Quickest Search Over Multiple Sequences

Lifeng Lai, *Member, IEEE*, H. Vincent Poor, *Fellow, IEEE*, Yan Xin, *Senior Member, IEEE*, and Georgios Georgiadis

***Abstract*—The problem of sequentially finding an independent and identically distributed sequence that is drawn from a probability distribution Q_1 by searching over multiple sequences, some of which are drawn from Q_1 and the others of which are drawn from a different distribution Q_0, is considered. In the problem considered, the number of sequences with distribution Q_1 is assumed to be a random variable whose value is unknown. Within a Bayesian formulation, a sequential decision rule is derived that optimizes a trade-off between the probability of false alarm and the number of samples needed for the decision. In the case in which one can observe one sequence at a time, it is shown that the cumulative sum (CUSUM) test, which is well-known to be optimal for a non-Bayesian statistical change-point detection formulation, is optimal for the problem under study. Specifically, the CUSUM test is run on the first sequence. If a reset event occurs in the CUSUM test, then the sequence under examination is abandoned and the rule switches to the next sequence. If the CUSUM test stops, then the rule declares that the sequence under examination when the test stops is generated by Q_1. The result is derived by assuming that there are infinitely many sequences so that a sequence that has been examined once is not retested. If there are finitely many sequences, the result is also valid under a memorylessness condition. Expressions for the performance of the optimal sequential decision rule are also developed. The general case in which multiple sequences can be examined simultaneously is considered. The optimal solution for this general scenario is derived.**

***Index Terms*—Bayesian, CUSUM, optimal stopping, quickest search, sequential analysis.**

I. Introduction

IN THE classical sequential testing problem, one sequentially observes an independent and identically distributed (i.i.d.) sequence generated by one of two distributions Q_0 or Q_1, and wishes to test hypothesis H_1 that the sequence is generated by Q_1 against hypothesis H_0 that the sequence is generated by Q_0 [2]. The goal is to find a decision rule that uses a minimal number of samples, on average, while satisfying certain error probability constraints, or that optimizes some other trade-off between error probabilities and the average number of samples. Under this model, the sequential probability ratio test (SPRT) is well-known to be optimal [3]. This basic setting was extended to the situation in which there are three or more hypotheses in [4] and [5]. Motivated by sensor network applications, decentralized sequential hypothesis testing, in which each of a set of sensors receives a sequence of samples sequentially, has also been considered [6]–[12]. In this paper, we consider another generalization of the sequential testing problem: sequential search over multiple sequences. In particular, we consider N sequences, each of which is generated by either Q_0 or Q_1. For different value of i, whether the i^{th} sequence is generated by Q_0 or Q_1 is independent of all other sequences. Here, we assume that for the i^{th} sequence, hypothesis H_1 occurs with prior probability π_0 and H_0 with prior probability $1-\pi_0$. As a result, the number of sequences that are generated from Q_1 is a random variable, whose value can be any number between 0 and N and is unknown *a priori*.

Assuming that one can observe only one sequence at a time,[1] our goal is to find *one* sequence that is generated by Q_1 in a way that minimizes an appropriate measure of error probability and sampling cost. This model is motivated by many applications. For example, in the detection of chemical or biological attacks using a large sensor network with a mobile data collector, the mobile data collector needs to locate the point of attack quickly *after knowing that an attack has occurred*. Due to the limited transmission range of each wireless sensor, the collector can observe the signal from a only limited set of sensors at each time. In this case, we can model the distribution of the observations from the sensor affected by the attack as Q_1, and the distribution of the observations from the sensors unaffected by the attack as Q_0. Hence, finding a sequence generated by Q_1 quickly means finding the point of attack quickly. This formulation is a suitable model for searching for an affected sensor with minimal delay. As another example, in cognitive radio systems [13], wireless communication devices need to find unoccupied frequency bands before transmitting information. Hence, a wireless device should listen to each possible frequency band to determine whether it is free or not. In this scenario, the observations from one frequency band is a sequence of received signal samples, Q_0 corresponds to the distribution of the received signal when there are other transmissions in the band, and Q_1 corresponds to the distribution of the received signal when the frequency band is free. The task of finding a free frequency channel clearly can be modeled as that of finding a sequence generated by Q_1. It is of interest to do so with minimal delay, in order to make optimal use of spectral resources. However, the device can typically examine only one band at a time due to

Manuscript received February 16, 2010; revised December 23, 2010; accepted April 22, 2011. Date of current version July 29, 2011. The work of L. Lai and H. V. Poor was supported by the Qatar National Research Fund under Grant NPRP-08-522-2-211. The material in this paper was presented in part at the International Workshop on Applied Probability, Madrid, Spain, July 2010.

L. Lai is with the Department of Systems Engineering, University of Arkansas, Little Rock, AR 72204 USA (e-mail: lxlai@ualr.edu).

H. V. Poor is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: poor@princeton.edu).

Y. Xin is with the NEC Laboratories America, Inc., Princeton, NJ 08540 USA (e-mail: yanxin@nec-labs.com).

G. Georgiadis is with the Computer Science Department, University of California, Los Angeles, CA 90095 USA (e-mail: giorgos@cs.ucla.edu).

Communicated by F. Hlawatsch, Associate Editor for Detection and Estimation.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIT.2011.2159038

[1]The extension to the case in which one can observe multiple sequences simultaneously is considered in Section VI.

hardware limitations. Thus, this problem fits the above model very well. Another example is quality monitoring in a factory with multiple manufacturing machines. The task of finding a malfunctioning machine sequentially can be formulated as a sequential testing in multiple sequences problem. Finally, the problem of sequentially searching multiple databases for a certain type of data can also be modeled by the above framework.

The problem considered in this paper belongs to the class of sequential decision problems [14]–[19]. In particular, the problem considered here is related to a class of scanning problems considered in [20]–[24]. In the scanning problem, there are n channels. The observations of channel i are drawn from either distribution Q_1 or Q_0. Furthermore, it is assumed [20], [24] that one and only one of these n channels is generated from Q_1. The goal of the scanning problem is to find the channel generated by Q_1 with the minimal average delay subject to a constraint that the error probability is below a threshold. Under this model, the optimal solution is obtained for the Brownian motion case in [20] and [24]. In addition to the Brownian motion case, [24] also considered the general discrete time i.i.d. case. The main differences between our model and the above mentioned work on the scanning problem are: 1) in the scenario considered here, the number of channels that are generated from Q_1 is a random variable, which can take any value from 0 to N, and we do not know this value *a priori*; 2) our optimal solution is obtained for general distributions in discrete time, as will be clear in the sequel; and 3) in the work of [20] and [24], a finite number of sequences are considered, and switch back is allowed. Our model assumes that N is infinite. This assumption allows us to derive an optimal solution with a particularly simple form. It also serves as a good approximation for applications in which there are large numbers of sensors or channels. A problem with a similar flavor has also been considered in [25]–[27], which assumes that the samples in each sequence are generated from an on-off process and the goal is to quickly detect a sequence that makes a transition. Compared with these works, there is no transition within each sequence in our model, since we assume that samples from the same sequence are i.i.d. Under this model, our work establishes the optimality of the CUSUM for the infinite horizon case and also provides optimal solutions for the finite horizon case, which models the situation in which there is a strict deadline. We also provide performance analysis, which is critical for providing guidelines for determining the parameters involved in the algorithm. In addition, we obtain the solution for the more general case in which one can observe more than one channel simultaneously. Furthermore, our solution is based on the framework of optimal stopping theory, while [26] and [27] rely on the partially observable Markov decision process (POMDP) framework.

In this paper, we show that the solution to the quickest search problem is the cumulative sum (CUSUM) test, which was initially developed for the statistical change-point detection problem in [28]. In particular, we run the CUSUM test on the first sequence. If a reset event occurs in the CUSUM test, we abandon the sequence under examination and switch to the next sequence. If the CUSUM test stops, we claim that the sequence under examination is generated by Q_1. It is well-known that the CUSUM test is optimal [29] for the non-Bayesian quickest detection problem formulated in [30]. It is interesting to see the optimality of the CUSUM test for this different problem.

We show the optimality of the CUSUM test in two steps. We first solve the optimization problem (1) for bounded stopping times in Section III. Using insights from Section III, we then solve the general problem in Section IV. The performance of the optimal solution is analyzed in Section V. The generalization to the scenario in which one can observe multiple sequences simultaneously is considered in Section VI. Finally, we provide some concluding remarks in Section VIII.

II. Model

We consider N sequences $\{Y_k^i; k=1,2,\ldots\}$, $i=1,\ldots,N$, where for each i, $\{Y_k^i; k=1,2,\ldots\}$ are i.i.d. observations taking values in a set Ω endowed with a σ-field $\mathcal{F}$ of events, that obey one of the two hypotheses

$$H_0: \quad Y_k^i \sim Q_0, \quad k=1,2,\ldots$$

versus

$$H_1: \quad Y_k^i \sim Q_1, \quad k=1,2,\ldots$$

where Q_0 and Q_1 are two distinct, but equivalent, distributions on $(\Omega,\mathcal{F})$. We use q_0 and q_1 to denote densities of Q_0 and Q_1, respectively, with respect to some common dominating measure. The sequences for different values of i are independent. Moreover, whether the $i^{\rm th}$ sequence $\{Y_k^i; k=1,2,\ldots\}$ is generated by Q_0 or Q_1 is independent of all other sequences. Here, we assume that for each i, hypothesis H_1 occurs with prior probability π_0 and H_0 with prior probability $1-\pi_0$. We use P to denote the probability measure defined as above. In addition, we will also use the probability measures P_0 and P_1 such that, under P_j, all the observations are i.i.d. with marginal distribution Q_j, for $j=0,1$. Furthermore, we will also use the probability measure $P_{\pi_0} = \pi_0 P_1 + (1-\pi_0)P_0$. As a result, the number of sequences that are generated from Q_1 is a random variable, whose value can be any number between 0 and N and is unknown *a priori*.

