

Quickest Change Detection with Controlled Sensing

Georgios Fellouris, *Member, IEEE*, and Venugopal V. Veeravalli, *Fellow, IEEE*

Abstract—In the problem of quickest change detection, a change occurs at some unknown time in the distribution of a sequence of random vectors that are monitored in real time, and the goal is to detect this change as quickly as possible subject to a certain false alarm constraint. In this work we consider this problem in the presence of parametric uncertainty in the post-change regime and controlled sensing. That is, the post-change distribution contains unknown parameters, and the distribution of each observation, before and after the change, is affected by a control action. In this context, in addition to a stopping rule that determines the time at which it is declared that the change has occurred, one also needs to determine a sequential control policy, which chooses the control action at each time based on the already collected observations that is “best” for the unknown post-change parameter. We formulate this problem mathematically using Lorden’s minimax criterion, and assuming that there are finitely many possible actions and post-change parameter values. We establish a universal lower bound on the worst-case detection delay, as the mean time to false alarm goes to infinity, which needs to be satisfied by any procedure for quickest change detection with controlled sensing. We then propose a specific procedure for this problem, which we call the Chernoff-CuSum procedure, for which the conditional expected detection delay, for any fixed value of the change-point, matches the universal lower bound up to a first-order asymptotic approximation as the mean time to false alarm goes to infinity.

Index Terms—Sequential change detection; Experimental design; Observation control; CuSum test.

I. INTRODUCTION

The problem of detecting changes or anomalies in stochastic systems and time series, often referred to as the quickest change detection (QCD) problem, arises in various branches of science and engineering. The observations of the system are assumed to undergo a change in distribution at the change-point, and the goal is detect this change as soon as possible, subject to false alarm constraints. See [1], [2], [3], [4] for books and survey articles on the topic.

In this paper we study an interesting variant of the QCD problem that was introduced in recent work [5], to which the authors refer as the “bandit” QCD problem. In this setting, the distribution of the observations is not only affected by the change but also through a control (action) variable that is chosen by the observer. As described in [5], a canonical example application for such a setting is in surveillance systems in which sensors can be switched and steered (controlled) to look for targets in different locations in space, and only a

subset of locations can be probed at any given time. The policy for controlling the sensors has to be designed jointly with the change detection algorithm to provide the best tradeoff between detection delay and false alarm rate. A number of other applications contexts for the “bandit” QCD problem are described in [5].

On a fundamental level, the QCD problem in which the distribution of the observations is affected by control actions falls squarely within the larger context of sequential decision-making problems with observation control (or *controlled sensing* [6]). Such controlled sensing problems have a rich history going back to the seminal work of Chernoff [7] on the sequential design of experiments, in which a sequential composite binary hypothesis testing problem with observation control is studied. Other works on sequential hypothesis testing with observation control include [8], [9], [10], [11], [12], [13] and more recently [14], [15], [16], [17]. There has also been considerable progress on the special “multi-channel” case of sequential hypothesis testing with observation control, which is also commonly referred to as sequential anomaly detection. In this context, there are multiple data streams, some of which are anomalous, and the goal is to accurately pick out the anomalous ones among them, while observing only a subset of the streams at each time-step [18], [19], [20], [21], [22], [23], [24]. The QCD problem in the multi-channel setting with observation control has been studied recently in [25], [26], [27]. In this context, an unknown subset of the streams undergo a change in their distributions at the unknown change-point, and only a subset of them can be observed at each time-step.

Our work is inspired by [5], in which a general setting for QCD with controlled sensing is considered. In particular, it is assumed that the pre- and post-change distributions can be affected by the choice of a control action, which takes values in a finite set. Furthermore, the post-change is allowed to have parametric uncertainty within a finite parameter set.

The formulation of the optimization problem to obtain the best tradeoff between detection delay and false alarm rate in [5] is not consistent with standard formulations of the QCD problem [3]. In particular, the false alarm constraint used in [5] is one where the probability of stopping (and declaring a change) before a *fixed time* m under the pre-change regime is constrained to some level α . Under such a constraint, if the change happens at any time after $m + 1$, then the delay can be made to be equal 0 by simply stopping at time $m + 1$, if the test did not stop before, while maintaining the false alarm constraint. The authors of [5] refer to [28] to justify their false alarm constraint, but the false alarm constraint proposed in [28] requires the probability of stopping in *any* time-window of size m in the pre-change regime be small (see [28, Section

G. Fellouris is with Department of Statistics, and V.V. Veeravalli is with the ECE Department, of the University of Illinois at Urbana-Champaign. Email: {fellouri, vvv}@illinois.edu.

This research was supported by the US Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196, and by the US National Science Foundation under grants ATD-1737962 and ECCS-2033900, through the University of Illinois at Urbana-Champaign.

II.B)), not just the window from 0 to m . Furthermore, it is more common in the QCD literature to constrain the *mean time to false alarm* [3], which is also a constraint used in [28, Section I].

