

QUICKEST CHANGE DETECTION UNDER TRANSIENT DYNAMICS

Georgios Rovatsos, Shaofeng Zou, and Venugopal V. Veeravalli

Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign

Email: rovatso2@illinois.edu, szou3@illinois.edu, vvv@illinois.edu

ABSTRACT

The problem of transient quickest change detection (QCD) is studied, in which the change from the initial to the final phase does not happen instantaneously, but after a series of cascading transient phases of finite durations, each one corresponding to a different probability distribution. The goal is to design a *stopping rule* to detect the change as quickly as possible, subject to false alarm constraints. In previous work, the D-CuSum algorithm was proposed for such a QCD problem. The D-CuSum does not incorporate any prior statistical information about the durations of the transient periods. In this work, we develop an algorithm, the D-S-R algorithm, which incorporates geometric priors on the durations of the transient periods. We compare the D-CuSum and D-S-R algorithms in numerical examples to develop some insights about the role of the prior on the transient durations on the performance.

Index Terms— Bayesian analysis, dynamic CuSum, dynamic Shiryaev-Roberts, quickest change detection, transient dynamics.

1. INTRODUCTION

The theory of quickest change detection (QCD) has seen numerous applications in systems where real-time decision making is crucial, ranging from the detection of subtle faults that may lead to catastrophic failures in large-scale systems to applications in financial surveillance [1]–[8]. In these applications, a change in the statistical properties of the observations will happen in response to an event in the system. The goal of QCD is to detect this change as quickly as possible subject to a tolerable false alarm constraint. In the classical QCD problem [9, 10], the statistical behavior of the observed process is characterized by the pre-change and the post-change distributions, that generate the data before and after the change point, respectively. Two formulations have been proposed for the classical single changepoint QCD problem: i) the *minimax* setting [11]–[13], where the changepoint is modeled as unknown and deterministic and the goal is to minimize a worst-case average detection delay (WADD)

subject to a lower bound on the mean time to false alarm (MTFA); ii) the *Bayesian* setting [14, 15], where the changepoint is a random variable with a known distribution and the goal is to minimize the average detection delay (ADD) subject to a bound on the probability of false alarm.

In this work, we study the transient QCD problem where the initial distribution does not change to the post-change distribution instantly but after a series of transient phases. Each transient phase corresponds to a different distribution on the observed data and each one starts from a respective changepoint. An algorithm for such a QCD problem was proposed in [2] in the context of line outage detection in power systems under transient dynamics. The proposed dynamic CuSum (D-CuSum) algorithm does not incorporate any prior information about the durations of the transient periods.

In practice, and in particular in the power system line outage detection problem, it is reasonable to assume that we have some prior statistical knowledge about the durations of the transient periods. To incorporate this prior knowledge, we formulate the transient QCD problem under the Bayesian setting where the (first) changepoint and the durations of the transient phases are modeled as independent random variables with geometric distributions. By studying the problem of minimizing the detection delay subject to false alarm constraints in a dynamic programming framework, we obtain the structure of the optimal stopping rule. We then propose a tractable dynamic Shiryaev-Roberts (D-S-R) algorithm for the case where the parameter of the geometric distribution of the changepoint approaches zero, which corresponds to the case without any prior knowledge of the changepoint.

In [16], the problem of detecting transient changes is studied, in which the system goes back to the initial distribution after one single transient phase. The goal is to detect the change within the transient period with high probability. This problem is fundamentally different from ours. In our problem, the system does not go back to the initial distribution after the transients, but to a distinct distribution from the initial one. Therefore, the samples after the transient periods still provide useful information about whether the change has happened. Furthermore, our goal here is to minimize the detection delay, subject to a constraint on the MTFA. Hence, a decision to stop after the transients subside is still valid.

This work was supported in part by the National Science Foundation (NSF) under grants CIF 1514245 and ECCS 146231.

