Detecting Changes in Hidden Markov Models

George V. Moustakides

Rutgers University and University of Patras

New Brunswick and Rion

NJ 08991, USA and 26500 Patras, Greece

Email: gm463@rutgers.edu and moustaki@upatras.gr

Abstract—We consider the problem of sequential detection of a change in the statistical behavior of a hidden Markov model. By adopting a worst-case analysis with respect to the time of change and by taking into account the data that can be accessed by the change-imposing mechanism we offer alternative formulations of the problem. For each formulation we derive the optimum Shewhart test that maximizes the worst-case detection probability while guaranteeing infrequent false alarms.

I. INTRODUCTION

We consider a hidden Markov model (HMM) where $\{\xi_t\}_{t \geq 1}$ is the observation process that is acquired sequentially and $\{z_t\}_{t \geq 0}$ is a Markov process that controls the statistical behavior of $\{\xi_t\}$ but its state is hidden. Let also $\tau \in \{0, 1, \ldots\}$ denote a *changetime* with the processes $\{(\xi_t, z_t)\}$ following a nominal probability measure P_∞ up to and including time τ while, *after* τ, the probability measure switches to an alternative regime P_0. This change induces a new measure which is denoted by P_τ with $E_\tau[\cdot]$ being reserved for the corresponding expectation.

To be more precise, for $t \leq \tau$, we make the simplifying assumptions that the observations $\{\xi_t\}$ are i.i.d. with a common pdf $f_\infty(\xi)$ and the Markov process having a transition pdf $g_\infty(z_t|z_{t-1})$. After the change the observations are conditionally independent and controlled by the Markov process. In particular ξ_t conditioned on z_t has a pdf $f_0(\xi_t|z_t)$ while the Markov process has transition pdf $g_0(z_t|z_{t-1})$. It is possible to have $g_0(z_t|z_{t-1}) = g_\infty(z_t|z_{t-1})$, with the Markov process not undergoing any change. As mentioned, under the nominal measure the observations are assumed not to be controlled by the Markov process.

We would like to detect the onset of the change in the statistical behavior using a sequential strategy. We are therefore interested in defining a stopping time T adapted to the filtration $\{\mathscr{F}_t^\xi\}_{t \geq 0}$ generated by the observations, that is, $\mathscr{F}_t^\xi = \sigma\{\xi_1, \ldots, \xi_t\}$, to perform the detection. In order to select optimally T we need to propose a suitable performance measure and properly optimize it. To derive our criterion we are going to extend the idea introduced in [6]. Even though we can access only the observation process $\{\xi_t\}$ to perform detection, there is a *change-imposing mechanism* that decides the time τ to impose the change. This mechanism may have access to *a completely different set of data* to make this decision.

II. PERFORMANCE MEASURE

As we suggested above T is an $\{\mathscr{F}_t^\xi\}$-adapted stopping time. In order to capture the fact that the change-imposing mechanism may have access to completely different information, we are going to assume that τ is also a stopping time (the time that the data stop following the nominal model) but adapted to a filtration $\{\mathscr{F}_t^w\}$ where $\mathscr{F}_t^w = \sigma\{w_1, \ldots, w_t\}$. In other words, the change-imposing mechanism sequentially consults the data sequence $\{w_t\}$ and at each time instant *makes a decision* as to whether it should impose a change or not. Both, ourselves that select T and the change-imposing mechanism that selects τ are bound by a *causality* constraint forbidding the use of any data from the future. Clearly process $\{w_t\}$, which is available to the change-imposing mechanism, may or may not include $\{\xi_t\}$ and it can be dependent or completely independent from the observations.

The change-imposing mechanism decides what is the best instant τ to impose the change while we decide what is the best time T to stop and declare that a change took place. If $\phi(t, s) \geq 0$ is a deterministic function expressing distance between or reward for the pair (t, s) then we can use the conditional expectation $E_\tau[\phi(T, \tau)|T > \tau]$ as a generic performance measure for the detection process. We condition on the event of no false alarms $\{T > \tau\}$ in order to compute the performance of T only during successes since we intend to take care of false alarms differently.

Most of the time the rule that defines the stopping time τ is unknown, therefore the proposed performance measure cannot be computed and, therefore, used to derive an optimum detection strategy. In such cases it is common to follow *a worst-case analysis* with respect to τ. In other words try to find the worst-case τ that will make the conditional expectation as unfavorable as possible to the detection goal. We have the following lemma that addresses this problem.

Lemma 1. *Suppose that T and τ are stopping times described as above, then*

$$\inf_\tau E_\tau[\phi(T, \tau)|T > \tau]$$
$$= \inf_{t \geq 0} \operatorname{ess\,inf} E_t[\phi(T, t)|T > t, \mathscr{F}_t^w], \quad (1)$$

The previous equality is also valid if we replace inf *and* ess inf *with* sup *and* ess sup *respectively.*

Proof: The proof is given in the Appendix. ∎

If we select $\phi(t,s) = (t-s)^+$ then we can define an extension of Lorden's measure [5] by computing the worst-case average detection delay as follows

$$\mathcal{J}(T) = \sup_{t \geq 0} \operatorname{ess\,sup} \mathsf{E}_t[T - t | T > t, \mathscr{F}_t^w].$$

We must emphasize that this is *not* the original Lorden measure since conditioning is with respect to the sigma-algebra $\mathscr{F}_t^w$ that controls τ and not $\mathscr{F}_t^\xi$ (which controls T) used in the original definition. Furthermore, in our approach the double maximization occurs naturally as a result of our worst-case analysis and not because of some arbitrary definition.

