Optimum Multi-Stream Sequential Change-Point Detection with Sampling Control

Qunzhi Xu, Yajun Mei, and George V. Moustakides, *Senior Member, IEEE*

***Abstract*—In multi-stream sequential change-point detection it is assumed that there are M processes in a system and at some unknown time, an occurring event changes the distribution of the samples of a particular process. In this article, we consider this problem under a sampling control constraint when one is allowed, at each point in time, to sample a single process. The objective is to raise an alarm as quickly as possible subject to a proper false alarm constraint. We show that under sampling control, a simple myopic-sampling-based sequential change-point detection strategy is second-order asymptotically optimal when the number M of processes is fixed. This means that the proposed detector, even by sampling with a rate $1/M$ of the full rate, enjoys the same detection delay, up to some additive finite constant, as the optimal procedure. Simulation experiments corroborate our theoretical results.**

***Index Terms*—Asymptotic optimality, change-point detection, myopic sampling, CUSUM, quickest detection.**

I. Introduction

SEQUENTIAL change-point detection for multi-stream data under sampling control has many important real-world applications such as quality control, surveillance or security, etc. Under a general setting, there are M processes or data streams available in a system, and at some unknown point in time, an occurring event impacts one of the available processes by changing the distribution of its samples. Unlike the conventional problem where one samples simultaneously all streams, under a sampling constraint scenario we are allowed to sample only one of the M local streams at each time. This constraint may be imposed due to sampling costs or limitations to the on-line processing power. We understand that in addition to the usual problem of developing a *stopping* strategy for signaling the detection of the change we must also provide a *sampling strategy* for the sampling of the available streams.

When we can sample all streams simultaneously without any constraint, this problem has been well studied in the sequential change-point detection literature [1]. An intuitive and efficient method is to monitor each individual process by a local CUSUM procedure and then raise a global alarm when any local CUSUM raises an alarm. This is equivalent to raising an alarm when the maximum of M local CUSUM statistics exceeds a threshold. In the sequel we will refer to this detection procedure as the full-sampling method.

When there is sampling control as to which process must be sampled at each time instant, literature is rather limited. The only existing result is [2] which proposes an algorithm that performs reasonably well in simulations but has no theoretical justification to support it. Related work is also [3] which proposes a data-efficient sampling technique but applies to a single data stream.

We should mention that sampling control has also been extensively used in two other well-known problems: sequential hypothesis testing and the multi-armed bandit problems. Pioneering work in sequential hypothesis testing for $M=2$ processes, when the data are Bernoulli distributed, can be found in [4]. Sequential hypothesis testing for the *homogeneous* case and for general M was considered in [5] but with the special requirement of identifying only a single process under the alternative hypothesis. The same switching and stopping strategy we propose in our current work, is optimum for this problem as well and, remarkably, the optimality turns out to be exact. Later in [6] we find an asymptotic optimality theory for sequential hypothesis testing problems when one is allowed to sample $K \geq 1$ out of M processes at each time. Finally, important results for the multi-armed bandit problem are offered in [7] where the first asymptotic optimality theory is developed.

The problem we are attempting to solve in this work has two major differences compared to sequential hypothesis testing and the multi-armed bandit problem: a) data collected from unaffected process provide no useful information for the affected stream, and b) data collected from the process to be affected before the change also provide no information about detection.

In this work we focus on an asymptotic optimality theory for sequential change-point detection for multi-stream data under sampling control. Our main result consists in proving that a simple myopic sampling scheme is second-order asymptotically optimum when the number M of processes is fixed. The main idea of the proposed detection strategy consists in exploring each local process periodically and decide whether or not a change took place. If we decide positively then we stop and raise a global alarm while if the decision is negative we switch to the next process. To the best of our knowledge, a second-order asymptotically optimum result is proved for the first time for the problem of multi-stream data monitoring under sampling control when time is discrete and the processes are inhomogeneous.

As one of our reviewers correctly pointed out, our proposed myopic sampling strategy is similar to the "cyclic-return system of observations over M directions" scheme proposed in [8] (see also [9]). We should emphasize that the mathematical formulation, models and technical details in [8] exhibit significant differences compared to the ones adopted here. Indeed, in [8] it is assumed that the change occurs with the same probability in any of the M local streams or directions (namely, there exists a prior). In our study we adopt a worst-case scenario which is consistent with Lorden's single stream min-max approach. In [8] the analysis is focused on continuous-time and continuous path (Wiener) processes with all local streams or directions being homogeneous. In our work we adopt discrete-time and inhomogeneous processes across different local streams. These essential dissimilarities require the employment of alternative analytical tools in order to be able to tackle problems that are not present in Wiener processes as, for example, overshoots. We must mention that the absence of overshoots in the continuous-time and continuous-path case allows for the derivation of fairly accurate expressions for the average detection delay and average time between false alarms, while in discrete time the analysis can only provide bounds of the appropriate order of magnitude. What is also interesting is that the proposed scheme in [8] turns out to be a repeated CUSUM test even though CUSUM, at the time, was not yet known for its optimality properties. Finally, we should add that we believe that under a more complicated setup where one is allowed to observe more than one local streams simultaneously, it is unclear how the methodology in [8] can be extended. On the other hand, our approach, as we discuss in Section IV, does not seem to have this problem.

The remainder of our paper is organized as follows. In Section II, we mathematically formulate the problem of interest, review existing methods and present our candidate sampling/stopping strategy. In Section III, we prove the second-order asymptotic optimality property of our scheme when M is fixed. In Section IV we provide certain remarks concerning possible extensions to more complicated scenarios and discuss the corresponding challenges. Numerical simulations are offered in Section V to illustrate the agreement between theory and practice. Finally, in Section VI we draw our conclusions and discuss future research topics. Technical proofs appear in the Appendix.

II. Problem Formulation and Background

To simplify our presentation we divide the current section into three parts. In Section II-A, we present the mathematical formulation of our problem. In Section II-B we review the CUSUM test for sequential change-point detection. We also recall the myopic and certain simple and well known sampling policies. Finally, in Section II-C, we present our candidate scheme that applies the myopic sampling policy to the problem of interest.

A. Mathematical Formulation

Suppose there are M statistically independent processes in a system, and denote with X_t^i the observation from the i-th process at time t, where $i = 1, \dots, M$ and $t = 1, 2, \dots$. Initially, the system is in the in-control state and the data stream $\{X_t^i\}$ from the ith process produces i.i.d. samples following the density $\mathsf{f}_i(X)$. At some unknown time τ, an event occurs which leads the system out-of-control with one of its M processes, say, the i-th, changing to i.i.d. samples following a new density $\mathsf{g}_i(X)$. Specifically, if the ith data stream is affected, then

$$X_t^i \sim \begin{cases} \mathsf{f}_i(X), & \text{if } t \leq \tau \\ \mathsf{g}_i(X), & \text{if } t > \tau, \end{cases} \tag{1}$$

while $X_t^j \sim \mathsf{f}_j(X)$ for $j \neq i$ and all $t > 0$.

Under the classical setup when the full data information is available at each time t we observe the complete set of M samples $\{X_t^1, \cdots, X_t^M\}$. However, when we adopt a sampling control policy then we are allowed to access only *one* of these M data points. This clearly requires the definitions of a sequence of *sampling indices* $\{R_t\}$ with $R_t \in \{1, \dots, M\}$. R_t is random and points to the process that must be sampled during the *next* time instant $t+1$. The sampling constraint can be expressed as

$$\mathbb{1}_{\{R_t=1\}} + \cdots + \mathbb{1}_{\{R_t=M\}} = 1, \qquad \text{for all times } t = 1, 2, \dots, \tag{2}$$

where $\mathbb{1}_A$ denotes the indicator function of the event A.

As we discuss next, due to the existence of various possibilities, we need to introduce several sequences of sigma-algebras (filtrations). With $\{\mathscr{F}_t^i\}$ we denote the filtration generated by the ith process, namely, $\mathscr{F}_t^i = \sigma\{X_1^i, \dots, X_t^i\}$. Then we define the filtration $\{\mathscr{F}_t\}$ containing the complete information where $\mathscr{F}_t = \mathscr{F}_t^1 \cup \cdots \cup \mathscr{F}_t^M$. Finally, by sampling one out of M processes at each time instant we generate the filtration $\{\mathscr{F}_t^R\}$ with $\mathscr{F}_t^R = \sigma\{X_1^{R_0}, \dots, X_t^{R_{t-1}}\}$. A sequential change-point detection procedure under sampling control contains two components, the sampling policy $\{R_t\}$ and the stopping time T. For the sampling policy each R_t is $\mathscr{F}_t^R$-measurable (we use the already available samples up to time t to decide which stream to sample at the next time instant $t+1$). The stopping time T is $\{\mathscr{F}_t^R\}$-adapted (uses all samples up to and including time t in order to decide whether to stop and raise an alarm at t or continue sampling according to R_t).

