Optimal Two-Stage Bayesian Sequential Change Diagnosis

Xiaochuan Ma and Lifeng Lai
Department of ECE
University of California, Davis
Davis, CA
Email: {xcma, lflai}@ucdavis.edu

Shuguang Cui
Shenzhen Research Institute of Big Data and
Future Network of Intelligence Institute (FNii),
Chinese University of Hong Kong, Shenzhen, China
Email: shuguangcui@cuhk.edu.cn

***Abstract*—In this paper, we formulate and solve a two-stage Bayesian sequential change diagnosis problem. Different from the one-stage sequential change diagnosis problem considered in the existing work, after a change has been detected, we can continue to collect samples so that we can identify the distribution after change more accurately. The goal is to minimize the total cost including delay, false alarm and mis-diagnosis probabilities. We first convert the two-stage sequential change diagnosis problem into a two-ordered optimal stopping time problem. Using tools from multiple optimal stopping time problems, we obtain the optimal change detection and distribution identification rules.**

I. Introduction

Abrupt changes detection and diagnosis problem using sequential observations has many applications, including network monitoring, outage detection and identification in power system, etc. [1]–[3]. These tasks can be formulated and generalized as a sequential change diagnosis (SCD) problem. In particular, an SCD problem can be viewed as a combination of change point detection (CPD) problem and sequential multiple hypothesis testing (SMHT) problem. In CPD problems, the goal is to detect the presence of change in the distribution quickly [4]–[10]. In SMHT problems, the distribution does not change, the focus is to identify the data distribution from K candidate distributions [11]–[14]. In SCD problem, the data distribution will change at an unknown time, from distribution f_0 to one of the K candidate distributions. We need to detect the change point as quickly as possible and identify the distribution after change as accurately as possible. [15] provides early results for SCD problem. [16] generalizes earlier work on SCD and provides more tractable and appropriate performance criteria. In addition, the optimal solution and asymptotically optimal solution of one-stage Bayesian SCD problem are derived in [17] and [18], respectively.

In the one-stage SCD problem [15]–[18], we must detect the change and identify the distribution after change at the same time. In practice, however, after we detect the change, we may still have opportunity to observe extra data samples, which may help us to make a more accurate identification decision. For example, we conduct quality test on a manufacturing process which includes multiple processing components. When a sudden fault occurs in one of the processing components, the goal is to detect the fault quickly and identify the faulted processing component accurately. If the quality testers have chance to observe some extra products after the fault is detected, the accuracy of fault diagnosis can be improved with a relatively low extra delay cost.

Motivated by this, we formulate a two-stage SCD problem. In this problem, we have two stopping times. The first stopping time is to raise an alarm once a change has been detected. After that, we can still collect more observations. The second stopping time will decide when we are ready to make the diagnosis decision. Therefore, in our problem formulation, change detection and distribution identification become two different stages of the whole SCD procedure. Taking advantage of samples of the second stage, it is possible to achieve a lower total cost by improving the identification accuracy.

In this paper, we characterize the optimal solution for the formulated Bayesian two-stage SCD problem. The main idea is to convert the two-stage SCD problem into two ordered optimal stopping time problems, one for change detection stage and the other for distribution identification stage. Then we solve them in reverse order. Firstly, we convert the distribution identification stage of the two-stage SCD problem into an optimal single stopping time problem. Afterwards, we study this problem under a finite-horizon dynamic programming (DP) framework, then expand it to infinite-horizon case and obtain the optimality equation. Applying DP method [19], we solve the optimality equation and get the optimal stopping rule for the distribution identification stage. Following the same method, we can also get the optimal stopping rule of the change detection stage. In addition, we investigate the properties of the cost function. Finally, we validate that the proposed optimal two-stage SCD rule generally outperforms the optimal one-stage SCD rule by simulation.

The remainder of the paper is organized as follows. In Section II, we provide our problem formulation. In Section III, we study the evolution of the posterior probability, and convert the two-stage SCD problem into two optimal single stopping time problems. In Section IV, we derive the optimal rules for the two optimal single stopping time problems. Simulation results are provided in Section V. Finally, we conclude this paper in Section VI. Due to space limitations, we omit the proof details. The proof details can be found in [20].