Another issue with the formulation in [5] is that the measure of delay being used is not specified clearly. There are standard ways to measure delay in QCD problems where no prior is assumed on the change-point, by taking the supremum of the expected delay over all possible change-points (see, e.g., the Lorden and Pollak formulations [3]), but these do not appear to be used in [5]. Furthermore, the asymptotic upper bound on the delay (for a fixed change-point) given in [5, Theorem 3] is larger than the corresponding lower bound given in [5, Theorem 2] (which assumes that the change happens at time 1) by a factor greater than 160. In contrast, first-order asymptotic optimality results in the QCD literature require the ratio of the upper and lower bounds on the delay metric to converge to 1, as the false alarm rate goes to 0 [3].

The ϵ -GCD procedure proposed in [5] for change detection with observation control uses at each time instance a maximum likelihood estimate (MLE) of the post-change parameter to determine the best action at each time step, except that with a fixed probability ϵ , the action is chosen uniformly at random. The MLE at each time instance is determined only by those samples that resulted from random exploration. The use of the current maximum likelihood estimate for determining the current action is in fact the key feature in Chernoff's proposed control policy in [7] for the problem of sequential composite binary hypothesis testing problem with observation control. The need for random exploration of actions in that context arises because the actual post-change parameter may not be distinguishable from other possible post-change parameter values under certain actions. However, exploring at random with probability $\epsilon > 0$ at each time instance will generally lead to substantial performance loss (by roughly a factor of $1/(1 - \epsilon)$) relative to that of an oracle that knows the post-change parameter. Another consideration with the use of the MLE that relies on data from the beginning to estimate the post-change parameter in the QCD setting (in contrast to the sequential hypothesis testing setting) is that it could potentially be biased away from the true post-change parameter due to the pre-change observations if the change-point is not small, as we discuss in Section V.

Our goal in this paper to precisely formulate the QCD problem with controlled sensing, as well as to propose and analyze a novel algorithm for this problem. As described in Section II, we use Lorden's metric [29] for the delay, and we pose the optimization problem as the minimization of this delay metric, under a constraint on the mean time to false alarm (MTFA). In Section III, we derive a universal lower bound on the delay of any procedure, under the MTFA constraint. In Section IV, we develop a procedure for change detection based on the CuSum rule for an arbitrary control policy, and derive some performance bounds for the procedure. In Section V, we specialize the CuSum-based procedure for a specific control policy, to which we refer as the Chernoff-

CuSum procedure, and we establish an upper bound on its conditional expected delay as the MTFA goes to infinity.

II. PROBLEM FORMULATION

Let $\{X_n : n \in \mathbb{N}\}$ be a sequence of random vectors whose values are observed sequentially, let $\{U_n : n \in \mathbb{N}\}$ be a sequence of random variables to be used for randomization purposes, and let $\{\mathcal{F}_n, : n \in \mathbb{N}\}$ the filtration generated by these two sequences, i.e.,

$$\mathcal{F}_n := \sigma(X_m, U_m : 1 \leq m \leq n), \quad n \in \mathbb{N}.$$

We also denote by $\mathcal{F}_0$ the trivial σ -algebra. We assume that, for any $n \in \mathbb{N}$, U_n is independent of $\mathcal{F}_{n-1}$ and uniformly distributed in $[0, 1]$, and that X_n is independent of U_n and conditionally independent of $\mathcal{F}_{n-1}$ given the value of a control A_n . The latter is assumed to be a measurable function of $(X_1, \dots, X_{n-1}, U_1, \dots, U_n)$, and to take values in a finite set $\mathbb{A}$. We refer to the sequence of actions $A := \{A_n : n \in \mathbb{N}\}$ as a *control policy* and we denote by $\mathcal{A}$ the *family* of all control policies, i.e., $\mathcal{A} = \{A_n : n \in \mathbb{N}\} \in \mathcal{A}$.

Let also $\{f_a^\theta : \theta \in \Theta\}$ be a set of densities with respect to a dominating measure, λ , where Θ is an arbitrary finite set, and let $\theta_0 \in \Theta$. We assume that, for any $n \in \mathbb{N}$ and $a \in \mathbb{A}$, X_n has conditional density f_a^θ given that $A_n = a$, where $\theta = \theta_0$ if $n < \nu$ and $\theta \in \Theta \setminus \{\theta_0\}$ if $n \geq \nu$. We refer to ν as the *change-point*, and we assume that it is a completely unknown and deterministic time.

