Bayesian Quickest Transient Change Detection

K. Premkumar[1], Anurag Kumar[1] and Venugopal V. Veeravalli[2]

[1] Department of Electrical Communication Engineering,
Indian Institute of Science, Bangalore, 560 012, INDIA,
(email: {kprem,anurag}@ece.iisc.ernet.in)

[2] Coordinated Science Lab (CSL),
University of Illinois at Urbana-Champaign, Illinois IL, USA,
(email: vvv@illinois.edu)

Abstract. We consider the problem of *quickest transient change detection* under a Bayesian setting. The change occurs at a random time Γ_1 and *disappears* at a random time $\Gamma_2 > \Gamma_1$. Thus, at any time k, the system can be in one of the following states, i) *prechange*, ii) *in–change*, and iii) *out–of–change*. We model the evolution of the state by a Markov chain. The state of the system can only be observed *partially* from the observations which are obtained sequentially. We formulate the quickest transient change detection problem as a *Partially Observable Markov Decision Process* (POMDP) and obtain the following detection rules for a target probability of false alarm $\mathsf{P}_{\mathsf{FA}} \leqslant \alpha$,

1. MinD (**Min**imum Detection **D**elay), which minimizes the mean detection delay E_{DD}
2. A–MinD (**A**symptotic **Min**imum Detection **D**elay), which is an asymptotic version of the procedure MinD when the mean time until the occurrence of change, $\mathsf{E}\left[\Gamma_1\right]$, goes to infinity
3. MaxP (**Max**imum **P**robability of Detection), which maximizes the probability of detection P_{D}.

We provide numerical examples to illustrate performance of the procedures we propose. We also compare our procedures with the CUSUM procedure for this problem which is known to be optimal under an appropriately defined minimax setting. Our numerical results show that while MinD achieves the least E_{DD} across all events, the CUSUM procedure outperforms MinD when we consider only the events that are stopped in the in–change state.

Keywords. Bayesian change detection, CUSUM

1 Introduction

In the classical problem of quickest change detection (Page (1954), Shiryaev (1978)), the change in the state of a system is modelled as a change in the distribution of an observation process which represents a noisy version of the state of the system. Also, the change is assumed to be persistent, i.e., the state of the system is 0 (pre–change state) before the change and is 1 (in–change state) forever after the change.

However, in some applications like intrusion detection, the change is transient and hence, when the change disappears, the system goes to an out–of–change state (also called state 2). Here, the distribution of the observations is the same, when the system is in state 0 or in state 2. Thus, making a decision about the state of the system as being pre–change or out–of–change, based on the observations, appears to be more challenging than in the case of persistent change.

Related Work: Polunchenko and Tartakovsky (2009) studied the non–Bayesian transient change detection problem and showed that the CUSUM procedure minimizes the supremum detection delay suject to a false alarm constraint. Kligys et al. (1998) formulated a target tracking problem as a sequence of track appearance and disappearance problems and studied the detection performance of CUSUM and Shiryaev–Roberts–Girshik–Rubin procedures.

2 Problem Formulation

We consider a discrete time system. A change occurs at a random time $\Gamma_1 \in \mathbb{Z}_+$ and disappears at a random time $\Gamma_2 \in \mathbb{Z}_+$ (where $\Gamma_2 > \Gamma_1$). This change–model is motivated by the behaviour of physical intrusion in a region under surveillance. Let Θ_k represent the state of the system at time k. We say that Θ_k is 0 before the change occurs (*pre–change*), 1 when the change is present in the system (*in–change*), and 2 after the change disappears (*out–of–change*) Thus, the state space of the system is $\mathbf{\Theta} = \{0, 1, 2\}$. We assume that the evolution of the state process $\{\Theta_k\}$ is Markovian and that the transition probabilities are given by $\mathsf{P}\{\Theta_k = j \mid \Theta_{k-1} = i\} = \rho_{ij},\ i, j \in \{0, 1, 2\}$ with the following conditions: $\rho_{02} = 0, \rho_{10} = 0$ and $\rho_{22} = 1$. Let the distribution of Θ_0 be given by $\mathsf{P}\{\Theta_0 = \theta\} = (1-\nu)\mathbf{1}_{\{\theta=0\}} + \nu\mathbf{1}_{\{\theta=1\}}$, for some $0 \leqslant \nu \leqslant 1$. Observations are obtained sequentially starting from time $k = 1$ onwards. Let $\mathbf{X}_k$ denote the noisy observation of state at time k. The distribution of $\mathbf{X}_k$ when $\Theta_k = 0$ or 2 is given by $F_0(\cdot)$, and that when $\Theta_k = 1$ is given by $F_1(\cdot) \neq F_0(\cdot)$. We assume that the corresponding pdfs f_0 and $f_1 \neq f_0$ exist. Conditioned on Γ_1 and Γ_2, the observations $\mathbf{X}_k$ are i.i.d. across time. This change model is an extension of what is considered in Shiryaev (1978), and the special case of $\Gamma_2 = \infty$ corresponds to the classical problem. Note that at any time k, the state Θ_k is observed only partially through the observation $\mathbf{X}_k$. Also, note that the states "0" and "2" are indistinguishable from the observations.

