# Asymptotically Optimal Anomaly Detection via Sequential Testing

Kobi Cohen and Qing Zhao, *Fellow, IEEE*

***Abstract*—Sequential detection of independent anomalous processes among $K$ processes is considered. At each time, only $M$ $(1 \leq M \leq K)$ processes can be observed, and the observations from each chosen process follow two different distributions, depending on whether the process is normal or abnormal. Each anomalous process incurs a cost per unit time until its anomaly is identified and fixed. Switching across processes and state declarations are allowed at all times, while decisions are based on all past observations and actions. 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 develop index-type algorithms for the case with both known observation distributions and the case when the observation distributions have unknown parameters. We show that the proposed algorithms are asymptotically optimal in terms of minimizing the total expected cost as the error constraints approach zero. Simulation results demonstrate strong performance in the finite regime.**

***Index Terms*—Anomaly detection, sequential hypothesis testing, sequential probability ratio test (SPRT), Wald's approximation.**

## I. Introduction

CONSIDER a system consisting of $K$ processes, which can be components (such as routers and paths) in a cyber system, channels in a communication network, potential locations of targets, and sensors monitoring certain events. The state of each process is either normal or abnormal (e.g., the busy/idle state of a channel, the presence or absence of a target or event). Process $k$ is abnormal with prior probability $\pi_k$, independent of other processes. Each abnormal process incurs a cost $c_k$ per unit time until its anomaly is identified and fixed. Normal processes incur no cost. Due to resource constraints, only $M$ $(1 \leq M \leq K)$ processes can be probed at a time, and the observations from a probed process follow distributions $f_k^{(0)}$ or $f_k^{(1)}$ depending on whether the process is normal or abnormal. The objective is a sequential search strategy that dynamically determines which processes 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.

The problem under study finds applications in intrusion detection in cyber systems, spectrum scanning in cognitive radio networks (for quickly catching and utilizing idle channels), target search, and event detection in sensor networks.

### A. Main Results

Since observations are drawn in a one-at-a-time manner, the above anomaly detection problem has a clear connection with the classic sequential hypothesis testing problem pioneered by Wald in [1]. 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 abnormal cost 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 Kullback-Leibler divergence 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.

This anomaly detection problem was first formulated and studied in our prior work [2], [3] under the restriction that each process must be probed continuously until its state is declared. In other words, switching across processes is allowed only when the state of the currently probed process is declared. It was shown in [3] that the optimal probing strategy is an open-loop strategy that probes processes in a decreasing order of $\frac{\pi_k c_k}{\mathbf{E}(N_k)}$ (referred to as the OL-$\pi cN$ rule), where $\mathbf{E}(N_k)$ is the expected detection time for process $k$. With the restriction that the test of the currently chosen process has to be completed before testing other processes, it is perhaps not surprising that the optimal probing strategy is open-loop: the probing order is predetermined based on prior information $\{\pi_k, c_k, f_k^{(0)}, f_k^{(1)}\}$, and $K$ uninterrupted sequential tests are carried out, one over each process.

In this paper we relax the restriction on switching across processes during the detection process. We are thus facing a full-blown dynamic problem where at any given time, the decision maker can choose any process whose state has not been declared and the optimal strategy hinges on fully utilizing the entire decision and observation history. In this case, the priority of each process in probing needs to be dynamically updated based on each newly obtained observation. In particular, the probability of each process being abnormal, a key factor in determining the

Manuscript received September 21, 2014; revised March 01, 2015; accepted March 05, 2015. Date of publication March 25, 2015; date of current version May 05, 2015. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Ami Wiesel. This work was supported in part by National Science Foundationunder Grants CCF-1320065 and CNS-1321115. Part of this work was presented at the Fifty-Second Annual Allerton Conference on Communication, Control, and Computing, 2014.

K. Cohen is with the Coordinated Science Lab, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: yscohen@illinois.edu).

Q. Zhao is with the Department of Electrical and Computer Engineering, University of California, Davis, CA 95616 USA (e-mail: qzhao@ucdavis.edu).

Digital Object Identifier 10.1109/TSP.2015.2416674

probing order as shown in our prior work [3], should be updated from the prior probability $\pi_k$ to the *a posteriori* probability $\pi_k(n)$ at time $n$ based on all past observations from this process. Consequently, the expected detection time of process $k$ will also dynamically change based on the *a posteriori* probability of being abnormal (see (11)). Building upon the insights obtained in our prior work [3], we thus propose the following closed-loop $\pi cN$ rule (referred to as CL-$\pi cN$). Let $\mathbf{E}^{(n)}(N_k)$ be the expected detection time of process $k$ based on $\pi_k(n)$. In general, it is difficult to obtain a closed-form expression for $\mathbf{E}^{(n)}(N_k)$ under the finite regime. Therefore, we propose to use the Wald's approximation of $\mathbf{E}^{(n)}(N_k)$, denoted by $\hat{\mathbf{E}}^{(n)}(N_k)$. At each given time $n$, each process is associated with an index $\gamma_k(n) \triangleq \frac{\pi_k(n)c_k}{\hat{\mathbf{E}}^{(n)}(N_k)}$. At each time (except at a sparse subsequence of time instants as detailed below), the process with the largest index is probed, and its state is detected via a sequential test using all past observations. The index of this process is also updated (based on the newly obtained observation) for comparison with other processes at the next time instant. To ensure that all processes are sufficiently probed so that the belief $\pi_k(n)$ (consequently the index $\gamma_k(n)$) is a sufficiently accurate indication of the process state, processes are probed in a round-robin fashion at a subsequence of time instants that grows exponentially sparse with time. In other words, a logarithmic order of time is used to explore the state of all processes to ensure the accuracy of the indices $\gamma_k(n)$ used in the remaining majority of time instants. We show that asymptotic (as the error constraints approach zero) optimality of the CL-$\pi cN$ strategy holds under $M = 1$ for both known and unknown observation models (i.e., whether $\{f_k^{(0)}, f_k^{(1)}\}$ are known or has unknown parameters). When $M > 1$, we show that CL-$\pi cN$ preserves its asymptotic optimality if processes incur the same cost when abnormal (i.e., $c_1 = c_2 = \cdots = c_K$). Asymptotic optimality of the algorithms holds even when the computation of the indices are based on the Wald's approximation $\hat{\mathbf{E}}^{(n)}(N_k)$ to the actual expected detection time $\mathbf{E}^{(n)}(N_k)$. It should be noted that the techniques used in proving the asymptotic optimality under the full-blown dynamic problem considered in this paper are fundamentally different from those used in [3] under the switching constraint. The proof for the optimality of the OL-$\pi cN$ policy under the restrictive model in [3] is mainly based on an interchange argument, which no longer holds in this fully dynamic problem. In proving the asymptotic optimality of the CL-$\pi cN$ rule under the general model, the key is to show that the average time spent on probing undesired processes (i.e., when noisy observations lead to an inaccurate indication of the process states) does not affect the asymptotic expected cost. This is done in two steps. First, we establish the asymptotic lower bound on the total cost that can be achieved by any policy. Second, by upper bounding the tail of the distribution of some ancillary random times, we show that CL-$\pi cN$ achieves the lower bound in the asymptotic regime.

### B. Related Work

Sequential hypothesis testing was pioneered by Wald in [1] 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 [4]–[8] 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.

There are a number of recent studies on sequential detection involving multiple independent processes for various applications (see, for example, [9]–[16] and references therein). Differing from this work (and our prior work [2], [3]), these studies focus on minimizing the total detection delay, which does not translate to minimizing the total system-wide cost in the anomaly detection problem at hand. The anomaly detection problem also shares similarities with the optimal search and target whereabouts problems as studied in [17]–[20] under a sequential setting and in [21]–[24] under a fixed sample size setting. The design objectives in these studies again differ from that in this paper. The problem of universal outlier hypothesis testing involving a vector of observations containing coordinates with a single outlier distribution was studied in [25], where at each time, all the vector is observed. In our model, however, only a subset of the processes can be observed at a time. Thus, a key parameter when designing a search strategy in our problem is the selection rule which determines which processes we should observe at each given time.

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 [26]. 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 [27]–[34]. The objective is to minimizing the detection delay in [27]–[29], [31]–[34]. A more general model was considered in [30], where sampling incurs a known non-uniform cost across processes and the objective is to minimize the total cost due to the sampling operation (in contrast to our model, where at each given time the cost incurred by abnormal processes whose states have not been identified and depends on the unknown system state). However, all these studies assumed different models than the model considered in this paper.

## II. System Model and Problem Formulation

Consider a system consisting of $K$ processes, where each process may be in a normal state (denoted by $H_0$) or abnormal state (denoted by $H_1$). Each process $k$ is abnormal with a priori probability $\pi_k$, independent of other processes. Each abnormal process $k$ incurs a cost $c_k$ $(0 \leq c_k < \infty)$ per unit time until it is tested and identified. Processes in a normal state do not incur cost. At each given time, only $M$ processes can be probed. We first consider $M = 1$. An extension to $M \geq 1$ is discussed in Section V.

When process $k$ is probed at time $n$, a measurement $y_k(n)$ is drawn independently in a one-at-a-time manner. 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 our results to the case where the densities have unknown parameters.

Let $\phi(n) \in \{1, 2, \ldots, K\}$ be a selection rule, indicating which process is chosen to be tested at time $n$. Let $\mathbf{y}(n) = \{\phi(t), y_{\phi(t)}(t)\}_{t=1}^{n}$ be the set of all the observations and actions up to time $n$. The selection rule $\phi(n)$ is a mapping from $\mathbf{y}(n-1)$ to $\{1, 2, \ldots, K\}$. The vector of selection rules over the time series is denoted by $\boldsymbol{\phi} = (\phi(1), \phi(2), \ldots)$. 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 $\tau_k$ be a stopping time (or a stopping rule), which is the time (counted from the beginning of the entire detection process) when the decision maker stops taking observations from process $k$ and declares its state. The vector of stopping times for the $K$ processes is denoted by $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_K)$. The random sample size required to make a decision regarding the state of process $k$ is denoted by $N_k$. Let $\delta_k \in \{0, 1\}$ be a decision rule, indicating the state declaration of process $k$ at time $\tau_k$. $\delta_k = 0$ 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. The vector of decision rules for the $K$ processes is denoted by $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_K)$.

