Remote Monitoring of Two-State Markov Sources in Random Access Channels: Joint Model and State Estimation

Houman Asgari*, Andrea Munari†, Gianluigi Liva†, Giuseppe Cocco‡

*Technical University of Munich (TUM), Germany, houman.asgari@tum.de
†German Aerospace Center (DLR), Germany, {andrea.munari, gianluigi.liva}@dlr.de
‡Polytechnic University of Catalonia (UPC), Spain, giuseppe.cocco@upc.edu

***Abstract*—This paper addresses the problem of remote monitoring of two-state Markov sources via a slotted ALOHA random access channel, where the source statistics are not known a priori to the receiver. We develop a joint model and state estimation method using the Baum-Welch algorithm for two different transmission strategies. In the first strategy, the nodes transmissions are independent of the underlying state evolution process (random policy). In the second strategy, the nodes transmit an update only upon observing a state transition (reactive policy). We show that the reactive approach is beneficial not only in terms of reducing the state estimation error probability (a result that was recently established under perfect knowledge of the source statistics), but that it allows a faster learning of the source statistics.**

***Index Terms*—Remote Monitoring, Random Access, Slotted Aloha, Expectation-Maximization, Markov Sources**

I. Introduction

In the context of Internet of Things (IoT) systems, monitoring the state of remotely located nodes is a challenging task, especially in wireless sensor networks. These networks often consist of numerous low-complexity, battery-powered devices that collect data and transmit updates to a central hub over a shared channel. The sporadic and irregular nature of these transmissions, combined with energy limitations and constraints of shared communication channels, pose significant challenges on how to efficiently use channel resources.

Traditional grant-based access methods, which rely on channel negotiation and reservation procedures, are often inefficient due to the large overhead. As a result, uncoordinated access protocols, such as those based on slotted ALOHA [1], are commonly employed to enable connectivity. However, the challenge remains: how to ensure timely and accurate state estimation of the sensed processes, especially when the shared channel leads to packet collisions and communication delays.

To characterize these aspects, performance metrics like age of information (AoI) have been widely adopted. AoI captures the freshness of information by measuring the time elapsed since the last update from a process was received [2]. While AoI has proven effective for evaluating timely delivery of status updates, it does not account for the accuracy or the relevance of a delivered message. Thus, it falls short in scenarios where the monitored sources exhibit memory or non-trivial dynamics, such as Markov processes.

To address these limitations, other metrics and strategies have been proposed to enhance the accuracy and efficiency of state estimation in various applications. Among them, state estimation entropy and state estimation error probability (SEEP) have gained attention [3]–[6]. State estimation entropy quantifies the uncertainty at the receiver regarding the current state of a monitored process, while SEEP directly measures the probability that the receiver's state estimate is incorrect. In [7] additional semantic-aware metrics such as the Age of Missed Alarm (AoMA) and Age of False Alarm (AoFA), accounting for the costs associated with different types of estimation errors were introduced. In turn, [8] explored the role of sampling rates in maintaining information freshness. The authors provided closed-form expressions for mean information freshness and optimized sampling strategies to ensure up-to-date and accurate state estimates. Similarly, [9] proposed a randomized stationary sampling policy to minimize actuation errors in Markov source monitoring. These results demonstrate the effectiveness of state-aware policies in optimizing sampling rates and reducing errors in real-time applications. Finally, [10] addresses the challenge of state estimation in networked control systems, studying the impact of packet loss and time delay on the transmission of sensory information from a stochastic source to a decoder.

In this work, we focus on the remote monitoring of two-state Markov sources over a random access channel, where destructive collisions occur, preventing successful packet delivery. This problem was initially introduced and analyzed in [6] where the receiver is assumed to have perfect knowledge of the sources' model. However, in practical scenarios, such a priori knowledge may not always be available. In this paper, we address this limitation by developing a joint model and state estimation approach, which extends the maximum a posteriori (MAP) detectors from [6]. Our approach alternates between a state estimation step and a model refinement step using the Baum-Welch algorithm [11], [12], enabling the receiver to dynamically learn both the state transition probabilities and

H. Asgari acknowledges the financial support of the Munich Aerospace scholarship within the group "Multi-access and Security Coding for Massive IoT Satellite Systems".

