

Quickest Change Detection With Controlled Sensing

Venugopal V. Veeravalli¹, *Fellow, IEEE*, Georgios Fellouris², *Member, IEEE*,
and George V. Moustakides³, *Life Senior Member, 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 an unknown parameter, 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. 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 then propose a specific procedure for this problem that employs an adaptive CuSum statistic in which (i) the estimate of the parameter is based on a fixed number of the more recent observations, and (ii) each action is selected to maximize the Kullback-Leibler divergence of the next observation based on the current parameter estimate, apart from a small number of exploration times. We show that this procedure, which we call the Windowed Chernoff-CuSum (WCC), is first-order asymptotically optimal under Lorden’s minimax criterion, for every possible value of the unknown post-change parameter, as the mean time to false alarm goes to infinity. We also provide simulation results to illustrate the performance of the WCC procedure.

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 to detect this change as soon as possible, subject to false alarm constraints. See [2], [3], [4], [5] for books and survey articles on the topic.

In this paper we study an interesting variant of the QCD problem described in [6], [7], which the authors call the “bandit” QCD problem. In this setting, the distribution of the observations is affected not only by the change but also through a control (action) variable that is chosen by the observer. As described in [7], 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 are described in [7].

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* [8]). Such controlled sensing problems have a rich history going back to the seminal work of Chernoff [9] 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 [10], [11], [12], [13], [14], [15] and more recently [16], [17], [18], [19].

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 [20], [21], [22], [23], [24], [25], [26]. The QCD problem in the multi-channel setting with observation control has been studied in [27], and more recently in [6], [28], [29], [30]. 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 [7], 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

Manuscript received 26 October 2023; revised 22 January 2024; accepted 30 January 2024. Date of publication 6 February 2024; date of current version 5 March 2024. This work was supported in part by the U.S. Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196, and in part by the U.S. National Science Foundation under Grant ATD-1737962 and Grant ECCS-2033900, through the University of Illinois at Urbana-Champaign. A preliminary version of this work was presented at the 2022 IEEE International Symposium on Information Theory (ISIT) [DOI: 10.1109/ISIT50566.2022.9834351]. (Corresponding author: Venugopal V. Veeravalli.)

Venugopal V. Veeravalli is with the ECE Department, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: vvv@illinois.edu).

Georgios Fellouris is with the Department of Statistics, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: fellouri@illinois.edu).

George V. Moustakides is with the ECE Department, University of Patras, 26500 Patras, Greece (e-mail: moustaki@upatras.gr).

Digital Object Identifier 10.1109/JSAT.2024.3362324

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 [7] is different from the standard formulations of the QCD problem [4]. In particular, the false alarm constraint used in [7] 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 . The authors of [7] refer to [31] to justify their false alarm constraint, but the false alarm constraint proposed in [31] requires the probability of stopping in *any* time-window of size m in the pre-change regime be small (see [31, Sec. 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* [4], which is also a constraint used in [31, Sec. I].

The measure of delay used in [7] is the expected delay conditioned on a fixed change-point. In QCD problems where no prior is assumed on the change-point, it is common to measure delay by taking the supremum of the expected delay over all possible change-points (see, e.g., the Lorden and Pollak formulations [4]). Furthermore, the asymptotic upper bound on the delay for a fixed change-point given in [7, Th. 3] is larger than the corresponding lower bound given in [7, Th. 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 [4].

The ϵ -Greedy Change Detector (ϵ -GCD) proposed in [7] for change detection with observation control uses at each time instant 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 instant 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 [9] 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 true 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 instant* 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. In contrast, as we show in this paper, we can achieve the performance of the latter to a first-order asymptotic approximation as the false alarm rate goes to 0, by performing random exploration at only $o(n)$ time-steps over a time-horizon of length n .

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. This means that it will take longer for the MLE to converge to the true post-change parameter in the post-change regime, thereby increasing the delay, as the change-point gets larger. We get around this problem by forcing the MLE to use only a window of past observations, and by using a windowed CuSum test (as is done in [32] for the problem of quickest change detection with post-change parametric uncertainty, without controlled sensing).

To sum up, our goal in this paper to precisely formulate the QCD problem with controlled sensing, as well as to propose and analyze a novel, asymptotically optimal algorithm for this problem. As described in Section II, we use Lorden's metric [33] 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, motivated by the discussion in the previous paragraphs, we develop a procedure, which we call the Windowed Chernoff-CuSum (WCC) procedure. We establish that the WCC procedure is asymptotically optimal under Lorden's criterion as the MTFA goes to infinity. In Section V, we provide some simulation results that illustrate the main points of the theoretical analysis.

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.,

$$\mathcal{F}_n := \sigma(X_1, \dots, X_n, U_1, \dots, U_n), \quad n \in \mathbb{N}. \quad (1)$$

We also denote by \mathcal{F}_0 the trivial σ -algebra, and for any $m, n \in \mathbb{N}$ with $m \leq n$ we set:

$$\mathcal{F}_{m:n} := \sigma(X_m, \dots, X_n, U_m, \dots, U_n). \quad (2)$$

Assumption 1: 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/action A_n . We assume that the action A_n takes values in a finite set \mathbb{A} , with at least two elements. Furthermore, for $n > 1$, A_n is a \mathcal{F}_{n-1} -measurable random variable, with A_1 being uniformly distributed on the 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., $A = \{A_n : n \in \mathbb{N}\} \in \mathcal{A}$.

Let $\{f_a^\theta : \theta \in \Theta, a \in \mathbb{A}\}$ be a set of densities with respect to a dominating measure, λ , where Θ is an arbitrary finite set. Furthermore, for some $\theta_0 \notin \Theta$, let $\{f_a^{\theta_0} : a \in \mathbb{A}\}$ be a set of densities with respect to λ .

Let the *change-point* be denoted by $\nu \in \mathbb{N}$, which we assume is completely unknown and deterministic. Given that $A_n = a$, where $a \in \mathbb{A}$, we assume that for $n < \nu$ (pre-change), X_n has conditional density $f_a^{\theta_0}$, and that for $n \geq \nu$ (post-change), X_n has conditional density f_a^θ , with $\theta \in \Theta$.

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

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

Moreover, we denote by $\mathbf{P}_{\infty,A}$ the underlying probability measure, and by $\mathbf{E}_{\infty,A}$ 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

$$\mathbf{P}_{\infty,A}(X_n \in B | \mathcal{F}_{n-1}) = \int_B f_{A_n}^{\theta_0} d\lambda.$$

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}_γ the subfamily of procedures for which the mean time to false alarm is at least γ , i.e.,

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

Delay Measure: We use a worst-case measure for delay, that is the commonly used Lorden's measure [33]. Specifically, for any $\theta \in \Theta$ and $(A, T) \in \mathcal{C}$ we set

$$\mathcal{J}_\theta(A, T) := \sup_{v \geq 1} \text{ess sup } \mathbf{E}_{v,A}^\theta[(T - v + 1)^+ | \mathcal{F}_{v-1}] \quad (3)$$

Optimization Problem: The optimization problem we consider is to find a procedure that can be designed to belong to \mathcal{C}_γ for every $\gamma > 1$, and achieve

$$\inf_{(A,T) \in \mathcal{C}_\gamma} \mathcal{J}_\theta(A, T) \quad (4)$$

to a first-order asymptotic approximation as $\gamma \rightarrow \infty$ simultaneously for every post-change parameter $\theta \in \Theta$.

We make the following further assumptions in our analysis:

Assumption 2: For every $\theta \in \Theta$ there exists an $a \in \mathbb{A}$ such that the Kullback-Leibler (KL) divergence between the densities f_a^θ and $f_a^{\theta_0}$ is positive, i.e.,

$$I_a^\theta := D(f_a^\theta \| f_a^{\theta_0}) = \int \log(f_a^\theta / f_a^{\theta_0}) f_a^\theta d\lambda > 0. \quad (5)$$

This means that the post-change distribution is distinguishable from the pre-change distribution for at least one choice of control.

If Assumption 2 does not hold for some $\theta \in \Theta$, then the change will not be detected efficiently if the true post-change parameter is θ . This assumption implies that

$$I^\theta := \max_{a \in \mathbb{A}} I_a^\theta > 0, \quad \forall \theta \in \Theta. \quad (6)$$

Assumption 3: For every $\theta \in \Theta$ and $a \in \mathbb{A}$ we have

$$\int \log(f_a^\theta / f_a^{\theta_0})^2 f_a^\theta d\lambda < V < \infty. \quad (7)$$

Assumption 3 is a technical condition that is needed in our analysis (see, e.g., Theorem 1 and Theorem 2).