To be more specific, we denote by $P_{\nu,A}^\theta$ the underlying probability measure, and by $E_{\nu,A}^\theta$ the corresponding expectation, when the change-point is ν , the post-change parameter $\theta \in \Theta \setminus \{\theta_0\}$, and the control policy A is used, which means that for any $n \in \mathbb{N}$ and any Borel set B we have

$$P_{\nu,A}^\theta(X_n \in B | \mathcal{F}_{n-1}) = \begin{cases} \int_B f_{A_n}^{\theta_0} d\lambda, & \text{if } n < \nu \\ \int_B f_{A_n}^\theta d\lambda, & \text{if } n \geq \nu. \end{cases}$$

Moreover, we denote by $P_{\infty,A}$ the underlying probability measure, and by $E_{\infty,A}$ the corresponding expectation, when the change never occurs and the control policy A is used, which means that for any $n \in \mathbb{N}$ and any Borel set B we have

$$P_{\infty,A}(X_n \in B | \mathcal{F}_{n-1}) = \prod_{i=1}^n \int_B f_{A_i}^{\theta_0} d\lambda.$$

A *procedure* for quickest change detection with controlled sensing consists of a pair (A, T) , where A is a control policy, i.e., $A \in \mathcal{A}$, and T is an $\{\mathcal{F}_n\}$ -stopping time, i.e., $\{T = n\} \in \mathcal{F}_n$ for every $n \in \mathbb{N}$. We denote by $\mathcal{C}$ the family of all procedures, i.e., $(A, T) \in \mathcal{C}$.

False Alarm Measure. We measure the false alarm performance of a procedure in terms of its mean time to false alarm, and we denote by $\mathcal{C}_\gamma$ the subfamily of procedures for which the mean time to false alarm is at least γ , i.e.,

$$\mathcal{C}_\gamma = \{(A, T) \in \mathcal{C} : E_{\infty,A}[T] \geq \gamma\}.$$

Delay Measure. We use a worst-case measure for delay, that is the commonly used Lorden's measure [29]. Specifically, for any $\theta \in \Theta \setminus \{\theta_0\}$ and $(A, T) \in \mathcal{C}$ we set

$$\mathcal{J}_\theta(A, T) := \sup_{\nu \geq 1} \text{ess sup } \mathbb{E}_{\nu, A}^\theta [(T - \nu + 1)^+ | \mathcal{F}_{\nu-1}]. \quad (1)$$

Optimization Problem. The optimization problem we consider is to find a test that can be designed to belong to $\mathcal{C}_\gamma$ for every $\gamma > 1$ and achieves

$$\inf_{(A, T) \in \mathcal{C}_\gamma} \mathcal{J}_\theta(A, T)$$

to a first-order asymptotic approximation as $\gamma \rightarrow \infty$ simultaneously for every $\theta \in \Theta \setminus \{\theta_0\}$.

We make the following assumptions in our analysis:

A1. For every $\theta \in \Theta \setminus \{\theta_0\}$ and $a \in \mathbb{A}$ we have

$$\int \log(f_a^\theta / f_a^{\theta_0})^2 f_a^\theta d\lambda < V < \infty. \quad (2)$$

A2. For every $\theta \in \Theta \setminus \{\theta_0\}$ there exists an $a \in \mathbb{A}$ so that

$$I_a^\theta := \int \log(f_a^\theta / f_a^{\theta_0}) f_a^\theta d\lambda > 0, \quad (3)$$

i.e., the post-change distribution is distinguishable from the pre-change distribution for at least one choice of control. This assumption implies that

$$I^\theta := \max_{a \in \mathbb{A}} I_a^\theta > 0. \quad (4)$$

A3. For every $\theta, \tilde{\theta} \in \Theta \setminus \{\theta_0\}$, such that $\theta \neq \tilde{\theta}$, there exists an $a \in \mathbb{A}$ so that:

$$\int \log(f_a^\theta / f_a^{\tilde{\theta}}) f_a^\theta d\lambda > 0, \quad (5)$$

i.e., the post-change distribution for two distinct values of the post-change parameter are distinguishable by at least one control.

We will use the two latter assumptions for the consistent estimation of θ in the post-change regime (see Lemma 2).

For any $\theta, \tilde{\theta} \in \Theta$ such that $\theta \neq \tilde{\theta}$ and $a \in \mathbb{A}$ we further define the Bhattacharya coefficient [30]:

$$\rho(\theta, \tilde{\theta}, a) := \int \sqrt{f_a^\theta f_a^{\tilde{\theta}}} d\lambda. \quad (6)$$

By assumptions **A2** and **A3** it follows (see, e.g. [30]) that for any $\theta, \tilde{\theta} \in \Theta$ such that $\theta \neq \tilde{\theta}$ there exists an $a \in \mathbb{A}$ such that $\rho(\theta, \tilde{\theta}, a) < 1$, and as a result

$$\rho(\theta, \tilde{\theta}) := \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} \rho(\theta, \tilde{\theta}, a) < 1. \quad (7)$$

We also denote by ρ the maximum of these quantities:

$$\rho := \max_{\theta \neq \tilde{\theta}} \rho(\theta, \tilde{\theta}). \quad (8)$$

III. UNIVERSAL LOWER BOUND

For any $A \in \mathcal{A}$ and $m \in \mathbb{N}$ we set

$$\Lambda_{m, A}^\theta := \log \left(\frac{f_{A_m}^\theta(X_m)}{f_{A_m}^{\theta_0}(X_m)} \right). \quad (9)$$

and we observe that for any $\nu, t \in \mathbb{N}$ we have

$$\frac{d\mathbb{P}_{\nu, A}^\theta}{d\mathbb{P}_{\infty, A}}(\mathcal{F}_{\nu+t}) = \exp \left\{ \sum_{m=\nu}^{\nu+t} \Lambda_{m, A}^\theta \right\}.$$

