Remote Monitoring of Markov Sources over Random Access Channels: False Alarm and Detection Probability

Andrea Munari, Giuseppe Cocco, Gianluigi Liva

***Abstract*—We study the problem of remote source monitoring in an Internet of Things (IoT) system where a set of devices share a wireless channel to a common receiver. Each device observes an independent two-state Markov chain, with one of the states visited sporadically (modeling a critical event), and may transmit the current source value following a slotted ALOHA contention. We focus on protocols that set the transmission probability over a slot based on the value of the monitored process over the current and past slot. In turn, the receiver estimates the source state leveraging the channel outputs leaning either on a simple *decode and hold* approach, which requires no knowledge of the source statistics, or a *maximum a posteriori* estimator. For both approaches, we derive an analytical characterization of the system behavior in terms of *false alarm* and *detection probability*, deriving interesting insights and highlighting protocol design hints that depart from those commonly employed for throughput or age of information optimization.**

I. Introduction

INTERNET of things (IoT) systems are often employed to remotely monitor the evolution of a processes of interest. An example is provided by wireless sensor networks, deploying a large number of battery-powered, low-complexity devices that report acquired measurements to a collecting unit for tracking, control or actuation. Applications are broad, including environmental and industrial monitoring, asset tracking, smart city and cyber-physical systems. In these settings, maintaining an accurate perception at the monitoring end is not trivial, as often devices transmit information relying on uncoordinated (random) access protocols over a shared channels. This is the case, for instance, of LoRa [1] or Sigfox [2] networks, implementing simple variation of the ALOHA scheme [3] that render communications prone to packet losses due to collisions.

In order to properly capture the performance of remote monitoring systems, a number of novel metrics have been introduced in the past few years. In this respect, a pioneering role has been played by the age of information (AoI) [4], [5], which evaluates how old is the latest information available at the receiver regarding a process of interest. AoI has proven to be an excellent proxy to characterize some fundamental design trade-offs [6], jointly depending on both the frequency at which sources generate updates and on the time it takes for such messages to cross the network and reach the final destination. On the other hand, the metric solely focuses on timeliness, and is oblivious on the actual state of a process. To overcome this gap, alternative indicators that capture the reliability of the estimate at the monitor can be considered. Among them, the age of incorrect information (AoII) [7] introduces a penalty which grows over time only if the knowledge at the receiver is not deemed accurate, and can be complemented by tackling different flavors of reconstruction errors [8]–[10]. Along a similar line, information theory-inspired metrics have also been considered [11]–[13], focusing on the uncertainty at the receiver side through the notion of entropy. Notably, while a good level of understanding has been reached leaning on these concepts for point-to-point links and coordinated multi-access networks, somewhat less attention has been devoted to the practically relevant setup of uncoordinated access channel contention. Early results have been obtained for AoI [14]–[23], and have recently been extended to AoII [24], [25] and to the notion of state-estimation entropy [13].

Taking the lead from this, we further explore in this paper the design of random access protocols for remote source monitoring, considering a set of devices that sense independent processes, modelled as two-state Markov chains, and report to a common receiver via a slotted ALOHA channel. In particular, we focus on the case in which one of the two states is a critical one and is visited sporadically. In practice, knowing when the system is in such state might be of paramount importance and possibly require some action in response. Accordingly, we consider as system performance the probability of *false alarm* events, in which the monitor erroneously assumes a critical condition for a source, and *detection* events, i.e., a visit to the critical state is correctly identified. For this setup, we study a family of access policies in which the probability that a terminal transmits an update at a given slot depends on the past and current state of the observed source. Moreover, at the receiver side we consider two different approaches to estimate the state of remote sources based on the channel output. On the one hand, we consider a simple approach, dubbed *decode and hold*, in which the current estimate is only updated upon reception of a new status update from the source of interest. On the other hand, a *maximum a posteriori* solution is considered, which leverages the knowledge of source statistics, channel access policy and network cardinality to perform an estimation based on past and present channel outputs. In both cases, we provide an analytical characterization of false alarm and detection probabilities, and optimize the channel access policies' parameters. The study reveals interesting and non-trivial insights, clarifying how tying the channel access policy

A. Munari and G. Liva are with the Institute of Communications and Navigation, German Aerospace Center (DLR), Wessling, Germany (email: {andrea.munari, gianluigi.liva}@dlr.de).

Giuseppe Cocco is with the Department of Signal Theory and Communications, Universitat Politècnica de Catalunya (UPC), 08034, Barcelona, Spain (e-mail: giuseppe.cocco@upc.edu).

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

G. Cocco is supported by the Ramon y Cajal fellowship program (grant RYC2021-033908-I) funded by MCIN/AEI/10.13039/501100011033 and by the European Union "NextGenerationEU" Recovery Plan for Europe.

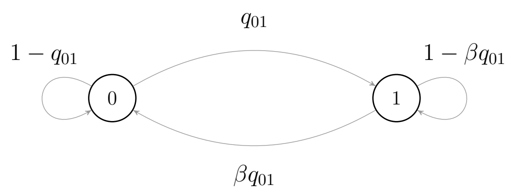

Fig. 1. Two-state Markov chain describing the evolution of a source.

to source evolution can significantly enhance the performance of remote monitoring systems, and thus providing design criteria that depart from the ones well-known for throughput and AoI optimization.

Notation:

We denote discrete random variables (r.v.s) and their realizations with capital and lowercase letters, respectively, e.g., X, x. The probability mass function of a r.v. is indicated as $P(x)$, and the probability of an event as $\mathsf{P}[X = x]$. For a discrete-time, finite state Markov chain, we denote the one-step transition probability between state i and j as q_{ij}, and the stationary probability of state i as π_i. Moreover, for a discrete-time stochastic process Y_n, $n \in \mathbb{N}$, the random vector $[Y_1 Y_2 \dots Y_n]$ is indicated as Y^n, and its realization as y^n.