An alternative measure can be generated by evaluating the performance of T using the probability of detecting the change *immediately* after it occurs. In other words we are interested in the probability of the event $\{T = \tau + 1\}$. For this reason we define $\phi(t,s) = \mathbb{1}_{\{t=s+1\}}$. If we apply Lemma 1 we can compute the worst-case detection probability

$$\mathcal{P}(T) = \inf_{t \geq 0} \operatorname{ess\,inf} \mathsf{P}_t(T = t+1 | T > t, \mathscr{F}_t^w), \qquad (2)$$

which in this work is the criterion we intend to adopt.

Returning to HMMs and using (2) we distinguish *four* different cases depending on how $\{w_t\}$ is related to the data of the problem.

i) The change-imposing mechanism accesses information that is independent from $\{(\xi_t, z_t)\}$. In this case in (2) there is no conditioning with respect to $\mathscr{F}_t^w$ since the probability does not depend on $\{w_t\}$. This yields

$$\mathcal{P}_{\mathrm{i}}(T) = \inf_{t \geq 0} \mathsf{P}_t(T = t+1 | T > t). \qquad (3)$$

This is the Pollak-like criterion proposed in [10].

ii) The change-imposing mechanism accesses only the observations, then we obtain the Lorden-like criterion

$$\mathcal{P}_{\mathrm{ii}}(T) = \inf_{t \geq 0} \operatorname{ess\,inf} \mathsf{P}_t(T = t+1 | T > t, \mathscr{F}_t^\xi), \qquad (4)$$

proposed in [7] corresponding to $w_t = \xi_t$.

iii) The change-imposing mechanism accesses only the state of the Markov process. This leads to

$$\mathcal{P}_{\mathrm{iii}}(T) = \inf_{t \geq 0} \operatorname{ess\,inf} \mathsf{P}_t(T = t+1 | T > t, \mathscr{F}_t^z), \qquad (5)$$

where $\mathscr{F}_t^z = \sigma\{z_1, \ldots, z_t\}$ corresponding to $w_t = z_t$.

iv) The change-imposing mechanism accesses both, the observations and the state of the Markov process resulting in

$$\mathcal{P}_{\mathrm{iv}}(T) = \inf_{t \geq 0} \operatorname{ess\,inf} \mathsf{P}_t(T = t+1 | T > t, \mathscr{F}_t^{\xi,z}), \qquad (6)$$

where $\mathscr{F}_t^{\xi,z} = \sigma\{\xi_1, \ldots, \xi_t, z_1, \ldots, z_t\}$ corresponding to $w_t = (\xi_t, z_t)$.

For each criterion we can define a constrained optimization problem whose solution will provide the optimum T:

$$\sup_T \mathcal{P}_l(T), \text{ subject to: } \mathsf{E}_\infty[T] \geq \gamma > 1, \qquad (7)$$

where $l = \mathrm{i}, \mathrm{ii}, \mathrm{iii}, \mathrm{iv}$. In other words we maximize the worst-case detection probability assuring at the same time that the

average period between false alarms is lower bounded by a constant $\gamma > 1$ which we control.

The idea of maximizing the detection probability was first introduced in [1] under Shiryaev's [12] Bayesian formulation. In [10] we have a variant of the original Pollak measure [9] while a variant of Lorden's measure [5] was adopted in [7] for independent processes and in [8] for Markov. In this work we address the case of HMMs. The problem of change-detection in HMMs has been considered in the past in [2]–[4] and from these results it is well understood that even the asymptotic analysis is extremely challenging, not always leading to outcomes that are practically implementable.

III. CANDIDATE TESTS

Let us first present the joint data pdf induced by a change occurring at some time $\tau = t \geq 0$. For $0 < s \leq t$ we have

$$\mathsf{f}_t(\xi_s, \ldots, \xi_1, z_s, \ldots, z_0) =$$
$$\mathsf{f}_\infty(\xi_s) \cdots \mathsf{f}_\infty(\xi_1) \times \mathsf{g}_\infty(z_s | z_{s-1}) \cdots \mathsf{g}_\infty(z_1 | z_0) \mathsf{g}_\infty(z_0),$$

while for $s > t$ the resulting pdf takes the form

$$\mathsf{f}_t(\xi_s, \ldots, \xi_1, z_s, \ldots, z_0) =$$
$$\mathsf{f}_0(\xi_s | z_s) \cdots \mathsf{f}_0(\xi_{t+1} | z_{t+1}) \times \mathsf{g}_0(z_s | z_{s-1}) \cdots \mathsf{g}_0(z_{t+1} | z_t)$$
$$\times \mathsf{f}_\infty(\xi_t) \cdots \mathsf{f}_\infty(\xi_1) \times \mathsf{g}_\infty(z_t | z_{t-1}) \cdots \mathsf{g}_\infty(z_1 | z_0) \mathsf{g}_\infty(z_0),$$

where $\mathsf{g}_\infty(z_0)$ is the marginal pdf of z_0. The pdfs $\mathsf{f}_\infty(\xi_t)$, $\mathsf{f}_0(\xi_t | z_t)$, $\mathsf{g}_\infty(z_t | z_{t-1})$, $\mathsf{g}_\infty(z_0)$ are assumed known.

