

Revisiting Estimation Quality: Significance-Aware Age of Consecutive Error

Jiping Luo and Nikolaos Pappas

Department of Computer and Information Science, Linköping University, Sweden

Email: {jiping.luo, nikolaos.pappas}@liu.se

Abstract—We study the semantics-aware remote state estimation of a Markov chain with *prioritized states*. The aim is to exploit the *significance of information* through the history of system realizations to determine the optimal timing of transmission, thereby reducing the amount of *uninformative data* transmitted in the network. To this end, we introduce the *significance-aware Age of Consecutive Error (AoCE)* that captures three semantic attributes: the *significance of estimation error*, the *cost of consecutive error* (or *lasting impact*, for short), and the *urgency of lasting impact*. We identify the optimal transmission problem as a countably infinite state Markov decision process (MDP) with unbounded costs. We give sufficient conditions under which an optimal policy exists to have bounded average costs. We show that the optimal policy exhibits a *switching structure* and, under certain conditions, degenerates into a simple *threshold policy*. A *structured policy iteration (SPI)* algorithm is proposed to compute an asymptotically optimal policy with reduced computation overhead. An important takeaway is that *the more semantic attributes we utilize, the fewer transmissions are needed*.

I. INTRODUCTION

Remote state estimation is a fundamental and significant problem in networked control systems (NCSs) [1]–[4]. Such systems often involve battery-powered devices sending local observations to remote ends over bandwidth-limited networks. Therefore, the transmitter can only transmit intermittently to *trade off estimation quality against resource utilization* [5]–[10]. An important question arises: How should the transmitter determine which measurements are valuable?

In classical remote estimation, estimation quality is measured by distortion metrics such as Hamming distortion or mean square error, where a measurement is valuable if it contributes to a more *accurate* estimate at the receiver [5]–[8]. The underlying assumption is that *all source states convey equally important information*, and the cost of an estimation error depends solely on the discrepancy between the source and the reconstructed signal. However, this assumption deserves a careful re-examination in many applications. For example, in manufacturing systems, a plant may either reside in a normal state or shift to an alarm state upon an abnormal change in the operation point [11], [12]. In this context, the alarm state is of greater importance, and consequently, missed alarms

This work has been supported in part by the Swedish Research Council (VR), ELLIIT, the Graduate School in Computer Science (CUGS), the European Union (ETHER, 101096526 and 6G-LEADER, 101192080), and the European Union’s Horizon Europe research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No 101131481 (SOVEREIGN).




Fig. 1. Remote estimation of a Markovian source with prioritized states.

incur significantly higher costs than false alarms. Similarly, connected autonomous vehicles demand more precise status information in critical situations (e.g., off-track or dense traffic). This motivates us to revisit the definition of “estimation quality” and incorporate *data significance* into system design.

This paper considers the remote estimation of a finite-state Markov chain $\{X_t\}_{t \geq 1}$ (see Fig. 1). The sensor decides when to send source updates to the receiver. The receiver is tasked with constructing the estimate process $\{\hat{X}_t\}_{t \geq 1}$ based on the received measurements. Given that some states convey more important information than others, we utilize three semantic attributes to capture data significance: the *significance of estimation error*, the *cost of consecutive error* (or *lasting impact*, for short), and the *urgency of lasting impact*. The first attribute is represented by the state-aware distortion (see, e.g., [10], [13]–[17]), where the cost of estimation error depends not only on the physical discrepancy but also on the potential costs or risks incurred without a correct estimate of the state. The second attribute is motivated by the observation that successive reception of an estimation error may have a disastrous impact on the system, not mere accumulated costs during this period [15], [18]–[20]. Moreover, the *urgency of lasting impact* depends on the significance of estimation error. For example, autonomous vehicles might tolerate a moderate number of consecutive status errors under relatively safe conditions. However, inaccurate information about urgent states, even for a few seconds, can lead to wrong operations or even crashes. In this paper, we propose using a *non-linear age function* $g_{X_t, \hat{X}_t}(\Delta_t)$ to model the cost of being in estimation error $(X_t, \hat{X}_t)$ for Δ_t consecutive time slots up until time t . Recently, information aging has received significant attention in status update systems. However, most existing studies have been devoted to content-agnostic age metrics. The design of efficient policies for optimizing significance-aware non-linear aging remains largely unexplored.

We present an overview of the literature in Section II. Section III formulates the optimization problem. Section IV presents our findings on the structure of an optimal policy, the asymptotic optimality of the truncated problem, and the SPI algorithm. Section V presents numerical results. We conclude our work in Section VI. The proofs can be found in [21].

II. RELATED WORK

In the remote estimation literature, the primary objective has been to minimize distortion over constraints on available resources, such as channel bandwidth or energy budget. Since distortion depends on the physical discrepancy between the original and reconstructed signals, the optimal transmission and estimation policies must be *signal-aware* and derived by accounting for the source's evolution pattern. Numerous studies have been devoted to *linear Gaussian systems*. Remote estimation of a scalar Gaussian source with communication costs was studied in [22], where the authors proved that a threshold transmission policy (i.e., the sensor transmits whenever the current error exceeds a threshold) and a Kalman-like estimator are jointly optimal. These results were further extended to systems with multidimensional Gaussian sources and energy-harvesting sensors [6], hard constraints on transmission frequency [8], unreliable channels with adaptive noise [23] and packet drops [24], to name a few.

Another line of research focuses on the use of smart sensors (e.g., with Kalman filters) to pre-estimate source states and then send the estimated states, rather than raw measurements, to the receiver [5], [7], [25]–[27]. Besides tractability, this approach establishes a connection between distortion and the *Age of Information (AoI)* [28], [29]. It has been shown that the error covariance is a monotonic non-decreasing function of AoI. Consequently, AoI serves as a sufficient statistic for decision-making, and the optimal policy initiates a transmission whenever the age exceeds a threshold. These results suggest that, in such systems, measurements are more valuable when they are *fresh*. However, this may not always be the case, as AoI ignores the source pattern and is, therefore, *signal-agnostic*. For instance, in the remote estimation of Wiener and Ornstein-Uhlenbeck processes [9], [30], the estimation error achieved by the distortion-optimal policy can be much smaller than that of the age-optimal policy.