2. PROBLEM FORMULATION AND THE DYNAMIC CUSUM ALGORITHM

Consider a process $\{Z_k\}_{k=1}^\infty$ which is observed sequentially by a centralized decision maker. At time instant γ_1 an event causes the initial distribution f_0 to undergo a change. It is assumed that this change occurs in multiple phases, the number of which is known, and is denoted by L . The first $L - 1$ phases correspond to transient intervals of finite durations, after which a persistent phase occurs. In this work, we assume that the observations are independent conditioned on the changepoints. The statistical model is described as follows:

$$Z_k \sim f_\ell, \quad \text{if } \gamma_\ell \leq k < \gamma_{\ell+1}, \quad (1)$$

for $\ell \in \{0, \dots, L\}$, where $\gamma_0 = 1$, $\gamma_{L+1} = \infty$ and the γ 's are unknown but deterministic. The goal in this problem is to detect the change that occurs at γ_1 as quickly as possible, subject to false alarm constraints. The performance of a stopping rule τ is evaluated by the following delay metric:

$$\text{WADD}(\tau) = \sup_{\gamma_1 \geq 1} \text{ess sup } \mathbb{E}_{\gamma_1} \left[(\tau - \gamma_1)^+ \middle| X_1, \dots, X_{\gamma_1-1} \right],$$

where \mathbb{E}_{γ_1} is the expectation when the underlying distribution is the one induced on the observations when a change occurs at γ_1 , and $(x)^+ \triangleq \max\{x, 0\}$. The frequency of the false alarm events is controlled by imposing a constraint on the MTFA. In particular we would like $\mathbb{E}_\infty[\tau] \geq \beta$, for $\beta > 0$, where \mathbb{E}_∞ is the expectation under the probability measure that no change has occurred.

The D-CuSum algorithm is a heuristic solution to the transient QCD problem with non-random changepoints. For the case of L post-change periods, the test statistic is given by:

$$W_k = \max \left\{ \Omega_k^{(1)}, \dots, \Omega_k^{(L)}, 0 \right\}, \quad (2)$$

where $\Omega_k^{(\ell)} = \max\{\Omega_{k-1}^{(\ell)}, \Omega_{k-1}^{(\ell-1)}\} + \log \frac{f_\ell(Z_k)}{f_0(Z_k)}$, for $\ell \in \{1, \dots, L\}$, $\Omega_k^{(0)} \triangleq 0$ for all $k \in \mathbb{Z}$ and $\Omega_0^{(\ell)} \triangleq 0$ for all ℓ . The corresponding stopping time is given by comparing W_k against a pre-determined positive threshold:

$$\tau = \min\{k \geq 1 : W_k > A\}. \quad (3)$$

3. BAYESIAN ANALYSIS

To incorporate prior knowledge of the statistics of the changepoints, we formulate the transient QCD problem under the Bayesian setting, where the changepoints in (1) are modeled as random with known statistics. We assume that the statistics of the changepoint process $\Gamma \triangleq [\Gamma_1 \dots \Gamma_L]$ are described by a *joint-geometric* model [17]. The distributions of Γ_1 and of the duration of each transient phase are modeled as independent random variables with geometric distributions. In particular, we have

$$\mathbb{P}(\Gamma_1 = m) = \rho(1 - \rho)^{m-1}, \quad m \in \mathbb{N}$$

and

$$\mathbb{P}(\Gamma_\ell = m_1 + m_2 | \Gamma_{\ell-1} = m_2) = \rho_{\ell-1, \ell} (1 - \rho_{\ell-1, \ell})^{m_1-1},$$

for $m_1 \in \mathbb{N}$ and $\ell \in \{2, \dots, L\}$.

At each time instant k , the information available to the decision maker is summarized by the information vector $I_k \triangleq \{Z_1, \dots, Z_k\}$, where I_0 denotes the empty *set*. The goal is to use this information to detect the change that occurs at Γ_1 as quickly as possible, subject to false alarm constraints. The change is declared at a stopping time τ and the delay-false alarm tradeoff is characterized by a Bayes risk as follows:

$$R(c) \triangleq \text{PFA} + c\text{ADD} = \mathbb{E}[\mathbb{1}(\{\tau < \Gamma_1\}) + c(\tau - \Gamma_1)^+], \quad (4)$$

where PFA denotes the probability of false alarm, ADD denotes the average detection delay, c is used to model the cost that we suffer for a delay of an extra time instant, and $\mathbb{1}(\{E\})$ is the indicator function of the event $\{E\}$. The goal is to design a stopping rule τ that minimizes the Bayes risk.