II. System Model and Preliminaries

We consider a system in which M terminals (nodes) share a wireless channel towards a common receiver. Each node monitors a stochastic process (source), and may transmit data packets containing the value of the current source state. Time is divided in slots of equal duration, tuned to fit a wireless message, and all terminals are assumed to be slot-synchronized. The tracked processes are modeled as two-state, i.i.d. discrete-time Markov chains, evolving at every slot and taking values in $\mathcal{X} = \{0, 1\}$. As reported in Fig. 1, the transition probabilities between different states are denoted by q_{01} and $q_{10} = \beta q_{01}$. Accordingly, the stationary distribution of the process can be expressed as

$$\pi_0 = \frac{\beta}{1+\beta}, \qquad \pi_1 = \frac{1}{1+\beta}.$$

In the remainder we will focus on the case $q_{01} \ll 1$, $\beta \gg 1$, corresponding to sporadic and short visits to state 1. This choice is representative of systems which may at times deviate from a normal condition (state 0) to enter a critical state, which may require attention at the receiver side.

Packet transmissions are performed following an uncoordinated medium access policy based on slotted ALOHA [3], so that a terminal sends an update over each slot with a given probability, irrespective of the behavior of other contenders. Details on the considered strategiess will be presented in Sec. II-A. Following the well established collision channel model, we assume that the receiver cannot decode any message from a slot in which two or more terminals transmitted concurrently (collision), whereas a packet sent over a singleton slot is always successfully received.

Without loss of generality we concentrate on a terminal of interest, and denote the value of the corresponding source at slot n as X_n. In turn, the receiver observes at time n the output Y_n, taking values in $\mathcal{Y} = \{0, 1, \mathtt{I}, \mathtt{C}\}$. Specifically, 0 and 1 denote the correct reception of a message from the reference terminal, communicating the corresponding state of the source; $\mathtt{C}$ indicates that a collision has taken place over the slot; and $\mathtt{I}$ accounts for the case in which either no message was transmitted (idle slot), or a message from a terminal different from the reference one was successfully decoded.

Leveraging the observed channel outputs, the receiver updates at each time slot an estimate of the state of the source of interest. We refer to the corresponding process estimates as $\widehat{X}_n$, and focus on two different estimators, introduced in the following.

Decode and hold (D&H) estimator: The receiver updates $\widehat{X}_n$ only upon receiving a status update from the source of interest, and retains the previous estimate otherwise:

$$\widehat{X}_n = \begin{cases} Y_n & \text{if } Y_n \in \{0,1\} \\ \widehat{X}_{n-1} & \text{otherwise}. \end{cases}$$

The solution does not require any knowledge of the source statistics, and can be particularly appealing in practical systems in view of its simplicity.

Maximum a posteriori (MAP) estimator: If transition probabilities, network cardinality, and channel access strategy are known, the receiver can exploit the whole observed sequence y^n to refine its estimate. To this aim, we introduce the a posteriori probability (APP) logarithmic ratio

$$\lambda_n := \ln \frac{\mathsf{P}[X_n = 0 \,|\, Y^n = y^n]}{\mathsf{P}[X_n = 1 \,|\, Y^n = y^n]} \tag{1}$$

and determine the estimate based on a threshold test:

$$\widehat{X}_n = \begin{cases} 0 & \text{if } \lambda_n \geq \theta \\ 1 & \text{if } \lambda_n < \theta \end{cases} \tag{2}$$

where ties are resolved in favor of the hypothesis $X_n = 0$. In both cases, we are interested in evaluating the *detection probability* $\mathsf{P}_{\mathsf{det}}$, i.e., the probability of correctly identifying the source being in state 1, as well as the *false alarm probability* P_{fa}, capturing the case in which the receiver erroneously estimates the source to be in the critical state. Specifically, we will characterize

$$\mathsf{P}_{\mathsf{det}} := \mathsf{P}\big[\widehat{X}_n = 1 \,|\, X_n = 1\big], \ \mathsf{P}_{\mathsf{fa}} := \mathsf{P}\big[\widehat{X}_n = 1 \,|\, X_n = 0\big].$$

Remark 1. *By properly selecting the threshold θ, the test on the APP ratio reported in* (2) *satisfies the Neyman-Pearson criterion. Thus, the considered MAP estimator maximizes the detection probability for a target false alarm probability [26].*

A. Medium Access Strategies

We focus on channel access strategies by which each node independently decides whether to send a status update based on the state of the monitored process on the current and previous slot. More formally, a terminal transmits the value of the source over slot n with probability $\tau_{x_{n-1}x_n}$, and the access strategy is fully specified by the vector $\boldsymbol{\tau} = [\,\tau_{00}, \tau_{01}, \tau_{10}, \tau_{11}\,]$.

This general definition embeds some special examples of particular interest, which we highlight in the following:

- *random transmission:* in this case, the access policy is agnostic of the source state, so that $\tau_{ij} = \tau$, $\forall i, j \in \{0,1\}$. The scheme offers a relevant benchmark, as it is representative of common implementations of slotted ALOHA in which the access probability does not depend on the content of the message to be transmitted;
- *state-based transmission:* the terminal behavior is solely determined by the current state of the source, regardless of the past, i.e. $\tau_{00} = \tau_{01} := \tau_0$, and $\tau_{10} = \tau_{11} := \tau_1$;
- *reactive transmission:* a status update is sent each time, and only when, the source transitions, i.e., $\tau_{00} = \tau_{11} = 0$, $\tau_{01} = \tau_{10} = 1$. This approach binds the access policy to the behavior of the monitored processes, and aims at reducing channel congestion (and thus collision probability) by only delivering messages in the presence of a source change;
- *balanced reactive transmission:* as a variation of the previous solution, transmissions can be performed only when the source changes state, although with a certain probability. The policy is described by the vector $\boldsymbol{\tau} = [\,0, \tau_{01}, \tau_{10}, 0\,]$, where in general $\tau_{01} \neq \tau_{10}$.