A. Munari and G. Liva acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the programme of "Souver an. Digital. Vernetzt." Joint project 6G-RIC, project identification number: 16KISK022.

G. Cocco acknowledges the financial support by the Ramon y Cajal fellowship program, grant RYC2021-033908-I, funded by the Spanish Ministry of Science and Innovation through MCIN/AEI/10.13039/501100011033 and by the European Union through the "NextGenerationEU" Recovery Plan for Europe.

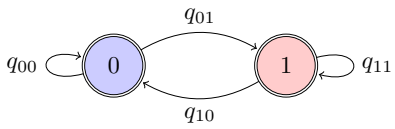

Fig. 1. Two-state Markov model for a generic source in the system.

the current state of the source. We investigate the performance of two transmission strategies: a random transmission policy, where nodes send updates independently of their state with a fixed probability, and a reactive transmission policy, where updates are only sent when a change in the source's state is detected. We consider both symmetric and asymmetric sources. The results indicate that the estimation performance for asymmetric sources is better than the symmetric ones in both transmission policies. Between the two policies, the reactive one shows superior performance in reducing state estimation errors and model mismatch.

II. System Model

A. Notation

We represent random variables using uppercase letters and their realizations with lowercase letters. The probability of the event $X = x$ is indicated as $\mathsf{P}[X = x]$, and the probability mass function for the random variable X is written as $P(x) = \mathsf{P}[X = x]$. For discrete-time, finite-state Markov chains, the one-step transition probability from state i to state j is represented by q_{ij}, while the stationary probability of state i is denoted by π_i. We finally denote by $Q = [q_{ij}]$ the state transition probability matrix.

B. System Model

We consider a system with M statistically independent sources (nodes) sharing a common wireless channel to communicate with a receiver. Time is slotted, and all nodes are synchronized to these slots. Each source $k \in \{0, 1, \ldots, M-1\}$ generates a random sequence of symbols

$$X_0^{(k)}\, X_1^{(k)}\, X_2^{(k)}\, \ldots$$

where $X_n^{(k)}$ belongs to the alphabet $\mathcal{X} = \{0, 1\}$, representing the random state of the source at time slot n. Each node operates as a homogeneous two-state stationary Markov chain, with transition probabilities

$$q_{ij}^{(k)} = \mathsf{P}\left[X_n^{(k)} = j \mid X_{n-1}^{(k)} = i\right]$$

for all $(i, j) \in \mathcal{X} \times \mathcal{X}$, as shown in Fig. 1. The steady-state distribution of the process is given by

$$\pi_0 = \frac{q_{10}}{q_{10} + q_{01}}, \quad \pi_1 = \frac{q_{01}}{q_{10} + q_{01}}. \tag{1}$$

Nodes transmit update packets over the shared channel, reporting their state to the receiver. We focus on random access policies, specifically examining two variations of the slotted ALOHA protocol [13], detailed in Sec. II-C.

Accordingly, three slot outcomes are possible at the receiver: i) idle, where no transmission occurs; ii) singleton, where only one source transmits; and iii) collision, where multiple packets are transmitted simultaneously. We follow the collision channel model [13], where a status update is received correctly in a singleton slot but cannot be decoded in the event of a collision.

Furthermore, we assume that the receiver can detect collisions but cannot determine how many packets were involved. Each transmission includes a source identifier, enabling the receiver to know the current state of the source upon decoding. Moreover, re-transmissions are not considered, as nodes do not receive feedback about their transmission attempts.