3.1. Dynamic Programming Formulation and Optimal Algorithm

In this subsection, we formulate the problem of minimization of the Bayes risk in a dynamic programming framework following similar steps to [17]. In [15], it is shown that for a stopping rule τ , the risk of (4) can be written as

$$R(c) = \mathbb{E}[\mathbb{1}(\{\tau < \Gamma_1\})] + c\mathbb{E}\left[\sum_{k=0}^{\tau-1} \mathbb{P}(\{\Gamma_1 \leq k\})\right]. \quad (5)$$

We first study the finite-horizon dynamic programming case, i.e., we restrict the stages in which decisions are made in the interval $[0, T]$. At each time instant k , we define the state of the system $S_k = \ell$, if the current phase of the system is ℓ . The state space is given by $\mathbb{S} = \{0, \dots, L, t\}$, where $S_k = \ell$ corresponds to the case that $\Gamma_\ell \leq k < \Gamma_{\ell+1}$, for $\ell \in \{0, \dots, L\}$, and $S_k = t$ means that a change has already been declared. The state of the system evolves according to the state evolution equation $S_k = f(S_{k-1}, \Gamma, \mathbb{1}(\{\tau \leq k\}))$, where f is given by

$$f(s, \gamma, \alpha) = \begin{cases} \ell, & \text{if } \Gamma_\ell \leq k < \Gamma_{\ell+1}, \quad \alpha = 0, \\ t, & \text{if } s = t \quad \text{or} \quad \alpha = 1, \end{cases} \quad (6)$$

for $\ell \in \{0, \dots, L\}$, $\Gamma_0 = 1$ and $\Gamma_{L+1} = \infty$. Note that the state information in the present setting is imperfect, i.e., at each time instant we observe a noisy version of the state. The observation equation is given by

$$Z_k = V_k^{(S_k)} \mathbb{1}(\{S_k \neq t\}) + \xi \mathbb{1}(\{S_k = t\}), \quad (7)$$

where $V_k^{(\ell)}$ is drawn from f_ℓ , for $\ell \in \{0, \dots, L\}$, and ξ is a dummy random variable. The Bayes risk of (4) can be written

as an expectation of an additive cost, thus, our problem fits the general dynamic programming setting with termination [18],[19]. The cost-to-go function is defined as follows:

$$\begin{aligned} J_T^T(I_T) &= \mathbb{P}(\{\Gamma_1 > T\}|I_T) \\ J_k^T(I_k) &= \min \left\{ \mathbb{P}(\{\Gamma_1 > k\}|I_k), c\mathbb{P}(\{\Gamma_1 \leq k\}|I_k) \right. \\ &\quad \left. + \mathbb{E} \left[J_{k+1}^T(I_{k+1})|I_k \right] \right\}, \quad 0 \leq k < T. \end{aligned}$$

At each time k there are two possible actions: stop and suffer a cost of stopping at time k , or continue sampling with an expected cost-to-go at time $k + 1$. The minimum expected cost for the finite-horizon optimization problem is $J_0^T(I_0)$.

The sufficient statistic for the solution of this problem is given by the conditional distribution of the state given I_k . The sufficient statistic can be explicitly defined by the vector $\mathbf{p}_k \triangleq [p_{k,0} \dots p_{k,L}]$, where $p_{k,\ell}$ is defined as

$$p_{k,\ell} \triangleq \mathbb{P}(\{S_k = \ell\}|I_k). \quad (8)$$

We now show that \mathbf{p}_k can be obtained from \mathbf{p}_{k-1} recursively. By applying Bayes's rule, it can be shown that

$$\begin{aligned} p_{k,\ell} &= \frac{\mathbb{P}(S_k = \ell|I_{k-1}, Z_k)f(Z_k|I_{k-1})}{f(Z_k|I_{k-1})} \\ &= \frac{\mathbb{P}(S_k = \ell|I_{k-1})f(Z_k|S_k = \ell, I_{k-1})}{\sum_{i=0}^L f(Z_k, S_k = i|I_{k-1})} \\ &= \frac{\mathbb{P}(S_k = \ell|I_{k-1})f(Z_k|S_k = \ell, I_{k-1})}{\sum_{i=0}^L \mathbb{P}(S_k = i|I_{k-1})f(Z_k|S_k = i, I_{k-1})} \\ &\triangleq \frac{A_{k,\ell}}{\sum_{i=0}^L A_{k,i}}, \end{aligned} \quad (9)$$

where $f(\cdot|\cdot)$ denotes the conditional probability density function of Z_k and $A_{k,i} \triangleq \mathbb{P}(S_k = i|I_{k-1})f(Z_k|S_k = i, I_{k-1})$. By the assumption that the observations are independent conditioned on the changepoints, i.e., $I_{k-1} \rightarrow S_k \rightarrow Z_k$,