Assumption 4: For every $\theta, \tilde{\theta} \in \Theta$, such that $\theta \neq \tilde{\theta}$, there exists an $a \in \mathbb{A}$ so that:

$$D(f_a^\theta \| f_a^{\tilde{\theta}}) = \int \log(f_a^\theta / f_a^{\tilde{\theta}}) f_a^\theta d\lambda > 0, \quad (8)$$

i.e., the post-change distribution for two distinct values of the post-change parameter are distinguishable by at least one control. If there is no control that distinguishes a particular pair of distinct values of the post-change parameter, then we keep only one of them in Θ and reduce the size of Θ by one.

Assumption 4 is crucial for the consistent estimation of θ in the post-change regime (see Lemma 3).

Remark 1: One could strengthen Assumption 4 by assuming that

$$D(f_a^\theta \| f_a^{\tilde{\theta}}) > 0,$$

for all $\theta, \tilde{\theta} \in \Theta$, such that $\theta \neq \tilde{\theta}$, and for all $a \in \mathbb{A}$. However, this stronger assumption fails to hold in many applications of interest.

For example, consider a multichannel setting with K streams (see, also, Section V) in which a single (unknown) stream of observations undergoes a change in distribution at the change-point, and only one of the streams is observed at each time step. Here $\Theta = \mathbb{A} = \{1, 2, \dots, K\}$. If $\theta = j$, then the j -th stream undergoes a change in distribution, and if $a = i$, the i -th channel is observed. In particular, if $a = 1$, then for $\theta = 2$ and $\tilde{\theta} = 3$, the observed channel (channel 1) has the pre-change distribution, which means that:

$$D(f_1^2 \| f_1^3) = 0.$$

Thus, this strengthened version of Assumption 4 fails to hold in this case.

For any $\theta, \tilde{\theta} \in \Theta$, such that $\tilde{\theta} \neq \theta$, and $a \in \mathbb{A}$, we further define the Bhattacharyya coefficient (see, e.g., [34]):

$$\rho(\theta, \tilde{\theta}, a) := \int \sqrt{f_a^\theta f_a^{\tilde{\theta}}} d\lambda. \quad (9)$$

By Assumptions 2 and 4, it follows that for any $\theta, \tilde{\theta} \in \Theta$, such that $\tilde{\theta} \neq \theta$, there exists an $a \in \mathbb{A}$ such that $\rho(\theta, \tilde{\theta}, a) < 1$, and as a result

$$\rho(\theta, \tilde{\theta}) := \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} \rho(\theta, \tilde{\theta}, a) < 1. \quad (10)$$

We also denote by ρ the maximum of these quantities:

$$\rho := \max_{\theta \neq \tilde{\theta}} \rho(\theta, \tilde{\theta}). \quad (11)$$

The quantity ρ , which is less than 1, will play a key role in the consistent estimation of θ in the post-change regime (see Lemma 3), as well as in lower bounding the post-change expected values of the increments of the proposed detection statistic (see Lemma 4).

Notation: The notation $f(x) \sim g(x)$ as $x \rightarrow x_0$ means that $\lim_{x \rightarrow x_0} \frac{f(x)}{g(x)} = 1$. The minimum of two numbers x and y is denoted as $x \wedge y$. For $k > \ell$, the set of variables $\{x_\ell, x_{\ell+1}, \dots, x_k\}$ is sometimes denoted compactly as $x_{[\ell,k]}$.

III. UNIVERSAL LOWER BOUND

For $\theta \in \Theta$, $A \in \mathcal{A}$, we define the log-likelihood ratio of the observation at time $m \in \mathbb{N}$ as

$$\Lambda_{m,A}^\theta := \log \left(\frac{f_{A_m}^\theta(X_m)}{f_{A_m}^{\theta_0}(X_m)} \right), \quad (12)$$

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

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

In the following result, we provide an asymptotic lower bound on the value of the optimization problem in (4) as the MTFA goes to infinity, i.e., as $\gamma \rightarrow \infty$.

Theorem 1: For any $\theta \in \Theta$, 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 $\theta \in \Theta$. By [31, Th. 1] it suffices to show that, for every $\delta > 0$, the sequence

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

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

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

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

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

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

$$\text{Var}_{v,A}^\theta[Y_{v:t,A}^\theta \mid \mathcal{F}_{v-1}] \leq Vt \quad \text{for all } t \in \mathbb{N}.$$

In the next section we propose a specific procedure for quickest change detection with controlled sensing, and establish that it achieves the lower bound in Theorem 1 asymptotically as $\gamma \rightarrow \infty$.

IV. THE WINDOWED CHERNOFF-CUSUM PROCEDURE

Consider any control policy $A \in \mathcal{A}$, and window of size $w \in \mathbb{N}$. At time $n > w$, the maximum likelihood estimator (MLE) of θ based on observations $\{X_m : m = n - w, \dots, n - 1\}$ is given by:

$$\hat{\theta}_n \in \arg \max_{\vartheta \in \Theta} \sum_{m=n-w}^{n-1} \log f_{A_m}^\vartheta(X_m). \quad (13)$$

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

$$\mathcal{N} \subseteq \{w + 1, w + 2, \dots\},$$

which contains at most q elements in any interval of length w , i.e.,

$$\sum_{i=w+1}^{w+n} \mathbb{I}\{i \in \mathcal{N}\} \leq \frac{n}{w} q, \quad \forall n \in \mathbb{N}, \quad (14)$$

for some positive integer q that is smaller than w . The sequence of times \mathcal{N} will be used for random exploration of controls, which will be crucial for consistent estimation of θ in the post-change regime, under Assumption 4.

Given such a sequence, we propose a control policy A^* as follows.

- For $n = 1, \dots, w$, A_n^* is selected uniformly at random on the set \mathbb{A} , using the randomization variable U_{n-1} .
- If $n > w$, and $n \in \mathcal{N}$, then A_n^* is selected uniformly at random on \mathbb{A} , using the randomization variable U_{n-1} .
- If $n > w$, and $n \notin \mathcal{N}$, A_n^* is selected to maximize the Kullback-Leibler divergence of the post-change versus the pre-change distribution based on $\hat{\theta}_n$, i.e.,

$$A_n^* \in \arg \max_{a \in \mathbb{A}} I_a^{\hat{\theta}_n} = \arg \max_{a \in \mathbb{A}} D(f_a^{\hat{\theta}_n} \| f_a^{\theta_0}). \quad (15)$$

The proposed stopping time is defined using a windowed CuSum test, which was introduced in [32]. Specifically, for any $A \in \mathcal{A}$, and for $n > w$, we define the following CuSum-like statistic recursively:

$$W_{n,A} = \max\{W_{n-1,A}, 0\} + \Lambda_{n,A}^{\hat{\theta}_n}, \quad W_{w,A} = 0, \quad (16)$$

where the log-likelihood ratio Λ is as defined in (12). The change is declared at the first time this statistic exceeds a threshold b , i.e.,

$$T_{b,A} := \inf\{n > w : W_{n,A} \geq b\}. \quad (17)$$

where b is to be selected to satisfy the false alarm constraint.

We refer to the pair (A^*, T_{b,A^*}) as the *Windowed Chernoff-CuSum* (WCC) procedure to acknowledge that Chernoff [9] was the first to suggest a control policy similar to A^* for the problem of sequential hypothesis testing with controlled sensing.

In the remainder of this section, we establish bounds on the performance of the WCC procedure. We begin by establishing a choice of the threshold b that guarantees that the MTFA constraint is met.

Lemma 1: For any $A \in \mathcal{A}$ and $\gamma > 1$, by setting $b = \log \gamma$, we can ensure that $(A, T_{b,A}) \in \mathcal{C}_\gamma$, i.e., that $\mathbf{E}_{\infty,A}[T_{b,A}] \geq \gamma$.