If our intention is to follow the min-max approach suggested by Lorden [10] it is, unfortunately, not obvious how to define our performance criterion due to the existence of these mutliple filtrations. In fact there are different possibilities we can end up with if we adopt the general performance criterion introduced in [11]. By modeling the change-point time τ as a stopping time (the time we stop using the nominal data) then, if τ is $\{\mathscr{F}_t\}$-adapted, namely has access to the whole information, and after following a worst-case scenario for the unknown change-time τ, we can define the performance criterion as

$$D(T) = \sup_{t \geq 0} \operatorname{ess\,sup} \mathsf{E}_t[T - t | \mathscr{F}_t, T > t], \tag{3}$$

where with $\mathsf{P}_t(\cdot), \mathsf{E}_t[\cdot]$ we denote the probability measure and the corresponding expectation induced by the change occurring at $\tau = t$. Alternatively, we can have change-point mechanisms that rely on individual streams and decide when to impose the change in each process by using the local samples. In this case we must define a criterion for each stream

$$D_i(T) = \sup_{t \geq 0} \operatorname{ess\,sup} \mathsf{E}_t^i[T - t | \mathscr{F}_t^i, T > t]. \tag{4}$$

where now $\mathsf{P}_t^i(\cdot), \mathsf{E}_t^i[\cdot]$ denote the probability measure and the corresponding expectation induced by the change occuring at Process i at time $\tau = t$. It is the second performance criterion in (4) which is mostly adopted in the literature and we also intend to follow here in our analysis. Regarding the detection part, we must point out that the stopping time T can be adapted to any *other* filtration which is not necessarily the same as the one used by the change-time τ. Indeed T depends on the information which becomes available to the scientist responsible for detection. This information does not have to be the same as the information employed by the change-point mechanism for deciding when to impose the change [11]. Regardless of this observation, the goal is always the same, namely, to solve the constrained min-max problem

$$\inf_{T,\{R_t\}} D_i(T) = \inf_{T,\{R_t\}} \sup_{t\geq 0} \text{ess sup } \mathsf{E}_t^i[T-t|\mathscr{F}_t^i, T>t], \quad (5)$$

subject to $\mathsf{E}_\infty[T] \geq \gamma > 1$, where $\mathsf{P}_\infty(\cdot), \mathsf{E}_\infty[\cdot]$ denote the probability measure and the corresponding expectation under the nominal probability measure (namely when the change occurs at ∞) and T and $\{R_t\}$ are defined over the appropriate filtration. As we can see the false alarm constraint requires the average false alarms period to be no less than some prescribed value $\gamma > 1$.

For the constrained problem in (5) it is very unrealistic to expect that we can find a *single* sampling/stopping policy capable of *exactly* optimizing it, simultaneously for all $i = 1, \ldots, M$. The goal of this work is to show that in fact such a possibility exists (by proposing a particular solution) but within the class of policies that are *second-order asymptotically optimum*.

B. Review of Change-Point Methods

Focusing on solving (5) with $D_i(T)$ defined in (4) we first consider the case where there is a genie that provides the index i of the process where the change occurs. If we know i then there is no reason to sample any other process, consequently $R_t = i$ at all times and we can limit T to be $\{\mathscr{F}_t^i\}$-adapted. In this case it is well known that the optimum stopping time is the CUSUM defined as

$$\mathsf{T}_i(A_i) = \inf\{t > 0 : W_t^i \geq A_i\}, \quad (6)$$

where W_t^i is the CUSUM statistic [12] that satisfies the recursion

$$W_t^i = \max\{W_{t-1}^i, 0\} + \log\frac{\mathsf{g}_i(X_t^i)}{\mathsf{f}_i(X_t^i)}, \quad (7)$$

for $t > 0$ and is initialized with $W_0^i = 0$. Threshold A_i is selected so that the false alarm constraint is met with equality. Proof for first-order asymptotic optimality was offered in Lorden [10] while exact optimality can be found in [13]. It is clear that there is no detection strategy that can outperform the CUSUM test that knows where (but not when) the change occurs.

A more practically interesting scenario consists in having access to the complete data set but not knowing where and when the change occurs. Again, there is no need to specify a sampling strategy $\{R_t\}$ since we sample all processes. Cleary T is now $\{\mathscr{F}_t\}$-adapted, namely the stopping time uses all the information up to time t to decide whether to stop at t or not. In this case we run local CUSUMs in parallel, one for each process and as it is proven in [1] we raise an alarm whenever one of the M stopping times stops. More specifically if T_i is defined as in (6) then the combination

$$\mathsf{T}_{\text{Full}} = \min\{\mathsf{T}_1(A), \cdots, \mathsf{T}_M(A)\}, \quad (8)$$

with all CUSUM having the same threshold A, is asymptotically optimum in the sense that is solves the problem defined in (5) asymptotically as $\gamma \to \infty$. In fact we can rewrite $\mathsf{T}_{\text{Full}}(A)$ as

$$\mathsf{T}_{\text{Full}}(A) = \inf\{t > 0 : \max_{1\leq i\leq M} W_t^i \geq A\}, \quad (9)$$

where, essentially, in the test we apply the generalized likelihood ratio with respect to index i. In Theorems 9.2.1 and 9.2.2 of [1], it is shown that $\mathsf{T}_{\text{Full}}(A)$ is second-order asymptotically optimum when the number M of processes is fixed. This means that we minimize each detection delay $D_i(T)$ up to an $O(1)$ quantity which is independent from the false alarm constraint parameter γ in (5).

C. Review of Sampling Policies

Under sampling control there is the need to define a sampling policy $\{R_t\}$ since, as we explained, at each time instant we are allowed to sample only one out of the M processes. Consequently let us review sampling possibilities that are widely adopted in the literature.

The most straightforward policy consists in sampling each process periodically meaning that $R_t = t \bmod M + 1$ where we visit each process deterministically every M samples. It is not difficult to show that this sampling strategy will lead to a detection delay which is M times larger than the optimum which, as mentioned, is enjoyed by the CUSUM stopping time T_i that knows where the change occurs. Clearly, this observation makes unrealistic any expectation for establishing second-order (in fact even first-order) asymptotic optimality with this form of sampling.

An alternative widely used policy is the myopic (or greedy) sampling policy (MSP) which samples the process exhibiting the maximal immediate reward. A general implementation of the myopic sampling policy is to define local statistics $\widetilde{W}_t^i, i = 1, \cdots, M$ that summarize the immediate sampling rewards for each process at time t and then sample the process with the largest local statistic $\widetilde{W}_t^i$ (with random sampling or pre-assigned order in case of ties). This type of sampling is frequent in the multi-armed bandit problem and there is extensive literature as to which is the most suitable selection for the local statistic $\widetilde{W}_t^i$.

Our intention is to use the myopic sampling policy for the problem of interest. It is clear that a natural candidate for the local statistic $\widetilde{W}_t^i$ is the CUSUM statistic defined in (7). However, in order for our sampling/stopping policy to be complete, we need to explicitly specify three points that are unclear: (i) How should we update the local statistic of a process not being sampled. (ii) How should we break ties when the largest local statistic occurs at multiple processes.

This is particularly important when all local CUSUM statistics become simultaneously 0 (which is a very frequent event under the nominal regime). (iii) When should we raise a global alarm. While it is possible to give intuitively meaningful answers to these three points, the challenge is to accompany them with theoretical justification capable of establishing the desired form of asymptotic optimality.

III. Main Results

In this section we introduce our candidate sampling/stopping strategy and establish its asymptotic optimality characteristics. Again we break our presentation into parts. In Section III-A we define our candidate stopping time $\mathsf{T}_{\mathrm{MSP}}$ implemented with the help of the myopic sampling policy. We also provide answers to the three points we mentioned in the previous section. In Section III-B, we study the non-asymptotic properties of the false-alarm and detection delay of $\mathsf{T}_{\mathrm{MSP}}$. We conclude the presentation of our main results by establishing the second-order asymptotic optimality of our candidate test, in Section III-C.