$$f(Z_k|S_k = i, I_{k-1}) = f(Z_k|S_k = i) = f_i(Z_k).$$

We then compute $A_{k,i}$ as follows:

$$\begin{aligned} A_{k,i} &= \mathbb{P}(S_k = i|I_{k-1})f(Z_k|S_k = i) \\ &= \sum_{j=0}^L \mathbb{P}(S_k = i, S_{k-1} = j|I_{k-1})f_i(Z_k) \\ &= \sum_{j=0}^L \mathbb{P}(S_{k-1} = j|I_{k-1})\mathbb{P}(S_k = i|S_{k-1} = j, I_{k-1})f_i(Z_k) \\ &= (p_{k-1,i-1}\rho_{i-1,i} + p_{k-1,i}(1 - \rho_{i,i+1}))f_i(Z_k), \end{aligned} \quad (10)$$

where the last step is due to the assumption that each phase contains at least one sample, i.e., $\mathbb{P}(S_k = i|S_{k-1} = j, I_{k-1}) = 0$ if $j \neq i$ or $i = 1$. We now study the optimal stopping rule τ_{opt} for the infinite horizon case by letting $T \rightarrow \infty$.

Theorem 1. Let $\mathbf{p} = [p_0 \dots p_L]$ be an element in the L -dimensional simplex $\mathcal{P} \triangleq \{\mathbf{p} : \sum_{j=0}^L p_j = 1\}$. The infinite-horizon cost-to-go for the DP has the following form

$$J(\mathbf{p}) = \min\{p_0, c(1 - p_0) + A_J(\mathbf{p})\},$$

where the function $A_J(\mathbf{p})$ is concave in \mathbf{p} over \mathcal{P} ; is bounded between 0 and 1; and satisfies $A_J(\mathbf{p}) = 0$ over the hyper-plane $\{\mathbf{p} : p_0 = 0\}$.

Proof. See [20]. \square

By Theorem 1, the optimal stopping time is

$$\tau_{opt} = \inf_{k \in \mathbb{N}} \{k : p_{k,0}(1 + c) - c < A_J(\mathbf{p}_k)\}. \quad (11)$$

3.2. The Dynamic Shiryaev-Roberts Algorithm

The form of $A_J(\mathbf{p})$ is not explicit, thus, τ_{opt} can only be computed numerically. We next construct a more simplified and tractable stopping rule. Note that we are only interested in detecting the change from phase 0 to phase 1 as quickly as possible, which is equivalent to detecting whether the system is still in state 0. Based on such an understanding, we compute $p_{k,0}$, which is the posteriori probability that the system is at state 0 given current observations, and consider the following alternative stopping time:

$$\tau_t = \inf_{k \in \mathbb{N}} \{k : p_{k,0} < B\}, \quad (12)$$

where B is an appropriate threshold. We are interested in the regime in which $\rho \rightarrow 0$, under which the stopping rule (12) converges to a non-Bayesian rule with respect to Γ_1 , the D-S-R algorithm. It should be noted that $\rho \rightarrow 0$ is equal to uniformizing the changepoint Γ_1 , which is equivalent to the case without any prior on the Γ_1 . The algorithm is named after the Shiryaev-Roberts test since the latter is the limit of the Shiryaev algorithm for single changepoint Bayesian QCD as the geometric parameter approaches zero [9].

We define the following invertible mapping:

$$q_{k,\ell} = \frac{p_{k,\ell}}{\rho p_{k,0}} \Leftrightarrow p_{k,\ell} = \frac{q_{k,\ell}}{\sum_{j=0}^L q_{k,j}}. \quad (13)$$

From the above, it is straightforward to show

$$p_{k,0} = \frac{1}{1 + \rho \sum_{j=1}^L q_{k,j}}, \quad (14)$$

and $q_{k,\ell}$ can be computed recursively by

$$q_{k,\ell} = \frac{(q_{k-1,\ell-1}\rho_{\ell-1,\ell} + q_{k-1,\ell}(1 - \rho_{\ell,\ell+1}))f_\ell(Z_k)}{(1 - \rho)f_0(Z_k)}, \quad (15)$$

with the following priors:

$$q_{0,0} = \frac{1}{\rho}, \quad \text{and} \quad q_{0,\ell} = 0, \quad \ell \in \{1, \dots, L\}. \quad (16)$$

