RESEARCH-ARTICLE

Minimizing the Age of Missed and False Alarms in Remote Estimation of Markov Sources

JIPING LUO, Linköping University, Linkoping, Östergötland, Sweden

NIKOLAOS PAPPAS, Linköping University, Linkoping, Östergötland, Sweden

Published: 14 October 2024

Citation in BibTeX format

MobiHoc '24: Twenty-fifth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing
October 14 - 17, 2024
Athens, Greece

Conference Sponsors:
SIGMOBILE

MOBIHOC '24: Proceedings of the Twenty-fifth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (October 2024)
https://doi.org/10.1145/3641512.3690161
ISBN: 9798400705212

Minimizing the Age of Missed and False Alarms in Remote Estimation of Markov Sources

Jiping Luo
Dept. of Computer and Information Science
Linköping University, Sweden
jiping.luo@liu.se

Nikolaos Pappas
Dept. of Computer and Information Science
Linköping University, Sweden
nikolaos.pappas@liu.se

ABSTRACT

We consider the remote estimation of a discrete-state Markov source with *normal* and *alarm* states. Data significance is revealed via two semantic attributes: 1) Erroneously announcing a normal state at the destination when the source is actually in an alarm state (i.e., missed alarm error) incurs a significantly higher cost than falsely announcing an alarm state when the source is in a normal state (i.e., false alarm error). 2) Successive reception of an estimation error may cause significant *lasting impact*, e.g., maintenance cost and wrong operations. Motivated by this, we assign different costs to different estimation errors and introduce two new age metrics, namely the *Age of Missed Alarm* (AoMA) and the *Age of False Alarm* (AoFA), to account for the *lasting impact* incurred by different estimation errors. We aim to achieve an optimal trade-off between the cost of estimation error, lasting impact, and communication utilization. The problem is formulated as an infinite-state Markov decision process (MDP). We show that the optimal policy exhibits a *switching* structure, i.e., triggering transmissions only when the AoMA or AoFA exceeds a threshold. Numerical results underscore that our approach significantly reduces the amount of less important information transmitted in the networks.

CCS CONCEPTS

• **Networks** → **Network performance modeling**; **Network performance analysis**.

KEYWORDS

Semantic communications, remote estimation, Markov source, age of information, data significance

ACM Reference Format:
Jiping Luo and Nikolaos Pappas. 2024. Minimizing the Age of Missed and False Alarms in Remote Estimation of Markov Sources. In *The Twenty-fifth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MOBIHOC '24), October 14–17, 2024, Athens, Greece.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3641512.3690161

1 INTRODUCTION

Efficient remote state estimation is the key to various networked control systems (NCSs) [7, 18, 27]. A fundamental question is *how to achieve an optimal trade-off between estimation performance and communication cost* [1, 2, 24]. This is often done through various metrics that capture the *value (semantics) of information*, thereby reducing the amount of inefficient data transmissions [5, 6, 10, 19]. *Accuracy* and *freshness*, usually measured by distortion and the Age of Information (AoI) [9, 26, 28], respectively, are two dominant semantic attributes in the remote estimation literature. Despite significant efforts, they are inefficient in Markovian models due to the ignorance of data significance and source evolution [10, 17].

Several metrics have been introduced to address the shortcomings of AoI. The Age of Incorrect Information (AoII) [3, 15] captures another semantic attribute — *lasting impact*, i.e., the cost of consecutive estimation errors. This attribute is significant in many applications. For instance, successive reception of erroneous status updates from a remotely controlled drone can lead to wrong operations or even crashes. However, AoII might not suffice in our problem since *treats all source states equally, i.e., content-agnostic.* This egalitarianism results in inadequate transmissions in alarm states but excessive transmissions in normal states. This gives incentives to *content-aware* metrics.

More recently, there have been some initial efforts to exploit content awareness (i.e., *state-dependent significance*) in remote estimation systems [14, 17, 20, 22, 25]. The authors in [25] assigned a quadratic age variable for the alarm state and a linear age variable for the normal state, thus punishing more on the age of the alarm state. However, this method inherits some of the shortcomings of AoI. The Version Age of Information (VAoI) [20] tracks content changes in the age process, which becomes more relevant in systems sensitive to error variations. The most relevant metric that directly captures the data significance is the *cost of actuation error* (CAE) [14, 17, 21, 22], which assigns *different costs to different estimation errors.* Intuitively, a missed alarm error incurs a significantly higher cost than a false alarm error. Theoretical results were presented in [14], which showed that the optimal policy specifies a deterministic mapping from estimation error to transmission decision. However, the lasting impact is ignored.

The main contributions of this work are as follows: 1) We introduce two new age metrics, namely the Age of Missed Alarm (AoMA) and the Age of False Alarm (AoFA), to account for the lasting impacts of missed alarm errors and false alarm errors, respectively. The AoMA and AoFA evolve dependently, allowing them to distinguish between estimation errors. 2) We show that the optimal policy is of a *switching-type*, i.e., triggering transmissions when the AoMA or AoFA exceeds a given threshold. Moreover, we give a sufficient condition under which the optimal policy degenerates from the switching-type to the simple threshold-type, i.e., identical thresholds. 3) For numerical tractability, we truncate the age

ACM ISBN 979-8-4007-0521-2/24/10
https://doi.org/10.1145/3641512.3690161

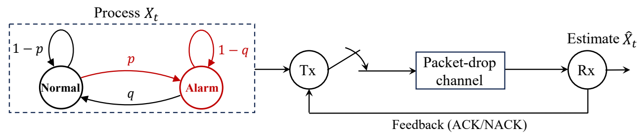

Figure 1: Remote estimation with prioritized states.

processes and show the asymptotical optimality of the truncated problem.

