

Timely Remote Estimation with Memory at the Receiver

Sirin Chakraborty

Dept. of ECE

Auburn University

Auburn, AL, USA

szc0260@auburn.edu

Yin Sun

Dept. of ECE

Auburn University

Auburn, AL, USA

yzs0078@auburn.edu

Abstract—In this study, we consider a remote estimation system that estimates a time-varying target based on sensor data transmitted over wireless channel. Due to transmission errors, some data packets fail to reach the receiver. To mitigate this, the receiver uses a buffer to store recently received data packets, which allows for more accurate estimation from the incomplete received data. Our research focuses on optimizing the transmission scheduling policy to minimize the estimation error, which is quantified as a function of the age of information vector associated with the buffered packets. Our results show that maintaining a buffer at the receiver results in better estimation performance for non-Markovian sources.




Fig. 1. System Model

I. INTRODUCTION

Timely status updates from sensors is a cornerstone for networked intelligent systems that rely on live data to make estimations and real-time decisions. It enables these systems to provide timely estimations, leading to intelligent and proactive actions. For instance, autonomous vehicles depend on real-time state estimations to make safety-critical decisions, ensuring accident avoidance and smoother traffic integration. Similarly, remote healthcare systems use timely inference to monitor vital signs, allowing quick responses to health emergencies of remote patients. In industrial IoT systems, real-time fault detection ensures operational efficiency by addressing malfunctions before they escalate into significant problems. But due to data loss, transmission error, transmission delay the status updates are not always fresh, which significantly affects the accuracy of timely predictions. To evaluate the freshness of data updates, we use the *Age of Information (AoI)* metric. Traditionally, AoI is defined as the time difference between the current time t and the generation time $U(t)$ of the freshest packet delivered to the receiver, such that $\Delta(t) = t - U(t)$. A smaller AoI indicates the presence of more recent information at the receiver.

In this paper, we consider a discrete-time remote estimation system for non-Markovian sources that consists of multiple sensors transmitting observations to a common receiver, as shown in Figure 1. Due to channel sharing among sensors and potential transmission errors, the data at the receiver may become outdated. To address the challenges posed by data staleness and the non-Markovian nature of the sources, we This work was supported in part by the NSF grant CNS-2239677.

propose a buffer-based remote estimation model. In this model, the receiver maintains a buffer of size b for each sensor n , storing the b most recently received data packets. After each successful packet delivery from sensor n , the newly received data packet is stored in the buffer, while the oldest packet is discarded to maintain the fixed buffer size b . This setup extends the conventional concept of AoI to an AoI vector. Specifically, let $\Delta_n(t) = (\Delta_{n,1}(t), \Delta_{n,2}(t), \dots, \Delta_{n,b}(t))$ denote the AoI vector of the buffered packets received from sensor n , where $\Delta_{n,i}(t)$ represents the age of the i -th most recently delivered packet from sensor n . Therefore, the first element $\Delta_{n,1}(t)$ corresponds to the AoI of the freshest packet, while the subsequent elements capture the AoI values of older packets stored in the buffer. When $b = 1$, the model reduces to the conventional remote estimation framework used for Markovian sources; see, e.g., [1]–[5]. The technical contributions of this paper are summarized as follows:

- We formulate a transmission scheduling problem for buffer-based remote estimation, aiming to minimize the average estimation error of the sensors. At each time t , the estimation error for sensor n is modeled as a function of its AoI vector $\Delta_n(t)$. This framework ensures accurate real-time estimation of non-Markovian sources, by dynamically prioritizing sensors for transmission based on their AoI vectors and their impact on reducing the overall system error.
- The transmission scheduling problem is cast as a Restless Multi-Armed Bandit (RMAB) and solved using Lagrangian

relaxation and dual decomposition. This method decomposes the relaxed problem into smaller sub-problems, which can be efficiently solved using dynamic programming. Leveraging the solution to the relaxed problem, a Maximum Gain First (MGF) scheduling policy is developed to address the original RMAB problem. Notably, this method eliminates the need to satisfy the indexability condition, which is typically required for Whittle index-based solutions.

