

QUICKEST CHANGE DETECTION IN MULTIPLE ON-OFF PROCESSES

Qing Zhao, Jia Ye

University of California, Davis, CA, 95616

Email: {qzhao,jiaye}@ucdavis.edu

ABSTRACT

A Bayesian formulation of quickest change detection in multiple on-off processes is obtained within a decision-theoretic framework. For geometrically distributed busy and idle times, we show that the optimal joint design of channel switching and change detection has a simple threshold structure under a mild condition. Extensions to arbitrarily distributed busy and idle times, in particular, heavy tail distributions, are discussed. We show that this problem presents a fresh twist to the classic problem of quickest change detection that considers only one stochastic process. We demonstrate that the key to quickest change detection in multiple processes is to abandon the current process when its state is unlikely to change in the near future (as indicated by the measurements obtained so far) and seek opportunities in a new process to avoid realizations of long busy periods. This problem arises in spectrum opportunity detection in cognitive radio networks where a secondary user searches for idle channels in the spectrum.

Index Terms— Quickest change detection, heavy tail distribution, spectrum opportunity detection, cognitive radio.

1. INTRODUCTION

The classic framework of quickest change detection dates back to 1931 [1]. In the conventional setting, the problem is to detect abrupt changes in the distribution of a *single* stochastic process. Specifically, it is assumed that the observations $X_1, X_2, \dots, X_{T_0-1}$ are i.i.d. according to a distribution f_0 . After a random change point T_0 , the observations $X_{T_0}, X_{T_0+1}, \dots$, are i.i.d. according to a different distribution f_1 . The objective is to detect the change point T_0 as quickly as possible subject to a reliability constraint, *i.e.*, a constraint on the probability of false alarm. The first optimal Bayesian change detection algorithm was developed by Shiryaev in 1960's [2], where the change point is assumed to have a geometric/exponential distribution. In the context of opportunity detection, this implies that the connection time (channel “on” time) of the primary system is geometrically/exponentially distributed. Generalizations of Shiryaev's algorithm to arbitrary prior distributions of the change point have been studied (see, for example, [3,4]).

In this paper, we formulate a new form of quickest detection by considering a large number of independent on-off processes. The objective is to catch as quickly as possible an idle/off period in any of the stochastic processes. This problem arises in cognitive radio systems for opportunistic spectrum access, where secondary users need to quickly and reliably detect channels temporarily unused by primary users in a spectrum consisting of multiple channels [5]. The objective is to detect, as soon as possible, whether the sensed channel

has become idle, in order to maximize the transmission time before primary users reclaim the channel. The design constraint is on the maximum probability of declaring a busy channel as idle in order to limit the interference to primary users.

In our previous work [6], we developed a Bayesian formulation of quickest change detection in multiple on-off processes within a decision-theoretic framework. We demonstrated that the key to quickest change detection in multiple processes is to abandon the current process when its state is unlikely to change in the near future (as indicated by the measurements obtained so far) and seek opportunities in a new process. The caveat here is the lost measurements obtained in the abandoned process. Given that change detection relies on the accumulation of “evidence” (measurements), the switching rule needs to be carefully chosen. *An analogy is climbing the corporate ladder: when a long-awaited promotion has yet to come, should one quit, abandon the established seniority, and look for a greener pasture?*

Built upon our previous work [6], this paper addresses the optimal joint design of channel switching and change detection. We show that when the busy and idle times of the on-off processes obey geometrical distributions, the optimal joint design of the switching rule and the detection rule has a simple threshold structure under a mild condition. The threshold structure is with respect to the *a posterior* probability λ_t (given the whole observation history) that the process currently being observed is idle at time t . Specifically, the user should switch to a new channel when $\lambda_t \in [0, \eta_s)$, should continue observing the current channel when $\lambda_t \in [\eta_s, \eta_d)$, and should declare that the current channel is idle when $\lambda_t \in [\eta_d, 1]$, where η_s and η_d are, respectively, the switching and detection thresholds. Furthermore, we show that when the channel switching time is negligible, the optimal switching threshold η_s is the *a priori* probability (before taking any measurements) that a channel is idle, *i.e.*, the average fraction of time that a channel is idle. Extensions to arbitrarily distributed busy and idle times, in particular, heavy tail distributions are discussed. For heavy-tailed busy time, we show that the persistency property of heavy tail distributions make it particularly important to adopt a channel switching strategy (rather than waiting faithfully in a single channel) to avoid realizations of exceptionally long busy periods. To our best knowledge, our previous work [6] and this paper are the first that consider quickest change detection in multiple stochastic processes.