Remote estimation of discrete-state *Markov chains* has gained significant interest in recent years [24], [31], [32]. A salient feature of Markov processes is that they evolve in a probabilistic manner; consequently, the estimation error does not necessarily evolve monotonically with AoI. Another reason traditional distortion and AoI metrics become inefficient in Markovian systems is that the states often convey richer information beyond simple physical amplitudes [10], [20]. Such motivated, the concept of semantics-aware estimation and a series of content-aware metrics that go beyond information accuracy and freshness have been proposed. State-aware distortion was first introduced in [13] to represent the significance of different estimation errors. Performance analysis and comparison of different policies were studied

in [10], [14]–[17]. Structural results of the optimal policy in resource-constrained systems were established in [10]. Various age metrics have been proposed to address the shortcomings of AoI. State-aware AoI [18] and the Uncertainty of Information (UoI) [33] reveal that the information quality depends on the source state (i.e., content) and evolves at different rates. The Version Age of Information (VAoI) tracks the number of content changes and is more relevant than AoI in Markovian systems [34]–[36]. The Age of Incorrect Information (AoII) is a signal-aware metric that counts only the time elapsed since the system was last synced [19], [37], [38]. However, these age metrics treat all estimation errors equally, leading to inadequate transmissions in alarm states but excessive transmissions in normal states.

The closest study to this paper is our previous work in [20], where we introduced the *Age of Missed Alarm (AoMA)* and the *Age of False Alarm (AoFA)* to account for the lasting impact of a binary Markov chain. We showed the existence of a switching optimal policy and derived closed-form expressions. The present work extends these results to a general Markov chain and significance-aware nonlinear age functions.

III. SYSTEM MODEL AND PROBLEM FORMULATION

A. Remote Estimation Model

Consider the remote state estimation model shown in Fig. 1. The process considered is a finite-state, discrete-time Markov chain (DTMC) $\{X_t\}_{t \geq 1}$ defined on the state space

$$\mathcal{X} = \{1, 2, \dots, M\}. \quad (1)$$

Here, state 1 is labeled as the “alarm” state. For later reference, we distinguish between the following types of errors:

- Missed alarms occur when the receiver falsely announces a normal state while the source is in an alarm state, i.e., $X_t = 1, \hat{X}_t \neq 1$. Timely detection of abnormalities is crucial for decision-making and system maintenance.
- False alarms refer to erroneously raising an alarm at the receiver when the source is in a normal state, i.e., $X_t \neq 1, \hat{X}_t = 1$. Although less critical than missed alarms, false alarms can lead to unnecessary expenditure on checking the system thus wasting resources.
- Other normal errors are considered indistinguishable and are not of primary interest.

Let Q denote the state transition probability matrix, where

$$Q = (Q_{i,j}, i, j \in \mathcal{X}), \quad Q_{i,j} = \Pr[X_{t+1} = j | X_t = i]. \quad (2)$$

To avoid pathological cases, we assume Q is *irreducible*. Let $\mathcal{X}_{\text{ap}}$ and $\mathcal{X}_{\text{p}}$ denote the sets of states with and without self-transitions, respectively, where

$$\mathcal{X}_{\text{ap}} = \{i \in \mathcal{X} : Q_{i,i} > 0\}, \quad \mathcal{X}_{\text{p}} = \mathcal{X} - \mathcal{X}_{\text{ap}}. \quad (3)$$

The chain is called *aperiodic* if $\mathcal{X}_{\text{ap}} \neq \emptyset$. Without otherwise stated, we assume Q is aperiodic.

The sensor observes the source state and decides whether or not to transmit a new measurement. Let $A_t \in \{0, 1\}$ denote the decision variable, where $A_t = 1$ means transmission while

$A_t = 0$ means no transmission. We consider an error-prone channel with *i.i.d.* packet drops. Let $H_t \in \{0, 1\}$ denote the packet dropout process, which is a Bernoulli process satisfying

$$\Pr[H_t = 1] = p_s, \Pr[H_t = 0] = 1 - p_s = p_f. \quad (4)$$

Here, $H_t = 1$ denotes that the packet will reach the destination by the end of time slot t (i.e., the decision epoch $t + 1$ in continuous time). In contrast, $H_t = 0$ indicates deep fading channel conditions, leading to transmission failure.

Upon successful reception, the receiver updates its estimate using the newly received measurement, i.e., $\hat{X}_{t+1} = X_t$, and sends an acknowledgment (ACK) packet to the sensor. Otherwise, a negative ACK (NACK) is feedback, and the remote estimate remains unchanged, i.e., $\hat{X}_{t+1} = \hat{X}_t$. We assume that ACK/NACK packets are delivered instantaneously and error-free. Therefore, the sensor knows precisely the remote estimate $\hat{X}_t$ at every decision epoch. The information available at the sensor up to time t is

$$I_t = (X_{1:t}, \hat{X}_{1:t}, A_{1:t-1}). \quad (5)$$

At decision epoch t , a decision A_t is taken according to a *transmission rule* π_t , i.e.,

$$A_t = \pi_t(I_t) = \pi_t(X_{1:t}, \hat{X}_{1:t}, A_{1:t-1}). \quad (6)$$

A *transmission policy* is a sequence of transmission rules, i.e., $\pi = (\pi_1, \pi_2, \dots)$. We call a policy *stationary* if it employs the same rule at every decision epoch. A policy is *deterministic* if, given the history I_t , it selects an action with certainty.

B. Significance of Estimation Errors

In classical remote state estimation, *whether a measurement is discarded or transmitted does not depend on the significance of the measurement*. A widely used metric for Markov sources is the Hamming distortion, i.e.,

$$d(X_t, \hat{X}_t) = \mathbb{1}\{X_t \neq \hat{X}_t\}, \quad (7)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function.

