Multi-Stream Quickest Detection with Unknown Post-Change Parameters Under Sampling Control

Qunzhi Xu
Georgia Institute of Technology
xuqunzhi@gatech.edu

Yajun Mei
Georgia Institute of Technology
ymei3@gatech.edu

***Abstract*—The multi-stream quickest detection problem with unknown post-change parameters is studied under the sampling control constraint, where there are M local processes in a system but one is only able to take observations from one of these M local processes at each time instant. The objective is to raise a correct alarm as quickly as possible once the change occurs subject to both false alarm and sampling control constraints. We propose an efficient myopic-sampling-based quickest detection algorithm under sampling control constraint, and show it is asymptotically optimal in the sense of minimizing the detection delay under our context when the number M of processes is fixed. Simulation studies are conducted to validate our theoretical results.**

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

I. Introduction

The quickest detection problem has a wide range of real-world applications ranging from quality control in manufacturing to environmental monitoring to biosurveillance, see Tartakovsky, Nikiforov, and Basseville [1]. When one monitors observations from $M = 1$ single local process whose distribution might change at some unknown time τ, this is the classical quickest detection problem and has been studied extensively in the past fifty years. In recent years, research of monitoring $M \geq 2$ local processes receives extensive attentions, e.g., Mei [2], Xie and Siegmund [3], and Chan [4].

In this paper, we investigate the quickest detection problem when monitoring $M \geq 2$ local processes, but with a new twist of the sampling control constraint where we are able to observe only one of these M local processes at each time instant. To be more concrete, under our setup, there are M possible local processes whose pre-change distributions are known. At some unknown time τ, the distribution of exactly one of these M local processes changes to an unknown post-change distribution which is assumed to belong to a general family of class of distributions. However, due to the communication or resource constraints, we are only able to take observations from exactly one of these M processes at each time instant. Our objective is how to take observations adaptively from these M local processes and how to use the observed data to raise a correct alarm as quickly as possible once the change occurs subject to both the false alarm and sampling control constraints.

This problem was first investigated in Shiryaev [5] which was motivated from signal/target detection when the radar sensor can only monitor one out of M directions at a time. Recently it was re-visited by Xu, Mei and Moustakides [6], [7] for asymptotic minimax optimality theories, and also extended to Gauss-Markov random fields by Heydari, Tajer and Poor [8]. All these research assume that the post-change distribution is completely specified, and here we assume that the post-change distribution includes unknown parameters. Meanwhile, when the sampling control is over the time domain, an efficient-data-sampling quickest detection technique is developed in Banerjee and Veeravalli [9] when monitoring a single local process. When the objective is to test hypothesis or optimize regrets, extensive research has been done in the literature of sequential hypothesis testing and multi-armed bandit problems, and many efficient sampling control algorithms have been developed. For a limited list of references, see Chernoff [10], Tsopelakos, Fellouris, and Veeravalli [11], Lai and Robbins [12], Nitinawarat, Atia and Veeravalli [13], Naghshvar and Javidi [14].

The main contribution of this paper is to propose an efficient myopic-sampling-based quickest detection algorithm under the sampling control constraint, and show that our proposed algorithm is asymptotically optimal when the number M of processes is fixed and the average run length to false alarm constraint γ goes to ∞. Our main idea is to explore each local process extensively to decide whether or not there is a local change, and then switch to new processes only when we are confident that the existing process does not involve any local change. Even with the sampling rate of $1/M$ at each time instant, our proposed algorithm is shown to have the same first-order detection delay performance as the oracle CUSUM procedure when the dimension $M << \log \gamma$ as the average run length to false-alarm constraint γ goes to ∞, but its performance degraded noticeably when the dimension M becomes larger. It remains an open problem to develop asymptotically optimal theories for high-dimensional streams under the sampling control.

The remainder of the paper is organized as follows. In Section II, we state the mathematical formulation of our problem and review some existing methods. In Section III, we present our proposed algorithm and provide its theoretical properties. Numerical simulations are presented in Section IV to illustrate the performance properties of our proposed algorithm. The technical proofs are presented in the Appendix.