At each time, we select a sequence, say sequence j, and make an observation from this sequence. After making each observation, we can take one of the following three actions: 1) stop sampling and claim that the sequence we are currently observing is generated by Q_1; 2) continue to the next observation from the same sequence to gather more evidence about its statistical behavior; or 3) abandon the sequence that we are currently observing and switch to another sequence. Hence if a sequence is abandoned, we will not come back and test it again. Without loss of generality, we start taking samples from the first sequence, and switch to the second sequence if we decide to abandon the first sequence. Similarly, we will switch to the $(i+1)^{\rm th}$ sequence if we decide to abandon the $i^{\rm th}$ sequence. To ensure that there is always a sequence to switch to, we consider the case $N=\infty$. When N is finite, our model is also applicable to the scenario in which when we switch back to a sequence that has been examined previously, it is treated as a new sequence with no memory of the observations that have been taken before. This assumption is valid for the case in which the controller has limited memory.

We use s_k to denote the index of the sequence that we are observing at time k. Hence, we observe $\{Y_k^{s_k}; k = 1, 2, \ldots\}$ sequentially. The observations generate the filtration $\{\mathcal{F}_k; k = 1, 2, \ldots\}$ with

$$\mathcal{F}_k = \sigma(Y_1^{s_1}, Y_2^{s_2}, \ldots, Y_k^{s_k}).$$

We use ϕ_k to denote the $\mathcal{F}_k$-measurable switching function at time k. Here, $\phi_k(\mathcal{F}_k) = 1$ if we decide to abandon sequence s_k and switch to the next sequence, that is, $s_{k+1} = 1 + s_k$. On the other hand $\phi_k(\mathcal{F}_k) = 0$ if we decide to continue observing sequence s_k, that is, $s_{k+1} = s_k$. Let $\mathcal{T}$ denote the set of all stopping times with respect to the filtration $\mathcal{F}_k$. Note that the sequence $s_1, s_2, \ldots$, and hence the filtration $\mathcal{F}_1, \mathcal{F}_2, \ldots$, depends on the sequence $\phi_1, \phi_2, \ldots$ of switching functions. A stopping time $\tau \in \mathcal{T}$ will decide when we should stop sampling and declare that the sequence we are currently observing is generated by Q_1. More specifically, if $\tau = k$, we should stop sampling at time k, and declare that sequence s_k is generated by Q_1. There are two performance indices: 1) the error probability that sequence s_τ is generated by Q_0, that is, $P(H^{s_\tau} = H_0)$, where H^j is the true hypothesis satisfied by sequence j; and 2) the average number of samples we take to make a decision, that is, $\mathbb{E}\{\tau\}$.

Our goal is to determine the stopping time τ and the switching rules $\phi = \{\phi_1, \phi_2, \ldots\}$ to solve the following optimization problem:

$$\inf_{\tau \in \mathcal{T}, \phi} \left[P(H^{s_\tau} = H_0) + c\mathbb{E}\{\tau\} \right]. \tag{1}$$

Here $c > 0$ is a constant that represents the cost of taking one sample. We assume $c < 1 - \pi_0$, as the case $c \geq 1 - \pi_0$ is trivial: we simply do not take any observations and choose a sequence at random as being generated by Q_1.

We note that other than the Bayesian formulation adopted in (1), one could also use a variational formulation to strike a balance between the error probability $P(H^{s_\tau} = H_0)$ and the average delay $\mathbb{E}\{\tau\}$. More specifically, in the variational formulation, one aims to solve the following optimization problem:

$$\begin{aligned} &\inf_{\tau \in \mathcal{T}, \phi} \mathbb{E}\{\tau\}, \\ &\text{s.t.} \quad P(H^{s_\tau} = H_0) \leq \alpha. \end{aligned} \tag{2}$$

That is, we want to minimize the average delay under the constraint that the error probability is less than a preset threshold α. However, following the same line of argument in [31, Sec. 4.3], one can obtain the solution to (2) once the solution of the Bayesian formulation is found. More specifically, for each α, there exists a constant $c(\alpha)$ such that the solution to (1) with this constant is the solution to (2).

Another problem formulation in sequential testing is a non-Bayesian formalism in which one does not assume the prior probability π_0. In the single sequence testing, we assume that the sequence is fixed to be either Q_0 or Q_1. We need to design a scheme that works well for all these two scenarios. In the current setup, the number of possible scenarios is 2^N, which grows without bound as N increases. Hence, developing a meaningful formulation for the non-Bayesian case is challenging.

III. Finite-Horizon Optimization

Before we solve the optimization problem (1), we study a finite-horizon version of it in which the stopping time τ is restricted to a finite interval $[0, T]$; that is, we must stop by time T. This finite-horizon problem has practical significance when there is a strict delay deadline.

We use $\pi_k = P(H^{s_k} = H_1 | \mathcal{F}_k)$ to denote the posterior probability that sequence s_k is generated by Q_1 after observing $\{Y_1^{s_1}, \ldots, Y_k^{s_k}\}$. From the independence assumptions mentioned in Section I and the fact that ϕ_k is $\mathcal{F}_k$ measurable for $k = 1, 2, \ldots$, we have the following recursive formula for π_k:

$$\begin{aligned} \pi_1 &= \frac{\pi_0 q_1(Y_1^1)}{\pi_0 q_1(Y_1^1) + (1-\pi_0) q_0(Y_1^1)} \\ &\vdots \\ \pi_{k+1} &= \frac{\pi_k q_1(Y_{k+1}^{s_{k+1}})}{\pi_k q_1(Y_{k+1}^{s_{k+1}}) + (1-\pi_k) q_0(Y_{k+1}^{s_{k+1}})} \mathbf{1}(\phi_k = 0) \\ &\quad + \frac{\pi_0 q_1(Y_{k+1}^{s_{k+1}})}{\pi_0 q_1(Y_{k+1}^{s_{k+1}}) + (1-\pi_0) q_0(Y_{k+1}^{s_{k+1}})} \mathbf{1}(\phi_k = 1), \end{aligned} \tag{3}$$

in which $\mathbf{1}(\cdot)$ is the indicator function. Note that π_k is not necessarily a sufficient statistic for this problem, unless the $\mathcal{F}_k$-measurable function ϕ_k depends only on π_k. We will show, however, that this is indeed the case.

At each time k, we need to decide whether to stop sampling or not based on $\mathcal{F}_k$. The minimal expected cost-to-go at time k, $0 \leq k \leq T$, is a function of $\mathcal{F}_k$, which we will denote by $\tilde{J}_k^T(\mathcal{F}_k)$. Obviously, we have

$$\tilde{J}_T^T(\mathcal{F}_T) = 1 - \pi_T.$$

Given $\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})$, we have the first equation at the bottom of the next page. The interpretation of each term in the equation is clear. Specifically, $1 - \pi_k$ is the cost incurred if we stop sampling at time k, and

$$c + \inf_{\phi_k} \mathbb{E}\left\{ \tilde{J}_{k+1}^T(\mathcal{F}_{k+1}) | \mathcal{F}_k, \phi_k \right\},$$

is the expected cost that we will incur if we continue sampling, which is the smaller of two costs: the expected cost that we will incur if we continue sampling in the same sequence, and the expected cost that we will incur if we switch to another sequence.

We first have the following lemma that converts this problem into a Markov optimal stopping problem.

Lemma 1: For each k, the function $\tilde{J}_k^T(\mathcal{F}_k)$ can be written as a function of only π_k, say $J_k^T(\pi_k)$, and the optimal switching rules ϕ_k can be restricted to a class of decision functions that depend only on π_k.

Proof: Clearly $\tilde{J}_T^T(\mathcal{F}_T) = 1 - \pi_T$ is a function of only π_T, and we write it as $J_T^T(\pi_T)$. For any k, $0 \le k \le T-1$, suppose that $\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})$ is a function of only π_{k+1} and we use $J_{k+1}^T(\pi_{k+1})$ to denote this function, then we have the second equation at the bottom of the page, in which $f_c(Y_{k+1}^{s_{k+1}}|\mathcal{F}_k)$ is the conditional density of $Y_{k+1}^{s_{k+1}}$ if we decide to stay in the same sequence to make more observations, that is

$$f_c(Y_{k+1}^{s_{k+1}}|\mathcal{F}_k) = \pi_k q_1(Y_{k+1}^{s_{k+1}}) + (1-\pi_k) q_0(Y_{k+1}^{s_{k+1}}). \quad (4)$$

Similarly, $f_s(Y_{k+1}^{s_{k+1}}|\mathcal{F}_k)$ is the conditional density of $Y_{k+1}^{s_{k+1}}$ if we decide to switch to another sequence to make observations, that is,

$$f_s(Y_{k+1}^{s_{k+1}}|\mathcal{F}_k) = \pi_0 q_1(Y_{k+1}^{s_{k+1}}) + (1-\pi_0) q_0(Y_{k+1}^{s_{k+1}}). \quad (5)$$

The relationship between π_{k+1} and $y_{k+1}^{s_{k+1}}$ is given in (3). Hence, f_c and f_s depend on $\mathcal{F}_k$ only through π_k. In fact, f_s is independent of $\mathcal{F}_k$. As a result, we can write $f_c(Y_{k+1}^{s_{k+1}}|\mathcal{F}_k)$ as $f_c(Y_{k+1}^{s_{k+1}}|\pi_k)$ and write $f_s(Y_{k+1}^{s_{k+1}}|\mathcal{F}_k)$ as $f_s(Y_{k+1}^{s_{k+1}}|\pi_k)$.

Hence, we have

$$\begin{aligned}&\int J_{k+1}^T(\pi_{k+1}) f_c(y_{k+1}^{s_{k+1}}|\mathcal{F}_k) dy_{k+1}^{s_{k+1}} \\ &= \int J_{k+1}^T\left(\frac{\pi_k q_1(y_{k+1}^{s_{k+1}})}{\pi_k q_1(y_{k+1}^{s_{k+1}}) + (1-\pi_k) q_0(y_{k+1}^{s_{k+1}})}\right) \\ &\qquad (\pi_k q_1(y_{k+1}^{s_{k+1}}) + (1-\pi_k) q_0(y_{k+1}^{s_{k+1}}))\, dy_{k+1}^{s_{k+1}},\end{aligned}$$

which is a function of π_k, and we will use $A_{k,c}^T(\pi_k)$ to denote this function.