To simplify our presentation we are going to assume that $\mathsf{g}_\infty(z_0)$ is the stationary pdf for the transition pdf $\mathsf{g}_\infty(z_t | z_{t-1})$, namely $\int \mathsf{g}_\infty(z_t | z_{t-1}) \mathsf{g}_\infty(z_{t-1}) dz_{t-1} = \mathsf{g}_\infty(z_t)$. We can then define the following average probability density

$$\bar{\mathsf{f}}_0^1(\xi_t) = \iint \mathsf{f}_0(\xi_t | z_t) \mathsf{g}_0(z_t | z_{t-1}) \mathsf{g}_\infty(z_{t-1}) dz_{t-1} dz_t, \qquad (8)$$

which, when $\mathsf{g}_0(z_t | z_{t-1}) = \mathsf{g}_\infty(z_t | z_{t-1})$, simplifies to

$$\bar{\mathsf{f}}_0^1(\xi_t) = \int \mathsf{f}_0(\xi_t | z_t) \mathsf{g}_\infty(z_t) dz_t, \qquad (9)$$

and will be used for Criteria i) and ii). For Criteria iii) and iv) we define

$$\bar{\mathsf{f}}_0^2(\xi_t) = \iint \mathsf{f}_0(\xi_t | z_t) \mathsf{g}_0(z_t | z_{t-1}) \pi(z_{t-1}) dz_{t-1} dz_t, \qquad (10)$$

where $\pi(z)$ is a pdf to be specified in the sequel.

With the help of the average pdfs $\bar{\mathsf{f}}_0^j(\xi_t)$, $j = 1, 2$ we can now introduce the candidate Shewhart stopping time as follows

$$\mathsf{L}_j(\xi_t) = \frac{\bar{\mathsf{f}}_0^j(\xi_t)}{\mathsf{f}_\infty(\xi_t)}, \quad \mathcal{S}_j = \inf\{t > 0 : \mathsf{L}_j(\xi_t) \geq \nu_j\}. \qquad (11)$$

Threshold $\nu_j > 0$ is selected to satisfy the false alarm constraint with equality, namely

$$\mathsf{E}_\infty[\mathcal{S}_j] = \frac{1}{\mathsf{P}_\infty(\mathsf{L}_j(\xi_t) \geq \nu_j)} = \gamma, \qquad (12)$$

Existence of ν_j is guaranteed since the equation $\mathsf{P}_\infty(\mathsf{L}_j(\xi_1) \geq \nu_j) = \frac{1}{\gamma}$ has always a solution if we assume that $\mathsf{L}_j(\xi_t)$ does not contain any atoms under P_∞. Otherwise, in order to satisfy (12) we may need randomization every time $\mathsf{L}_j(\xi_t) = \nu_j$.

We can also compute the corresponding worst-case detection probability of the two Shewhart schemes. For the first test, since there is no dependence on the past, we have

$$\beta_1 = \mathcal{P}_i(\mathcal{S}_1) = \mathcal{P}_{ii}(\mathcal{S}_1) = \int \bar{f}_0^1(\xi_t)\mathbb{1}_{\{L_1(\xi_t) \geq \nu_1\}}d\xi_t. \quad (13)$$

For the second Shewhart test the analysis for finding the worst-case detection probability is slightly more involved. Consider first the conditional pdf

$$f_0(\xi_t|z_{t-1}) = \int f_0(\xi_t|z_t)g_0(z_t|z_{t-1})dz_t$$

then the worst-case detection probability satisfies

$$\beta_2 = \mathcal{P}_{iii}(\mathcal{S}_2) = \mathcal{P}_{iv}(\mathcal{S}_2) = \\ \inf_{z_{t-1}} \int f_0(\xi_t|z_{t-1})\mathbb{1}_{\{L_2(\xi_t) \geq \nu_2\}}d\xi_t. \quad (14)$$

We recall that the second Shewhart test is defined in terms of an arbitrary probability density $\pi(z)$. This means that the stopping time $\mathcal{S}_2$ and also the worst-case detection probability β_2 are functions of $\pi(z)$ as well. To specify $\pi(z)$, let $\mathcal{Z}$ denote its support, then $\pi(z)$ must be such that

$$\int f_0(\xi_t|z_{t-1})\mathbb{1}_{\{L_2(\xi_t) \geq \nu_2\}}d\xi_t = \beta_2, \text{ for } z_{t-1} \in \mathcal{Z} \\ \int f_0(\xi_t|z_{t-1})\mathbb{1}_{\{L_2(\xi_t) \geq \nu_2\}}d\xi_t \geq \beta_2, \text{ for } z_{t-1} \notin \mathcal{Z}. \quad (15)$$

In other words, $\pi(z)$ must put all its probability mass onto points for which the Shewhart test exhibits its worst-case performance. In fact (15) is sufficient to define $\pi(z)$ uniquely.