Theorem 1. For any $\theta \in \Theta \setminus \{\theta_0\}$, as $\gamma \rightarrow \infty$ we have

$$\inf_{(A, T) \in \mathcal{C}_\gamma} \mathcal{J}_\theta(A, T) \geq \frac{\log \gamma}{I^\theta} (1 + o(1)).$$

Proof. Fix $(A, T) \in \mathcal{C}_\gamma$ and $\theta \in \Theta \setminus \{\theta_0\}$. By [28, Theorem 1] it suffices to show that, for every $\delta > 0$, the sequence

$$\sup_{\nu \in \mathbb{N}} \text{ess sup } \mathbb{P}_{\nu, A}^\theta \left(\max_{1 \leq t \leq n} \sum_{m=\nu}^{\nu+t} \Lambda_{m, A}^\theta > I^\theta (1 + \delta) n \mid \mathcal{F}_{\nu-1} \right)$$

converges to 0 as $n \rightarrow \infty$. Indeed, for every $n, \nu \in \mathbb{N}$ we have

$$\begin{aligned} & \mathbb{P}_{\nu, A}^\theta \left(\max_{1 \leq t \leq n} \sum_{m=\nu}^{\nu+t} \Lambda_{m, A}^\theta > I^\theta (1 + \delta) n \mid \mathcal{F}_{\nu-1} \right) \\ & \leq \mathbb{P}_{\nu, A}^\theta \left(\max_{1 \leq t \leq n} \sum_{m=\nu}^{\nu+t} (\Lambda_{m, A}^\theta - I_{A_m}^\theta) > I^\theta \delta n \mid \mathcal{F}_{\nu-1} \right) \\ & \leq \frac{nV}{(I^\theta \delta n)^2} = \frac{V}{(I^\theta \delta)^2 n}, \end{aligned}$$

Here, the first inequality follows from the fact that $I_a^\theta \leq I^\theta$ for every $a \in \mathbb{A}$, and the second one from a conditional version of Doob's submartingale inequality. We can apply the latter because

$$\left\{ Y_{\nu:t, A}^\theta := \sum_{m=\nu}^{\nu+t} (\Lambda_{m, A}^\theta - I_{A_m}^\theta), \mathcal{F}_{\nu+t}, t \in \mathbb{N} \right\}$$

is a $\mathbb{P}_{\nu, A}^\theta$ -martingale and by assumption **A1** it follows that

$$\text{Var}_{\nu, A}^\theta[Y_{\nu:t, A}^\theta] \leq V t \quad \text{for all } t \in \mathbb{N}.$$

□

IV. CuSUM PROCEDURE WITH CONTROLLED SENSING

In this section we design a stopping rule based on the CuSum test for quickest change detection. Specifically, for any $A \in \mathcal{A}$ and $\theta \in \Theta$ we define the CuSum statistic for detecting the change when the true post-change parameter is θ and control A is used, i.e.,

$$W_{n, A}^\theta := \max_{1 \leq m \leq n} \sum_{u=m}^n \Lambda_{u, A}^\theta,$$

where $\Lambda_{u, A}^\theta$ is defined in (9), which admits the following recursion:

$$W_{n, A}^\theta = \max\{W_{n-1, A}^\theta, 0\} + \Lambda_{n, A}^\theta, \quad n \in \mathbb{N}, \quad W_n^\theta = 0.$$

The change is declared at the first time any of these statistics exceeds a threshold b , i.e.,

$$T_{b,A} := \inf \left\{ n \in \mathbb{N} : \max_{\vartheta \in \Theta} W_{n,A}^{\vartheta} \geq b \right\}.$$

In what follows we simply write T_b instead of $T_{b,A}$, and we call the pair (A, T_b) a *CuSum procedure*. In Lemma 1 we show how to select b for (A, T_b) to satisfy the MTFA constraint. The proof of this result is based on upper bounding the CuSum stopping time by the stopping time of the corresponding Shiryaev-Roberts test (see, e.g. [3] for a similar argument for the QCD problem without observation control).

Lemma 1. *For any $A \in \mathcal{A}$ and $\gamma > 1$ we have $(A, T_b) \in \mathcal{C}_\gamma$ when $b = \log(|\Theta| \gamma)$.*

In the following theorem we give a general result regarding the delay for a CuSum procedure when $\nu = 1$ and the control policy satisfies certain conditions.

