($\varepsilon > 0$). Plugging (14) into (15) yields

$$\sigma = \lambda \beta \left[\left(\frac{\lambda}{\mu} \right)^{D_{\text{th}}-1} \left(\pi_0 - \frac{\varepsilon\sigma}{\lambda - \mu} \right) + \frac{\varepsilon\sigma}{\lambda - \mu} \right], \quad (16)$$

where $\beta \triangleq \mathbf{e}_1^\top (\mathbf{I} - \hat{\mathbf{P}})^{-1} \mathbf{1}$. Then summing over all states yields $\sum_{i=0}^{\infty} \pi_i = 1$, combining with (14) and (16) we obtain

$$\frac{1 - \left(\frac{\lambda}{\mu} \right)^{D_{\text{th}}}}{1 - \frac{\lambda}{\mu}} \left(\pi_0 - \frac{\varepsilon\sigma}{\lambda - \mu} \right) + D_{\text{th}} \frac{\varepsilon\sigma}{\lambda - \mu} + \sigma = 1,$$

and hence $\pi_0 = \left(\frac{\mu}{\lambda} \right)^{D_{\text{th}}-1} \left(\frac{1}{\lambda\beta} - \frac{\varepsilon}{\lambda - \mu} \right) + \frac{\varepsilon\sigma}{\lambda - \mu}$. It follows that (7) gives the solution of D_{th} . Next we will show that σ is monotonically non-increasing with D_{th} and hence $\sigma = \nu$ yields a unique solution (note that when $D_{\text{th}} = 0$, $\sigma = 1$; when $D_{\text{th}} \rightarrow \infty$, $\sigma \rightarrow 0$).

$$\begin{aligned} \frac{d\sigma^{-1}}{dD_{\text{th}}} &= \frac{-\log \frac{\lambda}{\mu}}{\lambda\beta \left(1 - \frac{\lambda}{\mu} \right)} \left(\frac{\lambda}{\mu} \right)^{1-D_{\text{th}}} \\ &\quad + \left(1 + \frac{\log \frac{\lambda}{\mu}}{1 - \frac{\lambda}{\mu}} \left(\frac{\lambda}{\mu} \right)^{1-D_{\text{th}}} \right) \frac{\varepsilon}{\lambda - \mu} \\ &\stackrel{(a)}{\geq} \left(1 - \left(\frac{\lambda}{\mu} \right)^{1-D_{\text{th}}} \right) \frac{\varepsilon}{\lambda - \mu} \geq 0, \end{aligned}$$

where the inequality (a) stems from $\log(x) \leq x - 1$, $\forall x > 0$. Thus far, we have shown that there is a unique solution to the threshold equation of (7), and thereby the index threshold can be derived accordingly based on Theorem 2. Since the top ν -fraction of nodes are allowed to compete for a transmission slot, which corresponds to approximately νN nodes, the transmission probability, i.e., p_{tx} , can be derived based on the p-CSMA results in, e.g., [21, Theorem 1].

REFERENCES

- [1] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?", in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, pp. 2731–2735, Mar 2012.
- [2] R. D. Yates, "The age of information in networks: Moments, distributions, and sampling," *arXiv preprint arXiv:1806.03487*, 2018.
- [3] C. Kam, S. Kompella, G. Nguyen, J. Wieselthiery, and A. Ephremides, "Towards an effective age of information: Remote estimation of a Markov source," in *Proc. IEEE INFOCOM, 1st Workshop on Age-of-Information*, pp. 1–9, Apr 2018.
- [4] Y. Sun, Y. Polyaniskiy, and E. Uysal-Biyikoglu, "Remote estimation of the Wiener process over a channel with random delay," in *Proc. IEEE Int'l Symp. Info. Theory*, pp. 321–325, Jun 2017.
- [5] Y. Sun, E. Uysal-Biyikoglu, R. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, pp. 1–9, Apr 2016.
- [6] J. Wu, Q. S. Jia, K. H. Johansson, and L. Shi, "Event-based sensor data scheduling: Trade-off between communication rate and estimation quality," *IEEE Trans. Autom. Control*, vol. 58, pp. 1041–1046, Apr 2013.
- [7] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," *arXiv preprint arXiv:1801.01803*, 2018.
- [8] Z. Jiang, B. Krishnamachari, X. Zheng, S. Zhou, and Z. Niu, "Decentralized status update for age-of-information optimization in wireless multiaccess channels," in *Proc. IEEE Int'l Symp. Info. Theory*, Jun 2018.
- [9] Z. Jiang, B. Krishnamachari, S. Zhou, and Z. Niu, "Can decentralized status update achieve universally near-optimal age-of-information in wireless multiaccess channels?", in *Proc. Int'l Teletraffic Congress*, Sep 2018.
- [10] I. Kadota, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Minimizing the age of information in broadcast wireless networks," in *Proc. Annu. Allerton Conf. Commun., Control, Comput.*, pp. 844–851, Sep 2016.
- [11] Y.-P. Hsu, "Age of information: Whittle index for scheduling stochastic arrivals," in *Proc. IEEE Int'l Symp. Info. Theory*, Jun 2018.
- [12] R. Talak, S. Karaman, and E. Modiano, "Optimizing information freshness in wireless networks under general interference constraints," in *Proc. ACM Int'l Symp. Mobile Ad Hoc Netw. Comput.*, pp. 61–70, Jun 2018.
- [13] F. V. Lu and R. F. Serfozo, "M/M/1 queueing decision processes with monotone hysteretic optimal policies," *Oper. Res.*, vol. 32, no. 5, pp. 1116–1132, 1984.
- [14] R. R. Weber and G. Weiss, "On an index policy for restless bandits," *Journal of Applied Probability*, vol. 27, no. 3, p. 637648, 1990.
- [15] R. G. Gallager, *Discrete stochastic processes*, vol. 321. Springer Science & Business Media, 2012.
- [16] R. Singh, X. Guo, and P. R. Kumar, "Index policies for optimal mean-variance trade-off of inter-delivery times in real-time sensor networks," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, pp. 505–512, Apr 2015.
- [17] J.-M. Lasry and P.-L. Lions, "Mean field games," *Japan. J. Math.*, vol. 2, pp. 229–260, Mar 2007.
- [18] U. Kamps, "Chebyshev polynomials and least squares estimation based on one-dependent random variables," *Linear Algebra and its Applications*, vol. 112, pp. 217 – 230, 1989.
- [19] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, "Age and value of information: Non-linear age case," in *Proc. IEEE Int'l Symp. Info. Theory*, pp. 326–330, Jun 2017.
- [20] D. P. Bertsekas, *Dynamic programming and optimal control*, vol. 1. Athena scientific Belmont, MA, 1995.
- [21] Y. Gai, S. Ganesan, and B. Krishnamachari, "The saturation throughput region of p-persistent CSMA," in *Information Theory and Applications Workshop*, pp. 1–4, Feb 2011.