IV. Max-Min Optimality

In this section we will demonstrate that the stopping times $\mathcal{S}_j$, $j = 1, 2$, defined in (11) solve the max-min constrained optimization problem defined in (7). In order to prove our claim we first need to find a suitable upper bound for $\mathcal{P}_l(T)$. The following theorem provides the necessary expressions.

Theorem 1. *For any stopping time T with $\mathsf{E}_\infty[T] < \infty$ we have*

$$\mathcal{P}_l(T) \leq \frac{\mathsf{E}_\infty[L_1(\xi_T)]}{\mathsf{E}_\infty[T]}, \, l = i, ii; \quad \mathcal{P}_l(T) \leq \frac{\mathsf{E}_\infty[L_2(\xi_T)]}{\mathsf{E}_\infty[T]}, \, l = iii, iv.$$

Additionally, if $T = \mathcal{S}_j$, $j = 1, 2$, then we have equality in the corresponding inequality.

Proof: The proof is given in the Appendix. ∎

The next theorem optimizes the upper bounds proposed in Theorem 1.

Theorem 2. *If T is any stopping time satisfying the false alarm constraint, then*

$$\frac{\mathsf{E}_\infty[L_j(\xi_T)]}{\mathsf{E}_\infty[T]} \leq \beta_j, \, j = 1, 2,$$

where β_1, β_2 are defined in (13), (14) respectively.

Proof: The proof is highlighted in the Appendix. ∎

Combining Theorems 1 and 2, immediately assures optimality of the Shewhart tests. In particular for $l = i, ii$ we have

$$\mathcal{P}_l(T) \leq \frac{\mathsf{E}_\infty[L_1(\xi_T)]}{\mathsf{E}_\infty[T]} \leq \beta_1 = \mathcal{P}_i(\mathcal{S}_1) = \mathcal{P}_{ii}(\mathcal{S}_1),$$

while for $l = iii, iv$ we conclude

$$\mathcal{P}_l(T) \leq \frac{\mathsf{E}_\infty[L_2(\xi_T)]}{\mathsf{E}_\infty[T]} \leq \beta_2 = \mathcal{P}_{iii}(\mathcal{S}_2) = \mathcal{P}_{iv}(\mathcal{S}_2).$$

These two relationships establish optimality of the two Shewhart tests. In the next section we offer an example involving an interesting HMM.

V. Example

We consider the case of a Gaussian process whose mean is controlled by a Gaussian Markov process. Specifically, let the observations $\{\xi_t\}$ before the change be i.i.d. with pdf $f_\infty(\xi_t) \sim \mathcal{N}(0, 1)$ and after the change assume $f_0(\xi_t|z_t) \sim \mathcal{N}(z_t, 1)$. The process $\{z_t\}$ is unobservable and of the form $z_t = \mu + v_t$ where μ is a constant denoting the mean of z_t and $\{v_t\}$ is an AR(1) Gaussian process with v_t being conditionally Gaussian of the form $v_t \sim \mathcal{N}(\alpha v_{t-1}, \sigma^2)$ with $|\alpha| < 1$.

For the stationary pdf we have $g_\infty(z) \sim \mathcal{N}(\mu, \frac{\sigma^2}{1-\alpha^2})$. Since in this example we assume that the Markov process $\{z_t\}$ does not change, if we focus on the solution for Criteria i) and ii), we use (9) to compute

$$\bar{f}_0^1(\xi) = \int f_0(\xi|z)g_\infty(z)dz \sim \mathcal{N}\left(\mu, 1 + \frac{\sigma^2}{1-\alpha^2}\right). \quad (16)$$

Following (11) we can easily establish that the optimal Shewhart test is equivalent to

$$\mathcal{S}_1 = \inf\left\{t > 0 : \left|\xi_t + \mu\frac{1-\alpha^2}{\sigma^2}\right| \geq \nu_1\right\}. \quad (17)$$

Threshold ν_1 is related to the average false alarm period γ through (12) which specializes to

$$\Phi\left(\mu\frac{1-\alpha^2}{\sigma^2} - \nu_1\right) + \Phi\left(-\mu\frac{1-\alpha^2}{\sigma^2} - \nu_1\right) = \frac{1}{\gamma}, \quad (18)$$

while the worst-case detection probability becomes

$$\beta_1 = \Phi\left(\frac{\mu\left(1 + \frac{1-\alpha^2}{\sigma^2}\right) - \nu_1}{\sqrt{1 + \frac{\sigma^2}{1-\alpha^2}}}\right) + \Phi\left(-\frac{\mu\left(1 + \frac{1-\alpha^2}{\sigma^2}\right) + \nu_1}{\sqrt{1 + \frac{\sigma^2}{1-\alpha^2}}}\right). \quad (19)$$