The receiver tries to maintain an accurate real-time understanding of the state of each source. To this aim, it is aware that nodes evolve as two-state Markovian processes, but it does not have any prior knowledge about the state-transition probabilities q_{ij} or their stationary distributions. For simplicity, we focus on the source with index $k = 0$, and represent the sequence of symbols it generates from time 0 to n by

$$X^n = X_0\, X_1\, X_2\, \ldots X_n.$$

During this time, the receiver observes the output sequence

$$Y^n = Y_0\ Y_1\ Y_2\ \ \ldots Y_n$$

where Y_n belongs to the alphabet $\mathcal{V} = \{0, 1, \mathtt{S}, \mathtt{C}\}$. Here, 0 and 1 represent the observed states of the reference source without collision. S instead, indicates a slot where the reference source did not transmit, i.e. it corresponds to either an idle slot, or a slot where a message from a node other than the reference one was decoded. Finally, C indicates a collision between transmissions from two or more sources, which may or may not involve the reference source.

C. Transmission Strategies

We explore two different transmission policies: random and reactive.

- **Random Transmission Policy**: In this approach, every node decides to transmit or refrain from transmission with probabilities α and $1 - \alpha$, respectively. Hence, the transmission decision is independent of how the source evolves. In this configuration, the number of messages transmitted within a generic time slot follows a binomial distribution with parameters M and α. Therefore, the probability for a source to deliver an update is:

  $$\omega = \alpha(1 - \alpha)^{M-1}. \tag{2}$$

  In this setting, the activation probability α is the same across all sources and is selected to be $\alpha = M^{-1}$ to maximize the successful decoding probability stated in equation (2).
- **Reactive Transmission Policy**: In this approach, transmissions are only triggered by a change in the status of the source. Accordingly, the probability for a source to

deliver an update over a generic slot can be captured by the following approximation:

$$\tilde{\omega} \simeq \tilde{\alpha}\,(1-\tilde{\alpha})^{M-1} \tag{3}$$

where

$$\tilde{\alpha} := \pi_0 q_{01} + \pi_1 q_{10} = \frac{2q_{01}q_{10}}{q_{01}+q_{10}} \tag{4}$$

is the average activation probability for a node under the reactive transmission strategy. Note that the expression will be exact in case of symmetric sources, i.e., $q_{01} = q_{10}$. Otherwise, the probability of success jointly depends on the state of all sources at this time slot. The approximations in (3)-(4) has anyways been shown to be very tight in the considered setting [6].

Such an event-triggered policy introduces an important trade-off. On the one hand, it can mitigate the channel congestion compared to a random transmission strategy by avoiding transmission in the absence of a state change. Furthermore, observations of slots that do not deliver a message from the reference source will still be informative. For instance, the channel observation $\mathtt{S}$ is indicator of no-state change for the reference source. On the other hand, when the source changes state infrequently, collisions can result in prolonged intervals without successfully transmitting updates, causing delays in the delivery of new information.

To effectively analyze the transmission strategies in our system, we formulate the problem using a hidden Markov model (HMM). In fact, the state of the source—whether it has changed or not—cannot be directly observed by the receiver; instead, the receiver can only infer it from the channel observations. For the random transmission strategy, the observations $\mathtt{S}, \mathtt{C}$ can both be regarded as non-informative. In other words, the occurrence of a collision or successful delivery of another source does not provide any information about the evolution of the reference source. Moreover, the probability of an output observation depends only on the current state and not on the prior ones. As a result, the emission probabilities under the random transmission strategy can be characterized as:

$$\mathsf{P}(Y_n|X^n) = \mathsf{P}(Y_n|X_n) = \begin{cases} \omega & \text{if } Y_n = X_n \\ 0 & \text{if } Y_n = |1-X_n| \\ 1-\omega & \text{Otherwise.} \end{cases} \tag{5}$$

In (5) the first condition corresponds to decoding the current state of the reference sources, which happens with probability ω. In turn, the second specifies that $\mathsf{P}(Y_n = 0|X_n = 1) = \mathsf{P}(Y_n = 1|X_n = 0) = 0$, as the source only transmits its current value.