A. Candidate Sampling/Stopping Strategy

At a high level, our algorithm is based on the myopic sampling policy and mimics the full-sampling method $\mathsf{T}_{\mathrm{Full}}(A)$ in (9) under the sampling control constraint. Here we exploit the prior knowledge that there is only one process which changes, and thus propose to sample each process until we are confident to decide whether a change has occurred or not. If we detect a change, then we stop and raise a global alarm. If we decide there is no change, then we deterministically switch to the next process to sample and repeat the previous step. Switching follows a periodic pattern starting from Process 1 and going to Process M. When we reach Process M and decide to switch again we simply restart from Process 1. We repeat these steps until we raise an alarm.

Let us now define our scheme rigorously through the recursive definition of the statistics $\widetilde{W}_t^i,\ i=1,\dots,M$ and the sampling sequence $\{R_t\}$. At time $t-1$ assume we already have available $\widetilde{W}_{t-1}^i, i=1,\dots,M$ and R_{t-1} with the latter pointing to the next process to be sampled at time t. For the local statistics we have the recursions

$$\begin{aligned}\widetilde{W}_t^i &= \max\{\widetilde{W}_{t-1}^i,0\}+\mathbb{1}_{\{i=R_{t-1}\}}\log\frac{\mathsf{g}_i(X_t^i)}{\mathsf{f}_i(X_t^i)}\\ &=\begin{cases}\max\{\widetilde{W}_{t-1}^i,0\}, & \text{if } i\neq R_{t-1}\\ \max\{\widetilde{W}_{t-1}^i,0\}+\log\frac{\mathsf{g}_i(X_t^i)}{\mathsf{f}_i(X_t^i)}, & \text{if } i=R_{t-1},\end{cases}\end{aligned} \tag{10}$$

and for the sampling sequence

$$R_t=\begin{cases}R_{t-1} & \text{if } \widetilde{W}_t^{R_{t-1}}>0\\ R_{t-1} \bmod M+1 & \text{if } \widetilde{W}_t^{R_{t-1}}\leq 0.\end{cases} \tag{11}$$

For all $i=1,\cdots,M$ we initialize with $\widetilde{W}_0^i=0$ and $R_0=1$. Similarly to the full-sampling method $\mathsf{T}_{\mathrm{Full}}(A)$ in (9), we propose the stopping time $\mathsf{T}_{\mathrm{MSP}}$ defined as

$$\mathsf{T}_{\mathrm{MSP}}(A)=\inf\left\{t>0:\max_{1\leq i\leq M}\widetilde{W}_t^i\geq A\right\}, \tag{12}$$

where threshold A is selected to meet the false alarm constraint.

We observe that we continue sampling Process i as long as $\widetilde{W}_t^i>0$ and switch to the next process when $\widetilde{W}_t^i\leq 0$. It is easy to see that we can equivalently define $\mathsf{T}_{\mathrm{MSP}}$ with a single test statistic that satisfies the update

$$\widetilde{W}_t=\max\{\widetilde{W}_{t-1},0\}+\log\frac{\mathsf{g}_{R_{t-1}}(X_t^{R_{t-1}})}{\mathsf{f}_{R_{t-1}}(X_t^{R_{t-1}})}, \tag{13}$$

with $\widetilde{W}_0=0$ and the sampling policy

$$R_t=\begin{cases}R_{t-1} & \text{if } \widetilde{W}_t>0\\ R_{t-1} \bmod M+1 & \text{if } \widetilde{W}_t\leq 0,\end{cases} \tag{14}$$

with $R_0=1$. The stopping time $\mathsf{T}_{\mathrm{MSP}}$ can then be equivalently written as

$$\mathsf{T}_{\mathrm{MSP}}(A)=\inf\left\{t>0:\widetilde{W}_t\geq A\right\}. \tag{15}$$

The recursion in (13) and the definition of our stopping time in (15) are clearly very CUSUM-like, the only difference being that instead of always sampling the same process which is the practice in the regular CUSUM, every time the test statistic $\widetilde{W}_t$ falls bellow 0, we switch to testing the next process by restarting and forgetting the whole past. Let us summarize the proposed scheme.

Step 1: Sample Process 1 until $\widetilde{W}_t\notin(0,A)$. If $\widetilde{W}_t\geq A$, we stop sampling and raise a global alarm; otherwise if $\widetilde{W}_t\leq 0$, we switch to sampling Process 2.

Step 2: Sample Process 2 until $\widetilde{W}_t\notin(0,A)$. If $\widetilde{W}_t\geq A$, we stop sampling and raise a global alarm; otherwise if $\widetilde{W}_t\leq 0$, we switch to sampling Process 3.

⋮

Step M: Sample Process M until $\widetilde{W}_t\notin(0,A)$. If $\widetilde{W}_t\geq A$, we stop sampling and raise a global alarm; otherwise if $\widetilde{W}_t\leq 0$, we switch to sampling Process 1.

Step $M+1$: Go back to Step 1.

The reason we expect that $\mathsf{T}_{\mathrm{MSP}}$ defined in (12) will enjoy second-order optimality properties is because when we start sampling a process we practically apply a sequential probability ratio test (SPRT) with the lower threshold set to 0. We recall that the classical CUSUM is also a *repeated* SPRT test with lower threshold equal to 0 only, as mentioned, it is always applied onto the *same* process. Here, what we propose is that every time we restart the SPRT we switch to the next process. Processes that do not change or the process which will change but is still under the pre-change state drive the SPRT to 0 very quickly with short random periods. When we hit post-change data then with high probability the corresponding SPRT will remain at this process and drive its statistic towards the high threshold to raise an alarm.

B. Finite-Sample Properties

To establish the desired asymptotic optimality characteristic for $\mathsf{T}_{\mathrm{MSP}}$ we first need to introduce certain finite-sample properties. We start by making some standard assumptions encountered in the classical sequential detection literature. For $i = 1, \dots, M$ we have

$$\text{(A1):}\quad \mathsf{I}_\infty^i = \int \log \frac{\mathsf{f}_i(X)}{\mathsf{g}_i(X)} \mathsf{f}_i(X) dX > 0,$$
$$\mathsf{I}_0^i = \int \log \frac{\mathsf{g}_i(X)}{\mathsf{f}_i(X)} \mathsf{g}_i(X) dX > 0,$$
$$\text{(A2):}\quad \mathsf{J}_\infty^i = \int \left(\log \frac{\mathsf{f}_i(X)}{\mathsf{g}_i(X)}\right)^2 \mathsf{f}_i(X) dX < \infty,$$
$$\mathsf{J}_0^i = \int \left(\log \frac{\mathsf{g}_i(X)}{\mathsf{f}_i(X)}\right)^2 \mathsf{g}_i(X) dX < \infty,$$

In other words we define the information numbers and make the Assumption (A1) that they are bounded away from 0 meaning that pre- and post-change densities must be essentially different. Assumption (A2) is technical and states that the second moments of the log-likelihood ratios are bounded away from ∞. Clealy (A2) implies that the information numbers are also bounded away from ∞.

To establish the second-order asymptotic optimality, we need to compare the performance of our proposed stopping time $\mathsf{T}_{\rm MSP}$ against the *optimum performance* delivered by the CUSUM stopping time $\mathsf{T}_i(A_i)$. Since there are no exact formulas for both schemes we present useful estimates that will allow us to achieve our goal. We start with the CUSUM test for which the next lemma provides the required estimates most of which are already established in the literature.

LEMMA 1 *Under Assumptions (A1), (A2), the CUSUM tests* $\mathsf{T}_i(A),\ i = 1, \dots, M$, *satisfy the following bounds for the average period of false alarms*

$$e^A \le \mathsf{E}_\infty[\mathsf{T}_i(A)] \le \mathcal{C}_i e^A, \tag{16}$$

while for the worst-case average detection delays

$$\frac{A}{\mathsf{I}_0^i} - \mathcal{L}_i \le D_i\big(\mathsf{T}_i(A)\big) \le \frac{A}{\mathsf{I}_0^i} + \mathcal{U}_i. \tag{17}$$

The quantities $\mathcal{C}_i, \mathcal{U}_i, \mathcal{L}_i$ *are positive constants that depend only on* $\mathsf{f}_i, \mathsf{g}_i$ *and not on* A.