2 SYSTEM MODEL

2.1 Remote Estimation Model

We consider a remote estimation system shown in Fig. 1. The model consists of the parts listed in the next subsections.

2.1.1 Source. The information source is modeled by a two-state, discrete-time Markov chain (DTMC) $\{X_t\}_{t\geq 1}$, where $X_t \in \{0,1\}$ is the state variable, $t \geq 1$ is the time index. At each slot t, the source resides in either state 0 (*normal*, or low-priority) or state 1 (*alarm*, or high-priority). It is worth mentioning that states here can represent either quantization levels of a physical process or abstract statuses (e.g., operation modes, component failures, or abrupt changes in system dynamics, etc.) of a system[1].

The state transition probability matrix of the DTMC is

$$Q = \begin{bmatrix} \bar{p} & p \\ q & \bar{q} \end{bmatrix} \tag{1}$$

where $\bar{p} = 1-p, \bar{q} = 1-q$, $Q_{i,j} = \Pr[X_{t+1} = j | X_t = i]$ is the probability of transitioning from state i to state j between two consecutive time slots. To avoid pathological cases, we assume Q is *irreducible*, i.e., $0 < p, q < 1$. A DTMC is called *symmetric* if $Q = Q^T$. Otherwise, the chain is called *asymmetric*. Moreover, we call a chain *positively correlated* if it is more likely to stay in the current state than to change state, i.e., $p < \bar{q}$. The chain is *negatively correlated* if $p > \bar{q}$ and is *i.i.d.* if $p = \bar{q}$.

2.1.2 Channel. The channel state $H_t \in \{0,1\}$ follows a Bernoulli distribution with a mean of p_s. Here, $H_t = 1$ denotes the successful reception of a data packet at time t, and $H_t = 0$ denotes a packet drop. Upon successful reception, the receiver sends a one-bit acknowledgment (ACK) packet to the sensor. Otherwise, a negative ACK (NACK) signal is feedback to indicate transmission failure. We assume that ACK/NACK packets are delivered instantaneously and error-free. Therefore, the sensor knows precisely the estimate $\hat{X}_t$ at the receiver.

2.1.3 Sensor. The sensor sequentially observes the process and decides whether or not to transmit the source state. We assume that *the decision made in the current time slot will take effect at the beginning of the next time slot, not immediately*. Specifically, if during the current time slot t, the sensor chooses to update the source state, a new packet X_{t+1} is sampled and transmitted at the beginning of slot $t+1$. This action delay accounts for several practical considerations such as measurement delay and system processing time [8]. Let $A_t \in \{0,1\}$ denote the decision variable, where $A_t = 1$ means transmission while $A_t = 0$ means no transmission.

The information available at the sensor at each decision epoch t is

$$I_t^1 = \{X_{1:t}, \hat{X}_{1:t}, A_{1:t-1}\}. \tag{2}$$

The sensor chooses an action A_t according to some *transmission rule* π_t, i.e.,

$$A_t = \pi_t(I_t^1) = \pi_t(X_{1:t}, \hat{X}_{1:t}, A_{1:t-1}). \tag{3}$$

Notice that π_t can be either *deterministic*, selecting an action in a given state with certainty, or *randomized*, specifying a probability distribution on the action space. The collection $\pi = \{\pi_t\}_{t=1}^{\infty}$ is called the *transmission policy*.

2.1.4 Receiver. Let $Y_t \in \{0,1,\mathcal{E}\}$ denote the output of the channel at time t, where the symbol $\mathcal{E}$ denotes no packet received. The information available at the receiver at each time t is

$$I_t^2 = \{Y_{1:t}, H_{1:t}\}. \tag{4}$$

However, it is challenging to reconstruct source states using undersampled measurements. The *estimation rule* g_t depends on the source statistics, the channel condition, and the transmission policy. In this paper, we assume that the receiver does not know the source pattern and updates its estimate using the latest received measurement[2][14, 15, 17], i.e.,

$$\hat{X}_t = g_t(I_t^2) = X_{U_t}, \tag{5}$$

where $U_t = \max\{1 \leq \tau \leq t : Y_\tau \neq \mathcal{E}\}$.

Remark 1. *The AoI in such systems is defined as $\Delta_t^{AoI} = t - U_t$. However, AoI ignores both the source evolution and the information content. Consequently, a large AoI does not necessarily mean poor estimation performance, as the system may remain synced for some time, and vice versa.*

2.2 Semantics-Aware Age Metrics

Since the source states are not equally important, we distinguish the estimation errors by the following two types.

- **False Alarm (FA):** It occurs when $X_t = 0$ and $\hat{X}_t = 1$, indicating an unnecessary alarm triggered by the receiver. While somewhat less important, FA errors may incur extra operation or maintenance costs.
- **Missed Alarm (MA):** It occurs when $X_t = 1$ and $\hat{X}_t = 0$, indicating a failure to detect an alarm by the receiver. MA errors are more crucial and thus higher penalties.