We further define $r_{k,\ell} = \lim_{\rho \rightarrow 0} q_{k,\ell}$, for $\ell \in \{1, \dots, L\}$. The recursive form of $r_{k,\ell}$ can be computed as follows:

$$r_{k,\ell} = \frac{(r_{k-1,\ell-1}\rho_{\ell-1,\ell} + r_{k-1,\ell}(1 - \rho_{\ell,\ell+1})) f_\ell(Z_k)}{f_0(Z_k)},$$

for $\ell \geq 2$, and

$$r_{k,1} = \frac{(1 + r_{k-1,\ell}(1 - \rho_{\ell,\ell+1})) f_\ell(Z_k)}{f_0(Z_k)},$$

with the following priors:

$$r_{0,\ell} = 0, \quad \text{for } \ell \in \{1, \dots, L\}.$$

Then, as $\rho \rightarrow 0$, by (14) and the definition of $r_{k,\ell}$, the stopping time (12) reduces to the D-S-R algorithm as follows:

$$\tau_A = \inf_{k \in \mathbb{N}} \left\{ k : V_k = \log \left(\sum_{\ell=1}^L r_{k,\ell} \right) \geq A \right\}, \quad (17)$$

where A is an appropriately chosen threshold.

4. SIMULATION RESULTS AND DISCUSSION

In this section, we compare the performance of the D-CuSum algorithm and the D-S-R algorithm under a minimax setting with respect to Γ_1 . The performance of each test is evaluated in terms of WADD and MTFA. For both algorithms, the WADD is simulated by considering the worst case in terms of delay, which corresponds to $\Gamma_1 = 1$. The MTFA is simulated by generating all the samples from f_0 , which corresponds to $\Gamma_1 = \infty$. The two metrics are computed through regular Monte Carlo simulations for a variety of threshold values.

We first compare the algorithms for the case of one transient phase ($L=2$). The parameters of the changepoint process are chosen to be $\rho_{1,2} = 0.1$. We choose the probability distributions to be Gaussian distributions with mean shift: $f_0 = \mathcal{N}(0, 1)$, $f_1 = \mathcal{N}(0.2, 1)$ and $f_2 = \mathcal{N}(0.4, 1)$. In Fig. 1, we plot WADD versus MTFA for both algorithms. In this case the D-CuSum algorithm outperforms the D-S-R algorithm.

We next compare the two algorithms for the case of one transient phase ($L=2$). We set a different $\rho_{1,2} = 0.001$, and choose different probability distributions: $f_0 = \mathcal{N}(0, 1)$, $f_1 = \mathcal{N}(3, 1)$ and $f_2 = \mathcal{N}(0.1, 1)$. In Fig. 2, we plot WADD versus MTFA for the two algorithms. It can be seen that in this case the performance of the D-S-R test is superior.

Next, we compare the two algorithms for the case of four transient periods ($L = 5$). We set the parameters of the changepoint process to be $\rho_{1,2} = \rho_{2,3} = \rho_{3,4} = \rho_{4,5} = \rho_{\text{trans}} = 0.1$. We choose the probability distributions to be Gaussian distributions: $f_0 = \mathcal{N}(0, 1)$, $f_1 = \mathcal{N}(1, 1)$, $f_2 =$

$\mathcal{N}(2, 1)$, $f_3 = \mathcal{N}(2.5, 1)$, $f_4 = \mathcal{N}(3, 1)$ and $f_5 = \mathcal{N}(3.5, 1)$. In Fig. 3, we plot the WADD as a function of MTFA for these two algorithms. It can be seen that in this case the D-CuSum algorithm outperforms the D-S-R algorithm.

From the three figures, we conclude that the two algorithms have similar performance, even though the D-CuSum algorithm does not exploit the prior information of the changepoint process. The reason may be due to the assumption of the geometric prior. It is therefore of interest to study the problem with a different prior, for which the Bayesian solution given will no longer be stationary in time.




Fig. 1. WADD versus mean time to false alarm for $L = 2, \rho_{1,2} = 0.1, f_0 = \mathcal{N}(0, 1), f_1 = \mathcal{N}(0.2, 1), f_2 = \mathcal{N}(0.4, 1)$.