It shall be noted that, for the class of transmission policies under study, the number of terminals that access the channel at a given slot depends in general on the state of the corresponding stochastic processes.[1] Due to this, an exact characterization of the probability to successfully deliver a message would require tracking the number of sources in a given state in the current and previous slot, with an increase in complexity as the network population grows. In this respect, we resort instead to an approximation, whose tightness will be verified and discussed in Sec. V. Namely, we assume that all terminals other than the reference one access the channel independently with probability α. This, in turn, is given by the average probability for a node to transmit over a slot, i.e.

$$\alpha = \pi_0[\,q_{00}\tau_{00} + q_{01}\tau_{01}\,] + \pi_1[\,q_{10}\tau_{10} + q_{11}\tau_{11}\,]. \tag{3}$$

We refer to this assumption as the *myopic approximation*. Following this model, the probability to deliver a message sent by the reference source can be expressed as

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

III. Decode and Hold Estimator

To gauge the performance of the D&H estimator, it is convetient to jointly track the evolution of the reference source and of its estimate at the receiver. The corresponding 4-state Markov chain $(X_n, \widehat{X}_n)$ is reported in Fig. 2, for the general transmission probability vector $\boldsymbol{\tau}$. The transition probabilities of the process can be readily derived for the system model under study. For instance, the chain transitions from $(0,0)$ to $(1,1)$ if the source evolves from 0 to 1 (probability q_{01}), and the terminal transmits an update which is correctly decoded at the receiver (overall probability $\tau_{01}\omega$). For the same source behavior, instead, the chain moves to $(0,1)$ if no message from the reference terminal is received (factor $1 - \tau_{01}\omega$). Finally, the process remains in state $(0,0)$ if the source does not transition (probability $1 - \tau_{01}$), regardless of the outcome of a potential transmission, since the receiver already has the correct estimate. Following a similar reasoning, all other probabilities can be derived, and are shown for completeness in Fig. 2.

[1] The only exceptions are the random transmission case, and the reactive transmission approach, under the assumption of symmetric sources ($\beta = 1$).

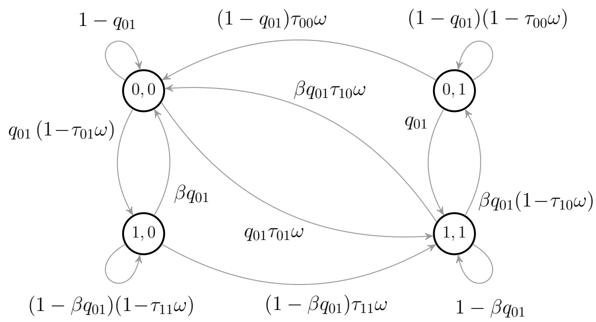

Fig. 2. Markov chain $(X_n, \widehat{X}_n)$ jointly tracking the evolution of the reference source and of the receiver estimate, under a D&H estimator.

The stationary distribution $\pi_{(i,j)}$, $i,j \in \{0,1\}$ for the irreducible, aperiodic chain can be computed directly from the balance equations. Of particular interest are the values obtained for $\pi_{(0,1)}$ and $\pi_{(1,1)}$, reported in (4) at the top of next page. Leaning on these, indeed, $\mathsf{P_{det}}$ and $\mathsf{P_{fa}}$ can be immediately computed from the definition of conditional probability as

$$\mathsf{P_{fa}} = \frac{\pi_{(0,1)}}{\pi_0}, \quad \mathsf{P_{det}} = \frac{\pi_{(1,1)}}{\pi_1}. \tag{5}$$

The presented formulations allow to characterize the behavior of the system in terms of false alarm and detection probability for any choice of the channel access vector $\boldsymbol{\tau}$. From this standpoint, we note that a specific configuration of the protocol parameters leads to a unique pair $(\mathsf{P_{fa}}, \mathsf{P_{det}})$ when the receiver relies on a D&H estimator.

Furthermore, while a thorough discussion of these trends will be provided in Sec. V, some useful initial insights are provided by considering the simple *random* and *reactive* transmission strategies. In the former case, the expressions in (5) simplify to

$$\mathsf{P_{fa}} = \frac{q_{01}(1-\tau\omega)}{\tau\omega + (1+\beta)q_{01}(1-\tau\omega)}, \quad \mathsf{P_{det}} = 1 - \beta\,\mathsf{P_{fa}} \tag{6}$$

prompting two remarks. On the one hand, recalling that for the random approach the probability for a node to transmit over a slot is independent of the current or past source state, we have $\alpha = \tau$, and $\omega = (1-\tau)^{M-1}$. Plugging this into (6), one can readily prove that the best performance (i.e., minimum false alarm probability and maximum detection rate) is attained by setting $\tau = 1/M$. In other words, the protocol configuration maximizing throughput and age of information [18] is also optimal for a D&H estimator if terminals behave ignoring the evolution of the monitored processes. On the other hand, a linear relationship, with proportionality factor β emerges

$$\pi_{(0,1)} = \frac{\beta q_{01}(\tau_{11} + \beta q_{01}\tau_{01} - \beta q_{01}\tau_{11})(1 - \tau_{10}\omega)}{\mathsf{z}}, \qquad \pi_{(1,1)} = \frac{\big(\tau_{11} + \beta q_{01}(\tau_{01} - \tau_{11})\big)\big(q_{01} + (1 - q_{01})\tau_{00}\omega\big)}{\mathsf{z}} \tag{4}$$

$$\mathsf{z} = (1 + \beta)\big\{\tau_{11}\left[q_{01} + \tau_{00}\omega(1 - q_{01})\right] + \beta q_{01}\left[q_{01}(\tau_{01} + \tau_{10} - \tau_{11} - \tau_{01}\tau_{10}\omega) + \tau_{00}(1 - q_{01})(1 - \tau_{11}\omega)\right]\big\}$$