- Our numerical results show that maintaining a buffer at the receiver results in better estimation for nonMarkovian sources. To the best of our knowledge, this is the first work to study Aol-vector-based scheduling for remote estimation with a receiver-side buffer.

II. RELATED WORKS

The Age of Information (Aol) has become a key metric for quantifying data freshness in networked systems since its introduction in [6]. Early research focused on optimizing average and peak Aol within communication networks [7]–[9]. More recent studies have applied Aol to enhance real-time applications, including remote inference [10], [11], edge computing [12], and control systems [3], [7], making it a central optimization tool. However, Aol focuses solely on timeliness and ignores the relevance of updates. To address this, complementary metrics such as Age of Incorrect Information (AolI) [13], Age of Synchronization (AoS) [14], and Value of Information (Vol) [15] have been developed. Additionally, [16] introduces the concept of credibility in real-time wireless networks, focusing on minimizing the system-wide Loss-ofCredibility (LoC) by ensuring timely packet deliveries. These metrics extend Aol by incorporating aspects like estimation error, context dependency, and the performance benefits of transmitting specific updates.

Previous works have extensively applied Kalman filters in remote estimation systems to address packet loss, delays, and communication constraints. [17] examined the stability of Kalman filtering over packet-erasure channels, identifying conditions for bounded estimation error despite intermittent observations. [18] analyzed Kalman filtering under random delays and packet drops, demonstrating that stability thresholds can be independent of delays under certain conditions. However, other studies, such as [1] and [4], highlight the interplay between Aol and remote estimation by designing optimal sampling strategies to minimize Aol-based estimation error. Building on this, our work incorporates the Aol as a vector into the estimation and scheduling process, enhancing accuracy under network constraints.

The optimization of both linear and non-linear Aol functions for multiple source scheduling can be framed as a restless multi-armed bandit (RMAB) problem, as explored in several works [11], [19]–[23]. Whittle’s pioneering work [24] introduced an index-based policy to solve RMAB problems with

binary actions. An extension of the Whittle index policy for handling multiple actions was presented in [25], though it requires satisfying a complex indexability condition. In addition to Whittle index-based policies, which rely on satisfying an indexability condition, non-indexable scheduling policies have also been explored in studies such as [20], [22], [23], [26]. Due to the complex buffer mechanism and channel fading conditions, Whittle index theory could not be applied to establish indexability for our multi-source scheduling problem. To address this, we adopt the “Net-gain Maximization” policy from our recent work [22], [23] and rename it as the “Maximum Gain First” policy. This policy does not rely on the indexability condition. It was also called the Optimal Lagrange Index policy in [27], LP-Index policy in [28], and Gain Index policy in [26].

III. MODEL AND FORMULATION

A. System Model

We consider the discrete-time remote estimation system depicted in Figure 1, where N sensors transmit status updates over M shared wireless channels to a common receiver. At each time slot $t \in \{0, 1, 2, \dots\}$, the receiver estimates a timevarying target $Y_{n,t} \in \mathcal{Y}$ using status update packets received from sensor n up to time t . A scheduler decides which sources to select for data transmission. The scheduler’s decision for sensor n at time t is represented by an indicator function

$$u_n(t) = \begin{cases} 1, & \text{if decides to transmit,} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