At the same time, we have

$$\begin{aligned}&\int J_{k+1}^T(\pi_{k+1}) f_s(y_{k+1}^{s_{k+1}}|\mathcal{F}_k) dy_{k+1}^{s_{k+1}} \\ &= \int J_{k+1}^T\left(\frac{\pi_0 q_1(y_{k+1}^{s_{k+1}})}{\pi_0 q_1(y_{k+1}^{s_{k+1}}) + (1-\pi_0) q_0(y_{k+1}^{s_{k+1}})}\right) \\ &\qquad (\pi_0 q_1(y_{k+1}^{s_{k+1}}) + (1-\pi_0) q_0(y_{k+1}^{s_{k+1}}))\, dy_{k+1}^{s_{k+1}},\end{aligned}$$

which is a constant independent of π_k, and we will use $A_{k,s}^T$ to denote this constant.

Thus, $\tilde{J}_k^T(\mathcal{F}_k) = \min\{1-\pi_k, c + \min\{A_{k,c}^T(\pi_k), A_{k,s}^T\}\}$ is a function of π_k, and we write it as $J_k^T(\pi_k)$. Continuing this recursive argument, we know that $\tilde{J}_k^T(\mathcal{F}_k)$ depends only on π_k for $k = 1, \ldots, T$.

Furthermore, since ϕ_k has only two values, the optimal switching rule ϕ_k is the following:

$$\phi_k = \begin{cases} 1, & \text{if } A_{k,c}^T(\pi_k) > A_{k,s}^T, \\ 0, & \text{otherwise,}\end{cases} \quad (6)$$

which depends only on π_k. This means that we should switch to another sequence, if the expected cost of continuing this sequence is larger than the expected cost that will be incurred if we switch to another sequence, although by doing this we lose all the evidence we have gathered to this point. Since ϕ_k depends only on π_k, from (3), we see that $\{\pi_k; k = 0, 1, \ldots\}$ forms a Markov process. ■

In summary, we have converted the finite-length version of problem (1) to a Markov optimal stopping problem. For finite T, we have the following recursive cost functions:

$$J_T^T(\pi_T) = 1 - \pi_T, \quad (7)$$

and for $k = 0, \ldots, T-1$,

$$J_k^T(\pi_k) = \min\left\{1 - \pi_k, c + \min\left\{A_{k,c}^T(\pi_k), A_{k,s}^T\right\}\right\}. \quad (8)$$

Regarding $J_k^T(\pi_k)$ and $A_{k,c}^T(\pi_k)$, we have the following result.

Lemma 2: The functions $J_k^T(\pi_k)$ and $A_{k,c}^T(\pi_k)$ are nonnegative concave functions of π_k, for $\pi_k \in [0,1]$. And $J_k^T(1) = A_{k,c}^T(1) = 0$.

Proof: The nonnegative property of these functions can be easily proved by simple inductive arguments. The fact that

$$\begin{aligned}\tilde{J}_k^T(\mathcal{F}_k) &= \min\left\{1-\pi_k, c + \inf_{\phi_k} \mathbb{E}\left\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k\right\}\right\} \\ &= \min\Big\{1-\pi_k, c + \min\Big\{\mathbb{E}\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k = 0\}, \\ &\qquad \mathbb{E}\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k = 1\}\Big\}\Big\}.\end{aligned}$$

$$\begin{aligned}\tilde{J}_k^T(\mathcal{F}_k) &= \min\Big\{1-\pi_k, c + \min\Big\{\mathbb{E}\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k = 0\}, \\ &\qquad \mathbb{E}\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k = 1\}\Big\}\Big\} \\ &= \min\Big\{1-\pi_k, c + \min\Big\{\int J_{k+1}^T(\pi_{k+1}) f_c(y_{k+1}^{s_{k+1}}|\mathcal{F}_k)\, dy_{k+1}^{s_{k+1}}, \\ &\qquad \int J_{k+1}^T(\pi_{k+1}) f_s(y_{k+1}^{s_{k+1}}|\mathcal{F}_k)\, dy_{k+1}^{s_{k+1}}\Big\}\Big\},\end{aligned}$$

$J_k^T(1) = A_{k,c}^T(1) = 0$ also follows from a simple inductive argument.

In the following we show the concavity of these two functions. First, we observe that $J_T^T(\pi_T) = 1 - \pi_T$ is a concave function of π_T. Now, assuming $J_{k+1}^T(\pi_{k+1})$ is concave in π_{k+1}, we prove that $A_{k,c}^T(\pi_k)$ is concave in π_k. Let π_k^1 and π_k^2 be two arbitrary points belonging to $[0,1]$. Consider $\lambda A_{k,c}^T(\pi_k^1) + (1-\lambda)A_{k,c}^T(\pi_k^2)$, with $0 \le \lambda \le 1$, we have

$$\begin{aligned}
&\lambda A_{k,c}^T(\pi_k^1) + (1-\lambda)A_{k,c}^T(\pi_k^2)\\
&\quad= \int \Big[\lambda J_{k+1}^T(\pi_{k+1}^1) f_c(y_{k+1}^{s_{k+1}}|\pi_k^1)\\
&\qquad +(1-\lambda)J_{k+1}^T(\pi_{k+1}^2) f_c(y_{k+1}^{s_{k+1}}|\pi_k^2)\Big]\, dy_{k+1}^{s_{k+1}}\\
&\quad= \int \left[\mu J_{k+1}^T(\pi_{k+1}^1) + (1-\mu)J_{k+1}^T(\pi_{k+1}^2)\right]\\
&\qquad (\lambda f_c(y_{k+1}^{s_{k+1}}|\pi_k^1) + (1-\lambda) f_c(y_{k+1}^{s_{k+1}}|\pi_k^2))\, dy_{k+1}^{s_{k+1}}\\
&\quad\le \int J_{k+1}^T(\mu\pi_{k+1}^1 + (1-\mu)\pi_{k+1}^2)\\
&\qquad (\lambda f_c(y_{k+1}^{s_{k+1}}|\pi_k^1) + (1-\lambda) f_c(y_{k+1}^{s_{k+1}}|\pi_k^2))\, dy_{k+1}^{s_{k+1}}, \qquad (9)
\end{aligned}$$

in which

$$\mu = \frac{\lambda f_c(y_{k+1}^{s_{k+1}}|\pi_k^1)}{\lambda f_c(y_{k+1}^{s_{k+1}}|\pi_k^1) + (1-\lambda) f_c(y_{k+1}^{s_{k+1}}|\pi_k^2)},$$

and where we have used the concavity of J_{k+1}^T in writing the inequality.

Now, on defining $\pi_k^3 = \lambda\pi_k^1 + (1-\lambda)\pi_k^2$, we have (10) at the bottom of the page, in which we have used (4) for (a).

At the same time, we have

$$f_c(Y_{k+1}^{s_{k+1}}|\pi_k^3) = \lambda f_c(Y_{k+1}^{s_{k+1}}|\pi_k^1) + (1-\lambda) f_c(Y_{k+1}^{s_{k+1}}|\pi_k^2). \quad (11)$$

Hence, continuing from (9), we have

$$\lambda A_{k,c}^T(\pi_k^1) + (1-\lambda)A_{k,c}^T(\pi_k^2) \le A_{k,c}^T(\pi_k^3), \qquad (12)$$

which means that $A_{k,c}^T(\pi_k)$ is a concave function of π_k. The concavity of $J_k^T(\pi_k)$ then follows from the fact that the minimum of concave functions is also concave. ■

It is also clear that $J_k^T(\pi_k) \le 1 - \pi_k$. Fig. 1 shows an illustration of the relationships of $J_k^T(\pi_k)$ and $1 - \pi_k$. With these supporting lemmas, we have the following solution for the finite-horizon optimization problem.

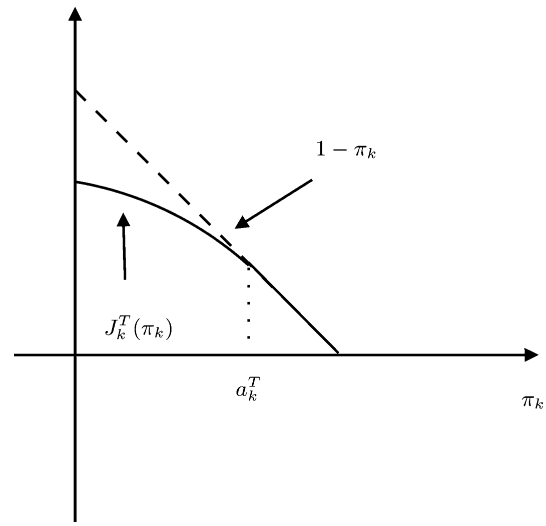

Fig. 1. An illustration of $J_k^T(\pi_k)$.

Theorem 3: For the finite-horizon version of problem (1), the optimal stopping time is $\tau_{\text{opt}} = \inf\{k : \pi_k > a_k^T\}$, in which a_k^T is given by the following equation:

$$1 - a_k^T = c + \min\{A_{k,c}^T(a_k^T), A_{k,s}^T\}.$$

And at time k, we switch to another sequence if, and only if, $A_{k,c}^T(\pi_k) > A_{k,s}^T$.