2. QUICKEST DETECTION IN A SINGLE CHANNEL

In this section, we illustrate the problem of quickest opportunity detection by first considering a single channel.

2.1. Problem Formulation

As shown in Fig. 1, suppose that sensing starts at $t = 0$, and the channel becomes idle at a random time $t = T_0$ unknown to the secondary user. The sensing measurements obtained before and after

This work was supported by the Army Research Laboratory CTA on Communication and Networks under Grant DAAD19-01-2-0011, by the Army Research Office under Grant W911NF-08-1-0467, and by the National Science Foundation under Grant CCF-0830685.

T_0 thus have different distributions. Specifically, the sensing measurements $\{X_1, X_2, \dots, X_{T_0-1}\}$ before the change point T_0 are i.i.d. random variables with distribution $f_0(x)$, and the sensing measurements $\{X_{T_0}, X_{T_0+1}, \dots\}$ after the change point T_0 are i.i.d. random variables with distribution $f_1(x)$. The time unit here is the secondary user's sampling period (the time for taking one channel measurement).




Fig. 1. Quickest detection of spectrum opportunities.

At each time instant t , the user aims to infer from measurements $\{X_1, X_2, \dots, X_t\}$ whether a change in the channel state has occurred, *i.e.*, whether to start transmitting or to continue monitoring the channel and taking another measurement X_{t+1} .

Suppose that at time $t = T_d$, the user is convinced that an opportunity has arisen and proceeds to transmit. The problem of quickest opportunity detection can be formulated as choosing a *stopping rule* T_d under the following objective and constraint:

$$\min \mathbb{E}[(T_d - T_0)^+] \quad \text{subject to} \quad \Pr[T_d < T_0] \leq \zeta, \quad (1)$$

where $(T_d - T_0)^+ \triangleq \max\{0, T_d - T_0\}$, and $\mathbb{E}[(T_d - T_0)^+]$ represents the expected detection delay. The constraint in (1) is on the probability that the secondary user starts transmitting when the channel is still busy. It should be capped below an interference constraint ζ . Clearly, the optimal stopping rule T_d should strike a balance between detection delay and detection reliability. Note that in the classic setting, once the change occurs, the process will never change back to distribution f_0 . In our formulation given in Sec. 3, we consider on-off processes with alternating busy and idle periods, a more realistic model of channel occupancy.

2.2. Shiryaev's Algorithm for Quickest Change Detection

Shiryaev's algorithm for quickest change detection was developed within the Bayesian framework [2], where the change point T_0 has a known geometric distribution with parameter p_B . Specifically,

$$\Pr[T_0 = k] = p_B(1 - p_B)^{k-1}(1 - \lambda_0), \quad \forall k > 0,$$

where $\lambda_0 \triangleq \Pr[T_0 = 0]$ is the probability that the change occurs before the observation starts (*i.e.*, sensing starts during an idle period of the channel in the context of opportunistic spectrum access).

It has been shown by Shiryaev that a sufficient statistic for quickest change detection for geometrically/exponentially distributed change point is the *a posteriori* probability λ_t that change has already occurred given the measurements obtained up to t ($t > 0$):

$$\lambda_t \triangleq \Pr[T_0 \leq t | X_1, X_2, \dots, X_t]. \quad (2)$$

Based on Bayes' rule, the sufficient statistic λ_t can be computed recursively at each time t using the new observation $X_t = x$.