Recall that in our problem, the alarm state (labeled as state 1) is of greater interest. Intuitively, missed alarms typically incur higher costs than other estimation errors. Therefore, we employ a *state-aware distortion* metric that assigns different costs to different estimation errors [10], [13], defined as

$$\bar{d}(X_t, \hat{X}_t) \triangleq \begin{cases} D_{i,j}, & \text{if } (X_t, \hat{X}_t) = (i, j), i \neq j, \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where $D_{i,j}, i \neq j$ are non-negative constants.

A notable shortcoming of distortion lies in its *history-independence*. Although the source evolution is Markovian, the value of information carried by the measurement depends on the history of past observations and decisions. For instance, not only does the instantaneous estimation error matter but also how long the system has been trapped in this error, i.e., the cost of consecutive error (or *lasting impact*, for short).

Definition 1. In this paper, we introduce a new age process, termed *Age of Consecutive Error (AoCE)*, to capture this history-dependent attribute, defined as

$$\Delta_t \triangleq \begin{cases} \Delta_{t-1} + 1, & \text{if } X_t \neq \hat{X}_t, (X_t, \hat{X}_t) = (X_{t-1}, \hat{X}_{t-1}), \\ 1, & \text{if } X_t \neq \hat{X}_t, (X_t, \hat{X}_t) \neq (X_{t-1}, \hat{X}_{t-1}), \\ 0, & \text{if } X_t = \hat{X}_t. \end{cases} \quad (9)$$

A comparison of AoCE to typical metrics is depicted in Fig. 2.

Remark 1. We note that AoCE (9) is change-aware, as it resets upon error changes, whereas AoII is change-agnostic and increments by 1 regardless of the error type. Moreover, AoCE offers the flexibility to handle the lasting impact of different estimation errors separately. However, age alone may not suffice, as it ignores the significance of the current estimation error. This gives incentives to significance-aware age processes.

Let S_t denote the system state at decision epoch t , where

$$S_t = (X_t, \hat{X}_t, \Delta_t). \quad (10)$$

The *significance-aware AoCE* for estimation error $(X_t, \hat{X}_t)$ at decision epoch t is defined as

$$c(S_t) = c(I_t) \triangleq \bar{d}(X_t, \hat{X}_t) \cdot g_{X_t, \hat{X}_t}(\Delta_t), \quad (11)$$

where $(g_{i,j}(\cdot), i, j \in \mathcal{X})$ are non-negative, non-decreasing, and possibly unbounded age functions, $g_{i,j}(\delta)$ represents the cost of being in error (i, j) for δ consecutive time slots. These age functions are quite general and may be discontinuous and non-convex. Given that no cost is incurred in synced states, we impose $g_{i,i}(\cdot) = 0$ for all $i \in \mathcal{X}$.

Remark 2. The significance of information is represented by the state-aware distortion $\bar{d}$, the history-dependent lasting impact Δ_t , and the non-linear age functions $g_{i,j}$. Notably, AoCE (9) is not a sufficient statistic [39] for decision-making unless Assumption 1 holds. We formally state this finding in Lemma 1. This result can be generalized to any age process whose evolution depends on the source pattern.

Assumption 1. The source is non-prioritized, i.e., $D_{i,j} = D$ and $g_{i,j} = g$ for all $i \neq j$, and is symmetric with equal state change probabilities, i.e.,

$$Q_{i,j} = \begin{cases} p, & \text{if } i \neq j, \\ \bar{p}, & \text{otherwise,} \end{cases} \quad (12)$$

where $\bar{p} + (M - 1)p = 1, 0 < p < 1$. In other words, the estimation errors contribute equally to the system in terms of both costs and occurrences.

Lemma 1. The AoCE (9) is a sufficient statistic for the MDP only if Assumption 1 holds.




Fig. 2. Illustration of typical distortion and age metrics, where (1, 2) is a missed alarm, and (3, 1) is a false alarm. New measurements are received at t_0 and t_3 ; however, the system remains erroneous due to changes in the source state. The system is automatically synced at t_5 . Distortion metrics (a)-(b) evaluate only the current estimation error, whereas age-based metrics (c)-(f) account for the history of past observations. The AoI (c) ignores the source evolution and continues to grow even when the system is synced. The AoII (d) increases by 1 whenever an estimation error occurs, while the AoCE (e) resets upon error variations. The significance-aware AoCE (f) assigns exponential age penalties to missed alarms and logarithmic penalties to false alarms, accounting for the urgency of the lasting impact of different estimation errors.

C. System Evolution

The system state $\{S_t\}_{t \geq 1}$ is a three-dimensional controlled Markov chain. It is possible to achieve the desired performance by controlling the chain's transition probabilities. Let

$$P = (P_{s,s'}(a), s, s' \in \mathcal{S}, a \in \mathcal{A}) \quad (13)$$

denote the transition probability matrix, where

$$P_{s,s'}(a) = \Pr[S_{t+1} = s' | S_t = s, A_t = a] \quad (14)$$

is the probability of transitioning to state s' at decision epoch $t+1$, given that the system is in state s and action a is taken at decision epoch t . The transition probabilities of a multi-state model are given as follows.

For any estimation error $(i, j), i \neq j$ with age $\delta \geq 1$, if the sensor decides not to transmit, the system will: (1) remain in this error, i.e., $s' = (i, j, \delta + 1)$, with probability (w.p.) $Q_{i,i}$, (2) become synced, i.e., $s' = (j, j, 0)$, w.p. $Q_{i,j}$, or (3) change to another error $(k, j, 1), k \neq i, j$ w.p. $Q_{i,k}$. Thus, we have

$$\Pr[s' | (i, j, \delta), 0] = \begin{cases} Q_{i,i}, & \text{if } s' = (i, j, \delta + 1), \\ Q_{i,j}, & \text{if } s' = (j, j, 0), \\ Q_{i,k}, & \text{if } s' = (k, j, 1), k \neq i, j, \\ 0, & \text{otherwise.} \end{cases} \quad (15)$$

If the sensor initiates a transmission and a packet is successfully received at the receiver w.p. p_s , the system may either enter a synced state $(i, i, 0)$ w.p. $Q_{i,i}p_s$ or change to an error $(k_1, i, 1), k_1 \neq i$ w.p. $Q_{i,k_1}p_s$. On the other hand, if the transmission is unsuccessful w.p. p_f , the system will: (1) stay in the current error w.p. $Q_{i,i}p_f$, (2) become synced w.p.