In response to the scheduler’s decision, each sensor n submits a time-stamped status update packet $(X_{n,t}, t)$ of its observation $X_{n,t} \in \mathcal{X}$ to a channel. However, due to transmission errors over wireless fading channels, some packets are lost during the transmissions. This transmission successful events can be expressed using indicator functions $c_n(t) \in \{0, 1\}$ that are i.i.d. across time and sensors. The transmission successful probability is denoted by $p_n = \Pr\{c_n(t) = 1\}$. The freshest information the receiver gets at time t is $X_{n,t-\Delta_n(t)}$, which was generated $\Delta_n(t)$ time slots ago. This time difference $\Delta_n(t)$ is the Aol of sensor n . If an update is sent at beginning of time slot t , it gets delivered at beginning of time slot $t+1$, then the Aol evolution of the n -th source follows

$$\Delta_n(t+1) = \begin{cases} 1, & \text{if } c_n(t)u_n(t) = 1, \\ \Delta_n(t) + 1, & \text{otherwise.} \end{cases} \quad (2)$$

We assume that at a given time, each channel can serve only one of the N sources and after every successful transmission, the receiver sends an error free acknowledgment back to the transmitter.

The receiver consists of N buffers, each of a size b , connected to the N estimators corresponding to the N sensors. Each buffer n can store at most b packets previously received from sensor n . The content of buffer n is $\{(X_{n,t-\Delta_{n,1}(t)}, \Delta_{n,1}(t)), (X_{n,t-\Delta_{n,2}(t)}, \Delta_{n,2}(t)), \dots, (X_{n,t-\Delta_{n,b}(t)}, \Delta_{n,b}(t))\}$, where the i -th packet consists of a stale observation $X_{n,t-\Delta_{n,i}(t)}$ from sensor n , generated $\Delta_{n,i}(t)$ time ago and its corresponding Aol $\Delta_{n,i}(t)$ such that $i \in \{1, 2, \dots, b\}$. At time $t+1$, if the latest transmission is successful, i.e. $u_n(t) = 1$, the buffer gets updated with the freshest packet $(X_{n,t}, 1)$ and discards the oldest data packet in order to maintain the fixed size of b . In contrast, if the transmission fails, i.e. $u_n(t) = 0$, the signal values in the buffer contents stay the same, but their age values get one time slot older, mathematically represented as $(X_{n,t-\Delta_{n,i}(t)}, \Delta_{n,i}(t)+1)$. The transition of the Aol vector is expressed as

$$\Delta_n(t+1) = \begin{cases} (1, \Delta_{n,1}(t) + 1, \dots, \Delta_{n,b-1}(t) + 1), & \text{if } c_n(t)u_n(t) = 1, (\Delta_{n,1}(t) + 1, \Delta_{n,2}(t) + 1, \dots, \Delta_{n,b}(t) + 1), & \text{otherwise.} \end{cases} \quad (3)$$

Estimator n takes the contents of buffer n as a feature vector of length b and generates an output $a = \phi_n(\mathbf{X}_{n,t-\Delta_n(t)}, \Delta_n(t)) \in A$, where $\mathbf{X}_{n,t-\Delta_n(t)} =$

$(X_{n,t-\Delta_{n,1}(t)}, X_{n,t-\Delta_{n,2}(t)}, \dots, X_{n,t-\Delta_{n,b}(t)})$ is the feature vector of length b , $\Delta_n(t) = (\Delta_{n,1}(t), \Delta_{n,2}(t), \dots, \Delta_{n,b}(t))$ is the age vector of the packets in the feature vector, and $\phi : (X^b \times Z^{+b}) \rightarrow A$ is the estimation function. We consider an estimator whose estimation performance for sensor n at time slot t is measured by a loss function $L_n : Y \times A \rightarrow \mathbb{R}$, where $L_n(y, a)$ indicates the incurred loss if the output a is used for estimation when $Y_{n,t} = y$. The estimation problem is formulated as $\text{err}_n(\Delta_n(t))$

$$= \min_{\phi \in \Phi} \mathbb{E} [L_n(Y_{n,t}, \phi_n(\mathbf{X}_{n,t-\Delta_n(t)}, \Delta_n(t))) | \Delta_n(t)]. \quad (4)$$

The loss function L_n is defined based on the objective of the remote estimation system [21]. For instance, in a neural network designed for minimum mean-squared error estimation, the loss function is given by $L_n(y, \hat{y}) = (y - \hat{y})^2$, where the action $a = \hat{y}$ represents the estimate of the target $Y_{n,t} = y$. In the case of softmax regression, which is a neural network-based approach for maximum likelihood classification, the action $a = Q_Y$ corresponds to a probability distribution over $Y_{n,t}$, and the loss function $L(y, Q_Y) = -\log(Q_Y(y))$ represents the negative log-likelihood of the target value $Y_{n,t} = y$.