At every time $k \in \mathbb{Z}_+$, after having observed $\mathbf{X}_k$, the decision maker has to make a decision as to whether the change has occurred (denoted by action "1") or to continue observing (denoted by action "0"). The decision procedure is allowed to depend on the prior information ν, the parameters ρ_{ij}, the observations so far $\mathbf{X}_1, \mathbf{X}_2, \cdots, \mathbf{X}_k$, and the pre–change and the in–change pdfs $f_0(\cdot)$ and $f_1(\cdot)$. Let τ be a stopping time with respect to the sequence $\mathbf{X}_1, \mathbf{X}_2, \cdots$. We now formulate the following transient change detection problems.

P1: Detect the change as early as possible subject to $\mathsf{P}_{\mathsf{FA}} \leqslant \alpha$:

$$\begin{aligned} \min_{\tau} \quad & \mathsf{E}\left[(\tau - \Gamma_1)^+\right] =: \mathsf{E}_{\mathsf{DD}} \\ \text{subject to} \quad & \mathsf{P}\{\tau < \Gamma_1\} \leqslant \alpha. \end{aligned} \tag{1}$$

P2: Detect the change such that the probability of detection is maximum subject to $\mathsf{P}_{\mathsf{FA}} \leqslant \alpha$:

$$\begin{aligned} \max_{\tau} \quad & \mathsf{P}\{\Gamma_1 \leqslant \tau < \Gamma_2\} =: \mathsf{P}_{\mathsf{D}} \\ \text{subject to} \quad & \mathsf{P}\{\tau < \Gamma_1\} \leqslant \alpha. \end{aligned} \tag{2}$$

3 Quickest Change Detection Procedures

We formulate the problems defined in (1) and (2) through a POMDP with the information state at time k being the vector $\mathbf{p}_k = [p_{k,0}, p_{k,1}, p_{k,2}]$ of posterior probabilities where $p_{k,\theta} := \mathsf{P}\{\Theta_k = \theta \mid \mathbf{X}_1, \mathbf{X}_2, \cdots, \mathbf{X}_k\}$, and the action space being $\{0, 1\}$. Optimal policies for the corresponding POMDPs can be obtained using Bellman's equation and are given by

$$\tau^{\mathsf{MinD}} = \inf\left\{k \geqslant 0 : p_{k,0} \leqslant c(1 - p_{k,0}) + \mathsf{E}\left[J^*\left(\mathbf{\Phi}\left(\mathbf{p}_k, \mathbf{X}_{k+1}\right)\right)\right]\right\} \tag{3}$$

$$\tau^{\mathsf{MaxP}} = \inf\left\{k \geqslant 0 : -p_{k,0} + g \cdot p_{k,1} \geqslant \mathsf{E}\left[J^*\left(\mathbf{\Phi}\left(\mathbf{p}_k, \mathbf{X}_{k+1}\right)\right)\right]\right\} \tag{4}$$

where $\mathbf{\Phi}(\mathbf{p}_k, \mathbf{X}_{k+1})$ is a function that yields the next information state $\mathbf{p}_{k+1}$ given $\mathbf{p}_k$ and $\mathbf{X}_{k+1}$, and c is the cost per unit delay and g is the gain per unit detection. Note that c and g are chosen such that $\mathsf{P}_{\mathsf{FA}} = \alpha$ is satisfied. We also obtain a policy $\tau^{\mathsf{A-MinD}}$ which is the limiting case of τ^{MinD} as $\rho_{01} \to 0$ and is given by

$$\tau^{\mathsf{A-MinD}} = \inf\left\{k : p_{k,1} + p_{k,2} \geqslant C\rho_{01}p_{k,0}\right\} \tag{5}$$

where C is a threshold chosen such that $\mathsf{P}_{\mathsf{FA}}(\tau^{\mathsf{A-MinD}}) = \alpha$ (see Raghavan and Veeravalli (2008)).

