Sequential Anomaly Detection under a Nonlinear System Cost

Andrey Gurevich, Kobi Cohen, Qing Zhao

***Abstract*— We consider the problem of anomaly detection among K heterogeneous processes. At each given time, one process is probed, and the random observations follow two different distributions, depending on whether the process is normal or abnormal. Each anomalous process incurs a cost until its anomaly is identified and fixed, and the cost is a nonlinear (specifically, polynomial with degree d) function of the duration of the anomalous state. The objective is a sequential search strategy that minimizes the total expected cost incurred by all the processes during the detection process under reliability constraints. We propose a search algorithm that consists of exploration, exploitation, and sequential testing phases. We establish its asymptotic optimality, and analyze the approximation ratio and the regret under computational constraints.**

***Index Terms*—** Anomaly detection, sequential hypothesis testing, Sequential Probability Ratio Test (SPRT).

I. Introduction

Consider a system consisting of K processes. Each process (say k) may be abnormal with an a-priori probability π_k independent of other processes. The processes can represent components (such as routers and paths) in a cyber system, channels in a communication network, potential locations of targets, or sensors monitoring certain events. Each abnormal process incurs a cost until the anomaly is identified and fixed, and the cost is polynomial $c_k(t) = \sum_{i=1}^{d} a_{k,i} t^i$ of degree d with the time that the process stays anomalous. At each given time, one process is probed, and the observation follows distributions $f_k^{(0)}$ or $f_k^{(1)}$ depending on whether the process is normal or anomalous, respectively. The objective is a sequential search strategy that dynamically determines which process to probe at each time and when to terminate the search so that the total expected cost incurred to the system during the entire detection process is minimized under reliability constraints.

A. Main Results

The above anomaly detection problem has a clear connection with the classic sequential hypothesis testing problem pioneered by Wald [2]. The presence of multiple processes and the objective of minimizing the total cost (rather than the detection delay), however, give the problem another dimension. In addition to quickly declaring the state of a process by fully utilizing past observations, the probing order is crucial in minimizing the total cost. It is intuitive that processes with a higher probability of being abnormal and a higher cost when abnormal should be probed first. At the same time, it may be desirable to probe processes that require more samples to detect their states (determined by the distance between $f_k^{(0)}$ and $f_k^{(1)}$) toward the end of the detection process to avoid long delays in catching other potentially abnormal processes.

There are several studies on sequential detection under linear cost functions, including a fixed cost per sample [3]–[5], and a fixed cost due to undetected anomalies per unit time [6], [7]. Unlike these studies, in this paper we tackle a general nonlinear cost setting (for a detailed comparison with related work see Section I-C). We develop an effective algorithm, referred to as Exploration, Exploitation, and Sequential Testing (EEST), and analyze its performance in terms of regret, approximation ratio, and asymptotic optimality. Specifically, since the system cost depends on the unknown system state, the EEST algorithm aims at learning the set of abnormal processes while at the same time probing the processes in a desired order so as to minimize the estimated system cost. To ensure the accuracy of the learning process, a logarithmic order of time is used *to explore* the state of *all processes*. In the remaining majority of time instants, the EEST algorithm *exploits* the information gathered up to this point to select the processes in a desired order based on the estimated states. The algorithm implements two exploitation phases, which are referred to as the *index selection phase*, and the *constrained combinatorial search phase*. Roughly speaking, the index selection phase performs low-complexity probing (only linear with K) with a bounded approximation ratio of the cost. The constrained combinatorial search phase minimizes the cost function directly (e.g., by performing branch and bound) but only over a small set of processes (determined by the required computational constraint). Then, in the sequential testing phase, the algorithm applies a sequential test using the collected observations to determine the state of the probed process subject to constraints on type I and type II errors.

We analyze the algorithm performance theoretically, and show that the EEST algorithm is asymptptically optimal in terms of minimizing the cost as the error probability constraints approach zero. Since that achieving asymptotic optimality requires to perform a combinatorial search over

Part of this work was presented at the Allerton Conference on Communication, Control, and Computing in 2017 [1].

Andrey Gurevich and Kobi Cohen are with the Department of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Beer-Sheva 84105, Israel. Email: {gurevian, yakovsec}@bgu.ac.il

Qing Zhao is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853, USA. Email: qz16@cornell.edu

The work of A. Gurevich and K. Cohen was supported by the Cyber Security Research Center at Ben-Gurion University of the Negev, and by the U.S.-Israel Binational Science Foundation (BSF) under grant 2017723. The work of Qing Zhao was supported by the National Science Foundation under Grant CCF-1815559.

the estimate abnormal set, we also establish performance guarantees when computational constraints apply. Specifically, we analyze the approximation ratio ρ, defined as the ratio between the cost under the EEST algorithm and the optimal cost, and the regret, defined as the performance loss of EEST algorithm against a benchmark, in which the optimal probing strategy with knowledge of the set of anomalous processes is employed. We show that the approximation ratio is bounded and the regret order is $O(t^{d-1}\log t)$ with time t for $d \geq 1$. We note that asymptotic optimality of the EEST algorithm (i.e., $\rho = 1$) for $d = 1$ (i.e., linear cost function) is obtained even when computational constraints apply.

B. Applications

The anomaly detection problem considered in this paper finds applications in systems where the loss incurred by abnormal processes can be modeled by a polynomial cost with time. In the following we give two specific examples.

Consider a cyber network consisting of K components (which can be routers, servers, domains, etc.). Due to resource constraints, only a subset of the components can be probed at a time. An intrusion detection system (IDS) monitors the traffic over the components to detect Denial of Service (DoS) attacks (such attacks aim to overwhelm the component with useless traffic to make it unavailable for its intended use until it is detected and fixed). An objective of minimizing the total loss in data rate during the DoS attacks translates to a cost $c_k(t)$ given by the expected data rate (or expected number of packets) per unit time. This setting results in minimizing a linear cost function with time. An alternative, and often more relevant objective is to minimize the total loss in Quality of Experience (QoE) of users during the DoS attacks, which is commonly modeled by a polynomial function of time [8]. In this case, the cost $c_k(t)$ is given by the QoE for using component k. This setting results in minimizing a polynomial cost function.

A similar example arises in spectrum sensing in cognitive radio systems, where a cognitive user aims at finding idle channels for transmission. Due to narrowband sensing, only a subset of the channels can be probed at a time. Similar to the previous example, the objective of minimizing the total loss in QoE during the sensing process leads to minimizing a polynomial cost function.

C. Related Work

Sequential hypothesis testing was pioneered by Wald [2], where he established the Sequential Probability Ratio Test (SPRT) for binary hypothesis testing. For simple hypothesis testing where the observation distributions are known, SPRT is optimal in terms of minimizing the expected sample size under given type I and type II error probability constraints. Various extensions to M-ary hypothesis testing and testing composite hypotheses have been studied in [9]–[13] for a single process. In these cases, asymptotically optimal performance can be obtained in terms of minimizing the expected sample size as the error probability approaches zero. In this paper, we focus on asymptotically optimal strategies with low computational complexity for sequential anomaly detection over multiple processes with a limited probing capacity. Different models have considered the case of searching for anomalies (or targets) without constraints on the probing capacity, where all processes are probed at each given time (i.e., $K = M$, which is a special case of the setting considered in this paper) [12], [14]–[16].

There are several studies on sequential detection under linear cost functions, including a fixed cost per sample [3]–[5], and a fixed cost due to undetected anomalies per unit time [6], [7]. Unlike these studies, in this paper we tackle a general nonlinear cost setting. It should be noted that there are fundamental differences between this paper and our previous work [7] as detailed next. First, in terms of the exploitation phase, in contrast to the linear cost case [7], the index selection phase alone cannot achieve asymptotic optimality under nonlinear costs (which can be shown by counterexamples). Thus, a combinatorial search is needed, and is a key feature to ensure asymptotic optimality. This, however, brings a new challenge: an efficient tradeoff between performance and computational complexity. Our approach to this challenge is to carry out the computationally expensive phase of the combinatorial search only when it has the potential of bringing the most information. More specifically, the combinatorial search is carried out only when the belief about the abnormal set is sufficiently high (i.e., the posterior probabilities of all processes are above the upper threshold or below the lower threshold), and the thresholds represent the tuning knob for trading off performance with computational complexity. These design principles constitute fundamental differences in terms of algorithm design as compared to $CL - \pi cN$ algorithm [7]. The theoretical analysis is thus fundamentally different in this respect. Specifically, proving asymptotic optimality and the approximation ratio requires analysis of two different regimes. One is the asymptotic optimality obtained by the combinatorial search phase (which is absent in [7]). The second is the approximation ratio of the index selection phase when computational constraints apply, which is not asymptotically optimal in the polynomial cost case in contrast to the linear case in [7]. Second, in terms of the exploration phase, in [7], we only showed that it is possible to achieve asymptotic optimality by exploring processes at rate $\omega_k \log(t)$, but the value of the leading constant ω_k remained unknown, which results in difficulties when implementing the algorithm. This issue is solved in this paper by establishing explicit sufficient conditions for the design of the exploration phase for general polynomial functions, which has significant implications in both practical implementation and theoretical study. Finally, regret analysis was not shown in previous works even for the linear cost case.

There are a number of recent studies on sequential detection involving multiple independent processes for various applications (see for example, [17]–[26] and the references therein). Unlike this work, these studies have focused on minimizing the total detection delay, which does not translate into minimizing the total cost in the problem considered here. Other related works have considered dynamic search and outlier detection, as studied in [14], [27]–[35] (and the references therein)

under a sequential setting and in [36]–[39] (and the references therein) under a fixed sample size setting, where the objective was to locate a single target among multiple processes quickly and reliably.

The anomaly detection problem studied in this paper can be considered as a variation of the sequential design of experiments problem first studied by Chernoff [40]. In this problem, a decision maker aims to infer the state of an underlying phenomenon by sequentially choosing the experiment (thus the observation model) to be conducted at each time among a set of available experiments. Classic and more recent studies of this problem can be found in [4], [29], [41]–[43]. In [29], [40]–[43], the objective of minimizing the total detection delay was considered. In [4], the problem was extended to the case where taking an observation under a specific experiment is associated with a deterministic known cost, and the objective is to minimize the total cost during the test. However, the setting of a random and unknown cost (depending on the system state) at each given time makes the problem considered in this paper fundamentally different from the problem considered in [4].

II. System Model and Problem Formulation

Consider a system consisting of K processes, where each process may be in an abnormal state (denoted by H_1) with probability π_k, or a normal state (denoted by H_0) with probability $1-\pi_k$, independent of other processes. Let $c_k(t) = \sum_{i=1}^{d} a_{k,i} t^i$, where $0 \leq a_{k,d} < \infty$ for $k = 1, ..., K$, be a polynomial cost function of degree d, monotonically increasing with $t \geq 0$. Each abnormal process k incurs a polynomial cost $c_k(t)$ until it is tested and identified. Processes in a normal state do not incur cost. At each given time, a single observation (or a fixed batch of observations) is collected from a chosen process, and a measurement (say process k observed at time n) $y_k(n)$ is drawn independently. If process k is in a normal state, $y_k(n)$ follows density $f_k^{(0)}$; if process k is abnormal, $y_k(n)$ follows density $f_k^{(1)}$. In Section III, we examine the case where the densities $f_k^{(0)}, f_k^{(1)}$ are known. In Section IV, we extend the results to the case where the densities have unknown parameters.

Let $\phi(n) \in \{1, 2, ..., K\}$ be a selection rule, indicating which process is chosen to be tested at time n. Let $\mathbf{y}(n) = \left\{\phi(t), y_{\phi(t)}(t)\right\}_{t=1}^{n}$ be the set of all the observations and actions taken by time n. The selection rule $\phi(n)$ is a mapping from $\mathbf{y}(n-1)$ to $\{1, 2, ..., K\}$. We also define the time series vector of selection rules by $\boldsymbol{\phi} = (\phi(1), \phi(2), ...)$. Let τ_k be a stopping time, which is the time (counted from the beginning of the entire detection process) when the decision maker stops collecting observations from process k and declares its state. We also define the stopping time vector for the K processes by $\boldsymbol{\tau} = (\tau_1, ..., \tau_K)$. Let N_k be the random sample size collected from process k until declaring its state. Let $\delta_k \in \{0, 1\}$ be a decision rule for process k at time τ_k. $\delta_k = 0$ if the decision maker declares that process k is in a normal state, and $\delta_k = 1$ if the decision maker declares that process k is in an abnormal state. We define the vector of decision rules for the K processes by $\boldsymbol{\delta} = (\delta_1, ..., \delta_K)$. An admissible strategy $\mathbf{s}$ for the anomaly detection problem is given by the tuple $\mathbf{s} = (\boldsymbol{\tau}, \boldsymbol{\delta}, \boldsymbol{\phi})$.