Let us now consider Criteria iii) and iv). We focus on the computation of (10) and perform it in two steps. The first involves the computation of the conditional pdf

$$f_0(\xi_t|z_{t-1}) = \int f_0(\xi_t|z_t)g_0(z_t|z_{t-1})dz_t \\ \sim \mathcal{N}\left((1-\alpha)\mu + \alpha z_{t-1}, 1 + \sigma^2\right). \quad (20)$$

The next step consists in finding the pdf $\pi(z)$. We are going to assume that $\pi(z)$ puts all its mass on the single point $z = -\mu\frac{1-\alpha}{\alpha}$. This implies that $\bar{f}_0^2(\xi) \sim \mathcal{N}(0, 1 + \sigma^2)$. We can then verify that the resulting Shewhart test is equivalent to

$$\mathcal{S}_2 = \inf\{t > 0 : |\xi_t| \geq \nu_2\} \quad (21)$$

with the threshold satisfying the false alarm constraint

$$2\Phi(-\nu_2) = \frac{1}{\gamma}. \quad (22)$$

Of course, in order for our selection of $\pi(z)$ to be correct we need to show validity of (15). Therefore we must prove that $\mathsf{P}_0(|\xi_t| \geq \nu_2|z_{t-1})$ has a minimum for $z_{t-1} = -\mu\frac{1-\alpha}{\alpha}$. Using (20) the desired probability is

$$\mathsf{P}_0(|\xi_t| \geq \nu_2|z_{t-1}) = \Phi\left(\frac{-\nu_2 + ((1-\alpha)\mu + \alpha z_{t-1})}{\sqrt{1+\sigma^2}}\right) \\ + \Phi\left(\frac{-\nu_2 - ((1-\alpha)\mu + \alpha z_{t-1})}{\sqrt{1+\sigma^2}}\right)$$

which is clearly minimized when $(1-\alpha)\mu + \alpha z_{t-1} = 0$ with the minimum being equal to

$$\beta_2 = 2\Phi\left(-\frac{\nu_2}{\sqrt{1+\sigma^2}}\right). \tag{23}$$

The latter also constitutes the optimum worst-case detection probability for the Shewhart test in (21). It is worth mentioning that the Shewhart stopping time $\mathcal{S}_2$ is UMP with respect to α, μ and σ^2 since, as we can see, it does not require knowledge of these parameters. What is equally interesting is that the optimum worst-case detection probability β_2 is only a function of σ^2 and not of α, μ. It is only $\pi(z)$ that depends on these two parameters.

Suppose now that we erroneously assume that the change-imposing mechanism does not access the state of the Markov process when in reality it does. In this case we will be using $\mathcal{S}_1$ from (17) instead of $\mathcal{S}_2$ from (21). For $\mathcal{S}_1$ it is not difficult to verify that the worst-case detection probability is equal to

$$\tilde{\beta}_1 = 2\Phi\left(-\frac{\nu_1}{\sqrt{1+\sigma^2}}\right). \tag{24}$$

A similar erroneous assumption may occur when we consider the change-imposing mechanism to be able to access the Markov state when in reality it does not. Consequently by using $\mathcal{S}_2$ from (21) we need to compute its performance under the pdf in (16). This yields

$$\tilde{\beta}_2 = \Phi\left(\frac{(-\nu_2+\mu)\sqrt{1-\alpha^2}}{\sqrt{1-\alpha^2+\sigma^2}}\right) + \Phi\left(-\frac{(\nu_2+\mu)\sqrt{1-\alpha^2}}{\sqrt{1-\alpha^2+\sigma^2}}\right). \tag{25}$$

Clearly (25) must be compared against the optimum (19) while (24) against the optimum (23).

For a numerical comparison, let $\alpha = 0.5$, $\mu = 1$, $\sigma^2 = 0.5$ with γ ranging from 1 to 1000. Fig. 1 depicts the corresponding detection probabilities. The graph in blue corresponds to the change-imposing mechanism having no access to the Markov process and we correctly assume that it does not. This means that we plot β_1 from (19) against γ computed from
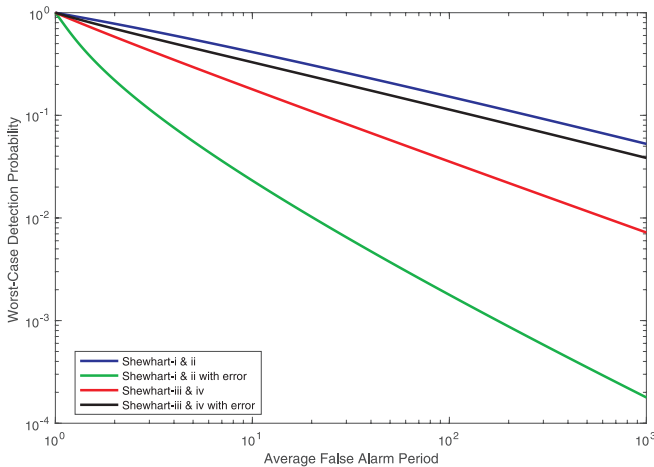

Fig. 1. Detection probability as a function of average false alarm period of Shewhart test when change-imposing mechanism does not access the Markov process state and we correctly assume it does not (blue); when it does not and we erroneously assume it does (black); when it does and we correctly assume it does (red) and finally when it does and we erroneously assume it does not (green).