Proof: We can rewrite the recursion for test statistic in (16) as:

$$e^{W_{n,A}} = \max\{e^{W_{n-1,A}}, 1\} \frac{f_{A_n}^{\hat{\theta}_n}(X_n)}{f_{A_n}^{\theta_0}(X_n)}, \quad n > w, \quad e^{W_{w,A}} = 1. \quad (18)$$

We now define a new statistic $\{R_{n,A}\}$ akin to the Shiryaev-Roberts statistic (see, e.g., [4]), which has the recursion:

$$R_{n,A} := (R_{n-1,A} + 1) \frac{f_{A_n}^{\hat{\theta}_n}(X_n)}{f_{A_n}^{\theta_0}(X_n)}, \quad n > w, \quad R_{w,A} := 1. \quad (19)$$

Clearly, $R_{n,A} \geq e^{W_{n,A}}$, for all $n \geq w$.

It is easily seen that $\{R_{n,A} - n : n \geq w\}$ is a martingale sequence under $\mathbf{P}_{\infty,A}$, with respect to the filtration $\{\mathcal{F}_n : n \in \mathbb{N}\}$. Indeed, since A_n is \mathcal{F}_{n-1} -measurable by Assumption 1, and $\hat{\theta}_n$ is \mathcal{F}_{n-1} -measurable by construction (see (13)), we have for $n > w$:

$$\mathbf{E}_{\infty,A}[R_{n,A} - n \mid \mathcal{F}_{n-1}] = R_{n-1,A} + 1 - n = R_{n-1,A} - (n-1).$$

Then, by applying the Optional Sampling Theorem [3]:

$$\mathbf{E}_{\infty,A}[R_{T_{b,A},A} - T_{b,A}] = R_{w,A} - w = 1 - w,$$

which implies that

$$\begin{aligned} \mathbf{E}_{\infty,A}[T_{b,A}] &= \mathbf{E}_{\infty,A}[R_{T_{b,A},A}] + w - 1 \\ &\geq \mathbf{E}_{\infty,A}[e^{W_{T_{b,A},A}}] \geq e^b. \end{aligned}$$

Therefore, setting $b = \log \gamma$ ensures that $(A, T_{b,A}) \in \mathcal{C}_\gamma$. ■

Note that Lemma 1 holds for any control policy $A \in \mathcal{A}$ satisfying Assumption 1, and not just the control policy A^* .

Our main goal in the remainder of this section is to analyze the delay of the WCC procedure as $\gamma \rightarrow \infty$. Towards this end, we first establish an upper bound on our measure of delay for any control policy $A \in \mathcal{A}$ (see (3)):

$$\mathcal{J}_\theta(A, T_{b,A}) = \sup_{\nu \geq 1} \text{ess sup } \mathbf{E}_{\nu,A}^\theta[(T - \nu + 1)^+ \mid \mathcal{F}_{\nu-1}].$$

Lemma 2: For any change-point $\nu \in \mathbb{N}$, control policy $A \in \mathcal{A}$, and threshold $b > 0$

$$\mathcal{J}_\theta(A, T_{b,A}) \leq \sup_{a_1, \dots, a_w} \mathbf{E}_{1,A}^\theta[T_{b,A} \mid A_1 = a_1, \dots, A_w = a_w].$$

Proof: For each $\nu \geq 1$, we denote by $T_{b,A}^\nu$ a version of the stopping time in (17) that starts computing the test statistic from time $\nu + w$, with initialization 0 at time $\nu + w - 1$, i.e.,

$$T_{b,A}^\nu := \inf\{n > \nu + w - 1 : W_{n,A}^\nu \geq b\}, \quad (20)$$

with

$$\begin{aligned} W_{n,A}^\nu &= \max\{W_{n-1,A}^\nu, 0\} + \hat{\Lambda}_{n,A}^\nu, \quad n > \nu + w - 1, \\ W_{\nu+w-1,A}^\nu &= 0. \end{aligned} \quad (21)$$

Then, for any $\nu \geq 1$ we have

$$\begin{aligned} &\mathbf{E}_{\nu,A}^\theta[(T_{b,A} - \nu + 1)^+ \mid \mathcal{F}_{\nu-1}] \\ &\stackrel{(a)}{\leq} w + \mathbf{E}_{\nu,A}^\theta[(T_{b,A} - \nu + 1 - w)^+ \mid \mathcal{F}_{\nu-1}] \\ &\stackrel{(b)}{\leq} w + \mathbf{E}_{\nu,A}^\theta[T_{b,A}^\nu - \nu + 1 - w \mid \mathcal{F}_{\nu-1}] \\ &\stackrel{(c)}{=} \mathbf{E}_{\nu,A}^\theta[\mathbf{E}_{\nu,A}^\theta[T_{b,A}^\nu - \nu + 1 \mid A_{[\nu, \nu+w-1]}, \mathcal{F}_{\nu-1}] \mid \mathcal{F}_{\nu-1}] \\ &\stackrel{(d)}{=} \mathbf{E}_{\nu,A}^\theta[\mathbf{E}_{\nu,A}^\theta[T_{b,A}^\nu - \nu + 1 \mid A_{[\nu, \nu+w-1]}] \mid \mathcal{F}_{\nu-1}] \\ &\leq \sup_{a_1, \dots, a_w} \mathbf{E}_{\nu,A}^\theta[T_{b,A}^\nu - \nu + 1 \mid A_\nu = a_1, \dots, A_{\nu+w-1} = a_w] \\ &\stackrel{(e)}{=} \sup_{a_1, \dots, a_w} \mathbf{E}_{1,A}^\theta[T_{b,A} \mid A_1 = a_1, \dots, A_w = a_w] \end{aligned}$$

where step (a) holds because $x^+ \leq w + (x - w)^+$ for any $x, w > 0$. Step (b) follows because for $n > \nu + w$, the test statistic $W_{n,A}$ (defined in (16)) is greater than or equal to $W_{n,A}^\nu$ (defined in (21)), thus, $T_{b,A} \leq T_{b,A}^\nu$. Also, by definition (see (20)), $T_{b,A}^\nu > \nu + w - 1$, which means that the

positive part can be dropped. Step (c) is due to the law of iterated expectation (and as explained at the end of Section II, $A_{[\nu, \nu+w-1]}$ denotes the set $\{A_\nu, \dots, A_{\nu+w-1}\}$). Step (d) holds because $T_{b,A}^\nu$ depends on the data before time ν only through the controls $A_{[\nu, \nu+w-1]}$, which determine the distributions of the data $X_{[\nu+1, \nu+w]}$. Finally, step (e) holds due to the definition of $T_{b,A}^\nu$ and the fact that $T_{b,A} = T_{b,A}^\nu$ for $\nu = 1$. ■

Next, we establish an auxiliary 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} . Specifically, at each time $n > w$, we upper bound the conditional probability that the MLE at some time instant makes an error in identifying the true post-change parameter θ , given all the already collected data apart from those at the w most recent time instants.

Lemma 3: If $A \in \mathcal{A}$ samples uniformly at random from \mathbb{A} at the subsequence of time instances \mathcal{N} , then:

$$\mathbf{P}_{1,A}^\theta\{\hat{\theta}_n \neq \theta \mid \mathcal{F}_{n-w-1}, A_{[1,w]}\} \leq |\Theta - 1| \rho^q, \quad \forall n > w. \quad (22)$$

Proof: Fix an arbitrary integer n with $n > w$. From the definition of the MLE in (13) it is clear that $\hat{\theta}_n = \theta$ if $S_n(\theta, \tilde{\theta}) > 0$ for all $\tilde{\theta} \neq \theta$, where

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

Therefore,

$$\begin{aligned} &\mathbf{P}_{1,A}^\theta\{\hat{\theta}_n \neq \theta \mid \mathcal{F}_{n-w-1}, A_{[1,w]}\} \\ &= \mathbf{P}_{1,A}^\theta \left(\bigcup_{\tilde{\theta} \neq \theta} \{S_n(\theta, \tilde{\theta}) \leq 0\} \mid \mathcal{F}_{n-w-1}, A_{[1,w]}\right), \end{aligned}$$

and by an application of the union bound we obtain

$$\begin{aligned} &\mathbf{P}_{1,A}^\theta\{\hat{\theta}_n \neq \theta \mid \mathcal{F}_{n-w-1}, A_{[1,w]}\} \\ &\leq \sum_{\tilde{\theta} \neq \theta} \mathbf{P}_{1,A}^\theta\{S_n(\theta, \tilde{\theta}) \leq 0 \mid \mathcal{F}_{n-w-1}, A_{[1,w]}\}. \end{aligned} \quad (23)$$