Theorem 2. *Let $\theta \in \Theta$. If $b = \log(\gamma|\Theta|)$ and $A \in \mathcal{A}$ is a control policy such that*

$$\sum_{n=1}^{\infty} \mathbb{P}_{1,A}^{\theta} \left\{ \sum_{k=1}^n 1\{I_{A_k}^{\theta} < I^{\theta}\} > \epsilon n \right\} < \infty \quad (10)$$

for every $\epsilon > 0$, then

$$\mathbb{E}_{1,A}^{\theta} [T_b] \leq \frac{\log \gamma}{I^{\theta}} (1 + o(1)) \quad \text{as } \gamma \rightarrow \infty.$$

Proof. In view of Lemma 1, it suffices to show that for every $\epsilon > 0$, we have

$$\limsup_{b \rightarrow \infty} \frac{1}{b} \mathbb{E}_{1,A}^{\theta} [T_b] \leq \frac{1}{I^{\theta} - \epsilon}.$$

For every $b > 0$ we have

$$\mathbb{E}_{1,A}^{\theta} [T_b] \leq \frac{b}{I^{\theta} - \epsilon} + \sum_{n > b/(I^{\theta} - \epsilon)} \mathbb{P}_{1,A}^{\theta} \{T_b > n\}. \quad (11)$$

By the definition of T_b it follows that for every $b > 0$ and $n \in \mathbb{N}$ we have

$$\{T_b > n\} \subseteq \left\{ \max_{\vartheta \in \Theta} \sum_{k=1}^n \Lambda_k^{\vartheta} < b \right\} \subseteq \left\{ \sum_{k=1}^n \Lambda_k^{\theta} < b \right\}.$$

Consequently, for every $b > 0$ and $n > b/(I^{\theta} - \epsilon)$ we have

$$\mathbb{P}_{1,A}^{\theta} \{T_b > n\} \leq \mathbb{P}_{1,A}^{\theta} \left\{ \sum_{k=1}^n \Lambda_k^{\theta} < n(I^{\theta} - \epsilon) \right\}. \quad (12)$$

Therefore, by (11) and (12) we have

$$\mathbb{E}_{1,A}^{\theta} [T_b] \leq \frac{b}{I^{\theta} - \epsilon} + \sum_{n=1}^{\infty} \mathbb{P}_{1,A}^{\theta} \left\{ \sum_{k=1}^n \Lambda_k^{\theta} < n(I^{\theta} - \epsilon) \right\}, \quad (13)$$

thus, it suffices to show that the series in the upper bound converges. This convergence follows by using steps similar to

those used in (21) (and those leading to equation (5.10) in [7]). Specifically, for every $n \in \mathbb{N}$

$$\begin{aligned} \sum_{k=1}^n \Lambda_k^{\theta} &= \sum_{k=1}^n [\Lambda_k^{\theta} - I_{A_k}^{\theta}] - \sum_{k=1}^n (I^{\theta} - I_{A_k}^{\theta}) + nI^{\theta} \\ &\geq \sum_{k=1}^n [\Lambda_k^{\theta} - I_{A_k}^{\theta}] + \Delta \sum_{k=1}^n 1\{I_{A_k}^{\theta} < I^{\theta}\} + nI^{\theta} \end{aligned} \quad (14)$$

where $\Delta := \max_{\vartheta \in \Theta} \max_{a \in \mathbb{A}} (I^{\vartheta} - I_a^{\vartheta})$.

By (11), (12), (14), the probability in (13) is bounded by

$$\mathbb{P}_{1,A}^{\theta} \left\{ \sum_{k=1}^n [\Lambda_k^{\theta} - I_{A_k}^{\theta}] < -\frac{\epsilon}{2} \right\} + \mathbb{P}_{1,A}^{\theta} \left\{ \sum_{k=1}^n 1\{I_{A_k}^{\theta} < I^{\theta}\} > \frac{\epsilon n}{2\Delta} \right\}$$

and it suffices to show that both of these sequences are summable. For the second one, this is trivial when $\Delta = 0$ and it follows by the assumption of the theorem when $\Delta > 0$. For the first one, this can be shown using the fact that

$$\mathbb{E}_{1,A}^{\theta} \left[\Lambda_k^{\theta} - I_{A_k}^{\theta} \middle| \mathcal{F}_{k-1} \right] = 0, \quad \forall k \in \mathbb{N}.$$

and following the steps leading to eq. (5.10) in [7]. $\square$

V. THE CHERNOFF-CUSUM PROCEDURE

We now specialize the CuSum procedure (A, T_b) introduced in the previous section by fixing the control policy. To this end, we introduce, for any $A \in \mathcal{A}$, the MLE of θ at time n , i.e.,

$$\hat{\theta}_n \in \arg \max_{\vartheta \in \Theta} \sum_{m=1}^n \log f_{A_m}^{\vartheta}(X_m). \quad (15)$$

To define the proposed policy, we also need to introduce a sequence of deterministic times,

$$\mathcal{N}_\eta := \{1\} \cup \{\lceil \eta^\ell \rceil : \ell \in \mathbb{N}\}, \quad \text{for some } \eta > 1.$$

Given such a sequence, we propose a control policy A^* according to which

- if $n \in \mathcal{N}_\eta$ or $\hat{\theta}_{n-1} = \theta_0$, then A_n^* is selected uniformly at random by $\mathbb{A}$, using the randomization variable U_n ,
- if $n \notin \mathcal{N}_\eta$ and $\hat{\theta}_{n-1} \neq \theta_0$, A_n^* is selected to maximize the Kullback-Leibler divergence of the post-change versus the pre-change distribution based on the estimate of θ at time $n-1$, i.e.,

$$A_n^* \in \arg \max_{a \in \mathbb{A}} I_a^{\hat{\theta}_{n-1}}.$$

We refer to (A^*, T_b) as the *Chernoff-CuSum* procedure to acknowledge that Chernoff [7] was the first to suggest such a control policy for the sequential design of experiments.