Shiryaev's change detection algorithm is given by the following stopping rule on the *a posteriori* probability λ_t .

$$T_d = \inf\{t : \lambda_t \geq \eta_d\}, \quad (3)$$

where the detection threshold η_d is determined by the reliability constraint ζ given in (1). Obtaining the detection threshold η_d in a

closed-form is generally difficult. Setting $\eta_d = 1 - \zeta$ has been shown to be asymptotically optimal as the reliability constraint becomes more strict ($\zeta \rightarrow 0$). In [3, 4], Shiryaev's algorithm has been shown to be asymptotically ($\zeta \rightarrow 0$) optimal when the change point has an arbitrary prior distribution.

3. QUICKEST DETECTION IN MULTIPLE CHANNELS

In this section, we present the Bayesian formulation of quickest opportunity detection in multiple channels developed in [6]. We then prove that the simple threshold policy proposed in [6] is, in fact, optimal for the joint design of channel switching and change detection.

3.1. A Bayesian Setup

We consider a spectrum consisting of a large number of homogeneous channels. In each channel, the channel usage of the primary users is an on-off process with alternating busy and idle periods. These on-off processes are stochastically independent and identical. Let $\{B_i\}_{i=-\infty}^{\infty}$ and $\{I_i\}_{i=-\infty}^{\infty}$ denote, respectively, the lengths of each busy and idle periods in a particular process. We assume that the busy periods $\{B_i\}_{i=-\infty}^{\infty}$ have an identical geometric distribution with parameter p_B , and the idle periods $\{I_i\}_{i=-\infty}^{\infty}$ have an identical geometric distribution with parameter p_I . The average busy and idle times are denoted by $m_B = 1/p_B$ and $m_I = 1/p_I$, respectively. Let λ_0 denote the fraction of channel idle time. It is given by

$$\lambda_0 \triangleq \frac{m_I}{m_B + m_I}. \quad (4)$$

A secondary user starts to sense a channel at $t = 0$. The objective is to catch an idle channel and start transmitting as quickly as possible subject to an interference constraint that caps the probability of transmitting over a busy channel below ζ . The user may switch to a different channel at any time. We assume that the number of channels is large enough so that the user can always switch to a channel that has not been visited. This is equivalent to the case that switching back to a channel is allowed but measurements obtained during previous visits to this channel are discarded. We also assume that channel switching time is negligible. The decision-theoretic formulation and the optimal policy can be easily extended to the general case with positive channel switching time (see Sec. 3.3).

Let L be the number of channels visited by the user before it declares, correctly or falsely, that an opportunity (an idle period) has arisen. It is a random variable depending on the switching and detection rules and the random observations in each channel. Let $T_s(l)$ ($l = 1, \dots, L-1$) denote the time spent in the l -th channel before switching to the $(l+1)$ -th channel. Let $T_d(L)$ denote the time spent in the last channel (the L -th channel) before declaring an opportunity. The problem of quickest change detection in multiple channels can be formulated as jointly choosing a sequence of switching rules $\{T_s(l)\}_{l=1}^{L-1}$ and a detection rule $T_d(L)$ under the following objective and constraint:

$$\begin{aligned} & \min \mathbb{E}[\sum_{l=1}^{L-1} T_s(l) + T_d(L)] \\ & \text{s.t.} \quad \Pr[Z_L(\sum_{l=1}^{L-1} T_s(l) + T_d(L)) = \text{busy}] \leq \zeta, \end{aligned} \quad (5)$$

where $\mathbb{E}[\sum_{l=1}^{L-1} T_s(l) + T_d(L)]$ represents the expected waiting time before catching an idle channel, and $Z_L(t)$ denotes the state of channel L at time t .

We can see from (5) that quickest change detection in multiple stochastic processes is fundamentally different from that in a single process, and is significantly more difficult in that a sequence of stopping rules $(T_s(1), T_s(2), \dots, T_s(L-1), T_d(L))$ need to be designed.

3.2. A Decision-Theoretic Formulation

In our previous work [6], the problem of quickest change detection in multiple on-off processes is formulated as a partially observable Markov decision process (POMDP) over a random horizon. Specifically, the underlying system has three states: 0, 1, and Δ , where 0 and 1 indicate, respectively, that the current process is busy and idle, Δ is an absorbing state, indicating the end of the decision horizon. There are three actions at each decision time: S (switch and take a measurement in a new process), C (continue taking measurements in the current process), and D (declare that a change has already happened in the current process, *i.e.*, the current process is idle). The transition probabilities under each action are given in Fig. 2.