Let

$$\mathcal{H}_0 \triangleq \{k : 1 \leq k \leq K \text{ , process } k \text{ is normal}\} \text{ ,}$$

$$\mathcal{H}_1 \triangleq \{k : 1 \leq k \leq K \text{ , process } k \text{ is abnormal}\} \text{ ,}$$

be the sets of the normal and abnormal processes, respectively. The objective is to find a strategy $\mathbf{s}$ that minimizes the total expected cost incurred by all the processes subject to type I (false-alarm) and type II (miss-detection) error constraints for each process:

$$\begin{aligned} \inf_{\mathbf{s}} \quad & \mathbf{E}\left\{\sum_{k\in\mathcal{H}_1} c_k\left(\tau_k\right)\right\} \\ s.t. \quad & P_k^{FA} \leq \alpha_k \qquad \forall k = 1, ..., K, \\ & P_k^{MD} \leq \beta_k \qquad \forall k = 1, ..., K \text{ ,} \end{aligned} \tag{1}$$

where P_k^{FA}, P_k^{MD} denote the false-alarm and miss-detect error probabilities for process k, respectively.

III. The Exploration, Exploitation, and Sequential Testing (EEST) Algorithm

Sequential detection problems involving multiple processes are partially-observed Markov decision processes (POMDP) [37] which have exponential complexity in general. For tractability, a commonly adopted performance measure is asymptotic optimality in terms of minimizing the objective function as the error probability approaches zero (see, for example, classic and recent results in [4], [9], [10], [15], [16], [29], [30], [40]–[42]). Hence, we are interested in developing a low-complexity algorithm that achieves asymptotically optimal performance in terms of minimizing (1) as the error constraints approach zero.

We first provide notations and definitions that will be used to describe the EEST algorithm. Let $\mathbf{1}_k(n)$ be the probing indicator function, where $\mathbf{1}_k(n) = 1$ if process k is probed at time n and $\mathbf{1}_k(n) = 0$ otherwise. Let

$$\ell_k(n) \triangleq \log \frac{f_k^{(1)}(y_k(n))}{f_k^{(0)}(y_k(n))} \text{ ,} \tag{2}$$

and

$$S_k(n) \triangleq \sum_{t=1}^{n} \ell_k(t) \mathbf{1}_k(t) \tag{3}$$

be the log-likelihood ratio (LLR) and the observed sum LLRs at time n of process k, respectively. Let $\mathcal{K}(n)$ be the set of processes whose states have not been declared up to time n. Let $\pi_k(n)$ denote the posterior probability of process k being abnormal at time n. Let $\mathbf{E}^{(n)}(N_k)$ be the expected sample size required for declaring the state of process k observed by time n (which dynamically changes due to the changes in the belief $\pi_k(n)$), where N_k is the random sample size required to declare its state. Define

$$\gamma_k(n) \triangleq \frac{\pi_k(n) a_{k,d}}{\mathbf{E}^{(n)}(N_k)} \text{ ,} \tag{4}$$

where $a_{k,d}$ is the dth-degree coefficient of the polynomial function $c_k(t)$ as defined in Section II. Let L_k, U_k be thresholds of process k, where $0 < L_k < U_k < 1$ (the setting of L_k, U_k is discussed later). Let $\mathcal{H}_1(n)$ be the set of processes in $\mathcal{K}(n)$ that satisfy $\pi_k(n) > U_k$ at time n with cardinality $|\mathcal{H}_1(n)| = N_{H_1}(n)$. The set $\mathcal{H}_1(n)$ can be viewed as the estimate of the set $\mathcal{H}_1$ at time n. Let $\mathbf{E}(N_k|H_1)$ be the expected sample size required to declare the state of process k under hypothesis H_1 (note that $\mathbf{E}(N_k|H_1)$ is fixed over time). The expected sample size $\mathbf{E}(N_k|H_1)$ is determined by the SPRT described later so that the error constraints are satisfied. Finally, let $N_C \in \{1, ..., N\}$ be a computational constraint value.

A. The Algorithm

Next, we describe the EEST algorithm with respect to the time index n. The algorithm parameters are described in Section III-B. Since processes whose states have been declared are no longer tested, we refer only to processes $k \in \mathcal{K}(n)$:

1) (Exploration phase:) If the number of observations that have been collected from some process $k \in \mathcal{K}(n)$ satisfies:
$$N_k(n) < \omega_k \log n,$$
then $\phi(n) = k$ and go to Step 1 again. Otherwise, go to Step 2.

2) If $N_{H_1}(n) > N_C$ or $N_{H_1}(n) = 0$ or there exists a process $k \in \mathcal{K}(n)$ that satisfies $L_k \leq \pi_k(n) \leq U_k$, then go to Step 3. Otherwise, go to Step 4.

3) (Index selection phase:) Probe the process with the highest index:
$$\phi(n) = \arg \max_{k \in \mathcal{K}(n)} \gamma_k(n),$$
and go to Step 5.

4) (Constrained combinatorial search phase:) Let
$$\boldsymbol{\sigma}(n) = (\sigma_1(n), ..., \sigma_{N_{H_1}(n)}(n))$$
be a permutation of $\mathcal{H}_1(n)$. Probe the next process $\phi(n)$ according to the permutation order that minimizes:
$$\sum_{i=1}^{N_{H_1}(n)} c_{\sigma_i(n)} \left(\sum_{j=\sigma_1(n)}^{\sigma_i(n)} \mathbf{E}(N_j|H_1) \right).$$
Go to Step 5.

5) Update $S_{\phi(n)}(n)$ based on the last observation. Following Wald's SPRT [2], $S_{\phi(n)}(n)$ is compared to boundary values $A_{\phi(n)}, B_{\phi(n)}$ as follows: If
$$S_{\phi(n)}(n) \in \big(A_{\phi(n)}, B_{\phi(n)}\big),$$
then $\phi(n) \in \mathcal{K}(n+1)$. Otherwise, if
$$S_{\phi(n)}(n) \geq B_{\phi(n)},$$
then stop taking observations from process $\phi(n)$ and declare it as abnormal (i.e., $\tau_{\phi(n)} = n$, $\delta_{\phi(n)} = 1$ and $\phi(n) \notin \mathcal{K}(n')$ for all $n' > n$). Otherwise, if
$$S_{\phi(n)}(n) \leq A_{\phi(n)},$$
then stop taking observations from process $\phi(n)$ and declare it as normal (i.e., $\tau_{\phi(n)} = n$, $\delta_{\phi(n)} = 0$ and $\phi(n) \notin \mathcal{K}(n')$ for all $n' > n$). Go to step 1.

The EEST algorithm can be intuitively explained as follows. During the index selection phase, the index $\gamma_k(n)$ gives a closed-form expression of how three key parameters—the cost coefficient $a_{k,d}$ of the dth degree term of the polynomial function that dominates the asymptotic cost, the posterior probability, and the difficulty in distinguishing the normal distribution $f_k^{(0)}$ from the abnormal distribution $f_k^{(1)}$—are traded off when choosing the observed process at each given time n. Sorting the indices can be done in $O(K \log K)$ time with a sorting algorithm. The algorithm moves to the constrained combinatorial search phase when the abnormal set $\mathcal{H}_1$ has been estimated with sufficient reliability, and $N_{H_1}(n)$ is sufficiently small (but not zero). This phase applies an exhaustive search over N_C processes at most, and more efficient algorithms such as branch and bound can be applied for small d. As a result, N_C applies a constraint on the combinatorial search used to guarantee the desired computational complexity. When the a priori probabilities for processes of being abnormal are small (so that $|\mathcal{H}_1| << K$ as in many anomaly detection applications), the constrained combinatorial search phase ensures an approximation ratio close to 1. The thresholds L_k, U_k are tuning parameters that trade off computational complexity with detection performance in the finite-sample regime. Typically, we set L_k, U_k close to $0, 1$, respectively. Note that setting L_k, U_k too close to 0 and 1 results in a longer index selection phase, which reduces complexity but also deteriorates performance in the finite-sample regime. On the other hand, setting L_k, U_k far from 0 and 1 results in a longer constrained combinatorial search, which improves performance, but increases the computational complexity in the finite-sample regime. We point out that switching between the index selection phase and constrained combinatorial search used to trade-off between performance and complexity, is not needed in the case of linear cost. The combinatorial search phase is used to guarantee asymptotically optimal performance, since the index selection phase is not asymptotically optimal in the case of nonlinear cost. The conditions $N_k(n) \geq \omega_k \log n$ for all k ensure that a logarithmic order of time is used to collect observations from all the processes (i.e., *exploration phase*). Setting ω_k judiciously (as detailed in Section III-C) guarantees a sufficiently accurate estimate of the abnormal set $\mathcal{H}_1$ used for achieving the theoretical performance. The setting of ω_k (see (11)) depends on the degree d of the polynomial cost function. The higher d the higher the required exploration rate.

B. Implementation

Implementing the EEST algorithm requires computing the posterior probability $\pi_k(n)$, the boundary values A_k, B_k, and the expected detection time $\mathbf{E}(N_k|H_i)$. The posterior probability $\pi_k(n)$ is updated according to Bayes' rule which is given by

$$\pi_k(n+1) = (1 - \mathbf{1}_k(n))\, \pi_k(n) + \frac{\mathbf{1}_k(n) \pi_k(n) f_k^{(1)}(y_k(n))}{\pi_k(n) f_k^{(1)}(y_k(n)) + (1 - \pi_k(n))\, f_k^{(0)}(y_k(n))}. \quad (5)$$

The boundary values A_k and B_k are set such that the error constraints are satisfied. In general, computing the exact boundary values is very laborious under the finite sample regime. Nevertheless, Wald's approximation can be applied to simplify the computation [2]:

$$\begin{aligned} A_k &\approx \log\left(\frac{\beta_k}{1-\alpha_k}\right) , \\ B_k &\approx \log\left(\frac{1-\beta_k}{\alpha_k}\right) . \end{aligned} \tag{6}$$

Wald's approximation performs well for small α_k, β_k and is asymptotically optimal as the error probability approaches zero. Since type I and type II errors are typically required to be small, Wald's approximation is widely used in practice [2].