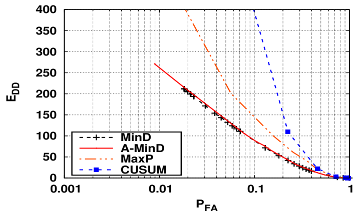

Fig. 1. E_{DD} vs P_{FA}

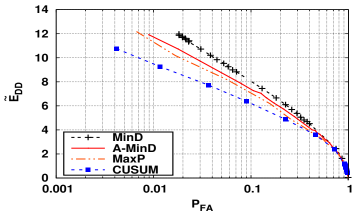

Fig. 2. $\tilde{\mathsf{E}}_{\mathsf{DD}}$ vs P_{FA}

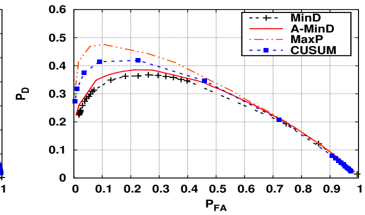

Fig. 3. P_{D} vs P_{FA}

4 Numerical Results and Conclusion

In this section, we study the E_{DD} and the P_{D} performance of the procedures, MinD, A-MinD, MaxP, and CUSUM. Note that $\mathsf{A}-\mathsf{MinD}$ and CUSUM are simple threshold rules whereas the procedures MinD and MaxP require the solution of the DPs defined by (3) and (4). We assume the following parameters for numerical studies, $f_0 \sim \mathcal{N}(0,1), f_1 \sim \mathcal{N}(1,1), \rho_{01} = 0.01, \rho_{12} = 0.1$, and $\nu = 0$. We study the E_{DD}, $\widetilde{\mathsf{E}}_{\mathsf{DD}} := \mathsf{E}\left[(\tau - \Gamma_1)^+ | \Theta_\tau = 1\right]$ and the P_{D} performance of all the detection procedures we propose, for various values of P_{FA}, and plot the results in Figs. 1–3.

Our numerical results show that while the MinD procedure achieves the least E_{DD} across all events (whether stopped in the in–change state or in the out–of–change state), the CUSUM procedure outperforms the MinD procedure when we consider only the events that are stopped in the in–change state. Also, we see from Figs. 1–3 that the E_{DD} of MinD is approximately equal to that of $\mathsf{A}-\mathsf{MinD}$.

Conclusion: We considered the transient change detection problem when the event that causes the change is transient (i.e., not persistent). For a given constraint on P_{FA}, we modeled the transient change detection problem as a POMDP and obtained the following Bayesian transient change detection procedures: 1) MinD, 2) $\mathsf{A}-\mathsf{MinD}$ and 3) MaxP. We showed in the numerical results that the E_{DD} of the MinD procedure is approximately equal to that of the $\mathsf{A}-\mathsf{MinD}$ procedure, the later procedure being much easier to implement. We also showed that the MaxP procedure results in the largest value of P_{D} as expected. Finally, we showed that $\tilde{\mathsf{E}}_{\mathsf{DD}}$ is the smallest for the CUSUM procedure.

Acknowelegements: The work of the first two authors has been supported under the Wireless Sensor Networks project (Project No. DRDO-571) funded by the Defence Research and Development Organization (DRDO), Govt. of India and that of the third author has been supported in part by the NSF grant CCF-0830169, and the U.S. Army Research Office MURI grant W911NF–06–1–0094, through a sub–contract from Brown University at the University of Illinois.

References

Kligys, S., Rozovsky, B., and Tartakovsky, A. G. (1998): "Detection Algorithms and Track Before Detect Architecture Based on Nonlinear Filtering for Infrared Search and Track Systems," Technical Report CAMS-98.9.1, USC.

Page, E. S. (1954): "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1/2, pp. 100–115.

Polunchenko, A. S. and Tartakovsky, A. S. (2009): "Sequential detection of transient changes in statistical distributions," *Technical Report.*

Raghavan, V., and Veeravalli, V. V. (2008): "Quickest Change Detection of a Markov Process Across a Sensor Array," *Submitted, IEEE Trans. Inform. Theory.*

Shiryaev, A. N. (1978): *Optimal Stopping Rules.* Springer, New York.

Veeravalli, V. V. (2001): "Decentralized quickest change detection," *IEEE Trans. Inf. Theory*, vol. 47, no. 4, pp. 1657–1665.