On the other hand, with the reactive transmission strategy, observation $Y_n \in \{\mathtt{S}, \mathtt{C}\}$ can provide side information about the state of the sources. For instance, by observing $Y_n = \mathtt{S}$ the receiver can infer that $X_n = X_{n-1}$. In other words, the probability of each emission depends on the evolution of all the sources. In fact, the number of other sources in a given state (e.g. state 0) along with the state of the reference source will form the state space of the HMM. This will lead to an HMM with $2M$ states and a size $4 \times (2M)^2$ emission probability matrix. To simplify the analysis of the HMM, we follow the approach mentioned in [4], [6] to approximate it with a surrogate *myopic* model where we assume that only the reference source employs the reactive transmission approach and the other $M-1$ sources adopt a random transmission policy with activation probability $\tilde{\alpha}$. This in turn will result in an HMM with 2 states and an emission probability matrix with 4×2^2 dimension. Using this myopic approximation, if the reference source does not change state (i.e., $X_n = X_{n-1}$), the only possible outcomes are an idle or collision observation. The probability of observing an idle slot is

$$\begin{aligned}\mathsf{P}(Y_n = \mathtt{S}|X_n = X_{n-1}) = &(1-\tilde{\alpha})^{M-1} \\ &+ (M-1)\tilde{\alpha}(1-\tilde{\alpha})^{M-2}\end{aligned} \tag{6}$$

while the probability of a collision is

$$\mathsf{P}(Y_n = \mathtt{C}|X_n = X_{n-1}) = 1 - \mathsf{P}(Y_n = \mathtt{S}|X_n = X_{n-1}). \tag{7}$$

Conversely, if the reference source changes state, the outcome is either a successful delivery with probability

$$\mathsf{P}(Y_n = X_n|X_n = |1-X_{n-1}|) = (1-\tilde{\alpha})^{M-1} \tag{8}$$

or a collision with probability

$$\begin{aligned}\mathsf{P}(Y_n = \mathtt{C}|X_n = &|1-X_{n-1}|) \\ &= 1 - \mathsf{P}(Y_n = X_n|X_n = |1-X_{n-1}|).\end{aligned} \tag{9}$$

D. Maximum A Posteriori Estimator

A *maximum a posteriori* (MAP) estimator is used to estimate the hidden state of a source based on observed data. In our case, the MAP estimator finds the most likely state of the source, $\hat{x}_n$, at time n given the sequence of channel observations y^n. The MAP estimate is defined as:

$$\hat{x}_n = \underset{x_n \in \{0,1\}}{\arg\max}\ \mathsf{P}(x_n \mid y^n). \tag{10}$$

Using this estimate, we also study the state estimation error probability (SEEP), which refers to the probability that the estimated state differs from the true state. The metric is defined as

$$\mathsf{SEEP} = \frac{1}{N}\sum_{n=0}^{N-1} \mathsf{P}\left[\hat{X}_n \neq X_n\right], \tag{11}$$

where N is the duration of the considered time horizon.

E. Kullback-Leibler Divergence

To assess the accuracy of the receiver's estimation of the transition probabilities, we use the Kullback-Leibler (KL) divergence [14]. More specifically, in our context, the KL divergence is used to quantify the discrepancy between the true transition probability matrix Q of the Markov source and the estimated transition probability matrix $\hat{Q}$ obtained from the Baum-Welch algorithm, as discussed in Sec. III.