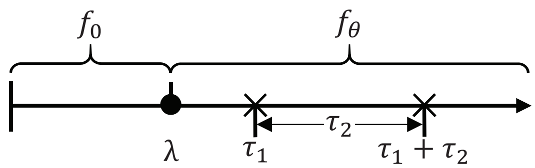

Figure 1. Time ordering of a two-stage SCD process

II. Problem Formulation

Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ that hosts a stochastic process $\{X_n\}_{n\geq 1}$. Let $\lambda : \Omega \mapsto \{0, 1, \ldots\}$ be the time when the distribution of X_n changes and $\theta : \Omega \mapsto \mathcal{I} \triangleq \{1, \ldots, I\}$ be the state after change. In particular, the distribution of X_n is f_0 when $n < \lambda$, and is f_θ when $n \geq \lambda$. λ and θ are independent random variables defined with the distributions

$$\mathbb{P}\{\lambda = t\} = \begin{cases} \rho_0, & \text{if} \quad \mathrm{t} = 0 \\ (1-\rho_0)(1-\rho)^{t-1}\rho, & \text{if} \quad \mathrm{t} \neq 0 \end{cases}$$

and $v_i = \mathbb{P}\{\theta = i\} > 0,\ i \in \mathcal{I}$, where ρ_0, ρ and $\{v_i\}_{i\in\mathcal{I}}$ are constants. Given λ and θ, random variables $\{X_n\}_{n\geq 1}$ are independent. In addition, $\mathbb{F} = (\mathcal{F}_n)_{n\geq 0}$ is the filtration generated by the stochastic process $\{X_n\}_{n\geq 1}$; namely, $\mathcal{F}_0 = \{\emptyset, \Omega\}$ and $\mathcal{F}_n = \sigma(X_1, X_2 \ldots X_n)$.

Our goal is to quickly raise an alarm when the change occurs and further accurately determine the state θ. Towards this goal, we employ a two-stage SCD rule (τ_1, τ_2, d) that includes two stopping times τ_1 and $\tau_1 + \tau_2$ and a decision rule d. Here, τ_1 is the time when we raise an alarm that a change has occurred. In our model, after τ_1, we can keep collecting more observations to make a more accurate diagnosis. Correspondingly, $\tau_1 + \tau_2$ is the time when we make the diagnosis decision d.

Let $\Delta := \{(\tau_1, \tau_2, d) | \tau_1, \tau_1 + \tau_2 \in \mathbb{F}, d \in \mathcal{I} \cup \{0\}\}$ be the set of all possible two-stage SCD rules. Here, $\tau \in \mathbb{F}$ means that τ is a stopping time associated to $\mathbb{F}$. The time ordering of a two-stage SCD process is shown in Fig.1. The possible costs of an SCD rule include costs of delay, false alarm and mis-diagnosis. The delay consists of two parts, $(\tau_1 - \lambda)_+$ and τ_2, which correspond to the change detection stage and the distribution identification stage respectively. The expected costs of them are $\mathbb{E}[c_1(\tau_1 - \lambda)_+]$ and $\mathbb{E}[c_2\tau_2]$, where c_1 and c_2 are per-unit costs associated with each stage. A false alarm occurs when a change alarm is raised before λ. The expected false alarm cost is $\mathbb{E}[a\mathbf{1}_{\{\tau_1<\lambda\}}]$, where a is the penalty factor of false alarm and $\mathbf{1}_{\{.\}}$ is the indicator function. Mis-diagnosis happens when a wrong distribution identification is made, i.e., when $d \neq \theta$. The expected mis-diagnosis cost is $\mathbb{E}[\sum_{i=1}^{I} b_{ij}\mathbf{1}_{\{\infty>\tau_1+\tau_2>\lambda,\theta=i,d=j\}} + b_{0j}\mathbf{1}_{\{\tau_1+\tau_2<\lambda,d=j\}}]$ for $d = j$, where b_{ij} is the penalty factor for wrong decision $d = j$ when $\theta = i$ and $b_{0,j}$ is penalty factor of false alarm of distribution identification stage. We set $b_{ij} = 0$ when $i = j$. Thus the Bayes cost function for an SCD rule $\delta \in \Delta$ is

$$C(\delta) = c_1\mathbb{E}\left[(\tau_1 - \lambda)_+\right] + c_2\mathbb{E}[\tau_2] + a\mathbb{E}[\mathbf{1}_{\{\tau_1<\lambda\}}] + \sum_{j=0}^{I}\mathbb{E}[\sum_{i=1}^{I} b_{ij}\mathbf{1}_{\{\infty>\tau_1+\tau_2>\lambda,\theta=i,d=j\}} + b_{0j}\mathbf{1}_{\{\tau_1+\tau_2<\lambda,d=j\}}]. \tag{1}$$

Our goal is to find the optimal rule $\delta \in \Delta$ that minimizes $C(\delta)$.

In a closely related one-stage SCD problem discussed in [17], the change detection and distribution identification must occur at the same time, and hence there is only one stopping time. We generalize the problem setup in [17] to allow the change detection and identification to occur at different times with the hope of improving the decision accuracy. We assume $c_1 > c_2$. Under this condition, we can improve the identification accuracy with lower per-unit costs in the distribution identification stage.