(18). If this assumption is wrong and the change-imposing mechanism can actually access the Markov state then we have a severe performance degradation depicted by the graph in green where we plot $\tilde{\beta}_1$ from (24) against γ from (18).

If we now use the test in (21) and the change-imposing mechanism can indeed access the Markov state then the red graph depicts the worst-case detection probability β_2 from (23) as a function of γ from (22). In case we made a mistake in our judgement and the change-imposing mechanism cannot access the Markov state then the same test has a performance depicted by the black curve where we plot $\tilde{\beta}_2$ from (23) in terms of γ from (22).

By using the Shewhart test $\mathcal{S}_2$ in (21), which is obtained under more severe assumptions we do not lose much as compared to the optimum (17) if our assumption about the access capabilities of the change-imposing mechanism is incorrect. On the other hand, we guard ourselves against a hostile change-imposing mechanism when the latter can access all the available information. If, however, we assume that the change-imposing mechanism cannot access the Markov state and use $\mathcal{S}_1$, this assumption can be catastrophic if it is wrong.

VI. CONCLUSION

We considered the sequential change-detection problem for HMM which is known for being challenging. By introducing a generalized version of Lorden's performance measure we were able to develop the optimum solution that maximizes the worst-case detection probability. This result is interesting since it constitutes the first time a solution is obtained for a performance measure that is different from the classical measures adopted so far in the literature.

ACKNOWLEDGEMENT

APPENDIX

Proof of Lemma 1: Since τ is a $\{\mathscr{F}_t^w\}$-adapted stopping time we have that $\{\tau = t\}$ is $\mathscr{F}_t^w$-measurable consequently we can write

$$\mathsf{E}_\tau[\phi(T,\tau)|T>\tau] =$$

$$\frac{\sum_{t=0}^{\infty}\mathsf{E}_\infty[\mathsf{E}_t[\phi(T,t)\mathbb{1}_{\{T>t\}}|\mathscr{F}_t^w]\mathbb{1}_{\{\tau=t\}}]}{\sum_{t=0}^{\infty}\mathsf{E}_\infty[\mathsf{P}_t(T>t|\mathscr{F}_t^w)\mathbb{1}_{\{\tau=t\}}]} \geq$$

$$\inf_{t\geq 0}\frac{\mathsf{E}_\infty[\mathsf{E}_t[\phi(T,t)\mathbb{1}_{\{T>t\}}|\mathscr{F}_t^w]\mathbb{1}_{\{\tau=t\}}]}{\mathsf{E}_\infty[\mathsf{P}_t(T>t|\mathscr{F}_t^w)\mathbb{1}_{\{\tau=t\}}]} \geq$$

$$\inf_{t\geq 0}\operatorname{ess\,inf}\frac{\mathsf{E}_t[\phi(T,t)\mathbb{1}_{\{T>t\}}|\mathscr{F}_t^w]}{\mathsf{P}_t(T>t|\mathscr{F}_t^w)} =$$

$$\inf_{t\geq 0}\operatorname{ess\,inf}\mathsf{E}_t[\phi(T,t)|T>t,\mathscr{F}_t^w].$$

This lower bound is in fact attainable. Suppose that the last double minimization is achieved by some t_0 (minimization over t) and realization $\{w_1,\ldots,w_{t_0}\}$ (minimization over the data), then the change-imposing mechanism can simply impose a change at $\tau = t_0$ when the specific combination of data occur. If there are more choices yielding the same lower bound then it can perform randomization between them. ∎

Proof of Theorem 1: Let us consider first Criterion i). We have

$$\mathsf{P}_t(T = t+1 | T > t) = \frac{\mathsf{P}_t(T=t+1)}{\mathsf{P}_t(T>t)}$$

$$= \mathsf{E}_\infty\left[\frac{\mathsf{f}_0(\xi_{t+1}|z_{t+1})\mathsf{g}_0(z_{t+1}|z_t)}{\mathsf{f}_\infty(\xi_{t+1})\mathsf{g}_\infty(z_{t+1}|z_t)}\mathbb{1}_{\{T=t+1\}}\right]\frac{1}{\mathsf{P}_\infty(T>t)},$$

where the denominator takes this specific form because the event $\{T > t\}$ is $\mathscr{F}_t^\xi$-measurable and refers to pre-change. Since $\{T = t+1\}$ is $\mathscr{F}_{t+1}^\xi$-measurable we need to average out $z_{t+1}, \ldots, z_0$ conditioned on $\mathscr{F}_{t+1}^\xi$. This is easy since under P_∞ the observations and the Markov process are independent. Indeed this conditional expectation becomes

$$\mathsf{E}_\infty\left[\frac{\mathsf{f}_0(\xi_{t+1}|z_{t+1})\mathsf{g}_0(z_{t+1}|z_t)}{\mathsf{f}_\infty(\xi_{t+1})\mathsf{g}_\infty(z_{t+1}|z_t)}\Big|\mathscr{F}_{t+1}^\xi\right] =$$