Then, we introduce two new age metrics, namely the Age of False Alarm (AoFA) and the Age of Missed Alarm (AoMA), to quantify the *lasting impact*[3] of the FA and MA errors, respectively. The age processes evolve as follows

$$\Delta_{t+1}^{\text{FA}} = \begin{cases} \Delta_t^{\text{FA}} + 1, & \text{if } X_{t+1} = 0, \hat{X}_{t+1} = 1, \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

[1]Examples include Markov jump systems[4], where the system has a finite number of operation modes governed by a DTMC. Each mode corresponds to a specific type of physical process or a particular group of system parameters.

[2]Assuming the availability of source statistics at the receiver side, one may consider other estimation rules such as MMSE and Maximum likelihood estimation (MLE)[12].

[3]Staying in an erroneous state for an extended period may lead to substantial ramifications, not mere accumulated costs during this period[15, 22].

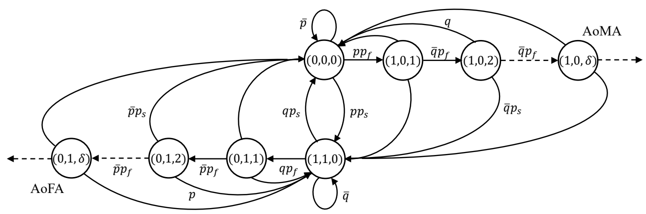

Figure 2: Three-dimensional DTMC representing $\{S_t\}_{t\geq 1}$.

$$\Delta_{t+1}^{\mathrm{MA}} = \begin{cases} \Delta_t^{\mathrm{MA}} + 1, & \text{if } X_{t+1} = 1, \hat{X}_{t+1} = 0, \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

Let $S_t = (X_t, \hat{X}_t, \Delta_t^{\mathrm{MA}}, \Delta_t^{\mathrm{FA}})$ denote the system state at time t. Since Δ_t^{FA} and Δ_t^{MA} cannot simultaneously be non-zero, they can be formed in a compact way as

$$\Delta_t = \mathbb{1}_{\{(X_t,\hat{X}_t)=(0,1)\}}\Delta_t^{\mathrm{FA}} + \mathbb{1}_{\{(X_t,\hat{X}_t)=(1,0)\}}\Delta_t^{\mathrm{MA}}, \tag{8}$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function of an event. For simplicity, in the rest of the paper, we shall use $S_t = (X_t, \hat{X}_t, \Delta_t)$ as the system state. The state space is given by

$$\mathcal{S} = \{(0,0,0),(1,1,0)\} \cup \{(0,1,\delta),(1,0,\delta) : \delta \geq 1\}, \tag{9}$$

where $\mathcal{S}_{\mathrm{synced}} = \{(0,0,0),(1,1,0)\}$, $\mathcal{S}_{\mathrm{FA}} = \{(0,1,\delta), \delta \geq 1\}$, and $\mathcal{S}_{\mathrm{MA}} = \{(1,0,\delta), \delta \geq 1\}$ are the set of synced states, the set of FA errors, and the set of MA errors, respectively. Note that $\mathcal{S}$ is a *countably infinite* set because the ages are possibly unbounded.

The cost of being in state S_t is defined as

$$\begin{aligned} c(S_t) &= \beta\Delta_t^{\mathrm{MA}} + (1-\beta)\Delta_t^{\mathrm{FA}} \\ &= \big(\beta\mathbb{1}_{\{(X_t,\hat{X}_t)=(1,0)\}} + (1-\beta)\mathbb{1}_{\{(X_t,\hat{X}_t)=(0,1)\}}\big)\Delta_t, \end{aligned} \tag{10}$$

where $\beta \in (0,1]$ represents the *significance* of MA errors. Recall that action A_t will take effect at the beginning of slot $t+1$. Therefore, we can only evaluate the effectiveness of an action after receiving ACK/NACK packets, i.e., knowing the outcome of S_{t+1}. The (per-step expected) cost of taking an action A_t in state S_t is

$$c(S_t, A_t) = \mathbb{E}\{c(S_{t+1})|S_t, A_t\}, \tag{11}$$

where the expectation is taken over the channel uncertainties, the source statistics, and the (possibly) random actions.

Remark 2. *The cost* (10) *assigns different costs and age penalties for different estimation errors. Notice that AoI and AoII [15] only have one age process. Therefore, they cannot distinguish between estimation errors and are not applicable in some applications such as quickest change detection [11] where a certain state holds more interest.*

2.3 System Dynamics

Take an action a in a certain state $s = (x, \hat{x}, \delta)$, the system state will transition to $s' = (x', \hat{x}', \delta')$ according to the following probabilities

$$P_{s,s'}(a) = \Pr\{(x', \hat{x}')|(x, \hat{x}), a\}\Pr\{\delta'|x', \hat{x}', \delta\}, \tag{12}$$

where

$$\Pr\{(x', \hat{x}')|(x, \hat{x}), a\} = \begin{cases} Q_{i,k}p_s, & a = 1, (x,\hat{x}) = (i,j), (x',\hat{x}') = (k,k), k \neq j, \\ Q_{i,k}p_f, & a = 1, (x,\hat{x}) = (i,j), (x',\hat{x}') = (k,j), k \neq j, \\ Q_{i,j}, & a = 1, (x,\hat{x}) = (i,j), (x',\hat{x}') = (j,j), \\ Q_{i,k}, & a = 0, (x,\hat{x}) = (i,j), (x',\hat{x}') = (k,j), \\ 0, & \text{otherwise.} \end{cases} \tag{13}$$

$$\Pr\{\delta'|x', \hat{x}', \delta\} = \begin{cases} 1, & x' \neq \hat{x}', \delta' = \delta + 1, \\ 1, & x' = \hat{x}', \delta' = 0, \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

where $p_f = 1 - p_s, i, j, k \in \{0, 1\}$. Fig. 2 shows the evolution of process S_t under the always-transmission policy, i.e., $A_t = 1, \forall t \geq 1$.