The expected sample size $\mathbf{E}^{(n)}(N_k)$ at time n depends on the current belief:

$$\mathbf{E}^{(n)}(N_k) = \pi_k(n)\mathbf{E}(N_k|H_1) + (1-\pi_k(n))\mathbf{E}(N_k|H_0) , \tag{7}$$

where $\mathbf{E}(N_k|H_i)$ is the expected sample size for process k conditioned on its state H_i. In general, it is difficult to obtain a closed-form expression for $\mathbf{E}^{(n)}(N_k|H_i)$ under the finite sample regime. However, Wald's approximation can be applied to simplify the computation [2]:

$$\begin{aligned} \hat{\mathbf{E}}(N_k|H_0) &\triangleq \frac{-(1-\alpha_k)A_k - \alpha_k B_k}{D(f_k^{(0)}||f_k^{(1)})} , \\ \hat{\mathbf{E}}(N_k|H_1) &\triangleq \frac{(1-\beta_k)B_k + \beta_k A_k}{D(f_k^{(1)}||f_k^{(0)})} , \end{aligned} \tag{8}$$

where $D(f_k^{(i)}||f_k^{(j)}) = \mathbf{E}_i\left(\log\frac{f_k^{(i)}(y_k(1))}{f_k^{(j)}(y_k(1))}\right)$ denotes the Kullback-Leibler (KL) divergence between the hypotheses H_i and H_j, where the expectation is with respect to $f_k^{(i)}$. This approximation approaches the exact expected sample size for small α_k, β_k. As a result, the approximation to the expected sample size is computed by:

$$\hat{\mathbf{E}}^{(n)}(N_k) = \pi_k(n)\hat{\mathbf{E}}(N_k|H_1) + (1-\pi_k(n))\hat{\mathbf{E}}(N_k|H_0) . \tag{9}$$

C. Performance Analysis

Next, we establish the asymptotic approximation ratio of the EEST algorithm. Let $P_C = Pr\left(|\mathcal{H}_1| > N_C\right)$, and let

$$P_e^{max} \triangleq \max\left(\alpha_1, \beta_1, ..., \alpha_K, \beta_K\right) . \tag{10}$$

Set

$$\omega_k > \frac{d+1}{\min\left\{I_{k|H_0}, I_{k|H_1}\right\}} \quad \forall k, \tag{11}$$

where

$$\begin{aligned} I_{k|H_0} &\triangleq \sup_{s>0}\left[-\log\mathbf{E}_{\sim f_k^{(0)}}\left(e^{s\ell_k}\right)\right] , \\ I_{k|H_1} &\triangleq \sup_{s>0}\left[-\log\mathbf{E}_{\sim f_k^{(1)}}\left(e^{-s\ell_k}\right)\right] \end{aligned} \tag{12}$$

are the Legendre-Fenchel transformation of ℓ_k at the origin with respect to distributions $f_k^{(0)}$ and $f_k^{(1)}$, respectively. In the analysis we show that setting the exploration rates according to (11) guarantees a sufficiently accurate estimate of the abnormal set (see the proof of Lemma 2 in the appendix). The following theorem establishes the asymptotic approximation ratio of the EEST algorithm.

Theorem 1. *Let $\boldsymbol{E}(C^{EEST})$ be the expected cost under the EEST algorithm, and $\inf_{\mathbf{s}} \boldsymbol{E}(C(\mathbf{s}))$ be the infimum expected cost over all algorithms that satisfy the error constraints. Then,*

a) (Asymptotic optimality of the EEST algorithm:) Set $N_C = K$ (i.e., computational constraints do not apply). Then[1]*,*

$$\boldsymbol{E}(C^{EEST}) \sim \inf_{\mathbf{s}} \boldsymbol{E}(C(\mathbf{s})) \quad as \quad P_e^{max} \to 0. \tag{13}$$

b) (Approximation ratio of the EEST algorithm under computational constraints:) Let

$$\rho \triangleq \frac{\boldsymbol{E}(C^{EEST})}{\inf_{\mathbf{s}} \boldsymbol{E}(C(\mathbf{s}))} \tag{14}$$

be the approximation ratio between the expected cost under the EEST algorithm, $\boldsymbol{E}(C^{EEST})$, and the infimum expected cost, $\inf_{\mathbf{s}} \boldsymbol{E}(C(\mathbf{s}))$, over all algorithms that satisfy the error constraints. Then,

$$\rho \leq (1-P_C) + P_C M(d) \quad as \quad P_e^{max} \to 0 , \tag{15}$$

where

$$M(d) \triangleq \max_{0\leq s\leq 1} \frac{(d+1)s + (1-s)^{d+1}}{1+ds^{d+1}}. \tag{16}$$

A detailed proof and regularity conditions are given in Appendix VII-A.

Remark 1. The function $M(d)$ equals 1 at $d = 1$ and increases with d. For instance, $M(1) = 1, M(2) \approx 1.3, M(3) \approx 1.8, M(4) \approx 2.3, M(5) \approx 3$. As a result, for the case of a linear cost with time (i.e., $d = 1$) the algorithm is asymptotically optimal, and good theoretical performance is guaranteed for small d. Simulation results demonstrate good performance for large d as well. Furthermore, when the probabilities of processes being abnormal are small or when the computational constraint is loose, then ρ is close to 1 since P_C is close to zero.

Next, we investigate how fast non-dominant terms vanish with time. Thus, we define the *regret* as the difference between the cost under the EEST algorithm and the cost obtained by a genie that has an accurate estimate of the abnormal set $\mathcal{H}_1(n) = \mathcal{H}_1$ for all n. Thus, the genie's beliefs are set to $\pi_i(n) = 0$ for all $i \in \mathcal{H}_0$, and $\pi_j(n) = 1$ for all $j \in \mathcal{H}_1$. Note that although the genie knows the abnormal set $\mathcal{H}_1$ accurately, it still performs the detection process so that we can measure how fast our algorithm learns the unknown side information that the genie has. Therefore, Steps $1, 2$ are replaced by the following step when applying the genie's strategy: Set $\mathcal{H}_1(n) = \mathcal{H}_1$. If $N_{H_1}(n) > N_C$, go to Step 3. Otherwise, go to Step 4. Note that we can set $N_C = N$ as a special case when comparing the algorithm with an optimal probing without computational constraints. Let C^* be the cost

[1]The notation $f \sim g$ as $P_e^{max} \to 0$ refers to $\lim_{P_e^{max}\to 0} f/g = 1$.

obtained under the genie's strategy, and C^{EEST} be the cost under the EEST algorithm.

Theorem 2. Let

$$R \triangleq \mathbf{E}\left[C^{EEST} - C^*\right] \tag{17}$$

be the regret of the EEST algorithm with respect to the genie's strategy, such that the error constraints are satisfied. Let $n = \max(\tau_1, ..., \tau_K)$ be the random termination time of the genie's strategy (which is a stopping time). Then,

$$R = O\left(\mathbf{E}[n^{d-1}\log n]\right) \text{ as } P_e^{max} \to 0\ . \tag{18}$$

The proof is given in Appendix VII-B.

Remark 2. Theorem 2 implies that the additional cost obtained under the EEST algorithm as compared to the genie's strategy is of order $O\left(\mathbf{E}[n^{d-1}\log n]\right)$. Since the cost under the genie's strategy is of order $O\left(\mathbf{E}[n^d]\right)$, the additional cost under the EEST algorithm vanishes with time. When $d = 1$, a logarithmic regret is obtained.

IV. Anomaly Detection Under Unknown Observation Models

In this section we consider the case where the observation distributions are governed by unknown parameters. Specifically, let θ_k be the unknown parameter (vector) of process k. The observation $y_k(n)$ follows density $f_k(y|\theta_k)$, $\theta_k \in \Theta_k$, where Θ_k is the parameter space of process k. If process k is in a normal state, then $\theta_k \in \Theta_k^{(0)}$; if process k is in an abnormal state, then $\theta_k \in (\Theta_k \backslash \Theta_k^{(0)})$. Let $\Theta_k^{(0)}$, $\Theta_k^{(1)}$ be disjoint subsets of Θ_k, where $R_k = \Theta_k \backslash (\Theta_k^{(0)} \cup \Theta_k^{(1)}) \neq \emptyset$ is an indifference region[2]. When $\theta_k \in R_k$, the detector is indifferent to the state of process k. As a result, there are no constraints on the error probabilities for all $\theta_k \in R_k$. The essence of the sequential testing step for process k is thus testing $\theta_k \in \Theta_k^{(0)}$ against $\theta_k \in \Theta_k^{(1)}$. Reducing R_k increases the sample size.

Although the SPRT algorithm used in the sequential testing step in the previous section under completely known distributions is optimal in terms of minimizing the required sample size of each process, it is highly sub-optimal when the distributions have unknown parameters. Therefore, our approach in this section is to incorporate an asymptotically optimal test (in terms of minimizing the sample size as the error probability approaches zero) in the sequential testing step. There are a number of asymptotically optimal sequential tests for a single process, where the basic idea is to use the maximum likelihood estimate (MLE) of the unknown parameters to perform a one-sided sequential test to reject H_0 and a one-sided sequential test to reject H_1. One way is to use the Adaptive Likelihood Ratio (ALR) statistics, which were first introduced by Robbins and Siegmund in [44]. Specifically, For $i, j \in \{0, 1\}$ and $i \neq j$, let

$$S_k^{(i)}(n) = \sum_{r=1}^{n} \log \frac{f_k(y_k(r)|\hat{\theta}_k(r-1))}{f_k(y_k(r)|\hat{\theta}_k^{(j)}(n))} \tag{19}$$

be the ALR statistics used to declare hypothesis H_i at stage n, where $\hat{\theta}_k(t) = \arg\max_{\theta_k \in \Theta_k} f_k\left(\mathbf{y}_k(t)|\theta_k\right)$ and $\hat{\theta}_k^{(j)}(t) = \arg\max_{\theta_k \in \Theta_k^{(j)}} f_k\left(\mathbf{y}_k(t)|\theta_k\right)$ are the maximum likelihood (ML) estimates of the parameters over the parameter spaces Θ_k and $\Theta_k^{(j)}$ at stage t, respectively. Let

$$N_k^{(i)} = \inf\left\{\ n\ : S_k^{(i)}(n) \geq B_k^{(i)}\right\} \tag{20}$$

be the stopping rule for declaring H_i, where $B_k^{(i)}$ is the boundary value. For each process k, the decision maker stops the sampling when $N_k = \min\left\{N_k^{(0)}, N_k^{(1)}\right\}$. If $N_k = N_k^{(0)}$, process k is declared as normal. If $N_k = N_k^{(1)}$, process k is declared as abnormal. We point out that another way is to use the Generalized Likelihood Ratio (GLR) statistics [10], which is able to improve performance by updating the estimate in the numerator for all n. Nevertheless, here we adopt the ALR statistics, since setting

$$\begin{aligned} B_k^{(0)} &= \log\frac{1}{\alpha_k}, \\ B_k^{(1)} &= \log\frac{1}{\beta_k}, \end{aligned} \tag{21}$$

satisfies the error probability constraints in (1), while such a simple setting cannot be applied when using the GLR statistics.

A. The algorithm

We modify the EEST algorithm so that it can apply to the case of the unknown observation model considered in this section. We replace the belief $\pi_k(n+1)$, and the index $\gamma_k(n)$ by their estimates $\hat{\pi}_k(n+1)$, and $\hat{\gamma}_k(n)$, respectively (see the explicit expressions in Section IV-B). We define $\hat{\mathcal{H}}_1(n)$ as the set of processes in $\mathcal{K}(n)$ that satisfy $\hat{\pi}_k(n) > U_k$ at time n with cardinality $|\hat{\mathcal{H}}_1(n)| = \hat{N}_{\hat{H}_1}(n)$. The set $\hat{\mathcal{H}}_1(n)$ can be viewed as the estimate of the set $\mathcal{H}_1$ at time n. Since processes whose states have been declared are no longer tested, we refer only to processes $k \in \mathcal{K}(n)$:

1) (Exploration phase:) If the number of observations that have been collected from some process $k \in \mathcal{K}(n)$ satisfies:
$$N_k(n) < \omega_k \log n,$$
then $\phi(n) = k$ and go to Step 1 again. Otherwise, go to Step 2.

2) If $\hat{N}_{\hat{H}_1}(n) > N_C$ or $\hat{N}_{\hat{H}_1}(n) = 0$ or there exists a process $k \in \mathcal{K}(n)$ that satisfies $L_k \leq \hat{\pi}_k(n) \leq U_k$, then go to Step 3. Otherwise, go to Step 4.

3) (Index selection phase:) Probe the process with the highest index:
$$\phi(n) = \arg\max_{k \in \mathcal{K}(n)}\ \hat{\gamma}_k(n),$$
and go to Step 5.

4) (Constrained combinatorial search phase:) Let
$$\boldsymbol{\sigma}(n) = (\sigma_1(n), ..., \sigma_{\hat{N}_{\hat{H}_1}(n)}(n))$$
be a permutation of $\hat{N}_{H_1}(n)$. Probe the next process $\phi(n)$ according to the permutation order that minimizes:

[2]We adopted the assumption of an indifference region as commonly used in the theory of sequential hypothesis testing for establishing asymptotic optimality. Nevertheless, in some cases this assumption can be relaxed as analyzed in [10].

$$\sum_{i=1}^{\hat{N}_{\hat{H}_1}(n)} c_{\sigma_i(n)} \left(\sum_{j=\sigma_1(n)}^{\sigma_i(n)} \hat{\mathbf{E}}(N_j) \right).$$

Go to Step 5.