Proof: The results of this lemma are already available in the literature. But, for completeness we decided to include them here with a parallel effort to provide explicit formulas for the constant parameters appearing in the estimates. All details are given in the Appendix. ■

The next theorem provides corresponding estimates for the proposed stopping time $\mathsf{T}_{\rm MSP}$.

THEOREM 1 *Under Assumptions (A1), (A2) the proposed stopping time* $\mathsf{T}_{\rm MSP}(A)$ *defined in* (12) *satisfies the following lower bound for the average false alarm period*

$$e^A \le \mathsf{E}_\infty[\mathsf{T}_{\rm MSP}], \tag{18}$$

while for the worst-case average detection delays for $i = 1, \dots, M$, *satisfy the upper bounds*

$$D_i(\mathsf{T}_{\rm MSP}) \le \frac{A}{\mathsf{I}_0^i} + \mathcal{U}_i + \mathcal{D}(M-1), \tag{19}$$

where $\mathcal{U}_i$ *are the constants from* (17) *in Lemma 1 and* $\mathcal{D}$ *is a constant that depends on all* $\{\mathsf{f}_i, \mathsf{g}_i\},\ i = 1, \dots, M$ *but not on* A.

Proof: At this point let us only provide a high-level sketch of the proof deferring the technical details to the Appendix. The key idea is to relate $\mathsf{T}_{\rm MSP}$ to the following M *prototype* SPRTs: For $i = 1, \cdots, M$, a prototype SPRT applied to the ith process is defined as

$$\mathcal{T}^i = \inf\left\{ t > 0 : S_t^i = \sum_{\ell=1}^t \log \frac{\mathsf{g}_i(X_\ell^i)}{\mathsf{f}_i(X_\ell^i)} \notin (0, A) \right\}. \tag{20}$$

Consider now the sequence $\{\mathcal{T}_\ell\}, \ell = 1, 2, \dots,$ of SPRTs applied to the data streams when we employ the periodic sampling policy. Then it is clear that each $\mathcal{T}_\ell$ has the same distribution as a particular prototype SPRT $\mathcal{T}^i$. In fact ℓ and i are related through the equation $i = (\ell - 1) \bmod M + 1$ suggesting that the distributions of $\{\mathcal{T}_1, \mathcal{T}_2, \dots\}$ change periodically with period M. In addition, if we define (a stopping time) k to be the first time the SPRT of the process being tested crosses the upper boundary A, then our proposed stopping time $\mathsf{T}_{\rm MSP}$ can be written as the sum

$$\mathsf{T}_{\rm MSP} = \mathcal{T}_1 + \mathcal{T}_2 + \cdots + \mathcal{T}_k = \sum_{\ell=1}^k \mathcal{T}_\ell. \tag{21}$$

The average of $\mathsf{T}_{\rm MSP}$ can therefore be computed by analyzing the sequence $\{\mathcal{T}_1, \mathcal{T}_2, \dots\}$ with each stopping time, as mentioned, matching in distribution one of the M prototype SPRTs $\{\mathcal{T}^1, \dots, \mathcal{T}^M\}$. ■

C. Second-Order Asymptotic Optimality

Using Lemma 1 and Theorem 1 we are now able to establish the second-order asymptotic optimality property for $\mathsf{T}_{\rm MSP}$. The optimum detection delay grows to infinity as the false alarm parameter $\gamma \to \infty$. Our intention is to show that $\mathsf{T}_{\rm MSP}$ has a detection delay which, for each i, grows to infinity at the same rate as the optimum CUSUM test that knows where the change occurs. More specifically we will show that the two performances can differ at most by a bounded constant which does not depend on γ.

If we know that the change is going to occur at Process i then the best detection delay performance, as we mentioned, is delivered by the CUSUM stopping time $\mathsf{T}_i(A_i)$ with A_i selected so that the false alarm constraint is safisfied with equality. In fact there is absolutely no other stopping time that can enjoy better performance since $\mathsf{T}_i(A_i)$ uses information which is absolutely relevant to the corresponding detection task while any other stopping time defined on a different filtration will use information that is not related to the change at Process i. We have the following corollary that establishes the second-order optimality of $\mathsf{T}_{\rm MSP}$ *simultaneously* for all $i = 1, \dots, M$.

COROLLARY 1 *Let* $A = \log\gamma$, *then our proposed stopping time* $\mathsf{T}_{\rm MSP}(A)$ *defined in (12) satisfies both the false alarm and the sampling control constraint. If for each* $i = 1, \dots, M$, *we have the optimum CUSUM tests* $\mathsf{T}_i(A_i)$ *with* A_i *selected to satisfy the false alarm constraint with equality, then*

$$0 \le D_i\big(\mathsf{T}_{\rm MSP}(A)\big) - D_i\big(\mathsf{T}_i(A_i)\big) \le CM, \tag{22}$$

for proper constant C.

Proof: Since we are interested in a change at Process i, in order to establish second-order asymptotic optimality we need to compare the performance of our test with the optimum performance delivered by CUSUM. The first step consists in estimating this performance. Note that $\mathsf{T}_i(A_i)$ and $\mathsf{T}_{\rm MSP}(A)$ are two different stopping times, therefore they do not necessarily share the same threshold. In fact we need to compute the threshold A_i of the CUSUM test that delivers the optimum performance and this is guaranteed only by the threshold that meets the false alarm constraint with equality[1]. Since for the average false alarm period we have bounds from (16) in Lemma 1, if we replace A with A_i and set $\mathsf{E}_\infty[\mathsf{T}_i(A_i)]=\gamma$ then we conclude that $\log\mathcal{C}_i+A_i\geq\log\gamma\geq A_i$ or equivalently $\log\gamma\geq A_i\geq\log\gamma-\log\mathcal{C}_i$, suggesting that the optimum threshold A_i that meets the false alarm constraint may differ from $\log\gamma$ at most by a constant independent from γ. Substituting the bounds for A_i in the bounds for the worst-case detection delay of CUSUM in (17), we obtain

$$\frac{\log\gamma}{\mathsf{I}_0^i}-\mathcal{L}_i-\frac{\log\mathcal{C}_i}{\mathsf{I}_0^i}\leq D_i\big(\mathsf{T}_i(A_i)\big)\leq\frac{\log\gamma}{\mathsf{I}_0^i}+\mathcal{U}_i. \quad (23)$$

As we can see the optimum performance can differ from $\frac{\log\gamma}{\mathsf{I}_0^i}$ at most by a constant.

We now focus on the proposed $\mathsf{T}_{\rm MSP}(A)$ then, from (18) in Theorem 1 if we select $A=\log\gamma$ we can see that $\mathsf{T}_{\rm MSP}$ satisfies the false alarm constraint. We must point out that this threshold *does not* deliver the best performance for $\mathsf{T}_{\rm MSP}$ which is reserved for the threshold that meets the false alarm constraint with equality. But if we can show that this non-best performance differs from the optimum (delivered by CUSUM) by only a constant the same will be true for the best performance. The proposed stopping time $\mathsf{T}_{\rm MSP}$ with threshold $A=\log\gamma$ exhibits a worst-case average detection delay from (19) that satisfies

$$D_i(\mathsf{T}_{\rm MSP})\leq\frac{A}{\mathsf{I}_0^i}+\mathcal{U}_i+\mathcal{D}(M-1)=\frac{\log\gamma}{\mathsf{I}_0^i}+\mathcal{U}_i+\mathcal{D}(M-1). \quad (24)$$

Since $\mathsf{T}_i(A_i)$ provides the smallest detection delay for a change occurring in Process i we can write using the lower bound from (23)

$$0\leq D_i(\mathsf{T}_{\rm MSP})-D_i(\mathsf{T}_i)\leq\mathcal{U}_i+\mathcal{L}_i+\frac{\log\mathcal{C}_i}{\mathsf{I}_0^i}+\mathcal{D}(M-1). \quad (25)$$

If we select $C=\max\{\mathcal{D},\max_{1\leq i\leq M}(\mathcal{U}_i+\mathcal{L}_i+\frac{\log\mathcal{C}_i}{\mathsf{I}_0^i})\}$ we can use it to strengthen (25) and prove (22). This completes the proof. ■

IV. Remarks

Let us now discuss possible extensions and corresponding challenges we may encounter.