The KL divergence between $\hat{Q}$ and Q is given by [15]:

$$D_{\mathrm{KL}}(Q \parallel \hat{Q}) = \sum_{i \in \mathcal{X}} \pi_i S\left(X_2 \mid X_1 = i\right), \tag{12}$$

where the conditional divergence $S(X_2 \mid X_1 = i)$ is:

$$S\left(X_2 \mid X_1 = i\right) := \sum_{j \in \mathcal{X}} q_{ij} \log \frac{q_{ij}}{\hat{q}_{ij}}. \tag{13}$$

III. Joint Model and State Estimation

The Baum-Welch (BW) algorithm [11], [12], a special case of the expectation-maximization (EM) algorithm [16], is employed to iteratively estimate the parameters of the reference Markov chain that best explain the observed data. Once these parameters are estimated, the algorithm also helps in computing the most likely state of the reference source, which in our case is the variable $\hat{X}_n$. The BW algorithm involves two main steps, performed in successive manner:

- **Expectation Step (E-Step)**: the algorithm calculates the expected value of the likelihood of the observed data, given the current estimates of the model parameters. This step requires calculating two key sets of values that capture the likelihood of being in a particular state at a given time, based on both recent and future observations. These values are computed using the forward-backward algorithm [12], [17].
- **Maximization Step (M-Step)**: Using the expected values computed in the E-Step, the algorithm re-estimates the model parameters to maximize the likelihood of the observed data. This includes re-estimating the transition probabilities between states and the emission probabilities that relate the hidden states to the observed data.

The algorithm begins with an initial guess for the transition probability matrix $\hat{Q}$ and the emission probabilities given the transmission strategy implemented for channel contention, and recursively performs the aforementioned steps to refine these estimates until convergence. At every slot, we run the BW algorithm using the entire output sequence up to that slot, in order to refine the model estimation and estimate the latest state of the reference source with such model.

IV. Simulation Results and Discussions

In this section, we present the results of simulations conducted to compare the performance of the two transmission strategies and assess the impact of model mismatch on state estimation. The estimator starts with arbitrary initial guess of $\hat{q}_{01} = \hat{q}_{10} = 0.5$. Monte Carlo simulations were run for a duration of 10^4 time slots and repeated across 100 realizations. We varied the number of sources (nodes) in the network, with node counts of $M = 10$ and $M = 100$ nodes, to explore the effects of network size on convergence and estimation accuracy. Furthermore, we consider both symmetric and asymmetric sources. We aim to provide insight into the interplay between these factors and determine which transmission strategy performs better under different network and source conditions.

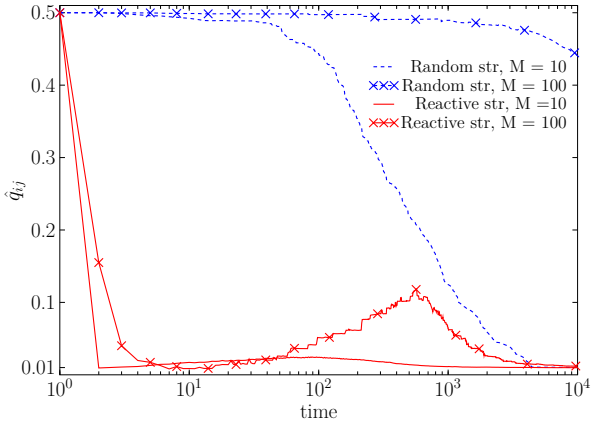

Fig. 2. The estimated state transition probability evolution when the source is symmetric with actual transition probability of $q_{01} = q_{10} = 0.01$. The random and reactive strategy convergence are showed with blue and red curves, respectively.