Fig. 2. The state transition diagram.

The observation at time t is X_t under actions S and C. The distribution of X_t is given by either $f_0(x)$ or $f_1(x)$ depending on the current state Z_t of the underlying system. Under action D, no observations are available. The actions of S and C have a unit cost that measures the delay in catching an idle period. Declaring a busy channel as idle incurs a cost of γ that models the tradeoff between detection delay and detection reliability. It is set to satisfy the interference constraint ζ given in (5). Note that it is not necessary to specify the value of γ based on ζ . As shown in Sec. 3.3, the optimal detection rule is specified by a detection threshold chosen to satisfy the interference constraint ζ .

The objective is to choose actions sequentially in time to minimize the expected total cost over an infinite horizon, or equivalently, over a random horizon defined by the hitting time of the absorbing state Δ . It is clear from the cost structure that the expected total cost (excluding the potential cost of γ at the end of the decision horizon) is the expected delay in catching an idle channel.

Since the underlying system state Z_t is not directly observable from the measurements $\{X_t\}$, what we have here is a POMDP. From the fundamental theory of stochastic control, we know that a sufficient statistic for choosing the optimal action at each time is the information state or the belief value: the *a posteriori* probability λ_t that $Z_t = 1$ (the current process is idle) given the measurements obtained up to t . As discussed in the previous section, the same statement was obtained by Shiryaev for quickest change detection in a single process.

It is easy to see that λ_t has the following recursive update depending on the action $a(t-1)$ and the observation $X(t)$.

$$\lambda_t = \begin{cases} \mathcal{T}(\lambda_0|x) & a(t-1) = S, X_t = x \\ \mathcal{T}(\lambda_{t-1}|x) & a(t-1) = C, X_t = x \end{cases}, \quad (6)$$

where $\mathcal{T}(\lambda|x)$ denotes the updated information state based on a new

measurement x . Let $\bar{p} \triangleq 1 - p$ for $p \in [0, 1]$. We have

$$\mathcal{T}(\lambda|x) \triangleq \frac{(\lambda\bar{p}_I + \bar{\lambda}p_B)f_1(x)}{(\lambda\bar{p}_I + \bar{\lambda}p_B)f_1(x) + (\lambda p_I + \bar{\lambda}\bar{p}_B)f_0(x)}. \quad (7)$$

A channel switching and change detection policy π specifies a function that maps an information state $\lambda_t \in [0, 1]$ to an action $a(t) \in \{S, C, D\}$ for each time t . Quickest change detection in multiple on-off processes can thus be formulated as the following stochastic optimization problem:

$$\pi^* = \arg \min_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} R_{\pi(\lambda_t)} | \lambda_0 = \frac{m_I}{m_B + m_I} \right], \quad (8)$$

where $\pi(\lambda_t)$ is the action specified by policy π in information state λ_t , and $R_{\pi(\lambda_t)}$ is the cost incurred under this action and can be easily obtained from the cost structure by averaging over the two possible values of Z_t with λ_t .

3.3. The Optimal Policy: A Threshold Policy

Referred to as the value function, $V(\lambda_t)$ denotes the minimum expected total remaining cost when the current information state is λ_t . It specifies the performance of the optimal policy π^* starting from the information state λ_t . Let $V_S(\lambda_t)$ denote the expected total remaining cost when we take action S at the current time and then follow the optimal policy π^* . Let $V_C(\lambda_t)$ and $V_D(\lambda_t)$ be similarly defined. We thus have

$$V(\lambda_t) = \min\{V_S(\lambda_t), V_C(\lambda_t), V_D(\lambda_t)\}. \quad (9)$$

From the cost structure, we obtain the following.