$Q_{i,j}p_f$, or (3) change to another error $(k_2, j), k_2 \neq i, j$ w.p. $Q_{i,k_2}p_f$. Thus, we obtain

$$\Pr[s' | (i, j, \delta), 1] = \begin{cases} Q_{i,i}p_s, & \text{if } s' = (i, i, 0), \\ Q_{i,k_1}p_s, & \text{if } s' = (k_1, i, 1), k_1 \neq i, \\ Q_{i,i}p_f, & \text{if } s' = (i, j, \delta + 1), \\ Q_{i,j}p_f, & \text{if } s' = (j, j, 0), \\ Q_{i,k_2}p_f, & \text{if } s' = (k_2, j, 1), k_2 \neq i, j, \\ 0, & \text{otherwise.} \end{cases} \quad (16)$$

For any synced state $(i, i, 0), \forall i \in \mathcal{X}$, the remote estimate remains unchanged, regardless of whether a new measurement is received or not. Therefore, the system will either remain synced w.p. $Q_{i,i}$, or enter an estimation error $(k, i), k \neq i$ w.p. $Q_{i,k}$. Formally, for each $a \in \{0, 1\}$,

$$\Pr[s' | (i, i, 0), a] = \begin{cases} Q_{i,i}, & \text{if } s' = (i, i, 0), \\ Q_{i,k}, & \text{if } s' = (k, i, 1), k \neq i, \\ 0, & \text{otherwise.} \end{cases} \quad (17)$$

D. Optimal Transmission Problem

The goal is to achieve a desired balance between estimation performance and communication cost. Given the cost of each transmission λ , the per-stage cost of taking an action A_t in state S_t is given by

$$l(S_t, A_t) = c(S_t) + \lambda \mathbf{1}\{A_t \neq 0\}. \quad (18)$$

The expected average cost of a transmission policy π over an infinite horizon is defined as

$$\mathcal{L}(\pi) \triangleq \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E}^\pi [l(S_t, A_t) | S_1 = s_1], \quad (19)$$

where $\mathbb{E}^\pi$ represents the conditional expectation, given that policy π is employed with initial state s_1 . The sensor aims to determine the optimal policy π^* to minimize (19), i.e.,

$$\mathcal{L}^* = \inf_{\pi \in \Pi} \mathcal{L}(\pi), \quad (20)$$

where Π is the set of all admissible policies.

Problem (20) is a Markov Decision Process (MDP) with an average cost optimality criterion. This MDP is described by the tuple $(\mathcal{S}, \mathcal{A}, P, l)$, where $\mathcal{S}$ is the state space of all possible values of S_t , $\mathcal{A} = \{0, 1\}$ is the action space, P is the transition probability matrix defined in (13), and l is the cost function given by (18). The state space $\mathcal{S}$ is the union of the set of synced states, $\mathcal{S}_0$, and $M(M-1)$ sets of error holding times, $\mathcal{S}_{i,j}$ for all $i \neq j$. These sets are defined as

$$\mathcal{S}_0 \triangleq \{(i, i, 0) : i \in \mathcal{X}\}, \quad (21)$$

$$\mathcal{S}_{i,j} \triangleq \begin{cases} \{(i, j, \delta) : \delta \geq 1\}, & \text{if } i \neq j, i \in \mathcal{X}_{\text{ap}}, \\ \{(i, j, 1)\}, & \text{if } i \neq j, i \in \mathcal{X}_{\text{p}}. \end{cases} \quad (22)$$

We note that $\mathcal{S}$ is a *countably infinite* set since the error holding time can grow indefinitely. Consequently, the problem encounters computing and memory challenges because classical dynamic programming methods cannot iterate over an infinite state space. Moreover, due to unbounded per-stage costs, an optimal policy may not exist [40]. The following questions relating to the optimal policy are of interest:

- 1) Under what conditions does an optimal policy exist?
- 2) Are there any special structures of the optimal policy that facilitate policy implementation and computation?
- 3) Is it possible to achieve asymptotic optimality by approximating the MDP with a finite state space?
- 4) How to compute the optimal policy with reduced computation?

IV. MAIN RESULTS

This section aims to answer the above questions.

A. Existence of an Optimal Policy

Recall that the AoCE is countably infinite, and the age functions are non-decreasing and unbounded. Consequently, the long-term average cost may never be bounded, no matter how we choose the transmission policy. Thus, we will be concerned with the conditions under which an optimal policy exists to achieve bounded long-term costs.

Assumption 2 gives sufficient conditions, based on which we show the existence of an optimal policy in Theorem 1. This result reveals that the policy space Π can be reduced to a small subset of Markovian (i.e., independent of $S_{1:t-1}$), stationary, and deterministic policies without losing optimality. Moreover, since the optimal policy depends only on the current

observation, all previous measurements can be discarded, thereby saving sensor memory.

Assumption 2. *The age functions, source pattern, and channel condition satisfy the following convergence conditions:*

$$\lim_{\delta \rightarrow \infty} \frac{g_{i,j}(\delta+1)}{g_{i,j}(\delta)} < \frac{1}{Q_{i,i}p_f}, \quad \forall i \in \mathcal{X}_{\text{ap}}, i \neq j, \quad (23)$$

where $Q_{i,i}p_f$ represents the probability of remaining in error (i, j) after each transmission attempt.

Theorem 1. *Suppose Assumption 2 holds. Then there exists an optimal stationary deterministic policy π^* such that for every $s \in \mathcal{S}$, $a = \pi^*(s)$ solves the following Bellman's optimality equation:*

$$\mathcal{L}^* + h(s) = \min_a \{l(s, a) + \sum_{s'} P_{s,s'}(a)h(s')\}, \quad (24)$$

where h is a bounded function, $\mathcal{L}^*$ is the minimal average cost independent of the initial state s_1 .