$$\int \frac{\mathsf{f}_0(\xi_{t+1}|z_{t+1})}{\mathsf{f}_\infty(\xi_{t+1})}\mathsf{g}_0(z_{t+1}|z_t)\mathsf{g}_\infty(z_t|z_{t-1})\cdots\mathsf{g}_\infty(z_0)dz_{t+1}\cdots dz_0$$

$$= \int \frac{\mathsf{f}_0(\xi_{t+1}|z_{t+1})}{\mathsf{f}_\infty(\xi_{t+1})}\mathsf{g}_0(z_{t+1}|z_t)\mathsf{g}_\infty(z_t)dz_{t+1}dz_t = \frac{\bar{\mathsf{f}}_0^1(\xi_{t+1})}{\mathsf{f}_\infty(\xi_{t+1})},$$

where we used the fact that $\mathsf{g}_\infty(z)$ is the stationary pdf. Since $\mathcal{P}_{\text{i}}(T) \leq \mathsf{P}_t(T = t+1 | T > t)$ we can conclude that

$$\mathcal{P}_{\text{i}}(T)\mathsf{P}_\infty(T > t) \leq \mathsf{E}_\infty\left[\frac{\bar{\mathsf{f}}_0^1(\xi_{t+1})}{\mathsf{f}_\infty(\xi_{t+1})}\mathbb{1}_{\{T=t+1\}}\right].$$

Summing over $t = 0, 1, \ldots$ yields the desired inequality. The previous inequality becomes an equality when $T = \mathcal{S}_1$ because the Shewhart test is an equalizer, namely, $\mathsf{P}_t(\mathcal{S}_1 = t+1 | \mathcal{S}_1 > t)$ is a constant independent from t.

For Criterion ii) derivations are similar. Indeed we can write

$$\mathcal{P}_{\text{ii}}(T)\mathsf{P}_\infty(T > t|\mathscr{F}_t^\xi) \leq \mathsf{E}_\infty\left[\frac{\bar{\mathsf{f}}_0^1(\xi_{t+1})}{\mathsf{f}_\infty(\xi_{t+1})}\mathbb{1}_{\{T=t+1\}}|\mathscr{F}_t^\xi\right].$$

Taking expectation on both sides with respect to the P_∞ measure and summing over t yields the desired result. Again we have equality when $T = \mathcal{S}_1$ because $\mathcal{S}_1$ is an equalizer.

Let us now consider Criterion iii), we have

$$\mathcal{P}_{\text{iii}}(T)\mathsf{E}_\infty[\mathbb{1}_{\{T>t\}}|\mathscr{F}_t^z] \leq \mathsf{E}_\infty\left[\frac{\mathsf{f}_0(\xi_{t+1}|z_{t+1})\mathsf{g}_0(z_{t+1}|z_t)}{\mathsf{f}_\infty(\xi_{t+1})\mathsf{g}_\infty(z_{t+1}|z_t)}|\mathscr{F}_t^z\right].$$

Multiplying both sides with $\varpi(z_t) \geq 0$ and averaging with respect to P_∞ yields

$$\mathcal{P}_{\text{iii}}(T)\mathsf{E}_\infty[\varpi(z_t)\mathbb{1}_{\{T>t\}}] \leq$$
$$\mathsf{E}_\infty\left[\frac{\mathsf{f}_0(\xi_{t+1}|z_{t+1})\mathsf{g}_0(z_{t+1}|z_t)}{\mathsf{f}_\infty(\xi_{t+1})\mathsf{g}_\infty(z_{t+1}|z_t)}\varpi(z_t)\mathbb{1}_{\{T=t+1\}}\right].$$

For the left hand side we have

$$\mathsf{E}_\infty[\varpi(z_t)\mathbb{1}_{\{T>t\}}] = \mathsf{E}_\infty[\mathsf{E}_\infty[\varpi(z_t)|\mathscr{F}_t^\xi]\mathbb{1}_{\{T>t\}}]$$

$$= \left(\int \varpi(z_t)\mathsf{g}_\infty(z_t)dz_t\right)\mathsf{E}_\infty[\mathbb{1}_{\{T>t\}}] = \mathsf{E}_\infty[\mathbb{1}_{\{T>t\}}],$$

where we define $\pi(z) = \varpi(z)\mathsf{g}_\infty(z)$ and, without loss of generality, we assume that $\int \pi(z_t)dz_t = 1$. For the right hand side we can similarly write

$$\mathsf{E}_\infty\left[\frac{\mathsf{f}_0(\xi_{t+1}|z_{t+1})\mathsf{g}_0(z_{t+1}|z_t)}{\mathsf{f}_\infty(\xi_{t+1})\mathsf{g}_\infty(z_{t+1}|z_t)}\varpi(z_t)\mathbb{1}_{\{T=t+1\}}\right] =$$

$$\mathsf{E}_\infty\left[\mathsf{E}_\infty\left[\frac{\mathsf{f}_0(\xi_{t+1}|z_{t+1})\mathsf{g}_0(z_{t+1}|z_t)}{\mathsf{f}_\infty(\xi_{t+1})\mathsf{g}_\infty(z_{t+1}|z_t)}\varpi(z_t)|\mathscr{F}_{t+1}^\xi\right]\mathbb{1}_{\{T=t+1\}}\right] =$$