Our main goal in this section is to analyze the delay of *Chernoff-CuSum* procedure as $\gamma \rightarrow \infty$. Towards this end, we first establish an auxiliary consistency result for the MLE, which applies to any control policy that samples uniformly at random from the set $\mathbb{A}$ at the subsequence of time instances $\mathcal{N}_\eta$. Specifically, for any $\theta \in \Theta \setminus \{\theta_0\}$ we bound the tail probability of the random time

$$N^\theta := \inf \left\{ n \in \mathbb{N} : \hat{\theta}_m = \theta \quad \text{for every } m \geq n \right\}, \quad (16)$$

at which the MLE fixes on θ . For this bound, we recall the definition of ρ in (8), and for any $\eta > 1$ we set

$$r_\eta := \frac{-\log \rho}{\log \eta}. \quad (17)$$

Since $\rho < 1$, we clearly have $r_\eta \in (0, \infty)$ for any $\eta > 1$.

Lemma 2. *If $A \in \mathcal{A}$ samples uniformly at random from $\mathbb{A}$ at the subsequence of time instances $\mathcal{N}_\eta$ and $\eta > 1$, then:*

$$\mathbb{P}_{1,A}^\theta \{N^\theta > n\} \leq (|\Theta|/r) n^{-(r_\eta-1)}, \quad \forall n > 1. \quad (18)$$

Proof. From the definition of the MLE in (15) it is clear that $\hat{\theta}_m = \theta$ if $S_m(\theta, \tilde{\theta}) > 0$ for all $\tilde{\theta} \neq \theta$, where

$$S_n(\theta, \tilde{\theta}) := \sum_{m=1}^n Z_m(\theta, \tilde{\theta}), \quad Z_n(\theta, \tilde{\theta}) := \log \left(\frac{f_{A_n}^\theta(X_n)}{f_{A_n}^{\tilde{\theta}}(X_n)} \right).$$

Therefore,

$$\mathbb{P}_{1,A}^\theta \{N^\theta > n\} = \mathbb{P}_{1,A}^\theta \{\exists m \geq n, \tilde{\theta} \neq \theta, \text{ s.t. } S_m(\theta, \tilde{\theta}) \leq 0\},$$

and by an application of the union bound we obtain

$$\mathbb{P}_{1,A}^\theta \{N^\theta > n\} \leq \sum_{\tilde{\theta} \neq \theta} \sum_{m \geq n} \mathbb{P}_{1,A}^\theta \{S_m(\theta, \tilde{\theta}) \leq 0\}. \quad (19)$$

Furthermore, by Markov's inequality it follows that for every $m \in \mathbb{N}$ we have:

$$\mathbb{P}_{1,A}^\theta \{S_m(\theta, \tilde{\theta}) \leq 0\} \leq \mathbb{E}_{1,A}^\theta \left[\exp \left\{ -\frac{1}{2} S_m(\theta, \tilde{\theta}) \right\} \right]. \quad (20)$$

By the conditional independence of $Z_m(\theta, \tilde{\theta})$ given $\mathcal{F}_{m-1}$ and (6)-(7) we have

$$\begin{aligned} & \mathbb{E}_{1,A}^\theta \left[\exp \left\{ -\frac{1}{2} Z_m(\theta, \tilde{\theta}) \right\} \middle| \mathcal{F}_{m-1} \right] \\ &= \begin{cases} \rho(\theta, \tilde{\theta}) & \text{if } m \in \mathcal{N} \\ \rho(\theta, \tilde{\theta}, A_m^*) & \text{otherwise} \end{cases} \leq \rho^{1\{m \in \mathcal{N}\}}. \end{aligned}$$

By this observation and the law of iterated expectation we obtain:

$$\begin{aligned} & \mathbb{E}_{1,A}^\theta \left[e^{-\frac{1}{2} S_m(\theta, \tilde{\theta})} \right] \\ &= \mathbb{E}_{1,A}^\theta \left[e^{-\frac{1}{2} S_{m-1}(\theta, \tilde{\theta})} \mathbb{E}_1^\theta \left[e^{-\frac{1}{2} Z_m(\theta, \tilde{\theta})} \middle| \mathcal{F}_{m-1} \right] \right] \\ &\leq \mathbb{E}_{1,A}^\theta \left[e^{-\frac{1}{2} S_{m-1}(\theta, \tilde{\theta})} \right] \rho^{1\{m \in \mathcal{N}\}}. \end{aligned}$$