5) Update $S^{(i)}_{\phi(n)}(n)$ for $i = 0, 1$ based on the last observation. Then, $S^{(i)}_{\phi(n)}(n)$ is compared to boundary values $B^{(0)}_{\phi(n)}, B^{(1)}_{\phi(n)}$ as follows:
If

$$S^{(0)}_k(n) < B^{(0)}_k \text{ and } S^{(1)}_k(n) < B^{(1)}_k,$$

then $\phi(n) \in \mathcal{K}(n+1)$. Otherwise, if

$$S^{(1)}_k(n) \geq B^{(1)}_k,$$

stop taking observations from process k and declare it as abnormal (i.e., $\tau_{\phi(n)} = n$, $\delta_{\phi(n)} = 1$ and $\phi(n) \notin \mathcal{K}(n')$ for all $n' > n$). Otherwise, if

$$S^{(0)}_k(n) \geq B^{(0)}_k,$$

stop taking observations from process k and declare it as normal (i.e., $\tau_{\phi(n)} = n$, $\delta_{\phi(n)} = 0$ and $\phi(n) \notin \mathcal{K}(n')$ for all $n' > n$). Go to step 1.

B. Implementation

We now detail the parameters used to implement the modified EEST algorithm. The modified belief is given by:

$$\hat{\pi}_k(n+1) = (1 - \mathbf{1}_k(n))\, \hat{\pi}_k(n) + \frac{\mathbf{1}_k(n) \hat{\pi}_k(n) \hat{f}^{(1)}_k(y_k(n))}{\hat{\pi}_k(n) \hat{f}^{(1)}_k(y_k(n)) + (1 - \hat{\pi}_k(n))\, \hat{f}^{(0)}_k(y_k(n))}, \tag{22}$$

where $\hat{\pi}_k(1) = \pi_k(1)$ and $\hat{f}^{(1)}_k(y_k(r)) \triangleq f_k(y_k(r)|\hat{\theta}^{(1)}_k(n))$, $\hat{f}^{(0)}_k(y_k(r)) \triangleq f_k(y_k(r)|\hat{\theta}^{(0)}_k(n))$ for all $1 \leq r \leq n$.

The modified index $\hat{\gamma}_k(n)$ is given by:

$$\hat{\gamma}_k(n) \triangleq \frac{\hat{\pi}_k(n) a_{k,d}}{\hat{\mathbf{E}}^{(n)}(N_k)}\,. \tag{23}$$

In general, it is difficult to obtain a closed-form expression for $\hat{\mathbf{E}}^{(n)}(N_k)$ under the finite sample regime. However, we can use the asymptotic property of the sequential tests to obtain a closed-form approximation for $\hat{\mathbf{E}}^{(n)}(N_k)$ based on the ML estimate of the parameter, which approaches the exact expected sample size as the error probability approaches zero. Let $D_k(\hat{\theta}_k(n)||\theta) \triangleq \mathbf{E}_{\hat{\theta}_k(n)} \left(\log \frac{f_k(y_k(n)|\hat{\theta}_k(n))}{f_k(y_k(n)|\theta)} \right)$ be the KL divergence between $f_k(y_k(n)|\hat{\theta}_k(n))$ and $f_k(y_k(n)|\theta)$, where the expectation is taken with respect to $f_k(y_k(n)|\hat{\theta}_k(n))$, and let $D_k(\hat{\theta}_k(n)||\Theta^{(i)}_k) = \inf_{\theta \in \Theta^{(i)}_k} D_k(\hat{\theta}_k(n)||\theta)$. Then, the estimated expected sample size required to make a decision regarding the state of process k is given by:

$$\hat{\mathbf{E}}^{(n)}(N_k) = \begin{cases} \dfrac{B^{(0)}_k}{D_k\left(\hat{\theta}_k(n)||\Theta^{(1)}_k\right)} & , \text{ if } \hat{\theta}_k(n)) \in \Theta^{(0)}_k\,, \\ \dfrac{B^{(1)}_k}{D_k\left(\hat{\theta}_k(n)||\Theta^{(0)}_k\right)} & , \text{ if } \hat{\theta}_k(n)) \in \Theta^{(1)}_k\,, \end{cases} \tag{24}$$

which is guaranteed to be the asymptotic sample size under various families of distributions with unknown parameters (e.g., exponential, multivariate distributions, and general distributions with unknown parameters that take a finite number of values) as the error probabilities approach zero [10]–[12], [40].

C. Performance Analysis

We next establish the approximation ratio and the regret under the case where Θ_k is discrete for all $k = 1, ..., K$. Set

$$\omega_k > \frac{d+1}{\min\left\{\tilde{I}_{k,0}, \tilde{I}_{k,1}\right\}} \;\; \forall k, \tag{25}$$

where

$$\tilde{I}_{k,0} \triangleq \inf_{\substack{\theta^{(0)} \neq \theta^{(1)}, \\ \theta^{(0)}, \theta^{(1)} \in \Theta_k}} \sup_{s>0} \left[-\log \mathbf{E}_{\sim f_k(y|\theta^{(0)})} \left(e^{s\ell_k(\theta^{(0)}, \theta^{(1)})} \right) \right],$$

$$\tilde{I}_{k,1} \triangleq \inf_{\substack{\theta^{(0)} \neq \theta^{(1)}, \\ \theta^{(0)}, \theta^{(1)} \in \Theta_k}} \sup_{s>0} \left[-\log \mathbf{E}_{\sim f_k(y|\theta^{(1)})} \left(e^{-s\ell_k(\theta^{(0)}, \theta^{(1)})} \right) \right], \tag{26}$$

are the Legendre-Fenchel transformations of

$$\ell_k(\theta^{(0)}, \theta^{(1)}) \triangleq \log \frac{f_k(y|\theta^{(1)})}{f_k(y|\theta^{(0)})}$$

at the origin with respect to distributions $f_k(y|\theta^{(0)})$ and $f_k(y|\theta^{(1)})$, respectively. In the analysis we show that setting the exploration rates according to (26) guarantees a sufficiently accurate estimate of the abnormal set (see the proof of Lemma 3 in the appendix). The following theorem establishes the asymptotic approximation ratio of the EEST algorithm for the unknown observation model.

Theorem 3. Consider the unknown observation model as detailed in this section. Let $\mathbf{E}(C^{EEST})$ be the expected cost under the EEST algorithm, and $\inf_{\mathbf{s}}\ \mathbf{E}(C(\mathbf{s}))$ be the infimum expected cost over all algorithms that satisfy the error constraints. Then,

a) (Asymptotic optimality of the EEST algorithm:) Set $N_C = K$ (i.e., computational constraints do not apply). Then,

$$\mathbf{E}(C^{EEST}) \sim \inf_{\mathbf{s}}\ \mathbf{E}(C(\mathbf{s})) \quad \text{as} \quad P_e^{max} \to 0. \tag{27}$$

b) (Approximation ratio of the EEST algorithm under computational constraints:) Let

$$\rho \triangleq \frac{\mathbf{E}(C^{EEST})}{\inf_{\mathbf{s}}\ \mathbf{E}(C(\mathbf{s}))} \tag{28}$$

be the approximation ratio between the expected cost under the modified EEST algorithm $\mathbf{E}(C^{EEST})$, described in Section IV-A, and the infimum expected cost $\inf_{\mathbf{s}}\ \mathbf{E}(C(\mathbf{s}))$ over all algorithms that satisfy the error constraints. Then,

$$\rho \leq (1 - P_C) + P_C M(d) \;\text{ as }\; P_e^{max} \to 0\,, \tag{29}$$

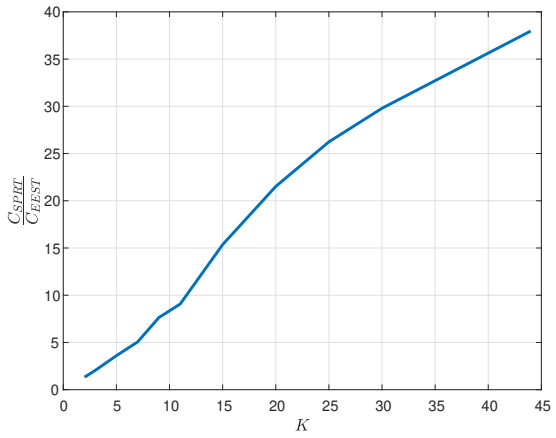

Fig. 1. The cost ratio between the EEST algorithm and the open loop SPRT as a function of the number of processes. The case of processes that follow a Poisson distribution.

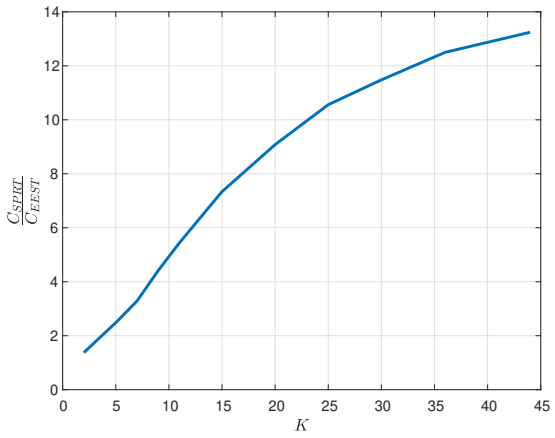

Fig. 2. The cost ratio between the EEST algorithm and the open loop SPRT as a function of the number of processes. The case of processes that follow an exponential distribution.

where $M(d) = \max\limits_{0\le s\le 1} \dfrac{(d+1)s + (1-s)^{d+1}}{1 + ds^{d+1}}$.

A detailed proof and regularity conditions are given in Appendix VII-C.

Next, we investigate how fast non-dominant terms vanish with time. Thus, we define the *regret* as the difference between the cost under the EEST algorithm and the cost obtained by a genie that has an accurate estimate of the abnormal set $\mathcal{H}_1(n) = \mathcal{H}_1$ for all n, and the true parameter θ. Thus, genie's beliefs are set to $\pi_i(n) = 0$ for all $i \in \mathcal{H}_0$, and $\pi_j(n) = 1$ for all $j \in \mathcal{H}_1$, and the estimate of θ is set to the true value. Note that although genie knows the abnormal set $\mathcal{H}_1$ and the parameter θ accurately, it still performs the same sequential testing step so that we can measure how fast our algorithm learns the unknown side information that genie has. Let C^* be the cost obtained under genie's strategy, and C^{EEST} be the cost under the modified EEST algorithm in this setting.

Theorem 4. Consider the unknown observation model as detailed in this section. Let

$$R \triangleq \mathbf{E}\left[C^{EEST} - C^*\right] \tag{30}$$

be the regret of the modified EEST algorithm, described in Section IV-A, with respect to genie's strategy, such that the error constraints are satisfied. Let $n = \max(\tau_1, ..., \tau_K)$ be the termination time of genie's strategy (which is a stopping time). Then,

$$R = O\left(\mathbf{E}[n^{d-1}\log n]\right) \text{ as } P_e^{max} \to 0\,. \tag{31}$$

The proof is given in Appendix VII-D.