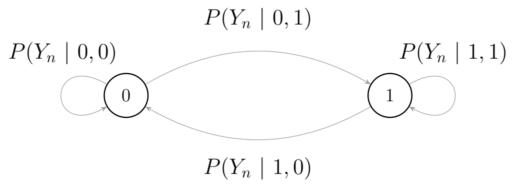

Fig. 3. Hidden Markov model considered for tracking the evolution of the APP logarithmic ratio. The hidden state is given by the reference source X_n, whereas the observation Y_n is driven by the source transition between times $n-1$ and n.

between false alarm and detection probability, regardless of the value chosen for τ.

When considering a *reactive* transmission strategy, instead, the performance indicators evaluate to

$$\mathsf{P}_{\mathsf{fa}} = \frac{1-\omega}{2-\omega}, \quad \mathsf{P}_{\mathsf{det}} = 1 - \mathsf{P}_{\mathsf{fa}}.$$

In this case, however, the channel access probability is purely dictated by the evolution of the sources ($\alpha = q_{01}(1+2\beta)/(1+\beta)$), so that worse performance will be attained when the tracked processes transition more frequently.

IV. MAP Estimator

In order to capture the performance of the threshold test introduced in (2), we characterize the joint distribution of the reference source and of the r.v. Λ_n, describing the time evolution of the APP logarithmic ratio. In this perspective, it is convenient to consider a hidden Markov model (HMM), where the underlying process X_n can be observed via the channel outputs Y_n. For the class of channel access policies under study, and leaning on the myopic approximation introduced in Sec. II, the probability of observing a specific slot outcome at the receiver depends solely on the current and past state of the reference source, leading to the HMM illustrated in Fig. 3. The conditional distribution $P(Y_n \,|\, X_{n-1}, X_n)$, $Y_n \in \mathcal{Y}$, $X_n \in \mathcal{X}$ follows thus from simple calculations. Specifically, the probability of observing an update from the source of interest is given by

$$P(y_n \,|\, x_{n-1}, x_n) = \tau_{x_{n-1}x_n}\,(1-\alpha)^{M-1} \quad \text{if } y_n = x_n$$

whereas the HMM generates an observation $\mathtt{I}$ with probability

$$P(\mathtt{I} \,|\, x_{n-1}, x_n) = (1 - \tau_{x_{n-1}x_n})\left[(1-\alpha)^{M-1} + \alpha(M-1)(1-\alpha)^{M-2}\right].$$

In the expression, the first factor accounts for the reference terminal not to access the channel, whereas the summation captures either the absence of other transmissions (idle slot) or the transmission of a single message from any of the other $M-1$ nodes. Finally, the probability of observing a collision can be expressed as

$$P(\mathtt{C} \,|\, x_{n-1}, x_n) = \tau_{x_{n-1}x_n}\left[1 - (1-\alpha)^{M-1}\right] + (1 - \tau_{x_{n-1}x_n})\left[(1-\alpha)^{M-1} + \alpha(M-1)(1-\alpha)^{M-2}\right].$$

Also in this case, the first addend captures the case in which the reference terminal transmits and at least one of the other nodes accesses the channel as well, whereas the second addend refers to a collision in which the node of interest remains silent, while two ore more other nodes transmit.

Leaning on the presented HMM, the evolution of the *approximated* APP logarithmic ratio can be computed iteratively, through the recursion

$$\lambda_n = \ln \frac{\sum_{x_{n-1}\in\mathcal{X}} P(y_n \,|\, x_{n-1}, 0)\, q_{x_{n-1}0} \cdot e^{-\lambda_{n-1}x_{n-1}}}{\sum_{x_{n-1}\in\mathcal{X}} P(y_n \,|\, x_{n-1}, 1)\, q_{x_{n-1}1} \cdot e^{-\lambda_{n-1}x_{n-1}}} := f(\lambda_{n-1}, y_n) \tag{7}$$

To derive the result, we start by considering the well-known forward equation of HMMs [27], which allows to express the joint distribution of the current source state at time n and the overall available observations up to that point as

$$P(x_n, y^n) = \sum_{x_{n-1}\in\mathcal{X}} P(y_n | x_{n-1}, x_n) P(x_n | x_{n-1}) P(x_{n-1}, y^{n-1}). \tag{8}$$

Furthermore, applying the exponential function to both sides of (1), and recalling that $\mathsf{P}[X_n = 0 \,|\, Y^n] = 1 - \mathsf{P}[X_n = 1 \,|\, Y^n]$ we readily obtain the relations $\mathsf{P}[X_n = 0 \,|\, Y^n] = 1/(1 + e^{-\lambda_n})$ as well as $\mathsf{P}[X_n = 1 \,|\, Y^n] = e^{-\lambda_n}/(1 + e^{-\lambda_n})$. Combining the two leads to the compact formulation

$$P(x_n \,|\, y^n) = \frac{e^{-\lambda_n x_n}}{1 + e^{-\lambda_n}}. \tag{9}$$

If we now observe that the definition of λ_n can be equivalently expressed in terms of the logarithmic ratio of the joint distribution of X_n and Y^n, the recursion (7) directly follows by writing both numerator and denominator of the ratio leaning on (8), and by plugging (9) into the obtained expressions.