*Definition 1:* An admissible strategy $\mathbf{s}$ for the sequential anomaly detection problem is given by the tuple $\mathbf{s} = (\boldsymbol{\tau}, \boldsymbol{\delta}, \boldsymbol{\phi})$.

Let

$$\mathcal{H}_0 \triangleq \{k : 1 \le k \le K, \text{process } k \text{ is normal}\},$$
$$\mathcal{H}_1 \triangleq \{k : 1 \le k \le K, \text{process } k \text{ is abnormal}\},$$

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

$$\inf_{\mathbf{s}} \mathbf{E}\left\{\sum_{k\in\mathcal{H}_1} c_k \tau_k\right\}$$
$$\text{subject to } P_k^{FA} \le \alpha_k \quad \forall k = 1, \ldots, K,$$
$$P_k^{MD} \le \beta_k \quad \forall k = 1, \ldots, K, \tag{1}$$

where $P_k^{FA}, P_k^{MD}$ denote the false-alarm and miss-detect error probabilities for process $k$, respectively. We point out that the total expected cost defined in (1) does not include the cost incurred by miss-detected abnormal processes. Since the error constraints are typically required to be small, (1) well approximates the actual loss in practice. For purposes of analysis in Section III.C we assume that the error constraints have the same order (for more details see Appendix VIII).

## III. Anomaly Detection Under Known Observation Models

In this section we derive an asymptotically optimal solution for the anomaly detection problem (1) under the case where the densities $f_k^{(0)}$, $f_k^{(1)}$ are known for all $k$. The proposed probing strategy has a simple closed-loop index form. The index of the currently probed process is updated based on the newly obtained measurement, and the process with the highest index is selected at each given time except a subsequence of time instants that grows exponentially sparse with time. In Section III.B we discuss the computation of the index in details. In Section III.C we show that the proposed algorithm minimizes the total expected cost as the error constraints approach zero.

### A. The CL-$\pi cN$ Policy

In this section we present the CL-$\pi cN$ policy. Let

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

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 $n$ (see (8) for the update of the belief based on a newly obtained measurement). Let $\mathbf{E}^{(n)}(N_k)$ be the expected detection time for process $k$ at time $n$ which dynamically changes due to the changes in the belief $\pi_k(n)$ (see (11)). Let $\hat{\mathbf{E}}^{(n)}(N_k)$ be an approximation to $\mathbf{E}^{(n)}(N_k)$ based on the Wald's approximation (for details see (11) in Section III.B). Define

$$\gamma_k(n) \triangleq \begin{cases} \dfrac{\pi_k(n)c_k}{\hat{\mathbf{E}}^{(n)}(N_k)}, & \text{if } k \in \mathcal{K}(n), \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

Let $\mathcal{N}_s = \{n_1, n_2, \ldots\}$ be a set of time instants that grows exponentially sparse with time (i.e., the cardinality of $\mathcal{N}_s$ grows at a logarithmic rate with time). The CL-$\pi cN$ policy selects the process with the highest index $\gamma_k(n)$ at all times except at time instants in $\mathcal{N}_s$. During the subsequence $\mathcal{N}_s$, all processes whose states have not been declared are probed in a round-robin fashion. Specifically,

$$\phi(n) = \begin{cases} \arg\max_k \gamma_k(n), & \text{if } n \notin \mathcal{N}_s, \\ r(n), & \text{if } n = n_i \quad \forall i = 2, 3, \ldots. \end{cases} \tag{5}$$

At time instants $\mathcal{N}_s$, the function $r(n)$ satisfies:

$$r(n_i) = [(\phi(n_{i-1}) + u(n_i)) \bmod K] + 1, \tag{6}$$

where $u(n_i) = \min(0, 1, \ldots, K-1)$ such that $r(n_i) \in \mathcal{K}(n_i)$, mod denotes the modulo operator, and $\phi(n_1) = 1$. Note that processes are no longer probed once their state has been declared. The round-robin probing subsequence $\mathcal{N}_s$ is to ensure all processes are sufficiently explored. We set[1] $\mathcal{N}_s = \{\lceil \zeta^\ell \rceil\}_{\ell=1}^{\infty}$, where $\zeta > 1$ is a design parameter (for details see Section III.B). We point out that this idea of introducing an exploration subsequence to ensure sufficient learning has also been used in [29], [35].

[1]Note that duplicate values in $\mathcal{N}_s$ are removed.

Following the Wald's SPRT [1], $S_{\phi(n)}(n)$ is compared to boundary values $A_{\phi(n)}, B_{\phi(n)}$ as follows:

- If $S_{\phi(n)}(n) \in (A_{\phi(n)}, B_{\phi(n)})$, then $\phi(n) \in \mathcal{K}(n+1)$ (i.e., continue to take observations from process $\phi(n)$ according to the selection rule (5) at time $n+1$).
- If $S_{\phi(n)}(n) \geq B_{\phi(n)}$, 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 $n' > n$).
- If $S_{\phi(n)}(n) \leq A_{\phi(n)}$, 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 $n' > n$).

The boundary values $A_k$ and $B_k$ are determined such that the error constraints are satisfied. In general, the exact computation of the boundary values is very laborious under the finite regime. Nevertheless, Wald's approximation can be applied to simplify the computation [1]:

$$A_k \approx \log\left(\frac{\beta_k}{1-\alpha_k}\right),$$
$$B_k \approx \log\left(\frac{1-\beta_k}{\alpha_k}\right). \tag{7}$$

Wald's approximation performs well for small $\alpha_k, \beta_k$ 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 [1].

Note that CL-$\pi cN$ is a closed-loop strategy, where the index $\gamma_k(n)$ is updated at each given time based on past observations and actions and the next process is selected accordingly. It can be seen that CL-$\pi cN$ handles the well-known trade-off between exploration and exploitation. The decision maker spends a logarithmic order of time by selecting the processes in a round-robin manner to explore their states and guard against miss-detected abnormal processes. On the other hand, at times $n \notin \mathcal{N}_s$, it exploits the information gathered so far to select the process according to the updated index $\gamma_k(n)$ at time $n$. The index form under the CL-$\pi cN$ policy which dynamically updates the priority of the processes is intuitively satisfying. We should prioritize processes that incur higher costs to the system when abnormal. Furthermore, the priority of a process should be increased as the updated belief of it being abnormal increases during the detection process. It is also desirable to place processes that require longer testing time toward the end of the testing process since their detection time contributes to the cost of every abnormal process that has not been identified. Thus, the priority of a process increases as the updated expected detection time decreases. Note that the sequential test uses an SPRT-based method with memory to minimize the expected sample size for every process. When switching back to a previously visited process (say $k$) at time $n$, the sequential test uses the sum LLRs $S_k(n)$ in decision making to exploit all past observations obtained during previous visits.

### B. Implementation

In this section we discuss the implementation of the proposed policy. At each time $n$, the decision maker updates the indices and the sum LLRs for the currently probed processes, and also sorts the indices for selecting the next process. Sorting the indices can be done in $O(K \log K)$ time via a sorting algorithm. Updating the indices and the sum LLRs (for the general case where $M$ processes are probed at a time) requires $O(M)$ time.

We now consider the computation of the index $\gamma_k(n) = \pi_k(n)c_k/\mathbf{E}^{(n)}(N_k)$. The posterior probability of process $k$ being abnormal can be updated at time $n+1$ based on the Bayes rule:

$$\pi_k(n+1) = (1 - \mathbf{1}_k(n))\pi_k(n) + \frac{\mathbf{1}_k(n)\pi_k(n)f_k^{(1)}(y_k(n))}{\pi_k(n)f_k^{(1)}(y_k(n)) + (1-\pi_k(n))f_k^{(0)}(y_k(n))}. \tag{8}$$

Note that at time $n+1$, only the index of the process that was probed at time $n$ needs to be updated. The expected sample size $\mathbf{E}^{(n)}(N_k)$ at time $n$ depends on the current belief value:

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

where $\mathbf{E}(N_k|H_i)$ is the expected detection time 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 regime. However, Wald's approximation can be applied to simplify the computation [1]:

$$\hat{\mathbf{E}}(N_k|H_0) \triangleq \frac{-(1-\alpha_k)A_k - \alpha_k B_k}{D(f_k^{(0)}||f_k^{(1)})},$$
$$\hat{\mathbf{E}}(N_k|H_1) \triangleq \frac{(1-\beta_k)B_k + \beta_k A_k}{D(f_k^{(1)}||f_k^{(0)})}, \tag{10}$$

where $D(f_k^{(i)}||f_k^{(j)}) = \mathbf{E}_i(\log\frac{f_k^{(i)}(y_k(1))}{f_k^{(j)}(y_k(1))})$ 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 $\alpha_k, \beta_k$. Thus, the approximation to the expected detection time under CL-$\pi cN$ is computed as follows:

$$\hat{\mathbf{E}}^{(n)}(N_k) = \pi_k(n)\hat{\mathbf{E}}(N_k|H_1) + (1-\pi_k(n))\hat{\mathbf{E}}(N_k|H_0). \tag{11}$$

We point out that asymptotic optimality of the probing strategy is preserved as long as the required *order* of the indices is preserved. Therefore, computing the exact expected remaining detection time of a process during a sequential test is not required. Using the Wald's approximation to the entire detection time when computing the indices at each given time is sufficient for obtaining asymptotic optimality.

Next, we discuss the design parameter $\zeta > 1$ used in the exploration subsequence $\mathcal{N}_s$. Note that as $\zeta$ approaches 1, the round-robin selection rule is executed more frequently. It is shown in Appendix VIII that asymptotic optimality of CL-$\pi cN$ holds when $\zeta$ is set sufficiently close to 1 to ensure that the round-robin probing gathers sufficient information so that the index $\gamma_k(n)$ is a sufficiently accurate indication of the process state. In the finite regime, however, $\zeta$ must be designed judiciously for better performance. Intuitively speaking, one should increase $\zeta$ as the sample sizes required to declare the process states decrease to reduce the time spent during the round-robin selection rule. For instance, consider the extreme case where only a single observation is required to declare the process states (i.e., the KL divergences between the observation distributions

are sufficiently large). Therefore, switching between processes is done only when the state of the currently probed process is declared. In this extreme case, the optimal probing strategy is to test the processes in decreasing order of $\pi_k c_k$. Hence, it is desirable to set $\zeta$ sufficiently high in that case so that only the first line in (5) will be executed to obtain optimal performance.