1) Our theoretical results still hold if one uses the slightly different form of the CUSUM statistic $\hat{W}_t=\max\{\widetilde{W}_t,0\}$. We have the following update of the test statistic for this version

$$\hat{W}_t=\max\left\{\hat{W}_{t-1}+\log\frac{\mathsf{g}_{R_{t-1}}(X_t^{R_{t-1}})}{\mathsf{f}_{R_{t-1}}(X_t^{R_{t-1}})},0\right\} \quad (26)$$

while the sampling policy becomes

$$R_t=\begin{cases}R_{t-1} & \text{if } \hat{W}_t>0\\ R_{t-1}\bmod M+1 & \text{if } \hat{W}_t=0,\end{cases} \quad (27)$$

As we can see by comparing (13) to (26) in this version the attractive linearity property no longer holds for $\hat{W}_t$.

2) The second-order asymptotic optimality property is assured under the assumption that the number M of processes is fixed. Since our finite-sampling estimates hold for any given M we can easily deduce that our proposed scheme can still enjoy *first-order* asymptotically optimality when $M\to\infty$ provided that the constants entering in the estimates of the CUSUMs and the proposed $\mathsf{T}_{\rm MSP}$ are uniformly bounded in M (for example when all processes follow the same pre- and post-change density) while the number of processes grows to ∞ as $M=o(\log\gamma)$.

3) A possible extension is to allow sampling of more than one processes at each time instant. This can be particularly advantageous when the number M of processes is large or when more than one processes may be affected (change) simultaneously. The sampling control in (2) can therefore be relaxed to

$$\mathbb{1}_{\{R_t=1\}}+\cdots+\mathbb{1}_{\{R_t=M\}}=Q, \quad (28)$$

for all times $t=1,2,\ldots$ and some integer $Q\in\{1,\ldots,M\}$. In particular, when $Q=M$, this corresponds to the full-sampling scenario. We can now extend the myopic sampling policy to this more general case by sampling the Q processes corresponding to the Q largest $\widetilde{W}_t^i$ values. Unfortunately, analyzing this more general sampling scheme is not as simple as the case $Q=1$ we already examined. For the analysis we note that we have now Q SPRTs running in parallel which, unfortunately, are not synchronized. It is in fact this last observation that makes the analysis challenging.

4) The previous general case enjoys a considerable simplification when we can have a change in only one process. A possible solution strategy consists in adopting a block sampling policy that satisfies the sampling control requirement (2) as follows: We divide the M processes into Q blocks, where each block contains roughly $[\frac{M}{Q}]$ local processes. In each block, we sample cyclically following the procedure proposed for the case $Q=1$ and we update the statistic of each process as in (13). We raise a global alarm when any statistic crosses the upper threshold.
Regarding the analysis, denote by $\widetilde{\mathsf{T}}_{\rm MSP}$ the corresponding block-sampling-based scheme subject to the sampling control in (28). Then it can be shown that

$$\frac{e^A}{Q}\leq\mathsf{E}_\infty[\widetilde{\mathsf{T}}_{\rm MSP}]\quad\text{and}\quad D_i(\widetilde{\mathsf{T}}_{\rm MSP})\leq\frac{A}{\mathsf{I}_0^i}+\mathcal{U}_i+\mathcal{D}\frac{M}{Q}. \quad (29)$$

[1]We would like to emphasize that for the optimum performance, which is our point of reference for asymptotic optimality, it is not enough to simply satisfy the false alarm constraint, *we need to satisfy it with equality*. Otherwise the performance we compute is clearly not the optimum.

Quantity $\mathcal{D}$ is the constant already introduced in (19) of Theorem 1 and it is related to all $\{\mathsf{f}_i, \mathsf{g}_i\},\ i = 1, \ldots, M$. To satisfy the false-alarm constraint in (5) we select $A = \log\gamma + \log Q$ and this yields

$$D_i(\widetilde{\mathsf{T}}_{\rm MSP}) \leq \frac{\log\gamma}{\mathsf{I}_0^i} + \frac{\log Q}{\mathsf{I}_0^i} + \mathcal{U}_i + \mathcal{D}\frac{M}{Q}, \tag{30}$$

which can ensure second-order asymptotic optimality when M is fixed. In the case where we allow $M, Q \to \infty$ it is possible to enjoy first-order asymptotic optimality if the following two rates are satisfied: $\log Q = o(\log\gamma)$ and $\frac{M}{Q} = o(\log\gamma)$. The latter is a clear improvement over the rate $M = o(\log\gamma)$ required when $Q = 1$.

5) Under a high-dimensional setting where $M \gg \log\gamma$ if multiple processes undergo a change, myopic sampling policy might turn out to be overly greedy and it might be necessary to encourage the exploration of processes that have not been sampled. A possible means to achieve this (see [2]) is to introduce a compensation coefficient $\Delta \geq 0$ to the local processes that are not being sampled. In other words modify the update in (10) as follows

$$\widetilde{W}_t^i = \begin{cases} \max\{\widetilde{W}_{t-1}^i, 0\} + \Delta, & \text{if } i \neq R_{t-1} \\ \max\{\widetilde{W}_{t-1}^i, 0\} + \log\frac{\mathsf{g}_i(X_t^i)}{\mathsf{f}_i(X_t^i)}, & \text{if } i = R_{t-1}, \end{cases} \tag{31}$$

which increases the chance of the myopic sampling policy to select unobserved processes. It is shown in [2] numerically that a suitable choice of $\Delta > 0$ can significantly improve performance, but it is still an open problem the rigorous theoretical analysis of this sampling scheme.

V. Simulations

In this section, we conduct Monte Carlo simulation studies to corroborate our theoretical results. Assume $\mathsf{f}_i = \mathsf{f} \sim \mathcal{N}(0,1)$, and $\mathsf{g}_i = \mathsf{g} \sim \mathcal{N}(\mu,1)$. We consider $\mu = 0.5, 1$ and $M = 2, 3, 5$ and simulate all six combinations of μ and M values. In each of these cases, we report the detection delay of our proposed stopping time $\mathsf{T}_{\rm MSP}$ under the sampling control $Q = 1$ (blue) and we compare it against the optimum CUSUM procedure T_i which knows where the change occurs (black) and $\mathsf{T}_{\rm Full}$ which has access to the full data set but does not know where the change occurs (red).

All competing schemes have either a CUSUM or a CUSUM-like update consequently we can safely claim that the worst case scenario for the detection delay is when the change occurs at $\tau = 0$. Furthermore, regarding our scheme, we consider the least favorable scenario of the change occuring in the Mth process which is the last to be sampled by our method. This clearly adds an extra initial delay until our test rejects the first $M-1$ processes and starts sampling the correct process. We compare the detection delay of the three tests as a function of the average false alarm period. This is achieved by performing $100,000$ Monte Carlo independent runs for different values of the threshold A.

From Fig. 1, it is clear that the gap between the three curves remains bounded for all values of γ (actually the curves

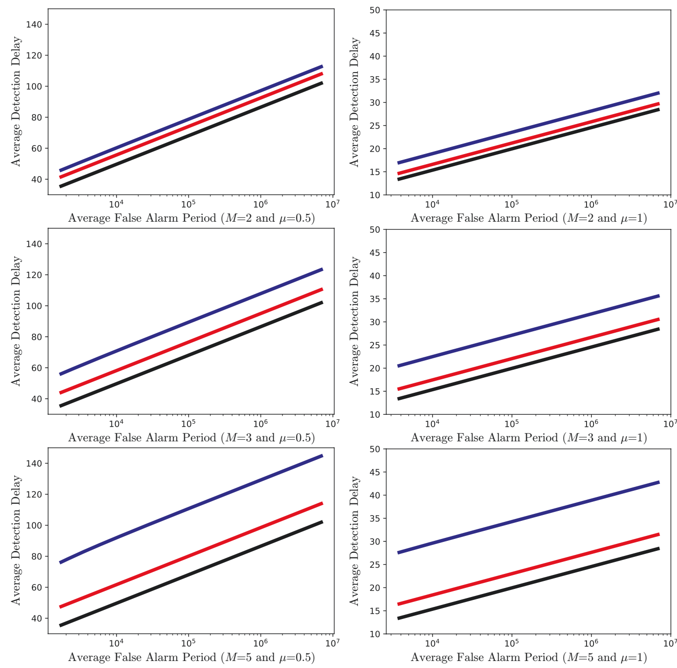

Fig. 1. Average detection delay as a function of average false alarm period for proposed $\mathsf{T}_{\rm MSP}$ (blue), full-sampling $\mathsf{T}_{\rm Full}$ (red) and optimum CUSUM T_M (black) for $M = 2, 3, 5$ for detecting a change in the mean of a Gaussian from 0 to $\mu > 0$. Top figures correspond to $\mu = 0.5$ and bottom to $\mu = 1$.

become parallel suggesting that the gap tends to a constant). This is consistent with our theoretical result that the proposed scheme is second-order asymptotically optimum when M is fixed. We also observe that the gap between the blue and the black curve increases with M which is again consistent with the upper bound recorded in (25) in Corollary 1. A similar phenomenon occurs for $\mathsf{T}_{\rm Full}$ but, as we can see, it is far less pronounced which is of course expected since for this test we sample all processes simultaneously and therefore there is no initial delay to start sampling the affected process.