Furthermore, by the conditional Markov inequality it follows that:

$$\begin{aligned} &\mathbf{P}_{1,A}^\theta\{S_n(\theta, \tilde{\theta}) \leq 0 \mid \mathcal{F}_{n-w-1}, A_{[1,w]}\} \\ &\leq \mathbf{E}_{1,A}^\theta \left[\exp \left\{ -\frac{1}{2} S_n(\theta, \tilde{\theta}) \right\} \mid \mathcal{F}_{n-w-1}, A_{[1,w]} \right]. \end{aligned} \quad (24)$$

By the conditional independence of $Z_m(\theta, \tilde{\theta})$ given \mathcal{F}_{m-1} and (9)-(10) we have

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

where ρ is defined in (11) and \mathbb{I} is the indicator function. Applying the previous inequality for $m = n - 1$ we obtain

$$\begin{aligned} & \mathbb{E}_{1,A}^\theta \left[e^{-\frac{1}{2} S_n(\theta, \hat{\theta})} \mid \mathcal{F}_{n-w-1}, A_{[1,w]} \right] \\ &= \mathbb{E}_{1,A}^\theta \left[e^{-\frac{1}{2} S_{n-1}(\theta, \hat{\theta})} \mathbb{E}_1^\theta \left[e^{-\frac{1}{2} Z_{n-1}(\theta, \hat{\theta})} \mid \mathcal{F}_{n-2} \right] \mid \mathcal{F}_{n-w-1}, A_{[1,w]} \right] \\ &\leq \mathbb{E}_{1,A}^\theta \left[e^{-\frac{1}{2} S_{n-1}(\theta, \hat{\theta})} \rho^{\mathbb{I}\{(n-1) \in \mathcal{N}\}} \mid \mathcal{F}_{n-w-1}, A_{[1,w]} \right]. \end{aligned}$$

Here, the equality follows by an application of the law of iterated expectation, since $w \geq 1$ implies that $\mathcal{F}_{n-2} \supseteq \mathcal{F}_{n-w-1}$, and $n > w$ implies that $A_{[1,w]}$, which is \mathcal{F}_{w-1} -measurable, is also \mathcal{F}_{n-2} -measurable.

Repeating the same argument w times we obtain:

$$\begin{aligned} & \mathbb{E}_{1,A}^\theta \left[\exp \left\{ -\frac{1}{2} S_n(\theta, \hat{\theta}) \right\} \mid \mathcal{F}_{n-w-1}, A_{[1,w]} \right] \\ &\leq \rho^{\sum_{m=n-w}^{n-1} \mathbb{I}\{m \in \mathcal{N}\}} \\ &\leq \rho^q. \end{aligned} \quad (25)$$

where the second inequality follows from (14), according to which at any time-period of length w there are at most q times in \mathcal{N} . Combining (25), along with (23) and (24) completes the proof. ■

We next establish a lower bound on the drift of the log-likelihood ratio process that drives the WCC procedure in the post-change regime. This lower bound will be useful in upper bounding the delay of the WCC procedure. For this, we need to introduce some additional notation. In particular, for every $u, \theta \in \Theta$, $a \in \mathbb{A}$, we set

$$\begin{aligned} B_a(\theta, u) &\equiv \int \log(f_a^u / f_a^{\theta_0}) f_a^\theta d\lambda \\ &= D(f_a^\theta \| f_a^{\theta_0}) - D(f_a^\theta \| f_a^u) = I_a^\theta - D(f_a^\theta \| f_a^u). \end{aligned} \quad (26)$$

Note that

$$B_a(\theta, \theta) = I_a^\theta.$$

Also, let

$$a^*(u) \in \arg \max_{a \in \mathbb{A}} I_a^u = \arg \max_{a \in \mathbb{A}} D(f_a^u \| f_a^{\theta_0}). \quad (27)$$

Note that $a^*(u)$ is an optimal control for the WCC procedure at time $n \notin \mathcal{N}$, when the post-change parameter estimate $\hat{\theta}_n = u$ (see (15)). Also, it is clear that:

$$B_{a^*(\theta)}(\theta, \theta) = I_{a^*(\theta)}^\theta = \max_{a \in \mathbb{A}} I_a^\theta = I^\theta. \quad (28)$$

Lemma 4: Let $n > w$. Then,

$$\begin{aligned} & \mathbb{E}_{1,A^*}^\theta \left[\Lambda_{n,A^*}^{\hat{\theta}_n} \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^* \right] \\ &\geq \begin{cases} I^\theta - J^\theta \rho^q, & n \notin \mathcal{N} \\ \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} (I_a^\theta - K_a^\theta \rho^q), & n \in \mathcal{N} \end{cases} \end{aligned} \quad (29)$$

where

$$\begin{aligned} J^\theta &\equiv (|\Theta| - 1) \cdot \left(I^\theta - \min_{u \neq \theta} B_{a^*(u)}(\theta, u) \right), \\ K_a^\theta &\equiv (|\Theta| - 1) \cdot \left(I_a^\theta - \min_{u \neq \theta} B_a(\theta, u) \right), \end{aligned}$$

with $a^*(u)$ as defined in (27). Moreover, for large enough q ,

both lower bounds in (29) are strictly positive.

Proof: When $n \notin \mathcal{N}$, A_n^* is chosen according to (15), and

$$\begin{aligned} & \mathbb{E}_{1,A^*}^\theta \left[\Lambda_{n,A^*}^{\hat{\theta}_n} \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^* \right] \\ &= \sum_{u \in \Theta} \mathbb{E}_{1,A^*}^\theta \left[\log \frac{f_{A_n^*}^{\hat{\theta}_n}(X_n)}{f_{A_n^*}^{\theta_0}(X_n)} \mid \hat{\theta}_n = u, \mathcal{F}_{n-w-1}, A_{[1,w]}^* \right] \\ &\quad \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n = u \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \\ &\stackrel{(a)}{=} \sum_{u \in \Theta} \left(\int \log \left(\frac{f_{a^*(u)}^u}{f_{a^*(u)}^{\theta_0}} \right) f_{a^*(u)}^\theta d\lambda \right) \\ &\quad \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n = u \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \\ &= \sum_{u \in \Theta} B_{a^*(u)}(\theta, u) \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n = u \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \\ &\stackrel{(b)}{\geq} I^\theta \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n = \theta \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) + \min_{u \neq \theta} B_{a^*(u)}(\theta, u) \\ &\quad \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n \neq \theta \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \\ &= I^\theta - \left(I^\theta - \min_{u \neq \theta} B_{a^*(u)}(\theta, u) \right) \\ &\quad \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n \neq \theta \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \\ &\stackrel{(c)}{\geq} I^\theta - \left(I^\theta - \min_{u \neq \theta} B_{a^*(u)}(\theta, u) \right) \cdot (|\Theta| - 1) \cdot \rho^q \end{aligned}$$

where (a) follows from (15) and (27), and from the fact that conditioned on $\hat{\theta}_n = u$, the optimal control $A_n^* = a^*(u)$ and X_n is independent of everything else. Step (b) follows from (28), and (c) from Lemma 3. This proves the first inequality in (29).

When $n \in \mathcal{N}$, A_n^* is chosen uniformly at random from the set \mathbb{A} , and following steps similar to those used in the case when $n \notin \mathcal{N}$, we obtain:

$$\begin{aligned} & \mathbb{E}_{1,A^*}^\theta \left[\Lambda_{n,A^*}^{\hat{\theta}_n} \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^* \right] \\ &= \sum_{u \in \Theta} \mathbb{E}_{1,A^*}^\theta \left[\log \left(\frac{f_{A_n^*}^{\hat{\theta}_n}(X_n)}{f_{A_n^*}^{\theta_0}(X_n)} \right) \mid \hat{\theta}_n = u, \mathcal{F}_{n-w-1}, A_{[1,w]}^* \right] \\ &\quad \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n = u \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \\ &= \sum_{u \in \Theta} \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} B_a(\theta, u) \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n = u \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \\ &= \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} \sum_{u \in \Theta} B_a(\theta, u) \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n = u \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \\ &\geq \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} \left[I_a^\theta \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n = \theta \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \right. \\ &\quad \left. + \min_{u \neq \theta} B_a(\theta, u) \cdot \mathbf{P}_{1,A^*}^\theta (\hat{\theta}_n \neq \theta \mid \mathcal{F}_{n-w-1}, A_{[1,w]}^*) \right] \\ &\geq \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} \left[I_a^\theta - \left(I_a^\theta - \min_{u \neq \theta} B_a(\theta, u) \right) \cdot (|\Theta| - 1) \cdot \rho^q \right], \end{aligned}$$

which proves the second inequality in (29).