The recursion in (7) can in turn be used to estimate the joint distribution of X_n and Λ_n in a computationally efficient way. To this aim, quantizing the APP logarithmic ratio so that it assumes values in a finite set, by the law of total proability we can write

$$P(\lambda_n, x_n) = \sum_{x_{n-1}\in\mathcal{X}} \sum_{\substack{y_n\in\mathcal{Y},\lambda_{n-1}: \\ f(y_n,\lambda_{n-1})=\lambda_n}} P(\lambda_{n-1}, y_n, x_n, x_{n-1}) \tag{10}$$

where the inner summation considers all the possible (y_n, λ_{n-1}) pairs that can lead to λ_n via (7). On the other hand,

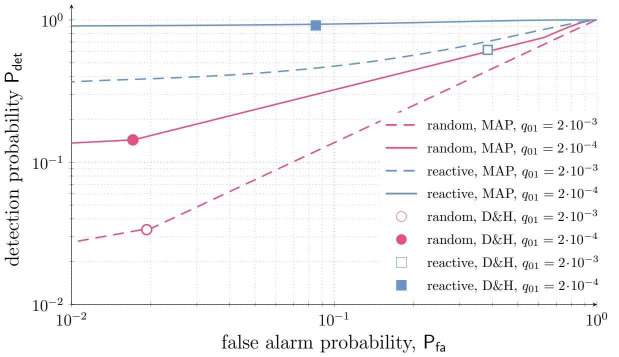

Fig. 4. Detection probability vs false alarm probability for the random and reactive transmission policies. Results for both D&H (markers) and MAP (lines) detectors are reported, for two different values of q_{01}.

the joint distribution within the summation can be rearranged using simple conditioning arguments as

$$P(\lambda_{n-1}, y_n, x_n, x_{n-1}) = \\ P(y_n \,|\, x_n, x_{n-1}) P(x_n \,|\, x_{n-1}) P(\lambda_{n-1}, x_{n-1}). \quad (11)$$

We thus obtain once more a simple recursion, which allows to compute $P(\lambda_n, x_n)$, with initial conditions set to $\mathsf{P}[\Lambda_{-1} = 0, X_{-1} = 0] = \mathsf{P}[\Lambda_{-1} = 0, X_{-1} = 1] = 0.5$. We remark that the presented approach implements a quantized density evolution, akin to the one reported in [28], and inspired by [29, Chapter 4].

Finally, the joint distribution of the APP logarithmic ratio and of the current source state provides the false alarm and detection probabilities attainable by the MAP estimator. Under the myopic approximation of Sec. II for a threshold θ, we can write

$$\mathsf{P}_{\mathsf{fa}} = \lim_{n\to\infty} \frac{1}{\pi_0} \sum_{\lambda_n < \theta} P(\lambda_n, 0)$$

$$\mathsf{P}_{\mathsf{det}} = \lim_{n\to\infty} \frac{1}{\pi_1} \sum_{\lambda_n < \theta} P(\lambda_n, 1)$$

and obtain an evaluation of the quantities by running the recursion in (10)-(11) over a sufficiently long time horizon.

V. Results and Discussion

To gather insights on the performance of the different channel access strategies under the two considered estimators, in the following we focus on a setting with $M = 250$ nodes. We considered asymmetric sources with $\beta = 50$, so that the fraction of time spent by a monitored process in the critical state is $\pi_1 = 0.02$.

To begin our study, we tackle the two simplest approaches, i.e., the random and reactive transmission strategies. In the former case, the transmission probability is set to $1/M$, following the discussion of Sec. IV. The corresponding behavior in terms of detection probability against false alarm rate is shown in Fig. 4. As discussed, performance is captured by a single point (square and circle markers) when the receiver relies on the D&H estimator, whereas a full curve can be obtained when using a MAP estimate by varying the decision threshold θ. The plot reports results for two different values of q_{01}, namely $2\cdot10^{-3}$ (dashed lines for MAP, empty markers for D&H) and $2\cdot10^{-4}$ (solid lines, filled markers). Consider first the D&H case. In this setting, the detection probability can be improved implementing reactive transmissions, although at a cost of a higher false alarm probability. Indeed, when updates are sent at random intervals, regardless of the source evolution, the average time between successive transmissions for a node is M slots, which is larger than the permanence time of the monitored process in state 1 (i.e., $1/(\beta q_{01})$). Therefore, a transition to the critical state may not be notified at all, leading to a low detection probability. In addition, the random strategy may lead nodes to access the channel attempting to deliver redundant information (e.g., multiple packets containing a source state the receiver is already aware of), increasing the chance of collision for updates signaling critical transitions.[2] The reactive solution circumvents this drawback by notifying only state changes. On the other hand, however, if a packet signaling that the source has gone back to state 0 is lost, a long time may elapse before the reactive protocol triggers another update (on average, $1/q_{01}$ slots), leaving the receiver with an erroneous estimate that increases the false alarm rate. From this standpoint, implementing random transmissions may be beneficial in some case, as (possibly repeated) packets notifying the source being in state 0 directly lead to a reduction of P_{fa}.

The benefit of using a reactive solution becomes apparent when the receiver implements a MAP estimator, as highlighted again by Fig. 4. In this case, binding the behavior of the channel access protocol to the evolution of the tracked source allows to better estimate the current state from channel observations. Indeed, if updates are triggered irrespective of the monitored processes, a slot outcome in $\{\mathtt{C}, \mathtt{I}\}$ does not bring any additional information to the receiver. This is reversed if nodes follow a reactive policy, as the receiver can for instance infer that the source of interest has not transitioned at the current slot upon observing channel output $\mathtt{I}$. We may observe that MAP curves in Fig. 4 exhibit a first-order discontinuity. We conjecture this effect to be the result of the choice of quantization intervals used by the density evolution analysis.