VI. Conclusion

Sequential change-point detection with sampling control is an important and challenging topic with many applications. In this work we developed a detection strategy based on the simple idea of myopic sampling. Interestingly with this simple sampling scheme we are able to establish second-order asymptotic optimality when the number M of processes is fixed and first-order when M increases to infinity but at a proper rate as compared to the false alarm parameter γ.

Future work may include extension to the case where we have multiple processes that simultaneously undergo a change and we are allowed to sample more than one processes simultaneously at each time instant.

Appendix

Proof of Lemma 1: Even though the estimates of Lemma 1 are already established elsewhere, for completeness we highlight the corresponding proofs referencing the original work where these proofs appear. The reason in doing so is to present explicit formulas for the various constants that appear in these estimates.

Important elements for demonstrating the necessary bounds constitute the *ladder variables*, see [14], Chapter VIII.4, defined as

$$\begin{aligned}\tau_-^i&=\inf\left\{t>0: S_t^i=\sum_{\ell=1}^t\log\frac{\mathsf{g}_i(X_\ell^i)}{\mathsf{f}_i(X_\ell^i)}\leq 0\right\},\\ \tau_+^i&=\inf\left\{t>0: S_t^i=\sum_{\ell=1}^t\log\frac{\mathsf{g}_i(X_\ell^i)}{\mathsf{f}_i(X_\ell^i)}> 0\right\}.\end{aligned}\tag{32}$$

We also need to define a third stopping time for $A\geq 0$

$$\tau_A^i=\inf\left\{t>0: S_t^i=\sum_{\ell=1}^t\log\frac{\mathsf{g}_i(X_\ell^i)}{\mathsf{f}_i(X_\ell^i)}\geq A\right\}.\tag{33}$$

From classical renewal theory and under Assumptions (A1), (A2), we have that $\mathsf{E}_0^i[\tau_+^i],\mathsf{E}_\infty[\tau_-^i]$ are both finite constants that depend only on $\{\mathsf{f}_i,\mathsf{g}_i\}$, see [17], Theorem C, p. 748 or [18]. This is true because the statistic S_t^i is a random walk with positive drift $\mathsf{I}_0^i>0$ under the post-change regime and negative drift $-\mathsf{I}_\infty^i<0$ under the nominal regime. We also have the following interesting equalities from [14], Corollary 8.39

$$\mathsf{E}_0^i[\tau_+^i]=\frac{1}{\mathsf{P}_0^i(\tau_-^i=\infty)},\quad \mathsf{E}_\infty[\tau_-^i]=\frac{1}{\mathsf{P}_\infty(\tau_+^i=\infty)},\tag{34}$$

while for τ_A^i a useful estimate [18], Theorem 1 regarding the overshoot of the threshold A

$$\mathsf{E}_0^i[S_{\tau_A^i}^i-A]\leq\frac{\mathsf{J}_0^i}{\mathsf{I}_0^i},\tag{35}$$

with the upper bound being constant and valid for all $A\geq 0$.

We would like to relate the three stopping times $\tau_-^i,\tau_+^i,\tau_A^i$ to the SPRT $\mathcal{T}^i$ defined in (20). The first observation is that the SPRT can be written as $\mathcal{T}^i=\min\{\tau_-^i,\tau_A^i\}$. Regarding the sequential hypothesis testing procedure implemented by the SPRT we have the two error probabilities $\alpha_i=\mathsf{P}_\infty(S_{\mathcal{T}^i}^i\geq A)$ which is the Type-I error and $\beta_i=\mathsf{P}_0^i(S_{\mathcal{T}^i}\leq 0)$ which is the Type-II. For τ_-^i we can immediately conclude that for $i=1,\dots,M$ we have $\mathcal{T}^i=\min\{\tau_-^i,\tau_A^i\}\leq\tau_-^i$. We also note that the event $\{S_{\mathcal{T}^i}^i\leq 0\}$ implies $\{\tau_-^i<\infty\}$, namely that τ_-^i will also stop. However, it is clear that we can have the latter event occurring without the former. This suggests for the Type-II error that $\beta_i=\mathsf{P}_0^i(S_{\mathcal{T}^i}^i\leq 0)\leq\mathsf{P}_0^i(\tau_-^i<\infty)$ and in combination with (34) we have

$$\frac{1}{1-\beta_i}\leq\frac{1}{1-\mathsf{P}_0^i(\tau_-^i<\infty)}=\frac{1}{\mathsf{P}_0^i(\tau_-^i=\infty)}=\mathsf{E}_0^i[\tau_+^i].\tag{36}$$

Since CUSUM is a repeated SPRT applied to the same process we have [14]

$$\begin{aligned}\mathsf{E}_\infty[\mathsf{T}_i(A)]&=\frac{\mathsf{E}_\infty[\mathcal{T}^i]}{\alpha_i},\\ D_i\big(\mathsf{T}_i(A)\big)&=\mathsf{E}_0^i[\mathsf{T}_i(A)]=\frac{\mathsf{E}_0^i[\mathcal{T}^i]}{1-\beta_i}.\end{aligned}\tag{37}$$

In (37) we also indicated the well known fact that the worst-case average detection delay for the classical CUSUM occurs for $\tau=0$ when its test statistic is initialized with $W_0^i=0$. Focusing on the average time to false alarm $\mathsf{E}_\infty[\mathsf{T}_i(A)]$, to establish the lower bound in (16) we observe that $\mathcal{T}^i\geq 1$ since we take at least one sample, consequently $\mathsf{E}_\infty[\mathcal{T}^i]\geq 1$. Furthermore, from the estimates provided by Wald [15] we know that $\alpha_i\leq e^{-A}$. Substituting in (37) proves the lower bound. The upper bound requires more work. From $\mathcal{T}^i\leq\tau_-^i$ we have $\mathsf{E}_\infty[\mathcal{T}^i]\leq\mathsf{E}_\infty[\tau_-^i]$ which, as we argued, is bounded. We now need a lower bound for α_i. We can write

$$\begin{aligned}\alpha_i&=\mathsf{P}_\infty(S_{\mathcal{T}^i}^i\geq A)=\mathsf{E}_0^i[e^{-S_{\mathcal{T}^i}^i}\mathbb{1}_{\{S_{\mathcal{T}^i}^i\geq A\}}]\\ &=e^{-A}\mathsf{E}_0^i[e^{-(S_{\mathcal{T}^i}^i-A)}|S_{\mathcal{T}^i}^i\geq A](1-\beta_i)\\ &\geq e^{-A}e^{-\mathsf{E}_0^i[S_{\mathcal{T}^i}^i-A|S_{\mathcal{T}^i}^i\geq A]}(1-\beta_i)\\ &=e^{-A}\exp\left\{-\frac{\mathsf{E}_0^i[(S_{\mathcal{T}^i}^i-A)\mathbb{1}_{\{S_{\mathcal{T}^i}^i\geq A\}}]}{1-\beta_i}\right\}(1-\beta_i),\end{aligned}\tag{38}$$