A. Symmetric Sources

We start with symmetric sources to clarify how the receiver's estimation of the transition probability q_{ij}s evolve. Considering sources with transition probabilities $q_{01} = q_{10} = 0.01$, Figure 2 reports the evolution of the mean estimates $\hat{q}_{ij}$ over time. The blue dashed lines represent the random transmission policy, while the red solid lines correspond to the reactive transmission policy. Plots with markers indicate $M = 100$ nodes, whereas unmarked curves represent $M = 10$ nodes. From this figure, it is evident that when the reactive transmission policy is employed, $\hat{q}_{ij}$ will immediately drop to a low value. In fact, with the initial guess $\hat{q}_{ij} = 0.5$ the receiver expects that the sources will transmit at every slot with a probability of 0.5. Such belief immediately changes when the receiver observes the earliest channel outputs, which are likely to be idle. As time evolves, the estimator gradually refines its underestimated $\hat{q}_{ij}$. The more successfully decoded packets and collisions are observed, the more the estimator tends to increase its estimate for the transition probabilities. On the other hand, under a random transmission strategy the observation of collisions or idle slots does not affect the belief about the $\hat{q}_{ij}$s and the estimator can improve its estimate only by having more successfully decoded updates of the reference source. This leads to a steadier, yet much slower, convergence of the $\hat{q}_{ij}$. Moreover, we see that by increasing the network size, the convergence speed of both policies decrease, due to reduced informative observations caused by the harsher contention.

As the BW algorithm converges to the true transition probabilities q_{ij}, the MAP estimator naturally achieves the same SEEP as it would if the receiver had direct access to the source's model. Let us know focus on the performance of the estimated models during the BW training process. We aim at understanding how discrepancies between the estimated and true models impact the SEEP of the state estimator while the

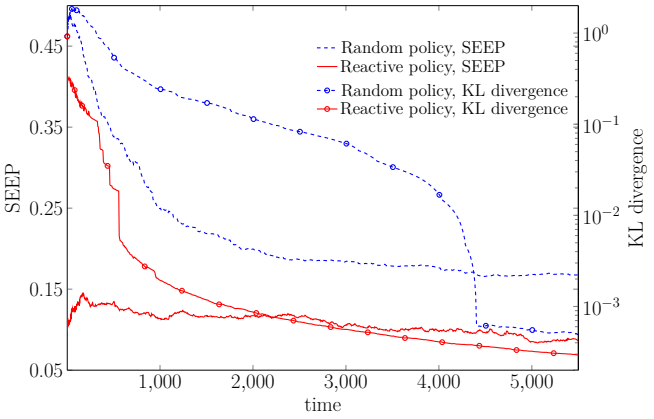

Fig. 3. SEEP and KL divergence of the approximated models when the source is symmetric with actual transition probability of $q_{01} = q_{10} = 0.01$. The random and reactive strategy convergence are shown with blue and red curves, respectively.

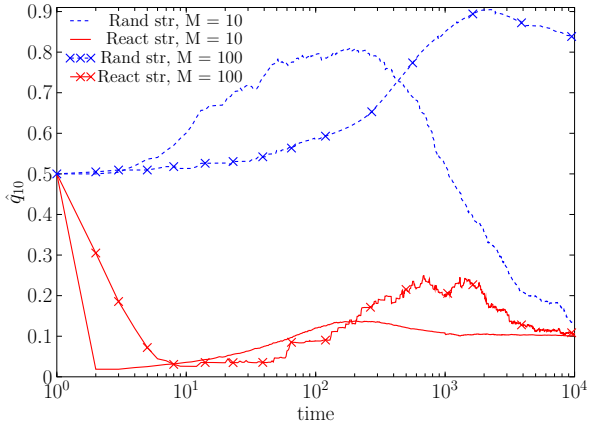

Fig. 4. Estimated value of q_{10} vs. time. Asymmetric sources with actual transition probability of $q_{01} = 0.01$ and $q_{10} = 0.1$. The reactive and random strategy convergence are shown with blue and red curves, respectively.

BW algorithm is still refining its estimates. To explore this, we propose the following approach: we halt the BW training at an arbitrary time n, even if the current estimate of the model may still differ from the true source dynamics. From time n onward, we use the model obtained at this intermediate stage (i.e. $\hat{Q}_n$) to continue estimating the reference source's state for $n+1, n+2, \ldots$. This approach allows us to examine how the SEEP of this mismatched model, obtained at time n, compares to that of the true model. By evaluating the SEEP during different stages of training, we can explore the sensitivity of the MAP estimator to model imperfections and determine the point at which further BW iterations yield diminishing returns in terms of state estimation accuracy, and accordingly gain insights into the trade-off between the length of the training period and the accuracy of state estimation.