II. Problem Formulation and Methods Review

In the multi-stream quickest detection problem, suppose there are M statistically independent lcoal processes in a system, and denote with X_t^i the observation from the i-th process at time t, where $i = 1, \ldots, M$ and $t = 1, 2, \ldots$. Let $f_\theta(x)$ be the density function of the normal distribution $N(\theta, \sigma^2)$ with mean θ and known variance σ^2. Initially, the system is in the in-control state and the data stream $\{X_t^i\}$ from the i-th process produces i.i.d. samples following the density $f_{\theta_i}(X)$ with known mean θ_i. At some unknown time τ, a triggering event occurs to the system and one of its M processes, say, the i-th, changes to i.i.d. samples following a new unknown density $f_{\phi_i}(X)$, i.e., with a new unknown mean ϕ_i. Specifically, if the i-th data stream is affected, then

$$X_t^i \sim \begin{cases} f_{\theta_i}(X), & \text{if } t \leq \tau \\ f_{\phi_i}(X), & \text{if } t > \tau, \end{cases} \tag{1}$$

whereas $X_t^j \sim f_{\theta_j}(X)$ for $j \neq i$ and all $t > 0$. Here we consider in detail the one-sided change-point problem where it is assumed that $\phi_i > \theta_i$ for all $i = 1, \cdots, M$.

Let us now discuss the sampling control constraints. To be mathematically rigorous, define a sequence of *sampling indices* $\{R_t\}$ with $R_t \in \{1, \ldots, M\}$, where R_t is a random variable and $\{R_t = m\}$ means that we will sample the m-th local process at the time instant t. Under our sampling constraint, we are allowed to access only one of these M local process at each time t, and this can be expressed as

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

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

Thus in the multi-stream quickest detection problem under sampling control, an algorithm consists two elements: one is the sampling policies, e.g., choose $\{R_t\}$ for all time instants t subject to (2), and the other is the decision policy that is defined as the stopping time T with respect to the observed data sequence $\{X_t^{i=R_t}\}_{t\geq 1}$. Note that the sampling decision R_t depends only on those observed data up to time $t-1$, and the stopping time $\{\mathsf{T} = t\}$ means that we raise an alarm at time t.

Following the classical minimax formulation for quickest detection proposed by Lorden [15], we are interested in finding a procedure $(\{R_t\}_{t=1,\cdots,\infty}, \mathsf{T})$ that minimizes the worst-case detection delay

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

for any $i = 1, 2, \cdots, M$ when the i-th process is affected by the change, subject to the average run length to false-alarm constraint

$$\mathsf{E}_\infty[\mathsf{T}] \geq \gamma > 1. \tag{4}$$

Here $\mathsf{P}_t^i(\cdot), \mathsf{E}_t^i[\cdot]$ denote the probability measure and the corresponding expectation induced by the change occuring at the i-th process at time $\tau = t$ and $\mathsf{P}_\infty(\cdot), \mathsf{E}_\infty[\cdot]$ denote the probability measure and the corresponding expectation induced by the change occuring at ∞. In addition, $\mathscr{F}_t^i$ is the σ-algebra generated by observed data/information on the i-th process up to time t.

Let us now review some existing research that is related to our problem and proposed algorithm. Under the simplest scenario where we had the prior knowledge on the index of i where change occurs to and the full information of post-change distribution f_{ϕ_i}, it is natural to always sample the i-th process, i.e., $R_t \equiv i$ for all t, and the corresponding optimal procedure will be the well-known CUSUM procedure:

$$\mathsf{T}_{\text{oracle}}(A) = \inf\{t > 0 : W_t^i \geq A\} \tag{5}$$

where W_t^i is the CUSUM statistics recursively defined as

$$W_t^i = \max\{W_t^{i-1}, 0\} + \log \frac{f_{\phi_i}(X_t^i)}{f_{\theta_i}(X_t^i)} \quad \text{for } t \geq 1 \tag{6}$$

and $W_0^i = 0$, see Moustakides [16]. Here we use $\mathsf{T}_{\text{oracle}}$ to emphasize that this CUSUM procedure makes an oracle assumption of known affected local process and known post-change distribution.