### C. Performance Analysis

In this section we analyze the performance of the CL-$\pi cN$ policy. Let

$$P_e^{\max} \triangleq \max(\alpha_1, \beta_1, \ldots, \alpha_K, \beta_K). \tag{12}$$

The following theorem shows that CL-$\pi cN$ is asymptotically optimal in terms of minimizing the expected cost as the error constraints approach zero. When deriving asymptotic we assume mild conditions on the error constraints, as discussed in Appendix VIII.

*Theorem 1:* Let $\mathbf{E}(C^*), \mathbf{E}(C(\mathbf{s}))$ be the expected costs under CL-$\pi cN$ and any other policy $\mathbf{s}$, respectively. Then[2] ,

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

*Proof:* See Appendix VIII.A. ■

It should be noted that asymptotic optimality does not require the exact computation of the detection times of the processes $\mathbf{E}^{(n)}(N_k)$ when evaluating the indices under CL-$\pi cN$. Computing the indices as defined in (4) using the Wald's approximation in (10), (11) is sufficient to achieve asymptotic optimality (see Appendix VIII.A for details).

## IV. Anomaly Detection Under Unknown Observation Models

In the previous section we focused on the case where the densities under both hypotheses are known. For that case, the sum LLRs was used by every process to design stopping and decision rules based on Wald's SPRT which minimizes the expected sample size for detection. In this section we consider the case where the densities have unknown parameters. While the SPRT applies to the latter case as well with minor modifications, it is highly sub-optimal in general. Therefore, in what follows we focus on asymptotically optimal tests in terms of minimizing the sample size as the error probability approaches zero.

Let $\theta_k$ be an unknown parameter (or a vector of unknown parameters) of process $k$. The observations $\{y_k(i)\}_{i\geq 1}$ are drawn from a common density $f_k(y|\theta_k)$, $\theta_k \in \Theta_k$, where $\Theta_k$ 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 $\Theta_k$, where $I_k = \Theta_k \backslash (\Theta_k^{(0)} \cup \Theta_k^{(1)}) \neq \emptyset$ is an indifference region[3] . When $\theta_k \in I_k$, the detector is indifferent regarding the state of process $k$. Hence, there are no constraints on the error probabilities for all $\theta_k \in I_k$. The hypothesis test regarding process $k$ is to test $\theta_k \in \Theta_k^{(0)}$ against $\theta_k \in \Theta_k^{(1)}$. Reducing $I_k$ increases the sample size.

Asymptotically optimal sequential tests for a single process have been widely studied in the literature, where the key 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 $H_1$. It is assumed that regularity conditions on the distribution hold to guarantee consistency of the MLE [36]. One way to perform the sequential test is to use the Generalized Likelihood Ratio (GLR) statistics. Let $\mathbf{y}_k(n) = (y_k(1), \ldots, y_k(n))$ be the vector of observations for process $k$ by time $n$. For $i, j \in \{0, 1\}$ and $i \neq j$, let

$$S_k^{(i),GLR}(n) = \sum_{r=1}^{n} \log \frac{f_k(y_k(r)|\hat{\theta}_k(n))}{f_k(y_k(r)|\hat{\theta}_k^{(j)}(n))} \tag{14}$$

be the GLR statistics used to declare hypothesis $H_i$ (i.e., reject hypothesis $H_j$) at stage $n$, where $\hat{\theta}_k(n) = \arg\max_{\theta_k \in \Theta_k} f_k(\mathbf{y}_k(n)|\theta_k)$ and $\hat{\theta}_k^{(j)}(n) = \arg\max_{\theta_k \in \Theta_k^{(j)}} f_k(\mathbf{y}_k(n)|\theta_k)$ are the Maximum-Likelihood (ML) estimates of the parameters over the parameter spaces $\Theta_k$ and $\Theta_k^{(j)}$ at stage $n$, respectively.

Another way is to use the Adaptive Likelihood Ratio (ALR) statistics. For $i, j \in \{0, 1\}$ and $i \neq j$, let

$$S_k^{(i),ALR}(n) = \sum_{r=1}^{n} \log \frac{f_k(y_k(r)|\hat{\theta}_k(r-1))}{f_k(y_k(r)|\hat{\theta}_k^{(j)}(n))} \tag{15}$$

be the ALR statistics used to declare hypothesis $H_i$ at stage $n$. Let $S_k^{(i)}(n)$ be the chosen statistics and let

$$N_k^{(i)} = \inf\left\{n : S_k^{(i)}(n) \geq B_k^{(i)}\right\} \tag{16}$$

be the stopping rule used to declare hypothesis $H_i$, where $B_k^{(i)}$ is the boundary value. For each process $k$, the decision maker stops the sampling when $N_k = \min\{N_k^{(0)}, N_k^{(1)}\}$. If $N_k = N_k^{(0)}$, process $k$ is declared as normal. If $N_k = N_k^{(1)}$, process $k$ is declared as abnormal. The advantage of using the ALR statistics is that setting $B_k^{(0)} = \log \frac{1}{\alpha_k}$, $B_k^{(1)} = \log \frac{1}{\beta_k}$ satisfies the error probability constraints in (1). However, such a simple setting cannot be applied when using the GLR statistics. Thus, implementing sequential tests using the ALR statistics is much simpler than using the GLR statistics. The disadvantage of using the ALR statistics is that poor early estimates (from a small number of observations) can never be revised even after a large number of observations have been collected. For more details on sequential tests involving densities with unknown parameters, the reader is referred to [4]–[7].

### A. The CL-$\pi cN$ Policy

With some modifications, the CL-$\pi cN$ policy proposed in Section III can be applied to the case with unknown observation models. Let $S_k^{(i)}(n)$ be the GLR (14) or ALR (15) statistics used in the test. Define

$$\hat{\gamma}_k(n) \triangleq \begin{cases} \dfrac{\hat{\pi}_k(n) c_k}{\hat{\mathbf{E}}^{(n)}(N_k)}, & \text{if } k \in \mathcal{K}(n), \\ 0, & \text{otherwise}, \end{cases} \tag{17}$$

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

[3]The assumption of an indifference region is widely used in the theory of sequential hypothesis testing to derive asymptotically optimal performance. Nevertheless, in some cases this assumption can be removed. For more details, the reader is referred to [5].

where $\hat{\pi}_k(n)$ denotes the estimated posterior probability of process $k$ being abnormal and $\hat{\mathbf{E}}^{(n)}(N_k)$ the updated expected detection time for process $k$ at time $n$ (see Section IV.B for the computation of the index). Similar to (5), the selection rule is given by:

$$\phi(n) = \begin{cases} \arg\max_k \hat{\gamma}_k(n), & \text{if } n \notin \mathcal{N}_s, \\ r(n), & \text{if } n = n_i \quad \forall i = 2, 3, \ldots, \end{cases} \tag{18}$$

where $r(n_i)$ is given in (6) and $\phi(n_1) = 1$. Then, $S^{(i)}_{\phi(n)}(n)$ is compared to boundary values $B^{(0)}_{\phi(n)}, B^{(1)}_{\phi(n)}$ as follows:

- If $S^{(0)}_{\phi(n)}(n) < B^{(0)}_{\phi(n)}$ and $S^{(1)}_{\phi(n)}(n) < B^{(1)}_{\phi(n)}$, then $\phi(n) \in \mathcal{K}(n+1)$ (i.e., continue to take observations from process $\phi(n)$ according to the selection rule (18) at time $n+1$).
- If $S^{(1)}_{\phi(n)}(n) \geq B^{(1)}_{\phi(n)}$, 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 $n' > n$).
- If $S^{(0)}_{\phi(n)}(n) \geq B^{(0)}_{\phi(n)}$, 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 $n' > n$).

### B. Implementation

In this section we discuss the implementation of the proposed policy when the densities have unknown parameters. At each time $n$, the decision maker updates the indices and the GLR/ALR statistics for the currently probed processes (i.e., $M$ processes in general), and also sorts the indices for selecting the next process. Sorting the indices can be done by $O(K \log K)$ time via a sorting algorithm. Note that when the densities have unknown parameters, the updated belief must be computed with respect to the current MLE. In cases where the unknown parameters can take a small number $L$ of values, the decision maker can update and store the beliefs for the $L$ values. Thus, $O(LM)$ time is required instead of $O(M)$. However, if the support has infinite values, then the index must be computed at each time $n$ using the past $n$ observations, which generally requires $O(Mn)$ time (unless a quantization on the support is applied). In general, the estimated belief of process $k$ can be updated at time $n+1$ as follows:

$$\hat{\pi}_k(n+1) = (1 - \mathbf{1}_k(n))\, \hat{\pi}_k(n) + \frac{\mathbf{1}_k(n) \hat{\pi}_k(n) \hat{f}^{(1)}_k(y_k(n))}{\hat{\pi}_k(n) \hat{f}^{(1)}_k(y_k(n)) + (1 - \hat{\pi}_k(n))\, \hat{f}^{(0)}_k(y_k(n))}, \tag{19}$$

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 $1 \leq r \leq n$. Note that computing $\hat{\pi}_k(n+1)$ at time $n+1$ requires $n$ computations with the current ML estimate of the parameter.

In general, it is difficult to obtain a closed-form expression for $\hat{\mathbf{E}}^{(n)}(N_k)$ under the finite regime. However, we can use the asymptotic property of the sequential tests to obtain a closed-form approximation to $\hat{\mathbf{E}}^{(n)}(N_k)$ 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:

$$\hat{\mathbf{E}}^{(n)}(N_k) = \begin{cases} \frac{B^{(0)}_k}{D_k\left(\hat{\theta}_k(n)||\Theta^{(1)}_k\right)}, & \text{if } \hat{\theta}_k(n)) \in \Theta^{(0)}_k, \\ \frac{B^{(1)}_k}{D_k\left(\hat{\theta}_k(n)||\Theta^{(0)}_k\right)}, & \text{if } \hat{\theta}_k(n)) \in \Theta^{(1)}_k, \end{cases} \tag{20}$$

which is guaranteed to be the asymptotic sample size under various families of distributions with unknown parameters (e.g., exponential, multi-variate distributions and general distributions when the unknown parameters can take a finite number of values) as the error probabilities approach zero [5]–[7], [26], [28].