V. Simulation Results

In this section we present simulation results to demonstrate the effectiveness of the proposed EEST algorithm. We start by demonstrating the significance of using the closed-loop selection rule of the proposed EEST algorithm. We compare the EEST algorithm to an algorithm that does not use memory (i.e., open loop) when selecting the next process (e.g., as in [14], [19]–[23]). Specifically, we simulated the algorithm in [19], which performs the SPRT algorithm in a predetermined order, which was designed to the case of independent anomalies among processes as considered in this paper, and was shown to be optimal in terms of minimizing the detection delay. We refer to this algorithm as the open-loop SPRT. It is well-known that there is often a gap between the sufficient conditions required by the theoretical asymptotic analysis (often due to union bounding events in analysis) and practical conditions used for efficient sequential statistical inference. Thus, for a practical implementation we set the exploration rate so that the total exploration time was approximately 5% of the total expected detection time. We first simulated a system with K processes, each (say k) of which follows a Poisson distribution $y_k(n) \sim$Pois$(\theta_k^{(i)})$. For the normal processes we set $\theta_k^{(0)} = 1$ for all k, and for the abnormal processes we set $\theta_k^{(1)}$ to be equally spaced in the interval $[1.2, 1.4]$. The cost for each abnormal process was set to $c_k(t) = a_k(t^3 + t)$. The coefficients a_k were equally spaced in the interval $[1, 4]$. The error probability constraints were set to $\alpha_k = \beta_k = 10^{-8}$ for each process $k \in \mathcal{K}$. Fig. 1 presents the tremendous gain observed for the EEST algorithm compared to the open-loop SPRT in terms of the system cost. We next changed the observation distributions so that processes (say k) followed an exponential distribution $y_k(n) \sim$ Exp$(\theta_k^{(i)})$. The cost for each abnormal process was set to $c_k(t) = a_k(t^4 + 3t^2)$. All other parameters remained the same. Fig. 2 once again presents the tremendous gain observed for the EEST algorithm compared to the open-loop SPRT in terms of the system cost. These results demonstrate the importance of using the closed-loop policy of the EEST algorithm when minimizing the system cost, since the open-loop strategy, though optimal in terms of minimizing the delay, is highly sub-optimal in terms of minimizing the system cost. Next, we examined the regret of the EEST algorithm with respect to the genie's policy, as described in Section III-C. The (normalized) regret is presented in Fig. 3 and supports the theoretical analysis in Theorem 2.

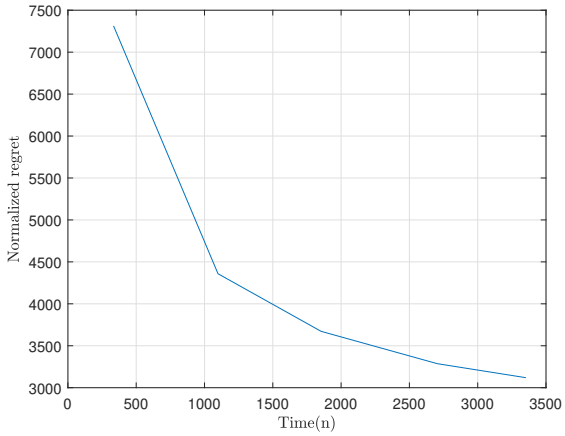

Fig. 3. The normalized regret $R/E[n^2\log(n)]$ with respect to the genie's policy

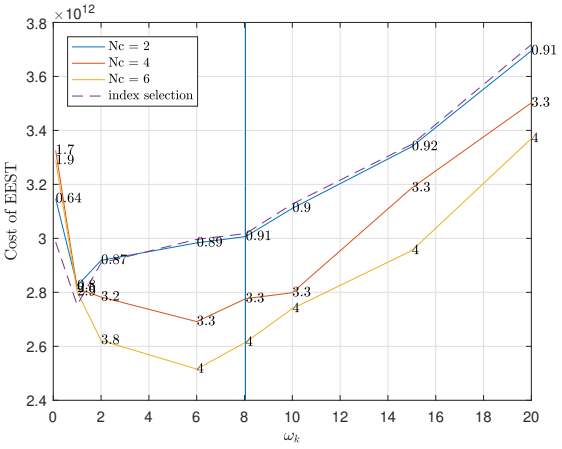

Fig. 4. The effect of N_c and ω_k on the total cost under the EEST algorithm. The vertical line represents the ω_k which is set according to the theoretical sufficient conditions in (11). The numbers along each point represent the number of times the combinatorial search was performed.

Finally, we examined the effect of the exploration rate ω_k and the computational constraint N_c on the algorithm's performance. Clearly, increasing N_c results in a higher probability to perform the combinatorial search phase. As a result, the algorithm performance is expected to be improved. Indeed, Figure 4 supports this insight. It can be seen that the system cost decreases as N_c increases with the price of a higher computational complexity. Next, we discuss the effect of ω_k on the algorithm performance. Setting ω_k too small implies that the algorithm is not able to infer the process states with a sufficiently high accuracy, which results in a high system cost. On the other hand, setting ω_k too high implies that the algorithm spends a long time on exploring the states of all processes instead of exploiting the gathered information for sampling the processes in the desired order. This insights are supported by Fig. 4. Furthermore, it can be shown that the sufficient conditions on the setting of ω_k given by the theoretical analysis are close to the ω_k that minimizes the system cost in this setting.

VI. Conclusion

We considered the problem of detecting independent anomalous processes among K processes quickly and reliably. The observations from each chosen process follow two different distributions, depending on whether the process is normal or abnormal. Each anomalous process incurs a polynomial cost with time until its anomaly is identified and fixed. The goal is to minimize the total expected cost incurred by all the processes during the entire detection process under error constraints. A search algorithm that consists of exploration, exploitation, and sequential testing phases was proposed and analyzed. For both known and unknown observation models the asymptotic approximation ratio and regret analysis were established.

It should be noted that the algorithm design and the analysis rely on the fact that the cost function is polynomial with degree d. An interesting future direction is to generalize the results for non-polynomial cost functions (e.g., exponential with time).

VII. Appendix

We start by establishing the lower bound on the asymptotic cost. We will then use it to prove the asymptotic optimality of the EEST algorithm.

We note that the analysis of the exploration rate in Lemma 2 is determined by ω_k which depends on the degree d of the polynomial function. Second, the proofs throughout the appendix are based on the fact that the cost function is polynomial for extracting the dominant terms in the asymptotic regime. Third, to analyze the approximation ratio in the proof of Theorem VII-A, the index γ_n depends explicitly on the coefficient $a_{k,d}$ of the dth degree term of the polynomial function. The coefficients $a_{k,1},...,a_{k,d-1}$ of terms with lower degree do not affect the asymptotic analysis.

Lemma 1. *Let $\mathcal{J}$ be the set of all possible permutations $\mathbf{j}=(j_1,j_2,...,j_{|\mathcal{H}_1|})$ of $\{1,2,...,|\mathcal{H}_1|\}$. Let $\boldsymbol{E}(C(s))$ be the total expected cost under policy s that satisfies the error constraints in (1). Then,*

$$\inf_s \boldsymbol{E}(C(s)) \geq \min_{\mathbf{j}\in\mathcal{J}} \sum_{k=1}^{|H_1|} a_{j_k,d} \sum_{i=1}^{k} \left(\frac{B_{j_i}}{D(f_{j_i}^{(1)}||f_{j_i}^{(0)})}\right)^d (1-o(1)), \tag{32}$$

where $o(1)\to 0$ as $P_e^{max}\to 0$.

Proof: Fix $\epsilon>0$, and let $E_\epsilon(s)$ be the event where the sample size of each observation vector taken from process k under strategy s satisfies $N_k > \frac{(1-\epsilon)B_k}{D(f_k^{(1)}||f_k^{(0)})}$ for all $k=1,...,K$. The complementary event is defined by $E_\epsilon^c(s)$. Therefore, the expected cost can be lower bounded by:

$$\begin{aligned}&\Pr(E_\epsilon(s))\cdot\mathbf{E}(C(s)|E_\epsilon(s))\\ &\leq \Pr(E_\epsilon(s))\cdot\mathbf{E}(C(s)|E_\epsilon(s))\\ &\quad+(1-\Pr(E_\epsilon(s)))\cdot\mathbf{E}(C(s)|E_\epsilon^c(s))=\mathbf{E}(C(s)).\end{aligned}$$

Therefore, it suffices to lower bound the term $\Pr(E_\epsilon(s))\cdot\mathbf{E}(C(s)|E_\epsilon(s))$. For this, we define a genie that declares the

state of each process k with detection probability 1 after taking only $\frac{(1-\epsilon)B_k}{D(f_k^{(1)}||f_k^{(0)})}$ samples. Since normal processes do not incur cost, we further assume that genie declares the states of the abnormal processes first. Clearly, the cost obtained under genie's optimal strategy lower bounds $\inf_s \mathbf{E}(C(s)|E_\epsilon(s))$. Next, we show that the optimal strategy for genie is to collect the samples consecutively one by one until declarations without switching between processes during the detection of each process. Assume by contradiction that the optimal sampling strategy for genie switches between processes during the detection of a process. Denote its cost and declaration order by C^S and $\mathbf{j} = (j_1, j_2, ..., j_{|H_1|})$, respectively. As a result, the achievable cost is given by:

$$C^S = \sum_{k=1}^{|H_1|} c_{j_k}\left(\tau_{j_k}^S\right),$$

where $\tau_{j_k}^S$ is the last time where a sample from process j_k was taken. Next, let C^{NS} be the cost obtained by a strategy which does not switch between processes during the detection of each process and declares the processes at the same order $\mathbf{j} = (j_1, j_2, ..., j_{|H_1|})$. As a result,

$$C^{NS} = \sum_{k=1}^{|H_1|} c_{j_k}\left(\tau_{j_k}^{NS}\right),$$

and $\tau_{j_k}^{NS} = \sum_{i=1}^{k} t_{j_i}$, where t_{j_i} is the number of samples needed to detect process j_i. Since C^{NS} is obtained by the same declaration order as C^S, but without switching during the detection of each process, we get $\tau_k^S \geq \tau_k^{NS}$ for all $k \in \{1, 2, ..., |\mathcal{H}_1|\}$. Finally, we use the fact that c_k is a monotonic increasing function resulting in $c_{j_k}(\tau_k^S) \geq c_{j_k}(\tau_k^{NS})$. Thus, $C^{NS} \leq C^S$ which contradicts the assumption. As a result, we next develop the optimal cost for genie by considering strategies that take samples consecutively from each process without switching during the detection of each process. Let $\tilde{\mathbf{j}} = (\tilde{j}_1, \tilde{j}_2, ..., \tilde{j}_{|\mathcal{H}_1|})$ be a permutation of $\{1, 2, ..., |\mathcal{H}_1|\}$ that determines the probing order (without switching during the detection of each process), and minimizes genie's cost. As a result, we have:

$$\begin{aligned}
&\inf_s \mathbf{E}(C(s)|E_\epsilon(s)) \\
&\geq \sum_{k=1}^{|\mathcal{H}_1|} \sum_{n=1}^{d} a_{\tilde{j}_k,n} \left(\sum_{i=1}^{k} \frac{(1-\epsilon)B_{\tilde{j}_i}}{D(f_{\tilde{j}_i}^{(1)}||f_{\tilde{j}_i}^{(0)})} \right)^n . \qquad (33)
\end{aligned}$$

Therefore, in the asymptotic regime we have:

$$\begin{aligned}
&\inf_s \mathbf{E}(C(s)|E_\epsilon(s)) \\
&\geq \sum_{k=1}^{|\mathcal{H}_1|} a_{\tilde{j}_k,d} \sum_{i=1}^{k} \left(\frac{(1-\epsilon)B_{\tilde{j}_i}}{D(f_{\tilde{j}_i}^{(1)}||f_{\tilde{j}_i}^{(0)})} \right)^d \qquad (34) \\
&\qquad\qquad \text{as } P_e^{max} \to 0 .
\end{aligned}$$

Finally, since that $\Pr(E_\epsilon(s))$ approaches 1 as P_e^{max} approaches 0 [45, Lemma 2.1], and (34) holds for every $\epsilon > 0$, (32) follows. □

Next, we present Lemma 2 that was used for proving the theorems in this paper. We start by defining $T_1(\epsilon)$ below.

Definition 1. Let $T_1(\epsilon)$ be the smallest integer such that for all $n \geq T_1$ the following events hold:

$$\begin{aligned}
&\pi_k(n) > 1-\epsilon \quad \forall k \in H_1, \\
&\pi_k(n) < \epsilon \quad \forall k \in H_0.
\end{aligned}$$

To simplify the presentation of the proofs, we often write T_1 without directly referring to ϵ.

T_1 can be viewed as the earliest time at which all the processes are categorized into two groups: one group belonging to the estimated normal set, and the second belonging to the estimated abnormal set. It should be noted that T_1 is not a stopping time. The decision maker does not know whether T_1 has occurred or not during the detection process.