Finally, we note that from (26) it follows that, for all $a \in \mathbb{A}$ and $u \in \Theta$,

$$B_a(\theta, u) \leq I_a^\theta \quad \text{and} \quad B_{a^*(u)}(\theta, u) \leq I^\theta.$$

Therefore, $J^\theta \geq 0$ and $K_a^\theta \geq 0$, for all $a \in \mathbb{A}$. Nevertheless, since $0 < \rho < 1$, as $q \rightarrow \infty$ we have:

$$I^\theta - J^\theta \rho^q \rightarrow I^\theta > 0$$

and

$$\frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} (I_a^\theta - K_a^\theta \rho^q) \rightarrow \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} I_a^\theta > 0,$$

which proves that both lower bounds in (29) are both positive for large enough $q > 0$. ■

Theorem 2: As $q, w, b \rightarrow \infty$ so that $q = o(w)$ and $w = o(b)$,

$$\sup_{a_1, \dots, a_w} \mathbb{E}_{1, A^*}^\theta [T_{b, A^*} | A_1^* = a_1, \dots, A_w^* = a_w] \leq \frac{b}{I^\theta} (1 + o(1)).$$

Proof: Define the stopping time

$$T'_{b, A^*} := \inf\{n > w : W'_{n, A^*} \geq b\}.$$

with

$$\begin{aligned} W'_{n, A^*} &= W'_{n-1, A^*} + \Lambda_{n, A^*}^{\hat{\theta}_n}, \quad n > w, \\ W'_{w, A^*} &= 0. \end{aligned} \quad (30)$$

Note from (16) that the difference between W'_{n, A^*} and W_{n, A^*} is that we do not take the positive part on the right-hand-side of the recursion in (30). This means that $T_{b, A^*} \leq T'_{b, A^*}$. Therefore, it suffices to show that as $q, w, b \rightarrow \infty$ so that $q = o(w)$ and $w = o(b)$,

$$\sup_{a_1, \dots, a_w} \mathbb{E}_{1, A^*}^\theta [T'_{b, A^*} | A_1^* = a_1, \dots, A_w^* = a_w] \leq \frac{b}{I^\theta} (1 + o(1)).$$

In what follows, we fix arbitrary a_1, \dots, a_w in \mathbb{A} , and we show that

$$\mathbb{E}_{1, A^*}^\theta [T'_{b, A^*} | A_1^* = a_1, \dots, A_w^* = a_w] \leq \frac{b}{I^\theta} (1 + o(1)).$$

To see this, we make the following identifications in Proposition 1 in the Appendix:

$$\begin{aligned} \mathbf{P} &\leftrightarrow \mathbf{P}_{1, A^*}^\theta (\cdot | A_1^* = a_1, \dots, A_w^* = a_w) \\ \mathcal{G}_n &\leftrightarrow \mathcal{F}_n \\ Y_n &\leftrightarrow \Lambda_{n, A^*}^{\hat{\theta}_n} \\ \tau_b &\leftrightarrow T'_{b, A^*} \end{aligned}$$

We also observe that by Lemma 4 we can set μ_*, μ^*, μ, ν in Proposition 1 as follows:

$$\begin{aligned} \mu_* &\leftrightarrow \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} (I_a^\theta - K_a^\theta \rho^q) \\ \mu^* &\leftrightarrow I^\theta - J^\theta \rho^q \\ \mu &\leftrightarrow I^\theta \\ V &\leftrightarrow \nu \end{aligned}$$

By Lemma 4 again we have $\mu^* \rightarrow I^\theta = \mu$ as $q \rightarrow \infty$. Thus, the condition in part (iii) of Proposition 1 is satisfied, which implies the result. ■

We then have the following result that establishes the first-order optimality of the WCC procedure.

Corollary 1: If $b = \log \gamma$, then as $q, w, b \rightarrow \infty$ so that $q = o(w)$ and $w = o(\log \gamma)$, then

$$\mathcal{J}_\theta(A^*, T_{b, A^*}) \sim \inf_{(A, T) \in \mathcal{C}_\gamma} \mathcal{J}_\theta(A, T) \sim \frac{\log \gamma}{I^\theta}$$

Proof: This follows by combining Theorem 1, Lemma 1, Lemma 2, and Theorem 2. ■

V. SIMULATIONS

We consider a simulation study in which, for each $n \in \mathbb{N}$, X_n is a K -dimensional random vector (X_n^1, \dots, X_n^K) with independent components. We assume that each of these components has variance 1 and that the mean changes in an unknown subset of components, θ . Specifically, if $k \notin \theta$, each X_n^k has mean 0, whereas if $k \in \theta$, then X_n^k has mean 0 for $n < \nu$ and mean μ_k for $n \geq \nu$. Therefore, in this context, the parameter space, Θ , consists of all non-empty subsets of $\{1, \dots, K\}$. We further assume that it is possible to sample only one of these components at each time instant. Therefore, the action space \mathbb{A} consists of all singletons $\{\{k\} : 1 \leq k \leq K\}$.

We set $K = 10$, $(\mu_1, \dots, \mu_k) = (0.5, 0.5, 1, 1, 1, \dots, 1)$, and we consider the case that only the first three components experience the change, i.e., the true value of θ is $\{1, 2, 3\}$. That is, among the three affected components, only one of them has the largest possible drift. Moreover, based on Lemma 2, we assume that the change occurs at time $\nu = 1$.

We consider two schemes.

- (i) The proposed *Windowed Chernoff-CuSum* procedure, with two choices for the size of the window, w , 10 and 20. With this procedure, we sample uniformly at random from all sources at the first w time instants, and an alarm can be raised only after this time. For both window choices, we set $q = 1$, that is, during each interval of the form $[mw, (m+1)w)$, where $m \in \mathbb{N}$, only once the component is selected at random. Apart from these times, at each time after w , we estimate the subset of components that have been affected by the change, and select to sample among those the one with the largest post-change mean. This means that whenever a component apart from the first two is estimated to be affected by the change, it can be selected for sampling next. In order to break such ties, we select the component whose observations in the last w time instants had the largest average. This choice was found to lead to much better performance than random sampling among the components that are currently estimated as affected by the change.
- (ii) The *greedy* algorithm considered in [6], [27], [28], [29], [30], according to which we keep sampling the same component until its cumulative log-likelihood ratio statistic hits either the threshold b or 0. In the former case, the alarm is raised and the process is terminated. In the latter, all past observations are forgotten and the next component is sampled. This algorithm has been

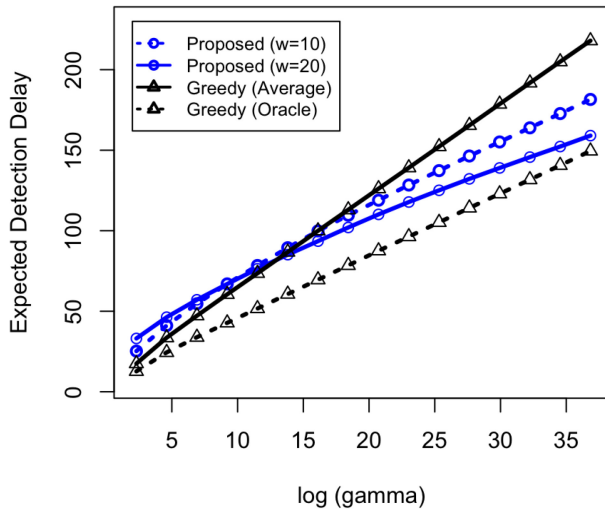


Fig. 1. The vertical axis represents expected detection delay and the horizontal axis $\log \gamma$. The blue lines/circles correspond to the proposed, window-limited Chernoff-CuSum scheme with windows $w = 10$ (dashed line) and $w = 20$ (solid line). The black lines/triangles correspond to the greedy algorithm. The solid line corresponds to its average behavior, and the dashed line to its best case scenario, which we refer to as “oracle”.

shown to enjoy certain asymptotic optimality properties when either only a single component is affected by the change [28] or all affected components have the same post-change distribution [29], but not in the setup of this simulation study. Moreover, its performance depends heavily on the component that is being sampled when the change occurs. In order to capture its average behavior, we simulate at random the component that is being sampled at the time of the change. For comparison purposes, we also consider the best case/oracle scenario, where it is component 3, i.e., the one with the largest post-change mean among those affected, that is being sampled at the time of the change in each simulation run.

In all cases, the threshold b is selected as $\log \gamma$, where γ takes values in $\{10^i : i = 1, \dots, 16\}$. Figure 1 presents the results, where in the vertical axis we plot the average detection delays of the above schemes, based on 16,000 simulation runs for each of them, and in the horizontal axis we plot $\log(\gamma)$.

Based on this graph, we can make the following observations. First, for the proposed Windowed Chernoff-CuSum, the larger window, i.e., setting w equal to 20 instead of 10, leads to smaller expected detection delay for large enough values of γ . This is consistent with our asymptotic theory, according to which the window size w needs to go to infinity as $\gamma \rightarrow \infty$ in order to achieve asymptotic optimality. Second, both these window choices lead to expected detection delays smaller than that of the greedy algorithm. In fact, the expected detection delay of the proposed scheme, with the larger window, approaches the oracle for the greedy algorithm as γ increases.

VI. CONCLUSION

We considered the problem of quickest change detection when there is parameter uncertainty regarding the post-change regime, and the distribution of the observations is affected

not only by the change but also by a control input that is chosen by the observer. We provided a precise mathematical formulation for a general setting of this problem. We then developed a procedure that is asymptotically optimal under Lorden’s criterion as the false alarm rate goes to zero.

Our asymptotic optimality result is established as the expected time to false alarm goes to infinity, under the assumption of a finite and fixed action space and a finite and fixed parameter space regarding the post-change distribution. Therefore, a natural direction for further research is the theoretical analysis of this algorithm as the sizes of the action space and the post-change parameter space go to infinity, or even in the case of an infinite action space and an infinite post-change parameter space. Another direction for further research is the consideration of more general stochastic models for the observations, in which the assumption of conditional independence is removed, in the spirit of [18]. Also, in our problem formulation, we implicitly assumed that all control actions incur the same cost. Extending the formulation and the analysis to allow for non-uniform costs across control actions would be of interest in certain applications. Finally, the proposed algorithm relies on a window of past observations during which the post-change parameter is estimated, and this is subsequently used to assign the appropriate control. Our asymptotic theory requires that the length of this window goes to infinity as the false alarm rate goes to 0. It would be of interest to obtain more precise bounds for the worst-case conditional expected detection delay that could inform the selection of this window size.

ACKNOWLEDGMENT

The first author owes a debt of gratitude to Prof. Toby Berger for encouraging him to work on quickest change detection when they were colleagues at Cornell University in the late 1990s, and for the engaging discussions they had on this and many other topics.

APPENDIX

In this Appendix we present a general, non-asymptotic upper bound for the expected hitting time of a sequence for which the conditional expectation of each increment given all past observations does not exceed a certain positive constant. This is the basis of the proof of Theorem 2, but it is stated in a more abstract form, as it can be of independent interest.

Proposition 1: Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space, and let \mathbf{E} denote expectation with respect to \mathbf{P} . Let $\{\mathcal{G}_n, n \geq 0\}$ be a filtration on this space, where \mathcal{G}_0 is the trivial σ -algebra. Let $\{Y_n, n \geq 1\}$ a sequence of random variables in \mathcal{L}^2 that is adapted to $\{\mathcal{G}_n, n \geq 1\}$. Let $\omega \in \mathbb{N}$, $\mathcal{N} \subseteq \{\omega + 1, \omega + 2, \dots\}$, $\mu_*, \mu^*, \mu, \nu > 0$, and $q \in \mathbb{N}$, and suppose that, for every $n > \omega$,

$$\mathbf{E}[Y_n | \mathcal{G}_{n-\omega-1}] \geq \begin{cases} \mu^*, & n \notin \mathcal{N} \\ \mu_*, & n \in \mathcal{N} \end{cases} \quad \text{a.s.} \quad (31)$$

$$\mathbf{E}[Y_n | \mathcal{G}_{n-1}] \leq \mu \quad \text{a.s.} \quad (32)$$

$$\mathbf{E}[Y_n^2 | \mathcal{G}_{n-1}] \leq \nu \quad \text{a.s.} \quad (33)$$

and

$$\sum_{i=w+1}^{w+n} \mathbb{I}\{i \in \mathcal{N}\} \leq \frac{n}{w} q. \quad (34)$$

For any $b > 0$, set

$$\tau_b \equiv \inf\{n > w : S_n \geq b\}, \quad \text{where} \quad S_n \equiv \sum_{i=w+1}^n Y_i.$$

Then:

- (i) $\mathbb{E}[\tau_b] < \infty$.
- (ii) There is a $C > 0$, that does not depend on b, q , or w , so that

$$\mathbb{E}[\tau_b] \leq \frac{b + w\mu + C(1 + \sqrt{b} + \sqrt{w})}{\mu^*(1 - q/w)}.$$

- (iii) If $\mu^* \rightarrow \mu$ as $q, w \rightarrow \infty$ so that $q = o(w)$, then

$$\mathbb{E}[\tau_b] \leq \frac{b}{\mu} (1 + o(1)),$$

where $o(1)$ is a vanishing term as $q, w, b \rightarrow \infty$ so that $q = o(w)$ and $w = o(b)$.

Remark 2: This proposition generalizes [14, Proposition 2.1], as in the latter condition (31) is assumed to hold for $w = 1$ and condition (34) for $q = 0$.

Proof: (i) For any $c > 0$ and $n > w$ we have

$$\begin{aligned} & \mathbb{E}[Y_n \cdot \mathbb{I}\{Y_n \geq c\} \mid \mathcal{G}_{n-w-1}] \\ & \stackrel{(a)}{\leq} \sqrt{\mathbb{E}[Y_n^2 \mid \mathcal{G}_{n-w-1}]} \sqrt{\mathbb{P}(Y_n \geq c \mid \mathcal{G}_{n-w-1})} \\ & \stackrel{(b)}{\leq} \frac{\mathbb{E}[Y_n^2 \mid \mathcal{G}_{n-w-1}]}{c} \stackrel{(c)}{\leq} \frac{\nu}{c} \quad \text{a.s.}, \end{aligned}$$

where (a) follows by the conditional Cauchy-Schwarz inequality, (b) by the conditional Markov inequality, (c) by the law of iterated expectation and (33). As a result, for c large enough we have

$$\mathbb{E}[Y_n \cdot \mathbb{I}\{Y_n \geq c\} \mid \mathcal{G}_{n-w-1}] \leq (\mu_* \wedge \mu^*)/2 \quad \text{a.s.}, \quad (35)$$

where we use $a \wedge b$ to denote $\min\{a, b\}$. Consequently, for c large enough we have

$$\begin{aligned} & \mathbb{E}[Y_n \wedge c \mid \mathcal{G}_{n-w-1}] \\ & \geq \mathbb{E}[Y_n \cdot \mathbb{I}\{Y_n < c\} \mid \mathcal{G}_{n-w-1}] \\ & = \mathbb{E}[Y_n \mid \mathcal{G}_{n-w-1}] - \mathbb{E}[Y_n \cdot \mathbb{I}\{Y_n \geq c\} \mid \mathcal{G}_{n-w-1}] \\ & \geq \mu_* \wedge \mu^* - (\mu_* \wedge \mu^*)/2 = (\mu_* \wedge \mu^*)/2 \quad \text{a.s.}, \end{aligned} \quad (36)$$

where the last inequality follows from (31) and (35). For every $n \in \mathbb{N}$ with $n > w$ we set

$$\tilde{Y}_n \equiv Y_n \wedge c \quad \text{and} \quad \tilde{S}_n \equiv \tilde{Y}_{w+1} + \dots + \tilde{Y}_n,$$

thus, by (36), we have

$$\mathbb{E}[\tilde{Y}_n \mid \mathcal{G}_{n-w-1}] \geq (\mu_* \wedge \mu^*)/2 \quad \text{a.s.} \quad (37)$$

For every $b > 0$ we set

$$\tilde{\tau} \equiv \inf\{n > w : \tilde{S}_n \geq b\}.$$

Note that $\tau_b \leq \tilde{\tau}$ almost surely. Let $m \in \mathbb{N}$ with $m > w$. Then:

$$\tilde{S}_{\tilde{\tau} \wedge m} \leq b + c. \quad (38)$$

Indeed, if $\tilde{\tau} > m$, then $\tilde{S}_{\tilde{\tau} \wedge m} = \tilde{S}_m < b$, whereas if $\tilde{\tau} \leq m$, then

$$\tilde{S}_{\tilde{\tau} \wedge m} = \tilde{S}_{\tilde{\tau}} \leq b + \tilde{Y}_{\tilde{\tau}} \leq b + c.$$

Moreover,

$$\tilde{S}_{\tilde{\tau} \wedge m} = \sum_{n=w+1}^{\tilde{\tau} \wedge m + w} \tilde{Y}_n - \sum_{n=\tilde{\tau} \wedge m + 1}^{\tilde{\tau} \wedge m + w} \tilde{Y}_n. \quad (39)$$

We start by lower bounding the conditional expectation of the first term in the right-hand side in (39). To this end, we first observe that

$$\begin{aligned} & \mathbb{E} \left[\sum_{n=w+1}^{\tilde{\tau} \wedge m + w} \tilde{Y}_n \right] \\ & = \mathbb{E} \left[\sum_{n=w+1}^{\infty} \tilde{Y}_n \cdot \mathbb{I}\{\tilde{\tau} \wedge m + w \geq n\} \right] \\ & \stackrel{(a)}{=} \sum_{n=w+1}^{\infty} \mathbb{E}[\tilde{Y}_n \cdot \mathbb{I}\{\tilde{\tau} \wedge m + w \geq n\}] \\ & \stackrel{(b)}{=} \sum_{n=w+1}^{\infty} \mathbb{E}[\mathbb{E}[\tilde{Y}_n \mid \mathcal{G}_{n-w-1}] \cdot \mathbb{I}\{\tilde{\tau} \wedge m + w \geq n\}]. \end{aligned} \quad (40)$$

The interchange of series and expectation in step (a) is justified because $\tilde{\tau} \wedge m$ is a bounded stopping time and the sequence (\tilde{Y}_n) is bounded above. Step (b) follows by an application of the law of iterated expectation, in view of the fact that $\{\tilde{\tau} \wedge m + w \geq n\}$ is \mathcal{F}_{n-w-1} -measurable for $n > w$.

Then, by (40) and (37) we conclude that

$$\begin{aligned} \mathbb{E} \left[\sum_{n=w+1}^{\tilde{\tau} \wedge m + w} \tilde{Y}_n \right] & \geq \frac{\mu_* \wedge \mu^*}{2} \mathbb{E} \left[\sum_{n=w+1}^{\tilde{\tau} \wedge m + w} 1 \right] \\ & = \frac{\mu_* \wedge \mu^*}{2} \mathbb{E}[\tilde{\tau} \wedge m]. \end{aligned} \quad (41)$$

We continue by upper bounding the conditional expectation of the second term in the right-hand side in (39). Since (\tilde{Y}_n) is a bounded above sequence and $\tilde{\tau} \wedge m$ is a bounded stopping time that takes values larger than w , by Wald's identity we have

$$\mathbb{E} \left[\sum_{n=\tilde{\tau} \wedge m + 1}^{\tilde{\tau} \wedge m + w + 1} \tilde{Y}_n \right] = \mathbb{E} \left[\sum_{n=\tilde{\tau} \wedge m + 1}^{\tilde{\tau} \wedge m + w + 1} \mathbb{E}[\tilde{Y}_n \mid \mathcal{G}_{n-1}] \right] \leq \mu w \quad (42)$$

where the inequality follows from (32).

Combining (38), (39), (41), and (42), we conclude that for every $m \in \mathbb{N}$ with $m > w$ we have

$$\mathbb{E}[\tilde{\tau} \wedge m] \leq \frac{b + c + w\mu}{(\mu_* \wedge \mu^*)/2}.$$

Letting $m \rightarrow \infty$, by the Monotone Convergence Theorem we obtain

$$\mathbb{E}[\tau_b] \leq \mathbb{E}[\tilde{\tau}] \leq \frac{b + c + w\mu}{(\mu_* \wedge \mu^*)/2} < \infty. \quad (43)$$

(ii) We observe that

$$\begin{aligned}
& \mathbb{E} \left[\sum_{n=w+1}^{\tau_b+w} Y_n \right] \\
& \stackrel{(a)}{=} \mathbb{E} \left[\sum_{n=w+1}^{\infty} \mathbb{E}[Y_n | \mathcal{G}_{n-w-1}] \cdot \mathbb{I}\{\tau_b + w \geq n\} \right] \\
& \stackrel{(b)}{\geq} \mathbb{E} \left[\sum_{n=w+1}^{\infty} \mathbb{E}[Y_n | \mathcal{G}_{n-w-1}] \cdot \mathbb{I}\{n \notin \mathcal{N}\} \cdot \mathbb{I}\{\tau_b + w \geq n\} \right] \\
& \stackrel{(c)}{\geq} \mu^* \mathbb{E} \left[\sum_{n=w+1}^{\tau_b+w} \mathbb{I}\{n \notin \mathcal{N}\} \right] \\
& \stackrel{(d)}{\geq} \mu^* (1 - q/w) \mathbb{E}[\tau_b].
\end{aligned} \tag{44}$$

Step (a) is obtained by following exactly the same steps as in (40), with \tilde{Y}_n replaced by Y_n and $\tilde{\tau} \wedge m$ replaced by τ_b , and it is justified because, by the result of part (i), τ_b is integrable, and, by assumption (33), $(\mathbb{E}[Y_n^2 | \mathcal{G}_{n-1}])$ is bounded. Step (b) and (c) follow by (31), and step (d) by (34).

Moreover,

$$\mathbb{E} \left[\sum_{n=w+1}^{w+\tau_b} Y_n \right] = \mathbb{E}[S_{\tau_b} | \mathcal{F}_w] + \mathbb{E} \left[\sum_{n=\tau_b+1}^{w+\tau_b} Y_n \right] \leq b + \mathbb{E}[R_b] + w\mu, \tag{45}$$

where $R_b \equiv S_{\tau_b} - b$ and the inequality follows in exactly the same way as (42), again with \tilde{Y}_n replaced by Y_n and $\tilde{\tau} \wedge m$ replaced by τ_b . From (44) and (45) we have

$$\mathbb{E}[\tau_b] \leq \frac{b + \mathbb{E}[R_b] + w\mu}{\mu^*(1 - q/w)}, \tag{46}$$

and it is clear that it suffices to show that there is a $C > 0$ so that

$$\mathbb{E}[R_b] \leq C(1 + \sqrt{b} + \sqrt{w}). \tag{47}$$

In order to do so, we first claim that

$$0 \leq 2 \int_0^b R_c dc \leq \sum_{n=w+1}^{\tau_b} Y_n^2 - R_b^2, \tag{48}$$

where R_c is defined as in R_b with b replaced by c . To prove this claim, we introduce the sequence of ladder variables

$$\begin{aligned}
\rho_m & \equiv \inf\{n > \rho_{m-1} : S_n > S_{\rho_{m-1}}\}, \quad m \in \mathbb{N}, \\
\rho_0 & \equiv \inf\{n > w : S_n > 0\}.
\end{aligned}$$

These can be shown to be integrable, in the same way as we show the integrability of τ_b in (i). Moreover, we set

$$M_b \equiv \inf\{m \in \mathbb{N} : S_{\rho_m} \geq b\},$$

thus, $\tau_b = \rho_{M_b}$ and $R_b = S_{\rho_{M_b}} - b$. Then, as in [35, eq. (1)], we have

$$0 \leq 2 \int_0^b R_c dc = \sum_{m=1}^{M_b} (S_{\rho_m} - S_{\rho_{m-1}})^2 - R_b^2. \tag{49}$$

By the definition of ladder times, we have

$$0 < S_{\rho_m} - S_{\rho_{m-1}} \leq Y_{\rho_m}^+, \quad \forall m \in \mathbb{N}.$$

Therefore

$$\begin{aligned}
\sum_{m=1}^{M_b} (S_{\rho_m} - S_{\rho_{m-1}})^2 & \leq \sum_{m=1}^{M_b} Y_{\rho_m}^2 \\
& \leq \sum_{m=1}^{M_b} \sum_{n=\rho_{m-1}+1}^{\rho_m} Y_n^2 \\
& = \sum_{n=w+1}^{\tau_b} Y_n^2
\end{aligned} \tag{50}$$

where the equality holds because

$$\tau_b = \rho_{M_b} = \sum_{m=1}^{M_b} (\rho_m - \rho_{m-1}).$$

Combining (49) and (50) proves (48).

Then, taking expectations in (48) and applying Jensen's inequality, we have

$$0 \leq \mathbb{E} \left[\sum_{n=w+1}^{\tau_b+w} Y_n^2 \right] - (\mathbb{E}[R_b])^2. \tag{51}$$

Moreover, by Wald's identity we have

$$\begin{aligned}
\mathbb{E} \left[\sum_{n=w+1}^{\tau_b} Y_n^2 \right] & = \mathbb{E} \left[\sum_{n=w+1}^{\tau_b} \mathbb{E}[Y_n^2 | \mathcal{G}_{n-1}] \right] \\
& \stackrel{(a)}{\leq} v \mathbb{E}[\tau_b] \\
& \stackrel{(b)}{\leq} K(b + \mathbb{E}[R_b] + w\mu),
\end{aligned} \tag{52}$$

where step (a) in (52) holds by assumption (33) of the theorem, step (b) by (46), and

$$K \equiv \frac{v}{\mu^*(1 - q/w)}.$$

Combining (51) and (52), we have

$$(\mathbb{E}[R_b])^2 - K(b + \mathbb{E}[R_b] + w\mu) \leq 0.$$

Solving this quadratic inequality, we conclude that

$$\begin{aligned}
\mathbb{E}[R_b] & \leq \frac{1}{2} \left(K + \sqrt{K^2 + 4K(b + w\mu)} \right) \\
& \leq K + \sqrt{Kb} + \sqrt{Kw\mu}
\end{aligned}$$

where the second inequality follows by the subadditivity of the square-root. This implies (47) with $C = \max\{K, \sqrt{K}, \sqrt{K\mu}\}$ and completes the proof.

(iii) This follows directly from (ii). ■

REFERENCES

- [1] G. Fellouris and V. V. Veeravalli, "Quickest change detection with controlled sensing," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2022, pp. 1921–1926.
- [2] H. V. Poor and O. Hadjiladis, *Quickest Detection*. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [3] A. G. Tartakovsky, I. V. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Change-Point Detection* (Statistics). Boca Raton, FL, USA: CRC Press, 2014.
- [4] V. V. Veeravalli and T. Banerjee, "Quickest change detection," in *Academic Press Library in Signal Processing: Array and Statistical Signal Processing*. Oxford, U.K.: Acad. Press, 2013.
- [5] L. Xie, S. Zou, Y. Xie, and V. V. Veeravalli, "Sequential (quickest) change detection: Classical results and new directions," *IEEE J. Select. Areas Inf. Theory*, vol. 2, no. 2, pp. 494–514, Jun. 2021.
- [6] W. Zhang and Y. Mei, "Bandit change-point detection for real-time monitoring high-dimensional data under sampling control," *Technometrics*, vol. 65, no. 1, pp. 33–43, 2023. [Online]. Available: <https://doi.org/10.1080/00401706.2022.2054861>
- [7] A. Gopalan, B. Lakshminarayanan, and V. Saligrama, "Bandit quickest changepoint detection," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, pp. 1–10.
- [8] V. Krishnamurthy, *Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing*. Cambridge, U.K.: Cambridge Univ. Press, 2016.
- [9] H. Chernoff, "Sequential design of experiments," *Ann. Math. Statist.*, vol. 30, no. 3, pp. 755–770, 1959.
- [10] S. A. Bessler, "Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments, part-I Theory," Dept. Statist., Stanford Univ., Rep. 55, 1960.
- [11] S. A. Bessler, "Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments, part II Applications," Dept. Statist., Stanford Univ., Rep. 56, 1960.
- [12] A. E. Albert, "The sequential design of experiments for infinitely many states of nature," *Ann. Math. Statist.*, vol. 32, no. 3, pp. 774–799, 1961.
- [13] J. Kiefer and J. Sacks, "Asymptotically optimum sequential inference and design," *Ann. Math. Statist.*, vol. 34, no. 3, pp. 705–750, 1963.
- [14] S. P. Lalley and G. Lorden, "A control problem arising in the sequential design of experiments," *Ann. Probab.*, vol. 14, no. 1, pp. 136–172, 1986.
- [15] R. Keener, "Second order efficiency in the sequential design of experiments," *Ann. Statist.*, vol. 12, no. 2, pp. 510–532, 1984.
- [16] S. Nitinawarat, G. Atia, and V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Trans. Aut. Control*, vol. 58, no. 10, pp. 2451–2464, Oct. 2013.
- [17] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *Ann. Statist.*, vol. 41, no. 6, pp. 2703–2738, 2013.
- [18] S. Nitinawarat and V. Veeravalli, "Controlled sensing for sequential multihypothesis testing with controlled Markovian observations and non-uniform control cost," *Sequential Anal.*, vol. 34, no. 1, pp. 1–24, Feb. 2015.
- [19] A. Deshmukh, V. V. Veeravalli, and S. Bhashyam, "Sequential controlled sensing for composite multihypothesis testing," *Sequential Anal.*, vol. 40, no. 2, pp. 1–38, 2021.
- [20] K. Cohen and Q. Zhao, "Asymptotically optimal anomaly detection via sequential testing," *IEEE Trans. Signal Process.*, vol. 63, no. 11, pp. 2929–2941, Jun. 2015.
- [21] A. Gurevich, K. Cohen, and Q. Zhao, "Sequential anomaly detection under a nonlinear system cost," *IEEE Trans. Signal Process.*, vol. 67, no. 14, pp. 3689–3703, Jul. 2019.
- [22] B. Hemo, T. Gafni, K. Cohen, and Q. Zhao, "Searching for anomalies over composite hypotheses," *IEEE Trans. Signal Process.*, vol. 68, no. 2, pp. 1181–1196, Feb. 2020.
- [23] N. K. Vaidhiyan and R. Sundaresan, "Learning to detect an odd-ball target," *IEEE Trans. Inf. Theory*, vol. 64, no. 2, pp. 831–852, Feb. 2018.
- [24] A. Tsopelakos, G. Fellouris, and V. V. Veeravalli, "Sequential anomaly detection with observation control," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2019, pp. 2389–2393.
- [25] A. Tsopelakos and G. Fellouris, "Sequential anomaly detection with observation control under a generalized error metric," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2020, pp. 1165–1170.
- [26] A. Tsopelakos and G. Fellouris, "Sequential anomaly detection with sampling constraints," *IEEE Trans. Inf. Theory*, vol. 69, no. 12, pp. 8126–8146, Dec. 2023.
- [27] V. Dragalin, "A simple and effective scanning rule for a multi-channel system," *Metrika*, vol. 43, no. 1, pp. 165–182, 1996.
- [28] Q. Xu, Y. Mei, and G. V. Moustakides, "Optimum multi-stream sequential change-point detection with sampling control," *IEEE Trans. Inf. Theory*, vol. 67, no. 11, pp. 7627–7636, Nov. 2021.
- [29] A. Chaudhuri, G. Fellouris, and A. Tajer, "Sequential change detection of a correlation structure under a sampling constraint," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2021, pp. 605–610.
- [30] Q. Xu and Y. Mei, "Asymptotic optimality theory for active quickest detection with unknown postchange parameters," *Sequential Anal.*, vol. 42, no. 2, pp. 150–181, 2023.
- [31] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Trans. Inf. Theory*, vol. 44, no. 7, pp. 2917–2929, Nov. 1998.
- [32] L. Xie, G. V. Moustakides, and Y. Xie, "Window-limited CUSUM for sequential change detection," *IEEE Trans. Inf. Theory*, vol. 69, no. 9, pp. 5990–6005, Sep. 2023.
- [33] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, no. 6, pp. 1897–1908, Dec. 1971.
- [34] P. Moulin and V. V. Veeravalli, *Statistical Inference for Engineers and Data Scientists*. Cambridge, U.K.: Cambridge Univ. Press, 2019.
- [35] G. Lorden, "On excess over the boundary," *Ann. Math. Statist.*, vol. 41, no. 2, pp. 520–527, 1970. [Online]. Available: <https://doi.org/10.1214/aoms/1177697092>