B. Scheduling Policy and Problem Formulation

A scheduling policy is denoted by $\pi = (\pi_n)_{n=1}^N$, where $\pi_n = (u_n(0), u_n(1), \dots)$. Let Π be the set of all possible causal scheduling policies, in which every decision $u_n(t)$ is made using the current and history information available at the scheduler.

As our system consists of M different channels, the condition $\sum_{n=1}^N u_n(t) \leq M$ is required to hold for all t .

Our goal is to find the optimal scheduling policy that minimizes the discounted sum of the expected estimation errors among all the N sources over an infinite time-horizon. So the optimization problem can be formulated as

$$\inf_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \sum_{n=1}^N \sum_{t=0}^{\infty} \mathbb{E}_{\pi} [\gamma^t \text{err}_n(\Delta_n(t))] \quad (5) \text{ s.t. } X$$

$$u_n(t) \leq M, u_n(t) \in \{0, 1\}, t = 0, 1, \dots, \quad (6) \quad n=1$$

where $\gamma \in [0, 1]$ is the discount factor and $\mathbb{E}_{\pi}[\cdot]$ is the expectation under policy π .

IV. RESTLESS MULTI-ARMED BANDIT SOLUTION

Problem (5)-(6) is an RMAB problem, considering each source n as an arm with $\Delta_n(t)$ as the state. To solve RMAB problems, the Whittle index policy in [24] is the most commonly used method. However, a key challenge in applying the Whittle index is satisfying the problem's indexability condition. Due to the complex buffer state transitions, proving indexability is challenging in our problem. Therefore, we use a different algorithm, following [22], which does not require satisfying the indexability condition.

A. Lagrangian Relaxation and Dual Decomposition

To address (5)-(6), we first relax the per-time-slot channel constraints (6) into a discounted time-summation form, as expressed in constraint (8) of the following relaxed problem:

$$\begin{aligned} & \inf_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \sum_{n=1}^N \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{\pi} [\text{err}_n(\Delta_n(t))], \quad (7) \\ & \limsup_{T \rightarrow \infty} \sum_{n=1}^N \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{\pi} [u_n(t)] \leq \frac{M}{1-\gamma} \quad (8) \\ & \text{s.t.} \end{aligned}$$

To solve this relaxed problem, we apply the Lagrange dual decomposition method [24], using a Lagrange multiplier $\lambda \geq 0$. The dual problem is formulated as

$$\lambda^* = \arg\max_{\lambda \geq 0} q(\lambda), \quad (9)$$

where the dual function $q(\lambda)$ is defined as

$$q(\lambda) = \inf_{\pi \in \Pi} \limsup_{T \rightarrow \infty} \sum_{n=1}^N \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{\pi} [\text{err}_n(\Delta_n(t)) + \lambda u_n(t)] - \frac{\lambda M}{1-\gamma}. \quad (10)$$

Given λ , problem (10) can be decomposed into N subproblems, and the sub-problem for sensor n is formulated as

$$\inf_{\pi_n \in \Pi_n} \limsup_{T \rightarrow \infty} \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{\pi_n} [\text{err}_n(\Delta_n(t)) + \lambda u_n(t)], \quad (11)$$

where Π_n is the set of all causal scheduling policies $\pi_n = (u_n(0), u_n(1), \dots)$ of sensor n . The Lagrange multiplier $\lambda \geq 0$ can be considered as a transmission cost, which sensor n has to pay for using a channel resource.