$$\mathsf{E}_\infty\left[\frac{\iint \mathsf{f}_0(\xi_{t+1}|z_{t+1})\mathsf{g}_0(z_{t+1}|z_t)\varpi(z_t)\mathsf{g}_\infty(z_t)dz_tdz_{t+1}}{\mathsf{f}_\infty(\xi_{t+1})}\mathbb{1}_{\{T=t+1\}}\right]$$

$$= \mathsf{E}_\infty\left[\frac{\bar{\mathsf{f}}_0^2(\xi_{t+1})}{\mathsf{f}_\infty(\xi_{t+1})}\mathbb{1}_{\{T=t+1\}}\right],$$

where in the last equality we use the definition in (10). The desired inequality can be shown as in the previous cases. Finally, when $T = \mathcal{S}_2$ we have equality because $\pi(z)$ puts all its mass on values of z_t where the essential infimum is attained and because the resulting value is independent from t (equilizer). Similarly we can prove the upper bound for Criterion iv). ∎

Proof of Theorem 2: The first step in the proof consists in observing that we can limit ourselves to stopping times T that satisfy the false alarm constraint with equality, that is, $\mathsf{E}_\infty[T] = \gamma$. Indeed if $\mathsf{E}[T] > \gamma$ then we can perform a randomization before taking any observations as to whether we should stop at time 0 with probability p or continue according to the stopping time T with probability $1 - p$. This generates a new stopping $\tilde{T}$ that satisfies $\mathsf{E}_\infty[\tilde{T}] = (1-p)\mathsf{E}_\infty[T]$ and therefore we can select p so that $\tilde{T}$ satisfies $\mathsf{E}_\infty[\tilde{T}] = \gamma$. On the other hand we can verify that

$$\frac{\mathsf{E}_\infty[\mathsf{L}_j(\xi_{\tilde{T}})]}{\mathsf{E}_\infty[\tilde{T}]} = \frac{\mathsf{E}_\infty[\mathsf{L}_j(\xi_T)]}{\mathsf{E}_\infty[T]}.$$

Because of the previous observation we need to prove that $\mathsf{E}_\infty[\beta_j T - \mathsf{L}_j(\xi_T)] \geq 0$ over all T satisfying the false alarm constraint with equality. In fact it will be sufficient if we consider the unconstrained problem

$$\mathsf{E}_\infty\left[(\beta_j - \tfrac{\nu_j}{\gamma})T - \mathsf{L}_j(\xi_T)\right] \geq -\nu_j \qquad (26)$$

obtained by subtracting $\frac{\nu_j}{\gamma}\mathsf{E}_\infty[T]$ from the left and ν_j from the right side. We can now assume that there is no constraint on T and minimize the left hand side in (26) over T. Since T is $\{\mathscr{F}_t^\xi\}$-adapted and $\{\xi_t\}$ under P_∞ is i.i.d., this optimal stopping problem can be easily solved and we can show that the optimum stopping time is $\mathcal{S}_j$ defined in (11). By direct computation we can also verify that the minimum value of the left hand side in (26) is indeed equal to $-\nu_j$. ∎

REFERENCES

[1] T. Bojdecki, "Probability maximizing approach to optimal stopping and its application to a disorder problem," *Stochastics*, vol. 3, pp. 61–71, 1979.

[2] C.-D. Fuh, "SPRT and CUSUM in hidden Markov models," *Ann. Stat.*, vol. 31, no. 3, pp. 942–997, 2003.

[3] C.-D. Fuh and Y. Mei, "Quickest change detection and Kullback-Leibler divergence for two-state hidden Markov models," *Trans. Inf. Theory*, vol. 63, no. 18, pp. 4866–4878, 2015.

[4] C.-D. Fuh and A. G. Tartakovsky, "Asymptotic Bayesian theory of quickest change detection for hidden Markov models," *Trans. Inf. Theory*, vol. 65, no. 1, pp. 511–529, 2019.

[5] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Stat.*, vol. 42, pp. 1897–1908, 1971.

[6] G. V. Moustakides, "Sequential change detection revisited," *Ann. Stat.*, vol. 36, no. 2, pp. 787–807, 2008.

[7] G. V. Moustakides, "Multiple optimality properties of the Shewhart test," *Seq. Anal.*, vol. 33, pp. 318–344, 2014.

[8] G. V. Moustakides, "Optimum Shewhart tests for Markovian data," 53rd Annual Allerton Conference on Communication, Control and Computing, pp. 822–826, 2015.

[9] M. Pollak, "Optimal detection of a change in distribution," *Ann. Stat.*, vol. 13, pp. 206–227, 1985.

[10] M. Pollak and A. M. Krieger, "Shewhart revisited," *Seq. Anal.*, vol. 32, pp. 230–242, 2013.

[11] W. A. Shewhart, *Economic Control of Quality of Manufactured Product*. New York: D. Van Nostrand Company, 1931.

[12] A. N. Shiryaev, "On optimal methods in quickest detection problems," *Theory Probab. Applic.* vol. 8, pp. 22–46, 1963.