Note that it is highly non-trivial to develop an efficient algorithm under our setup due to two challenges: the first is that the post-change distributions are unknown. This challenge has been addressed in Lorden and Pollak [17] when monitoring $M = 1$ local process. Their idea is to estimate the post-change parameter by the average of recent observations after the candidate change-point, and update the local statistics as in the classical CUSUM statistic in (6). To be more specific, when monitoring the i-th local process, the estimator $\hat{\theta}_{t,i}$ for the i-th process at time instant t can be defined as

$$\hat{\theta}_{t,i} = \max(\theta_{l,i}, \frac{\sum_{n=g(t)}^{t-1} X_n^i}{t - g(t)}), \tag{7}$$

where $\theta_{l,i}$ is the lower bound of the post-change parameter and $g(t) \geq 0$ is the candidate change-point that is defined as the last time up to time t when local CUSUM-type statistics are zero. The local statistics $\widetilde{W}_t^i$ can then be defined as in the recursion (6) with $\phi_i = \hat{\theta}_t^i$, which yields to Lorden and Pollak's procedure

$$\mathsf{T}_{\text{LP}}(A) = \inf\{t > 0 : \widetilde{W}_t^i \geq A\}. \tag{8}$$

The second challenge is that the index of i of affected local process is unknown, and thus it is unclear how to choose sampling indices $\{R_t\}$ suitably so as to detect the change quickly. A naive sampling idea is to sample each local process periodically, i.e., $R_t = t \bmod M + 1$ for all time instants $t = 1, 2, \cdots$, and each process is visited only once during each M time instants. Combing this naive sampling policy with Lorden and Pollak's procedure in (8) yields the following quickest detection algorithm:

$$\mathsf{T}_{\text{naive}}(A) = \inf\{t > 0 : \max\{\widetilde{W}_t^1, \cdots, \widetilde{W}_t^M\} > A\} \tag{9}$$

where $\widetilde{W}_t^i (i = 1, \cdots, M)$ is only updated when $R_t = i$. In the sequel we will refer (9) as the naive algorithm with the naive periodic sampling policy.

Clearly, the naive algorithm in (9) seems to be inefficient, as it might spend too much time on those $M-1$ unaffected processes. To the best of our knowledge, no efficient algorithms have been developed in the quickest detection literature to simultaneously address these two challenges of unknown post-change distribution and unknown index of affected local process.

III. Our Proposed Algorithm

In this section we present an efficient quickest detection algorithm under the sampling control constraint (2). Our key idea is based on a myopic sampling policy where we continue to sample a process until we are confident to make one of the following two decisions: either a change has occurred or no changes have occurred. For better presentation, we split this section into two subsections. In Section III-A, we propose our algorithm including both the myopic sampling policy and the stopping time $\mathsf{T}_{\mathrm{MSP}}$. In Section III-B we investigate the theoretical properties of our proposed algorithm $\mathsf{T}_{\mathrm{MSP}}$.

A. Algorithm Development

There are three essential components in our proposed algorithm: (i) how to construct and update local statistics for all local processes including the process not being sampled; (ii) how to decide which local process to be sampled based on these local statistics, and (iii) when to raise a global alarm.

At the high level, our proposed algorithm exploits the prior knowledge that there is only one process affected by the change, and adopts the myopic sampling policy that samples the local process with the maximum local statistics. By choosing local statistics as those in Lorden and Pollak's procedure (8), the myopic sampling policy implies that we sample each process until we are confident to decide whether a change has occurred or not. If there is a local change, then we stop and raise a global alarm. If there are no local changes, then we switch to sample the next local process. We repeat these steps until we raise a global alarm.