Some further remarks can be extracted from Fig. 4. As expected, all policies perform better under any estimator for lower values of q_{01}, as stochastic processes with fewer and more sporadic transitions can be tracked more easily. Moreover, the MAP criterion always improves over a D&H tracking, exploiting additional knowledge on the source and channel access statistics. The gap is larger for higher values of q_{01}, as the stronger channel contention induced by the reactive scheme can beset the D&H performance. On the other hand, the improvement in terms of detection probability is rather limited for low q_{01}, prompting the simple estimator as an interesting solution for many practical settings. Nonetheless,

[2] Note that, with the random strategy, the channel is operated at load of $\mathsf{G} = 1$ pkt/slot, with a probability of successfully delivering a packet approximated by $e^{-1} = 0.36$. In turn, for the considered values of q_{01}, the channel load of the reactive policy can easily be estimated based on (3) as $\mathsf{G} \simeq \alpha M = 2\beta q_{01} M/(1+\beta)$, leading to $\mathsf{G} \simeq 0.1$ (success probability ~ 0.9) and $\mathsf{G} \simeq 1$ for $q_{01} = 2\cdot10^{-4}$ and $q_{01} = 2\cdot10^{-3}$, respectively.

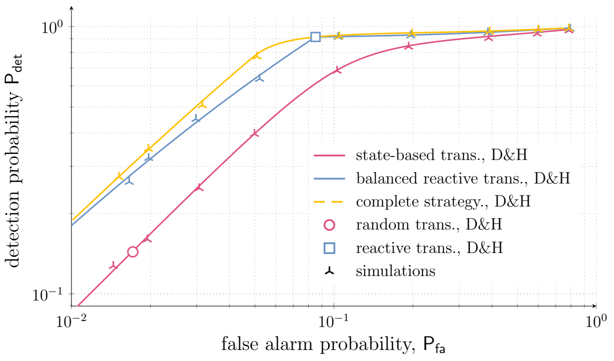

Fig. 5. Maximum detection probability achievable for a target false alarm probability for different channel access strategies when the receiver employs a D&H estimator. Solid lines denote analytical results under the myopic approximation, whereas markers (⅄) were obtained by means of Montecarlo simulations. In all cases, $q_{01} = 2 \cdot 10^{-4}$.

as noted, relying on a MAP estimate offers the ability to pick the desired operating point over a complete receiver operating characteristic (ROC) curve, so that a small loss in terms of detection probability can lead to a significant reduction of the false alarm rate.

This observation triggers the natural question of how the performance of the D&H estimator can be changed by tuning the vector of channel access probabilities $\boldsymbol{\tau}$, going beyond the basic random and reactive transmission schemes. Specifically, for a target false alarm probability $\mathsf{P}_{\mathsf{fa}}^*$, an interesting problem is to determine the protocol configuration that maximizes the detection probability in the D&H case. We thus tackle the constrained optimization problem

$$\begin{aligned} &\max_{\boldsymbol{\tau}} \ \mathsf{P}_{\mathsf{det}}(\boldsymbol{\tau}) \\ &\text{s.t. } \mathsf{P}_{\mathsf{fa}}(\boldsymbol{\tau}) = \mathsf{P}_{\mathsf{fa}}^* \end{aligned} \tag{12}$$

where $\mathsf{P}_{\mathsf{det}}$ and P_{fa} are obtained via (4) and (5), and which can easily be solved via standard methods. The outcome of this study is reported in Fig. 5 for the case $q_{01} = 2 \cdot 10^{-4}$. In the plot, solid lines refer to the solution of the problem in (12), whose expressions rely on the myopic approximation. Instead, the starred markers (⅄) show results obtained via Montecarlo simulations that account for the proper behavior of all terminals, i.e., the number of transmissions in a slot is jointly driven by the state of all the sources in the system. Different colors identify the state-based, balanced reactive, and the complete strategy. Finally, the circle and square markers show for reference the results obtained with the random and reactive solutions (single point, same values shown already in Fig. 4).

The plot highlights how different operating points in terms of $(\mathsf{P}_{\mathsf{fa}}, \mathsf{P}_{\mathsf{det}})$ can indeed be targeted with the use of a D&H estimator by properly operating the underlying access protocol. For completeness, the values of the transmission probability vector $\boldsymbol{\tau}$ leading to the maximum $\mathsf{P}_{\mathsf{det}}$ for a specific P_{fa} are shown in Fig. 6 at the top of next page, for all considered strategies. As a first remark, state-based policies can improve over the pure random scheme, leading to higher detection probabilities. This is obtained by sending updates more likely when the monitored process lies in the critical state, and by reducing the transmission frequency otherwise, so to avoid channel congestion. The intuitive approach comes though at the risk of leaving the receiver with a wrong estimate after the process has reverted to 0. In this perspective, a simple take-away emerges: transmission policies that do not tackle directly state changes fail to deliver good performance. For example, a reasonable detection rate (e.g., larger than 0.9) can only be attained with the state-based strategy when tolerating a very large false alarm rate (> 0.3 in the considered setting). Further insights are offered by considering the balanced-reactive strategy, as little improvement can be obtained over the simple reactive solution. Indeed, higher values of $\mathsf{P}_{\mathsf{det}}$ are achieved by significantly increasing P_{fa}, and, in turn, a reduced false alarm rate leads to a rather steep decrese in terms of detection capabilities. Notably, the optimal protocol configuration in this case (Fig. 6b) foresees to notify with probability 1 all source transitions in one direction, and setting to a strictly positive, yet lower, value the probability of update transmission for the complementary event.