$$\begin{aligned} V_S(\lambda_t) &= 1 + \int_x P(x; \lambda_0) V(\mathcal{T}(\lambda_0|x)) dx, \\ V_C(\lambda_t) &= 1 + \int_x P(x; \lambda_t) V(\mathcal{T}(\lambda_t|x)) dx, \\ V_D(\lambda_t) &= (1 - \lambda_t)\gamma, \end{aligned} \quad (10)$$

where $P(x; \lambda) = (\lambda\bar{p}_I + \bar{\lambda}p_B)f_1(x) + (\lambda p_I + \bar{\lambda}\bar{p}_B)f_0(x)$ is the probability of observing x when the process has probability λ to be idle. It is easy to see that $V_S(\lambda_t) = V_C(\lambda_0)$ and is independent of λ_t . Furthermore, $V_D(\lambda_t)$ is linearly decreasing with λ_t (see Fig. 3).

Theorem 1 *When $p_B + p_I \leq 1$, the optimal joint design π^* of channel switching and change detection is given by two thresholds η_s and $\eta_d \in (\eta_s, 1]$: switch to a new process whenever $\lambda_t \leq \eta_s$, continue in the current process whenever $\lambda_t \in (\eta_s, \eta_d)$, and declare whenever $\lambda_t \geq \eta_d$. Furthermore, the optimal switching threshold $\eta_s = \lambda_0 \triangleq \frac{m_I}{m_B + m_I}$, the fraction of the idle time.*

Proof: Omitted due to the space limit.

This simple threshold policy agrees with our intuition: switch to a new channel when the prospect of catching an opportunity in a new channel is better than staying in the current channel (*i.e.*, $\lambda_t \leq \lambda_0$). The condition of $p_B + p_I \leq 1$ generally holds. For example, if the average busy and idle times are more than two sample periods, the condition is satisfied.

We point out that for the general case with an arbitrary channel switching time τ_s , the threshold structure of the optimal policy still holds. The only difference is that the optimal switching threshold η_s is smaller than λ_0 when $\tau_s > 0$. This can be shown by noticing that $V_S(\lambda_t) = \tau_s + V_C(\lambda_0)$, *i.e.*, the horizontal line in Fig. 3 is raised up by τ_s and intersect with $V_C(\lambda_t)$ at a point smaller than λ_0 . It




Fig. 3. The threshold structure of the optimal policy.

is possible that when the channel switching time τ_s is sufficiently large, the optimal policy is to never switch channels.

In Fig. 4, we compare the single-channel and multi-channel strategies for different on-off processes. Specifically, we increase both the average busy time m_B and the average idle time m_I while keeping the fraction λ_0 of idle time unchanged. In this case, we observe that the average detection time of the single-channel strategy increases linearly with m_B , as suggested by our intuition. On the other hand, the multi-channel strategy can maintain the same small average detection time regardless of the increase in the length of busy periods in every channel. The performance improvement is thus dramatic when the average busy time is large. This is due to the channel switching strategy that avoids large realizations of busy time and fully exploits the presence of multiple channels.




Fig. 4. Detection delay ($\lambda_0 = 0.7, \zeta = 0.1, \eta_d = 1 - \zeta$).

4. EXTENSION TO ARBITRARY DISTRIBUTIONS

We now consider the case that the busy period B has an arbitrary distribution $\{g_t\}_{t>0}$. We are particularly interested in scenarios where $\{g_t\}_{t>0}$ is a heavy tail distribution. A commonly used heavy tail distribution is the Pareto distribution:

$$g_t \triangleq \Pr[B = t] = \begin{cases} 0 & t < a \\ a^\alpha (t^{-\alpha} - (t+1)^{-\alpha}) & t \geq a \end{cases},$$

where $a > 0$ is the minimum connection time. We consider $\alpha > 1$ so that the connection time has a finite mean but potentially infinite variance. It is easy to show that for a Pareto distributed connection time $B, \forall s > 0$,

$$\Pr[B > \tau + s \mid O > \tau] \nearrow 1 \text{ as } \tau \rightarrow \infty.$$

This is the persistency property of heavy tail distributions. In other words, for heavy-tailed connection time, a connection that has lasted longer than a certain threshold is more likely to persist into the future, and such exceptionally large realizations, albeit rare, dominate the average behavior. This persistency property of heavy-tailed distribution makes it crucial to design an optimal channel switching rule to avoid exceptionally long busy periods.