III. Posterior Analysis

Let $\Pi_n = (\Pi_n^{(0)}, \ldots, \Pi_n^{(I)})_{n\geq 0} \in \mathcal{Z}$ be the posterior probability process defined as

$$\begin{cases} \Pi_n^{(i)} := \mathbb{P}\{\lambda \leq n, \theta = i | \mathcal{F}_n\}, i \in \mathcal{I} \\ \Pi_n^{(0)} := \mathbb{P}\{\lambda > n | \mathcal{F}_n\} \end{cases}$$

where $\mathcal{Z} \triangleq \{\Pi \in [0,1]^{I+1} | \sum_{i\in\mathcal{I}\cup\{0\}} \Pi^{(i)} = 1\}$.

It is easy to check that $\{\Pi_n\}_{n\geq 0}$ is a Markov process satisfying

$$\Pi_n^{(i)} = \frac{D_i(\Pi_{n-1}, X_n)}{\sum_{j\in\mathcal{I}\cup\{0\}} D_j(\Pi_{n-1}, X_n)} \tag{2}$$

where

$$D_i(\Pi, x) := \begin{cases} (1-\rho)\Pi^{(0)} f_0(x) & i = 0 \\ (\Pi^{(i)} + \Pi^{(0)}\rho v_i) f_i(x) & i \in \mathcal{I}. \end{cases}$$

The initial state is

$$\begin{cases} \Pi_0^{(0)} = 1 - \rho_0 & i = 0 \\ \Pi_0^{(i)} = \rho_0 v_i & i \in \mathcal{I} \end{cases}.$$

Proposition 1. *With the process Π_n, we can express* (1) *as*

$$C(\delta) = \mathbb{E}\left[\sum_{n=0}^{\tau_1-1} c_1\left(1 - \Pi_n^{(0)}\right) + c_2\tau_2 + \mathbf{1}_{\{\tau_1<\infty\}} a\Pi_{\tau_1}^{(0)} + \mathbf{1}_{\{\tau_1+\tau_2<\infty\}} \sum_{j=0}^{I} \mathbf{1}_{\{d=j\}} \sum_{i=0}^{I} b_{ij}\Pi_{\tau_1+\tau_2}^{(i)}\right].$$

Define $B_j(\Pi)=\sum_{i=0}^{I}\Pi^{(i)}b_{ij}$, which is the mis-diagnosis cost associated with the decision $d=j$. We have

$$\begin{aligned}
C(\delta) &= \mathbb{E}\left[\sum_{n=0}^{\tau_1-1} c_1\left(1-\Pi_n^{(0)}\right)+c_2\tau_2+\mathbf{1}_{\{\tau_1<\infty\}}a\Pi_{\tau_1}^{(0)}\right.\\
&\qquad\left.+\mathbf{1}_{\{\tau_1+\tau_2<\infty\}}\sum_{j=0}^{I}\mathbf{1}_{\{d=j\}}B_j(\Pi_{\tau_1+\tau_2})\right]\\
&\geq \mathbb{E}\left[\underbrace{\sum_{n=0}^{\tau_1-1} c_1\left(1-\Pi_n^{(0)}\right)+\mathbf{1}_{\{\tau_1<\infty\}}a\Pi_{\tau_1}^{(0)}}_{\text{part 1}}\right.\\
&\qquad\left.+\underbrace{c_2\tau_2+\mathbf{1}_{\{\tau_1+\tau_2<\infty\}}B\left(\Pi_{\tau_1+\tau_2}\right)}_{\text{part 2}}\right]\\
&= C(\tau_1,\tau_2,d^*).
\end{aligned} \tag{3}$$

in which $B(\Pi)=\min_{j\in\mathcal{I}\cup\{0\}}B_j(\Pi)$, the smallest mis-diagnosis cost. From (3), we can see that the optimal decision d^* is the choice that achieves $B(\Pi)$. Then we only need to find the optimal stopping times τ_1 and τ_2, which means that the SCD problem becomes an optimal ordered two-stopping problem. [21] showed that the ordered multiple stopping time problem can be reduced to a sequence of optimal single stopping time problems defined by backward induction. Here we use the same method and reduce the two-stage stopping problem to two optimal single stopping time problems. According to equation (3), the cost can be divided into two parts. The first part is the expected cost of the change detection stage, while the second part corresponds to the distribution identification stage. The first part depends on τ_1 while the second part depends both on τ_1 and τ_2. Let the cost functions of change detection stage and distribution identification stage be

$$C_1(\tau_1)=\sum_{n=0}^{\tau_1-1} c_1\left(1-\Pi_n^{(0)}\right)+\mathbf{1}_{\{\tau_1<\infty\}}a\Pi_{\tau_1}^{(0)} \tag{4}$$

and

$$C_2(\Pi_{\tau_1},\tau_2)=c_2\tau_2+\mathbf{1}_{\{\tau_1+\tau_2<\infty\}}B\left(\Pi_{\tau_1+\tau_2}\right). \tag{5}$$