It should be noted that implementing the open-loop policy OL-$\pi cN$ [3] when the densities have unknown parameters requires a priori knowledge of the parameter's distribution (since the testing order is predetermined and switching between processes is allowed only when the state of the currently probed process is declared). However, under CL-$\pi cN$, the testing order is updated dynamically depending on all past observations and actions. As a result, estimating the detection time at time $n$ does not require a priori knowledge of $\theta_k$ since $\hat{\theta}_k(n)$ converges to its true value.

### C. Performance Analysis

The following theorem shows that the proposed policy is asymptotically optimal in terms of minimizing the expected cost as the error probability approaches zero. For purposes of analysis we consider the model in [26], where $\theta_k$ can take only a finite number of values.

*Theorem 2:* Let $\mathbf{E}(C^*), \mathbf{E}(C(\mathbf{s}))$ be the expected costs under CL-$\pi cN$ and any other policy $\mathbf{s}$, respectively. Then,

$$\mathbf{E}(C^*) \sim \inf_{\mathbf{s}} \mathbf{E}\left(C(\mathbf{s})\right) \text{ as } P_e^{\max} \to 0. \tag{21}$$

*Proof:* See Appendix VIII.B. ∎

## V. Extension to Multi-Process Probing

In this section we extend the results reported in the previous sections to the case where more than one process can be probed simultaneously (i.e., $M \geq 1$). For the ease of presentation, we will focus on the case where the observation models are known. However, the results apply to the case where the densities have unknown parameters.

Let $\boldsymbol{\sigma}(n) = (\sigma_1(n), \ldots, \sigma_K(n))$ be a permutation of $\{1, \ldots, K\}$ at time $n$ such that:

$$\gamma_{\sigma_1(n)}(n) \geq \gamma_{\sigma_2(n)}(n) \geq \cdots \geq \gamma_{\sigma_K(n)}(n). \tag{22}$$

The CL-$\pi cN$ policy selects the processes with the $M$ highest indices at all times except times $\mathcal{N}_s$ at which processes are probed in a round-robin manner, i.e.,

$$\phi(n) = \begin{cases} (\sigma_1(n), \ldots, \sigma_M(n)), & \text{if } n \notin \mathcal{N}_s, \\ (r_1(n), \ldots, r_M(n)), & \text{if } n = n_i \quad \forall i = 2, 3, \ldots. \end{cases} \tag{23}$$

At time instants $\mathcal{N}_s$, the functions $(r_1(n), \ldots, r_M(n))$ select the processes whose states have not been declared by time $n$ in a around-robin manner and are given recursively by:

$$r_1(n_i) = [(r_M(n_{i-1}) + u_1(n_i)) \mod K] + 1,$$
$$r_m(n_i) = [(r_{m-1}(n_i) + u_m(n_i)) \mod K] + 1, \quad i = 2, \ldots, M, \tag{24}$$

where $u_m(n_i) = \min(0, 1, \ldots, K - m)$ such that $r_m(n_i) \in \mathcal{K}(n_i)$, mod denotes the modulo operator, and $r_m(n_1) = m$. If there is no solution to $r_m(n_i)$ (i.e., when $|\mathcal{K}(n_i)| < M$), then $r_m(n_i)$ remains empty. Then, sequential tests with memory are executed for the selected processes as described in the previous sections. The following theorem shows that if $c_k = c_{k'}$ holds for all $1 \leq k, k' \leq K$, then CL-$\pi cN$ is asymptotically optimal.

*Theorem 3:* Assume that $c_k = c_{k'}$ holds for all $1 \leq k, k' \leq K$. Let $\mathbf{E}(C^*), \mathbf{E}(C(\mathbf{s}))$ be the expected costs under CL-$\pi cN$ and any other policy $\mathbf{s}$, respectively. Then,

$$\mathbf{E}(C^*) \sim \inf_{\mathbf{s}} \mathbf{E}(C(\mathbf{s})) \text{ as } P_e^{\max} \to 0. \tag{25}$$

*Proof:* See Appendix VIII.C. ■

## VI. Numerical Examples

In this section we present numerical examples to illustrate the performance of the proposed CL-$\pi cN$ policy. We test the following hypotheses: under normal state, the observations from process $k$ follow Poisson distribution $y_k(n) \sim \text{Poi}(\theta_k^{(0)})$, where under abnormal state the observations follow Poisson distribution $y_k(n) \sim \text{Poi}(\theta_k^{(1)})$. This model applies to cyber-systems, where the observations from a probed component represent packet arrival rate under normal state or under reduction of quality attacks as in [37]. We compare the optimal open-loop probing strategy OL-$\pi cN$ developed in [3] with CL-$\pi cN$. We set the following parameters unless otherwise specified: $c_k = \theta_k^{(0)}$ (i.e., the cost represents the normal expected traffic over the component). Thus, in this setting minimizing the total expected cost minimizes the maximal damage to the network in terms of the expected number of failed packets during a denial of service attack. Only a single component is probed at a time (i.e., $M = 1$). The design parameter for the round-robin exploration is set to $\zeta = 1.7$. The error constraints are set to $P_k^{FA} = 10^{-3}, P_k^{MD} = 10^{-6}$ and the a priori probabilities of the components being abnormal are set to $\pi_k = 0.5$ for all $k$.

First, we simulate the case where $\theta_k^{(0)}$ are equally spaced in the interval $[10, 20]$, where $\theta_k^{(1)} = 1.5 \cdot \theta_k^{(0)}$ with probability 0.5 and $\theta_k^{(1)} = 1.2 \cdot \theta_k^{(0)}$ with probability 0.5. This models the situation where both strong and weak deviations from the normal state may occur. We implemented CL-$\pi cN$ under densities with unknown parameters (i.e., the level of deviation from the normal state in this scenario) as described in Section IV. The performance of the algorithms is presented in Fig. 2. It can be seen that CL-$\pi cN$ saves roughly 40% of the average total cost as compared to OL-$\pi cN$. Second, we simulate the case where $M = 5$ components are probed at a time. We set $\theta_k^{(0)} = 10$ for $k = 1, 2, \ldots, K/2, \theta_k^{(0)} = 20$ for $k = K/2 + 1, K/2 + 2, \ldots, K$

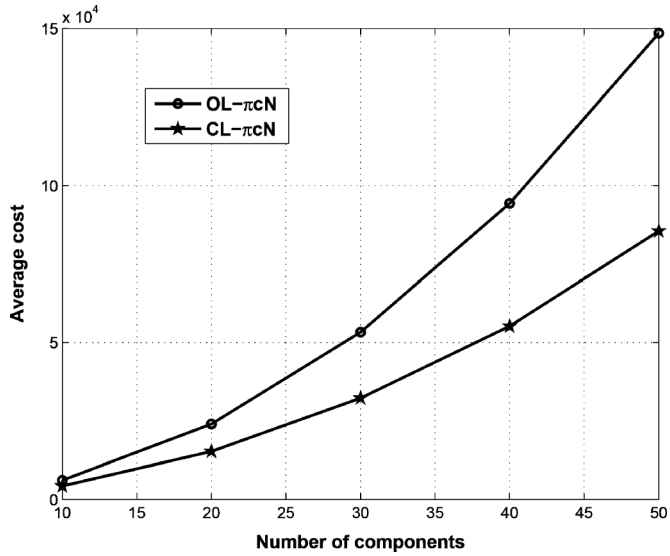


Fig. 1. The average total cost as a function of the number of components. A case where both strong and weak deviations from the normal state may occur with equal probability.

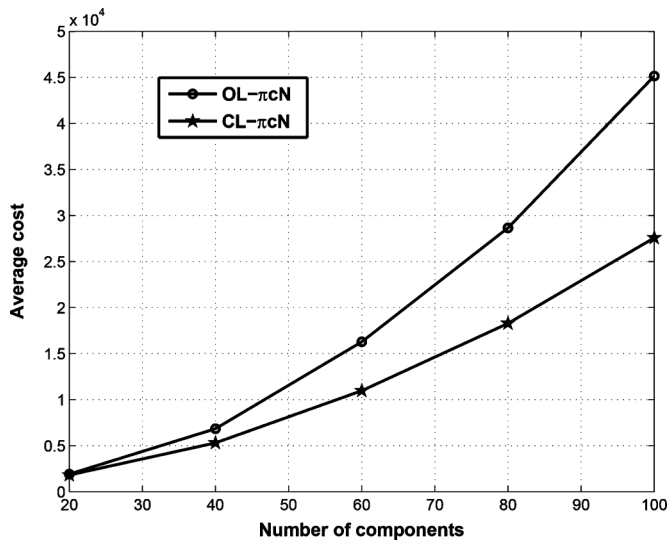


Fig. 2. The average total cost as a function of the number of components. A case where $M = 5$ components are probed at a time.

and $\theta_k^{(1)} = 1.5 \cdot \theta_k^{(0)}$. Note that in that case, asymptotic optimality is an open question due to different costs across the processes. The CL-$\pi cN$ is implemented via multi-process probing as described in Section V. The performance of the algorithms is presented in Fig. 1. It can be seen that CL-$\pi cN$ significantly outperforms OL-$\pi cN$ under this setting as well.

Next, we examine the interesting case where any switching to components $k = 1, \ldots, K/2$ adds a delay $d_1$, while any switching to components $k = K/2 + 1, \ldots, K$ adds a delay $d_2$. This models the situation (as in power systems or communication networks for instance) where monitoring different components requires an initialization process which results in different delays. Note that for any fixed delay incurred by switching among components, the CL-$\pi cN$ preserves its optimality in the asymptotic regime. This can be verified by

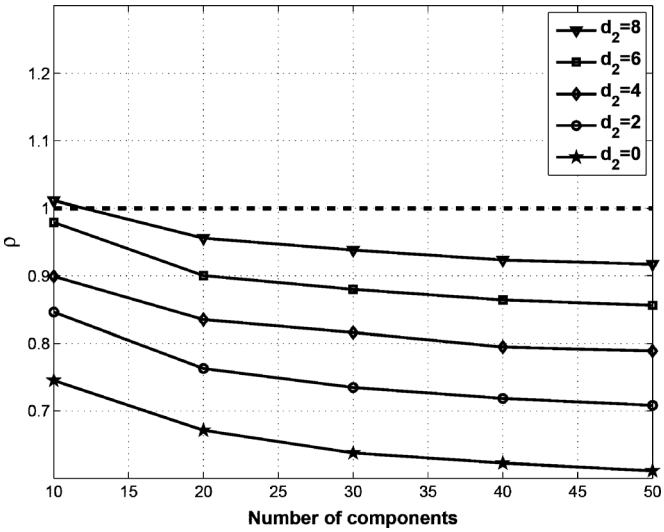


Fig. 3. The gain $\rho = \frac{C_{CL}}{C_{OL}}$ as a function of the number of components and the delay incurred by switching. Switching to components $1, \ldots, K/2$ adds delay $d_1 = 1$ time unit, while switching to components $K/2+1, \ldots, K$ adds delay $d_2$, which ranges between 0 to 8 time units. The CL-$\pi cN$ policy outperforms the OL-$\pi cN$ policy for all $\rho \leq 1$.