The threshold policy given in Theorem 1 can also be applied to cases with arbitrarily distributed busy time. In this case, however, the *a posteriori* probability λ_t can no longer be computed recursively as given in (6). Similar to [4], we can obtain a recursive implementation by considering the likelihood ratio $\hat{\lambda}_t$ defined similarly to the case of quickest detection in a single process¹. Specifically, we have

$$\hat{\lambda}_t = \begin{cases} \frac{\lambda_0}{1-\lambda_0}, & t = 0 \\ \left(\frac{P_t^c}{P_t^c} \hat{\lambda}_{t-1} + \frac{p_t}{P_t^c} \frac{f_1(X_t)}{f_0(X_t)} \right), & t > 0 \end{cases}, \quad (11)$$

where $\{p_t\}_{t>0}$ is the distribution of the residual busy time when the user starts sensing a particular channel (after channel switching), $P_t^c = \sum_{k=t+1}^{\infty} p_k$ is the complement cumulative distribution of the residual busy time. Note that the distribution $\{p_t\}_{t>0}$ of the residual busy time is different from $\{g_t\}_{t>0}$, since the time instant at which the user starts sensing a particular channel is not synchronized with the starting point of a busy period (see Fig. 1). The only exception is when $\{g_t\}_{t>0}$ is a geometric distribution. Based on the distribution of the so-called forward renewal time or the residual life of a renewal interval, we have

$$p_t = \frac{1-\lambda_0}{m_B} \sum_{l=t}^{\infty} g_l, \quad t > 0 \quad (12)$$

For Pareto distributed busy time, we have

$$p_t = \begin{cases} \frac{\alpha-1}{a\alpha} (1-\lambda_0) & 0 < t \leq a \\ \frac{a^{\alpha-1}}{\alpha} ((t-1)^{-(\alpha-1)} - t^{-(\alpha-1)}) (1-\lambda_0) & t > a \end{cases},$$

which remains to be a heavy tail distribution with a tail index of $\alpha-1$ (a heavier tail).

With the distribution $\{p_t\}_{t>0}$, the user can recursively update the likelihood ratio $\hat{\lambda}_t$ according to (11). The threshold channel switching and change detection strategy given in Theorem 1 is equivalent to a threshold policy on $\hat{\lambda}_t$: switch to a new process whenever $\hat{\lambda}_t \leq \hat{\eta}_s$, continue in the current process whenever $\hat{\lambda}_t \in (\hat{\eta}_s, \hat{\eta}_d)$, and declare whenever $\hat{\lambda}_t \geq \hat{\eta}_d$, where the switching threshold $\hat{\eta}_s = \frac{\lambda_0}{1-\lambda_0}$, and the detection threshold can be set to $\hat{\eta}_d = \frac{1-\zeta}{\zeta}$.

5. REFERENCES

- [1] W. Shewhart, *Economic Control of Quality of Manufactured Product*. D. Van Nostrand Company, 1931.
- [2] A. N. Shiryaev, "On optimum methods in quickest detection problems," *Theory Prob. Applications*, vol. 8, pp. 22–46, June 1963.
- [3] A. A. Borovkov, "Asymptotically optimal solutions in the change-point problem," *Theory Prob. Applications*, vol. 43, pp. 539–561, Oct. 1997.
- [4] A. G. Tartakovsky and V. V. Veeravalli, "General asymptotic bayesian theory of quickest change detection," *Theory Prob. Applications*, vol. 49, no. 3, pp. 458–497, 2005.
- [5] Q. Zhao and B. Sadler, "A Survey of Dynamic Spectrum Access," *IEEE Signal Processing magazine*, vol. 24, pp. 79–89, May 2007.
- [6] Q. Zhao and J. Ye, "When to Quit for a New Job: Quickest Detection of Spectrum Opportunities in Multiple Channels," in *Proc. of IEEE MILCOM*, Nov., 2008.

¹Here we assume that the user can catch an idle period before it changes back to a busy period. This assumption holds approximately when channel idle periods are longer than the detection time with high probability, which is the scenario when opportunistic spectrum access is feasible.