Let us now define our proposed algorithm under the sampling control constraint (2). We start with the construction and update of local statistics $\widetilde{W}_t^i$. There are two cases, depending on whether a local process is sampled or not. If a local process is not sampled, then we update $\widetilde{W}_t^i$ as $\max\{\widetilde{W}_{t-1}^i, 0\}$. If a local process is sampled, then we update $\widetilde{W}_t^i$ as in Lorden and Pollak's procedure (8). To be more specific, let $\hat{\theta}_{t,i}$ be the estimate of the post-change parameter for the i-th process at time t, which will be defined in a little bit later, and the local statistics $\widetilde{W}_t^i$ can be defined recursively as

$$\begin{aligned} \widetilde{W}_t^i &= \max\{\widetilde{W}_{t-1}^i, 0\} + \mathbb{1}_{\{i=R_t\}} \log \frac{f_{\hat{\theta}_{t,i}}(X_t^i)}{f_{\theta_i}(X_t^i)} \\ &= \begin{cases} \max\{\widetilde{W}_{t-1}^i, 0\}, & \text{if } i \neq R_t \\ \max\{\widetilde{W}_{t-1}^i, 0\} + \log \frac{f_{\hat{\theta}_{t,i}}(X_t^i)}{f_{\theta_i}(X_t^i)}, & \text{if } i = R_t, \end{cases} \end{aligned} \tag{10}$$

with the initial values $\widetilde{W}_0^i = 0$ for all $i = 1, \cdots, M$.

Our proposed stopping time $\mathsf{T}_{\mathrm{MSP}}$ is then defined as

$$\mathsf{T}_{\mathrm{MSP}}(A) = \inf\left\{ t > 0 : \max_{1 \leq i \leq M} \widetilde{W}_t^i \geq A \right\}, \tag{11}$$

for some pre-specified constant A.

As for the post-change parameter estimators $\theta_{t,i}$, by (10), we only need to pay attention to the sampled local process and thus adopt the same idea as in Lorden and Pollak's procedure (8). To be more concrete, at time instant t, assume that we sample at the i-th process, and denote by $M(t)$ the total time instants in which we have consecutively sampled at the i-th process. In other words, we observed data from the i-th process during the time period of $t-M(t)+1$ to t. Here we propose to estimate the post-change parameter based on the observed data from the i-th process during the time period of $t-M(t)+1$ to $t-1$, as we save the data at the time instant t for quickest detection, not for parameter estimation. Mathematically, we can define the estimator $\hat{\theta}_{t,i}$ based on the method of moments (MOM) estimator of the distribution in (1):

$$\hat{\theta}_{t,i} = \max\left\{ \theta_{l,i}, \frac{\sum_{j=t-M(t)+1}^{t-1} X_j^i}{M(t)-1} \right\} \tag{12}$$

with 0/0=$-\infty$. Here $\theta_{l,i}$ is the lower bound of the post-change parameter.

It remains to define the sampling policies $\{R_t\}$. At the high level, the sampling policy $\{R_t\}$ at time instant t can be defined by the local statistics $\widetilde{W}_{t-1}^i$ at time instant $t-1$, which in turn depends on the sampling policy $\{R_{t-1}\}$ at time instant $t-1$. Here our proposed algorithm keeps sampling on Process i as long as the local statistics $\widetilde{W}_t^i > 0$ and switches to the next process when $\widetilde{W}_t^i \leq 0$. Mathematically, in our proposed algorithm, the sampling policy $\{R_t\}$ at time instant t can be defined by the local statistics $\widetilde{W}_{t-1}^i$ at time instant $t-1$:

$$R_t = \begin{cases} R_{t-1} & \text{if } \widetilde{W}_{t-1}^{R_{t-1}} > 0 \\ R_{t-1} \bmod M + 1 & \text{if } \widetilde{W}_{t-1}^{R_{t-1}} \leq 0. \end{cases} \tag{13}$$

with initial values $R_0 = 1$.

In summary, our proposed algorithm can be summarized as follows:

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

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

$\vdots$

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

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

B. Theoretical Properties

In this subsection, we will investigate the theoretical properties of our proposed algorithm $\mathsf{T}_{\mathrm{MSP}}$ in (11). First, we need to make necessary assumptions. As in Lorden and Pollak [17], we assume that the post-change parameter ϕ_i and the pre-change parameter θ_i is separable by the cutoff value $\theta_{l,i}(i = 1, \cdots, M)$, i.e., $\theta_i < \theta_{l,i} < \phi_i$. for all $i = 1, \cdots, M$. Moreover, we assume that Kullback-Leibler information numbers are positive and finite for all $i = 1, \cdots, M$:

$$\begin{aligned} I(\theta_i, \phi_i) &= \int \log \frac{f_{\theta_i}(X)}{f_{\phi_i}(X)} f_{\theta_i}(X) dX \in (0, \infty), \\ I(\phi_i, \theta_i) &= \int \log \frac{f_{\phi_i}(X)}{f_{\theta_i}(X)} f_{\phi_i}(X) dX \in (0, \infty). \quad (14) \end{aligned}$$