Lemma 2. Set the exploration rates as follows:

$$\omega_k > \frac{d+1}{\min\left\{I_{k|H_0}, I_{k|H_1}\right\}} \quad \forall k. \qquad (35)$$

Then, the d^{th} moment of T_1 is finite,

$$\mathbf{E}[T_1^d] = O(1). \qquad (36)$$

Proof: We start by bounding the d^{th} moment of T_1 as follows:

$$\mathbf{E}[T_1^d] = \sum_{n=0}^{\infty} \Pr(T_1^d > n) \leq \sum_{n=0}^{\infty} \Pr(T_1 > \left\lfloor n^{1/d} \right\rfloor). \qquad (37)$$

Using the definition of T_1 we have:

$$\begin{aligned}
&\Pr\left(T_1 > \left\lfloor n^{1/d} \right\rfloor\right) \\
&= \Pr\Big(\exists\ k \in \mathcal{H}_1\ s.t.\ \pi_k(t) \leq 1-\epsilon \text{ for some } t > \left\lfloor n^{1/d} \right\rfloor \\
&\qquad \text{or } \exists\ k \in \mathcal{H}_0\ s.t.\ \pi_k(t) \geq \epsilon \text{ for some } t > \left\lfloor n^{1/d} \right\rfloor\Big). \qquad (38)
\end{aligned}$$

Define $d_k \triangleq \frac{1-\pi_k(1)}{\pi_k(1)}$ and

$$\begin{aligned}
M_k^{(1)} &\triangleq -\log\left(\frac{\epsilon}{d_k(1-\epsilon)}\right), \\
M_k^{(0)} &\triangleq -\log\left(\frac{d_k\epsilon}{(1-\epsilon)}\right).
\end{aligned}$$

By rewriting the update equation of $\pi_k(n)$, it can be shown that:

$$\pi_k(n) = \left(d_k e^{-S_k(n)} + 1\right)^{-1}. \qquad (39)$$

As a result, $\pi_k(n) \geq 1-\epsilon$ iff $S_k(n) \geq M_k^{(1)}$and $\pi_k(n) \leq \epsilon$ iff $S_k(n) \leq -M_k^{(0)}$. Substituting these two identities into (38)

yields:

$$\begin{aligned}&\Pr\left(T_1 > \left\lfloor n^{1/d}\right\rfloor\right)\\ &= \Pr\left(\exists\, k\in\mathcal{H}_1 \; s.t.\; S_k(t)\le M_k^{(1)} \;\text{ for some }\; t > \left\lfloor n^{1/d}\right\rfloor\right)\\ &\quad+\Pr\left(\exists\, k\in\mathcal{H}_0 \; s.t.\; S_k(t)\ge M_k^{(0)} \;\text{ for some }\; t > \left\lfloor n^{1/d}\right\rfloor\right)\\ &\le \sum_{k\in\mathcal{H}_1}\Pr\left(S_k(t)\le M_k^{(1)} \;\text{ for some }\; t > \left\lfloor n^{1/d}\right\rfloor\right)\\ &\quad+\sum_{k\in\mathcal{H}_0}\Pr\left(S_k(t)\ge M_k^{(0)} \;\text{ for some }\; t > \left\lfloor n^{1/d}\right\rfloor\right).\end{aligned} \tag{40}$$

Summing over all possible values of t yields:

$$\begin{aligned}&\Pr\left(T_1 > \left\lfloor n^{1/d}\right\rfloor\right)\\ &\le \sum_{k\in\mathcal{H}_1}\sum_{t=\left\lfloor n^{1/d}\right\rfloor}^{\infty}\Pr\left(S_k(t)\le M_k^{(1)})\right)\\ &\quad+\sum_{k\in\mathcal{H}_0}\sum_{t=\left\lfloor n^{1/d}\right\rfloor}^{\infty}\Pr\left(S_k(t)\ge M_k^{(0)}\right).\end{aligned} \tag{41}$$

Next, we bound each term in the summations. Let $N_k(t)$ be the number of samples taken from process k up to time t. Let $\widetilde{\ell_k}(i) \triangleq -\ell_k(i) + \frac{M_k^{(1)}}{N_k(t)}$ be a modified log likelihood ratio. Applying the Chernoff bound and using the i.i.d. property yields:

$$\begin{aligned}\Pr(S_k(t)\le M_k^{(1)}) &\le \min_{s>0}\left[\mathbf{E}\left[e^{s\widetilde{\ell_k}(i)}\right]\right]^{N_k(t)} =\\ &= \min_{s>0}\left[e^{-N_k(t)\left[-\log\mathbf{E}\left[e^{s\widetilde{\ell_k}(i)}\right]\right]}\right]\\ &= e^{-N_k(t)\left(\sup_{s>0}\left[-\log\mathbf{E}\left[e^{s\widetilde{\ell_k}(i)}\right]\right]\right)}.\end{aligned}$$

Since $\widetilde{\ell_k}(i)\to -\ell_k(i)$ as $N_k(t)\to\infty$, we have:

$$\sup_{s>0} -\log\mathbf{E}\left[e^{s\widetilde{\ell_k}(i)}\right] \underset{N_k(t)\to\infty}{\longrightarrow} \sup_{s>0} -\log\mathbf{E}\left[e^{-s\ell_k(i)}\right] \triangleq I_k^*,$$

where I_k^* is the Legendre-Fenchel transformation of $-\ell_k(i)$. Thus, we can choose $I_{k,\epsilon_I}^* \triangleq I_k^* - \epsilon_I$ for any $\epsilon_I > 0$ such that

$$\Pr(S_k(t)\le M_k^{(1)}) \le C(\epsilon_I)e^{-N_k(t)I_{k,\epsilon_I}^*},$$

where $C(\epsilon_I)$ is a finite constant independent of the sample size. Since $N_k(t) > \omega_k\log(t)$ by the construction of the algorithm, for all t we have:

$$\begin{aligned}\sum_{t=\left\lfloor n^{1/d}\right\rfloor}^{\infty}\Pr(S_k(t)\le M_k^{(1)}) &\le \sum_{t=\left\lfloor n^{1/d}\right\rfloor}^{\infty} C(\epsilon_I)t^{-\omega_k I_{k,\epsilon_I}^*}\\ &\le \int_{t=\left\lfloor n^{1/d}\right\rfloor-1}^{\infty} C(\epsilon_I)t^{-\omega_k I_{k,\epsilon_I}^*}dt\\ &\le \tilde{C}(\epsilon_I)\left(\left\lfloor n^{1/d}\right\rfloor-1\right)^{1-\omega_k I_{k,\epsilon_I}^*}\\ &\le \tilde{C}(\epsilon_I)n^{\frac{1-\omega_k I_{k,\epsilon_I}^*}{d}},\end{aligned}$$

where $\tilde{C}(\epsilon_I)$ is a finite constant independent of n. Using the same method, we can derive the bound for $\sum_{t=\left\lfloor n^{1/d}\right\rfloor}^{\infty}\Pr(S_k(t)\ge M_k^{(0)})$ where I_{k,ϵ_I}^* can be determined with respect to H_0 instead of H_1. Substituting these bounds into (37) yields:

$$\begin{aligned}\mathbf{E}[T_1^d] &= \sum_{n=0}^{\infty}\Pr(T_1^d > n) \le\\ &\le \sum_{k\in H_1}\sum_{n=0}^{\infty}\tilde{C}(\epsilon_I)n^{\frac{1-\omega_k I_{k,\epsilon_I}^*}{d}}\\ &\quad+\sum_{k\in H_0}\sum_{n=0}^{\infty}\tilde{C}(\epsilon_I)n^{\frac{1-\omega_k I_{k,\epsilon_I}^*}{d}}.\end{aligned} \tag{42}$$

Since we want (42) to be bounded, we require that:

$$-\frac{1-\omega_k I_{k,\epsilon_I}^*}{d} > 1 \;\; \forall k.$$

Since $\epsilon_I > 0$ can be arbitrarily small, it is sufficient to require:

$$\omega_k > \frac{d+1}{I_k^*} \quad \forall k.$$

However, since the process states are unknown, sufficient conditions can be derived by requiring:

$$\omega_k > \frac{d+1}{\min\{I_{k|H_1}, I_{k|H_0}\}} \quad \forall k,$$

where $I_{k|H_0} = \sup_{s>0}\left[-\log\mathbf{E}_{\sim f_k^{(0)}}\left(e^{s\ell_k}\right)\right]$, $I_{k|H_1} = \sup_{s>0}\left[-\log\mathbf{E}_{\sim f_k^{(1)}}\left(e^{-s\ell_k}\right)\right]$ are the Legendre-Fenchel transformation of ℓ_k, and $-\ell_k$, at the origin with respect to distributions $f_k^{(0)}$ and $f_k^{(1)}$, respectively, which proves the lemma. □

A. Proof of Theorem 1

Define by

$$\rho_1 \triangleq \frac{\mathbf{E}(C^{EEST}||\mathcal{H}_1| > N_c)}{\inf_{\mathbf{s}}\;\mathbf{E}(C(\mathbf{s}))}, \tag{43}$$

and

$$\rho_2 \triangleq \frac{\mathbf{E}(C^{EEST}||\mathcal{H}_1| \le N_c)}{\inf_{\mathbf{s}}\;\mathbf{E}(C(\mathbf{s}))}, \tag{44}$$

the approximation ratios between the expected cost under the EEST algorithm given that $|\mathcal{H}_1| > N_c$ and $|\mathcal{H}_1| \le N_c$, respectively, and the infimum expected cost, $\inf_{\mathbf{s}}\;\mathbf{E}(C(\mathbf{s}))$, over all algorithms that satisfy the error constraints.

In order to prove theorem 1.a we show that

$$\rho_2 = 1 \;\text{ as }\; P_e^{max} \to 0. \tag{45}$$

For theorem 1.b we first rewrite the approximation ratio as follows:

$$\rho = P_c\rho_1 + (1-P_c)\rho_2, \tag{46}$$

where $P_c = P(|H_1| > N_c)$. For the proof of theorem 1.b we show that (46) is equal to (15).

We start by evaluating ρ_1. Without loss of generality, we assume that the EEST algorithm has declared the process states in the following order: $1, 2, 3, ..., K$. Thus, $\tau_i \triangleq \sum_{k=1}^{i} N_k$ is

defined as the detection time of process i. Since τ_k increases as P_e^{max} decreases, the cost that process k incurs is asymptotically equal to[3] $c_k(\tau_k) \sim a_{k,d}\tau_k^d$.

Let $\tilde{N}_k$ be the remaining sample size required to declare the state of process k after T_1 has elapsed. Let N_k^{EX} be the sample size taken during the exploration phase. Using these definitions, we can bound the stopping time of process i by:

$$\tau_i \leq T_1 + \sum_{k=1}^{K} N_k^{\text{EX}} + \sum_{k=1}^{i} \tilde{N}_k.$$

Hence,

$$\begin{aligned} C^{EEST} &\sim \sum_{i=1}^{K} a_{i,d}\tau_i^d \leq \\ &\leq \sum_{i=1}^{K} a_{i,d}\left[T_1 + \sum_{k=1}^{K} N_k^{\text{EX}} + \sum_{k=1}^{i} \tilde{N}_k\right]^d \leq \\ &\leq \sum_{i=1}^{K} a_{i,d}\Bigg[T_1^d + \left(\sum_{k=1}^{K} N_k^{\text{EX}}\right)^d + \left(\sum_{k=1}^{i} \tilde{N}_k\right)^d \\ &\quad + f(T_1, \sum_{k=1}^{K} N_k^{\text{EX}}, \sum_{k=1}^{i} \tilde{N}_k)\Bigg], \end{aligned} \tag{47}$$

where $f(T_1, \sum_{k=1}^{K} N_k^{\text{EX}}, \sum_{k=1}^{i} \tilde{N}_k)$ is a function that contains all the cross products between $\{T_1, \sum_{k=1}^{K} N_k^{\text{EX}}, \sum_{k=1}^{i} \tilde{N}_k\}$ up to an order $(d-1)$. Next, we upper bound each term in the summation on the RHS of (47). From [46, Theorem 4.1], we know that the d^{th} moment of the sample size of the sequential test is asymptotically equal to:

$$\mathbf{E}[N_k^d] \sim \left(\frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}\right)^d. \tag{48}$$

Let $n \triangleq \tau_K$ be the total detection time. Since $\sum_{k=1}^{K} N_k^{\text{EX}} < c\log(n)$ for some constant $c > 0$ by the construction of the algorithm, we have:

$$\left(\sum_{k=1}^{K} N_k^{\text{EX}}\right)^d = O(\log^d(n)). \tag{49}$$

Next, since that $n \sim T_1 + \sum_{k=1}^{K} \tilde{N}_k$, taking expectation, and applying Lemma 2 yields:

$$\mathbf{E}\left[\left(\sum_{k=1}^{K} \tilde{N}_k\right)^d\right] = O(\mathbf{E}[n^d]).$$

Similarly, since $f(T_1, \sum_{k=1}^{K} N_k^{\text{EX}}, \sum_{k=1}^{i} \tilde{N}_k)$ contains cross products of all the terms up to order $d-1$, its expectation is of order $o(\mathbf{E}[n]^d)$.