Figure 3 illustrates this comparative performance across different stages of model refinement. The curves without markers show for each training duration n the SEEP obtained using the corresponding model, whereas the ones with markers delineate the KL divergence of the estimated models from the true source's model. The figure highlights how the SEEP sensitivity to model mismatch differs under reactive and random transmission policies. The reactive strategy's KL divergence stabilizes quite quickly. In the case of the random strategy, the approximated transition probabilities $\hat{q}_{10}$ and $\hat{q}_{01}$ might be far from the actual ones, but relatively close to each other (e.g., 0.4 for both transitions), which yields a SEEP close to the one with exact model knowledge.

B. Asymmetric Sources

Consider now a setting where the sources are asymmetric, with $q_{10} = 0.1$ and $q_{01} = 0.01$. Figure 4 shows the evolution over time of the estimate $\hat{q}_{10}$ using the BW algorithm. The blue dashed lines represent the random transmission policy, and the red solid lines correspond to the reactive transmission strategy. The plots with markers indicate $M = 100$ nodes, while the unmarked curves represent $M = 10$ nodes. Similar to the symmetric case, the reactive transmission strategy consistently leads to faster convergence and performs better in handling larger networks compared to the random transmission strategy. A notable fact can be seen in the random transmission strategy. Since the source spends more time in state 0—due to the higher likelihood of transitions from state 1 to state 0—the BW estimator initially attributes the more frequent 0 observations to a higher transition probability from state 1 to 0. This leads to inflated estimates of $\hat{q}_{10}$. As more data is observed, the overestimation gradually decreases, and the estimator converges toward the true transition probabilities.

Finally, Figure 5 shows the evolution of SEEP and KL divergence for 10 asymmetric sources. Curves and markers are akin to those described for Figure 3. The results indicate that while the reactive transmission strategy leads to faster convergence of the estimated parameters, it takes longer for the BW estimator to achieve the SEEP that would be expected with exact model knowledge. This is primarily due to the fact that, under the reactive transmission policy, discrepancies in the estimated model also affect the emission probabilities. Since these probabilities depend on the source's state transition probability matrix, any inaccuracies in the model parameter estimate introduce errors in the approximated transition probabilities as well.

In contrast, under the random transmission strategy, the emission probability matrix is independent of the state transition parameters. Therefore, even in the presence of model mismatches, the emission probabilities remain accurate. This allows the BW algorithm to perform more consistently during the estimation process. This fundamental difference in how the emission matrix is affected by model discrepancies explains why, despite the reactive strategy's faster parameter convergence, it can initially result in higher SEEP compared to the random strategy until model estimates are more accurate.

Furthermore, compared to the case of symmetric sources,

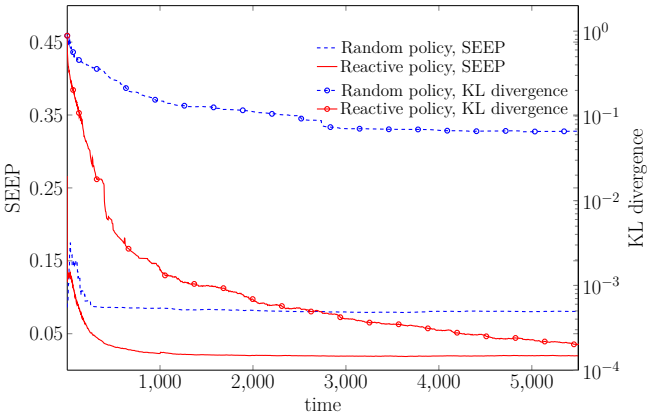

Fig. 5. SEEP and KL divergence of the approximated models when the source is asymmetric with actual transition probability of $q_{01} = 0.01$ and $q_{10} = 0.1$. The random and reactive strategy convergence are shown with blue and red curves, respectively.