For normal distributions, both Kullback-Leibler information numbers are the same and become $(\theta_i - \phi_i)^2/(2\sigma^2)$. Here we adopt general notations, as we feel our results below hold for more general distributions.

Now we are ready to present the theoretical properties of our proposed algorithm T_{MSP} in (11). The main results are summarized in the following theorem and its corollary, whose high-level proofs are presented in the appendix.

Theorem 1 *For our proposed algorithm* T_{MSP} *in (11), we have*

$$\mathsf{E}_\infty[\mathsf{T}_{\text{MSP}}] \geq e^A. \quad (15)$$

Moreover, its detection delay satisfies

$$D_i(\mathsf{T}_{\text{MSP}}) \leq \frac{A}{I(\phi_i, \theta_i)} + C_0 \log A + C_1(M-1) \quad (16)$$

as $A \to \infty$ *for any* $i \in 1, \cdots, M$. *Here* C_0, C_1 *are constants depending only on the distributions, not on* A *and* M..

Corollary 1 *Let* $A = \log \gamma$, *then our proposed algorithm* $\mathsf{T}_{\text{MSP}}(A)$ *in (11) satisfies both the false alarm constraint in (4) and the sampling control constraint in (2). Moreover, for each* $i = 1, \ldots, M$, *its detection delay satisfies*

$$0 \leq D_i(\mathsf{T}_{\text{MSP}}) - D_i^{\text{orc}} \leq C_0 \log \log \gamma + C_2 M \quad (17)$$

where D_i^{orc} *is the oracle detection delay achieved by the classical CUSUM procedure for monitoring a change in distribution of the* i*-th process subject to the false alarm constraint in (4):*

$$D_i^{\text{orc}} = \frac{\log \gamma}{I(\phi_i, \theta_i)} + C_3 \quad (18)$$

and the parameters C_2, *and* C_3 *are constants depending only on the distributions, not on* γ *and* M.

It is useful to add some remarks. Note that relationship (17) holds for every M and γ. On one hand, our proposed algorithm T_{MSP} has the same detection delay of the oracle or CUSUM procedure up to $O(\log \log \gamma)$ when M is fixed as $\gamma \to \infty$, or when $M = O(\log \log \gamma)$. On the other hand when M is large but γ is moderately large, the additional term $C_0 \log \log \gamma + C_2 M$ can be comparable to or even larger than D_i^{orc}, and thus the performance of our proposed algorithm will be much worse than the oracle or CUSUM procedure. This is not surprising for high-dimensional setting, as the sampling control in (2) is too restrictive for large M and we should not be able to detect the change quickly if we only sample one out of M processes at each time instant. In other contexts, we can evaluate the constants C_0, C_1 and C_2 to see the effects of the dimension M on the performance of our proposed algorithm, also see Wang and Mei [18] for similar contexts. It remains an open problem to develop a general asymptotic optimality theory for high-dimensional streams under the sampling control.

TABLE I
Comparison of Detection Delay of $\mathsf{T}_{\text{naive}}$ and T_{MSP}

$\gamma = 50000$		$M = 2$		$M = 10$	
μ	Oracle	Naive	T_{MSP}	Naive	T_{MSP}
0.5	61.87	144.01	90.56	701.23	234.10
0.75	29.62	64.13	39.07	308.52	100.06
1.0	17.20	36.45	22.65	174.67	60.85
1.25	11.35	23.40	15.46	112.12	43.33
1.5	7.93	16.60	11.21	80.28	35.03

IV. Simulation

In this section, we conduct Monte Carlo simulations to demonstrate the performance properties of our proposed algorithm T_{MSP} in (11). In our simulation, we consider two choices on the number M of processes: $M = 2$ or $M = 10$. For each choice of M processes, we consider the mean shift in normal distribution from $N(0, 1)$ to $N(\mu, 1)$ with $\mu \geq 0.5$. Due to the page constraints, here we only present the homogeneous setting (i.e., $f_{\theta_i} \equiv f$ and $f_{\phi_i} \equiv g$ for all $i = 1, \cdots, M$).