On the other hand, noticeable gains can be reaped with the complete policy (yellow line). For instance, an improvement of a factor ~ 20% and ~ 85% in terms of $\mathsf{P}_{\mathsf{det}}$ with respect to the balanced reactive and state-based policies is obtained for a false alarm probability of 0.06 (not achievable in this setting with the plain reactive policy). Furthermore, the complete policy presents a performance curve with a smoother behavior around values of $\mathsf{P}_{\mathsf{det}} \sim 0.9$, allowing some flexibility in matching a target false alarm rate without compromising too much the detection capabilities. Fig. 6c illustrates how the benefits are achieved, pinpointing that the optimal channel access strategy when a D&H estimator is implemented consists in notifying with high probability state changes, and complementing this by sporadically transmitting even in case of no change, so as to reduce the probability for the receiver to remain stuck with a wrong estimate. It is important to observe that such protocol operation mode is profoundly different from the one maximizing throughput or minimizing age of information, which foresees transmitting independently of the source state or evolution. In Fig. 5 we also plot the outcome of a full blown simulation which does not rely on the myopic approximation introduced in Section II for tracking the state of the reference source. The plot shows that simulations are in excellent agreement with the analytical results, supporting the validity of the myopic approximation for the considered setup.

To conclude our study, we explore whether protocol settings that are optimal when the receiver implements a D&H estimator offer good performance also when a MAP criterion is employed. To this aim, we select two reference values of P_{fa}, namely 0.05 and 0.1, and denote by $\boldsymbol{\tau}_1$ and $\boldsymbol{\tau}_2$ the transmission probability vectors that solve problem (12) for the complete strategy, i.e., they lead to the points of abscissa 0.05 and 0.1 for the yellow curve in Fig. 5. We then test the behavior of the MAP estimator under such protocol configurations, computing the $(\mathsf{P}_{\mathsf{det}}, \mathsf{P}_{\mathsf{fa}})$ pairs that are obtained by varying the decision threshold θ. The results are reported by the solid curves in Fig. 7. Always for the MAP estimator, the dashed

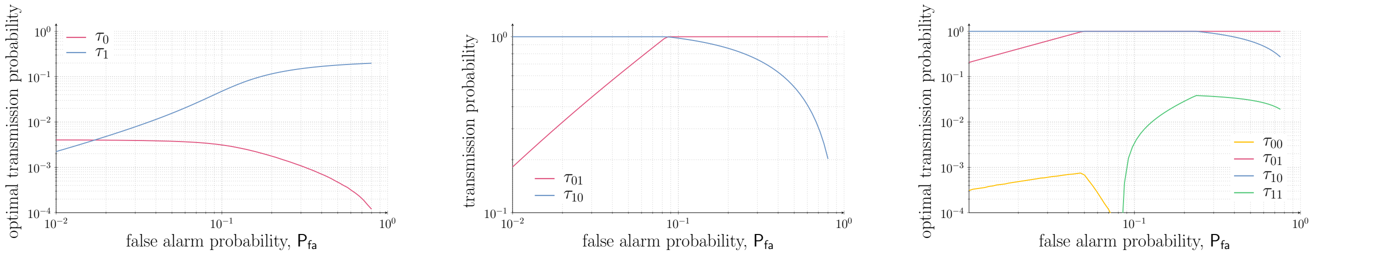

(a) state-based (b) balanced-reactive (c) complete

Fig. 6. Transmission probabilities for a state-based (a), balanced reactive (b) and complete (c) transmission strategy in order to maximize the detection probability for a target false alarm rate (x-axis). The results are obtained by solving the optimization problem in (12) under the myopic approximation.

curve shows the performance attained when nodes implement the simple reactive policy, i.e., $\tau_{10} = \tau_{01} = 1$, $\tau_{11} = \tau_{00} = 0$. Finally, square and circle markers indicate the best operating points already discussed for the D&H estimator when targeting the two considered values of P_{fa}, whereas starred markers ($\curlywedge$) show results obtained via Montecarlo simulations that account for the proper behavior of all terminals. Some key messages emerge from the plot. First, a significant degradation is experienced when the protocol operates with probability vector $\boldsymbol{\tau}_1$ if compared to the plain reactive approach. From this standpoint, a configuration which allowed to improve detection probability under the D&H estimator is detrimental when the receiver implements the MAP-based one. As highlighted from Fig. 6c, $\boldsymbol{\tau}_1$ differentiates from a reactive solution by setting $\tau_{00} \simeq 7 \cdot 10^{-3}$. This choice is beneficial in the D&H case, as notifying the permanence in state 0 can reduce the false alarm rate if a message signaling the transition $1 \to 0$ was lost. On the other hand, when tracking the APP logarithmic ratio, the receiver can update its estimate without the need for such messages, leaning on the source evolution statistics which predicts a rather quick return to the normal operating state. In this perspective, the higher number of transmissions become detrimental, increasing the overall collision probability. The situation changes when the protocol is configured with the probability vector $\boldsymbol{\tau}_2 = [0, 1, 1, 3.5 \cdot 10^{-3}]$. In this case, a small improvement over the reactive case is experienced also for the MAP estimator, benefiting from some additional transmissions when in state 1. We further observe that an excellent agreement is obtained between simulations and analytical results obtained leaning on the myopic approximation also for the MAP case.

More importantly, the result paves the way for further studies, aimed at deriving the optimal strategy to employ under a MAP-based estimate criterion, and which we regard as part of our future work. The reported results stress in any case the positive impact that connecting the channel access policy to the dynamic and transitions of monitored processes can have on remote source monitoring, in contrast to what commonly observed when targeting information freshness metrics.

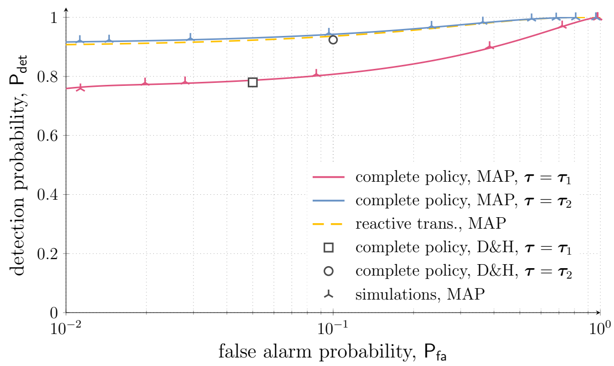

Fig. 7. (P_{det}, P_{fa}) pairs obtained using a MAP estimator for three different protocol configurations. $\boldsymbol{\tau}_1$ and $\boldsymbol{\tau}_2$ correspond to the probability vectors that maximize the detection probability under the D&H estimator for a target P_{fa} of 0.05 and 0.1, respectively. In addition, performance under a reactive strategy using a MAP estimator is also reported (dashed line). In all cases, $q_{01} = 2 \cdot 10^{-4}$. rategies when the receiver employs a D&H estimator. For MAP, lines denote analytical results under the myopic approximation, whereas markers ($\curlywedge$) were obtained by means of Montecarlo simulations.