B. Optimal Solution for the Relaxed Problem (7)-(8)

The sub-problem (11) for each sensor n is a discounted infinite-horizon MDP. Suppose that the current state $\Delta_n(t)$ and action $u_n(t)$ are denoted as s and a , respectively. Then, the Bellman optimality equation for the MDP (11) is

$$V_{n,\lambda}^*(s) = \min_{a \in \{0,1\}} Q_{n,\lambda}^*(s,a), \quad (12)$$

$$Q_{n,\lambda}^*(s,a) = \text{err}_n(s) + \lambda a + \gamma \sum_{s'} p_{ss'}^a V_{n,\lambda}^*(s'), \quad (13)$$

where $V_{n,\lambda}^*(\cdot)$ and $Q_{n,\lambda}^*(\cdot, \cdot)$ are the optimal value function and the optimal action-value function of the MDP (11), respectively. The action-value function $Q_{n,\lambda}^*(s,a)$ is also known as Algorithm 1 Maximum Gain First Scheduling Policy

- 1: Input: Optimal Lagrange multiplier λ^* obtained by solving (7)-(8) using (12)-(14).
- 2: for each time step $t = 1, 2, \dots$ do 3:
- for each sensor $n = 1, 2, \dots, N$ do 4:
- Update state $\Delta_n(t)$ based on (2).
- 5: Update the gain $\alpha_{n,\lambda^*}(\Delta_n(t))$ based on (15).
- 6: end for
- 7: Select at most M sensors with the highest nonnegative gains $\alpha_{n,\lambda^*}(\Delta_n(t))$.
- 8: end for

the Q-function or Q-value. Given any $\lambda \geq 0$, we use a dynamic programming algorithm to solve (12)-(13).

Next, we use the stochastic sub-gradient ascent method to solve (9) and obtain the optimal dual variable λ^* . At each iteration k , the dual variable is updated as follows:

$$\lambda(k+1) = \lambda(k) + \frac{\beta}{k} \left[\sum_{n=1}^N \sum_{t=0}^{\infty} \gamma^t u_{n,\lambda(k)}^*(t) - \frac{M}{1-\gamma} \right], \quad (14)$$

where $\beta > 0$ is a constant, and $u_{n,\lambda(k)}^*(t)$ represents whether sensor n is scheduled at time slot t in the optimal solution to the sub-problem (11) when $\lambda = \lambda(k)$.

C. Maximum Gain First Policy for Problem (5)-(6)

Now we present the Maximum Gain First (MGF) scheduling policy, which is a feasible solution to the original problem (5)-(6). Using the optimal dual variable λ^* of the relaxed problem (7)-(8), the gain $\alpha_{n,\lambda^*}(\Delta_n(t))$ of scheduling sensor n at time slot

t is defined as the difference of Q-values between not transmitting and transmitting [22], [23]:

$$\alpha_{n,\lambda^*}(\Delta_n(t)) = Q_{n,\lambda^*}^*(\Delta_n(t), 0) - Q_{n,\lambda^*}^*(\Delta_n(t), 1). \quad (15)$$

At each time slot t , the MGF policy selects no more than M sensors that have the highest non-negative gains $\alpha_{n,\lambda^*}(\Delta_n(t))$, as described in Algorithm 1. We note that the MGF policy is also the solution to the following problem:

$$\begin{aligned} \max_{u_n(t), n=1, \dots, N} \quad & \sum_{n=1}^N \alpha_{n,\lambda^*}(\Delta_n(t)) u_n(t) \\ \text{s.t.} \quad & \sum_{n=1}^N u_n(t) \leq M, \quad u_n(t) \in \{0, 1\} \end{aligned} \quad (16)$$

V. EVALUATION

We evaluate the performance of buffer-based remote estimation using the following fourth-order autoregressive system:

$X_{n,t} = 0.1X_{n,t-1} + 0.8X_{n,t-4} + W_{n,t}$ (18) where $W_{n,t} \in \mathbb{R}$ represents i.i.d. Gaussian noise with zero mean and unit variance. The receiver uses a Kalman filter to estimate the current system state $X_{n,t}$, based on previously received data packets that are stored in the buffer. The mean squared error (MSE) of the Kalman filter is a function of AoI.