C_2 is a function of Π_{τ_1} and τ_2 because Π_{τ_1} and the observations from τ_1 to $\tau_1+\tau_2$ are sufficient to calculate $\Pi_{\tau_1+\tau_2}$. Then we have the minimal expected cost for the SCD process,

$$\begin{aligned}
C(\tau_1^*,\tau_2^*,d^*) &= \min_{\tau_1,\tau_1+\tau_2\in\mathbb{F}}\mathbb{E}\left[C_1(\tau_1)+C_2(\tau_1,\tau_2)\right]\\
&= \min_{\tau_1,\tau_1+\tau_2\in\mathbb{F}}\mathbb{E}\left[C_1(\tau_1)+\mathbb{E}\left[C_2(\tau_2)|\Pi_{\tau_1}\right]\right]\\
&= \min_{\tau_1\in\mathbb{F}}\mathbb{E}\left[C_1(\tau_1)+\min_{\tau_1+\tau_2\in\mathbb{F}}\mathbb{E}\left[C_2(\tau_2)|\Pi_{\tau_1}\right]\right].
\end{aligned} \tag{6}$$

By equation (6), we can see that the two-stage stopping time problem becomes two optimal single stopping time problems. The first one is for the identification stage, its goal is finding the optimal τ_2 which minimizes $\mathbb{E}[C_2(\tau_2)|\Pi_{\tau_1}]$ for any given τ_1 and Π_{τ_1}. The second single stopping time problem is to find the best stopping rule for the detection stage, i.e., selecting the optimal τ_1 to minimize the expected cost of the whole SCD process, $C(\tau_1,\tau_2,d^*)$. From the last line of (6), it is easy to see that we can find an optimal τ_1 to minimize the expected cost for the whole SCD process if the optimal rule for τ_2 is known. Therefore, we will solve the SCD problem in a reversed order, i.e., find the optimal rule for the identification stage first, then select the optimal stopping time for the detection stage.

IV. Optimal Solution

A. Finite-Horizon Case

To solve the two-stage SCD problem, we first restrict attention to the finite-horizon case. In particular, we can spend at most T_1 amount of time in the detection stage, i.e., $\tau_1\leq T_1$, and we can spend at most T_2 amount of time in the identification stage, i.e., $\tau_2\leq T_2$. Here, T_1 and T_2 are fixed positive integers. After establishing dynamic programming framework for the finite-horizon case, we will then extend it and further obtain the optimal rule for infinite-case.

In the distribution identification stage, τ_1 and Π_{τ_1} are already given. After we get the optimal τ_2^* and minimum expected cost, $C_2(\Pi_{\tau_1},\tau_2^*)$, for any τ_1 and Π_{τ_1}, we will further introduce the optimal stopping rule for the change detection stage later.

Now we consider the optimal single stopping time problem under a DP framework. Let $S_k^{(2)}$ denote the state of the system at time $k\in[\tau_1,T_2+\tau_1]$. $S_k^{(2)}$ can take $\theta\in\mathcal{I}$, 0 and E (End). Here, $S_k^{(2)}=\theta$ means that the change has happened before k and the distribution after the change is f_θ. $S_k^{(2)}=0$ means that no change has happened before k, which implies a false alarm was made at time τ_1. Once the result of distribution identification is declared, the state of system becomes E. The state evolves as:

$$S_k^{(2)}=g_2(S_{k-1}^{(2)},\lambda,\mathbf{1}_{\{\tau_1+\tau_2\leq k\}})$$

where the transition function g_2 is

$$g_2(s,\lambda,a)=\begin{cases}0, & \text{if } \lambda>k, s\neq E, a=0\\ \theta, & \text{if } \lambda\leq k, s\neq E, a=0\\ E, & \text{if } s=E \text{ or } a=1\end{cases}.$$

The initial state $S_{\tau_1}^{(2)}=0$ if $\lambda>\tau_1$, otherwise $S_{\tau_1}^{(2)}=\theta$. In addition, the observations in this DP framework are the data samples $\{X_n\}_{n\geq1}$.

Under this DP framework, we can see that $\Pi_k^{(i)}=P(S_k^{(2)}=i|\mathcal{F}_k)$. Then the expected cost of the distribution identification stage can be expressed in terms of Π_k as

$$C_2(\Pi_k,k)=c_2(k-\tau_1)+\mathbf{1}_{\{k-\tau_1<\infty\}}B(\Pi_k).$$

Therefore, Π_k is the sufficient statistics for the DP process. Furthermore, we can express the minimum cost-to-go function at time k for this DP problem as

$$V_k^{T_2+\tau_1}(\Pi_k)=B(\Pi_k),\ \text{if } k=T_2+\tau_1 \tag{7}$$

$$V_k^{T_2+\tau_1}(\Pi_k) = \min(B(\Pi_k), \\ c_2 + G_k^{T_2+\tau_1}(\Pi_k)), \text{ if } k < T_2 + \tau_1 \quad (8)$$

where

$$\begin{aligned} G_k^{T_2+\tau_1}(\Pi_k) &= \mathbb{E}[V_{k+1}^{T_2+\tau_1}(\Pi_{k+1})|\mathcal{F}_k] \\ &= \int \left[V_{k+1}^{T_2+\tau_1}(\Pi_{k+1})f(X_{k+1}|\mathcal{F}_k)\right]|_{X_{k+1}=x}dx \\ &= \int \left[V_{k+1}^{T_2+\tau_1}(\Pi_{k+1}(\Pi_k, x))f(x|\Pi_k)\right]dx. \end{aligned} \quad (9)$$

The first item of the minimization in equations (8) is the mis-diagnosis cost for stopping at time k, while the second item corresponds to the cost of proceeding to time $k+1$. In this way, we know that the minimum expected cost for the finite-horizon DP problem is $V_{\tau_1}^{T_2+\tau_1}(\Pi_{\tau_1})$. By solving the optimality equations (7) and (8), we can get the cost $V_{\tau_1}^{T_2+\tau_1}(\Pi_{\tau_1})$ for any given $\tau_1 \le T_1$ and $\Pi_{\tau_1} \in \mathcal{Z}$.

Here, we explain why $G_k^{T_2+\tau_1}$ is a function of Π_k. Firstly, we already know that Π_{k+1} is a function of Π_k and the data sample X_{k+1}. Besides, for any given value of X_{k+1}, $f(X_{k+1}|\mathcal{F}_k)$ is also a function of Π_k because

$$f(X_{k+1}|\mathcal{F}_k) = \sum\nolimits_{j\in\mathcal{I}\cup\{0\}} D_j(\Pi_k, X_{k+1}). \quad (10)$$

Therefore, $G_k^{T_2+\tau_1}$ is a function of Π_k and so is $V_k^{T_2+\tau_1}$.

In addition, there are some useful properties of $G_k^{T_2+\tau_1}(\Pi)$.

Lemma 1. *For $k \in [\tau_1, T_2+\tau_1]$ and $\Pi \in \mathcal{Z}$, (1). $0 \le G_k^{T_2+\tau_1}(\Pi) \le \max_{i,j\in I\cup\{0\}} b_{i,j}$; (2). $G_k^{T_2+\tau_1}(\Pi)$ is concave.*

By Lemma 1, we can conclude that $V_k^{T_2+\tau_1}$ is concave for $k \in [\tau_1, T_2+\tau_1]$. This is because $B(\Pi)$ is concave and the minimization of concave functions will still be concave.

After knowing the optimal stopping rule of the distribution identification stage and the minimum expected cost $V_{\tau_1}^{T_2+\tau_1}(\Pi_{\tau_1})$ for any given τ_1 and Π_{τ_1}, selecting an optimal τ_1 to minimize the whole Bayes cost becomes a single stopping time problem. The method to solve this problem is similar with the distribution identification stage.

Now we consider the optimal stopping problem of the change detection stage under a finite-horizon DP framework. Let $S_k^{(1)}$ denotes the state of the system of the change detection stage at time $k \in [0, T_1]$. $S_k^{(1)}$ can take value 1 (post-change), 0 (pre-change) and E (End). Once a change alarm is raised, the state of system becomes E. The state evolves as:

$$S_k^{(1)} = g_1(S_{k-1}^{(1)}, \lambda, \mathbf{1}_{\{\tau_1\le k\}})$$

with $S_0^{(1)} = 0$, where the transition function g_1 is

$$g_1(s,\lambda,a) = \begin{cases} 0, & \text{if } \lambda > k, s \ne E, a = 0 \\ 1, & \text{if } \lambda \le k, s \ne E, a = 0 \\ E, & \text{if } s = E \text{ or } a = 1 \end{cases}.$$

In addition, the observations of this DP framework are the data samples $\{X_n\}_{n\ge1}$.

Under this DP framework, we can see that $\Pi_k^{(0)} = P(S_k^{(2)} = 0|\mathcal{F}_k)$ and $1 - \Pi_k^{(0)} = P(S_k^{(2)} = 1|\mathcal{F}_k)$. Then the expected cost of the whole SCD process can be expressed in terms of $\{\Pi_n\}_{n\le k}$ as

$$C(k,\tau_2,d^*) = V_k^{T_2+k}(\Pi_k) + \sum_{n=0}^{k-1} c_1\left(1-\Pi_n^{(0)}\right) + \mathbf{1}_{\{k<\infty\}} a\Pi_k^{(0)}.$$

Therefore, $\{\Pi_n\}_{n\le k}$ is the sufficient statistics for the DP process. Furthermore, we can express the minimum cost-to-go function at time k for this DP problem as

$$W_k^{T_1}(\Pi_k) = a\Pi_k^{(0)} + V_k^{T_2+k}(\Pi_k), \text{ if } k = T_1 \quad (11)$$

$$W_k^{T_1}(\Pi_k) = \min(a\Pi_k^{(0)} + V_k^{T_2+k}(\Pi_k), \\ c_1(1-\Pi_k^{(0)}) + U_k^{T_1}(\Pi_k)), \text{ if } k < T_1 \quad (12)$$

where

$$\begin{aligned} U_k^{T_1}(\Pi_k) &= \mathbb{E}[W_{k+1}^{T_1}(\Pi_{k+1})|\mathcal{F}_k] \\ &= \int \left[W_{k+1}^{T_1}(\Pi_{k+1})f(X_{k+1}|\mathcal{F}_k)\right]|_{X_{k+1}=x}dx \\ &= \int \left[W_{k+1}^{T_1}(\Pi_{k+1}(\Pi_k,x))f(x|\Pi_k)\right]dx. \end{aligned} \quad (13)$$

The first item of the minimization in equation (12) is the cost for stopping at time k, while the second item corresponds to the cost of proceeding to time $k+1$. In this way, we know that the minimum expected cost for the finite-horizon DP problem is $W_0^{T_1}(\Pi_0)$.

Similar to $G_k^{T_2+\tau_1}$, $U_k^{T_1}$ is also a function of Π_k. Firstly, we already know that Π_{k+1} is a function of Π_k and the data sample X_{k+1}. Besides, for any given value of X_{k+1}, $f(X_{k+1}|\mathcal{F}_k)$ is also a function of Π_k because of equation (10). Therefore, $U_k^{T_1}$ is a function of Π_k and so is $W_k^{T_1}$.

In addition, there are some useful properties of $U_k^{T_1}(\Pi)$.

Lemma 2. *For $k \in [0, T_1]$ and $\Pi \in \mathcal{Z}$, (1). $0 \le U_k^{T_1}(\Pi) \le a + \max_{i,j\in I\cup\{0\}} b_{i,j}$; (2). $U_k^{T_1}(\Pi)$ is concave.*

By Lemma 2, we can conclude that $W_k^{T_1}$ is concave for $k \in [0, T_1]$. This is because $\Pi^{(0)}(\Pi)$ and $V_k^{T_2+k}(\Pi)$ are also concave and minimization operation preserves concavity.

B. Infinite-horizon Case

After establishing the DP frameworks for the two stages of the finite-horizon SCD problem, we can study the infinite-horizon case by extending the frameworks to infinite-horizon case, i.e., letting T_1 and T_2 go to infinity.

Theorem 1. *For any $\Pi \in \mathcal{Z}$, the infinite-horizon cost-to-go function for the DP process of the identification stage is*

$$V(\Pi) = \lim_{T_2\to\infty} V_k^{T_2+\tau_1}(\Pi) = \min(B(\Pi), c_2 + G_V(\Pi)) \quad (14)$$

where

$$\begin{aligned} G_V(\Pi) &= \mathbb{E}[V(\tilde{\Pi})|\mathcal{F}] \\ &= \int \left[V(\tilde{\Pi}(\Pi,x))f(x|\Pi)\right]dx. \end{aligned} \quad (15)$$

In (15), $\tilde{\Pi}$ *denotes the posterior probability at time next to the time of* Π *and* $\mathcal{F}$*. In addition, the function* $V(\Pi)$ *is concave and bounded as* $0 \leq V(\Pi) \leq \max_{i,j \in I \cup \{0\}} b_{i,j}$.

From optimality equation (14), we know that the optimal rule for this single optimal stopping time problem is

$$\tau_2^* = \inf_{k \geq \tau_1} \{B(\Pi_k) < c_2 + G_V(\Pi_k)\}. \tag{16}$$

That is, we should make an identification when the hyper-plane on the right side exceeds $B(\Pi)$. In addition, the expected cost of the distribution identification stage is $V(\Pi_{\tau_1})$.

Heuristically, at any time $k \geq \tau_1$, we have two choices: (i) Making an identification; (ii) Waiting for the next data sample. With the posterior probability Π_k, the expected cost values of the two choices are $B(\Pi_k)$ and $c_2 + G_V(\Pi_k)$, respectively. The optimal stopping rule (15) tells us that when $B(\Pi_k) < c_2 + G_V(\Pi_k)$, the optimal option is making identification immediately. If $B(\Pi_k) \geq c_2 + G_V(\Pi_k)$, observing more data samples is a better choice.

Based on (12), we can study the infinite-horizon DP process of change detection stage by letting $T_1 \to \infty$.