Repeating the same argument $m-1$ times we obtain:

$$\begin{aligned} \mathbb{E}_{1,A}^\theta \left[\exp \left\{ -\frac{1}{2} S_m(\theta, \tilde{\theta}) \right\} \right] &\leq \rho^{|\{k \in \mathcal{N}: k \leq m\}|} \\ &\leq \rho^{(\log m)/(\log \eta)} = m^{-r_\eta}, \end{aligned} \quad (21)$$

Using (21), along with (19) and (20), we obtain

$$\mathbb{P}_{1,A}^\theta \{N^\theta > n\} \leq |\Theta| \sum_{m=n}^{\infty} \frac{1}{m^{r_\eta}} \quad \forall n > 1.$$

The result now follows by upper bounding the series in the upper bound by the corresponding integral. $\square$

Theorem 3. *If $\eta \in (1, 1/\sqrt{\rho})$ and $b = \log(\gamma |\Theta|)$, then for every $\theta \in \Theta \setminus \{\theta_0\}$ we have as $\gamma \rightarrow \infty$ that*

$$\mathbb{E}_{1,A^*}^\theta [T_b] \leq \frac{\log \gamma}{I^\theta} (1 + o(1)).$$

Proof. We fix $\theta \in \Theta \setminus \{\theta_0\}$, $a \in \mathbb{A}$, and $\epsilon > 0$. By Theorem 2 it then suffices to show that (10) holds. We also fix $\eta \in (1, 1/\sqrt{\rho})$, and observe that by the definition of r_η in (17) it follows that $r_\eta > 2$. To lighten the notation, in the rest of the proof, we write r instead of r_η .

Since the MLE fixes on θ at time N^θ , defined in (16), we have $\hat{\theta}_k = \theta$ for every $N^\theta \leq k \leq n$ except for at most $\log n / \log \eta$ time instances after N^θ . Therefore,

$$\sum_{k=1}^n 1\{I_{A_k}^\theta < I^\theta\} \leq N^\theta + \frac{\log n}{\log \eta} \quad \forall n \in \mathbb{N}.$$

Moreover, for n large enough we have

$$\begin{aligned} & \mathbb{P}_{1,A^*}^\theta \left\{ \sum_{k=1}^n 1\{I_{A_k}^\theta < I^\theta\} > \epsilon n \right\} \\ &\leq \mathbb{P}_{1,A^*}^\theta \left\{ N^\theta > \epsilon n - \frac{\log n}{\log \eta} \right\} \leq \frac{|\Theta|/r}{(\epsilon n)^{r/2}} \end{aligned}$$

where the second inequality follows from Lemma 2 and the last one because $r > 2$ implies $r-1 > r/2$. The upper bound is clearly summable, and this completes the proof. $\square$

Following the same steps that led to Theorem 3, we can bound the worst-case delay when the changepoint is ν . We state this result without proof in the following theorem.

Theorem 4. *If $\eta \in (1, 1/\sqrt{\rho})$ and $b = \log(\gamma |\Theta|)$, then for every $\theta \in \Theta \setminus \{\theta_0\}$, and fixed $\nu \in \mathbb{N}$, we have as $\gamma \rightarrow \infty$ that*

$$\text{ess sup } \mathbb{E}_{\nu,A^*}^\theta [(T - \nu + 1)^+ | \mathcal{F}_{\nu-1}] \leq \frac{\log \gamma}{I^\theta} (1 + o(1)).$$

However, the $o(1)$ term in Theorem 4 depends on ν and could possibly grow without bound as $\nu \rightarrow \infty$. Therefore, this theorem does not establish the first-order asymptotic optimality of the Chernoff-CuSum test with respect to Lorden's criterion (1). As we alluded to in Section I, one of the reasons the $o(1)$ term in Theorem 4 can grow with ν is that the MLE as defined in (15) can become heavily biased away from θ in the worst case due to the pre-change observations. A possible way to fix this issue might be to modify the MLE so that it forgets past observations whenever there is strong evidence of being in the pre-change regime (e.g., whenever $\hat{\theta}_n = \theta_0$ or $\max_{\theta \in \Theta} W_{n,A}^\theta \leq 0$). However, the analysis of the delay with this modification appears to be considerably more challenging. Even with the modification, the esssup in Theorem 4 may still grow without bound as ν goes to infinity, and perhaps a different delay measure, such as the one suggested by Pollak [31], where one averages over the observations in the pre-change regime, may be required to establish a first-order asymptotic optimality result.