Fig. 2. AoI $\Delta_1(t)$ vs estimation error with fixed $\Delta_2(t) = 3$.

Figure 2 depicts the MSE as a function of AoI under two scenarios: (i) buffer size $b = 1$, where the MSE is represented as $\text{err}(\Delta_1(t))$, and (ii) buffer size $b = 2$, where the MSE is represented as $\text{err}(\Delta_1(t), \Delta_2(t))$. In both cases, the MSE is found to be a non-monotonic function of AoI, consistent with the results reported in [21]. One can observe that storing two data packets in the buffer reduces the MSE compared to the single-packet case in traditional remote estimation framework.

We evaluate the performance of our proposed scheduling algorithm in minimizing the overall estimation error. The

system consists of two sensors ($N = 2$) and one communication channel ($M = 1$). To assess the impact of buffer size at the receiver, we analyze two scenarios: (i) a buffer that holds a single data packet for each sensor ($b = 1$) and (ii) a buffer storing two data packets for each sensor ($b = 2$). Figure 3 presents the average estimation error as a function of the transmission probability p_n , assuming equal transmission success probabilities for both sensors, i.e., $p_1 = p_2$. The results demonstrate that, utilizing a buffer of size 2 reduces the optimized average estimation error compared to a buffer of size 1.

VI. CONCLUSION

In this paper, we design a remote estimation system with buffers at the receiver storing historically received data. Our results show that maintaining a buffer of multiple packets reduces the estimation error for non-Markovian sources.

VII. ACKNOWLEDGEMENT

The authors appreciate the encouragement from I-Hong Hou, Roy D. Yates, and Shaolei Ren to consider incorporating a buffer at the receiver.

REFERENCES

- [1] Y. Sun, Y. Polyanskiy, and E. Uysal-Biyikoglu, "Remote estimation of the wiener process over a channel with random delay," in *IEEE ISIT*, 2017, pp. 321–325.
- [2] Y. Sun and B. Cyr, "Sampling for data freshness optimization: Non-linear age functions," *Journal of Communications and Networks*, vol. 21, no. 3, pp. 204–219, 2019.
- [3] M. Klugel, M. H. Mamduhi, S. Hirche, and W. Kellerer, "AoI-penalty minimization for networked control systems with packet loss," in *IEEE INFOCOM Aol Workshop*, 2019, pp. 189–196.
- [4] T. Z. Ornee and Y. Sun, "Sampling and remote estimation for the Ornstein-Uhlenbeck process through queues: Age of




Fig. 3. Average estimation error vs transmission probability p_n .