IV. Infinite-Horizon Optimization

Now, we remove the finiteness restriction on the stopping time τ by letting $T \longrightarrow \infty$. First, we have

$$J_k^{T+1}(\pi) \le J_k^T(\pi), \quad 0 \le \pi \le 1, \qquad (13)$$

since the set of allowed stopping time is enlarged if we allow the horizon T to increase. Further we have $0 \le J_k^T \le 1$ for any k and T, and hence the following limit is well-defined:

$$\lim_{T\longrightarrow\infty} J_k^T(\pi) = \inf_{T>k} J_k^T(\pi) = J_k^\infty(\pi). \qquad (14)$$

Also, we have $J_k^\infty(\pi) = J_{k+1}^\infty(\pi)$, due to the i.i.d. nature of the observations in each sequence. We will use $J(\pi)$ to denote this common function. It is easy to check that $J(\pi)$ is a concave function in π.

$$\begin{aligned}
\pi_{k+1}^3 &= \frac{\pi_k^3 q_1(Y_{k+1}^{s_{k+1}})}{\pi_k^3 q_1(Y_{k+1}^{s_{k+1}}) + (1-\pi_k^3) q_0(Y_{k+1}^{s_{k+1}})}\\
&= \frac{(\lambda\pi_k^1 + (1-\lambda)\pi_k^2) q_1(Y_{k+1}^{s_{k+1}})}{(\lambda\pi_k^1 + (1-\lambda)\pi_k^2) q_1(Y_{k+1}^{s_{k+1}}) + (1 - \lambda\pi_k^1 - (1-\lambda)\pi_k^2) q_0(Y_{k+1}^{s_{k+1}})}\\
&\overset{(a)}{=} \frac{(\lambda\pi_k^1 + (1-\lambda)\pi_k^2) q_1(Y_{k+1}^{s_{k+1}})}{\lambda f_c(Y_{k+1}^{s_{k+1}}|\pi_k^1) + (1-\lambda) f_c(Y_{k+1}^{s_{k+1}}|\pi_k^2)}\\
&= \mu\pi_{k+1}^1 + (1-\mu)\pi_{k+1}^2, \qquad (10)
\end{aligned}$$

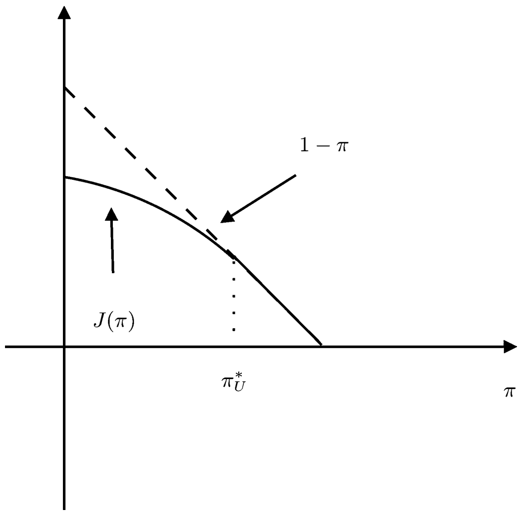

Fig. 2. An illustration of $J(\pi)$.

Lemma 4: The function $J(\pi)$ is unique.

Proof: The proof of this result follows the argument used in [17, Proposition 7.4]. ∎

By the dominated convergence theorem, the limit in (15) at the bottom of the page is well defined, which is a constant independent of k.

Similarly, we have

$$\begin{aligned} A_c(\pi) &:= \lim_{T\longrightarrow\infty} A_{k,c}^T(\pi) \qquad (16)\\ &= \int J\left(\frac{\pi q_1(y_{k+1}^{s_{k+1}})}{\pi q_1(y_{k+1}^{s_{k+1}})+(1-\pi)q_0(y_{k+1}^{s_{k+1}})}\right) \\ &\qquad (\pi q_1(y_{k+1}^{s_{k+1}})+(1-\pi)q_0(y_{k+1}^{s_{k+1}}))\, dy_{k+1}^{s_{k+1}}. \end{aligned}$$

From the fact that $A_{k,c}^T$ is concave for each k and T, it is easy to check that $A_c(\pi)$ is concave in π.

Hence,

$$J(\pi) = \min\{1-\pi, c+\min\{A_s, A_c(\pi)\}\}. \qquad (17)$$

Fig. 2 illustrates the relationship between $J(\pi)$ and π.

Now, we have the following lemma regarding the relationship between A_s and $A_c(\pi)$.

Lemma 5:

$$A_c(\pi)\begin{cases} > A_s, & \text{if } \pi < \pi_0,\\ = A_s, & \text{if } \pi = \pi_0,\\ < A_s, & \text{if } \pi > \pi_0. \end{cases} \qquad (18)$$

Proof: We first show that $A_c(0) > A_s$. From (17), we have

$$J(0) = \min\{1, c+\min\{A_s, A_c(0)\}\}. \qquad (19)$$

From (16), we have

$$A_c(0) = \int J(0) q_0(y_{k+1}^{s_{k+1}})\, dy_{k+1}^{s_{k+1}} = J(0). \qquad (20)$$

Hence (19) becomes

$$J(0) = \min\{1, c+\min\{A_s, J(0)\}\}. \qquad (21)$$

As a result, $J(0)$ is either 1 or $c+A_s$.

We consider these two cases separately.

1) If $J(0) = c+A_s$, then from (20), we have that $A_c(0) = J(0) = c+A_s > A_s$.
2) If $J(0) = 1$, then from the facts that $J(1) = 0$, $J(\pi) \leq 1-\pi$, and $J(\pi)$ is a concave function of π, we know that $J(\pi) = 1-\pi$. Substituting this function into (15), we have

$$\begin{aligned} A_s &= \int J\left(\frac{\pi_0 q_1(y_{k+1}^{s_{k+1}})}{\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}})}\right) \\ &\qquad (\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}}))\, dy_{k+1}^{s_{k+1}} \\ &= \int \left(1-\frac{\pi_0 q_1(y_{k+1}^{s_{k+1}})}{\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}})}\right) \\ &\qquad (\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}}))\, dy_{k+1}^{s_{k+1}} \\ &= \int (1-\pi_0) q_0(y_{k+1}^{s_{k+1}})\, dy_{k+1}^{s_{k+1}} \\ &= 1-\pi_0. \end{aligned} \qquad (22)$$

Hence, in this case, we still have $A_c(0) = J(0) = 1 > 1-\pi_0 = A_s$.

$$\begin{aligned} \lim_{T\longrightarrow\infty} A_{k,s}^T &= \lim_{T\longrightarrow\infty} \int J_{k+1}^T\left(\frac{\pi_0 q_1(y_{k+1}^{s_{k+1}})}{\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}})}\right) \\ &\qquad (\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}}))\, dy_{k+1}^{s_{k+1}} \\ &= \int \lim_{T\longrightarrow\infty} J_{k+1}^T\left(\frac{\pi_0 q_1(y_{k+1}^{s_{k+1}})}{\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}})}\right) \\ &\qquad (\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}}))\, dy_{k+1}^{s_{k+1}} \\ &= \int J\left(\frac{\pi_0 q_1(y_{k+1}^{s_{k+1}})}{\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}})}\right) \\ &\qquad (\pi_0 q_1(y_{k+1}^{s_{k+1}})+(1-\pi_0)q_0(y_{k+1}^{s_{k+1}}))\, dy_{k+1}^{s_{k+1}} \\ &:= A_s, \end{aligned} \qquad (15)$$

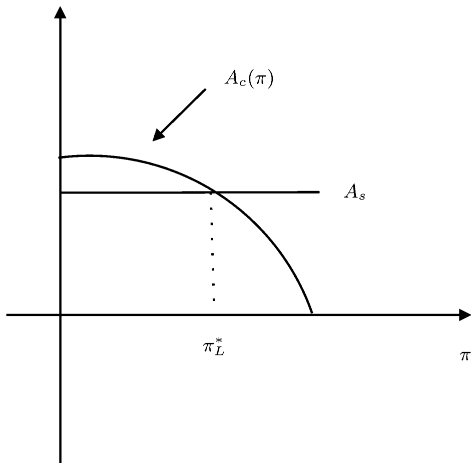

Fig. 3. An illustration of $A_c(\pi)$ and A_s.

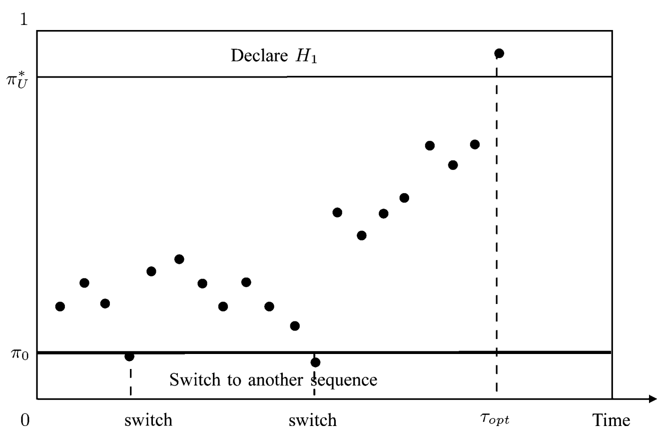

Fig. 4. An illustration of the optimal procedure.

Fig. 3 illustrates the relationship between $A_c(\pi)$ and A_s.

Since $A_c(\pi)$ is concave, $A_c(0) > A_s$ and $0 = A_c(1) \leq A_s$, we know that there exists a unique π_L^* such that $0 < \pi_L^* \leq 1$ and $A_c(\pi_L^*) = A_s$. At the same time, from (15) and (16), we know that $A_c(\pi_0) = A_s$. Hence the unique point, in which $A_c(\pi)$ and A_s are equal, is π_0. That is $\pi_L^* = \pi_0$. As a result we know that $A_c(\pi) > A_s$ when $\pi < \pi_0$, and $A_c(\pi) < A_s$ when $\pi > \pi_0$. ∎

Lemma 5 implies that we should switch to another sequence once π_k is less than the prior probability π_0.

By examining Figs. 2 and 3, and applying [17, Th. 3.7] and Lemma 5 above, we have the following solution for the infinite-horizon optimization problem.

Theorem 6: The optimal stopping time for (1) is given by $\tau_{\text{opt}} = \inf\{k : \pi_k > \pi_U^*\}$ in which

$$1 - \pi_U^* = c + \min\{A_c(\pi_U^*), A_s\}.$$

And at time k, we switch to another sequence if, and only if, $\pi_k < \pi_0$.

Fig. 4 illustrates the operation of the optimal test.

Now we briefly review the CUSUM test, which was developed for the quickest detection problem [17], before discussing the connection between this test and our algorithm. In the quickest detection problem, one observes a random sequence $Y_k; k = 1, 2, \ldots$. There is a change point $t \geq 1$ such that, given t, $Y_1, \ldots, Y_{t-1}$ are drawn i.i.d. from distribution Q_0, and $Y_t, Y_{t+1}, \ldots$ are drawn i.i.d. from distribution Q_1. In the non-Bayesian formulation, the change point t is assumed to be a fixed, nonrandom quantity that can be either ∞ or any value in the positive integers. One aims to detect the occurrence of this change with minimal delay subject to a certain false alarm constraint. In a formulation initiated by Lorden [30], the following optimization problem is considered:

$$\inf_{T \in \mathcal{T}} \ \operatorname{ess\,sup} \mathbb{E}_t \left\{ (T - t + 1)^+ | \mathcal{F}_{t-1} \right\} \tag{23}$$

$$\text{s.t.} \quad \mathbb{E}_\infty\{T\} \geq \gamma, \tag{24}$$

in which $\mathcal{F}_t$ denotes the smallest σ-field with respect to which $Y_1, \ldots, Y_t$ are measurable, $\mathcal{T}$ is the set of all stopping times with respect to the filtration $\mathcal{F}_t; t \geq 0$, and $\mathbb{E}_t$ denotes expectation assuming that the change time is t.