3 PROBLEM FORMULATION

Given a transmission policy π, the average content-aware estimation error is defined as

$$C(\pi) \triangleq \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}^{\pi}\big[c(S_t, A_t)\big|S_1 = s_1\big], \tag{15}$$

where $\mathbb{E}^{\pi}$ represents the conditional expectation, given that policy π is employed with initial state $s_1 = (0,0,0)$.

Similarly, the transmission frequency (average number of transmissions) is defined as

$$F(\pi) \triangleq \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}^{\pi}\big[f(S_t, A_t)\big|S_1 = s_1\big], \tag{16}$$

where $f(S_t, A_t) = \mathbb{1}_{\{A_t \neq 0\}}$.

Given the communication cost (i.e., the cost of joint sampling and transmission) λ for each transmission, the sensor aims to achieve a desired balance between the communication cost and the estimation performance. Hence, the sensor's global objective function is

$$\begin{aligned} \mathcal{L}^{\lambda}(\pi) &\triangleq \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}^{\pi}\big[l^{\lambda}(S_t, A_t)\big|S_1 = s_1\big], \\ &= C(\pi) + \lambda F(\pi), \end{aligned} \tag{17}$$

where $l^{\lambda}(S_t, A_t) \triangleq c(S_t, A_t) + \lambda f(S_t, A_t)$.

The sensor aims to determine the λ-optimal transmission policy π_{λ}^* to minimize (17), i.e.,

$$\mathcal{L}^{\lambda} = \inf_{\pi\in\Pi} \mathcal{L}^{\lambda}(\pi), \tag{18}$$

where Π is the set of all admissible policies.

Remark 3. *Problem* (18) *is an average-cost Markov Decision Process (MDP). However, it encounters computing and memory challenges due to the infinite state space and (possibly) unbounded per-step costs. The main approach to remedy these difficulties is to reveal favorable properties of the optimal policy, thereby restricting the policy searching space.*

Definition 1. A policy is called λ-optimal if, for a given $\lambda \geq 0$, it attains the infimum in (18).

4 MAIN RESULTS

4.1 Optimality of Switching-Type Policy

In the sequel, we develop structural results of the λ-optimal policy. Problem (18) can be characterized by an MDP with a 4-tuple

$(\mathcal{S}, \mathcal{A}, P, l^{\lambda})$. Herein, $l^{\lambda}(s,a)$ is the (per-step) cost function, and P is the system transition probability function given in (12). To address the "curse of dimensionality" and the "curse of memory", we intend to derive the structure of the λ-optimal policy and characterize the limiting behavior of the induced Markov chain.

We first show in the following proposition the existence of a λ-optimal policy. Moreover, it allows us to restrict attention to stationary deterministic policies.

Proposition 1. *Suppose that Q is irreducible. Then, there exists a function h^{λ} such that*

$$\mathcal{L}^{\lambda} + h^{\lambda}(i) = \min_a \{l^{\lambda}(i,a) + \sum\nolimits_j P_{i,j}(a) h^{\lambda}(j)\} \tag{19}$$

for all $i \in \mathcal{S}$, where $\mathcal{L}^{\lambda}$ is the minimum average cost. Furthermore, the λ-optimal policy π^*_{λ} is stationary deterministic given by

$$\pi^*_{\lambda}(i) = \arg\min_a \{l^{\lambda}(i,a) + \sum\nolimits_j P_{i,j}(a) h^{\lambda}(j)\}, i \in \mathcal{S}. \tag{20}$$

Proof. Please see Proposition 2 in [13]. □

Remark 4. *One may try to apply value iteration methods to solve the Bellman's optimality equation* (19)*. However, it is not applicable in practice as we cannot iterate over infinitely many states. Furthermore, even with such an optimal policy at hand, it is impossible to store all state-action pairs due to memory constraints at the sensor. Deep reinforcement learning (DRL) techniques can partially address these challenges, however, at the cost of suboptimality and lack of interpretability.*

We first address the "curse of memory". Fortunately, the following important theorem shows that the λ-optimal policy has a simple *switching-type structure* facilitating the policy storage and algorithm design.

Assumption 1. *The source is positively correlated*[4].

Definition 2. We define the ordering $s_1 \leq s_2$ for $s_1, s_2 \in \mathcal{S}_{\text{AoMA}} \triangleq \mathcal{S}_{\text{synced}} \cup \mathcal{S}_{\text{MA}}$ if the age satisfies $0 \leq \delta_1 \leq \delta_2$. Similarly, the ordering for $s_1, s_2 \in \mathcal{S}_{\text{AoFA}} \triangleq \mathcal{S}_{\text{synced}} \cup \mathcal{S}_{\text{FA}}$ is $s_1 \leq s_2$ if $0 \leq \delta_1 \leq \delta_2$.