Fig. 2. WADD versus mean time to false alarm for $L = 2, \rho_{1,2} = 0.001, f_0 = \mathcal{N}(0, 1), f_1 = \mathcal{N}(3, 1), f_2 = \mathcal{N}(0.1, 1)$.




Fig. 3. WADD versus mean time to false alarm for $L = 5, \rho_{\text{trans}} = 0.1, f_0 = \mathcal{N}(0, 1), f_1 = \mathcal{N}(1, 1), f_2 = \mathcal{N}(2, 1), f_3 = \mathcal{N}(2.5, 1), f_4 = \mathcal{N}(3, 1), f_5 = \mathcal{N}(3.5, 1)$.

5. REFERENCES

- [1] G. Rovatsos, X. Jiang, A. D. Domínguez-García, and V. V. Veeravalli, "Comparison of statistical algorithms for power system line outage detection," in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Apr. 2016.
- [2] G. Rovatsos, J. Jiang, A. D. Domínguez-García, and V. V. Veeravalli, "Statistical power system line outage detection under transient dynamics," *submitted to IEEE Transactions on Signal Processing*, 2016.
- [3] K. Mechitov, W. Kim, G. Agha, and T. Nagayama, "High-frequency distributed sensing for structure monitoring," in *Proc. First Intl. Workshop on Networked Sensing Systems (INSS)*, 2004.
- [4] J. A. Rice, K. Mechitov, S. H. Sim, T. Nagayama, S. Jang, R. Kim, B. F. Spencer, G. Agha, and Y. Fujino, "Flexible smart sensor framework for autonomous structural health monitoring," *Smart Structures and Systems*, vol. 6, no. 5-6, pp. 423–438, July 2010.
- [5] A. G. Tartakovsky, B. L. Rozovskii, R. B. Blazek, and H. Kim, "A novel approach to detection of intrusions in computer networks via adaptive sequential and batch-sequential change-point detection methods," *IEEE Transactions on Signal Processing*, vol. 54, no. 9, pp. 3372–3382, Sept. 2006.
- [6] S. E. Fienberg and G. Shmueli, "Statistical issues and challenges associated with rapid detection of bio-terrorist attacks," *Statistics in Medicine*, vol. 24, no. 4, pp. 513–529, 2005.
- [7] M. Frisn, "Optimal sequential surveillance for finance, public health, and other areas," *Sequential Analysis*, vol. 28, no. 3, pp. 310–337, 2009.
- [8] L. Lai, Y. Fan, and H. V. Poor, "Quickest detection in cognitive radio: A sequential change detection framework," in *Global Telecommunications Conference (GLOBECOM)*, Nov. 2008, pp. 1–5.
- [9] V. V. Veeravalli and T. Banerjee, "Quickest change detection," *Academic press library in signal processing: Array and statistical signal processing*, vol. 3, pp. 209–256, 2013.
- [10] H. V. Poor and O. Hadjiladis, *Quickest Detection*, Cambridge University Press, 2009.
- [11] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, no. 6, pp. 1897–1908, Dec. 1971.
- [12] M. Pollak, "Optimal detection of a change in distribution," *Annals of Statistics*, vol. 13, no. 1, pp. 206–227, Mar. 1985.
- [13] G. V. Moustakides, "Optimal stopping times for detecting changes in distributions," *Ann. Statist.*, vol. 14, no. 4, pp. 1379–1387, Dec. 1986.
- [14] A. N. Shiryaev, "On optimum methods in quickest detection problems," *Theory of Probability & Its Applications*, vol. 8, no. 1, pp. 22–46, 1963.
- [15] A. N. Shiryaev, *Optimal Stopping Rules*, New York: Springer-Verlag, 1978.
- [16] B. K. Gupi, L. Fillatre, and I. Nikiforov, "Sequential detection of transient changes," *Sequential Analysis*, vol. 31, no. 4, pp. 528–547, 2012.
- [17] V. Raghavan and V. V. Veeravalli, "Quickest change detection of a Markov process across a sensor array," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1961–1981, 2010.
- [18] D. P. Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*, Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1987.
- [19] V. V. Veeravalli, "Decentralized quickest change detection," *IEEE Transactions on Information theory*, vol. 47, no. 4, pp. 1657–1665, 2001.
- [20] G. Rovatsos, S. Zou, and V. V. Veeravalli, "Quickest change detection under transient dynamics," *in preparation*, 2017.