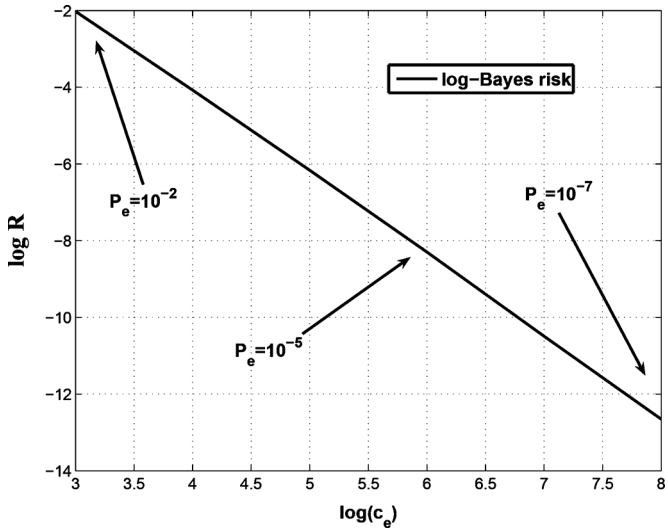


Fig. 4. The tradeoff curve between the average total cost and the error probabilities (i.e., Bayes risk) as a function of the cost for a wrong declaration.

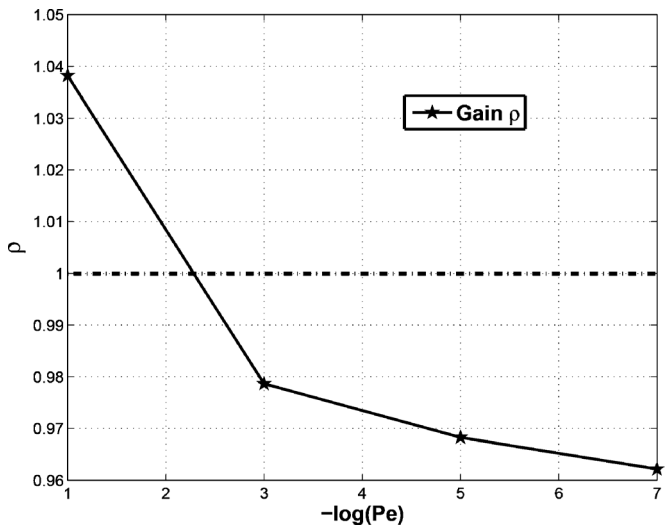


Fig. 5. The gain $\rho = \frac{C_{CL}(\zeta=1.005)}{C_{CL}(\zeta\to\infty)}$ as a function of the error probability for process 1. The CL-$\pi cN$ policy under $\zeta = 1.005$ outperforms the CL-$\pi cN$ policy under $\zeta \to \infty$ for all $\rho \leq 1$.

Lemmas 3, 4 showing that the time spent until the desired asymptotic order is preserved (where switching no longer occurs) is small enough and does not affect the asymptotic expected cost. In the finite regime, however, one should reduce the number of switchings as the delay incurred in switching increases. As discussed in [3], the advantage of OL-$\pi cN$ is that only $K-1$ switchings among components are required. Hence, we expect OL-$\pi cN$ to outperform CL-$\pi cN$ in the finite regime as the delay incurred in switching increases. We set $\theta_k^{(0)} = 10$ for $k = 1, 2, \ldots, K/2, \theta_k^{(0)} = 20$ for $k = K/2+1, K/2+2, \ldots, K$ and $\theta_k^{(1)} = 1.5 \cdot \theta_k^{(0)}$. We set $d_1 = 1$. Let $\rho = \frac{C_{CL}}{C_{OL}}$, where $C_{CL}, C_{OL}$, are the average total costs under CL-$\pi cN$ and OL-$\pi cN$, respectively. The performance of the algorithms is presented in Fig. 3, where $d_2$ ranges between 0 to 8 time units. It can be seen that CL-$\pi cN$ saves roughly 30%–40% of the average total cost as compared to OL-$\pi cN$ when $d_2 = 0$. On the other hand, OL-$\pi cN$ may be preferred for $d_2 > 8$.

The next numerical example demonstrates the trade-off curve between the average total cost and the error probabilities (i.e., a Bayes risk) to quantify the threshold effects of the sequential tests. We set $K = 10$ and $\theta_k^{(0)} = 10, \theta_k^{(1)} = 15, c_k = 1, \pi_k = 0.5$ for all $k$. We assign a cost $c_e$ for a wrong declaration and examine the following normalized (by $c_e$) Bayes risk: $R \triangleq \sum_{k\in\mathcal{H}_1} \left[\frac{1}{c_e}\tau_k + \left(P_k^{FA} + P_k^{MD}\right)\right]$. The log-Bayes risk is presented in Fig. 4 as a function of $\log c_e$, with the corresponding error probabilities $P_e$. As expected, as the cost for a wrong declaration $c_e$ increases, the error probability decreases. Note also that the Bayes risk decreases as $c_e$ increases. Intuitively speaking, this result follows from the fact that the minimal sample size under a sequential testing has the order of $\log(c_e)$, and $P_e$ has the order of $1/c_e$ [26]. Thus, the log-Bayes risk decreases approximately linearly with $\log c_e$ as $c_e$ increases.

Finally, we demonstrate the loss of optimality in the asymptotic regime when the round-robin selection rule is not executed. We set $K = 2, \theta_1^{(0)} = \theta_2^{(0)} = 10, \theta_1^{(1)} = 10.1, \theta_2^{(1)} = 10.3$ (i.e., small deviations from normal states are required to be detected), $\pi_1 = 0.9, \pi_2 = 0.1, c_1 = c_2 = 1$. We simulated CL-$\pi cN$ under $\zeta = 1.005$ (i.e., the round-robin scheduling is executed very frequently) and $\zeta \to \infty$ (i.e., the round-robin scheduling is not executed). Let $\rho = \frac{C_{CL}(\zeta=1.005)}{C_{CL}(\zeta\to\infty)}$, where $C_{CL}(\zeta = 1.005)$ and $C_{CL}(\zeta \to \infty)$ are the average total costs under CL-$\pi cN$ with $\zeta = 1.005$ and $\zeta \to \infty$, respectively. The performance of the algorithms as a function of the error probability for process 1 is presented in Fig. 5. The error probability for process 2 was set such that $\gamma_1(1) = 2\gamma_2(1)$ holds. It can be seen that setting $\zeta = 1.005$ outperforms $\zeta \to \infty$ as the error probability decreases. This result demonstrates the significance of the round-robin selection rule to guarantee optimality in the asymptotic regime. It should be noted, however, that the loss

by removing the round-robin scheduling (i.e., always setting $\zeta \to \infty$) is small and CL-$\pi cN$ may perform well with $\zeta \to \infty$ under typical error probabilities.

## VII. Conclusion

The problem of sequential detection of independent anomalous processes among $K$ processes was considered. At each time, only a subset of the processes can be observed, and the observations from each chosen process follow two different distributions, depending on whether the process is normal or abnormal. Each anomalous process incurs a cost per unit time until it is identified. The objective is a sequential search strategy that minimizes the total expected cost incurred by all the processes during the entire detection process, under reliability constraints. Asymptotically optimal closed-loop policies were developed and strong performance in finite regime was demonstrated via simulations as compared to the optimal open-loop policies when the cost incurred by switching across processes is not too high.

## Appendix

In this Appendix we prove the asymptotic optimality of the proposed tests as the error constraints approach zero. For purposes of analysis, we assume that the asymptotic expected sample sizes $\mathbf{E}(N_k|H_0), \mathbf{E}(N_{k'}|H_1)$ have the same order for all $k, k'$. This condition implies that $\log(P_k^{FA})/\log(P_{k'}^{MD})$ is bounded away from zero and infinity for every pair $k, k'$. Throughout the proof, we use the fact that the round-robin selection rule (i.e., second line in (5)) observes all the processes according to a predetermined order at times $n = \lceil \zeta^\ell \rceil$, for $\ell = 1, 2, \ldots$, where $\zeta$ is a design parameter. We will show that asymptotic optimality holds when $\zeta$ is set sufficiently close to 1.

Deriving asymptotic optimality is done in two steps. First, we establish the asymptotic lower bound on the total cost that can be achieved by any policy. Second, we show that CL-$\pi cN$ achieves the lower bound in the asymptotic regime. The key in proving the second step is to upper bound the tail of the distribution of some ancillary random times. Specifically, when CL-$\pi cN$ is implemented indefinitely (i.e., CL-$\pi cN$ probes the processes indefinitely according to its selection rule, while the stopping rules and decision rules are disregarded), we can define an event $T_1$ in which for all $n \geq T_1$, the index $\gamma_k(n)$ is a sufficient indication to the process state. The event $T_1$ depends on the future and the true state, and is not a stopping time. The decision maker does not know whether it has arrived. However, we show that $T_1$ is sufficiently small. As a result, we show that when CL-$\pi cN$ is implemented in the asymptotic regime ($P_e^{\max} \to 0$ and thus the detection time approaches infinity), the cost incurred by abnormal processes during the first $T_1$ time units does not affect the asymptotic total expected cost.