In each case, we set the false alarm constraint $\gamma = 50,000$. For our proposed algorithm $\mathsf{T}_{\text{MSP}}(A)$ and the naive method $\mathsf{T}_{\text{naive}}(A)$ in (9), we first use the bisection method to find suitable threshold A to attain the false alarm constraint, and then simulate the worst-case detection delay under different post-change scenarios where the change occurs to the M-th process (as our algorithm start to sample at the first process).

Tables I report the detection delay of our proposed algorithm T_{MSP} and the naive method $\mathsf{T}_{\text{naive}}$ in (9). In addition, we also report the oracle detection delay of the CUSUM procedure in (5). All numerical results are based on $50,000$ Monte Carlo runs. From the tables, it is clear that our proposed algorithm T_{MSP} is much better than the naive method T_{base} and can reduce the detection delay by at least 25% when $M = 2$ and 50% when $M = 10$. In other words, as compared to the naive periodic sampling, our proposed myopic sampling policy can lead to a significantly improvement on the detection delay performance.

Moreover, our results also shows that as the dimension M increases from $M = 2$ to $M = 10$, the detection delays of both our proposed algorithm T_{MSP} and the naive method $\mathsf{T}_{\text{naive}}$ in (9) increase significantly. We conjecture that the oracle bound of the CUSUM procedure is unattainable for high-dimensional monitoring under the sampling control, but we are unable to provide a rigorous proof.

Appendix

In this Appendix, we only provide a high-level sketch of the proof for Theorem 1, as the proof of the corollary follows directly from Theorem 1 and the optimality property of the classical CUSUM procedure is in (5). Without lost of generality, we assume the variance $\sigma^2 = 1$.

The main idea in the proof is to present an equivalent definition of our proposed algorithm T_{MSP} by the sequential probability ratio tests (SPRTs) and their extensions. For each process $i (i = 1, \cdots, M)$, we define sequential tests $\mathcal{T}^i$:

$$\mathcal{T}^i = \inf \left\{ t > 0 : \sum_{j=1}^{t} \log \frac{f_{\hat{\theta}_{j,i}}(X_j^i)}{f_{\theta_i}(X_j^i)} \notin (0, A) \right\} \quad (19)$$

Then our proposed algorithm T_{MSP} can be written as the sum of k independent sequential tests, where k is the first time the sequential test being tested crosses the upper boundary A:

$$\mathsf{T}_{\text{MSP}} = \mathcal{T}_1 + \cdots + \mathcal{T}_k, \quad (20)$$

Due to the periodic nature of T_{MSP}, it is easy to verify that $\mathcal{T}_l$ follows the same distribution with $\mathcal{T}^i$ where $i = (l-1) \bmod M + 1$. The theoretical properties of T_{MSP} can then be analyzed by considering the corresponding sequential tests.

Due to the page limit, below we focus on the proof of detection delay relationship (16). Without loss of generality, we assume the change occurs to the M-th process at time $\tau = 0$, since our test needs to go through all other $M-1$ processes before sampling the one that has changed. Hence, it suffices to show that $D_M(\mathsf{T}_{\text{MSP}}) = \mathsf{E}_0^M[\mathsf{T}_{\text{MSP}}]$ satisfies relationship (16).