Theorem 3. *The λ-optimal policy has a switching-type structure. That is, for any given ordering $s_- \leq s \leq s_+, s_-, s, s_+ \in \mathcal{S}_{\text{AoMA}}$ (or $\mathcal{S}_{\text{AoFA}}$), if $\pi^*_{\lambda}(s) = 1$, then $\pi^*_{\lambda}(s_+) = 1$ for all $s_+ \geq s$. Also, if $\pi^*_{\lambda}(s) = 0$, then $\pi^*_{\lambda}(s_-) = 0$ for all $s_- \leq s$.*

Proof. Please see Theorem 1 in [13]. □

The sensor initiates a transmission only when the current AoMA exceeds the threshold $\bar{\delta}$ or the AoFA exceeds the threshold $\underline{\delta}$. Therefore, the sensor only needs to store two threshold values instead of all possible state-action pairs, thus remedying the "curse of memory". We note that the threshold $\bar{\delta}$ for AoMA does not necessarily equal the threshold $\underline{\delta}$ for AoFA. We will show that identical thresholds are generally suboptimal (see Section 4.3).

[4]This is a common assumption in the literature [16]. Analysis without this assumption can be found in [13].

4.2 Performance of Switching-Type Policy

In the sequel, we present analytical results of the switching-type policy. It is worth mentioning that *our approach can differentiate between the synced states*–a feature rarely considered in the literature. This feature offers significant performance benefits since the *synced states are not equally important.* For instance, if the sensor triggers a transmission in the synced state $(0,0,0)$, then there is a high probability that the system will directly transition to another synced state $(1,1,0)$, thereby forcing the system into low-cost FA errors. In outline, our approach allows us to distinguish between the following four cases: (i) always transmission, i.e., $\bar{\delta} = \underline{\delta} = 0$, (ii) no transmission in synced states, i.e., $\bar{\delta}, \underline{\delta} \geq 1$, (iii) transmitting in state $(0,0,0)$, i.e., $\bar{\delta} = 0, \underline{\delta} \geq 1$, and (iv) transmitting in state $(1,1,0)$, i.e., $\bar{\delta} \geq 1, \underline{\delta} = 0$.

The main result of this section follows. Proposition 2 establishes the stationary distribution of the switching-type policy for the dominating case $\bar{\delta}, \underline{\delta} \geq 1$. Further, theorem 4 provides the closed-form expressions for the average cost associated with this case. The analysis for the other three cases is carried out similarly (see [13]).

Proposition 2. *Suppose that the thresholds for AoMA and AoFA are $\bar{\delta}, \underline{\delta} \geq 1$. Then, the DTMC in question is irreducible and admits a stationary distribution given by*

$$\nu_{0,0,0} = \Gamma_0(a,b), \ \ \nu_{1,1,0} = \Gamma_1(a,b), \tag{21}$$

$$\nu_{1,0,k} = \begin{cases} p\bar{q}^{k-1}\nu_{0,0,0}, & 1 \leq k \leq \bar{\delta}, \\ p\bar{q}^{\bar{\delta}-1}(\bar{q}p_f)^{k-\bar{\delta}}\nu_{0,0,0}, & k > \bar{\delta}. \end{cases} \tag{22}$$

$$\nu_{0,1,k} = \begin{cases} q\bar{p}^{k-1}\nu_{1,1,0}, & 1 \leq k \leq \underline{\delta}, \\ q\bar{p}^{\underline{\delta}-1}(\bar{p}p_f)^{k-\underline{\delta}}\nu_{1,1,0}, & k > \underline{\delta}. \end{cases} \tag{23}$$

where

$$\Gamma_0(a,b) = \frac{\bar{p}b(\underline{\delta})}{(1+a(1))\bar{p}b(\underline{\delta}) + (1+b(1))\bar{q}a(\bar{\delta})}, \tag{24}$$

$$\Gamma_1(a,b) = \frac{\bar{q}a(\bar{\delta})}{(1+a(1))\bar{p}b(\underline{\delta}) + (1+b(1))\bar{q}a(\bar{\delta})}, \tag{25}$$

$$a(1) = \frac{p(1-\bar{q}^{\bar{\delta}-1})}{1-\bar{q}} + a(\bar{\delta}), \ \ a(\bar{\delta}) = \frac{p\bar{q}^{\bar{\delta}-1}}{1-\bar{q}p_f}, \tag{26}$$

$$b(1) = \frac{q(1-\bar{p}^{\underline{\delta}-1})}{1-\bar{p}} + b(\underline{\delta}), \ \ b(\underline{\delta}) = \frac{q\bar{p}^{\underline{\delta}-1}}{1-\bar{p}p_f}. \tag{27}$$

Proof. Please see Proposition 3 in [13]. □

Theorem 4. *For a given switching-type policy π with thresholds $\bar{\delta}, \underline{\delta} \geq 1$, the average costs are*

$$C(\pi) = \beta p\psi(\bar{q}, \bar{\delta})\nu_{0,0,0} + (1-\beta)q\psi(\bar{p}, \underline{\delta})\nu_{1,1,0}, \tag{28}$$

$$F(\pi) = a(\bar{\delta})\nu_{0,0,0} + b(\underline{\delta})\nu_{1,1,0}, \tag{29}$$

$$\mathcal{L}^{\lambda}(\pi) = C^{\lambda}(\pi) + \lambda F^{\lambda}(\pi), \tag{30}$$

where $a(\bar{\delta}), b(\underline{\delta}), \nu_{0,0,0}$, and $\nu_{1,1,0}$ are given in Proposition 2, $\psi(p,\delta) = \frac{1-[p+(1-p)\delta]p^{\delta-1}}{(1-p)^2} + \frac{[pp_f+(1-pp_f)\delta]p^{\delta-1}}{(1-pp_f)^2}$.