Remark 3. *The conditions in Assumption 2 are relatively relaxed if the source evolves rapidly and the channel condition is good, i.e., when $Q_{i,i}$ and p_f are small. The intuition behind this is that frequent state changes in the source lead to more frequent error variations (age drops), and good channel conditions increase the likelihood of successful transmissions. Note that in the remote estimation of linear Gaussian processes, the system is easier to stabilize if the source evolves slowly [7], [26]. This notable contradiction arises because the error covariance of linear systems increases monotonically with time and only drops when a new measurement is received. From this perspective, the Markovian nature helps resist lasting impact.*

As a result, we provide several age functions and derive their convergence conditions. It reveals that an optimal policy trivially exists when the age functions are bounded, linear, or logarithmic; it also holds for exponential functions with rate constraints. Similar analyses can be applied to other age metrics, e.g., [18]–[20], [36].

Corollary 1. *The assertion of Theorem 1 holds if:*

- i. $g_{i,j}(\delta)$ is an upper bounded clipping function, i.e.,

$$g_{i,j}(\delta) = \begin{cases} g_{i,j}(\delta), & \text{if } \delta \leq \Delta_{\max}, \\ g_{i,j}(\Delta_{\max}), & \text{otherwise,} \end{cases}$$

where $\Delta_{\max}$ is a finite constant.

- ii. $g_{i,j}(\delta) = \alpha\delta + \beta$ is linear, where $\alpha \geq 0$.
- iii. $g_{i,j}(\delta) = \log(\alpha\delta) + \beta$ is logarithmic, where $\alpha \geq 0$.
- iv. $g_{i,j}(\delta) = \alpha^{\beta\delta} + \zeta$ is an exponential function, where the base and power satisfy $\alpha^\beta < 1/(Q_{i,i}p_f)$. For example, if $g_{i,j}(\delta) = e^{\beta\delta} + \zeta$, we impose $0 \leq \beta < -\ln(Q_{i,i}p_f)$; if $g_{i,j}(\delta) = \alpha^\delta + \zeta$, we impose $1 \leq \alpha < 1/(Q_{i,i}p_f)$.

B. Structure of an Optimal Policy

Next, we show that the optimal policy has a switching structure. This result significantly remedies the ‘‘curse of memory’’ and the ‘‘curse of dimensionality.’’ Moreover, under

Assumption 1, we show that the optimal policy degenerates into a threshold policy, i.e., featuring identical thresholds.

Theorem 2. *The optimal policy, if it exists, exhibits a switching structure. That is, for any given estimation error (i, j) with age $\delta \geq 1$, the sensor initiates a transmission only when δ exceeds a certain threshold $\tau_{i,j}^* \geq 1$. Formally, we write*

$$\pi^*(s) = \begin{cases} 1, & \text{if } \delta \geq \tau_{i,j}^*, i \neq j, \\ 0, & \text{otherwise,} \end{cases} \quad (25)$$

where $\tau_{i,j}^* = 1$ means always transmitting in estimation error (i, j) , whereas $\tau_{i,j}^* \rightarrow \infty$ means no transmission.

Corollary 2. *Under Assumption 1, the optimal policy degenerates to a simple threshold policy, i.e., $\tau_{i,j}^* = \tau^*, \forall i \neq j$.*

Remark 4. *The nice property of the switching policy (25) is its simplicity in implementation and computation: i) to embed the optimal policy in the sensor memory, one only needs to store at most $M(M-1)$ estimation errors and their corresponding threshold values instead of all possible state-action pairs; ii) to implement the policy in a real-time sensor scheduler, trigger a transmission only when the age of the current estimation error exceeds the corresponding threshold and remain silent otherwise; iii) to compute the optimal policy, it suffices to search for a small number of optimal thresholds as opposed to solving a high-dimensional dynamic programming recursion.*

Remark 5. *Theorem 2 answers the fundamental question of “what and when to transmit”. Distortion-optimal policies specify a deterministic mapping from estimation error to transmission decision [10], guiding us on “whether to transmit when a certain error occurs”. By exploiting the lasting impact, our approach further determines the optimal timing to initiate a transmission for each estimation error. Therefore, existing results on distortion can be viewed as special cases where $\tau_{i,j}^* = 1$. An important takeaway is that the more semantic attributes we utilize, the fewer transmissions are needed.*

C. Structured Policy Iteration Algorithm

While Theorem 2 considerably reduces the policy searching space, finding the optimal thresholds is still challenging as classical dynamic programming methods cannot iterate over infinitely many states. For numerical tractability, we truncate the state space and propose a finite-state approximate MDP. The truncated AoCE is defined as

$$\Delta_t(N) \triangleq \min\{\Delta_t, N\}, \quad (26)$$

where Δ_t is the original age process defined in (9), $N > 0$ is the truncation size. In this way, the age associated with each estimation error (i, j) is confined within

$$\mathcal{S}_{i,j}^N = \begin{cases} \{(i, j, \delta) : 1 \leq \delta \leq N\}, & \text{if } i \neq j, i \in \mathcal{X}_{\text{ap}}, \\ \{(i, j, 1)\}, & \text{if } i \neq j, i \in \mathcal{X}_{\text{p}}. \end{cases} \quad (27)$$

The truncated state space is $\mathcal{S}^N = \mathcal{S}_0 \cup_{i \neq j} \mathcal{S}_{i,j}^N$.

The state space truncation, however, inevitably changes the stationary distribution of the induced Markov chain, yielding

inconsistent system performance. Consequently, an optimal policy for the truncated MDP may be suboptimal for the original problem. Thus, we will be concerned with the performance loss caused by state space truncation. Fortunately, the following result establishes the asymptotic optimality of the truncated MDP, and we shall feel safe to truncate the AoCE with an appropriately chosen N .