Based on the periodic nature of our proposed algorithm, by (20), we have

$$\begin{aligned} D_M[\mathsf{T}_{\text{MSP}}(A)] &= \mathsf{E}_0^M\left[\sum_{\ell=1}^{k} \mathcal{T}_\ell\right] \\ &= \sum_{\ell=1}^{\infty} \mathsf{E}_0^M[\mathcal{T}_\ell]\mathsf{P}_0^M(k \geq \ell) \\ &= \Omega_M + \omega_M \Omega_M + \omega_M^2 \Omega_M + \cdots \\ &= \frac{\Omega_M}{1 - \omega_M}, \end{aligned}$$

where

$$\begin{aligned} \omega_M &= \beta_M \prod_{i=1}^{M-1} (1 - \alpha_i) \\ \Omega_M &= \mathsf{E}_0^M[\mathcal{T}^1] + \mathsf{E}_0^M[\mathcal{T}^2](1-\alpha_1) + \cdots + \\ &\quad + \mathsf{E}_0^M[\mathcal{T}^M] \prod_{i=1}^{M-1} (1 - \alpha_i) \end{aligned}$$

where β_M denotes the type-II error probability of the sequential test $\mathcal{T}^M$ and α_i denotes the type-I error probability of the sequential test $\mathcal{T}^i (i = 1, \cdots, M-1)$. Since $0 \leq \alpha_i, \beta_i \leq 1$, we have the following upper bound

$$D_M(\mathsf{T}_{\text{MSP}}(A)) \leq \frac{\mathsf{E}_0^M[\mathcal{T}^M]}{1 - \beta_M} + \frac{1}{1 - \beta_M} \sum_{i=1}^{M-1} \mathsf{E}_\infty[\mathcal{T}^i]. \quad (21)$$

Since the M-th process is the only local process that is affected by the change and the remaining $M-1$ process that are unaffected, we need to investigate two terms on the right-hand sides of (21) separately. Let us first consider the easier one, which is related to the unaffected $M-1$ processes. A key step is to introduce the following "*ladder variables*" for $i = 1, \cdots, M$,

$$\tau_-^i = \inf \left\{ t > 0 : S_t^i = \sum_{j=1}^{t} \log \frac{f_{\hat{\theta}_{j,i}}(X_j^i)}{f_{\theta_i}(X_j^i)} \leq 0 \right\}. \quad (22)$$

Then the standard arguments on sequential tests shows that

$$\begin{aligned} 1 - \beta_M &\geq \mathsf{P}_0^M(\tau_-^M = \infty) > 0 \\ \mathsf{E}_\infty[\mathcal{T}^i] &\leq \mathsf{E}_\infty[\tau_-^i] < +\infty, \end{aligned}$$

which provide the upper bound on the detection delays of the second term on the right-hand sides of (21).

It is slightly more difficult to investigate the first term on the right-hand sides of (21), which is the detection delay of $\mathcal{T}^M$ for the M-th process: while its form is similar to the classical SPRT, it updates the (post-change) parameter recursively, and thus it cannot be written as the sum of i.i.d. under the probability measure P_0^M. Nevertheless, we can use the SPRT to derive its properties by comparing the estimated (post-change) parameter with the true post-change parameter.

To be more specific, for $\mathcal{T}^M$, by Wald's equation, we have

$$\begin{aligned} \mathsf{E}_0^M[\mathcal{T}^M] &= \frac{\mathsf{E}_0^M[\sum_{t=1}^{\mathcal{T}^M}(\phi_M - \theta_M)X_t^M - (\phi_M^2 - \theta_M^2)/2)]}{I(\phi_M, \theta_M)} \\ &= \frac{\mathsf{E}_0^M[S_{\mathcal{T}^M}^M]}{I(\phi_M, \theta_M)} + \frac{\mathsf{E}_0^M[\sum_{t=1}^{\mathcal{T}^M}(\phi_M - \hat{\theta}_{t,M})X_t^M]}{I(\phi_M, \theta_M)} \\ &\quad + \frac{\mathsf{E}_0^M[\sum_{t=1}^{\mathcal{T}^M}(\hat{\theta}_{t,M}^2 - \phi_M^2)/2]}{I(\phi_M, \theta_M)} \\ &= D_1 + D_2 + D_3 \end{aligned}$$

where $S_{\mathcal{T}^M}^M$ is the log-likelihood S_t^i in (19) with $i = M$ and $t = \mathcal{T}^M$.