Proof. Please see Theorem 2 in [13]. □

4.3 Suboptimality of Threshold-Type Policy

The following theorem shows that, when the source is symmetric and the states are equally important, the optimal policy degenerates from a switching-type to a simple threshold-type. *This result underscores that existing results on AoII can be viewed as special cases of our approach (see [15]).*

Theorem 5. *Suppose that the source is symmetric and the states are equally important, i.e., $p = q, \beta = 0.5$. Then, the optimal policy has identical thresholds.*

Proof. Please see Theorem 3 in [13]. □

Remark 5. *The assertion of Theorem 5 fails to hold when the source is asymmetric or the states are prioritized. Since our approach covers the single age process case, it implies that the threshold-type policy is generally suboptimal. Intuitively, a single threshold cannot distinguish between different types of errors and different synced states. Also, a threshold policy with $\delta = 0$ is nothing but a continuous-transmission policy.*

4.4 Asymptotical Optimality

As mentioned earlier, it is impractical to iterate over an infinite state space. For numerical tractability, we truncate the state space and propose a finite-state approximate MDP[23, Chapter 8]. The truncated age processes evolve as follows

$$\Delta_{t+1}^{\mathrm{FA}}(N) = \begin{cases} \left[\Delta_t^{\mathrm{FA}} + 1\right]_N^+, & \text{if } X_{t+1} = 0, \hat{X}_{t+1} = 1, \\ 0, & \text{otherwise.} \end{cases} \quad (31)$$

$$\Delta_{t+1}^{\mathrm{MA}}(N) = \begin{cases} \left[\Delta_t^{\mathrm{MA}} + 1\right]_N^+, & \text{if } X_{t+1} = 1, \hat{X}_{t+1} = 0, \\ 0, & \text{otherwise.} \end{cases} \quad (32)$$

where $[x]_m^+ = x$ if $x \le m$ and $[x]_m^+ = m$ otherwise. That is, the system is confined within

$$\mathcal{S}_N = \{(0,0,0), (1,1,0), (1,0,\delta), (0,1,\delta) : \forall \delta \le N\}. \quad (33)$$

All other states are considered indistinguishable from the boundary states $(1,0,N)$ and $(0,1,N)$.

We now show that the truncated MDP converges to the original MDP exponentially fast in N. This result significantly mitigates the "curse of dimensionality", and we shall feel safe to truncate the state space with an appropriately chosen N.

Theorem 6. *Let $\mathcal{L}^\lambda(N)$ be the minimum average cost for the truncated MDP. Then, $\lim_{N\to\infty} \mathcal{L}^\lambda(N) \to \mathcal{L}^\lambda$.*

Proof. Please see Theorem 4 in [13]. □

The optimal thresholds can be obtained by finding the minima of the cost function $\mathcal{L}^\lambda(\pi, N)$ within the policy domain $0 \le \bar{\delta}, \underline{\delta} \le N$. One may also apply the relative value iteration algorithm[14] in the state domain $\mathcal{S}_N$, but at the cost of high complexity. According to our analysis, the impact of state space truncation vanishes for a considerably large N. Therefore, the state space truncation method entails a trade-off between optimality and complexity.

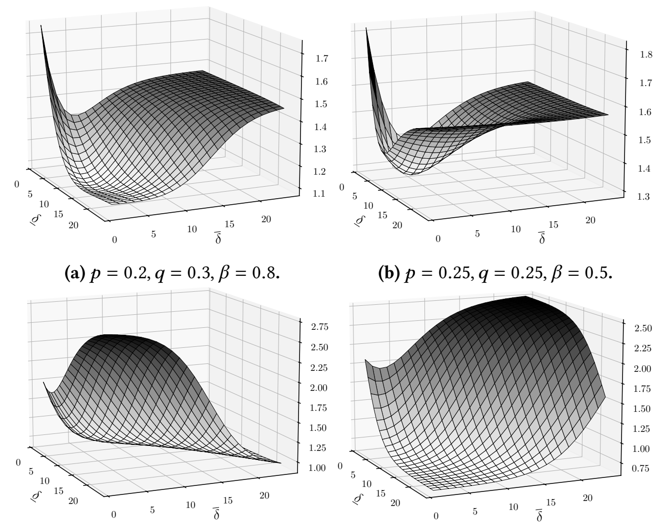

(a) $p = 0.2, q = 0.3, \beta = 0.8$. (b) $p = 0.25, q = 0.25, \beta = 0.5$. (c) $p = 0.2, q = 0.3, \beta = 0.5$. (d) $p = 0.25, q = 0.25, \beta = 0.8$.

Figure 3: The average cost as a function of different thresholds when $p_s = 0.9, \lambda = 8, N = 100$. The minimum is found at (a) $(\bar{\delta}^*, \underline{\delta}^*) = (3, 13)$, (b) $(\bar{\delta}^*, \underline{\delta}^*) = (5, 5)$, (c) $(\bar{\delta}^*, \underline{\delta}^*) = (11, 1)$, and (d) $(\bar{\delta}^*, \underline{\delta}^*) = (1, 24)$, respectively.

5 NUMERICAL RESULTS

In this section, we present numerical results for the remote estimation problem (18).