Theorem 3. *Suppose Assumption 2 holds. Let $\mathcal{L}^*(N)$ denote the minimal cost of the truncated MDP. Then, $\mathcal{L}^*(N)$ converges to the optimal value $\mathcal{L}^*$, i.e., $\lim_{N \rightarrow \infty} \mathcal{L}^*(N) \rightarrow \mathcal{L}^*$, at an exponential rate of $\max_{i \in \mathcal{X}_{\text{ap}}} \{(Q_{i,i} p_f)^N\}$.*

Classical unstructured policy iteration methods can be applied to solve the Bellman equation (24) and obtain an optimal stationary deterministic policy [41]. Let Π^{SD} and Π^* denote the space of stationary deterministic policies and the space of switching policies, respectively, where $\Pi^* \subset \Pi^{\text{SD}}$. Assume the AoCE comprises m ($m = |\mathcal{X}_{\text{ap}}|(M-1) \leq M(M-1)$) sub age processes (see Remark 2), each with a truncation size of N ($m \ll N$). Recall that there are $|\mathcal{S}^N| = M^2 - m + mN$ states and $|\mathcal{A}| = 2$ possible actions. The numbers of possible policies in Π^* and Π^{SD} are given by

$$|\Pi^*| = N^m, \quad |\Pi^{\text{SD}}| = |\mathcal{A}|^{|\mathcal{S}^N|} = 2^{M^2 - m + mN} \approx 2^{mN}.$$

Taking a logarithm of these numbers yields

$$\log_2(|\Pi^*|) = m \log_2(N) \ll \log_2(|\Pi^{\text{SD}}|) = mN.$$

Therefore, the reduction in the size of the structured policy space is quite significant compared to the set of all deterministic policies. Moreover, unstructured iteration methods evaluate all possible policies in a stochastic manner, which makes the convergence rate even slower.

Next, we propose a *structured policy iteration* (SPI) algorithm that exploits the structure results established in Theorem 2. The algorithm proceeds as follows:

- 1) *Initialization:* Initialize policy π^k , reference state s_{ref} , and set $k = 0$; Choose a truncation size N such that for arbitrarily small constant ϵ , $(Q_{i,i} p_f)^N < \epsilon, \forall i \in \mathcal{X}_{\text{ap}}$.
- 2) *Policy Evaluation:* Find $\mathcal{L}^k$ and h^k by solving

$$\mathcal{L}^k + h^k(s) = l(s, \pi^k(s)) + \sum_{s' \in \mathcal{S}^N} P_{s,s'}(\pi^k(s)) h^k(s')$$

for all $s \in \mathcal{S}^N$ such that $h^k(s_{\text{ref}}) = 0$.

- 3) *Policy Improvement:* Find π^{k+1} that satisfies

$$\pi^{k+1}(s) = \min_{a \in \mathcal{A}} \left[l(s, a) + \sum_{s' \in \mathcal{S}^N} P_{s,s'}(a) h^k(s') \right].$$

To ensure the known policy structure, $\pi^{k+1}(s)$ is calculated in the increasing order of δ . If an optimal action $a^* = \pi^k(s)$ for state $s = (i, j, \delta)$ is 1 (transmit), then the optimal action for the remaining states (i, j, δ') , $\delta' > \delta$ is to transmit without further computation.

- 4) *Stopping Criterion:* If $\pi^{k+1} = \pi^k$, the algorithm terminates with $\mathcal{L}^* = \mathcal{L}^k$ and $\pi^* = \pi^k$; otherwise increase $k = k + 1$ and return to step 2.

V. NUMERICAL RESULTS

In the first example, we show the impact of significance-aware non-linear aging on the switching curve. To isolate the effect of other factors on the policy, we consider a symmetric source of the form in (12) with the following parameters: $M = 4, p = 0.1, \bar{p} = 0.7$, and $D_{i,j} = 1, \forall i \neq j$. We assign exponential age functions to missed alarms and logarithmic age functions to false alarms, i.e.,

$$g_{i,j}(\delta) = \begin{cases} e^{0.3\delta}, & \text{if } i = 1, j \neq 1, \\ \log(\delta) + 1, & \text{if } i \neq 1, j = 1, \\ 1, & \text{otherwise.} \end{cases}$$

This assignment satisfies the convergence condition (23). By Theorem 2, there exists an optimal switching policy.

The optimal thresholds for different communication cost λ with a fixed success probability p_s are shown in Table I. Similarly, the thresholds for different p_s with a fixed λ are presented in Table II. Herein, (*) represents the distortion-optimal policy. We observe that for both policies, the thresholds are non-decreasing in λ for fixed p_s and are non-increasing in p_s for fixed λ . When communication is costly, or the channel condition is relatively bad, the optimal switching policy is to transmit less frequently (or never transmit) in false alarms and normal errors while consistently prioritizing missed alarms. In contrast, the distortion-optimal policy either initiates a transmission in all errors (i.e., $\tau_{i,j} = 1, \forall i \neq j$) when the communication is inexpensive ($\lambda \leq 1$), or remains silent otherwise (i.e., $\tau_{i,j} = \infty, \forall i \neq j$). *This highlights the effectiveness of exploiting data significance in such systems.*

We continue to consider a general asymmetric source with

$$Q = \begin{bmatrix} 0.7 & 0.1 & 0.1 & 0.1 \\ 0.05 & 0.7 & 0.15 & 0.1 \\ 0.1 & 0.1 & 0.6 & 0.2 \\ 0.05 & 0.1 & 0.05 & 0.8 \end{bmatrix}.$$

The other parameters are the same as in the first example. For comparison purposes, we consider several benchmark policies: i) **Randomized policy**, which transmits at every slot with a fixed probability $p_\alpha \in [0, 1]$; ii) **periodic policy**, which transmits packets every $t_h \geq 1$ slots; iii-iv) **AoI/AoII policy**, which attains an optimal trade-off between (linear) AoI/AoII minimization and communication utilization; v) **threshold policy**, which initiates a transmission whenever the AoCE exceeds a threshold $\delta_{th} \geq 1$ independent of the error. We numerically search for the optimal values $p_\alpha^*, t_h^*, \delta_{th}^*$, as well as the optimal thresholds $\tau_{i,j}^*$ for the distortion, AoI/AoII, and switching policies. The performance of these policies is summarized in Table III.