### A. Proof of Theorem 1

In this section we prove the asymptotic optimality of CL-$\pi cN$ under the case where the densities are completely known. Note that the SPRT's boundary values (used to test every process) satisfy $B_k = -\log(\alpha_k), A_k = -\log(\beta_k)$ in the asymptotic regime [8]. Let $\mathbf{E}^*(N_k|H_i)$ be the expected sample size for process $k$ under the SPRT. Without loss of generality we assume that $\mathcal{H}_1 = \{1, 2, \ldots, K_1\}, \mathcal{H}_0 = \{K_1+1, K_1+2, \ldots, K\}$ and[4]

$$\frac{c_1}{\mathbf{E}^*(N_1|H_1)} > \frac{c_2}{\mathbf{E}^*(N_2|H_1)} > \cdots > \frac{c_{K_1}}{\mathbf{E}^*(N_{K_1}|H_1)}. \tag{26}$$

where $\mathbf{E}^*(N_k|H_1) \to \frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}$ as $P_e^{\max} \to 0$. Note that the Wald's approximation to the expected detection time in (10) satisfies $\hat{\mathbf{E}}(N_k|H_0) \to \frac{-A_k}{D(f_k^{(0)}||f_k^{(1)})}, \hat{\mathbf{E}}(N_k|H_1) \to \frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}$ as $P_e^{\max} \to 0$. Thus, the approximation to the expected detection time used in CL-$\pi cN$ in (11) approaches

$$\hat{\mathbf{E}}^{(n)}(N_k) \to \pi_k(n)\frac{B_k}{D\left(f_k^{(1)}||f_k^{(0)}\right)} + (1-\pi_k(n))\frac{-A_k}{D\left(f_k^{(0)}||f_k^{(1)}\right)} \tag{27}$$

as $P_e^{\max} \to 0$.

Since the indices under CL-$\pi cN$ are given by $\gamma_k(n) = \pi_k(n)c_k/\hat{\mathbf{E}}^{(n)}(N_k)$ and we are interested in establishing optimality as $P_e^{\max} \to 0$, it suffices to prove the theorem when the indices are evaluated as:

$$\gamma_k(n) = \frac{\pi_k(n)c_k}{\pi_k(n)\frac{B_k}{D\left(f_k^{(1)}||f_k^{(0)}\right)} + (1-\pi_k(n))\frac{-A_k}{D\left(f_k^{(0)}||f_k^{(1)}\right)}}, \tag{28}$$

It should be noted that the proof holds under any computation of the indices that approaches (28) as $P_e^{\max} \to 0$. Throughout the paper, we proposed to use the Wald's approximation specifically since it performs well in the finite regime and approaches (28) as $P_e^{\max} \to 0$.

The proof is mainly based on Lemmas 1, 4. In lemma 1, we establish the asymptotic lower bound on the expected cost that can be achieved by any policy. Then, Lemma 4 shows that CL-$\pi cN$ achieves the lower bound in the asymptotic regime.

*Lemma 1:* Let $\mathbf{E}(C(s))$ be the total expected cost under policy $s$ that satisfies the error constraints in (1). Then,

$$\inf_s \mathbf{E}\left(C(s)\right) \geq (1-o(1))\sum_{i=1}^{K_1} c_i \sum_{k=1}^{i} \frac{B_k}{D\left(f_k^{(1)}||f_k^{(0)}\right)}, \tag{29}$$

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

*Proof:* Note that observing normal processes before declaring the states of abnormal processes can only increase the total expected cost. Hence, for establishing the lower bound on the actual cost we assume that all the abnormal processes are tested before those in a normal state.

Let $\mathbf{y}_k$ be the vector of observations taken from process $k$ and $\mathbf{y} = (\mathbf{y}_1, \ldots, \mathbf{y}_K)$ be the collection of the observation vectors. Let

$$\mathcal{Y}_\epsilon(s) = \left\{\mathbf{y} : N_k > (1-\epsilon)\frac{B_k}{D(f_k^{(1)}||f_k^{(0)})} \forall k\right\} \tag{30}$$

[4]In cases where processes have the same $c_i/E^*(N_i|H_1)$, we can arbitrarily order them (by computing their index using a modified cost with an additive small noise $\tilde{c}_k = c_k + \epsilon_k$) without affecting the objective function in the asymptotic regime.

be the set of all possible observations collected from the processes with sample sizes satisfying $N_k > (1-\epsilon)\frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}$ for all $k$ under policy $s$. Let $C_{\mathcal{Y}_\epsilon(s)}(\mathbf{y})$ be the total cost incurred by the processes when observations $\mathbf{y} \in \mathcal{Y}_\epsilon(s)$ were taken under policy $s$.

Next, we lower bound $C_{\mathcal{Y}_\epsilon(s)}(\mathbf{y})$. We define a modified vector of observations for process $k, \tilde{\mathbf{y}}_k$ with length $\tilde{N}_k \triangleq (1-\epsilon)\frac{B_k}{D(f_k^{(1)}||f_k^{(0)})} \leq N_k$ by removing observations $\tilde{N}_k+1, \tilde{N}_k+2, \ldots, N_k$ for all $k$. The set $\tilde{\mathcal{Y}}_\epsilon(s)$ is defined accordingly as the set of the modified vectors of observations. Let $C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}})$ be the total cost incurred by the modified vectors of observations, where the selection rule under $s$ skips the time indices that have been removed. As a result, $C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}}) \leq C_{\mathcal{Y}_\epsilon(s)}(\mathbf{y})$.

Following the Smith rule [38], minimizing $C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}})$ is done by ordering the processes in decreasing order of $c_k/\tilde{N}_k$. Since $\mathbf{E}^*(N_k|H_1) \to \tilde{N}_k/(1-\epsilon)$ as $P_e^{\max} \to 0$ [1], we have:

$$\inf_s C_{\mathcal{Y}_\epsilon(s)}(\mathbf{y}) \geq (1-\epsilon)\sum_{i=1}^{K_1} c_i \sum_{k=1}^{i} \frac{B_k}{D\left(f_k^{(1)}||f_k^{(0)}\right)} \quad \text{as } P_e^{\max} \to 0. \tag{31}$$

Finally, we apply ([39], Lemma 2.1), where an asymptotic probabilistic lower bound on the sample size achieved by any test (for a single process) that satisfies specific error constraints was established. The lemma was originally stated for a more general case of $M$-ary hypothesis testing and non-i.i.d. observations. It requires a weaker condition on the convergence of a (variation of) the average LLR than the strong law of large numbers. Therefore, it directly applies to the case of binary hypothesis and i.i.d. observations (i.e., the strong law of large numbers implies the convergence of the average LLR to the corresponding KL divergence), considered in this paper. Specifically, applying ([39], Lemma 2.1, (2.13)) to our model yields:

$$\inf_s \Pr\left(N_k > \frac{(1-\epsilon)B_k}{D\left(f_k^{(1)}||f_k^{(0)}\right)}\right) = 1 \text{ as } P_e^{\max} \to 0 \quad \forall k \in \mathcal{H}_1. \tag{32}$$

Hence, $\Pr(\mathbf{y} \in \mathcal{Y}_\epsilon(s)) = 1$ as $P_e^{\max} \to 0$ for every $\epsilon > 0$, which completes the proof.

For the next lemmas, we assume that CL-$\pi cN$ is implemented and show that CL-$\pi cN$ achieves the asymptotic lower bound on the expected total cost (29) as $P_e^{\max} \to 0$.

*Definition 2:* For every $0 < \epsilon < 1, T_1(\epsilon)$ is defined as the smallest integer such that $\pi_k(n) \geq 1-\epsilon$ for all $k \in \mathcal{H}_1$ and $\pi_k(n) \leq \epsilon$ for all $k \in \mathcal{H}_0$ for all $n \geq T_1(\epsilon)$.

In the following lemma we show that $T_1(\epsilon)$ is sufficiently small.

*Lemma 2:* Assume that CL-$\pi cN$ is implemented indefinitely. Then, for every fixed $0 < \epsilon < 1$ and $\nu > 0$, there exists $\delta > 0$ such that for all $1 < \zeta \leq 1+\delta$ the following holds:

$$\Pr\left(T_1(\epsilon) > n\right) \leq O(n^{-\nu}). \tag{33}$$

*Proof:* Let $d_k \triangleq \frac{1-\pi_k(1)}{\pi_k(1)}$ and

$$M_k^{(1)} \triangleq -\log\left(\frac{\epsilon}{d_k(1-\epsilon)}\right), \quad M_k^{(0)} \triangleq -\log\left(\frac{d_k\epsilon}{1-\epsilon}\right). \tag{34}$$

By rewriting the update formula in (8), it can be shown that:

$$\pi_k(n) = \left(d_k e^{-S_k(n)} + 1\right)^{-1}. \tag{35}$$

As a result, $\pi_k(n) \geq 1-\epsilon$ iff $S_k(n) \geq M_k^{(1)}$ and $\pi_k(n) \leq \epsilon$ iff $S_k(n) \leq -M_k^{(0)}$, where $S_k(n)$ is the sum of i.i.d. r.v (i.e., LLR) with mean $\mathbf{E}(\ell_k(n)) = D(f_k^{(1)}||f_k^{(0)}) > 0$ for all $k \in \mathcal{H}_1$ and $\mathbf{E}(\ell_k(n)) = -D(f_k^{(0)}||f_k^{(1)}) < 0$ for all $k \in \mathcal{H}_0$. Since the round-robin selection guarantees that for large $n, \log n/(K \log\zeta)$ samples are taken from every process up to time $n$, (33) follows for an arbitrarily large $\nu$ following the same argument as in [29] when $\zeta$ is set sufficiently close to 1. ∎

*Definition 3:* $T_1$ is defined as the smallest integer such that $\gamma_1(n) > \gamma_2(n) > \cdots > \gamma_{K_1}(n) > \max_{k\in\mathcal{H}_0}\gamma_k$ for all $n \geq T_1$.

Before presenting the next lemma, we provide an intuition for the definition of $T_1$. Assume that no state has been declared by time $T_1$. Then, $T_1$ represents the earliest time where the testing order required to achieve the asymptotic lower bound (i.e., the order: $1, 2, \ldots, K_1$) is preserved for all $n \geq T_1$. In the following lemma we show that $T_1$ is sufficiently small, such that the cost incurred by abnormal processes during $T_1$ does not affect the asymptotic expected total cost.