the SEEP is more robust against model imperfections in the case of asymmetric sources. This robustness stems from the imbalance between the transition probabilities, which allows the MAP estimator to distinguish the states more effectively. As a result, the state estimator yields a performance closer to that of the MAP estimator with perfect model knowledge, even before the model estimation algorithm fully converges.

Conclusions

In this paper, we investigated the performance of random and reactive transmission policies for joint model and state estimation of a two-state Markov source with unknown transition probabilities over slotted ALOHA channels, focusing on the SEEP and KL divergence of the estimated states and models. The results indicate the superiority of the reactive strategy in both tasks. Future studies will consider higher-order Markov sources and enhanced transmission protocols for this scenario.

References

[1] N. Abramson, "The ALOHA System - Another Alternative for Computer Communications," in *Proc. 1970 Fall Joint Computer Conference*. AFIPS Press, 1970.

[2] R. Yates and S. Kaul, "The age of information: Real-time status updating by multiple sources," *IEEE Trans. Inf. Theory*, vol. 65, no. 3, pp. 1807–1827, 2019.

[3] M. Rezaeian, B. Vo, and J. Evans, "The optimal observability of partially observable Markov decision processes: Discrete state space," *IEEE Trans. Autom. Control*, vol. 55, no. 12, pp. 2793–2798, 2010.

[4] G. Cocco, A. Munari, and G. Liva, "State estimation entropy for two-state Markov sources in slotted ALOHA random access channels," in *Proc. IEEE Info. Theo. Workshop (ITW)*, France, Apr. 2023.

[5] 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.

[6] G. Cocco, A. Munari, and G. Liva, "Remote monitoring of two-state Markov sources via random access channels: An information freshness vs. state estimation entropy perspective," *IEEE J. Sel. Topics Inf. Theory.*, vol. 4, pp. 651–666, 2023.

[7] J. Luo and N. Pappas, "Exploiting data significance in remote estimation of discrete-state Markov sources," *arXiv preprint arXiv:2406.18270*, 2024.

[8] N. Akar and S. Ulukus, "Timely monitoring of Markov chains under sampling rate constraints," in *Proc. IEEE ICC*, Denver, CO, USA, June 2024, pp. 189–194.

[9] M. Salimnejad, M. Kountouris, and N. Pappas, "State-aware real-time tracking and remote reconstruction of a Markov source," *Journal of Communications and Networks*, vol. 25, no. 5, pp. 657–669, 2023.

[10] T. Soleymani, J. Baras, and K. Johansson, "State estimation over delayed and lossy channels: An encoder-decoder synthesis," *IEEE Transactions on Automatic Control*, vol. 69, pp. 1568–1583, 2023.

[11] L. E. Baum, T. Petrie, G. Soules, and N. Weiss, "A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains," *The Annals of Mathematical Statistics*, vol. 41, no. 1, pp. 164–171, 1970.

[12] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, Feb. 1989.

[13] N. Abramson, "The throughput of packet broadcasting channels," *IEEE Trans. Commun.*, vol. COM-25, no. 1, pp. 117–128, 1977.

[14] S. Kullback and R. A. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 1951.

[15] Z. Rached, F. Alajaji, and L. Campbell, "The Kullback-Leibler divergence rate between Markov sources," *IEEE Trans. Inf. Theory*, vol. 50, no. 5, pp. 917–921, 2004.

[16] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, no. 1, pp. 1–38, 1977.

[17] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv, "Optimal decoding of linear codes for minimizing symbol error rate (corresp.)," *IEEE Trans. Inf. Theory*, vol. 20, no. 2, pp. 284–287, 1974.