REFERENCES

- [1] H. V. Poor and O. Hadjiladis, *Quickest detection*. Cambridge University Press, 2009.
- [2] A. G. Tartakovsky, I. V. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Change-Point Detection*, ser. Statistics. CRC Press, 2014.
- [3] V. V. Veeravalli and T. Banerjee, “Quickest change detection,” in *Academic press library in signal processing: Array and statistical signal processing*. Academic Press, 2013.
- [4] L. Xie, S. Zou, Y. Xie, and V. V. Veeravalli, “Sequential (quickest) change detection: Classical results and new directions,” *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 2, pp. 494–514, 2021.
- [5] A. Gopalan, B. Lakshminarayanan, and V. Saligrama, “Bandit quickest changepoint detection,” *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [6] V. Krishnamurthy, *Partially observed Markov decision processes: From filtering to controlled sensing*. Cambridge university press, 2016.
- [7] H. Chernoff, “Sequential design of experiments,” *Ann. Math. Statist.*, vol. 30, pp. 755–770, 1959.
- [8] S. A. Bessler, “Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments, Part I: Theory.” Department of Statistics, Stanford University, Technical Report 55, 1960.
- [9] —, “Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments, Part II: Applications.” Department of Statistics, Stanford University, Technical Report 56, 1960.
- [10] A. E. Albert, “The Sequential Design of Experiments for Infinitely Many States of Nature,” *The Annals of Mathematical Statistics*, vol. 32, no. 3, pp. 774 – 799, 1961.
- [11] J. Kiefer and J. Sacks, “Asymptotically Optimum Sequential Inference and Design,” *The Annals of Mathematical Statistics*, vol. 34, no. 3, pp. 705 – 750, 1963.
- [12] S. P. Lalley and G. Lorden, “A Control Problem Arising in the Sequential Design of Experiments,” *The Annals of Probability*, vol. 14, no. 1, pp. 136 – 172, 1986.
- [13] R. Keener, “Second Order Efficiency in the Sequential Design of Experiments,” *The Annals of Statistics*, vol. 12, no. 2, pp. 510 – 532, 1984.
- [14] S. Nitinawarat, G. Atia, and V. Veeravalli, “Controlled sensing for multihypothesis testing,” *IEEE Trans. Aut. Contr.*, vol. 58, pp. 2451–2464, 2013.
- [15] M. Naghshvar and T. Javidi, “Active sequential hypothesis testing,” *The Annals of Statistics*, vol. 41, no. 6, pp. 2703–2738, 2013.
- [16] S. Nitinawarat and V. Veeravalli, “Controlled sensing for sequential multihypothesis testing with controlled markovian observations and non-uniform control cost,” *Sequential Analysis*, vol. 34, no. 1, pp. 1–24, Feb. 2015.
- [17] A. Deshmukh, V. V. Veeravalli, and S. Bhashyam, “Sequential controlled sensing for composite multihypothesis testing,” *Sequential Analysis*, pp. 1–38, 2021.
- [18] K. Cohen and Q. Zhao, “Asymptotically optimal anomaly detection via sequential testing,” *IEEE Transactions on Signal Processing*, vol. 63, no. 11, pp. 2929–2941, 2015.
- [19] A. Gurevich, K. Cohen, and Q. Zhao, “Sequential anomaly detection under a nonlinear system cost,” *IEEE Transactions on Signal Processing*, vol. 67, no. 14, pp. 3689–3703, 2019.
- [20] B. Hemo, T. Gafni, K. Cohen, and Q. Zhao, “Searching for anomalies over composite hypotheses,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 1181–1196, 2020.
- [21] N. K. Vaidhiyan and R. Sundaresan, “Learning to detect an oddball target,” *IEEE Transactions on Information Theory*, vol. 64, no. 2, pp. 831–852, 2018.
- [22] A. Tsopelakos, G. Fellouris, and V. V. Veeravalli, “Sequential anomaly detection with observation control,” in *2019 IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 2389–2393.
- [23] A. Tsopelakos and G. Fellouris, “Sequential anomaly detection with observation control under a generalized error metric,” in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 1165–1170.
- [24] —, “Sequential anomaly detection with sampling constraints,” 2022.
- [25] W. Zhang and Y. Mei, “Bandit change-point detection for real-time monitoring high-dimensional data under sampling control,” *arXiv preprint arXiv:2009.11891*, 2020.
- [26] Q. Xu, Y. Mei, and G. V. Moustakides, “Optimum multi-stream sequential change-point detection with sampling control,” *IEEE Transactions on Information Theory*, vol. 67, no. 11, pp. 7627–7636, 2021.
- [27] A. Chaudhuri, G. Fellouris, and A. Tajer, “Sequential change detection of a correlation structure under a sampling constraint,” in *2021 IEEE International Symposium on Information Theory (ISIT)*, 2021, pp. 605–610.
- [28] T. L. Lai, “Information bounds and quick detection of parameter changes in stochastic systems,” *IEEE Transactions on Information Theory*, vol. 44, no. 7, pp. 2917–2929, November 1998.
- [29] G. Lorden, “Procedures for reacting to a change in distribution,” *The Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1897–1908, Dec. 1971.
- [30] P. Moulin and V. V. Veeravalli, *Statistical Inference for Engineers and Data Scientists*. Cambridge, UK: Cambridge University Press, 2019.
- [31] M. Pollak, “Optimal Detection of a Change in Distribution,” *The Annals of Statistics*, vol. 13, no. 1, pp. 206 – 227, 1985.