VI. Conclusions

In this paper, we have studied an IoT system in which a large number of devices share a wireless channel to a common receiver in an uncoordinated manner. Each device monitors an independent stochastic process, modeled as a two-state Markov chain where one of the states is visited sporadically (capturing a critical or alarm event). Time is slotted, and at every slot a device decides whether to transmit a packet containing the value of the corresponding process, with a probability that depends on the present and past state of the source. The receiver estimates the state of each process across time, leaning on the slot outputs it observes. Specifically, we have considered two approaches: a simple decode and hold estimator, which does not require any knowledge of the source statistics, and a more efficient maximum a posteriori solution. For both, we have presented an analytical characterization of the system performance in terms of false alarm and detection probability, resorting to tools based on (hidden) Markov models. The behavior of different channel access strategies has been discussed, ranging from approaches in which transmission probabilities are oblivious of the source evolution, to solutions that tie the access to source transitions. Furthermore, an optimization of the protocol parameters when a decode and hold estimator is employed has been presented, also discussing the effectiveness of this configuration under a MAP estimator. The study reveals several non-trivial take-aways, clarifying how a tuning of the random access policy to the evolution of tracked processes can have a profound impact in improving detection probability and reducing false alarm rates. Notably, these remarks depart from well-known design criteria that are used to optimize ALOA-based systems relying on metrics such as throughput or age of information.

REFERENCES

[1] LoRa Alliance, "The LoRa Alliance Wide Area Networks for Internet of Things," www.lora-alliance.org.

[2] Sigfox, "SIGFOX: The Global Communications Service Provider for the Internet of Things," www.sigfox.com.

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

[4] S. Kaul, R. Yates, and M. Gruteser, "On piggybacking in vehicular networks," in *Proc. IEEE GLOBECOM*, Dec 2011.

[5] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing age of information in vehicular networks," in *Proc. IEEE SECON*, June 2011.

[6] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1183–1210, May 2021.

[7] 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, 2020.

[8] M. Salimnejad, M. Kountouris, and N. Pappas, "State-aware real-time tracking and remote reconstruction of a Markov source," 2023. [Online]. Available: http://arxiv.org/abs/2309.11950

[9] T. Soleymani, J. Baras, and S. Hirche, "Value of information in feedback control," *IEEE Trans. Autom. Control*, 2020.

[10] Y. 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, Feb. 2020.

[11] Y. Sun and B. Cyr, "Information aging through queues: A mutual information perspective," in *Proc. IEEE SPAWC*, 2018.

[12] 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.

[13] 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 Journal on Selected Areas in Information Theory*, 2023.

[14] R. Yates and S. K. Kaul, "Status updates over unreliable multiaccess channels," in *Proc. IEEE ISIT*, 2017.

[15] R. Yates and S. Kaul, "Age of information in uncoordinated unslotted updating," in *Proc. IEEE ISIT*, 2020.

[16] O. T. Yavaskan and E. Uysal, "Analysis of slotted ALOHA with an age threshold," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, May 2021.

[17] X. Chen, K. Gatsis, H. Hassani, and S. S. Bidokhti, "Age of information in random access channels," *IEEE Trans. Inf. Theory*, vol. 68, no. 10, 2022.

[18] A. Munari, "Modern random access: an age of information perspective on irregular repetition slotted ALOHA," *IEEE Trans. Commun.*, vol. 69, no. 6, Jun. 2021.

[19] A. Maatouk, M. Assaad, and A. Ephremides, "On the age of information in a CSMA environment," *IEEE/ACM Trans. Netw.*, vol. 28, no. 2, pp. 818–831, 2020.

[20] L. Badia, A. Zanella, and M. Zorzi, "A game of ages for slotted ALOHA with capture," *IEEE Transactions on Mobile Computing*, 2023.

[21] A. Munari, F. Lázaro, G. Durisi, and G. Liva, "The dynamic behavior of frameless ALOHA: Drift analysis, throughput, and age of information," *IEEE Trans. Commun.*, 2023.

[22] A. Maatouk, M. Assaad, and A. Ephremides, "Minimizing the age of information: NOMA or OMA?" in *Proc. IEEE INFOCOM Workshops*, April 2019.

[23] A. Munari and L. Badia, "The role of feedback on AoI under limited transmission opportunities," in *Proc. IEEE Globecom*, 2022.

[24] A. Nayak, A. Kalør, F. Chiariotti, and P. Popovski, "A decentralized policy for minimization of age of incorrect information in slotted ALOHA systems," in *Proc. IEEE ICC.*, 2023.

[25] A. Munari, "Age of incorrect information in random access channels without feedback," 2023. [Online]. Available: http://arxiv.org/abs/2308.06034

[26] H. V. Trees, *Detection, Estimation, and Modulation Theory*. New York: Wiley, 2001.

[27] 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.

[28] G. Cocco, A. Munari, and G. Liva, "State estimation entropy for two-state markov sources in slotted ALOHA random access channels," in *Proc. IEEE ITW*, 2023.

[29] H. D. Pfister, "On the capacity of finite state channels and the analysis of convolutional accumulate-m codes," Ph.D. dissertation, University of California, San Diego, 2003.