It remains to investigate the terms D_1, D_2, D_3 in the above relationship. Since they are based on the sequential tests, a tedious but standard overshoot analysis shows that

$$\begin{aligned} D_1 &\leq \frac{(1 - \beta_M)(A + C_4)}{I(\phi_M, \theta_M)}, \\ D_2 + D_3 &\leq \frac{(1 - \beta_M)C_5 \log A}{I(\phi_M, \theta_M)} \end{aligned}$$

Combing the results above, it follows that

$$\begin{aligned} D_M(\mathsf{T}_{\text{MSP}}(A)) &\leq \frac{A + C_4}{I(\phi_M, \theta_M)} + \frac{C_5 \log A}{I(\phi_M, \theta_M)} \\ &\quad + \frac{\sum_{i=1}^{M-1} \mathsf{E}_\infty[\tau_-^i]}{\mathsf{P}_0^M(\tau_-^M = \infty)} \\ &\leq \frac{A}{I(\phi_M, \theta_M)} + C_0 \log A + C_1(M-1) \end{aligned}$$

where $C_0 = \frac{C_5}{I(\phi_M, \theta_M)}$ and $C_1 = \frac{\max_{i=1,\cdots,M-1} \mathsf{E}_\infty[\tau_-^i]}{\mathsf{P}_0^M(\tau_-^M = \infty)} + \frac{C_4}{I(\phi_M, \theta_M)}$. □

ACKNOWLEDGMENT

This work was supported in part by the US National Science Foundation under Grant DMS-1830344 and DMS-2015405 through Georgia Institute of Technology.

REFERENCES

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

[2] Y. Mei, “Efficient scalable schemes for monitoring a large number of data streams,” *Biometrika*, vol. 97, no. 2, pp. 419-433, 2010.

[3] Y. Xie and D. Siegmund, “Sequential multi-sensor change-point detection,” *The Annals of Statistics*, vol. 41, no. 2, pp. 670-692, 2013.

[4] H. P. Chan, “Optimal sequential detection in multi-stream data,” *The Annals of Statistics*, vol. 45, no. 6, pp. 2736-2763, 2017.

[5] 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.

[6] Q. Xu, Y. Mei and G. V. Moustakides, “Second-order asymptotically optimal change-point detection algorithm with sampling control,” 2020 IEEE International Symposium on Information Theory (ISIT), pp. 1136-1140, 2020.

[7] Q. Xu, Y. Mei and G. V. Moustakides, “Optimum multi-stream sequential change-point detection with sampling control,” *IEEE Transactions on Information Theory,* 2021 (Accepted). doi: 10.1109/TIT.2021.3074961.

[8] J. Heydari, A. Tajer, and H. V. Poor. “Quickest detection of Gauss-Markov random fields.” *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton) IEEE,* pp. 808-814, 2015.

[9] 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.

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

[11] A. Tsopelakos, G. Fellouris, and V. V. Veeravalli, “Sequential anomaly detection with observation control,” 2019 IEEE International Symposium on Information Theory (ISIT), pp. 2389-2393, 2019.

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

[13] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli. “Controlled sensing for multihypothesis testing.” *IEEE Transactions on Automatic Control,* vol. 58, no. 10, pp. 2451-2464, 2013.

[14] M. Naghshvar and T. Javidi, “Active sequential hypothesis testing,” *The Annals of Statistics,* vol. 41, no. 6, pp. 2703-2738, 2013.

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

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

[17] G. Lorden and M. Pollak, “Sequential change-point detection procedures that are nearly optimal and computationally simple, ” *Sequential Analysis*, vol. 27, no. 4, pp. 476–512, 2008.

[18] Y. Wang and Y. Mei, “Large-scale multi-stream quickest change detection via shrinkage post-change estimation,” *IEEE Transactions on Information Theory*, vol. 6, no. 12, pp. 6926-6938, 2015.

[19] G. Lorden and M. Pollak, “Nonanticipating estimation applied to sequential analysis and changepoint detection, ” *The Annals of Statistics*, vol. 33, no. 3, pp. 1422–1454, 2005.

[20] D. Siegmund, *Sequential Analysis: Tests and Confidence Intervals*, Springer, New York, 1985.