Theorem 2. *For any* $\Pi \in \mathcal{Z}$*, the infinite-horizon cost-to-go function for the detection stage is*

$$\begin{aligned} W(\Pi) &= \lim_{T_1 \to \infty} W_k^{T_1}(\Pi) \\ &= \min(a\Pi^{(0)} + V(\Pi), c_1(1 - \Pi^{(0)}) + U_W(\Pi)) \end{aligned} \tag{17}$$

where

$$\begin{aligned} U_W(\Pi) &= \mathbb{E}[W(\tilde{\Pi})|\mathcal{F}] \\ &= \int \left[W(\tilde{\Pi}(\Pi, x)) f(x|\Pi) \right] dx. \end{aligned} \tag{18}$$

In (18), $\tilde{\Pi}$ *denotes the posterior probability at time next to the time of* Π *and* $\mathcal{F}$*. In addition, the function* $W(\Pi)$ *is concave and bounded as* $0 \leq W(\Pi) \leq a + \max_{i,j \in \mathcal{I} \cup \{0\}} b_{i,j}$.

From optimality equation (17), we can see that the optimal rule for this problem is

$$\tau_{1,opt} = \inf_{k \geq 0} \{a\Pi_k^{(0)} + V(\Pi_k) < c_1(1 - \Pi_k^{(0)}) + U_W(\Pi_k)\}. \tag{19}$$

That is, we should raise a change alarm when the hyper-plane on the right side exceeds $a\Pi^{(0)} + V(\Pi)$. Finally, we get the minimal expected cost of the two-stage SCD problem, $W(\Pi_0)$.

Similar to the identification stage, at any time k of the change detection stage, we have two choices: (i) Raise a change alarm and enter the distribution identification stage; (ii) Wait for the next data sample. The expected costs of the two choices are $a\Pi_k^{(0)} + V(\Pi_k)$ and $c_1(1 - \Pi_k^{(0)}) + U_W(\Pi_k)$, respectively. The optimal stopping rule (19) tells us that when $a\Pi_k^{(0)} + V(\Pi_k) < c_1(1 - \Pi_k^{(0)}) + U_W(\Pi_k)$, the optimal option is to raise change alarm immediately. If $a\Pi_k^{(0)} + V(\Pi_k) \geq c_1(1 - \Pi_k^{(0)}) + U_W(\Pi_k)$, it is better to wait and observe more samples.

Table I
COMPARISON OF THE BAYESIAN COSTS WITH DIFFERENT c_1 AND r

c_1 \ r	0.02	0.05	0.2	0.5	1
0.005	0.0720	0.0798	0.1009	0.1309	0.1580
0.02	0.2352	0.2511	0.3115	0.3695	0.4016
0.05	0.4763	0.5086	0.6123	0.6853	0.6980
0.2	0.9392	0.9892	1.0021	1.0023	1.0023
0.5	1.0059	1.0062	1.0058	1.0064	1.0067

V. NUMERICAL EXAMPLE

In this section, we provide a numerical example to illustrate the performance of the proposed SCD rule. In our simulation, the observed data samples independently follow a two-dimensional normal distribution, $\mathcal{N}(\mu, \mathbf{I}_2)$. The mean vector μ changes at the change point, from μ_0 to two possible mean vectors μ_1 and μ_2. Denote the mean vectors as $\mu_0 = (\mu_0^{(1)}, \mu_0^{(0)})$ and $\mu_i = (\mu_i^{(1)}, \mu_i^{(0)})$. The elements of these mean vectors satisfy

$$\begin{cases} \mu_1^{(1)} = \mu_0^{(1)} + 1 \\ \mu_1^{(2)} = \mu_0^{(2)} \end{cases} \text{and} \begin{cases} \mu_2^{(1)} = \mu_0^{(1)} + 1 \\ \mu_2^{(2)} = \mu_0^{(2)} + 0.5 \end{cases}.$$

In addition, we set $I = 2$, $\rho_0 = 0$, $\rho = 0.01$, $(v_1, v_2) = (0.3, 0.7)$. Finally, all the penalty factors of false alarm and mis-diagnosis are 1. Table I presents the expected costs of the optimal two-stage SCD rule with different delay penalty factor settings, i.e., with different c_1 and r. Here, $r = c_1/c_2$. The results are estimated by Monte-Carlo simulations.

From Table I, we can see that the performance of the optimal two-stage SCD rule becomes better as c_1 and r get smaller. In particular, with identical c_1, the optimal two-stage SCD rules with $r < 1$ generally outperform the rules with $r = 1$. This result validates that the optimal two-stage SCD rule generally outperforms the optimal one-stage SCD rule. Furthermore, with smaller c_1, the performance improvement brought by reducing r is more significant. The reason is, with a small c_1, we can use more data to improve the accuracy of change detection and identification without a significant increment of the delay cost. On the contrary, when c_1 is large enough, the performance can still be very poor even with a very small r. This result implies that when the unit delay cost is too large, the improvement on diagnosis accuracy becomes too expensive and also negligible.