*Lemma 3:* Assume that CL-$\pi cN$ is implemented indefinitely. Then, for every fixed $\nu > 0$, there exists $\delta > 0$ such that for all $1 < \zeta \leq 1+\delta$ the following holds:

$$\Pr(T_1 > n) \leq O(n^{-\nu}). \tag{36}$$

*Proof:* Note that Lemma 2 holds for any $0 < \epsilon < 1$ and it is assumed that $\frac{c_1}{\mathbf{E}^*(N_1|H_1)} > \frac{c_2}{\mathbf{E}^*(N_2|H_1)} > \cdots > \frac{c_{K_1}}{\mathbf{E}^*(N_{K_1}|H_1)}$ holds, where $\mathbf{E}^*(N_k|H_1) \to \frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}$ as $P_e^{\max} \to 0$. Since $\gamma_k(n) = \frac{\pi_k(n)c_k}{\pi_k(n)\hat{\mathbf{E}}(N_k|H_1)+(1-\pi_k(n))\hat{\mathbf{E}}(N|H_0)}$, where $\hat{\mathbf{E}}(N_k|H_1) \to \frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}, \hat{\mathbf{E}}(N_k|H_0) \to \frac{-A_k}{D(f_k^{(0)}||f_k^{(1)})}$ have the same order by assumption, we can choose a sufficiently small $\epsilon > 0$ that satisfies the lemma. ∎

In the following lemma we show that the total expected cost under CL-$\pi cN$ approaches the lower bound (29) as $P_e^{\max} \to 0$.

*Lemma 4:* Let $\mathbf{E}(C^*)$ be the total expected cost under CL-$\pi cN$. Then,

$$\mathbf{E}(C^*) \sim \sum_{i=1}^{K_1} c_i \sum_{k=1}^{i} \frac{B_k}{D\left(f_k^{(1)}||f_k^{(0)}\right)} \text{ as } P_e^{\max} \to 0. \tag{37}$$

*Proof:* Without loss of generality, assume that no state has been declared by time $T_1$ (otherwise, the resulting cost is even smaller than the cost computed below). Thus, for all $n \geq T_1$, CL-$\pi cN$ tests the processes in the following order: $1, 2, \ldots, K_1$ and then test the normal ones. Let $\bar{c} = \max_k c_k$. Since the total

cost incurred up to time $T_1$ is upper bounded by $K\bar{c}T_1$, the total cost $C^*$ under CL-$\pi cN$ is upper bounded by

$$C^* \leq K\bar{c}T_1 + K\bar{c}\sum_{k=1}^{K} N_k^s + \sum_{i=1}^{K_1} c_i \sum_{k=1}^{i} \tilde{N}_k, \qquad (38)$$

The term $K\bar{c}T_1$ upper bounds the total cost incurred up to time $T_1$, the term $K\bar{c}\sum_{k=1}^{K} N_k^s$ upper bounds the total cost incurred due to the round-robin scheduling, where $N_k^s$ is the observation sample size due to the round-robin selection rule for process $k$ (i.e., $\mathbf{E}(N_k^s) \leq O(\log B_1)$ in the asymptotic regime since the error probabilities have the same order by assumption). The term $\sum_{i=1}^{K_1} c_i \sum_{k=1}^{i} \tilde{N}_k$ is the total cost incurred for all $n \geq T_1$ (since by the definition of $T_1$ the processes are tested in the following order: $1, 2, \ldots, K_1$), where $\tilde{N}_k$ is the remaining sample size required to declare the state for process $k$ for all $n \geq T_1$. Therefore, applying Lemma 3 and using the fact that $\mathbf{E}(\tilde{N}_k|H_1) \leq \mathbf{E}^*(N_k|H_1) \to \frac{B_k}{D(f_k^{(1)}||f_k^{(0)})}$ as $P_e^{\max} \to 0$ yields:

$$\mathbf{E}(C^*) \leq O(\log B_1) + (1+o(1)) \sum_{i=1}^{K_1} c_i \sum_{k=1}^{i} \frac{B_k}{D\left(f_k^{(1)}||f_k^{(0)}\right)}, \qquad (39)$$

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

Combining (39) and (29) completes the proof. ∎

### B. Proof of Theorem 2

In this section we prove the asymptotic optimality of the proposed policy when the densities have unknown parameters. For purposes of analysis we consider the model in [26], where $\theta_k$ can take only a finite number of values. Throughout the proof we omit steps that use similar arguments as in the proof under the case of completely known densities.

Using a similar argument as in Lemma 1, it can be shown that

$$\inf_s \mathbf{E}(C(s)) \sim \sum_{i=1}^{K_1} c_i \sum_{k=1}^{i} \frac{B_k^{(0)}}{D_k^*\left(\theta_k||\Theta_k^{(0)}\right)} \text{ as } P_e^{\max} \to 0. \qquad (40)$$

Next, we show that CL-$\pi cN$ achieves this bound.

*Definition 4:* $T_{ML}$ is defined as the smallest integer such that $\hat{\theta}_k(n) = \theta_k$ for all $k$ for all $n \geq T_{ML}$.

In the following lemma we show that $T_{ML}$ is sufficiently small.

*Lemma 5:* Assume that CL-$\pi cN$ is implemented indefinitely. Then, for every fixed $\nu > 0$, there exists $\delta > 0$ such that for all $1 < \zeta \leq 1 + \delta$ the following holds:

$$\Pr(T_{ML} > n) \leq O(n^{-\nu}). \qquad (41)$$

*Proof:* Note that when $K = 1$ (i.e., all the observations are taken from a single process), $\Pr(T_{ML} > n)$ decays exponentially with $n$ following the same argument as in [26]. Furthermore, for large $n$, at least $\log n/(K \log \zeta)$ samples are taken from every process by time $n$. Thus, (41) follows when $\zeta$ is set sufficiently close to 1. ∎

*Definition 5:* For every $0 < \epsilon < 1, T_1(\epsilon)$ is defined as the smallest integer such that $\hat{\pi}_k(n) \geq 1 - \epsilon$ for all $k \in \mathcal{H}_1$ and $\hat{\pi}_k(n) \leq \epsilon$ for all $k \in \mathcal{H}_0$ for all $n \geq T_1(\epsilon)$.

In the following lemma we show that $T_1(\epsilon)$ is sufficiently small.

*Lemma 6:* Assume that CL-$\pi cN$ is implemented indefinitely. Then, for every fixed $0 < \epsilon < 1$ and $\nu > 0$, there exists $\delta > 0$ such that for all $1 < \zeta \leq 1 + \delta$ the following holds:

$$\Pr(T_1(\epsilon) > n) \leq O(n^{-\nu}). \qquad (42)$$

*Proof:* Note that:

$$\Pr(T_1(\epsilon) > n) \leq \Pr(T_1(\epsilon) > n, T_{ML} \leq n) + \Pr(T_{ML} > n). \qquad (43)$$

The term $\Pr(T_{ML} > n)$ decays polynomially with $n$ by applying Lemma 5. Thus, it suffices to show that $\Pr(T_1(\epsilon) > n, T_{ML} \leq n)$ decays polynomially with $n$.

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

$$M_k^{(1)} \triangleq -\log\left(\frac{\epsilon}{d_k(1-\epsilon)}\right), \quad M_k^{(0)} \triangleq -\log\left(\frac{d_k\epsilon}{1-\epsilon}\right). \qquad (44)$$

By rewriting the update formula in (8), it can be shown that:

$$\hat{\pi}_k(n) = \left(d_k e^{-S_k^{(1),GLR}(n)} + 1\right)^{-1}, \qquad (45)$$

for all $k \in \mathcal{H}_1$ for all $n \geq T_{ML}$, and

$$\hat{\pi}_k(n) = \left(d_k e^{S_k^{(0),GLR}(n)} + 1\right)^{-1}, \qquad (46)$$

for all $k \in \mathcal{H}_0$ for all $n \geq T_{ML}$.

As a result, $\hat{\pi}_k(n) \geq 1-\epsilon$ iff $S_k^{(1),GLR}(n) \geq M_k^{(1)}$ for all $k \in \mathcal{H}_1$ and $\hat{\pi}_k(n) \leq \epsilon$ iff $S_k^{(0),GLR}(n) \geq M_k^{(0)}$ for all $k \in \mathcal{H}_0$ for all $n \geq T_{ML}$. Thus, it suffices to show that $\Pr(S_k^{(1),GLR}(n) \leq M_k^{(1)}|n \geq T_{ML})$ for all $k \in \mathcal{H}_1$ and $\Pr(S_k^{(0),GLR}(n) \leq M_k^{(0)}|n \geq T_{ML})$ for all $k \in \mathcal{H}_0$ decay polynomially with $n$. Note that when $T_{ML} \leq n$ occurs, $S_k^{(0),GLR}(n)$ for all $k \in \mathcal{H}_1$ and $S_k^{(1),GLR}(n)$ for all $k \in \mathcal{H}_0$ are sums of i.i.d. r.v. with positive KL divergence (since $\hat{\theta}_k(n) = \theta_k$ for all $n \geq T_{ML}$). Since at least $\log n/(K \log \zeta)$ samples are taken from every process by time $n$, the lemma follows. ∎

The rest of the proof follows with minor modifications to the proof under the case of completely known densities.

### C. Proof of Theorem 3

In this Appendix we prove the asymptotic optimality of CL-$\pi cN$ under multi-process probing when $c \triangleq c_1 = c_2 = \cdots = c_K$. Throughout the proof we omit steps that use similar arguments as in the proof under single-process probing. We also use similar notations as in Appendix VIII.A.

First, we establish the asymptotic lower bound on the expected cost that can be achieved by any policy. Using the same notations as in the proof of Lemma 1, we aim to lower-bound $C_{\mathcal{Y}_\epsilon(s)}(\mathbf{y})$ using the definition of $C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}})$. Recall that

$C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}})$ is the total cost incurred by the modified vectors of observations with a fixed sample size.