Fig. 3 plots the average cost of the switching-type policy as a function of different threshold values, i.e., $\mathcal{L}^\lambda(\bar{\delta}, \underline{\delta})$. Note that we only plot the results when $\bar{\delta}, \underline{\delta} \ge 1$ since transmitting in synced states is costly in the considered scenarios. Fig. 3a shows the results of an asymmetric source with prioritized states. The surface is asymmetric and the low-cost region is where $\bar{\delta}$ is small but $\underline{\delta}$ is large. This implies that *the alarm state is significantly more important, thus forcing the sensor to trigger transmissions more frequently in MA errors.* On the other hand, Fig. 3b shows a special case where the source is symmetric and the states are equally important. In this case, the sensor has no preference for different states, and the source behaves in a balanced manner. Consequently, the surface is symmetric, and the optimal threshold-type policy $(\bar{\delta}^*, \underline{\delta}^*) = (5, 5)$ is globally optimal.

Fig. 3c and Fig. 3d plot the results for an asymmetric source with equally important states and for a symmetric source with prioritized states, respectively. In case (c), although the states are equally important, the source will stay in state 0 (normal state) for a larger fraction of time. Hence, a smaller $\underline{\delta}^*$ for AoFA is expected. In case (d), the source is symmetric but state 1 (alarm) is of great interest. It is observed that the threshold-type policy is suboptimal since it treats the states equally. By contrast, the optimal policy, i.e., $(\bar{\delta}^*, \underline{\delta}^*) = (1, 24)$, updates more frequently in MA errors. In outline, *our approach offers the flexibility to schedule transmissions according to data significance, which depends on both the state importance and the source pattern.*

Fig. 4 compares the performance gap between the threshold-type and switching-type policies, i.e., $\mathcal{L}^\lambda(\delta^*) - \mathcal{L}^\lambda(\bar{\delta}^*, \underline{\delta}^*)$. From Fig. 4a, we observe that, when the source states are equally important, the threshold-type policy demonstrates comparable performance if the difference between the stationary probabilities of state 0 and state 1, i.e., $\frac{|p-q|}{p+q}$, is relatively small. However, when the states are prioritized, as shown in Fig. 4b, the largest performance gap occurs when $\frac{|p-q|}{p+q}$ is small. This happens because in this case, the information significance depends mainly on the importance of the states. Consequently, the threshold-type policy can perform arbitrarily poorly as it treats all states equally.

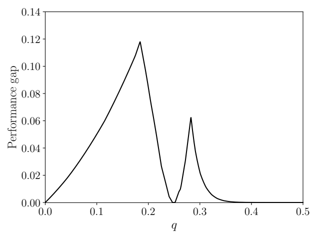

(a) $\beta = 0.5$.

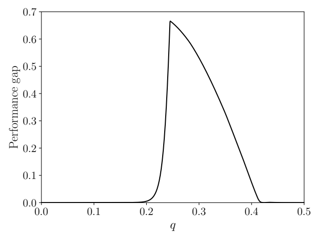

(b) $\beta = 0.8$.

Figure 4: The performance gap as a function of q when $p = 0.25, p_s = 0.9, \lambda = 8, N = 100$.

6 CONCLUSION

This paper studied the semantic-aware remote estimation of a discrete-state Markov source with prioritized states. We introduced two new age metrics (i.e., AoMA and AoFA) to account for the costs of different estimation errors. We identified the problem as a countably infinite state MDP with unbounded costs. We showed the existence of a switching-type optimal policy and derived analytical results. For numerical traceability, we proposed a finite-state approximate MDP and proved its asymptotical optimality. Numerical results underscored the effectiveness of exploiting data significance in such systems.

ACKNOWLEDGMENTS

This work has been supported in part by the Swedish Research Council (VR), ELLIIT, the European Union (ETHER) under Grant 101096526; the European Union's Horizon Europe Research and Innovation Programme under the Marie Skłodowska-Curie Grant under Agreement 101131481 (SOVEREIGN); by the Horizon Europe/JU SNS Project ROBUST-6G under Grant 101139068, and the Graduate School in Computer Science (CUGS).

REFERENCES

[1] John Baillieul and Panos J. Antsaklis. 2007. Control and Communication Challenges in Networked Real-Time Systems. *Proc. IEEE* 95, 1 (2007), 9–28. https://doi.org/10.1109/JPROC.2006.887290

[2] Jhelum Chakravorty and Aditya Mahajan. 2017. Fundamental Limits of Remote Estimation of Autoregressive Markov Processes Under Communication Constraints. *IEEE Trans. Autom. Control* 62, 3 (2017), 1109–1124. https://doi.org/10.1109/TAC.2016.2580589

[3] Yutao Chen and Anthony Ephremides. 2024. Minimizing Age of Incorrect Information Over a Channel With Random Delay. *IEEE/ACM Trans. Netw.* 32, 4 (2024), 2752–2764. https://doi.org/10.1109/TNET.2024.3389964

[4] Oswaldo Luiz Valle Costa, Marcelo Dutra Fragoso, and Ricardo Paulino Marques. 2005. *Discrete-time Markov jump linear systems.* Springer.

[5] Jennifer Gielis, Ajay Shankar, and Amanda Prorok. 2022. A critical review of communications in multi-robot systems. *Curr. Robot. Rep.* 3, 4 (2022), 213–225.

[6] Wilhelmus PMH Heemels, Karl Henrik Johansson, and Paulo Tabuada. 2012. An introduction to event-triggered and self-triggered control. In *Proc. IEEE Conf. Decision and Control.* 3270–3285.