VI. CONCLUSION

In this paper, we have formulated the Bayesian two-stage sequential change diagnosis problem. We have converted the problem into two optimal single stopping time problems and obtained the optimality equations of them. After solving these equations using dynamic programming, we have obtained the optimal rule for the Bayesian two-stage SCD problem.

ACKNOWLEDGMENT

This work was supported in part by the National Science Foundation under Grants ECCS-1711468, CCF-1717943, CNS-1824553 and CCF-1908258.

REFERENCES

[1] L. Fillatre and I. Nikiforov, "Asymptotically uniformly minimax detection and isolation in network monitoring," *IEEE Transactions on Signal Processing*, vol. 60, no. 7, pp. 3357–3371, Apr. 2012.

[2] Y. C. Chen, T. Banerjee, A. D. Domínguez-García, and V. V. Veeravalli, "Quickest line outage detection and identification," *IEEE Transactions on Power Systems*, vol. 31, no. 1, pp. 749–758, Feb. 2015.

[3] J. Heydari, Z. Sun, and A. Tajer, "Quickest line outage localization under unknown model," in *Proc. IEEE Global Conference on Signal and Information Processing*. Montreal, Canada, Nov. 2017, pp. 1065–1069.

[4] S. Aminikhanghahi and D. J. Cook, "A survey of methods for time series change point detection," *Knowledge and information systems*, vol. 51, no. 2, pp. 339–367, May. 2017.

[5] T. Banerjee and V. V. Veeravalli, "Data-efficient quickest change detection in sensor networks," *IEEE Transactions on Signal Processing*, vol. 63, no. 14, pp. 3727–3735, May. 2015.

[6] Y. Xie and D. Siegmund, "Sequential multi-sensor change-point detection," in *Proc. Information Theory and Applications Workshop*. San Diego, CA, May. 2013, pp. 1–20.

[7] S. Zou, G. Fellouris, and V. V. Veeravalli, "Quickest change detection under transient dynamics: Theory and asymptotic analysis," *IEEE Transactions on Information Theory*, vol. 65, no. 3, pp. 1397–1412, Oct. 2018.

[8] H. V. Poor and O. Hadjiliadis, *Quickest detection*, Cambridge University Press, 2008.

[9] Y. Wang and Y. Mei, "Large-scale multi-stream quickest change detection via shrinkage post-change estimation," *IEEE Transactions on Information Theory*, vol. 61, no. 12, pp. 6926–6938, Dec. 2015.

[10] O. Hadjiliadis, H. Zhang, and H. V. Poor, "One shot schemes for decentralized quickest change detection," *IEEE Transactions on Information Theory*, vol. 55, no. 7, pp. 3346–3359, June. 2009.

[11] C. W. Baum and V. V. Veeravalli, "A sequential procedure for multihypothesis testing," *IEEE Transactions on Information Theory*, vol. 40, no. 6, Nov. 1994.

[12] A. Novikov, "Optimal sequential multiple hypothesis tests," *Kybernetika*, vol. 45, no. 2, pp. 309–330, 2009.

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

[14] J. Heydari and A. Tajer, "Controlled sensing for multi-hypothesis testing with co-dependent actions," in *Proc. IEEE International Symposium on Information Theory*. Vail, CO, June. 2018, pp. 2321–2325.

[15] I. V. Nikiforov, "A generalized change detection problem," *IEEE Transactions on Information theory*, vol. 41, no. 1, pp. 171–187, Jan. 1995.

[16] T. L. Lai, "Sequential multiple hypothesis testing and efficient fault detection-isolation in stochastic systems," *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 595–608, Mar. 2000.

[17] S. Dayanik, C. Goulding, and H. V. Poor, "Bayesian sequential change diagnosis," *Mathematics of Operations Research*, vol. 33, no. 2, pp. 475–496, May. 2008.

[18] S. Dayanik, W. B. Powell, and K. Yamazaki, "Asymptotically optimal bayesian sequential change detection and identification rules," *Annals of Operations Research*, vol. 208, no. 1, pp. 337–370, Sept. 2013.

[19] D. P. Bertsekas, *Dynamic programming and optimal control*, vol. 1, Athena scientific Belmont, MA, 1995.

[20] X. Ma, L. Lai, and S. Cui, "Two-stage bayesian sequential change diagnosis," *IEEE Transactions on Information Theory*, Submitted, 2020.

[21] M. Kobylanski, M. C. Quenez, and E. Rouy-Mironescu, "Optimal multiple stopping time problem," *The Annals of Applied Probability*, vol. 21, no. 4, pp. 1365–1399, Nov. 2011.