where we used Jensen's inequality. As mentioned the SPRT satisfies $\mathcal{T}^i=\min\{\tau_-^i,\tau_A^i\}$. From this equality we conclude that $(S_{\mathcal{T}^i}^i-A)\mathbb{1}_{\{S_{\mathcal{T}^i}^i\geq A\}}=(S_{\tau_A^i}^i-A)\mathbb{1}_{\{\tau_A^i<\tau_-\}}\leq(S_{\tau_A^i}^i-A)$ and therefore, using (35), we obtain $\mathsf{E}_0^i[(S_{\mathcal{T}^i}^i-A)\mathbb{1}_{\{S_{\mathcal{T}^i}^i\geq A\}}]\leq\mathsf{E}_0^i[S_{\tau_A^i}^i-A]\leq\frac{\mathsf{J}_0^i}{\mathsf{I}_0^i}$. Additionally, as we argued above, we have $1-\beta_i\geq\mathsf{P}_0^i(\tau_-^i=\infty)=\frac{1}{\mathsf{E}_0^i[\tau_+^i]}$. If we substitute in (38) we strengthen the inequality and we obtain a lower bound for α_i. Substituting this lower bound in (37) provides the desired upper bound

$$\mathsf{E}_\infty[\mathsf{T}_i(A)]\leq\left(\mathsf{E}_\infty[\tau_-^i]\mathsf{E}_0^i[\tau_+^i]e^{-\mathsf{E}_0^i[\tau_+^i]\frac{\mathsf{J}_0^i}{\mathsf{I}_0^i}}\right)e^A,\tag{39}$$

and therefore we can define $\mathcal{C}_i=\mathsf{E}_\infty[\tau_-^i]\mathsf{E}_0^i[\tau_+^i]e^{-\mathsf{E}_0^i[\tau_+^i]\frac{\mathsf{J}_0^i}{\mathsf{I}_0^i}}$.

Consider now the worst-case average detection delay $D_i\big(\mathsf{T}_i(A)\big)$. To find the upper bound in (17), we combine Corollary 1 from [18] with (36) and obtain

$$D_i\big(\mathsf{T}_i(A)\big)=\mathsf{E}_0^i[\mathsf{T}_i(A)]=\frac{\mathsf{E}_0^i[\mathcal{T}^i]}{1-\beta_i}\leq\frac{A}{\mathsf{I}_0^i}+\mathsf{E}_0^i[\tau_+^i]\frac{\mathsf{J}_0^i}{\mathsf{I}_0^i},\tag{40}$$

suggesting that $\mathcal{U}_i=\mathsf{E}_0^i[\tau_+^i]\frac{\mathsf{J}_0^i}{\mathsf{I}_0^i}$, see also [10] and [16]. To find an equivalent lower bound we use Wald's lower bound for the expectation of the SPRT [15]

$$\mathsf{E}_0^i[\mathcal{T}^i]\geq\frac{1}{\mathsf{I}_0^i}\left\{(1-\beta_i)\log\frac{1-\beta_i}{\alpha_i}+\beta_i\log\frac{\beta_i}{1-\alpha_i}\right\},\tag{41}$$

from which we conclude

$$\begin{aligned}D_i\big(\mathsf{T}_i(A)\big)&\geq\frac{1}{\mathsf{I}_0^i}\left\{\log\frac{1-\beta_i}{\alpha_i}+\frac{\beta_i}{1-\beta_i}\log\frac{\beta_i}{1-\alpha_i}\right\}\\ &\geq\frac{1}{\mathsf{I}_0^i}\left\{\log\frac{1}{\alpha_i}+\log(1-\beta_i)+\frac{\beta_i}{1-\beta_i}\log\beta_i\right\},\end{aligned}\tag{42}$$

where the last inequality resulted by removing the nonnegative term $\frac{\beta_i}{1-\beta_i}\log\frac{1}{1-\alpha_i}$. Recall that $\alpha_i\leq e^{-A}$ and $\beta_i\leq 1-\frac{1}{\mathsf{E}_0^i[\tau_+^i]}$, furthermore by taking the derivative we can prove that the function $\log(1-\beta)+\frac{\beta}{1-\beta}\log\beta$ is decreasing in $\beta\in[0,1]$, consequently

$$D_i\big(\mathsf{T}_i(A)\big)\geq\frac{A}{\mathsf{I}_0^i}-\mathcal{L}_i\tag{43}$$

with $\mathcal{L}_i = \frac{1}{\mathsf{I}_0^i}\{\log(\mathsf{E}_0^i[\tau_+^i]) + (\mathsf{E}_0^i[\tau_+^i]-1)\log(\frac{\mathsf{E}_0^i[\tau_+^i]}{1-\mathsf{E}_0^i[\tau_+^i]})\}$ which is clearly nonnegative since $\tau_+^i \geq 1$. This concludes the proof of the lemma. ∎

PROOF OF THEOREM 1: Let us start by considering the lower bound for the average false alarm period in (18). As we mentioned, we intend to to analyze the proposed scheme $\mathsf{T}_{\rm MSP}$ by using (21). The main challenge comes from the fact that these SPRTs require different time steps. However, due to the periodic nature of their corresponding distributions it is possible to come up with interesting formulas. In particular recalling that with $\mathsf{P}_\infty(\cdot), \mathsf{E}_\infty[\cdot]$ we denote the probability measure and the corresponding expectation under the pre-change regime then, using (21) we can write

$$\begin{aligned}\mathsf{E}_\infty[\mathsf{T}_{\rm MSP}] &= \mathsf{E}_\infty\left[\sum_{\ell=1}^{k}\mathcal{T}_\ell\right]\\ &= \sum_{\ell=1}^{\infty}\mathsf{E}_\infty[\mathcal{T}_\ell]\mathsf{P}_\infty(k\geq\ell)\\ &= \varOmega_\infty + \omega_\infty\varOmega_\infty + \omega_\infty^2\varOmega_\infty + \cdots = \frac{\varOmega_\infty}{1-\omega_\infty}\end{aligned} \tag{44}$$

where, as in Lemma 1, $\alpha_i = \mathsf{P}_\infty(S_{\mathcal{T}^i} \geq A)$ denotes the Type-I error probability of the SPRT $\mathcal{T}^i$ when applied onto the ith process as in (20) and where

$$\begin{aligned}\omega_\infty &= \mathsf{P}_\infty(k > M) = \prod_{m=1}^{M}(1-\alpha_m),\\ \varOmega_\infty &= \sum_{j=1}^{M}\mathsf{E}_\infty[\mathcal{T}^i]\prod_{m=1}^{j-1}(1-\alpha_m),\end{aligned} \tag{45}$$

with $\prod_a^b = 1$ when $a > b$. The formulas in (44) and (45) are a consequence of the fact that the delays $\mathcal{T}_\ell,\ \ell = 1, 2, \ldots$ of the SPRTs comprising $\mathsf{T}_{\rm MSP}$ in (21), have the same distribution as the $\mathcal{T}^i,\ i = 1, \ldots, M$ defined in (20) with the correspondence between the two sets of SPRTs being periodic with period M and of the form $\ell \to i = (\ell-1) \bmod M + 1$. Finally, we must mention that in (44) the most crucial point is the fourth equality which is true because $\mathcal{T}_\ell$ and $\mathbb{1}_{\{k\geq\ell\}} = \mathbb{1}_{\{k>\ell-1\}}$ are functions of *non-overlapping* data (the event $\{k > \ell - 1\}$ depends on data used by $\mathcal{T}_1, \ldots, \mathcal{T}_{\ell-1}$) therefore, due to independence accros space and time, and the fact that each $\mathcal{T}_\ell$ has no memory of past data we conclude that $\mathcal{T}_\ell$ and $\mathbb{1}_{\{k\geq\ell\}}$ are independent. As in Lemma 1 we note that $\mathcal{T}_\ell \geq 1$ and $\alpha_i \leq e^{-A}$, therefore (45) implies

$$\begin{aligned}\omega_\infty &\geq (1-e^{-A})^M,\\ \varOmega_\infty &\geq \sum_{j=1}^{M}(1-e^{-A})^{j-1}\\ &= \Big(1-(1-e^{-A})^M\Big)e^A.\end{aligned} \tag{46}$$