Finally, fixing $\mathcal{H}_1$, taking expectation (with respect to the randomness of the observations) on both sides of (47), and using the fact that normal processes do not incur costs yields:

$$\mathbf{E}\left[C^{EEST}\right] \leq (1+o(1)) \sum_{i=1}^{|\mathcal{H}_1|} a_{i,d}\left(\sum_{k=1}^{i} \frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}\right)^d. \tag{50}$$

Finally, since the sample size of each process is asymptotically equal to (48), it is sufficient to evaluate the ratio between $\sum_{i=1}^{|\mathcal{H}_1|} a_{i,d}\mathbf{E}\left[\left(\sum_{k=1}^{i} \frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}\right)^d\right]$, and the infimum cost obtained over all possible orders of processes $1, ..., |\mathcal{H}_1|$ with sample sizes $\left\{\frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}\right\}_{k=1}^{|\mathcal{H}_1|}$, which are deterministic values. Since $|\mathcal{H}_1| > N_C$ when evaluating ρ_1, we can choose $\epsilon > 0$ sufficiently small in Definition 1 so that $\frac{a_{1,d}}{B_1/D(f_1^{(1)}||f_1^{(0)})} \geq \frac{a_{2,d}}{B_2/D(f_2^{(1)}||f_2^{(0)})} \geq \cdots \geq \frac{a_{\mathcal{H}_1,d}}{B_{\mathcal{H}_1}/D(f_{\mathcal{H}_1}^{(1)}||f_{\mathcal{H}_1}^{(0)})}$. Finally, we can apply the bound in [47, Corollary 3.9] for deterministic sequences to upper bound the target approximation ratio by:

$$\max_{0\leq s\leq 1} \frac{(d+1)s + (1-s)^{d+1}}{1 + ds^{d+1}}, \tag{51}$$

for any $d \geq 1$. As a result, ρ_1 is upper bounded by:

$$\rho_1 \leq M(d) = \max_{0\leq s\leq 1} \frac{(d+1)s + (1-s)^{d+1}}{1 + ds^{d+1}} \text{ as } P_e^{max} \to 0. \tag{52}$$

Next, we evaluate ρ_2. Since $\epsilon > 0$ in Definition 1 can be arbitrarily small, let

$$\epsilon < \min\{L_1, L_2, ..., L_K, 1-U_1, 1-U_2, ..., 1-U_K\}.$$

As a result, for all $t > T_1$, we have: $\pi_k(t) > U_k$ for all $k \in \mathcal{H}_1$, and $\pi_k(t) < L_k$ for all $k \in \mathcal{H}_0$. Next, we use similar steps as in the evaluation of ρ_1. Since the expected cost up to time T_1 is finite independent of t by Lemma 2, and the algorithm performs a combinatorial search to test all the abnormal processes in the optimal order for all $t > T_1$, we have:

$$\rho_2 = 1 \text{ as } P_e^{max} \to 0. \tag{53}$$

Note that this is the condition of theorem 1.a as given in (45). Substituting (52) and (53) into (46) completes the proof of part b of the theorem. □

B. Proof of Theorem 2

Without loss of generality, we assume that the order in which the processes are declared under the genie's policy is $1, 2, 3, ..., K$. Recall that under the genie's policy the beliefs are set as follows:

$$\pi_k(n) = \begin{cases} 1, & k \in \mathcal{H}_1, \\ 0, & k \in \mathcal{H}_0, \end{cases}$$

for all n. As a result, the genie performs the sequential testing step for all processes following the order given by the index selection phase if $|\mathcal{H}_1| > N_C$ or the combinatorial search phase if $|\mathcal{H}_1| \leq N_C$. As a result, the cost under the genie's strategy is given by:

$$C^* = \sum_{i=1}^{K} a_{i,d}\left(\sum_{k=1}^{i} N_k\right)^d + O(n^{d-1}).$$

[3]The notation $f(x) \sim g(x)$ refers to $\lim_{x\to\infty} \frac{f(x)}{g(x)} = 1$.

Using the same sample path, the cost under the EEST algorithm is given by

$$C^{EEST} = \sum_{i=1}^{K} a_{i,d}\tau_i^d + O(n^{d-1}) + O(T_1^d), \qquad (54)$$

where the last term is due to a wrong detection order by time $n \leq T_1$. The declaration time of process i under the EEST algorithm can be written as

$$\tau_i = \sum_{k=1}^{i} N_k + O(\log n),$$

where $O(\log n)$ is due to the exploration phase. Thus,

$$\tau_i^d = \left(\sum_{k=1}^{i} N_k\right)^d + O(n^{d-1}\log n).$$

Therefore, the cost under the EEST algorithm in (54) can be rewritten as

$$C^{EEST} = \sum_{i=1}^{K} a_{i,d}\left(\sum_{k=1}^{i} N_k\right)^d + O(n^{d-1}\log n) + O(T_1^d).$$

As a result, we have:

$$C^{EEST} - C^* = O(n^{d-1}\log n) + O(T_1^d). \qquad (55)$$

Finally, taking the expectation and applying Lemma 2 yields:

$$R = O(\mathbf{E}[n^{d-1}\log n]), \qquad (56)$$

as required. □

C. Proof of Theorem 3

In this section we show that asymptotic optimality and the approximation ratio given in (29) hold under the composite hypothesis case. We consider the composite hypothesis setting in [40], in which the unknown parameter θ_k takes a value from a finite set Θ_k. The true value of the unknown parameter for process k is denoted by θ_k. The lower bound on the asymptotic cost satisfies (32) with the corresponding parametric distributions using similar arguments as in Lemma 1, where the KL divergence is measured between the true parameter value and the closest alternative as in standard composite hypothesis testing.

We first develop the exploration rate required for achieving the desired estimation consistency.

Definition 2. T_{ML} is defined as the smallest integer such that $\hat{\theta}_k(n) = \theta_k$ for all $n > T_{ML}$ for all k.

In the following lemma we introduce a sufficient condition on the exploration rate ω_k such that the d^{th} moment of T_{ML} is finite.

Lemma 3. Set the exploration rates as defined in (25). Then, the d^{th} moment of T_{ML} is finite:

$$\mathbf{E}\left[T_{ML}^d\right] = O(1).$$

Proof: Let $N_k(t)$ be the number of samples taken from process k, and $\mathbf{1}_k(t)$ be the probing indicator function of process k, where $\mathbf{1}_k(t) = 1$ if process k was probed at time t, and $\mathbf{1}_k(t) = 0$ otherwise. We also define:

$$S_k(\theta_k, \theta, n) \triangleq \sum_{t=1}^{n} \mathbf{1}_k(t)\ell_k(\theta_k, \theta, t),$$

as the sum of LLRs with respect to the parameters θ_k and θ, where $\ell_k(\theta_k, \theta, t) \triangleq \log \frac{f_k(y_k(t)|\theta_k)}{f_k(y_k(t)|\theta)}$.

By the definition of the MLE, when $\hat{\theta}_k = \theta_k$, we have: $S_k(\theta_k, \theta, n) > 0$ for all $\theta \neq \theta_k$. Hence, we can write

$$\begin{aligned}
\Pr(T_{ML}^d > n) &\leq \Pr\left(T_{ML} > \left\lfloor n^{1/d}\right\rfloor\right) \\
&= \Pr\left(\exists\ t > \left\lfloor n^{1/d}\right\rfloor, k, \theta \neq \theta_k\ s.t.\ S_k(\theta_k, \theta, t) < 0\right) \\
&\leq \sum_{t=\left\lfloor n^{1/d}\right\rfloor+1}^{\infty} \sum_{k=1}^{K} \sum_{\theta\in\Theta_k\setminus\theta_k} \Pr\left(S_k(\theta_k, \theta, t) < 0\right) \\
&\leq \sum_{t=\left\lfloor n^{1/d}\right\rfloor+1}^{\infty} \sum_{k=1}^{K} \sum_{\theta\in\Theta_k\setminus\theta_k} \left[\mathbf{E}\left[e^{s\ell_k(\theta_k, \theta, t)}\right]\right]^{N_k(t)},
\end{aligned}$$

where the last inequality follows by applying the Chernoff bound. Since the state of process k is unknown, we choose the minimum rate function over the whole parameter set when deriving conditions on ω_k, and similar to the steps in Lemma 2 we have:

$$\omega_k > \frac{d+1}{\min\left\{\tilde{I}_{k,0}, \tilde{I}_{k,1}\right\}}\ \ \forall k, \qquad (57)$$

where

$$\begin{aligned}
&\tilde{I}_{k,0} \triangleq \\
&\inf_{\substack{\theta^{(0)}\neq\theta^{(1)},\\ \theta^{(0)},\theta^{(1)}\in\Theta_k}} \sup_{s>0} \left[-\log \mathbf{E}_{\sim f_k(y|\theta^{(0)})}\left(e^{s\ell_k(\theta^{(0)},\theta^{(1)})}\right)\right], \\
&\tilde{I}_{k,1} \triangleq \\
&\inf_{\substack{\theta^{(0)}\neq\theta^{(1)},\\ \theta^{(0)},\theta^{(1)}\in\Theta_k}} \sup_{s>0} \left[-\log \mathbf{E}_{\sim f_k(y|\theta^{(1)})}\left(e^{-s\ell_k(\theta^{(0)},\theta^{(1)})}\right)\right]
\end{aligned} \qquad (58)$$

are the Legendre-Fenchel transformation of

$$\ell_k(\theta^{(0)}, \theta^{(1)}) \triangleq \log \frac{f_k(y|\theta^{(1)})}{f_k(y|\theta^{(0)})}$$

at the origin with respect to distributions $f_k(y|\theta^{(0)})$ and $f_k(y|\theta^{(1)})$, respectively. □

We next modify Definition 1 and Lemma 2 to prove the theorem under unknown parameters.

Definition 3. Let $T_1(\epsilon)$ be the smallest integer such that for all $n \geq T_1$ the following events hold:

$$\begin{aligned}
&\hat{\pi}_k(n) > 1-\epsilon \quad \forall k \in H_1, \\
&\hat{\pi}_k(n) < \epsilon \quad \forall k \in H_0.
\end{aligned}$$

To simplify the presentation of the proofs, we often write T_1 without directly referring to ϵ.

Lemma 4. Set the exploration rates as in (25). Then, the d^{th} moment of T_1 is finite,

$$\mathbf{E}[T_1^d] = O(1). \qquad (59)$$

Proof: Note that:

$$\mathbf{E}[T_1^d] = \sum_{n=0}^{\infty} \Pr(T_1^d > n). \tag{60}$$

Thus, it suffices to show that $\Pr(T_1^d > n)$ decreases polynomially with time. Note that

$$\begin{aligned}\Pr(T_1^d > n) &= \Pr(T_1^d > n, T_{ML}^d \leq n) \\ &\quad + \Pr(T_1^d > n, T_{ML}^d > n) \\ &\leq \Pr(T_1^d > n, T_{ML}^d \leq n) + \Pr(T_{ML}^d > n). \end{aligned} \tag{61}$$

The first term on the RHS decreases polynomially with n following similar steps as in Lemma 2, and the fact that $\hat{\theta}_k(n) = \theta_k$ for all $n \geq T_{ML}$ for all k. The second term decreases polynomially by applying Lemma 3. □

To complete the proof of Theorem 3, note that the cost is given by:

$$C^{EEST} \leq \sum_{i=1}^{K} a_{i,d} \left[T_{ML} + T_1 + \sum_{k=1}^{K} N_k^{\text{EX}} + \sum_{k=1}^{i} \tilde{N}_k \right]^d, \tag{62}$$

since that

$$\tau_i \leq T_{ML} + T_1 + \sum_{k=1}^{K} N_k^{\text{EX}} + \sum_{k=1}^{i} \tilde{N}_k.$$

Applying Lemmas 3, 4 yields $\mathbf{E}[T_{ML}^d] = O(1)$ and $\mathbf{E}[T_1^d] = O(1)$. Therefore, the rest of the proof follows by similar steps as in Theorem 1. □

D. Proof of Theorem 4

Following the same arguments as in the proof of Theorem 2, the cost under the genie's policy is given by:

$$C^* = \sum_{i=1}^{K} a_{i,d} \left(\sum_{k=1}^{i} N_k \right)^d + O(n^{d-1}). \tag{63}$$

The cost under the EEST algorithm is given by:

$$C^{EEST} = \sum_{i=1}^{K} a_{i,d} \tau_i^d + O(n^{d-1}) + O(T_1^d) + O(T_{ML}^d). \tag{64}$$

By applying Lemma 3, the d^{th} moment of T_{ML} is finite. Thus, The rest of the proof is similar to the proof of Theorem 2. □

REFERENCES

[1] A. Gurevich, K. Cohen, and Q. Zhao, "Anomaly detection under a nonlinear system cost objective function," in *Proc. of the 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 634–638, 2017.

[2] A. Wald, "Sequential analysis," *New York: Wiley*, 1947.

[3] V. Katewa and V. Gupta, "On the optimality of sequential test with multiple sensors," in *American Control Conference (ACC)*, pp. 282–287, 2012.

[4] S. Nitinawarat and V. 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, 2015.

[5] C.-Z. Bai, V. Katewa, V. Gupta, and Y.-F. Huang, "A stochastic sensor selection scheme for sequential hypothesis testing with multiple sensors," *IEEE transactions on signal processing*, vol. 63, no. 14, pp. 3687–3699, 2015.

[6] K. Cohen, Q. Zhao, and A. Swami, "Optimal index policies for anomaly localization in resource-constrained cyber systems," *IEEE Transaction on Signal Processing*, vol. 62, pp. 4224–4236, Aug 2014.

[7] K. Cohen and Q. Zhao, "Asymptotically optimal anomaly detection via sequential testing," *IEEE Trans. Signal Processing*, vol. 63, no. 11, pp. 2929–2941, 2015.

[8] H. Rifaï, S. Mohammed, and A. Mellouk, "A brief synthesis of QoS-QoE methodologies," in *10th IEEE International Symposium on Programming and Systems (ISPS)*, pp. 32–38, 2011.

[9] G. Schwarz, "Asymptotic shapes of Bayes sequential testing regions," *The Annals of mathematical statistics*, pp. 224–236, 1962.

[10] T. L. Lai, "Nearly optimal sequential tests of composite hypotheses," *The Annals of Statistics*, pp. 856–886, 1988.

[11] I. V. Pavlov, "Sequential procedure of testing composite hypotheses with applications to the Kiefer-Weiss problem," *Theory of Probability and Its Applications*, vol. 35, no. 2, pp. 280–292, 1990.

[12] A. G. Tartakovsky, "An efficient adaptive sequential procedure for detecting targets," in *IEEE Aerospace Conference Proceedings, 2002*, vol. 4, pp. 1581–1596, 2002.

[13] V. Dragalin, A. G. Tartakovsky, and V. V. Veeravalli, "Multihypothesis sequential probability ratio tests - part I: Asymptotic optimality," *IEEE Transactions on Information Theory*, vol. 45, no. 7, pp. 2448–2461, 1999.

[14] V. Dragalin, "A simple and effective scanning rule for a multi-channel system," *Metrika*, vol. 43, no. 1, pp. 165–182, 1996.

[15] Y. Song and G. Fellouris, "Sequential multiple testing with generalized error control: an asymptotic optimality theory," *arXiv preprint arXiv:1608.07014*, 2016.

[16] Y. Song and G. Fellouris, "Asymptotically optimal, sequential, multiple testing procedures with prior information on the number of signals," *Electronic Journal of Statistics*, vol. 11, no. 1, pp. 338–363, 2017.

[17] Q. Zhao and J. Ye, "Quickest detection in multiple on–off processes," *IEEE Transactions on Signal Processing*, vol. 58, no. 12, pp. 5994–6006, 2010.

[18] H. Li, "Restless watchdog: Selective quickest spectrum sensing in multichannel cognitive radio systems," *EURASIP Journal on Advances in Signal Processing*, vol. 2009, 2009.

[19] R. Caromi, Y. Xin, and L. Lai, "Fast multiband spectrum scanning for cognitive radio systems," *IEEE Transaction on Communications*, vol. 61, no. 1, pp. 63–75, 2013.

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

[21] M. L. Malloy and R. D. Nowak, "Sequential testing for sparse recovery," *IEEE Transactions on Information Theory*, vol. 60, no. 12, pp. 7862–7873, 2014.

[22] M. L. Malloy, G. Tang, and R. D. Nowak, "Quickest search for a rare distribution," *IEEE Annual Conference on Information Sciences and Systems*, pp. 1–6, 2012.

[23] A. Tajer and H. V. Poor, "Quick search for rare events," *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4462–4481, 2013.

[24] J. Geng, W. Xu, and L. Lai, "Quickest search over multiple sequences with mixed observation," *available at http://arxiv.org/abs/1312.2287*, 2014.

[25] J. Heydari, A. Tajer, and H. V. Poor, "Quickest linear search over correlated sequences," *IEEE Transactions on Information Theory*, vol. 62, no. 10, pp. 5786–5808, 2016.

[26] G. Fellouris, G. V. Moustakides, and V. V. Veeravalli, "Multistream quickest change detection: Asymptotic optimality under a sparse signal," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6444–6447, 2017.

[27] K. S. Zigangirov, "On a problem in optimal scanning," *Theory of Probability and Its Applications*, vol. 11, no. 2, pp. 294–298, 1966.

[28] K. Cohen and Q. Zhao, "Quickest anomaly detection: a case of active hypothesis testing," *in Proc. of Information Theory and Applications (ITA) Workshop*, 2014.

[29] K. Cohen and Q. Zhao, "Active hypothesis testing for anomaly detection," *IEEE Transactions on Information Theory*, vol. 61, no. 3, pp. 1432–1450, 2015.

[30] S. Nitinawarat and V. V. Veeravalli, "Universal scheme for optimal search and stop," in *Information Theory and Applications Workshop (ITA)*, pp. 322–328, 2015.

[31] K. Leahy and M. Schwager, "Always choose second best: Tracking a moving target on a graph with a noisy binary sensor," in *European Control Conference (ECC), 2016*, pp. 1715–1721, 2016.

[32] J. Heydari, A. Tajer, and H. V. Poor, "Quickest detection of markov networks," in *IEEE International Symposium on Information Theory (ISIT), pp. 1341–1345, 2016, arXiv version: arXiv:1711.04268v2 , 2017.*

[33] B. Hemo, K. Cohen, and Q. Zhao, "Asymptotically optimal search of unknown anomalies," in *IEEE International Symposium on Signal Processing and Information Technology (ISSPIT)*, pp. 75–80, 2016.

[34] B. Huang, K. Cohen, and Q. Zhao, "Active anomaly detection in heterogeneous processes," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3924–3928, April 2018.

[35] B. Huang, K. Cohen, and Q. Zhao, "Active anomaly detection in heterogeneous processes," *IEEE Transactions on Information Theory*, vol. 65, no. 4, pp. 2284–2301, 2019.

[36] J. B. Kadane, "Optimal whereabouts search," *Operations Research*, vol. 19, no. 4, pp. 894–904, 1971.

[37] D. A. Castanon, "Optimal search strategies in dynamic hypothesis testing," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 25, no. 7, pp. 1130–1138, 1995.

[38] Y. Zhai and Q. Zhao, "Dynamic search under false alarms," *in Proc. IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2013.

[39] Y. Li, S. Nitinawarat, and V. V. Veeravalli, "Universal outlier hypothesis testing," *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 4066–4082, 2014.

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

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

[42] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *The Annals of Statistics*, vol. 41, no. 6, pp. 2703–2738, 2013.

[43] M. Naghshvar and T. Javidi, "Sequentiality and adaptivity gains in active hypothesis testing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 5, pp. 768–782, 2013.

[44] H. Robbins and D. Siegmund, "A class of stopping rules for testing parametric hypotheses," *Proceeding of the Sixth Berkeley Symposium on Theory of Probability and Math. Statistics*, pp. 37–41, 1972.

[45] A. G. Tartakovsky, "Asymptotic optimality of certain multihypothesis sequential tests: Non-iid case," *Statistical Inference for Stochastic Processes*, vol. 1, no. 3, pp. 265–295, 1998.

[46] V. Dragalin, A. G. Tartakovsky, and V. V. Veeravalli, "Multihypothesis sequential probability ratio tests. i. asymptotic optimality," *IEEE Transactions on Information Theory*, vol. 45, no. 7, pp. 2448–2461, 1999.

[47] W. Höhn and T. Jacobs, "On the performance of Smith's rule in single-machine scheduling with nonlinear cost," *ACM Transactions on Algorithms (TALG)*, vol. 11, no. 4, p. 25, 2015.

Andrey Gurevich has an M.Sc. degree in Electrical and Computer Engineering from Ben-Gurion University of the Negev. His thesis focused on sequential anomaly detection via statistical signal processing and machine learning. He is currently working as an algorithm developer as part of the Autonomous Vehicle project at Mobileye Vision Technology Ltd, Israel.

Kobi Cohen received the B.Sc. and Ph.D. degrees in electrical engineering from Bar-Ilan University, Ramat Gan, Israel, in 2007 and 2013, respectively. In October 2015, he joined the Department of Electrical and Computer Engineering at Ben-Gurion University of the Negev (BGU), Beer Sheva, Israel, as a Senior Lecturer (a.k.a. Assistant Professor in the US). He is also a member of the Cyber Security Research Center, and the Data Science Research Center at BGU. Before joining BGU, he was with the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign (08/2014–07/2015) and the Department of Electrical and Computer Engineering at the University of California, Davis (11/2012–07/2014) as a postdoctoral research associate. His main research interests include decision theory, stochastic optimization, and statistical inference and learning, with applications in large-scale systems, cyber systems, wireless and wireline networks. Dr. Cohen received several awards including the Best Paper Award in the International Symposium on Modeling and Optimization in Mobile, Ad hoc and Wireless Networks (WiOpt) 2015, the Feder Family Award (second prize), granted by the Advanced Communication Center at Tel Aviv University (2011), and President Fellowship (2008-2012) and top honor list's prizes (2006, 2010, 2011) from Bar-Ilan University.

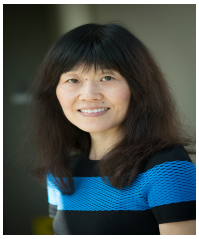

Qing Zhao (F'13) is a Joseph C. Ford Professor of Engineering in the School of Electrical and Computer Engineering at Cornell University, Ithaca, NY. Prior to joining Cornell University in 2015, she was a Professor at University of California, Davis. She received her Ph.D. degree in Electrical Engineering from Cornell University in 2001. Her research interests include decision theory and stochastic optimization, machine learning, statistical inference, and algorithmic theory with applications in infrastructure and communication systems and social economic networks. She is a Fellow of IEEE. She received the 2010 IEEE Signal Processing Magazine Best Paper Award and the 2000 Young Author Best Paper Award from the IEEE Signal Processing Society. While on the faculty of UC Davis, she held the title of UC Davis Chancellor's Fellow. She is a Marie Skłodowska-Curie Fellow awarded by the European Union's Horizon 2020 research and innovation program, and a Jubilee Chair Professor of Chalmers University, Sweden, during sabbatical leave in 2018-2019.