Let $L_k = q_1(Y_k)/q_0(Y_k)$, and $m_{k+1} = \max\{m_k, 0\} + \log(L_{k+1})$ with $m_0 = 0$; then the CUSUM stopping time is $T_h = \inf\{k \geq 0 | m_k \geq h\}$ where h is a threshold. That is one stops whenever the statistic m_k is larger than a given threshold, and resets the statistic to 0 once it is smaller than 0. It was shown in [29] that the CUSUM test with h chosen to satisfy (24) with equality is the optimal solution to the problem (24) for all values of $\gamma > 0$.

It is now easy to see the equivalence between the optimal test in Theorem 6 and the CUSUM test. More specifically, let

$$L_k = \frac{q_1(Y_k^{s_k})}{q_0(Y_k^{s_k})};$$

then under the condition that $\phi_k = 1$ if $\pi_k < \pi_0$ and $\phi_k = 0$ if $\pi_k \geq \pi_0$, the recursive formula in (3) is the equivalent to the following recursive formula:

$$R_1 = \log(L_1) \tag{25}$$

$$\vdots$$

$$\begin{aligned} R_{k+1} &= (R_k + \log(L_{k+1}))\,\mathbf{1}(R_k \geq 0) + \log(L_{k+1})\mathbf{1}(R_k < 0) \\ &= \max\{R_k, 0\} + \log(L_{k+1}). \end{aligned}$$

In terms of R_k, the optimal solution is to switch to the next sequence if $R_k < 0$ (this corresponds to a reset event in the CUSUM test, which is to reset R_k to zero, if $R_k < 0$), and to stop when

$$R_k \geq \log\left(\frac{(1-\pi_0)\pi_U^*}{\pi_0(1-\pi_U^*)}\right).$$

Hence the test in Theorem 6 is equivalent to a CUSUM test with parameter $(1-\pi_0)\pi_U^*/(\pi_0(1-\pi_U^*))$, in which we switch to another sequence if a reset event occurs in the CUSUM test, and we stop and claim that the sequence under examination is generated by Q_1 when the CUSUM test stops.

V. Performance Analysis

From Section IV, it is clear that the optimal solution can be parameterized by one threshold π_U^*, whose value depends on the cost per sample c. In this section, we analyze the average delay $\mathbb{E}\{\tau_{\text{opt}}\}$ and the error probability $P(H^{s_{\tau_{\text{opt}}}} = H_0)$ in terms of

π_U^*. The analysis will provide further insight into the structure of the optimal solution and give guidance on how to set the parameter. Since the optimal test is the same as the CUSUM test, we can use techniques similar to those used in the performance analysis of the CUSUM test (see [32] and [33] and references therein) with proper modifications to take the Bayesian framework into consideration.

It is clear that $\tau_{\rm opt}$ arises from a renewal process [34], with renewals occurring whenever π_k is reset to π_0 (this occurs when π_k exits from the lower end of $[\pi_0, \pi_U^*]$), and with a termination when π_k exits from upper end of $[\pi_0, \pi_U^*]$. It follows that we can write

$$\tau_{\rm opt} = \sum_{l=1}^{L} t_l, \tag{26}$$

where $t_1, t_2, \dots$ are i.i.d. repetitions (under $P_{\pi_0} = \pi_0 P_1 + (1-\pi_0)P_0$) of the random variable

$$t = \min\{k \geq 1 | \pi_k \notin (\pi_0, \pi_U^*)\}, \tag{27}$$

and where L denotes the number of repetitions of t that occur before the posterior probability results in an exit at the upper boundary. Hence, the analysis of $\mathbb{E}\{\tau_{\rm opt}\}$ and $P(H^{s_{\tau_{\rm opt}}} = H_0)$ can be carried out by analyzing this renewal process under P_{π_0}.

We define F_0 and F_1 to be the events $\{\pi_t < \pi_0\}$ and $\{\pi_t \geq \pi_U^*\}$, respectively. We also define $\alpha = P_0(F_1)$ and $\beta = P_1(F_0)$. It is clear that α is the probability that we will claim that the sequence is generated by Q_1, when hypothesis H_0 is true. Similarly, β is the probability that we will make a switch while hypothesis H_1 is true. We have the following theorem regarding the performance of the optimal solution of Theorem 6.

Theorem 7: The average number of samples of the optimal solution for the sequential testing with multiple sequences is characterized by

$$\mathbb{E}_{\pi_0}\{\tau_{\rm opt}\} = \frac{\mathbb{E}_{\pi_0}\{t\}}{1 - P_{\pi_0}(F_0)} = \frac{\mathbb{E}_{\pi_0}\{t\}}{1 - \pi_0\beta - (1-\pi_0)(1-\alpha)}, \tag{28}$$

and the error probability is characterized by

$$P_{\pi_0}(H^{\tau_{\rm opt}} = H_0) = \frac{(1-\pi_0)\alpha}{(1-\pi_0)\alpha + \pi_0(1-\beta)}. \tag{29}$$

Proof: Let M_l denote the indicator of the event that the $l^{\rm th}$ repetition of t exits at the upper boundary. Then L is a stopping time with respect to the sequence $(t_1, M_1), (t_2, M_2), \dots$, which is i.i.d. under P_{π_0}. From Wald's identity, we have

$$\mathbb{E}_{\pi_0}\left\{\sum_{l=1}^{L} t_l\right\} = \mathbb{E}_{\pi_0}\{L\}\mathbb{E}_{\pi_0}\{t\}. \tag{30}$$

It is easy to see that, under $P_{\pi_0} = \pi_0 P_1 + (1-\pi_0)P_0$, L is a geometric random variable with

$$P_{\pi_0}(L = l) = [1 - P_{\pi_0}(F_0)][P_{\pi_0}(F_0)]^{l-1}, \quad l = 1, 2, \dots, \tag{31}$$

where $P_{\pi_0}(F_0) = \pi_0\beta + (1-\pi_0)(1-\alpha)$. Hence

$$\mathbb{E}_{\pi_0}\{L\} = \frac{1}{1 - P_{\pi_0}(F_0)}.$$

Substituting this into (30), we have (28).

At the same time, we have

$$\begin{aligned} P_{\pi_0}(H^{\tau_{\rm opt}} = H_0) &= \sum_{L=l} P_{\pi_0}(H^l = H_0 | L = l) P_{\pi_0}(L = l) \\ &= P_{\pi_0}(H^1 = H_0 | L = 1) \\ &= \frac{(1-\pi_0)\alpha}{(1-\pi_0)\alpha + \pi_0(1-\beta)}. \end{aligned} \tag{32}$$

■

The analysis of $\mathbb{E}_{\pi_0}\{t\}$, α and β in terms of the parameters (π_0, π_U^*) follows from the standard SPRT analysis, which can be found, for example, in [33].

VI. Multiple Simultaneous Observations

In this section, we consider the general case in which we can observe multiple sequences simultaneously. Our goal is still to find one sequence that is generated by Q_1. We use M to denote the number of sequences that we can observe at each time. We will discuss the case in which $M = 2$ in detail. (The case in which $M > 2$ is similar.) The development of the optimal solution follows that of the single sequence case closely.

We use s_k^a and s_k^b with $s_k^a \neq s_k^b$ to denote the indices of the two sequences that we are observing at time k. Hence, we observe $\{Y_k^{s_k^a}, Y_k^{s_k^b}; k = 1, 2, \dots\}$ sequentially. The observations generate the filtration $\{\mathcal{F}_k; k = 1, 2, \dots\}$ with $\mathcal{F}_k = \sigma(Y_1^{s_1^a}, Y_1^{s_1^b}, Y_2^{s_2^a}, Y_2^{s_2^b}, \dots, Y_k^{s_k^a}, Y_k^{s_k^b})$. Now, at each time k, we can 1) stop sampling and claim that one of the two sequences under examination is generated by Q_1, or 2) continue to the next observation from both sequences to gather more evidence about their statistical behavior, or 3) abandon one or both of the sequences under examination and switch to new sequence(s). We use ϕ_k^a to denote the $\mathcal{F}_k$-measurable switching function at time k that will decide whether we should abandon sequence s_k^a. Similarly, we use ϕ_k^b to denote the $\mathcal{F}_k$-measurable switching function at time k that will decide whether we should abandon sequence s_k^b. In the following, we will also use $\phi_k = (\phi_k^a, \phi_k^b)$. Again, let $\mathcal{T}$ denote the set of all stopping times with respect to the filtration $\mathcal{F}_k$. A stopping time $\tau \in \mathcal{T}$ will decide when we should stop sampling and declare that one of the two sequences that we are currently observing is generated by Q_1. As before, there are two performance indices. The first one is the error probability that the sequence selected is generated by Q_0. Obviously, when we stop at time k, we will select the sequence that has a higher probability of being generated by Q_1, and hence the error probability is $P_e = \min\{P(H^{s_\tau^a} = H_0), P(H^{s_\tau^b} = H_0)\}$. The second performance metric is the average number of samples we take to make a decision, that is, $\mathbb{E}\{\tau\}$.