Next, we apply ([40], Theorem 5.4.2) to minimize $C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}})$. In [40], the problem of ordering jobs with fixed processing times over $M$ parallel machines was considered. It was shown that scheduling the jobs in decreasing order of $1/\tilde{N}_k$, where $\tilde{N}_k$ is the processing time for job $k$, minimizes the sum completion times of the jobs. When applying ([40], Theorem 5.4.2) to our case, the sum completion times for the modified observation vectors is $\frac{1}{c}C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}})$ when all the abnormal processes incur the same cost $c$ per unit time. Since $c = c_1 = \cdots = c_K$ by assumption (and in particular $c = c_1 = \cdots = c_{K_1}$ for any realization of the true system state), we can apply ([40], Theorem 5.4.2). As a result, minimizing $C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}})$ is done by ordering the processes in decreasing order of $1/\tilde{N}_k$. Let

$$\tilde{c}_k = \begin{cases} c, & \text{if } k \in \mathcal{H}_1, \\ 0, & \text{otherwise.} \end{cases} \tag{47}$$

Note that minimizing $C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}})$ by ordering the modified observation vectors in decreasing order of $1/\tilde{N}_k$ implies that at each given time the $M$ vectors with the smallest sample sizes among the remaining vectors contribute to the total cost. As a result, Similar to (31), for any $\epsilon > 0$, we can lower bound the actual cost by the cost achieved by minimizing $C_{\tilde{\mathcal{Y}}_\epsilon(s)}(\tilde{\mathbf{y}})$:

$$\begin{aligned} \inf_s C_{\mathcal{Y}_\epsilon(s)}(\mathbf{y}) \geq (1-\epsilon) \sum_{m=1}^{M} \sum_{i=1}^{\lceil K_1/M \rceil} \tilde{c}_{m+(i-1)M} \\ \times \sum_{k=1}^{i} \frac{B_{m+(k-1)M}}{D\left(f^{(1)}_{m+(k-1)M} \| f^{(0)}_{m+(k-1)M}\right)} \\ \times \text{as } P_e^{\max} \to 0 \end{aligned} \tag{48}$$

, Hence, following the same argument as in Lemma 1, we obtain:

$$\begin{aligned} \inf_s \mathbf{E}(C(s)) \geq (1-o(1)) \sum_{m=1}^{M} \sum_{i=1}^{\lceil K_1/M \rceil} \tilde{c}_{m+(i-1)M} \\ \times \sum_{k=1}^{i} \frac{B_{m+(k-1)M}}{D\left(f^{(1)}_{m+(k-1)M} \| f^{(0)}_{m+(k-1)M}\right)}, \end{aligned} \tag{49}$$

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

Next, we show that CL-$\pi cN$ achieves the lower bound (49) in the asymptotic regime. Following the definition of $T_1$, for all $n \geq T_1$, CL-$\pi cN$ tests the processes in the desired order required to obtain the lower bound as specified in (49). Note that by applying Lemma 3, we can set $\zeta > 1$ sufficiently close to 1, such that $\Pr(T_1 > n) \leq O(n^{-\nu})$ for an arbitrarily large $\nu > 0$. Therefore, similar to (38), (39), we have:

$$\begin{aligned} \mathbf{E}(C^*) \leq (1+o(1)) \sum_{m=1}^{M} \sum_{i=1}^{\lceil K_1/M \rceil} \tilde{c}_{m+(i-1)M} \\ \times \sum_{k=1}^{i} \frac{B_{m+(k-1)M}}{D\left(f^{(1)}_{m+(k-1)M} \| f^{(0)}_{m+(k-1)M}\right)} + O(\log B_1), \end{aligned} \tag{50}$$

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

Combining (49) and (50) completes the proof.

## REFERENCES


[1] A. Wald, *Sequential Analysis*. New York, NY, USA: Wiley, 1947.

[2] K. Cohen, Q. Zhao, and A. Swami, "Optimal index policies for quickest localization of anomaly in cyber networks," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Dec. 2013, pp. 221–224.

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

[4] G. Schwarz, "Asymptotic shapes of Bayes sequential testing regions," *Ann. Math. Statist.*, pp. 224–236, 1962.

[5] T. L. Lai, "Nearly optimal sequential tests of composite hypotheses," *Ann. Statist.*, pp. 856–886, 1988.

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

[7] A. G. Tartakovsky, "An efficient adaptive sequential procedure for detecting targets," in *Proc. IEEE Aerosp. Conf.*, 2002, vol. 4, pp. 1581–1596.

[8] V. Draglin, A. G. Tartakovsky, and V. V. Veeravalli, "Multihypothesis sequential probability ratio tests—Part I: Asymptotic optimality," *IEEE Trans. Inf. Theory*, vol. 45, no. 7, pp. 2448–2461, 1999.

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

[10] H. Li, "Restless watchdog: Selective quickest spectrum sensing in multichannel cognitive radio systems," *EURASIP J. Adv. Signal Process.*, vol. 2009, 2009, DOI:10.1155/2009/417457.

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

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

[13] M. L. Malloy and R. D. Nowak, "Sequential testing for sparse recovery," [Online]. Available: http://arxiv.org/abs/1212.1801 [arXiv:1212.1801 [cs.IT]

[14] M. L. Malloy, G. Tang, and R. D. Nowak, "Quickest search for a rare distribution," in *Proc. IEEE Ann. Conf. Inf. Sci. Syst.*, 2012, pp. 1–6.

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

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

[17] K. S. Zigangirov, "On a problem in optimal scanning," *Theory Probabil. Appl.*, vol. 11, no. 2, pp. 294–298, 1966.

[18] E. Klimko and J. Yackel, "Optimal search strategies for Wiener processes," *Stoch. Process. .Appl.*, vol. 3, no. 1, pp. 19–33, 1975.

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

[20] L. D. Stone and J. A. Stanshine, "Optimal search using uninterrupted contact investigation," *SIAM J. Appl. Math.*, vol. 20, no. 2, pp. 241–263, 1971.

[21] K. P. Tognetti, "An optimal strategy for a whereabouts search," *Oper. Res.*, vol. 16, no. 1, pp. 209–211, 1968.

[22] J. B. Kadane, "Optimal whereabouts search," *Oper. Res.*, vol. 19, no. 4, pp. 894–904, 1971.

[23] Y. Zhai and Q. Zhao, "Dynamic search under false alarms," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Dec. 2013, pp. 201–204.

[24] D. A. Castanon, "Optimal search strategies in dynamic hypothesis testing," *IEEE Trans. Syst., Man, Cybern.*, vol. 25, no. 7, pp. 1130–1138, 1995.

[25] Y. Li, S. Nitinawarat, and V. V. Veeravalli, "Universal outlier hypothesis testing," 2013 [Online]. Available: http://arxiv.org/abs/1302.4776

[26] H. Chernoff, "Sequential design of experiments," *Ann. Math. Statist.*, vol. 30, no. 3, pp. 755–770, 1959.

[27] S. Bessler, "Theory and applications of the sequential design of experiments, K-actions and infinitely many experiments: Part I—Theory," Appl. Math. Statist. Lab., Stanford Univ., Stanford, CA, USA, Tech. Rep., 1960, vol. 55.

[28] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for hypothesis testing," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2012, pp. 5277–5280.

[29] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Trans. Autom. Control*, vol. 58, no. 10, pp. 2451–2464, 2013.
[30] S. Nitinawarat and V. V. Veeravalli, "Controlled sensing for sequential multihypothesis testing with controlled Markovian observations and non-uniform control cost," 2013 [Online]. Available: http://arxiv.org/abs/1310.1844 [arXiv:1310.1844]
[31] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *Ann. Statist.*, vol. 41, no. 6, pp. 2703–2738, 2013.
[32] M. Naghshvar and T. Javidi, "Sequentiality and adaptivity gains in active hypothesis testing," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 5, pp. 768–782, 2013.
[33] K. Cohen and Q. Zhao, "Active hypothesis testing for anomaly detection," *IEEE Trans. Inf. Theory*, vol. 61, no. 3, pp. 1432–1450, 2015.
[34] K. Cohen and Q. Zhao, "Quickest anomaly detection: A case of active hypothesis testing," in *Proc. Inf. Theory Appl. (ITA) Workshop*, Feb. 2014, DOI: 10.1109/ITA.2014.6804268.
[35] S. Vakili, K. Liu, and Q. Zhao, "Deterministic sequencing of exploration and exploitation for multi-armed bandit problems," *IEEE J. Sel. Topics Signal Process. (JSTSP)*, vol. 7, pp. 759–767, Oct. 2013.
[36] S. M. Kay*, Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.
[37] I. Onat and A. Miri, "An intrusion detection system for wireless sensor networks," in *Proc. IEEE Int. Conf. Wireless Mobile Comput. Netw. Commun.*, 2005, vol. 3, pp. 253–259.
[38] W. E. Smith, "Various optimizers for single-stage production," *Naval Res. Logist. Quart.*, vol. 3, no. 1–2, pp. 59–66, 1956.
[39] A. G. Tartakovsky, "Asymptotic optimality of certain multihypothesis sequential tests: Non-iid case," *Statist. Inf. Stoch. Process.*, vol. 1, no. 3, pp. 265–295, 1998.
[40] M. L. Pinedo*, Scheduling: Theory, Algorithms, and Systems*. New York, NY, USA: Springer, 2012.

**Kobi Cohen** received the B.Sc. (*cum laude*) and Ph.D. degrees in electrical engineering from Bar-Ilan University, Ramat Gan, Israel, in 2007 and 2013, respectively.

He is currently a Postdoctoral Research Associate at the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, USA. From 2012 to 2014, he was with the Department of Electrical and Computer Engineering, University of California, Davis, USA, as a Postdoctoral Researcher. His main research interests include decision theory, stochastic optimization, game theory and statistical inference, with applications in large-scale systems, and wireless and wireline networks.

**Qing Zhao** (F'13) received the Ph.D. degree in electrical engineering in 2001 from Cornell University, Ithaca, NY, USA.

In August 2004, she joined the Department of Electrical and Computer Engineering, University of California, Davis (UC Davis), USA, where she is currently a Professor. She is also a Professor with the Graduate Group of Applied Mathematics at UC Davis. Her research interests are in the general area of stochastic optimization, decision theory, machine learning, and algorithmic theory in dynamic systems and communication and social-economic networks.

Dr. Zhao received the 2010 IEEE SIGNAL PROCESSING MAGAZINE Best Paper Award and the 2000 Young Author Best Paper Award from the IEEE Signal Processing Society. She holds the title of UC Davis Chancellor's Fellow and received the 2014 Outstanding Mid-Career Faculty Research Award and the 2008 Outstanding Junior Faculty Award from the UC Davis College of Engineering. She was a plenary speaker at the 11th IEEE Workshop on Signal Processing Advances in Wireless Communications (SPAWC), 2010. She is also a coauthor of two papers that received student paper awards at ICASSP 2006 and the IEEE Asilomar Conference 2006.