- information and beyond," *IEEE/ACM Transactions on Networking*, vol. 29, no. 5, pp. 1962–1975, 2021.
- [5] T. Z. Ornee and Y. Sun, "A Whittle index policy for the remote estimation of multiple continuous Gauss-Markov processes over parallel channels," in *ACM MobiHoc*, 2023, pp. 91–100.
- [6] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" In *2012 Proceedings IEEE INFOCOM*, 2012, pp. 2731–2735.
- [7] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, pp. 1183–1210, 2021.
- [8] R. D. Yates, "Lazy is timely: Status updates by an energy harvesting source," in *IEEE ISIT*, 2015, pp. 3008–3012.
- [9] 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 Transactions on Information Theory*, vol. 63, no. 11, pp. 7492–7508, 2017.
- [10] M. K. C. Shisher, H. Qin, L. Yang, F. Yan, and Y. Sun, "The age of correlated features in supervised learning based forecasting," in *IEEE INFOCOM Aol Workshop*, 2021, pp. 1–8.
- [11] M. K. C. Shisher and Y. Sun, "How does data freshness affect real-time supervised learning?" In *ACM MobiHoc*, 2022, pp. 31–40.
- [12] A. Muhammad, I. Sorkhoh, M. Samir, D. Ebrahimi, and C. Assi, "Minimizing age of information in multiaccess-edgecomputing-assisted iot networks," *IEEE Internet of Things Journal*, vol. 9, no. 15, pp. 13052–13066, 2021.
- [13] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides, "The age of incorrect information: A new performance metric for status updates," *IEEE/ACM Transactions on Networking*, vol. 28, no. 5, pp. 2215–2228, 2020.
- [14] J. Zhong, R. D. Yates, and E. Soljanin, "Two freshness metrics for local cache refresh," in *IEEE ISIT*, 2018, pp. 1924–1928.
- [15] T. Soleymani, J. S. Baras, S. Wang, S. Hirche, and K. H. Johansson, "Consistency of value of information: Effects of packet loss and time delay in networked control systems tasks," *arXiv preprint arXiv:2403.11932*, 2024.
- [16] D. Guo and I.-H. Hou, "On the credibility of information flows in real-time wireless networks," in *2019 International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*, 2019, pp. 1–8.
- [17] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. I. Jordan, and S. S. Sastry, "Kalman filtering with intermittent observations," *IEEE transactions on Automatic Control*, vol. 49, no. 9, pp. 1453–1464, 2004.
- [18] L. Schenato, "Optimal estimation in networked control systems subject to random delay and packet drop," *IEEE transactions on automatic control*, vol. 53, no. 5, pp. 1311–1317, 2008.
- [19] V. Tripathi and E. Modiano, "A Whittle index approach to minimizing functions of age of information," in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2019, pp. 1160–1167.
- [20] Y. Chen and A. Ephremides, "Scheduling to minimize age of incorrect information with imperfect channel state information," *Entropy*, vol. 23, no. 12, p. 1572, 2021.
- [21] M. K. C. Shisher, Y. Sun, and I.-H. Hou, "Timely communications for remote inference," *IEEE/ACM Transactions on Networking*, vol. 32, no. 5, pp. 3824–3839, 2024.
- [22] M. K. C. Shisher, B. Ji, I.-H. Hou, and Y. Sun, "Learning and communications co-design for remote inference systems: Feature length selection and transmission scheduling," *IEEE*

Journal on Selected Areas in Information Theory, vol. 4, pp. 524–538, 2023.

- [23] T. Z. Ornee, M. K. C. Shisher, C. Kam, and Y. Sun, "Contextaware status updating: Wireless scheduling for maximizing situational awareness in safety-critical systems," in *IEEE MILCOM 2023*, 2023, pp. 194–200.
- [24] P. Whittle, "Restless bandits: Activity allocation in a changing world," *Journal of applied probability*, vol. 25, no. A, pp. 287–298, 1988.
- [25] D. J. Hodge and K. D. Glazebrook, "On the asymptotic optimality of greedy index heuristics for multi-action restless bandits," *Adv. Appl. Probab.*, vol. 47, no. 3, pp. 652–667, 2015.
- [26] G. Chen and S. C. Liew, "An index policy for minimizing the uncertainty-of-information of Markov sources," *IEEE Transactions on Information Theory*, vol. 70, no. 1, pp. 698–721, 2024.
- [27] D. B. Brown and J. E. Smith, "Index policies and performance bounds for dynamic selection problems," *Management Science*, vol. 66, no. 7, pp. 3029–3050, 2020.
- [28] N. Gast, B. Gaujal, and C. Yan, "Linear program-based policies for restless bandits: Necessary and sufficient conditions for (exponentially fast) asymptotic optimality," *Mathematics of Operations Research*, 2023.