It is observed that when communication is cost-free ($\lambda = 0$), all the policies employ the always-transmission strategy. The distortion and switching policies achieve optimal performance when communication is inexpensive ($\lambda \leq 1$). In contrast, when communication is expensive ($\lambda \geq 3$), the optimal randomized and periodic policies adopt the non-transmission strategy, resulting in significant lasting costs. An interesting observation

is that the distortion- and AoII-optimal policies perform poorly compared to the others when $\lambda \geq 3$. This occurs because they transmit measurements in less important errors, leading to poor estimation performance and unnecessary transmission costs. The AoI metric, however, shows obvious disadvantages in our problem since it completely ignores the source evolution and sends information even when the system is in synced states. The reason accounting for the performance losses of the distortion and threshold policies is that they make decisions based on insufficient statistics. Specifically, the distortion policy relies on $(X_t, \hat{X}_t)$, while the threshold policy utilizes only the AoCE Δ_t .

TABLE I
NUMERICAL RESULTS WITH DIFFERENT COMMUNICATION COSTS λ FOR $p_s = 0.9, N = 20$.

λ	Optimal thresholds				Average cost	
	Missed alarms	False alarms	Normal errors	*	Switching	*
0	1	1	1	1	0.35	0.35
1	1	1	1	1	0.67	0.67
2	1	1	∞	∞	0.88	1.63
3	2	3	∞	∞	0.99	1.63
4	3	11	∞	∞	1.04	1.63
5	3	∞	∞	∞	1.05	1.63

TABLE II
NUMERICAL RESULTS WITH DIFFERENT CHANNEL CONDITIONS p_s FOR $\lambda = 2, N = 20$.

p_s	Optimal thresholds				Average cost	
	Missed alarms	False alarms	Normal errors	*	Switching	*
0.4	2	∞	∞	∞	1.05	1.63
0.5	2	8	∞	∞	1.04	1.63
0.6	2	3	∞	∞	0.99	1.63
0.7	2	2	∞	∞	0.96	1.63
0.8	1	2	∞	∞	0.92	1.63
0.9	1	1	∞	∞	0.88	1.63
1.0	1	1	∞	∞	0.85	1.63

TABLE III
PERFORMANCE COMPARISONS WITH DIFFERENT POLICIES FOR $p_s = 0.9, N = 20$.

λ	Achievable minimal average costs						
	Randomized	Periodic	distortion	AoI	AoII	Threshold	Switching
0	0.31	0.31	0.31	0.31	0.31	0.31	0.31
1	0.93	0.85	0.59	0.84	0.62	0.62	0.59
2	1.46	1.05	1.32	1.14	0.91	0.88	0.73
3	1.20	1.20	1.34	1.42	1.14	0.97	0.80
4	1.20	1.20	1.34	1.65	1.34	1.01	0.85
5	1.20	1.20	1.34	1.84	1.34	1.04	0.88

VI. CONCLUSION

In this paper, we have investigated the semantics-aware remote estimation of an asymmetric Markov chain through the significance-aware AoCE. We first give sufficient conditions for the existence of an optimal policy. We prove that a switching policy is optimal and develop a structure-aware algorithm to find the optimal thresholds with reduced computation. Our numerical comparisons show that the optimal policy can be much better than existing rule-, distortion- and age-based policies. The results in this paper generalize recent research on distortion and information aging.