[7] Joo P. Hespanha, Payam Naghshtabrizi, and Yonggang Xu. 2007. A Survey of Recent Results in Networked Control Systems. *Proc. IEEE* 95, 1 (2007), 138–162. https://doi.org/10.1109/JPROC.2006.887288

[8] K.V. Katsikopoulos and S.E. Engelbrecht. 2003. Markov decision processes with delays and asynchronous cost collection. *IEEE Trans. Autom. Control* 48, 4 (2003), 568–574. https://doi.org/10.1109/TAC.2003.809799

[9] Antzela Kosta, Nikolaos Pappas, Vangelis Angelakis, et al. 2017. Age of information: A new concept, metric, and tool. *Found. Trends Netw.* 12, 3 (2017), 162–259.

[10] Marios Kountouris and Nikolaos Pappas. 2021. Semantics-Empowered Communication for Networked Intelligent Systems. *IEEE Commun. Mag.* 59, 6 (2021), 96–102. https://doi.org/10.1109/MCOM.001.2000604

[11] Vikram Krishnamurthy. 2013. How to Schedule Measurements of a Noisy Markov Chain in Decision Making? *IEEE Trans. Inf. Theory* 59, 7 (2013), 4440–4461. https://doi.org/10.1109/TIT.2013.2253352

[12] Vikram Krishnamurthy. 2016. *Partially observed Markov decision processes: from filtering to controlled sensing.* Cambridge Univ. Press.

[13] Jiping Luo and Nikolaos Pappas. 2024. Exploiting Data Significance in Remote Estimation of Discrete-State Markov Sources. *arXiv preprint arXiv:2406.18270* (2024).

[14] Jiping Luo and Nikolaos Pappas. 2024. Semantic-Aware Remote Estimation of Multiple Markov Sources Under Constraints. *arXiv preprint arXiv:2403.16855* (2024).

[15] Ali Maatouk, Saad Kriouile, Mohamad Assaad, and Anthony Ephremides. 2020. The Age of Incorrect Information: A New Performance Metric for Status Updates. *IEEE/ACM Trans. Netw.* 28, 5 (2020), 2215–2228. https://doi.org/10.1109/TNET.2020.3005549

[16] Wenzhuo Ouyang, Atilla Eryilmaz, and Ness B. Shroff. 2016. Downlink Scheduling Over Markovian Fading Channels. *IEEE/ACM Trans. Netw.* 24, 3 (2016), 1801–1812. https://doi.org/10.1109/TNET.2015.2438009

[17] Nikolaos Pappas and Marios Kountouris. 2021. Goal-Oriented Communication For Real-Time Tracking In Autonomous Systems. In *Proc. IEEE Int. Conf. Auto. Syst.* https://doi.org/10.1109/ICAS49788.2021.9551200

[18] Pangun Park, Sinem Coleri Ergen, Carlo Fischione, Chenyang Lu, and Karl Henrik Johansson. 2018. Wireless Network Design for Control Systems: A Survey. *IEEE Commun. Surveys Tuts.* 20, 2 (2018), 978–1013. https://doi.org/10.1109/COMST.2017.2780114

[19] Petar Popovski, Federico Chiariotti, Kaibin Huang, Anders E. Kalør, Marios Kountouris, Nikolaos Pappas, and Beatriz Soret. 2022. A Perspective on Time Toward Wireless 6G. *Proc. IEEE* 110, 8 (2022), 1116–1146. https://doi.org/10.1109/JPROC.2022.3190205

[20] Mehrdad Salimnejad, Marios Kountouris, Anthony Ephremides, and Nikolaos Pappas. 2024. Version Innovation Age and Age of Incorrect Version for Monitoring Markovian Sources. *arXiv preprint arXiv:2401.17691* (2024).

[21] Mehrdad Salimnejad, Marios Kountouris, and Nikolaos Pappas. 2023. State-aware real-time tracking and remote reconstruction of a Markov source. *Journal Commun. and Netw.* 25, 5 (2023), 657–669. https://doi.org/10.23919/JCN.2023.000035

[22] Mehrdad Salimnejad, Marios Kountouris, and Nikolaos Pappas. 2024. Real-Time Reconstruction of Markov Sources and Remote Actuation Over Wireless Channels. *IEEE Trans. Commun.* 72, 5 (2024), 2701–2715. https://doi.org/10.1109/TCOMM.2024.3356456

[23] Linn I Sennott. 1998. *Stochastic dynamic programming and the control of queueing systems.* John Wiley & Sons.

[24] Ling Shi, Peng Cheng, and Jiming Chen. 2011. Sensor data scheduling for optimal state estimation with communication energy constraint. *Automatica* 47, 8 (2011), 1693–1698. https://doi.org/10.1016/j.automatica.2011.02.037

[25] George Stamatakis, Nikolaos Pappas, and Apostolos Traganitis. 2019. Control of Status Updates for Energy Harvesting Devices That Monitor Processes with Alarms. In *Proc. IEEE Global Commun. Conf. Workshops.* 1–6. https://doi.org/10.1109/GCWkshps45667.2019.9024463

[26] Yin Sun, Igor Kadota, Rajat Talak, and Eytan Modiano. 2019. Age of information: A new metric for information freshness. *Synth. Lectures Commun. Netw.* 12, 2 (2019), 1–224.

[27] G.C. Walsh and Hong Ye. 2001. Scheduling of networked control systems. *IEEE Control Syst. Mag.* 21, 1 (2001), 57–65. https://doi.org/10.1109/37.898792

[28] Roy D. Yates, Yin Sun, D. Richard Brown, Sanjit K. Kaul, Eytan Modiano, and Sennur Ulukus. 2021. Age of Information: An Introduction and Survey. *IEEE J. Sel. Areas Commun.* 39, 5 (2021), 1183–1210. https://doi.org/10.1109/JSAC.2021.3065072