Our goal is to design the stopping time τ and the switching rules $\phi = \{\phi_1^a, \phi_1^b, \phi_2^a, \phi_2^b, \dots\}$ to solve the following optimization problem:

$$\inf_{\tau \in \mathcal{T}, \phi} \left[P_e + c\mathbb{E}\{\tau\}\right]. \tag{33}$$

We use $\pi_k^a = P(H^{s_k^a} = H_1|\mathcal{F}_k)$ and $\pi_k^b = P(H^{s_k^b} = H_1|\mathcal{F}_k)$ to denote the respective posterior probability that sequences s_k^a and s_k^b are generated by Q_1 after observing $\{Y_1^{s_1^a}, Y_1^{s_1^b}, \ldots, Y_k^{s_k^a}, Y_k^{s_k^b}\}$. Similarly to (3), we have the following recursive formula for $\boldsymbol{\pi}_k = (\pi_k^a, \pi_k^b)$:

$$\begin{aligned}\pi_1^a &= \frac{\pi_0 q_1(Y_1^1)}{\pi_0 q_1(Y_1^1) + (1-\pi_0) q_0(Y_1^1)} \qquad (34)\\ \pi_1^b &= \frac{\pi_0 q_1(Y_1^2)}{\pi_0 q_1(Y_1^2) + (1-\pi_0) q_0(Y_1^2)}\\ &\vdots\\ \pi_{k+1}^a &= \frac{\pi_k^a q_1(Y_{k+1}^{s_{k+1}^a})}{\pi_k^a q_1(Y_{k+1}^{s_{k+1}^a}) + (1-\pi_k^a) q_0(Y_{k+1}^{s_{k+1}^a})}\mathbf{1}(\phi_k^a = 0)\\ &\quad + \frac{\pi_0 q_1(Y_{k+1}^{s_{k+1}^a})}{\pi_0 q_1(Y_{k+1}^{s_{k+1}^a}) + (1-\pi_0) q_0(Y_{k+1}^{s_{k+1}^a})}\mathbf{1}(\phi_k^a = 1),\\ \pi_{k+1}^b &= \frac{\pi_k^b q_1(Y_{k+1}^{s_{k+1}^b})}{\pi_k^b q_1(Y_{k+1}^{s_{k+1}^b}) + (1-\pi_k^b) q_0(Y_{k+1}^{s_{k+1}^b})}\mathbf{1}(\phi_k^b = 0)\\ &\quad + \frac{\pi_0 q_1(Y_{k+1}^{s_{k+1}^b})}{\pi_0 q_1(Y_{k+1}^{s_{k+1}^b}) + (1-\pi_0) q_0(Y_{k+1}^{s_{k+1}^b})}\mathbf{1}(\phi_k^b = 1).\end{aligned}$$

As in the single sequence case, we first consider the situation in which the stopping time τ is restricted to a finite interval $[0, T]$.

At each time k, we need to decide whether to stop sampling or not based on $\mathcal{F}_k$. The minimal expected cost-to-go at time k, $0 \le k \le T$, is a function of $\mathcal{F}_k$, which we will denote by $\tilde{J}_k^T(\mathcal{F}_k)$. Obviously, we have

$$\tilde{J}_T^T(\mathcal{F}_T) = \min\{1-\pi_T^a, 1-\pi_T^b\}.$$

And given $\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})$, we have the following equation:

$$\tilde{J}_k^T(\mathcal{F}_k) = \min\Big\{\min\{1-\pi_k^a, 1-\pi_k^b\},\\ c + \inf_{\phi_k^a, \phi_k^b} \mathbb{E}\left\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k^a, \phi_k^b\right\}\Big\}.$$

We first have the following lemma that converts the finite-horizon version of problem (33) into a Markov optimal stopping problem.

Lemma 8: For each k, the function $\tilde{J}_k^T(\mathcal{F}_k)$ can be written as a function of only $\boldsymbol{\pi}_k$, say $J_k^T(\boldsymbol{\pi}_k)$, and the optimal switching rules $\phi_k = (\phi_k^a, \phi_k^b)$ can be restricted to a class of decision functions that depend only on $\boldsymbol{\pi}_k$.

Proof: Clearly $\tilde{J}_T^T(\mathcal{F}_T) = \min\{1-\pi_T^a, 1-\pi_T^b\}$ is a function of only $\boldsymbol{\pi}_T = (\pi_T^a, \pi_T^b)$, and we write it as $J_T^T(\boldsymbol{\pi}_T)$. For any k, $0 \le k \le T-1$, suppose that $\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})$ is a function of only $\boldsymbol{\pi}_{k+1}$ and we write it as $J_{k+1}^T(\boldsymbol{\pi}_{k+1})$. We have

$$\tilde{J}_k^T(\mathcal{F}_k) = \min\Big\{\min\{1-\pi_k^a, 1-\pi_k^b\},\\ c + \min_{\phi_k^a, \phi_k^b}\left\{\mathbb{E}\left\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k^a, \phi_k^b\right\}\right\}\Big\}.$$

Now, we examine the term $\min_{\phi_k^a, \phi_k^b}\left\{\mathbb{E}\left\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k^a, \phi_k^b\right\}\right\}$. If $\phi_k^a = 0$ and $\phi_k^b = 0$, we have

$$\mathbb{E}\left\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k^a = 0, \phi_k^b = 0\right\}\\ = \int J_{k+1}^T(\boldsymbol{\pi}_{k+1}) f_c(Y_{k+1}^{s_{k+1}^a}|\mathcal{F}_k) f_c(Y_{k+1}^{s_{k+1}^b}|\mathcal{F}_k)\, dy_{k+1}^{s_{k+1}^a} dy_{k+1}^{s_{k+1}^b}, \qquad (35)$$

in which $f_c(\cdot|\mathcal{F}_k)$ is given in (4), while the relationship among $\boldsymbol{\pi}_{k+1}, \boldsymbol{\pi}_k$ and $(y_{k+1}^{s_{k+1}^a}, y_{k+1}^{s_{k+1}^b})$ is given in (34). From (4), it is clear that

$$\mathbb{E}\left\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k^a = 0, \phi_k^b = 0\right\}$$

is a function of only $\boldsymbol{\pi}_k$.

Using the same argument as above, we know that $\mathbb{E}\left\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k^a, \phi_k^b\right\}$ is a function of only $\boldsymbol{\pi}_k$ for other values of ϕ_k^a and ϕ_k^b. As a result, $c + \min_{\phi_k^a, \phi_k^b}\left\{\mathbb{E}\left\{\tilde{J}_{k+1}^T(\mathcal{F}_{k+1})|\mathcal{F}_k, \phi_k^a, \phi_k^b\right\}\right\}$ is a function of only $\boldsymbol{\pi}_k$, and we will use $A_k^T(\boldsymbol{\pi}_k)$ to denote it.

Hence, we know that $\tilde{J}_{k+1}^T(\mathcal{F}_k)$ is a function of only $\boldsymbol{\pi}_k$, and the switching function $\phi_k = (\phi_k^a, \phi_k^b)$ can be limited to functions of $\boldsymbol{\pi}_k$. Moreover, $\{\boldsymbol{\pi}_k : k = 0, 1, \ldots,\}$ forms a Markov process. ■

Similarly to Lemma 2, we also have the following result.

Lemma 9: The functions $J_k^T(\boldsymbol{\pi}_k)$ and $A_k^T(\boldsymbol{\pi}_k)$ are nonnegative concave functions of $\boldsymbol{\pi}_k$. And $J_k^T(1, \pi_k^b) = J_k^T(\pi_k^a, 1) = 0$.

Proof: This result follows from a similar argument to that used in the proof of Lemma 2. ■

Now, we remove the finiteness restriction on the stopping time τ by letting $T \longrightarrow \infty$. Similarly to the single sequence case, we know that the following functions are well-defined:

$$\lim_{T \longrightarrow \infty} J_k^T(\boldsymbol{\pi}) = \inf_{T > k} J_k^T(\boldsymbol{\pi}) = J_k^\infty(\boldsymbol{\pi}). \qquad (36)$$

Also, we have $J_k^\infty(\boldsymbol{\pi}) = J_{k+1}^\infty(\boldsymbol{\pi})$, due to the i.i.d. nature of the observations in each sequence. We will use $J(\boldsymbol{\pi})$ to denote this common function.

Similarly $A_k(\boldsymbol{\pi}) = \lim_{T \longrightarrow \infty} A_k^T(\boldsymbol{\pi})$ is well defined. Moreover, it is independent of k, and we will use $A(\boldsymbol{\pi})$ to denote this common function. Hence, we have

$$J(\boldsymbol{\pi}) = \min\{\min\{1-\pi^a, 1-\pi^b\}, c + A(\boldsymbol{\pi})\}.$$

From the concavity property, the fact that $J(1, \pi^b) = J(\pi^a, 1) = 0$, and the fact that $J(\boldsymbol{\pi}) \le 1 - \max\{\pi^a, \pi^b\}$, we have the following solution for problem (33).

Theorem 10: The optimal stopping time for (33) is given by $\tau_{\text{opt}} = \inf\{k : \max\{\pi_k^a, \pi_k^b\} > 1 - A(\boldsymbol{\pi}_k)\}$. And the switching rule is a function of only $\boldsymbol{\pi}_k$.

VII. Numerical Results

Here, we give an example to illustrate the analytical results of the previous sections. In the example, we assume that under H_0, the observations are i.i.d. Gaussian random variables with means 0 and variances $P + \sigma^2$. Under H_1, the observations are i.i.d. Gaussian random variables with means 0 and variances σ^2.

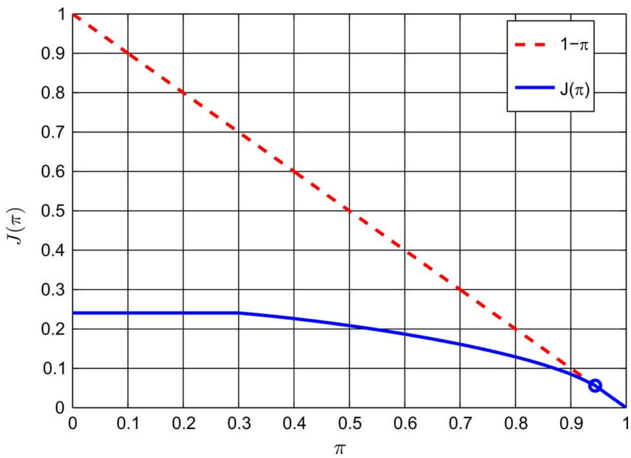

Fig. 5. The cost-to-go function $J(\pi)$.

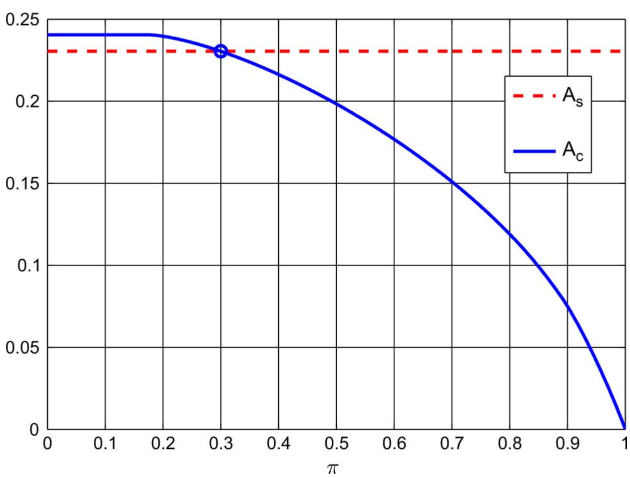

Fig. 6. The relationship between $A_c(\pi)$ and A_s.

We first present numerical results for the single channel case. The cost-to-go function $J(\pi)$ is shown in Fig. 5 for the case of $P = 3$, $\sigma^2 = 1$, $\pi_0 = 0.3$ and $c = 0.01$. The cost-to-go function is computed by recursively using (8). We stop the recursion when the $\mathcal{L}_2$ distance between $J_{k-1}^T(\pi)$ and $J_k^T(\pi)$ is less than 10^{-5}. For this scenario, we find that $\pi_U^* = 0.944$. Fig. 6 shows the relationship between the cost of continuing on the same sequence $A_c(\pi)$ and the cost of switching to another sequence $A_s(\pi)$. It confirms our analysis that we should switch to another sequence when π_k is less than π_0.

After obtaining π_U^*, the optimal algorithm is fixed. In Fig. 7, we show the relationship between the average number of samples one needs to take before the test stops for various values of signal-to-noise ratio (SNR). In generating this figure, we set π_0 to be 0.3 and c to be 0.01. We note that for different values of SNR, the value of π_U^* is different. These results, and the ones in Figs. 8–12 below, were obtained via simulations.

The function generally follows an expected trend. The higher the SNR, the easier it is to distinguish between different channels. Hence fewer steps are required to make a decision. But as it can be seen from Fig. 7, this is a nonmonotonic function. There is an intuitive explanation for this. When the SNR is very low, the information provided by taking more steps does not justify the cost required to take these steps. So a low SNR creates uncertainty about the value of further information. Fig. 8 shows the relationship between the decision delay and SNR, with $c = 0.001$ and all other parameters remain the same as above. From Fig. 8, one can see that the number of steps is generally larger than the case with $c = 0.01$. This is mainly due to the fact that the cost of taking more samples is smaller here.

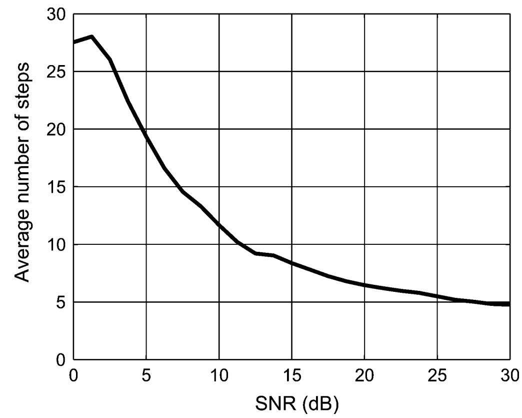

Fig. 7. Average number of steps vs SNR with $c = 0.01$.

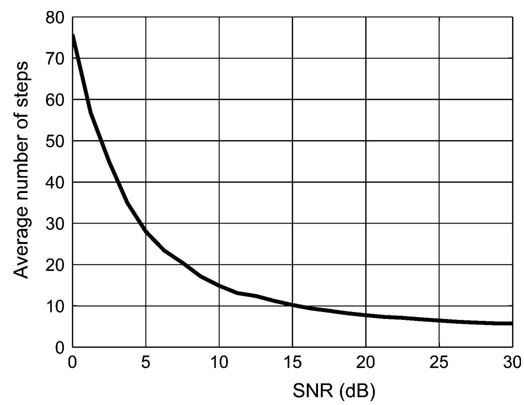

Fig. 8. Average number of steps vs SNR with $c = 0.001$.

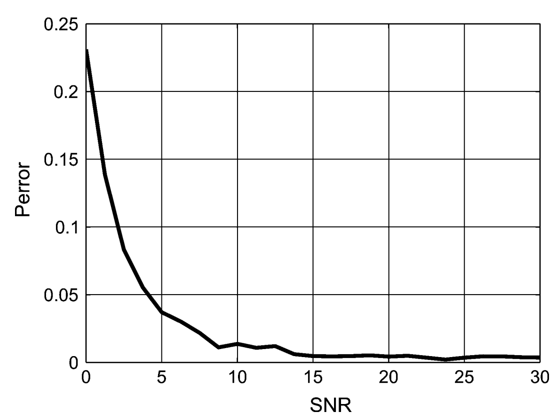

Fig. 9. Error probability $P(H^{s_\tau} = H_0)$ vs SNR with $c = 0.01$.

The probability of error, $P(H^{s_\tau} = H_0)$, is also an important statistic. Fig. 9 shows typical values for the case of $c = 0.01$. We can compare this curve with that for the same parameters except with $c = 0.1$. We show the second function in Fig. 10.

There are a few fundamental differences between Figs. 9 and 10. In Fig. 9, $P(H^{s_\tau} = H_0)$ is rapidly decreasing at the beginning and stabilizes at higher SNR values. On the other hand,

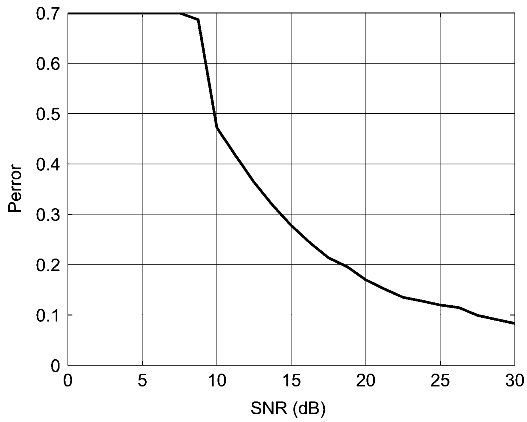

Fig. 10. Error probability $P(H^{s_\tau} = H_0)$ vs SNR with $c = 0.1$.

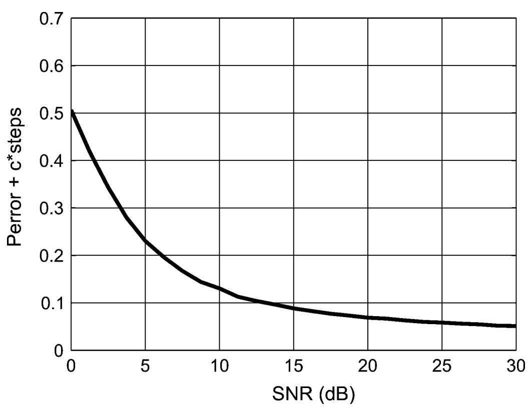

Fig. 11. $P(H^{s_\tau} = H_0) + c\mathbb{E}\{\tau\}$ vs SNR.

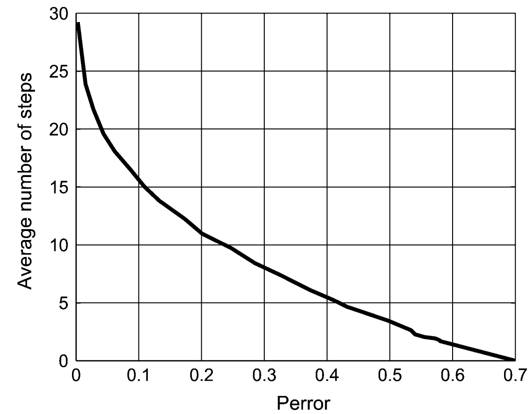

Fig. 12. $\mathbb{E}\{\tau\}$ against $P(H^{s_\tau} = H_0)$.

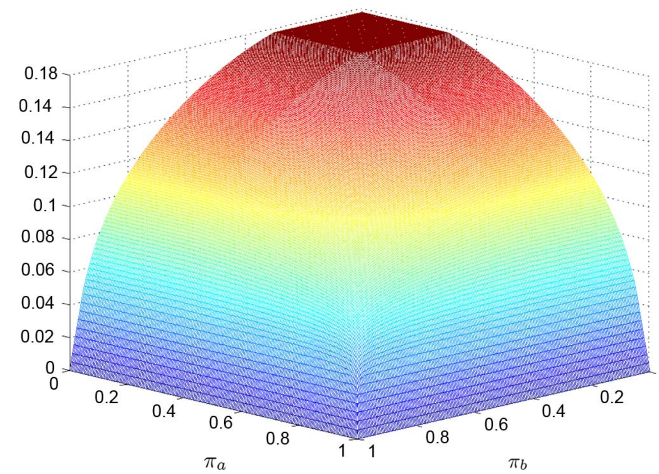

Fig. 13. The two dimensional cost-to-go function $J(\pi_a, \pi_b)$.

in Fig. 10, $P(H^{s_\tau} = H_0)$ is constant for small values of SNR and then decreases after a threshold value. In the second case, the cost of taking any samples for small SNR values is so large, that it does not justify taking any samples at all. In fact the probability of error is equal to the prior probability that H_0 is true in this range. When SNR increases, there is greater justification for more samples and hence the probability of error decreases. In addition, in Fig. 9 the cost of taking samples is smaller. As a result, the number of samples taken is larger, and hence the probability of error is smaller than that of the curve shown in Fig. 10 under the same SNR.

Fig. 11 shows the cost function (1). As expected this is a monotonically decreasing function of SNR. The higher the SNR the lower the objective value.

Fig. 12 shows the average number of steps as a function of the probability of error. This figure was generated for a given SNR value (4.77 dB) and fixed value of c (0.01). As shown in the graph, the higher the probability of error, the lower the number of steps taken. The lower the probability of error required the more the average number of steps taken.

We now present results for the case of two simultaneous observations. Fig. 13 shows $J(\boldsymbol{\pi})$, when $P = 3, \sigma^2 = 1, \pi_0 = 0.3$ and $c = 0.01$. In the two simultaneous observations case, the boundary for stopping is a curve; that is, for any given value of π^b, there exists a π_U^a, such that we stop sampling once π_k^a exceeds π_U^a. The same is true, if we reverse the role of π^a and π^b. In our simulation, we find that if $\pi^b = 0.5$, then $\pi_U^a = 0.955$. And, if $\pi^b = 0.895$, then $\pi_U^a = 0.965$.

VIII. Conclusion

We have considered the problem of quickest sequential search over multiple sequences, in which the goal is to find a sequence drawn from a particular distribution Q_1 among infinitely many sequences in such a manner that a properly defined cost is minimized. We have shown that if one can observe only one sequence at each time, the optimal solution under a Bayesian formulation, which strikes a balance between the cost of sampling and the false alarm probability, is the CUSUM test. The result is derived by assuming that there are infinitely many sequences and one will not switch back to a sequence that has been tested previously. If there are finitely many sequences, the result is also valid if there is no memory of previously collected samples if we switch back to a sequence that has been tested before. We have also investigated the performance of the optimal solution and found that the performance can be written in terms of the performance of the classical SPRT. We have also considered the general case in which one can observe multiple sequences simultaneously and have developed an optimal solution for this general case.

In terms of future work, it is of interest to extend this study to the corresponding problems in continuous time. In this case, the case of Brownian observations will be the first problem to consider. It is also of interest to study the nonhomogeneous case in which the distribution or the prior probability of each sequence is different. It is also of practical interest to study the case in which the number of sequences is finite and one allows memory.

References

[1] L. Lai, H. V. Poor, Y. Xin, and G. Georgiadis, "Quickest sequential opportunity search in multichannel systems," in *Proc. Int. Workshop Applied Probability*, Madrid, Spain, Jul. 2010.

[2] A. Wald, "Sequential tests of statistical hypotheses," *Ann. Math. Stat.*, vol. 16, no. 2, pp. 117–186, 1945.

[3] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," *Ann. Math. Stat.*, vol. 19, no. 3, pp. 326–339, 1948.

[4] M. Sobel and A. Wald, "A sequential decision procedure for choosing one of three hypotheses concerning the unknown means of normal distribution," *Ann. Math. Stat.*, vol. 20, no. 4, pp. 502–522, 1949.

[5] G. Simons, "Lower bounds for average sample number of sequential multihypothesis tests," *Ann. Math. Stat.*, vol. 38, no. 5, pp. 1343–1364, 1967.

[6] D. Teneketzis and Y. C. Ho, "The decentralized wald problem," *Inf. Comput.*, vol. 73, no. 1, pp. 23–44, 1987.

[7] V. V. Veeravalli, T. Basar, and H. V. Poor, "Decentralized sequential detection with a fusion center performing the sequential test," *IEEE Trans. Inf. Theory*, vol. 39, no. 2, pp. 433–442, Mar. 1993.

[8] J. Chamberland and V. V. Veeravalli, "Asymptotic results for decentralized detection in power constrainted wireless sensor networks," *IEEE J. Sel. Areas Commun.*, vol. 22, pp. 1007–1015, Aug. 2004.

[9] J. Chamberland and V. V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Trans. Signal Process.*, vol. 51, pp. 407–416, Feb. 2003.

[10] Y. Yang, R. S. Blum, and B. M. Sadler, "Energy-efficient routing for signal detection in wireless sensor networks," *IEEE Trans. Signal Process.*, vol. 57, pp. 2050–2063, Jun. 2009.

[11] R. S. Blum, S. A. Kassam, and H. V. Poor, "Distributed detection with multiple sensors—Part II: Advanced topics," *Proc. IEEE*, vol. 85, pp. 64–79, Jan. 1997.

[12] H. Chen, B. Chen, and P. K. Varshney, "Further results on the optimality of the likelihood-ratio test for local sensor decision rules in the presence of nonideal channels," *IEEE Trans. Inf. Theory*, vol. 55, no. 2, pp. 828–832, Feb. 2009.

[13] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 23, pp. 201–220, Feb. 2005.

[14] M. DeGroot, *Optimal Statistical Decisions*. New York: Wiley, 1970.

[15] G. V. Haggstrom, "Optimal stopping and experimental design," *Ann. Math. Stat.*, vol. 37, no. 1, pp. 7–29, 1966.

[16] J. Kiefer and J. Sacks, "Asymptotically optimum sequential inference and design," *Ann. Math. Stat.*, vol. 34, no. 3, pp. 705–750, 1963.

[17] H. V. Poor and O. Hadjiliadis, *Quickest Detection*. Cambridge, U.K.: Cambridge Univ. Press, 2009.

[18] A. N. Shiryaev, "On the theory of decision functions and control by observation from incomplete data," in *Proc. 3rd Prague Conf. Inf. Theory*, Prague, Czechoslovakia, 1964, pp. 657–687.

[19] A. Wald, *Statistical Decision Functions*. New York: Wiley, 1950.

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

[21] E. C. Posner and H. Rumsey, Jr, "Continuous sequential decision in the presence of a finite number of hypothesis," *IEEE Trans. Inf. Theory*, vol. IT-12, no. 2, pp. 248–255, Apr. 1966.

[22] E. C. Posner and H. Rumsey, Jr, "An uncertainty function arising in sequential search theory," *J. Appl. Probab.*, vol. 4, pp. 137–143, Apr. 1967.

[23] E. M. Klimko and J. A. Yackel, "Some aspects of search strategies for Wiener process," in *Proc. Symp. Stat. Decision Theory Related Topics*, West Lafayette, IN, Nov. 23–25, 1970.

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

[25] H. Li, "Restless watchdog: Selective quickest spectrum sensing in multichannel cognitive radio systems," *EURASIP J. Adv. Signal Process.*, vol. 2009, pp. 1–12, 2009.

[26] Q. Zhao and J. Ye, "Quickest change detection in multiple on-off processes," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Taipei, Taiwan, Apr. 2009, pp. 3411–3415.

[27] Q. Zhao and J. Ye, "Quickest detection in multiple on-off processes," *IEEE Trans. Signal Process.*, vol. 58, pp. 5994–6006, Dec. 2010.

[28] E. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1–2, pp. 100–115, 1954.

[29] G. Moustakides, "Optimal stopping times for detecting changes in distributions," *Ann. Statist.*, vol. 14, no. 4, pp. 1379–1387, 1986.

[30] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Stat.*, vol. 42, no. 6, pp. 1897–1908, 1971.

[31] A. N. Shiryaev, *Optimal Stopping Rules*. New York: Springer, 1978.

[32] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes—Theory and Application*. Englewood Cliffs, NJ: Prentice-Hall, 1993.

[33] D. Siegmund, *Sequential Analysis*. New York: Springer, 1985.

[34] S. M. Ross, *Stochastic Process*. New York: Wiley, 1996.

Lifeng Lai (M'07) received the B.E. and M.E. degrees from Zhejiang University, Hangzhou, China, in 2001 and 2004, respectively, and the Ph.D. degree from the The Ohio State University at Columbus, OH, in 2007. He was a postdoctoral research associate at Princeton University from 2007 to 2009. Since Aug. 2009, he has been an assistant professor at University of Arkansas, Little Rock. His current research interests include network information theory, information theoretic security, statistical analysis of wireless networks, and biometric security systems.

Dr. Lai was a Distinguished University Fellow of the Ohio State University from 2004 to 2007. He received the Best Paper Award from IEEE Global Communications Conference (Globecom) in 2008, the Best Paper Award from IEEE International Conference on Communications (ICC) in 2011, and the Faculty Early Career Development (CAREER) Award from the National Science Foundation in 2011.

H. Vincent Poor (S'72–M'77–SM'82–F'87) received the Ph.D. degree in electrical engineering and computer science from Princeton University in 1977. From 1977 until 1990, he was on the faculty of the University of Illinois at Urbana-Champaign. Since 1990 he has been on the faculty at Princeton, where he is the Dean of Engineering and Applied Science, and the Michael Henry Strater University Professor of Electrical Engineering. Dr. Poor's research interests are in the areas of stochastic analysis, statistical signal processing and information theory, and their applications in wireless networks and related fields. Among his publications in these areas are *Quickest Detection* (Cambridge University Press, 2009), coauthored with Olympia Hadjiliadis, and *Information Theoretic Security* (Now Publishers, 2009), coauthored with Yingbin Liang and Shlomo Shamai.

Dr. Poor is a member of the National Academy of Engineering and of the National Academy of Sciences, a Fellow of the American Academy of Arts and Sciences, and an International Fellow of the Royal Academy of Engineering (U.K.). He is also a Fellow of the Institute of Mathematical Statistics, the Optical Society of America, and other organizations. In 1990, he served as President of the IEEE Information Theory Society, in 2004–07 as the Editor-in-Chief of these *Transactions*, and in 2009 as General Co-chair of the IEEE International Symposium on Information Theory, held in Seoul, South Korea. He received a Guggenheim Fellowship in 2002 and the IEEE Education Medal in 2005. Recent recognition of his work includes the 2009 Edwin Howard Armstrong Achievement Award of the IEEE Communications Society, the 2010 IET Ambrose Fleming Medal for Achievement in Communications, the 2011 IEEE Eric E. Sumner Award, and an honorary D.Sc. from the University of Edinburgh, awarded in June 2011.

Yan Xin (S'00–M'03–SM'10) received the B.E. degree in electronics engineering from Beijing University of Technology, Beijing, China, in 1992, the M.Sc. degree in mathematics, the M.Sc. degree in electrical engineering, and the Ph.D. degree in electrical engineering from University of Minnesota, Minneapolis, in 1998, 2000, and 2003, respectively. From 2004 to 2008, he was an assistant professor in the Department of Electrical and Computer Engineering, National University of Singapore. He is now a research staff member at NEC Laboratories America Inc., Princeton, New Jersey. His current research interests include cognitive radio, MIMO communication, and network information theory.

Dr. Xin was the co-recipient of the 2004 IEEE Marconi Prize Paper Award in Wireless Communications.

Georgios Georgiadis received the M.S. degree in Electrical Engineering from Stanford University in 2010. He previously received the M.Eng. degree in Electrical and Electronic Engineering from Imperial College, London, U.K., in 2008. He has also been a Visiting Student Research Collaborator (VSRC) at Princeton University. As of September 2010, he is a Ph.D. student in the Computer Science Department at UCLA. He is also a graduate student researcher at the UCLA Vision Lab. His research interests span the fields of computer vision, machine learning, and information theory.