REFERENCES

- [1] W. S. Wong and R. Brockett, "Systems with finite communication bandwidth constraints. i. state estimation problems," *IEEE Trans. Autom. Control*, vol. 42, no. 9, pp. 1294–1299, 1997.
- [2] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu, "A survey of recent results in networked control systems," *Proc. IEEE*, vol. 95, no. 1, pp. 138–162, 2007.
- [3] L. Schenato, B. Sinopoli, M. Franceschetti, K. Poolla, and S. S. Sastry, "Foundations of control and estimation over lossy networks," *Proc. IEEE*, vol. 95, no. 1, pp. 163–187, 2007.
- [4] P. Park, S. Coleri Ergen, C. Fischione, C. Lu, and K. H. Johansson, "Wireless network design for control systems: A survey," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 2, pp. 978–1013, 2018.
- [5] 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, no. 4, pp. 1041–1046, 2013.
- [6] A. Nayyar, T. Başar, D. Teneketzis, and V. V. Veeravalli, "Optimal strategies for communication and remote estimation with an energy harvesting sensor," *IEEE Trans. Autom. Control*, vol. 58, no. 9, pp. 2246–2260, 2013.
- [7] A. S. Leong, S. Dey, and D. E. Quevedo, "Sensor scheduling in variance based event triggered estimation with packet drops," *IEEE Trans. Autom. Control*, vol. 62, no. 4, pp. 1880–1895, 2017.
- [8] J. Chakravorty and A. Mahajan, "Fundamental limits of remote estimation of autoregressive markov processes under communication constraints," *IEEE Trans. Autom. Control*, vol. 62, no. 3, pp. 1109–1124, 2017.
- [9] J. Sun, Y. Polyanskiy, and E. Uysal, "Sampling of the wiener process for remote estimation over a channel with random delay," *IEEE Trans. Inf. Theory*, vol. 66, no. 2, pp. 1118–1135, 2020.
- [10] J. Luo and N. Pappas, "Semantic-aware remote estimation of multiple markov sources under constraints," in *Proc. Int. Symp. Modeling Optim. Mobile, Ad hoc, Wireless Netw.*, 2024, pp. 15–21.
- [11] L. Xie, S. Zou, Y. Xie, and V. V. Veeravalli, "Sequential (quickest) change detection: Classical results and new directions," *IEEE J. Sel. Areas Inf. Theory*, vol. 2, no. 2, pp. 494–514, 2021.
- [12] V. Krishnamurthy, "How to schedule measurements of a noisy markov chain in decision making?" *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4440–4461, 2013.
- [13] N. Pappas and M. Kountouris, "Goal-oriented communication for real-time tracking in autonomous systems," in *Proc. IEEE Int. Conf. Auto. Syst.*, 2021.
- [14] M. Salimnejad, M. Kountouris, and N. Pappas, "Real-time reconstruction of markov sources and remote actuation over wireless channels," *IEEE Trans. Commun.*, vol. 72, no. 5, pp. 2701–2715, 2024.
- [15] —, "State-aware real-time tracking and remote reconstruction of a markov source," *Journal Commun. and Netw.*, vol. 25, no. 5, pp. 657–669, 2023.
- [16] A. Zakeri, M. Moltafet, and M. Codreanu, "Semantic-aware real-time tracking of a markov source under sampling and transmission costs," in *Proc. Asilomar Conf. Signals, Syst. Comput.*, 2023, pp. 694–698.
- [17] J. Luo and N. Pappas, "Goal-oriented estimation of multiple markov sources in resource-constrained systems," in *Proc. IEEE PIMRC*, 2024, pp. 1–6.
- [18] G. Stamatakis, N. Pappas, and A. Traganitis, "Control of status updates for energy harvesting devices that monitor processes with alarms," in *Proc. IEEE Global Commun. Conf. Workshops*, 2019, pp. 1–6.
- [19] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides, "The age of incorrect information: A new performance metric for status updates," *IEEE/ACM Trans. Netw.*, vol. 28, no. 5, pp. 2215–2228, 2020.
- [20] J. Luo and N. Pappas, "Minimizing the age of missed and false alarms in remote estimation of markov sources," in *Proc. ACM MobiHoc*, 2024, p. 381–386.
- [21] —, "On the cost of consecutive estimation error: Significance-aware non-linear aging," *arXiv preprint arXiv:2410.03637*, 2024.
- [22] G. M. Lipsa and N. C. Martins, "Remote state estimation with communication costs for first-order lti systems," *IEEE Trans. Autom. Control*, vol. 56, no. 9, pp. 2013–2025, 2011.
- [23] X. Gao, E. Akyol, and T. Başar, "Optimal communication scheduling and remote estimation over an additive noise channel," *Automatica*, vol. 88, pp. 57–69, 2018.
- [24] J. Chakravorty and A. Mahajan, "Remote estimation over a packet-drop channel with markovian state," *IEEE Trans. Autom. Control*, vol. 65, no. 5, pp. 2016–2031, 2020.
- [25] S. Wu, X. Ren, S. Dey, and L. Shi, "Optimal scheduling of multiple sensors over shared channels with packet transmission constraint," *Automatica*, vol. 96, pp. 22–31, 2018.
- [26] S. Wu, X. Ren, Q.-S. Jia, K. H. Johansson, and L. Shi, "Learning optimal scheduling policy for remote state estimation under uncertain channel condition," *IEEE Trans. Control Netw. Syst.*, vol. 7, no. 2, pp. 579–591, 2020.
- [27] W. Liu, D. E. Quevedo, Y. Li, K. H. Johansson, and B. Vucetic, "Remote state estimation with smart sensors over markov fading channels," *IEEE Trans. Autom. Control*, vol. 67, no. 6, pp. 2743–2757, 2022.
- [28] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. IEEE Int. Conf. Comput. Commun.*, 2012, pp. 2731–2735.
- [29] A. Kosta, N. Pappas, V. Angelakis *et al.*, "Age of information: A new concept, metric, and tool," *Found. Trends Netw.*, vol. 12, no. 3, pp. 162–259, 2017.
- [30] T. Z. Ornee and Y. Sun, "Sampling and remote estimation for the ornstein-uhlenbeck process through queues: Age of information and beyond," *IEEE/ACM Trans. Netw.*, vol. 29, no. 5, pp. 1962–1975, 2021.
- [31] W. Chen, J. Wang, D. Shi, and L. Shi, "Event-based state estimation of hidden markov models through a gilbert-elliott channel," *IEEE Trans. Autom. Control*, vol. 62, no. 7, pp. 3626–3633, 2017.
- [32] X. Liu, M. Cheng, D. Shi, and L. Shi, "Toward event-based state estimation for neuromorphic event cameras," *IEEE Trans. Autom. Control*, vol. 68, no. 7, pp. 4281–4288, 2023.
- [33] G. Chen, S. C. Liew, and Y. Shao, "Uncertainty-of-information scheduling: A restless multiarmed bandit framework," *IEEE Trans. Inf. Theory*, vol. 68, no. 9, pp. 6151–6173, 2022.
- [34] R. D. Yates, "The age of gossip in networks," in *Proc. Int. Symp. Inf. Theory*, 2021, pp. 2984–2989.
- [35] B. Buyukates, M. Bastopcu, and S. Ulukus, "Version age of information in clustered gossip networks," *IEEE J. Sel. Areas Inf. Theory*, vol. 3, no. 1, pp. 85–97, 2022.
- [36] M. Salimnejad, M. Kountouris, A. Ephremides, and N. Pappas, "Version innovation age and age of incorrect version for monitoring markovian sources," in *Proc. Int. Symp. Modeling Optim. Mobile, Ad hoc, Wireless Netw.*, 2024, pp. 162–169.
- [37] Y. Chen and A. Ephremides, "Minimizing age of incorrect information over a channel with random delay," *IEEE/ACM Trans. Netw.*, pp. 1–13, 2024.
- [38] I. Cosandal, N. Akar, and S. Ulukus, "Aoii-optimum sampling of ctmc information sources under sampling rate constraints," in *Proc. Int. Symp. Inf. Theory*, 2024, pp. 1391–1396.
- [39] A. Mahajan and M. Mannan, "Decentralized stochastic control," *Annals of Operations Research*, vol. 241, no. 1, pp. 109–126, 2016.
- [40] L. I. Sennott, *Stochastic dynamic programming and the control of queueing systems*. John Wiley & Sons, 1998.
- [41] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 1994.