Using these lower bounds in (44) yields (18) and proves the desired false alarm estimate. Consider now the second estimate depicted in (19) for the detection delay. As we mentioned in Section V, the test statistic we are employing is CUSUM-like, therefore the worst-case scenario occurs when the change is at time $\tau = 0$. Furthermore we experience the longest initial delay in our scheme when the change occurs at Process M, since then our test has to first go through all the other processes before sampling the one that has changed. We recall that $\mathsf{P}_0^M(\cdot)$ denotes the probability measure induced by the change occurring at the Mth process at time $\tau = 0$ and $\mathsf{E}_0^M[\cdot]$ the corresponding expectation. It is then clear that all processes for $i = 1, \ldots, M-1$ are under the nominal regime while only the last process is under the alternative. Similarly to the previous case we can write

$$\begin{aligned}D_M(\mathsf{T}_{\rm MSP}) &= \mathsf{E}_0^M[\mathsf{T}_{\rm MSP}]\\ &= \varOmega_M + \omega_M\varOmega_M + \omega_M^2\varOmega_M + \cdots = \frac{\varOmega_M}{1-\omega_M}\end{aligned} \tag{47}$$

where

$$\begin{aligned}\omega_M &= \mathsf{P}_0^M(k > M) = \beta_M\prod_{m=1}^{M-1}(1-\alpha_m)\\ \varOmega_M &= \mathsf{E}_0^M[\mathcal{T}^M]\prod_{m=1}^{M-1}(1-\alpha_m) + \sum_{j=1}^{M-1}\mathsf{E}_\infty[\mathcal{T}^j]\prod_{m=1}^{j-1}(1-\alpha_m),\end{aligned} \tag{48}$$

where, as in Lemma 1, $\beta_i = \mathsf{P}_0^i(S_{\mathcal{T}^i} \leq 0)$ denotes the Type-II error probability. Since $0 \leq \alpha_i, \beta_i \leq 1$, using (47) and (48) we have the following upper bound

$$D_M(\mathsf{T}_{\rm MSP}) \leq \frac{\mathsf{E}_0^M[\mathcal{T}^M]}{1-\beta_M} + \frac{1}{1-\beta_M}\sum_{j=1}^{M-1}\mathsf{E}_\infty[\mathcal{T}^j]. \tag{49}$$

In (49), in the first term of the right hand side, we recognize the detection delay of a CUSUM with threshold A applied solely to the Mth process. For this quantity we have thee upper bound from (17) of Lemma 1 yielding

$$D_M(\mathsf{T}_{\rm MSP}) \leq \frac{A}{\mathsf{I}_0^M} + \mathcal{U}_M + \frac{1}{1-\beta_M}\sum_{j=1}^{M-1}\mathsf{E}_\infty[\mathcal{T}^j]. \tag{50}$$

As pointed out in Lemma 1, $\mathcal{U}_M$ is a constant that depends only on f_M and g_M and not on A.

Since the change may occur at any process we can write similar estimates for any Process i, specifically

$$D_i(\mathsf{T}_{\rm MSP}) \leq \frac{A}{\mathsf{I}_0^i} + \mathcal{U}_i + \frac{1}{1-\beta_i}\sum_{j=1,j\neq i}^{M}\mathsf{E}_\infty[\mathcal{T}^j]. \tag{51}$$

To bound the $M-1$ terms in the sum in (51) we need to bound β_i and $\mathsf{E}_\infty[\mathcal{T}^i]$ which can be accomplished by employing the ladder variables defined in (32). Combining (51) and (36) we can write

$$D_i(\mathsf{T}_{\rm MSP}) \leq \frac{A}{\mathsf{I}_0^i} + \mathcal{U}_i + \mathsf{E}_0^i[\tau_+^i]\sum_{j=1,j\neq i}^{M}\mathsf{E}_\infty[\tau_-^j]. \tag{52}$$

Since, as we argued in the proof of Lemma 1, both averages $\mathsf{E}_0^i[\tau_+^i], \mathsf{E}_\infty[\tau_-^i]$ are finite under (A1), (A2), if we define

$$\mathcal{D} = \left(\max_{1\leq i\leq M}\mathsf{E}_0^i[\tau_+^i]\right)\left(\max_{1\leq i\leq M}\mathsf{E}_\infty[\tau_-^i]\right) \tag{53}$$

and use it to strengthen the inequality in (52) we obtain (19). This completes the proof. ∎

ACKNOWLEDGMENT

The authors are grateful to the Associate Editor and two anonymous reviewers for the detailed and constructive comments that greatly improve the quality and presentation of the article.

REFERENCES

[1] A. Tartakovsky, I. Nikiforov and M. Basseville, *Sequential Analysis: Hypothesis Testing and Changepoint Detection*, Chapman and Hall/CRC, 2014.

[2] K. Liu, Y. Mei and J. Shi, "An adaptive sampling strategy for online high-dimensional process monitoring," *Technometrics*, vol. 57, no. 3, pp. 305–319, 2015.

[3] T. Banerjee and V. V. Veeravalli,"Data-efficient quickest change detection in minimax settings," *IEEE Transactions on Information Theory*, vol. 59, no. 10, pp. 6917-6931, 2013.

[4] H. Chernoff, "Sequential design of experiments," *The Annals of Mathematical Statistics* , vol. 30, no. 3, pp. 755–770, 1959.

[5] L. Lai, H. V. Poor, Y. Xin and G. Georgiadis, "Quickest search over multiple sequences," *IEEE Transactions on Information Theory*, vol. 57, no. 8, pp. 5375-5386, 2011.

[6] A. Tsopelakos, G. Fellouris, and V. V. Veeravalli, "Sequential anomaly detection with observation control," 2019 IEEE International Symposium on Information Theory (ISIT), 2019.

[7] T. Lai and H. Robbins, "Asymtptocally efficient adaptive allocation rules," *Advances in Applied Mathematics*, vol. 6, pp. 4-22, 1985.

[8] A. N. Shiryaev, "On the Detection of Disorder in a Manufacturing Process. I," *Theory of Probability and its Applications*, vol. 8, no. 3, pp. 247-265, 1963.

[9] A. N. Shiryaev, "Quickest detection problems: Fifty years later," *Sequential Analysis*, vol. 29, no. 4, pp. 345-385, 2010.

[10] G. Lorden, "Procedures for reacting to a change in distribution," *The Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1897–1908, 1971.

[11] A. Tartakovsky and G. V. Moustakides, "State-of-the-art in Bayesian changepoint detection," *Sequential Analysis*, vol. 29: pp. 125-145, 2010.

[12] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol.41, pp. 100-115, 1954.

[13] G. V. Moustakides, "Optimal stopping times for detecting changes in distributions," *The Annals of Statistics*, vol. 14, no. 4, pp. 1379–1387, 1986.

[14] D. Siegmund, *Sequential analysis: Tests and confidence intervals*, Springer, 1985.

[15] A. Wald, "Sequential tests of statistical hypotheses," *Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117-186, 1945.

[16] Y. Mei, "Sequential change-point detection when unknown parameters are present in the pre-change distribution," *The Annals of Statistics*, vol. 34, no. 1, pp. 92–122, 2006.

[17] J. Kiefer and J. Sack, "Asymptotically optimum sequential inference and design," *The Annals of Mathematical Statistics*, vol. 34, pp. 705–750, 1963.

[18] G. Lorden, "On excess over the boundary," *The Annals of Mathematical Statistics*, vol. 41, no. 2, pp. 520-527, 1970.

Qunzhi Xu received the B.S. degree in Statistics from University of Science and Technology of China, Hefei, P. R. China in 2019. Currently he is a Ph.D. candidate with specialization in Statistics in The H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology, Atlanta, GA. His research interests include sequential change-point detection, multi-armed bandit problems and machine learning.

Yajun Mei received the B.S. degree in Mathematics from Peking University, Beijing, P. R. China in 1996, and the Ph. D. degree in Mathematics with a minor in Electrical Engineering in 2003 from the California Institute of Technology, Pasadena, CA, USA. He is currently a Professor in The H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology, Atlanta, GA. His research interests include streaming data analysis, quickest change detection, sensor networks, sequential analysis, and their applications in engineering and biomedical sciences.

George V. Moustakides received the diploma in Electrical and Mechanical Engineering from the National Technical University of Athens, Greece in 1979, the MSE in Systems Engineering from the University of Pennsylvania in 1980, and the M.Sc and Ph.D in Electrical Engineering and Computer Science from Princeton University in 1983. Since 1988 he is with the University of Patras, Greece initially with the Computer Engineering and Informatics department and afterwards with the department of Electrical and Computer Engineering. In 2017 he also joined the Computer Science department at Rutgers University. In the past he held various visiting or long- term appointments with INRIA, Princeton University, University of Pennsylvania, Columbia University, University of Maryland, Georgia Institute of Technology, University of Southern California, University of Illinois at Urbana-Champaign and Rutgers University. His interests include Sequential Detection, Statistical Signal Processing and Machine Learning. From 2011 to 2014 and from 2016 to 2018 he served as Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY.