

Quickest Detection of Dynamic Events in Networks

Shaofeng Zou^{1b}, *Member, IEEE*, Venugopal V. Veeravalli^{2b}, *Fellow, IEEE*,
Jian Li^{3b}, *Member, IEEE*, and Don Towsley^{4b}, *Life Fellow, IEEE*

Abstract—The problem of quickest detection of dynamic events in networks is studied. At some unknown time, an event occurs, and a number of nodes in the network are affected by the event, in that they undergo a change in the statistics of their observations. It is assumed that the event is dynamic, in that it can propagate along the edges in the network, and affect more and more nodes with time. The event propagation dynamics is assumed to be unknown. The goal is to design a sequential algorithm that can detect a “significant” event, i.e., when the event has affected no fewer than η nodes, as quickly as possible, while controlling the false alarm rate. Fully connected networks are studied first, and the results are then extended to arbitrarily connected networks. The designed algorithms are shown to be adaptive to the unknown propagation dynamics, and their first-order asymptotic optimality is demonstrated as the false alarm rate goes to zero. The algorithms can be implemented with linear computational complexity in the network size at each time step, which is critical for online implementation. Numerical simulations are provided to validate the theoretical results.

Index Terms—Adaptive algorithms, asymptotic optimality, event propagation, Network-CuSum (N-CuSum), Spartan-CuSum (S-CuSum), structured network.

I. INTRODUCTION

IN THE problem of quickest change detection (QCD), a stochastic system is observed sequentially. At some unknown time, a change occurs that changes the data generating process. Observations are taken sequentially with time, and the objective is to detect the change as quickly as

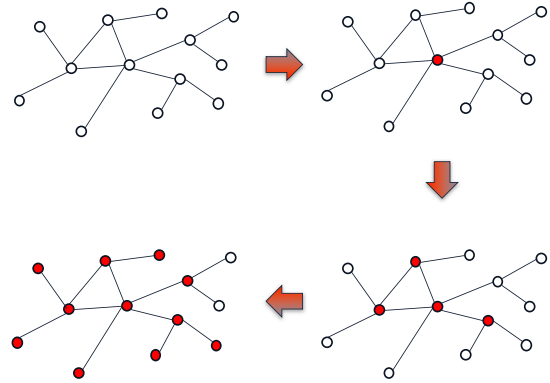


Fig. 1. A dynamic event propagates in a network with time.

possible subject to false alarm constraints (see [3]–[6] for an overview). The QCD framework models a wide range of applications, e.g., fraud detection, intrusion detection, environmental monitoring, line outage detection in power systems, quality control in online manufacturing systems and spectrum monitoring in wireless communications. However, in many applications, e.g., epidemic detection [7], [8], opinion mining in social networks [9], anomalous event detection in sensor networks (e.g., internet of battlefield things) [10], detection of malicious code spreading in computer networks [11], the data are usually collected from networks with certain underlying topologies. Following the occurrence of an event, it propagates dynamically across the network, affects more and more nodes, and changes their data generating behaviors with time (see Fig. 1). The propagation dynamics is usually unknown in practice, and depends on the underlying network topology.

Motivated by these applications, we study the problem of quickest detection of dynamic events in networks. Suppose a network is monitored in real time by a set of L nodes that communicate with a fusion center. At some unknown time, an event occurs in the network that causes eventual changes in the observations of a connected subset of nodes. The event occurs at a connected subset of nodes and then dynamically propagates along the edges in the network, and the affected nodes form a connected sub-graph, the size of which grows with time. The propagation dynamics are assumed to be unknown, i.e., the set of nodes and the order in which they are affected are unknown. We are interested in detecting a “significant” event, i.e., one that affects $\eta \geq 1$ nodes as quickly as possible, subject to false alarm constraints.

Manuscript received July 17, 2018; revised October 10, 2019; accepted October 13, 2019. Date of publication October 18, 2019; date of current version March 17, 2020. This work was supported in part by the Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196. S. Zou and V. V. Veeravalli were supported in part by the National Science Foundation (NSF) under Grant CCF 16-18658 and in part by the Air Force Office of Scientific Research (AFOSR) under Grant FA9550-16-1-0077. This article was presented in part at the 2018 IEEE International Conference on Acoustics, Speech, and Signal Processing, and in part at the Asilomar Conference on Signals, Systems, and Computers.

S. Zou is with the Department of Electrical Engineering, University at Buffalo, The State University of New York, Buffalo, NY 14221 USA (e-mail: szou3@buffalo.edu).

V. V. Veeravalli is with the Coordinated Science Laboratory, ECE Department, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA, and also with the Department of Statistics, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: vvv@illinois.edu).

J. Li is with the Department of Electrical and Computer Engineering, Binghamton University, The State University of New York, Binghamton, NY 13902 USA (e-mail: lij@binghamton.edu).

D. Towsley is with the College of Information and Computer Sciences, University of Massachusetts Amherst, Amherst, MA 01003 USA (e-mail: towsley@cs.umass.edu).

Communicated by L. Ying, Associate Editor for Communication Networks.

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

Digital Object Identifier 10.1109/TIT.2019.2948350

A. Related Works

The problem in this paper is closely related to the problem of QCD under the multi-channel setup, in which one or multiple unknown nodes perceive a change simultaneously [12]–[16], or alternatively, at different times [17]–[19]. The major differences from these previous works lie in that: (i) we are interested in detecting whether the event has affected at least η nodes, i.e., the event is “significant” enough, whereas previous works focus on the special case with $\eta = 1$, i.e., whether the event has occurred or not; (ii) instead of considering the worst-case performance over all possible times that the nodes are affected [19] or taking a Bayesian approach [17], [18], we assume that the times that the nodes are affected are deterministic and unknown, and we are interested in designing algorithms that adapt to unknown propagation dynamics; and (iii) we consider *structured* networks, over which events can only propagate along network edges.

On a temporal scale, the data generating distribution of the whole network dynamically changes over time. As the event affects more nodes with time, the network goes through multiple transient phases in which the sets of affected nodes are different. This is related to the problem of QCD under transient dynamics [20]–[23], where after an event occurs, the pre-change distribution does not change to a persistent post-change distribution instantaneously, but only after a number of transient phases. Each transient phase is associated with a distinct data generating distribution. In [20]–[23], it is assumed that the number of transient phases and the data generating distributions associated with each phase are known. In this paper, event propagation dynamics are unknown. Therefore, the results in [20]–[23] cannot be directly applied to solve the problem here. Moreover, in [22] a Bayesian approach is employed, where it is assumed that the transient durations are geometrically distributed with statistics known to the decision maker. In contrast to [22], in this paper, we make no probabilistic assumptions on the times at which the nodes are affected.

The data generating distributions for the whole network before and after η nodes are affected are both composite, i.e., belong to a set of distributions, as they are determined by the unknown subset of affected nodes with unknown change times. Therefore, the problem in this paper is related to the problem of QCD with composite pre-/post-change distributions [24]–[27]. However, our problem differs from these works in the following ways. First, the data generating distribution before η nodes are affected is composite, whereas in [24], [25], [27], only the post-change distribution is composite. Second, our test statistics can be computed efficiently at each time step with a computational complexity linear in the network size (number of nodes), whereas a sliding window approach is usually used to control the computational cost in [24]–[27]. Third, in our problem, the distribution of samples before and after the nodes are affected by the event are arbitrary (not necessarily belonging to an exponential family), and the parameters of the data generating distribution for the whole network (times at which the nodes are affected by the event) are discrete, and do not necessarily belong to a compact parameter space.

The offline setting of our problem has been extensively studied in the literature [28]–[36]. The aim of these works is to detect whether there exists a subset of nodes in the network, which have certain geometric structures (e.g., connected subgraphs), and the observations received by this subset of nodes are generated from a distribution different from the one that generates the samples for the rest of the nodes in the network. Our problem is the online dynamic version with a growing anomalous geometric structure. Hence, a small computational complexity at each time step is important in order for making timely decisions. However, many previous works are based on the scan statistic that scans over all connected subgraphs, which is computationally inefficient for large networks, and thus cannot be directly applied to our online setting. As will be shown later, our algorithms can be updated recursively with computational complexity linear in the network size at each time step, and we do not reprocess the previous data over and over again.

B. Contributions

In this paper, we start with fully connected networks, and then extend to arbitrarily connected networks.

For fully connected networks, the event can propagate from any node to any other node. Then, the algorithm design does not need to account for the fact that the event only propagates along the edges in the network, and the network structure does not matter. This simplifies the problem to one where we are simply interested in detecting when an arbitrary subset of η nodes has been affected by the event.

For fully connected networks, we solve the QCD problem by reformulating it as a dynamic composite hypothesis testing problem, where we distinguish between two hypotheses at each time instant. The null hypothesis corresponds to the case that less than η nodes are affected; and the alternative hypothesis corresponds to the case that at least η nodes are affected. The data generating distributions for the whole network before and after η nodes are affected are both composite, as they are determined by the unknown subset of affected nodes with unknown change times. We take the generalized log-likelihood ratio between the two composite hypotheses as the detection statistic, and compare it to a positive threshold. If it is greater than the threshold, then we stop and raise an alarm; otherwise, we take another sample from each node in the network.

We show that the generalized log-likelihood ratio test is equivalent to one that compares the sum of the smallest $L - \eta + 1$ local Cumulative Sum (CuSum) statistics [37] to the same threshold. The resulting algorithm, which we refer to as Spartan-CuSum (S-CuSum), is computationally efficient with $O(L)$ complexity at each time step. This guarantees that the algorithm can be implemented efficiently in an online fashion for large networks. We further show that the S-CuSum algorithm satisfies the false alarm constraints with a properly chosen threshold for all scenarios with fewer than η affected nodes, and adapts to unknown event propagation dynamics. We then establish the asymptotic optimality of the S-CuSum algorithm up to a first-order approximation as the false alarm rate goes to zero.

For arbitrarily connected networks, the S-CuSum is still applicable. However, it does not account for the fact that the event only propagates along the edges in the network, and thus will trigger more false alarms. A direct generalization of the generalized log-likelihood ratio test involves a complicated statistic that scans over all connected sub-graphs and propagation dynamics at each time step. This is computationally intractable for a large arbitrarily connected network, especially under the online setting. We then construct an algorithm based on a thresholding approach, the breadth-first search (BFS) algorithm and the S-CuSum algorithm, which we refer to as the Network-CuSum (N-CuSum) algorithm. We show that the computational complexity of the N-CuSum algorithm is linear in the network size at each time step, and that it is asymptotically optimal up to a first-order approximation as the false alarm rate goes to zero. Moreover, we show through our numerical results that the N-CuSum algorithm, which accounts for the network structure, has a better performance than the S-CuSum algorithm. Both the S-CuSum algorithm and the N-CuSum algorithm are better than the generalized multi-chart CuSum algorithm [38], which stops when at least η local CuSum statistics cross their individual thresholds, and the network generalized multi-chart CuSum algorithm, which stops when the local CuSum statistics of at least η connected nodes cross their individual thresholds simultaneously.

C. Paper Organization

This paper is organized as follows. In Section II, we formulate the problem mathematically. In Section III, we focus on fully connected networks, present the S-CuSum algorithm, and present a method to choose the threshold to satisfy the false alarm constraints. In Section IV, we demonstrate the asymptotic optimality of the S-CuSum algorithm. In Section V, we study arbitrarily connected networks, present the N-CuSum algorithm and demonstrate its asymptotic optimality. In Section VI, we present numerical results. Finally, in Section VII, we discuss some potential extensions.

II. PROBLEM FORMULATION

Consider a network monitored in real time by a set of L nodes. We use an unweighted, undirected graph $G = (V, E)$ to denote the underlying structure of the network. Here, $L = |V|$. In practice, an edge connecting two nodes may be due to the fact that two nodes communicate with each other, or are geometrically close to each other.

Before an event occurs, node $i \in \{1, 2, \dots, L\}$ receives independent and identically distributed (i.i.d.) samples from distribution f_0 . If an event occurs, and node i is affected by the event at an unknown time v_i , then it starts to receive i.i.d. samples from distribution f_1 , i.e., v_i is the change-point at node i . If $v_i = \infty$, node i will not be affected by the event ever. More specifically, if we denote the observation received by node i at time k by $X_i[k]$, then

$$X_i[k] \sim \begin{cases} f_0, & \text{if } k < v_i, \\ f_1, & \text{if } k \geq v_i. \end{cases} \quad (1)$$

We assume that the event first affects a *connected* subset of nodes, which might be due to the locality of the event, and

then dynamically propagates along the edges in the network (see Fig. 1 for an example). Equivalently, at every time step, the induced sub-graph on all the affected nodes is *connected*.

We consider a centralized setting in which a fusion center obtains the samples of all the nodes without delay. We are interested in sequentially detecting a “significant” event, i.e., one that affects at least $\eta \geq 1$ affected nodes. If an alarm is triggered at a time when fewer than η nodes are affected, it is then considered as a false alarm event.

Let $\mathbf{v} = \{v_1, \dots, v_L\}$, which is unknown in advance. Without loss of generality, we assume that $v_1 \leq v_2 \leq \dots \leq v_L$, with the ordering being unknown to the decision maker in advance. We note that v_i can be equal to v_{i+1} , i.e., one node can affect more than one of its neighbors simultaneously. Then v_η is the first time when at least η nodes are affected by the event. Thus, our problem is to detect the change at v_η as quickly as possible subject to false alarm constraints.

For any \mathbf{v} , denote by $C(\mathbf{v}) = \{i : v_i < \infty\}$ the set of all the indices of the nodes that will eventually be affected by the event. Then

$$|C(\mathbf{v})| = \sum_{1 \leq i \leq L} \mathbb{1}_{\{v_i < \infty\}} \quad (2)$$

is the total number of affected nodes. If $|C(\mathbf{v})| = L$, then all the nodes will be affected by the event eventually.

We use $\mathbb{P}_{\mathbf{v}}$ to denote the probability measure of the samples with the set of change-points being \mathbf{v} , and let $\mathbb{E}_{\mathbf{v}}$ denote the corresponding expectation. For a given \mathbf{v} , if $|C(\mathbf{v})| < \eta$, i.e., $v_\eta = \infty$, then there are fewer than η nodes that will be affected under $\mathbb{P}_{\mathbf{v}}$. In this case, if an alarm is triggered, it is a false alarm. For any stopping time τ , to measure how frequently false alarms occur, we define the worst-case average run length (WARL) to false alarm as follows:

$$\text{WARL}(\tau) = \inf_{\mathbf{v}: |C(\mathbf{v})| < \eta} \mathbb{E}_{\mathbf{v}}[\tau]. \quad (3)$$

Then with a larger WARL, we have fewer false alarms.

Let $d_i = v_{i+1} - v_i$ denote the time it takes for the event to propagate from node i to node $i + 1$, for $1 \leq i \leq L - 1$. If $d_i = 0$, then node i and node $i + 1$ are affected simultaneously. Denote $\mathbf{D} := \{d_\eta, d_{\eta+1}, \dots, d_{L-1}\}$. For a fixed \mathbf{D} , to measure how quickly we can detect when at least η nodes are affected, we define the worst-case average detection delay (WADD) using a criterion based on Pollak’s criterion [39] as follows:

$$J_{\mathbf{D}}(\tau) = \sup_{v_1 \leq \dots \leq v_\eta < \infty} \mathbb{E}_{\mathbf{v}}[\tau - v_\eta | \tau \geq v_\eta]. \quad (4)$$

We note that the supremum in (4) is only taken over $v_1 \leq \dots \leq v_\eta < \infty$. Therefore, $J_{\mathbf{D}}[\tau]$ is a function of \mathbf{D} .

We denote by \mathcal{F}_k the σ -algebra generated by the observations of all the nodes up to time k , for $k = 1, 2, \dots$. We wish to find a $\{\mathcal{F}_k\}_{k \in \mathbb{N}}$ -stopping time that achieves “small” detection delay, while controlling the false alarm rate. More specifically, for any \mathbf{D} , the goal is to minimize $J_{\mathbf{D}}[\tau]$ subject to a constraint on the WARL:

$$\inf_{\tau: \text{WARL}(\tau) \geq \gamma} J_{\mathbf{D}}(\tau). \quad (5)$$

To describe the objective in words, we want to find stopping rules so that for all possible scenarios with fewer than η affected nodes, the average run length to false alarm is at least γ . At the same time, among those stopping rules that satisfy the false alarm requirement, we want to find the one that minimizes the WADD for all propagation dynamics after η nodes are affected. There is no guarantee that the optimization problem in (5) has a solution, since we require the same stopping rule to simultaneously minimize the WADD for all propagation dynamics after η nodes are affected. What we will show in the following sections is that such a “uniformly” optimum solution can be found up to a first-order approximation in an appropriately defined asymptotic setting.

Notation: We denote the samples across all the nodes at time k by $X[k] = \{X_1[k], \dots, X_L[k]\}$, and the samples across all the nodes from time k_1 to k_2 by $X[k_1, k_2] = \{X[k_1], \dots, X[k_2]\}$. We further define

$$Z_i[k_1, k_2] = \sum_{k=k_1}^{k_2} \log \frac{f_1(X_i[k])}{f_0(X_i[k])}, \quad (6)$$

which is the log-likelihood ratio for the samples at node i from time k_1 to k_2 . We use the following conventions: $\sum_{j=k_1}^{k_2} A_j = 0$ and $\prod_{j=k_1}^{k_2} A_j = 1$ if $k_1 > k_2$. We use X^+ to denote the positive part of X , i.e., $X^+ = \max\{X, 0\}$. We denote the Kullback-Leibler (KL) divergence between f_1 and f_0 as

$$I = \int f_1(x) \log \frac{f_1(x)}{f_0(x)} dx, \quad (7)$$

which is assumed to be positive and finite. We denote $x = o(1)$, as $c \rightarrow c_0$, if $\forall \epsilon > 0, \exists \delta > 0$, s.t., $|x| \leq \epsilon$ if $|c - c_0| < \delta$. We denote $g(c) \sim h(c)$, as $c \rightarrow c_0$, if $\lim_{c \rightarrow c_0} \frac{h(c)}{g(c)} = 1$.

III. FULLY CONNECTED NETWORKS

In this section, we study fully connected networks, for which G is a complete graph. In this case, the event can propagate from any node to any other node, and the induced sub-graph on any subset of nodes is connected. We present the Spartan-CuSum (S-CuSum) algorithm, and show that it can be implemented efficiently with complexity that is linear in L at each time step. We then establish a lower bound on the WARL for the S-CuSum algorithm, and show how to choose the parameter to satisfy the false alarm constraint.

A. The S-CuSum Algorithm

We reformulate the quickest detection problem in Section II when G is a complete graph as a dynamic composite hypothesis testing problem, i.e., to distinguish the following two hypotheses at each time k :

$$\mathcal{H}_0[k] : \sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} < \eta, \quad (8)$$

$$\mathcal{H}_1[k] : \sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} \geq \eta. \quad (9)$$

Both the null and alternative hypotheses are composite, since the data generating distribution depends on unknown \mathbf{v} under each hypothesis. This hypothesis testing procedure stops once a decision in favor of the alternative hypothesis is reached; otherwise, a new sample is taken from each node in the network.

To distinguish between the two hypotheses, we consider the log-likelihood ratio between them. Since \mathbf{v} under each hypothesis is unknown, we take a maximum likelihood approach with respect to the unknown \mathbf{v} , and construct the following generalized log-likelihood ratio statistic:

$$W[k] = \log \left(\frac{\max_{\mathbf{v}: \sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} \geq \eta} \mathbb{P}_{\mathbf{v}}(X[1, k])}{\max_{\mathbf{v}: \sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} < \eta} \mathbb{P}_{\mathbf{v}}(X[1, k])} \right). \quad (10)$$

The max in the numerator in (10) is taken over all \mathbf{v} such that $\sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} \geq \eta$, where at time k , there are no fewer than η affected nodes. Likewise, the max in the denominator in (10) is taken over all \mathbf{v} such that $\sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} < \eta$, where at time k , there are fewer than η affected nodes.

Since the network is fully connected, the induced sub-graph on any subset of nodes is connected. Therefore, the max in both the numerator and denominator in (10) is taken without explicitly enforcing connectivity and propagation dynamics.

The corresponding stopping time is then given by comparing $W[k]$ against a pre-determined positive threshold:

$$\tilde{\tau}(b) = \inf\{k \geq 1 : W[k] > b\}, \quad (11)$$

where b will be selected according to the false alarm constraint.

B. A Simpler but Equivalent Form

In this subsection, we develop an equivalent but much simpler form of (11), which can be computed with complexity $O(L)$ at each time step.

Let \mathbb{P}_{∞} denote the probability measure with $v_i = \infty$, $\forall 1 \leq i \leq L$. Then, it can be easily shown that

$$W[k] = \max_{\mathbf{v}: \sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} \geq \eta} \log \left(\frac{\mathbb{P}_{\mathbf{v}}(X[1, k])}{\mathbb{P}_{\infty}(X[1, k])} \right) - \max_{\mathbf{v}: \sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} < \eta} \log \left(\frac{\mathbb{P}_{\mathbf{v}}(X[1, k])}{\mathbb{P}_{\infty}(X[1, k])} \right). \quad (12)$$

Due to the fact that

$$\begin{aligned} & \log \left(\frac{\mathbb{P}_{\mathbf{v}}(X[1, k])}{\mathbb{P}_{\infty}(X[1, k])} \right) \\ &= \log \left(\prod_{i=1}^L \frac{\prod_{j=1}^{\min\{v_i-1, k\}} f_0(X_i[j]) \prod_{j=v_i}^k f_1(X_i[j])}{\prod_{j=1}^k f_0(X_i[j])} \right) \\ &= \sum_{i=1}^L \sum_{j=v_i}^k \log \frac{f_1(X_i[j])}{f_0(X_i[j])}, \end{aligned} \quad (13)$$

the first term in (12) is equivalent to

$$\max_{\mathbf{v}: \sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} \geq \eta} \sum_{i=1}^L \sum_{j=v_i}^k \log \frac{f_1(X_i[j])}{f_0(X_i[j])}. \quad (14)$$

Similarly, the second term in (12) is equivalent to

$$\max_{\mathbf{v}: \sum_{i=1}^L \mathbb{1}_{\{v_i \leq k\}} < \eta} \sum_{i=1}^L \sum_{j=v_i}^k \log \frac{f_1(X_i[j])}{f_0(X_i[j])}. \quad (15)$$

If we denote the individual CuSum statistic [37] at node i (testing a change from f_0 to f_1) at time k as

$$W_i[k] = \max_{1 \leq v_i \leq k} \sum_{j=v_i}^k \log \frac{f_1(X_i[j])}{f_0(X_i[j])}, \quad (16)$$

and define a permutation $\mu(\cdot)$ such that

$$W_{\mu(1)}[k] \geq W_{\mu(2)}[k] \geq \dots \geq W_{\mu(L)}[k], \quad (17)$$

then, $\tilde{\tau}(b)$ is equivalent to

$$\tilde{\tau}(b) = \inf \left\{ k \geq 1 : \sum_{i=\eta}^L (W_{\mu(i)}[k])^+ \geq b \right\}, \quad (18)$$

which we refer to as the S-CuSum algorithm. Such an equivalence can be established as follows.

- 1) If $W_{\mu(\eta)}[k] \geq 0$, then (14) is equal to $\sum_{i=1}^L (W_{\mu(i)}[k])^+$, and (15) is equal to $\sum_{i=1}^{\eta-1} W_{\mu(i)}[k]$. It then follows that $W[k] = \sum_{i=\eta}^L (W_{\mu(i)}[k])^+$.
- 2) If $W_{\mu(\eta)}[k] < 0$, then (14) is equal to $\sum_{i=1}^{\eta} W_{\mu(i)}[k]$, and (15) is equal to $\sum_{i=1}^{\eta-1} (W_{\mu(i)}[k])^+$. In this case, $W[k]$ is non-positive, and $\sum_{i=\eta}^L (W_{\mu(i)}[k])^+ = 0$. Since b is positive, the test in (11) is equivalent to comparing $\sum_{i=\eta}^L (W_{\mu(i)}[k])^+$ to b .

The test in (18) can be implemented efficiently. First of all, for each node i , $W_i[k]$ can be updated recursively:

$$W_i[k] = (W_i[k-1])^+ + \log \frac{f_1(X_i[k])}{f_0(X_i[k])}. \quad (19)$$

Second, we do not need to sort all $W_i[k]$ at each time k . We only need to find the smallest $L - \eta + 1$ numbers from L numbers, which can be solved with $O(L)$ computational cost using the algorithm in [40] instead of $O(L \log L)$. Thus, the total computational cost at each time k is $O(L)$. Here, we note that, at time k , each node i may choose to send $X_i[k]$, $\log \frac{f_1(X_i[k])}{f_0(X_i[k])}$, or $W_i[k]$ to the fusion center.

C. Lower Bound on the WARL

The following theorem provides a lower bound on the WARL for the S-CuSum algorithm.

Theorem 1. *The WARL for the S-CuSum algorithm in (18) is lower bounded as follows:*

$$\text{WARL}(\hat{\tau}(b)) \geq \frac{1}{\text{poly}(b)} e^b, \quad (20)$$

where $\text{poly}(b)$ denotes a polynomial of b .

Proof: To show the lower bound on $\text{WARL}(\hat{\tau}(b))$, it suffices to show that for any \mathbf{v} with $v_i = \infty, \forall \eta \leq i \leq L$,

$$\mathbb{E}_{\mathbf{v}}[\hat{\tau}(b)] \geq \frac{1}{\text{poly}(b)} e^b. \quad (21)$$

For any $t \in \mathbb{N}$ and $b > 0$, it follows that

$$\begin{aligned} & \mathbb{P}_{\mathbf{v}}(\hat{\tau}(b) \leq t) \\ &= \mathbb{P}_{\mathbf{v}} \left(\max_{1 \leq k \leq t} \sum_{i=\eta}^L (W_{\mu(i)}[k])^+ > b \right) \\ &\leq \sum_{k=1}^t \mathbb{P}_{\mathbf{v}} \left(\sum_{i=\eta}^L (W_{\mu(i)}[k])^+ > b \right). \end{aligned} \quad (22)$$

Since $\sum_{i=\eta}^L (W_{\mu(i)}[k])^+ \leq \sum_{i=\eta}^L (W_i[k])^+$, we have

$$\begin{aligned} & \mathbb{P}_{\mathbf{v}} \left(\sum_{i=\eta}^L (W_{\mu(i)}[k])^+ > b \right) \\ &\leq \mathbb{P}_{\mathbf{v}} \left(\sum_{i=\eta}^L (W_i[k])^+ > b \right). \end{aligned} \quad (23)$$

By [14, Lemma B1], it then follows that

$$\mathbb{P}_{\mathbf{v}} \left(\sum_{i=\eta}^L (W_i[k])^+ > b \right) \leq \text{poly}(b) e^{-b}, \quad (24)$$

where $\text{poly}(b)$ has an order of $L - \eta$. Therefore, $\mathbb{P}_{\mathbf{v}}(\hat{\tau}(b) \leq t) \leq t \cdot \text{poly}(b) e^{-b}$, which implies that

$$\begin{aligned} \mathbb{E}_{\mathbf{v}}[\hat{\tau}(b)] &= \sum_{t=0}^{\infty} \mathbb{P}_{\mathbf{v}}(\hat{\tau}(b) \geq t) \\ &\geq \sum_{t=0}^{\infty} (1 - t \cdot \text{poly}(b) e^{-b})^+ \\ &= \sum_{t=0}^{e^b/\text{poly}(b)} (1 - t \cdot \text{poly}(b) e^{-b}) \\ &= \frac{1}{\text{poly}(b)} e^b. \end{aligned} \quad (25)$$

□

Corollary 1. *To guarantee $\text{WARL}(\hat{\tau}(b)) \geq \gamma$, it suffices to choose b such that*

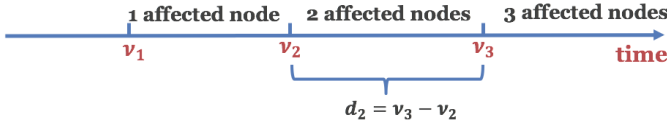
$$\frac{1}{\text{poly}(b)} e^b = \gamma, \quad (26)$$

and $b \sim \log \gamma$, as $\gamma \rightarrow \infty$.

Proof: The result follows from Theorem 1. □

IV. ASYMPTOTIC ANALYSIS

In this section we study the asymptotic performance of the proposed S-CuSum algorithm in (18) and demonstrate its asymptotic optimality. For our asymptotic analysis to be non-trivial, we let both the prescribed lower bound on the WARL, γ , and $d_\eta, d_{\eta+1}, \dots, d_{L-1}$, go to infinity. Indeed, if the event propagates fast, i.e., $d_\eta, d_{\eta+1}, \dots, d_{L-1}$ are small compared to $\log \gamma$, then the event propagation dynamics can be ignored, and only considering the number of nodes that will be affected eventually is sufficient for approximating the performance well. This case corresponds to the asymptotic scenario where

Fig. 2. Example with $L = 3$.

$d_\eta, d_{\eta+1}, \dots, d_{L-1}$ have a smaller order than $\log \gamma$. However, if the event propagates slowly, and $d_\eta, d_{\eta+1}, \dots, d_{L-1}$ are large and cannot be ignored compared to $\log \gamma$, then the event can be detected even before affecting all nodes that will eventually be affected. In this case, we would like to characterize how the fundamental performance limit depends on the underlying event propagation dynamics instead of only considering the number of nodes that will be eventually affected. Therefore, in order to perform a general and relevant asymptotic analysis, we let $d_\eta, d_{\eta+1}, \dots, d_{L-1}$ go to infinity with γ . Without loss of generality, suppose that

$$d_{\eta+i-1} \sim c_i \frac{\log \gamma}{iI}, \quad (27)$$

as $\gamma \rightarrow \infty$, where $c_i \in [0, \infty]$ for $i = 1, \dots, L - \eta$ are unknown. We further assume that $d_L = \infty$, $c_{L-\eta+1} = \infty$. Here if $d_{\eta+i-1}$ has an order less than $\log \gamma$, then $c_i = 0$, and if $d_{\eta+i-1}$ has an order greater than $\log \gamma$, then $c_i = \infty$. This asymptotic setup covers the two cases mentioned above corresponding to fast and slow event propagation. Such an asymptotic setting is only used for the convenience of analysis, and as an approximation of the performance when $d_\eta, d_{\eta+1}, \dots, d_{L-1}$ and γ are large. Applying the algorithm in practice does not rely on this assumption. We further note that $d_\eta, d_{\eta+1}, \dots, d_{L-1}$ do not depend on γ , but as will be shown later, the ratio between $d_\eta, d_{\eta+1}, \dots, d_{L-1}$ and $\log \gamma$ will affect the fundamental performance limits¹.

In the following, we first present an example with $L = 3$ and $\eta = 2$ to understand the results, and then move on to the general results.

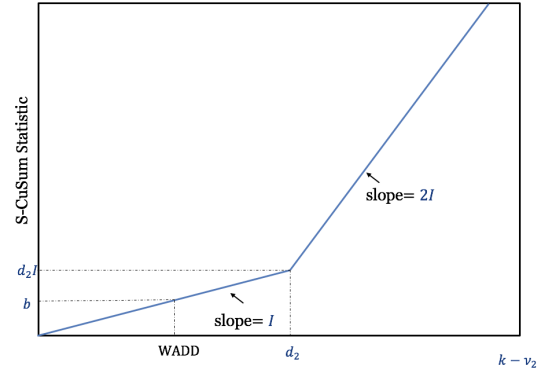
A. Example: $L = 3$ and $\eta = 2$

Consider a fully connected network with three nodes, i.e., $L = 3$. Our goal is to detect when at least two nodes are affected, i.e., $\eta = 2$. Then the S-CuSum algorithm is equivalent to comparing the sum of the smallest two individual CuSum statistics to a threshold b :

$$\hat{\tau}(b) = \inf \left\{ k \geq 1 : \min_{1 \leq i < j \leq 3} (W_i[k])^+ + (W_j[k])^+ \geq b \right\}. \quad (28)$$

For simplicity, we use the notion of phase i to denote the phase in which there are i affected nodes. Then after an event occurs, the network first changes from phase 0 to phase 1, then to phase 2, and eventually stabilizes in phase 3 (see Fig. 2).

¹In this paper, we focus on first order asymptotics, for which considering the scaling between $d_\eta, d_{\eta+1}, \dots, d_{L-1}$ and $\log \gamma$ is sufficient. For a more accurate characterization (e.g., second-order and/or third-order analyses), a more accurate asymptotic setup needs to be considered, which is a topic for future research.

Fig. 3. Scenario 1: d_2 is large.

For this example, the (first-order) asymptotic optimality of the S-CuSum algorithm is characterized in the following proposition.

Proposition 1. *Let the threshold $b \sim \log \gamma$ so that $\text{WARL}(\hat{\tau}(b)) \geq \gamma$. Assume that d_2 and γ go to infinity as in (27), then the S-CuSum is asymptotically optimal:*

$$J_D(\hat{\tau}(b)) \sim \inf_{\tau: \text{WARL}(\tau) \geq \gamma} J_D(\tau) \sim \begin{cases} \frac{\log \gamma}{I}, & \text{if } c_1 > 1, \\ \log \gamma \left(\frac{c_1}{I} + \frac{1 - c_1}{2I} \right), & \text{if } c_1 \leq 1. \end{cases} \quad (29)$$

Proof: The proof for this special case is omitted. See detailed proof for the general case in Theorem 4. \square

The optimal performance shows a dichotomy depending on whether $c_1 > 1$ or not. In the following, we will first provide a heuristic explanation for the dichotomy, and then provide the results for the general case together with rigorous proofs.

Roughly speaking, if we consider the CuSum statistic at node i , before the change-point v_i , it is small and close to 0, and after the change-point, it grows with a positive slope of I . Therefore, by (28), the S-CuSum is close to 0 in phases 0 and 1, and grows with slope I in phase 2 and with slope $2I$ in phase 3.

If $d_2 I > b$, i.e., $c_1 > 1$, then the S-CuSum statistic crosses the threshold b within phase 2 (see Fig. 3). The detection delay for this case is b/I . If $d_2 I \leq b$, i.e., $c_1 \leq 1$, then the S-CuSum statistic is not large enough to cross the threshold b within phase 2, and it needs more samples from phase 3 (see Fig. 4). The detection delay is then equal to the sum of the duration of phase 2 and the number of samples needed from phase 3:

$$d_2 + \frac{b - d_2 I}{2I} \sim b \left(\frac{c_1}{I} + \frac{1 - c_1}{2I} \right). \quad (30)$$

Depending on whether or not phase 2 is long enough, the performance of the S-CuSum algorithm shows a dichotomy.

B. Asymptotic Universal Lower Bound on the WADD

In this subsection, we study the universal lower bound on the WADD for any stopping rule with the WARL no smaller

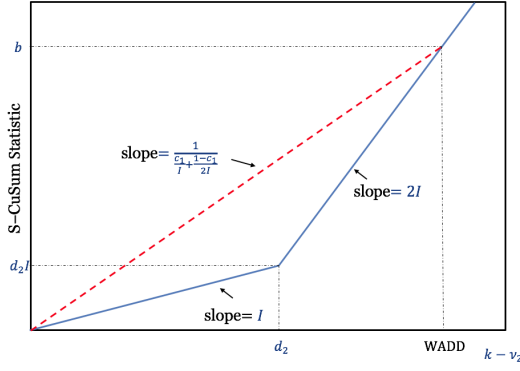


Fig. 4. Scenario 2: d_2 is small.

than γ . We denote

$$h = \inf\{1 \leq j \leq L - \eta + 1 : \sum_{i=1}^j c_i \geq 1\}. \quad (31)$$

Theorem 2. Suppose (27) holds. Then as $\gamma \rightarrow \infty$,

$$\begin{aligned} & \inf_{\tau: \text{WARL}(\tau) \geq \gamma} J_D(\tau) \\ & \geq \log \gamma \left(\sum_{i=1}^{h-1} \frac{c_i}{iI} + \frac{1 - \sum_{i=1}^{h-1} c_i}{hI} \right) (1 - o(1)). \end{aligned} \quad (32)$$

Proof: See Appendix A. \square

Theorem 2 suggests that to meet the asymptotic universal lower bound, an algorithm should be adaptive to the unknown $d_\eta, d_{\eta+1}, \dots, d_{L-1}$. An intuitive understanding of h is that the algorithm shall stop within phase $h + \eta$, when there are $h + \eta$ affected nodes.

The proof is based on a change-of-measure argument and a Law of Large Numbers argument for the log-likelihood ratio statistics, similar to those in [24]. However, a major difference in the change-of-measure argument compared to [24] is that the “pre-change” mode is composite, i.e., there are multiple possible scenarios with fewer than η affected nodes. Furthermore, the post-change statistic is more complicated, since the propagation dynamics is unknown, and the number of affected nodes is changing with time.

C. Asymptotic Upper Bound on the WADD

Recall that we can choose $b \sim \log \gamma$ such that the false alarm constraint is satisfied. Then, by (27), it follows that

$$d_{\eta+i-1} \sim c_i \frac{b}{iI}, \quad (33)$$

as $\gamma \rightarrow \infty$, where $c_i \in [0, \infty]$, for every $i = 1, \dots, L - \eta$.

An asymptotic upper bound on the WADD for the S-CuSum algorithm in (18) is characterized in the following theorem.

Theorem 3. Suppose (33) holds. Then as $b \rightarrow \infty$,

$$J_D(\hat{\tau}(b)) \leq b \left(\sum_{i=1}^{h-1} \frac{c_i}{iI} + \frac{1 - \sum_{i=1}^{h-1} c_i}{hI} \right) (1 + o(1)). \quad (34)$$

Proof: See Appendix A. \square

From Theorem 3, it is clear that although the S-CuSum algorithm does not exploit the knowledge of

$d_\eta, d_{\eta+1}, \dots, d_{L-1}$, the performance is still adaptive to the unknown $d_\eta, d_{\eta+1}, \dots, d_{L-1}$. This is consistent with the insights from the asymptotic universal lower bound in Theorem 2.

The proof of the asymptotic upper bound on WADD is based on partitioning the samples into independent blocks and applying the Law of Large Numbers for the log-likelihood ratio statistics, as in [24, Theorem 4]. The major difficulty here is due to the more complicated test statistic, in which the number of affected nodes changes with time.

D. Asymptotic Optimality of S-CuSum

We are now ready to establish the asymptotic optimality of the S-CuSum algorithm, which is presented in the following theorem.

Theorem 4 (S-CuSum, Asymptotic Optimality). Let the threshold $b \sim \log \gamma$ so that $\text{WARL}(\hat{\tau}(b)) \geq \gamma$. Assume that $d_\eta, d_{\eta+1}, \dots, d_{L-1}$ and γ go to infinity as in (27), then the S-CuSum algorithm is asymptotically optimal:

$$\begin{aligned} J_D(\hat{\tau}(b)) & \sim \inf_{\tau: \text{WARL}(\tau) \geq \gamma} J_D(\tau) \\ & \sim \log \gamma \left(\sum_{i=1}^{h-1} \frac{c_i}{iI} + \frac{1 - \sum_{i=1}^{h-1} c_i}{hI} \right). \end{aligned} \quad (35)$$

Proof: This result follows from Theorems 1, 2 and 3. \square

V. ARBITRARILY CONNECTED NETWORKS

In this section, we extend results of the previous section to arbitrarily connected networks.

Since we assume that the event propagates along edges in the network, the induced sub-graph on the affected nodes is then connected at each time step. Here, for an arbitrarily connected network, the induced sub-graph on some subset of nodes may not be connected. Therefore, although the S-CuSum algorithm still applies, and it is asymptotically optimal (up to a first-order approximation), it will have more false alarms due to the fact that it may raise alarms when it detects η nodes that are not connected (as we will show numerically in Section VI).

To exploit knowledge of the network structure, one can directly adapt the generalized log-likelihood ratio test in (10). However, this will involve a scan statistic over all possible propagation dynamics along the edges in the network, which leads to a combinatorial problem over a large search space. Therefore, it is computationally infeasible, especially for large networks.

In the following, we present the Network-CuSum (N-CuSum) algorithm, which not only employs knowledge of the network structure, but does so in a computationally efficient way. We then establish the first-order asymptotic optimality of the N-CuSum algorithm. Also, as will be shown in the numerical results in Section VI, for an arbitrarily connected network, the N-CuSum algorithm performs much better than the S-CuSum.

A. The N-CuSum Algorithm

At each time step k , we update the local CuSum statistics $W_i[k]$, $\forall 1 \leq i \leq L$. We then compare each local CuSum statistic to a threshold $\log b$, and delete this node if its CuSum statistic is less than $\log b$. The resulting graph is then denoted by $G'[k]$. For this step, the computational complexity is $O(L)$. These deleted nodes are highly likely to be not affected by the event.

We run the BFS algorithm on $G'[k]$ to recover all connected components of $G'[k]$: $C_1[k], C_2[k], \dots$. The computational complexity for this step is at most $O(L + |E|)$. We then run the S-CuSum algorithm on each connected component, and use $\text{S-CuSum}_i[k]$ to denote the test statistic value of the S-CuSum algorithm on $C_i[k]$:

$$\text{S-CuSum}_i[k] = \min_{\substack{C' \subseteq C_i[k]: \\ |C'| = |C_i[k]| - \eta + 1}} \sum_{i \in C'} (W_i[k])^+. \quad (36)$$

If any of these statistics crosses the threshold b , we stop and raise an alarm. For this step, the computational complexity is less than $O(L)$. Therefore, the overall computational complexity at each time step is $O(L + |E|)$, which means that this algorithm scales well as the network size grows, assuming that the network is not dense. The N-CuSum algorithm is described in detail in Algorithm 1.

B. Performance Analysis of the N-CuSum Algorithm

We next provide theoretical analysis for the N-CuSum algorithm. The following theorem provides a lower bound on the WARL for the N-CuSum algorithm.

Theorem 5. *The WARL for the N-CuSum algorithm is lower bounded as follows:*

$$\text{WARL}(\bar{\tau}(b)) \geq \frac{1}{\text{poly}(b)} e^b. \quad (37)$$

Proof: It can be shown that $\text{S-CuSum}_i[k]$ is less than the S-CuSum statistic applied on the whole network, for any i and k . Therefore,

$$\text{WARL}(\bar{\tau}(b)) \geq \text{WARL}(\hat{\tau}(b)). \quad (38)$$

Together with Theorem 1, this completes the proof. \square

To guarantee $\text{WARL}(\bar{\tau}(b)) \geq \gamma$, it suffices to choose b such that

$$\frac{1}{\text{poly}(b)} e^b = \gamma, \quad (39)$$

and $b \sim \log \gamma$.

For the asymptotic analysis, we choose the same asymptotic setting as in (27) in Section IV. By choosing $b \sim \log \gamma$, we also have (33). We then have the asymptotic upper bound on the WADD for the N-CuSum algorithm as characterized in the following theorem.

Theorem 6. *Suppose (33) holds. Then as $b \rightarrow \infty$,*

$$J_D(\bar{\tau}(b)) \leq b \left(\sum_{i=1}^{h-1} \frac{c_i}{iI} + \frac{1 - \sum_{i=1}^{h-1} c_i}{hI} \right) (1 + o(1)). \quad (40)$$

Algorithm 1 N-CuSum

Input:

G : graph
 η : size of sub-graph of interest
 f_0, f_1 : distributions before and after change
 b : threshold

Output:

$\bar{\tau}$: stopping time

Initialization:

$W_i[0] \leftarrow 0$, for $1 \leq i \leq L$
 $k = 0$

Method:

while 1 do

$k \leftarrow k + 1$

 Observe $X_i[k]$, for $1 \leq i \leq L$

$W_i[k] \leftarrow (W_i[k-1])^+ + \log \frac{f_1(X_i[k])}{f_0(X_i[k])}$

$G'[k] \leftarrow G$

for $i = 1$ to L **do**

if $W_i[k] \leq \log b$ **then**

 Delete node i and all edges connected to node i in $G'[k]$

end if

end for

 All connected components of $G'[k]$: $C_1[k], C_2[k], \dots \leftarrow$
 run BFS on $G'[k]$

for $i=1,2,\dots$ **do**

$\text{S-CuSum}_i[k] \leftarrow$ Run S-CuSum on $C_i[k]$

if $\text{S-CuSum}_i[k] \geq b$ **then**

$\bar{\tau} \leftarrow k$

 Break

end if

end for

end while

Return $\bar{\tau}$

Proof: See Appendix A. \square

The proof of the asymptotic upper bound is similar to that of Theorem 3, but requires a more careful construction. This is due to the dependency between the individual S-CuSum statistics and the partition of the graph $G'[k]$ into connected components. The asymptotic universal lower bound in Theorem 2 also applies to the arbitrarily connected network here. We then establish the asymptotic optimality of N-CuSum in the following theorem.

Theorem 7 (N-CuSum, Asymptotic Optimality). *Let threshold $b \sim \log \gamma$ so that $\text{WARL}(\bar{\tau}(b)) \geq \gamma$. Assume that $d_\eta, d_{\eta+1}, \dots, d_{L-1}$ and γ go to infinity as in (27), then the N-CuSum algorithm is asymptotically optimal:*

$$\begin{aligned} J_D(\bar{\tau}(b)) &\sim \inf_{\tau: \text{WARL}(\tau) \geq \gamma} J_D(\tau) \\ &\sim \log \gamma \left(\sum_{i=1}^{h-1} \frac{c_i}{iI} + \frac{1 - \sum_{i=1}^{h-1} c_i}{hI} \right). \end{aligned} \quad (41)$$

Proof: This result follows from Theorems 2, 5 and 6. \square

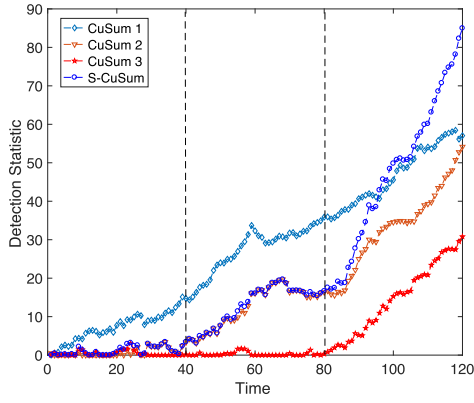


Fig. 5. Sample evolution paths of all individual CuSums and the S-CuSum. CuSum i denotes the individual CuSum statistic at node i , for $i = 1, 2, 3$.

VI. NUMERICAL RESULTS

In this section, we present some numerical results. We start with an example to demonstrate a typical evolution path of the S-CuSum algorithm. We then study a fully connected network with three nodes. We compare the S-CuSum algorithm to a generalization of the multi-chart CuSum algorithm in [38] which stops when at least η local CuSums have crossed their individual thresholds. We further study two networks that are not fully connected: a lattice network with 36 nodes (see Fig. 7) and an arbitrary network with 16 nodes (see Fig. 10). We also consider a network generalized multi-chart CuSum algorithm, which is to wait until η local CuSums have crossed their individual thresholds simultaneously, and those nodes form a connected subgraph. We compare the generalized multi-chart CuSum, the network generalized multi-chart CuSum, the S-CuSum and the N-CuSum algorithms. For all four algorithms, the communication complexity at each time step is L , since we are considering a centralized setting.

A. A Sample Evolution Path of S-CuSum

In Fig. 5, we plot the evolution paths of the S-CuSum statistic and all the individual CuSum statistics. We consider a fully connected network with $L = 3$ and $\eta = 2$. We choose $f_0 = \mathcal{N}(0, 1)$, and $f_1 = \mathcal{N}(1, 1)$. We set $\mathbf{v} = \{1, 40, 80\}$. There are in total three phases, depending on the number of affected nodes.

In phase 1, i.e., $k < 40$, only the statistic of CuSum 1 grows with a positive slope, and all the other statistics are small and close to zero. Then, in phase 2, i.e., $40 \leq k < 80$, the statistics of CuSum 1, CuSum 2 and S-CuSum grow with a positive slope. In this phase, the S-CuSum statistic is almost the same as the CuSum 2 statistic, which is due to the fact that the S-CuSum statistic is the sum of the smallest two individual CuSum statistics, i.e. CuSum 2 and CuSum 3. Since the CuSum 3 statistic is small and close to zero, the S-CuSum statistic is almost the same as the CuSum 2 statistic in this phase. Eventually, in phase 3, i.e., $k \geq 80$, the statistics of CuSum 1, CuSum 2 and CuSum 3 all increase with a positive slope. In this phase, the S-CuSum statistic is the sum of the CuSum 2 and CuSum 3 statistics. Therefore, the slope of the S-CuSum statistic is larger than that in phase 2.

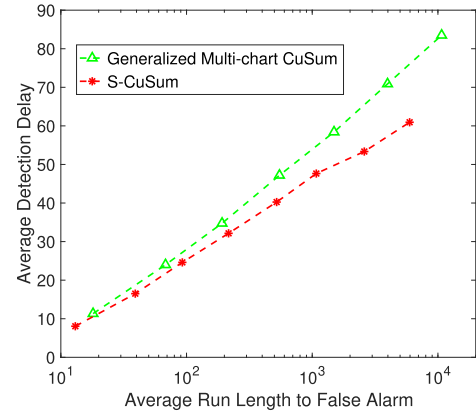


Fig. 6. Comparison between the S-CuSum algorithm and the generalized multi-chart CuSum algorithm for a fully connected network.

In summary, the S-CuSum statistic is small and close to zero with only one affected node, and gradually grows but with different slopes with two and three affected nodes. Thus, the S-CuSum algorithm is adaptive to the unknown propagation dynamics.

B. Comparison between S-CuSum and Generalized Multi-Chart CuSum

We study the performance of the S-CuSum algorithm, validate our theoretical assertions, and compare it with a generalization of the multi-chart CuSum algorithm in [38], which stops when at least η local CuSum algorithms have crossed their individual thresholds simultaneously. Here also we consider a fully connected network with $L = 3$ and $\eta = 2$. We choose $f_0 = \mathcal{N}(0, 1)$, and $f_1 = \mathcal{N}(0.4, 1)$. Here we choose $v_1 = v_2 = 1$, and $d_2 = 40$ to simulate the average detection delay, and choose $v_1 = 1$, $v_2 = v_3 = \infty$ to simulate the average run length to false alarm. The plot is averaged over 1000 runs. We plot WADD versus WARL for the S-CuSum algorithm and the generalized multi-chart CuSum algorithm in Fig. 6.

In Fig. 6, the slope of the curve corresponding to the S-CuSum algorithm gradually changes after WADD = 40, which validates our theoretical results in Proposition 1. Furthermore, the S-CuSum algorithm performs better than the generalized multi-chart CuSum algorithm, especially when WADD $\geq d_2$. This is because when WADD $\geq d_2$, there are three affected nodes. The samples from the third affected node also contain information about whether there are no fewer than η affected nodes (although the local CuSum statistic at the third affected node is not large), and this information is used by the S-CuSum algorithm but not the multichart CuSum algorithm.

C. Comparison on a Lattice Network

We consider a lattice network with 36 nodes (see Fig. 7). We set $\eta = 4$, $f_0 = \mathcal{N}(0, 1)$, and $f_1 = \mathcal{N}(1, 1)$. To simulate the average detection delay, we assume that the event first affects nodes 14, 15, 16, 22 at time 1, then propagates to nodes 9 and 17 at time 10, and no other node is affected by

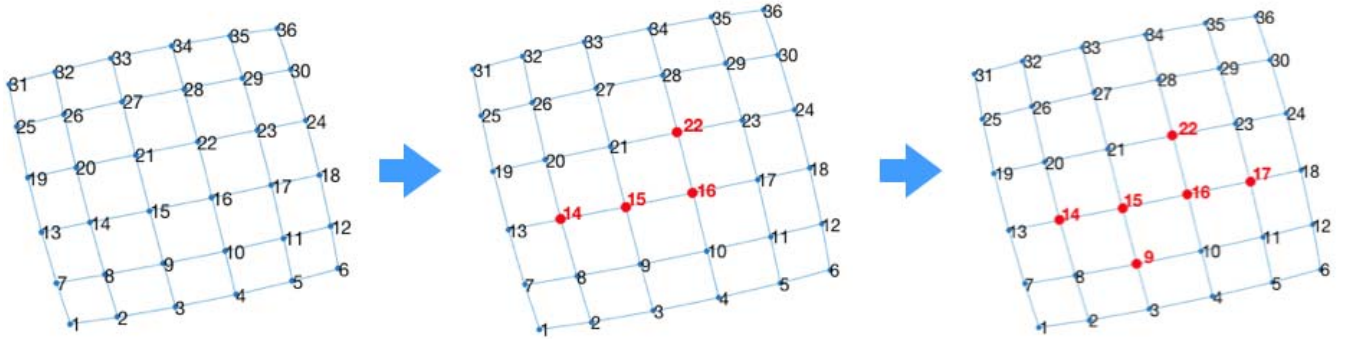


Fig. 7. A dynamic event propagates in a lattice network.

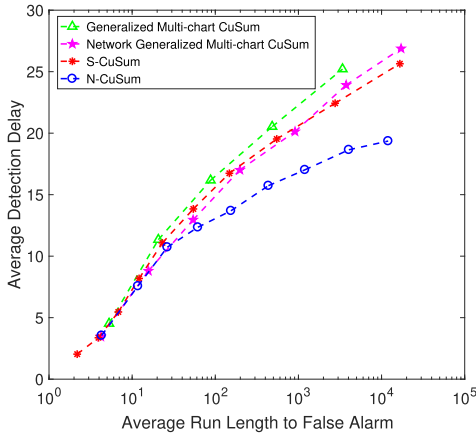
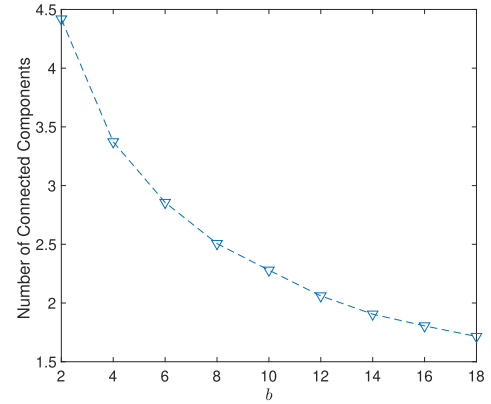


Fig. 8. Comparison among the generalized multi-chart CuSum algorithm, the N-generalized multi-chart CuSum algorithm, the S-CuSum algorithm and the N-CuSum algorithm for a lattice network.

the event. To simulate the average run length to false alarm, we assume that nodes 14, 15, 16 are affected at time 1, and no other node is affected by the event. We repeat the simulation for 1000 times. As we can see from Fig. 8, the N-CuSum algorithm has the best performance among the four algorithms. Compared to the S-CuSum algorithm, the N-CuSum algorithm has significantly reduced the average detection delay, which is due to its effective exploitation of the network structure.

Under the same setting used to simulate the average detection delay as shown in Fig. 8, we plot the average number of connected components in $G'[k]$ when N-CuSum crosses the threshold b , as a function of the threshold b in Fig. 9. We observe that as b increases, the average number of connected components decreases, and its value is between 1.5 and 4.5. This is because for large b , the unaffected nodes are eliminated with high probability, and the resulting graph $G'[k]$ contains mostly the affected nodes.

We also compare the computational complexity of these four algorithms. We consider the case where all nodes are not affected. We ran the algorithms for 10000 steps without stopping ($b = \infty$), and repeated the experiment for 100 times. We ran the experiments on a 2.0GHz Intel core i5 CPU using Matlab. The average time consumption for 10000 steps (in seconds) is given in Table I. We observe that all four

Fig. 9. Number of connected components in $G'[k]$ when N-CuSum crosses the threshold.TABLE I
COMPARISON OF TIME CONSUMPTION

Generalized Multi-chart CuSum	Network Generalized Multi-chart CuSum	S-CuSum	N-CuSum
0.15	0.77	0.12	0.94

algorithms take less than one second. Also, not surprisingly, the algorithms that exploit the network structure consume more computational power.

D. Comparison on an Arbitrary Network

We further consider an arbitrary network with 16 nodes (see Fig. 10). We set $\eta = 4$, $f_0 = \mathcal{N}(0, 1)$ and $f_1 = (1, 1)$. To simulate the average detection delay, we assume that the event first affects nodes 1, 4, 12, 14, and then propagates to nodes 3, 5, 8 at time 10. To simulate the average run length to false alarm, we assume that nodes 5, 7, 15 are affected at time 1, and no other node is affected by the event. We repeat the simulation for 1000 times.

Similar conclusions to the ones for Fig. 8 can be obtained from Fig. 11. In particular, we observe that the N-CuSum algorithm has the best performance among these four algorithms. Compared to other algorithms, the N-CuSum algorithm has a much smaller average detection delay, which is due to its effective exploitation of network structure and event propagation dynamics. Moreover, we can observe a slope

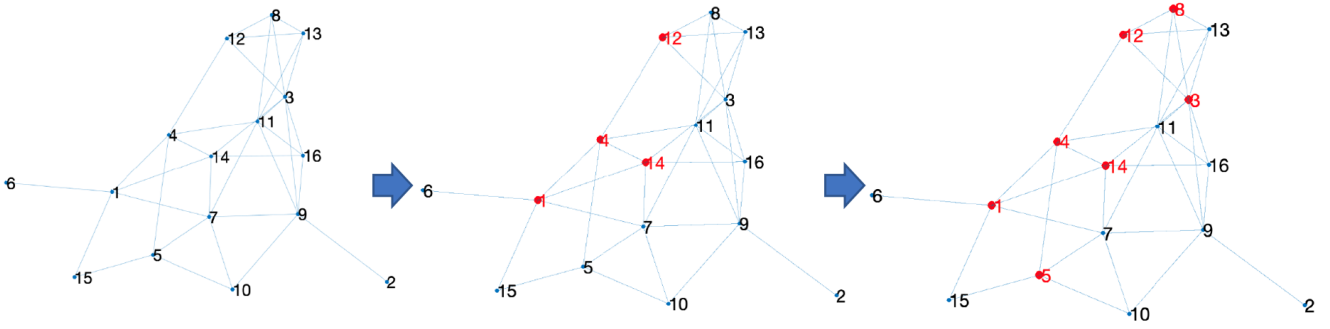


Fig. 10. Dynamic event propagation in an arbitrary network.

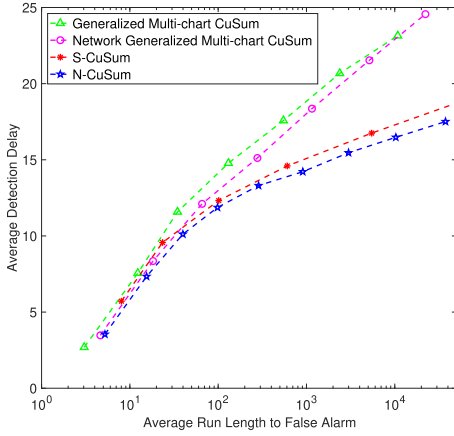


Fig. 11. Comparison of the generalized multi-chart CuSum algorithm, the N-generalized multi-chart CuSum algorithm, the S-CuSum algorithm and the N-CuSum algorithm for a general network.

change in the S-CuSum and N-CuSum algorithms, which verifies the theoretical assertions in Theorems 4 and 7.

VII. DISCUSSION

The results in this paper can be easily generalized to the case in which the distributions of the samples are different across the different nodes (heterogeneous sensors). For example, the generalized S-CuSum algorithm for this case is also constructed by comparing the sum of the smallest $L - \eta + 1$ local CuSums to a threshold b , and hence can be implemented efficiently. The asymptotic optimality of this algorithm can also be established similarly, but the optimal performance takes on a more complicated form.

In our modeling and analysis we have assumed that only one connected subgraph in the network is affected. In an arbitrarily connected network, if two or more events occur simultaneously, two or more connected subgraphs could be affected. Here the goal might be to detect whether there exists *one* event that has affected at least η nodes; two small events are not of interest. Then the S-CuSum algorithm is clearly not going to be asymptotically optimal in this case, because it clubs all the affected nodes together ignoring the network structure. The N-CuSum algorithm is also not going to be asymptotically optimal for the following reason. The N-CuSum algorithm may not be able to correctly delete the unaffected nodes, and this could result in bridges between small affected connected

subgraphs (with less than η affected nodes) to form a big one (with more than η nodes), thus triggering false alarms. Developing an efficient algorithm to handle the case of two or more events occurring simultaneously is an interesting open problem.

In this paper, the observations are assumed to be i.i.d. before and after the change-point at each node. It is clearly of interest to generalize to the case with non i.i.d. observations using tools described in [6]. Furthermore, we assumed that the data generating distributions before and after a node is affected by the event are known. In many practical applications, this assumption may not hold, and it is of further interest to construct algorithms that do not rely on complete knowledge of the distributions and still provide good performance [41], [42]. In this paper, we consider a centralized setting, in which all the samples are available at a fusion center. We also extended our results to the distributed setting where this is no fusion center and the sensors communicate directly with each other [43]. The adversarial setting is also worth exploring, in which some nodes may be comprised by an adversarial party.

APPENDIX

We recall the following useful lemma, which is a slight generalization of the Weak Law of Large Numbers.

Lemma 1. [44, Lemma A.1] Suppose random variables Y_1, Y_2, \dots, Y_k are i.i.d. on $(\Omega, \mathcal{F}, \mathbb{P})$ with $\mathbb{E}[Y_i] = \mu > 0$, and denote $S_k = \sum_{i=1}^k Y_i$, then for any $\epsilon > 0$, as $n \rightarrow \infty$,

$$\mathbb{P} \left(\frac{\max_{1 \leq k \leq n} S_k}{n} - \mu > \epsilon \right) \rightarrow 0. \quad (42)$$

For a given tuple of $\{d_\eta, d_{\eta+1}, \dots, d_{L-1}\}$, it can be shown that

$$\begin{aligned} J_D[\tau] &= \sup_{v_1 \leq \dots \leq v_\eta < \infty} \mathbb{E}_{\mathbf{v}}[\tau - v_\eta | \tau \geq v_\eta] \\ &\geq \sup_{\substack{v_1 = v_2 = \dots = v_{\eta-1} = 1 \\ 1 \leq v_\eta < \infty}} \mathbb{E}_{\mathbf{v}}[\tau - v_\eta | \tau \geq v_\eta]. \end{aligned} \quad (43)$$

It then suffices to lower bound (43) for any stopping rule τ that satisfies the false alarm constraint. In the following of the proof, \mathbf{v} is specified by $v_1 = v_2 = \dots = v_{\eta-1} = 1, v_\eta$, and $d_\eta, d_{\eta+1}, \dots, d_{L-1}$.

For simplicity, for any $\epsilon > 0$, denote

$$\alpha_\gamma = \log \gamma \left(\sum_{i=1}^{h-1} \frac{c_i}{iI} + \frac{1 - \sum_{i=1}^{h-1} c_i}{hI} \right) (1 - \epsilon). \quad (44)$$

By the Markov inequality,

$$\begin{aligned} \mathbb{E}_{\mathbf{v}}[\tau - v_\eta | \tau \geq v_\eta] \\ \geq \mathbb{P}_{\mathbf{v}}(\tau - v_\eta \geq \alpha_\gamma | \tau \geq v_\eta) \alpha_\gamma. \end{aligned} \quad (45)$$

It then suffices to show that

$$\mathbb{P}_{\mathbf{v}}(\tau - v_\eta \geq \alpha_\gamma | \tau \geq v_\eta) \rightarrow 1, \quad (46)$$

as $\gamma \rightarrow \infty$.

We denote $\bar{\mathbf{v}} = \{1, \dots, 1, \infty, \dots, \infty\}$ with the first $\eta - 1$ elements being 1, and all the remaining elements being infinity. Clearly, under $\mathbb{P}_{\bar{\mathbf{v}}}$, there are $\eta - 1$ affected nodes. For any stopping time τ that satisfies the false alarm constraint, we have

$$\mathbb{E}_{\bar{\mathbf{v}}}[\tau] \geq \gamma, \quad (47)$$

which implies that for each $m < \gamma$, there exists some $v \geq 1$, such that

$$\mathbb{P}_{\bar{\mathbf{v}}}(\tau \geq v) > 0 \text{ and } \mathbb{P}_{\bar{\mathbf{v}}}(\tau < v + m | \tau \geq v) \leq \frac{m}{\gamma}. \quad (48)$$

This can be shown by contradiction as in [24, Theorem 1].

By a change of measure argument, it follows that

$$\begin{aligned} \mathbb{P}_{\bar{\mathbf{v}}} (v_\eta \leq \tau < v_\eta + \alpha_\gamma) \\ = \mathbb{E}_{\bar{\mathbf{v}}} (\mathbb{1}_{\{v_\eta \leq \tau < v_\eta + \alpha_\gamma\}}) \\ = \mathbb{E}_{\mathbf{v}} \left(\mathbb{1}_{\{v_\eta \leq \tau < v_\eta + \alpha_\gamma\}} \frac{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, \tau])}{\mathbb{P}_{\mathbf{v}}(X[v_\eta, \tau])} \right) \\ \geq \mathbb{E}_{\mathbf{v}} \left(\mathbb{1}_{\{v_\eta \leq \tau < v_\eta + \alpha_\gamma, \log \frac{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, \tau])}{\mathbb{P}_{\mathbf{v}}(X[v_\eta, \tau])} \geq -a\}} \frac{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, \tau])}{\mathbb{P}_{\mathbf{v}}(X[v_\eta, \tau])} \right) \\ \geq e^{-a} \mathbb{P}_{\mathbf{v}} \left(v_\eta \leq \tau < v_\eta + \alpha_\gamma, \log \frac{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, \tau])}{\mathbb{P}_{\mathbf{v}}(X[v_\eta, \tau])} \geq -a \right) \\ = e^{-a} \mathbb{P}_{\mathbf{v}} \left(v_\eta \leq \tau < v_\eta + \alpha_\gamma, \log \frac{\mathbb{P}_{\mathbf{v}}(X[v_\eta, \tau])}{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, \tau])} \leq a \right) \\ \geq e^{-a} \mathbb{P}_{\mathbf{v}} \left(v_\eta \leq \tau < v_\eta + \alpha_\gamma, \right. \\ \left. \max_{v_\eta \leq j \leq v_\eta + \alpha_\gamma} \log \frac{\mathbb{P}_{\mathbf{v}}(X[v_\eta, j])}{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, j])} \leq a \right), \end{aligned} \quad (49)$$

where a will be specified later.

The event $\{\tau \geq v_\eta\}$ only depends on $X[1, v_\eta - 1]$, which follows the same distribution under $\mathbb{P}_{\mathbf{v}}$ and $\mathbb{P}_{\bar{\mathbf{v}}}$. This implies that

$$\mathbb{P}_{\mathbf{v}}(\tau \geq v_\eta) = \mathbb{P}_{\bar{\mathbf{v}}}(\tau \geq v_\eta). \quad (50)$$

It then follows that

$$\begin{aligned} \mathbb{P}_{\bar{\mathbf{v}}} (v_\eta \leq \tau < v_\eta + \alpha_\gamma | \tau \geq v_\eta) \\ \geq e^{-a} \mathbb{P}_{\mathbf{v}} \left(v_\eta \leq \tau < v_\eta + \alpha_\gamma, \right. \\ \left. \max_{v_\eta \leq j \leq v_\eta + \alpha_\gamma} \log \frac{\mathbb{P}_{\mathbf{v}}(X[v_\eta, j])}{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, j])} \leq a \middle| \tau \geq v_\eta \right). \end{aligned} \quad (51)$$

Due to the fact that for any events A and B , $\mathbb{P}(A \cap B) \geq \mathbb{P}(A) - \mathbb{P}(B^c)$, it follows that

$$\begin{aligned} \mathbb{P}_{\mathbf{v}} \left(v_\eta \leq \tau < v_\eta + \alpha_\gamma, \right. \\ \left. \max_{v_\eta \leq j \leq v_\eta + \alpha_\gamma} \log \frac{\mathbb{P}_{\mathbf{v}}(X[v_\eta, j])}{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, j])} \leq a \middle| \tau \geq v_\eta \right) \\ \geq \mathbb{P}_{\mathbf{v}} \left(v_\eta \leq \tau < v_\eta + \alpha_\gamma \middle| \tau \geq v_\eta \right) \\ - \mathbb{P}_{\mathbf{v}} \left(\max_{v_\eta \leq j \leq v_\eta + \alpha_\gamma} \log \frac{\mathbb{P}_{\mathbf{v}}(X[v_\eta, j])}{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, j])} > a \middle| \tau \geq v_\eta \right) \\ \stackrel{(a)}{=} \mathbb{P}_{\mathbf{v}} \left(v_\eta \leq \tau < v_\eta + \alpha_\gamma \middle| \tau \geq v_\eta \right) \\ - \mathbb{P}_{\mathbf{v}} \left(\max_{v_\eta \leq j \leq v_\eta + \alpha_\gamma} \log \frac{\mathbb{P}_{\mathbf{v}}(X[v_\eta, j])}{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, j])} > a \right), \end{aligned} \quad (52)$$

where (a) is due to the fact that the event $\{\tau \geq v_\eta\}$ only depends on $X[1, v_\eta - 1]$, which is independent from $X[v_\eta, j]$, $\forall v_\eta \leq j \leq v_\eta + \alpha_\gamma$.

Combining (51) and (52), we obtain

$$\begin{aligned} \mathbb{P}_{\mathbf{v}} \left(v_\eta \leq \tau < v_\eta + \alpha_\gamma \middle| \tau \geq v_\eta \right) \\ \leq e^a \mathbb{P}_{\bar{\mathbf{v}}} (v_\eta \leq \tau < v_\eta + \alpha_\gamma | \tau \geq v_\eta) \\ + \mathbb{P}_{\mathbf{v}} \left(\max_{v_\eta \leq j \leq v_\eta + \alpha_\gamma} \log \frac{\mathbb{P}_{\mathbf{v}}(X[v_\eta, j])}{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, j])} > a \right). \end{aligned} \quad (53)$$

By (48), it follows that for $m = \alpha_\gamma$, there exists v_η , such that

$$\mathbb{P}_{\bar{\mathbf{v}}} (v_\eta \leq \tau < v_\eta + \alpha_\gamma | \tau \geq v_\eta) \leq \frac{\alpha_\gamma}{\gamma}. \quad (54)$$

Let $a = (1 - \epsilon^2) \log \gamma$, then

$$\begin{aligned} e^a \mathbb{P}_{\bar{\mathbf{v}}} (v_\eta \leq \tau < v_\eta + \alpha_\gamma | \tau \geq v_\eta) \\ \leq \gamma^{1 - \epsilon^2} \frac{\alpha_\gamma}{\gamma} \\ \rightarrow 0, \text{ as } \gamma \rightarrow \infty. \end{aligned} \quad (55)$$

We then show the second term in (52) also converges to 0 as $\gamma \rightarrow \infty$. It can be shown that

$$\begin{aligned} \mathbb{P}_{\mathbf{v}} \left(\max_{v_\eta \leq j \leq v_\eta + \alpha_\gamma} \log \frac{\mathbb{P}_{\mathbf{v}}(X[v_\eta, j])}{\mathbb{P}_{\bar{\mathbf{v}}}(X[v_\eta, j])} > a \right) \\ = \mathbb{P}_{\mathbf{v}} \left(\max_{v_\eta \leq j \leq v_\eta + \alpha_\gamma} \sum_{i \in \mathcal{C}(\mathbf{v})} \sum_{k=v_i}^j \log \frac{f_1(X_i[k])}{f_0(X_i[k])} > a \right) \\ \rightarrow 0, \text{ as } \gamma \rightarrow \infty, \end{aligned} \quad (56)$$

where the last step follows by applying Lemma 1.

Combining (55) and (56), it follows that

$$\mathbb{P}_{\mathbf{v}} (v_\eta \leq \tau < v_\eta + \alpha_\gamma | \tau \geq v_\eta) \rightarrow 0, \quad (57)$$

as $\gamma \rightarrow \infty$. This concludes the proof.

By the recursive structure of $(W_i[k])^+$, and the fact that it is always non-negative, and is zero when $k = 0$, then for a given \mathbf{D} , the worst-case of the average detection delay is achieved when $v_1 = \dots = v_\eta = 1$.

Now

$$\begin{aligned} & \sum_{i=\eta}^L (W_{\mu(i)}[k])^+ \\ &= \min_{C': |C'|=L-\eta+1} \sum_{i \in C'} (W_i[k])^+ \\ &\stackrel{(a)}{\geq} \min_{C' \subseteq C(\mathbf{v}): |C'|=|C(\mathbf{v})|-\eta+1} \sum_{i \in C'} (W_i[k])^+, \end{aligned} \quad (58)$$

where (a) is due to the fact that $(W_i[k])^+$ is always non-negative, especially for $i \notin C(\mathbf{v})$.

Define the stopping rule $N(b)$:

$$N(b) = \inf \left\{ k : \min_{C' \subseteq C(\mathbf{v}): |C'|=|C(\mathbf{v})|-\eta+1} \sum_{i \in C'} (W_i[k])^+ > b \right\}. \quad (59)$$

It then follows that $\hat{\tau}(b) \leq N(b)$. Therefore, it suffices to establish an upper bound on the $\mathbb{E}_{\mathbf{v}}[N(b)]$.

For simplicity, for any $\epsilon > 0$, we denote

$$\alpha_b = b \left(\sum_{i=1}^{h-1} \frac{c_i}{iI} + \frac{1 - \sum_{i=1}^{h-1} c_i}{hI} \right) (1 + \epsilon). \quad (60)$$

Now

$$\begin{aligned} \mathbb{E}_{\mathbf{v}}[N(b)/\alpha_b] &\leq \sum_{\ell=0}^{\infty} \mathbb{P}_{\mathbf{v}}(N(b)/\alpha_b > \ell) \\ &\leq 1 + \sum_{\ell=1}^{\infty} \mathbb{P}_{\mathbf{v}}(N(b)/\alpha_b > \ell). \end{aligned} \quad (61)$$

We then bound $\mathbb{P}_{\mathbf{v}}(N(b)/\alpha_b > \ell)$ for every $\ell \geq 1$ as in (62), shown at the bottom of this page, where (a) is

by the definition of $W_i[k]$; (b) follows by the independency among the random variables: $X[1, \alpha_b], X[\alpha_b + 1, 2\alpha_b], \dots, X[(\ell - 1)\alpha_b + 1, \ell\alpha_b]$.

We then bound the first term in (62). It follows that

$$\begin{aligned} & \mathbb{P}_{\mathbf{v}} \left(\left(\min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} Z_i[\max\{1, v_i\}, \alpha_b] < b \right) \right) \\ &\leq \sum_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \mathbb{P}_{\mathbf{v}} \left(\sum_{i \in C'} Z_i[\max\{1, v_i\}, \alpha_b] < b \right). \end{aligned} \quad (63)$$

It is clear that for large b ,

$$\alpha_b > \sum_{i=1}^{h-1} d_{\eta+i-1}. \quad (64)$$

Moreover, $\sum_{i \in C'} Z_i[\max\{1, v_i\}, \alpha_b]$ is the summation of the log-likelihood ratios of the samples from f_1 . Therefore, for any $C' \subseteq C(\mathbf{v})$ such that $|C'| = |C(\mathbf{v})| - \eta + 1$, $\sum_{i \in C'} Z_i[\max\{1, v_i\}, \alpha_b]$ is the sum of the log likelihood ratio between f_1 and f_0 of at least

$$d_{\eta} + 2d_{\eta+1} + \dots + (h-1)d_{\eta+h-2} + h(\alpha_b - \sum_{i=1}^{h-1} d_{\eta+i-1}) \quad (65)$$

number of samples generated by f_1 . Then by the Weak Law of Large Numbers, it follows that

$$\frac{\sum_{i \in C'} Z_i[\max\{1, v_i\}, \alpha_b]}{b} \rightarrow \beta, \quad (66)$$

$$\begin{aligned} & \mathbb{P}_{\mathbf{v}}(N(b)/\alpha_b > \ell) \\ &= \mathbb{P}_{\mathbf{v}} \left(\forall 1 \leq k \leq \ell\alpha_b : \min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} (W_i[k])^+ < b \right) \\ &\leq \mathbb{P}_{\mathbf{v}} \left(\left(k = \alpha_b : \min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} (W_i[k])^+ < b \right) \cap \left(k = 2\alpha_b : \min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} (W_i[k])^+ < b \right) \right. \\ &\quad \left. \cap \dots \cap \left(k = \ell\alpha_b : \min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} (W_i[k])^+ < b \right) \right) \\ &\stackrel{(a)}{\leq} \mathbb{P}_{\mathbf{v}} \left(\left(\min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} Z_i[\max\{1, v_i\}, \alpha_b] < b \right) \cap \left(\min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} Z_i[\max\{\alpha_b + 1, v_i\}, 2\alpha_b] < b \right) \right. \\ &\quad \left. \cap \dots \cap \left(\min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} Z_i[\max\{(\ell-1)\alpha_b + 1, v_i\}, \ell\alpha_b] < b \right) \right) \\ &\stackrel{(b)}{=} \mathbb{P}_{\mathbf{v}} \left(\min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} Z_i[\max\{1, v_i\}, \alpha_b] < b \right) \times \mathbb{P}_{\mathbf{v}} \left(\min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} Z_i[\max\{\alpha_b + 1, v_i\}, 2\alpha_b] < b \right) \\ &\quad \times \dots \times \mathbb{P}_{\mathbf{v}} \left(\min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'|=|C(\mathbf{v})|-\eta+1}} \sum_{i \in C'} Z_i[\max\{(\ell-1)\alpha_b + 1, v_i\}, \ell\alpha_b] < b \right). \end{aligned} \quad (62)$$

in probability, where $\beta > 1$. Therefore, as $b \rightarrow \infty$,

$$\mathbb{P}_{\mathbf{v}}\left(\sum_{i \in C'} Z_i[\max\{1, v_i\}, \alpha_b] < b\right) \rightarrow 0. \quad (67)$$

Together with (63), this further implies that

$$\begin{aligned} & \mathbb{P}_{\mathbf{v}}\left(\min_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'| = |C(\mathbf{v})| - \eta + 1}} \sum_{i \in C'} Z_i[\max\{1, v_i\}, \alpha_b] < b\right) \\ & \leq \sum_{\substack{C' \subseteq C(\mathbf{v}): \\ |C'| = |C(\mathbf{v})| - \eta + 1}} \delta' \\ & \triangleq \delta, \end{aligned} \quad (68)$$

where δ' and δ can be arbitrarily small for large b .

Following similar steps, we can also show that each term in (62) is upper bounded by δ for large b . Therefore,

$$\mathbb{P}_{\mathbf{v}}(N(b)/\alpha_b > \ell) \leq \delta^\ell, \quad (69)$$

and

$$\begin{aligned} \mathbb{E}_{\mathbf{v}}[N(b)/\alpha_b] & \leq 1 + \sum_{\ell=1}^{\infty} \delta^\ell \\ & = \frac{1}{1 - \delta}. \end{aligned} \quad (70)$$

This implies that

$$\begin{aligned} \mathbb{E}_{\mathbf{v}}[N(b)] & \leq \frac{\alpha_b}{1 - \delta} \\ & = b \left(\sum_{i=1}^{h-1} \frac{c_i}{iI} + \frac{1 - \sum_{i=1}^{h-1} c_i}{hI} \right) \frac{1 + \epsilon}{1 - \delta}. \end{aligned} \quad (71)$$

Due to the fact that ϵ is chosen arbitrarily and δ can be arbitrarily small for large b , as $b \rightarrow \infty$,

$$\mathbb{E}_{\mathbf{v}}[N(b)] \leq b \left(\sum_{i=1}^{h-1} \frac{c_i}{iI} + \frac{1 - \sum_{i=1}^{h-1} c_i}{hI} \right) (1 + o(1)). \quad (72)$$

This concludes the proof.

By the recursive structure of $(W_i[k])^+$, and the fact that it is always non-negative, and is zero when $k = 0$, then for a given \mathbf{D} , the worst-case of the average detection delay is achieved when $v_1 = \dots = v_\eta = 1$.

We use the same notation of α_b as in (60). It can be shown that

$$\begin{aligned} \mathbb{E}_{\mathbf{v}}[\bar{\tau}(b)/\alpha_b] & \leq \sum_{\ell=0}^{\infty} \mathbb{P}_{\mathbf{v}}(\bar{\tau}(b)/\alpha_b > \ell) \\ & \leq 1 + \sum_{\ell=1}^{\infty} \mathbb{P}_{\mathbf{v}}(\bar{\tau}(b)/\alpha_b > \ell). \end{aligned} \quad (73)$$

Recall that we assume that $v_1 \leq v_2 \leq \dots \leq v_L$, and for large b , we have (64). We then denote $H = \{1, 2, \dots, h\}$ as the set of indices of nodes that have changed their distribution by the time α_b . For simplicity of notation, we use $\text{S-CuSum}_H[k]$ to denote the test statistic value of the S-CuSum on H at time k :

$$\text{S-CuSum}_H[k] = \min_{\substack{C' \subseteq H: \\ |C'| = |H| - \eta + 1}} \sum_{i \in C'} (W_i[k])^+. \quad (74)$$

We next bound $\mathbb{P}_{\mathbf{v}}(\bar{\tau}(b)/\alpha_b > \ell)$ for every $\ell \geq 1$ as follows:

$$\begin{aligned} & \mathbb{P}_{\mathbf{v}}(\bar{\tau}(b)/\alpha_b > \ell) \\ & = \mathbb{P}_{\mathbf{v}}\left(\forall 1 \leq k \leq \ell \alpha_b : \max_i \text{S-CuSum}_i[k] < b\right) \\ & = \mathbb{P}_{\mathbf{v}}\left(\forall 1 \leq k \leq \ell \alpha_b : A[k] \cup B[k]\right) \\ & \leq \mathbb{P}_{\mathbf{v}}\left(\forall k \in \{\alpha_b, 2\alpha_b, \dots, \ell \alpha_b\} : A[k] \cup B[k]\right) \end{aligned} \quad (75)$$

where $A[k]$ and $B[k]$ are two events defined as follows:

$$\begin{aligned} A[k] & = \left\{ \left\{ \max_i \text{S-CuSum}_i[k] < b \right\} \cap \left\{ \exists i : H \subseteq C_i[k] \right\} \right\}, \\ B[k] & = \left\{ \left\{ \max_i \text{S-CuSum}_i[k] < b \right\} \cap \left\{ \forall i : H \not\subseteq C_i[k] \right\} \right\}. \end{aligned} \quad (76)$$

We first analyze $A[k]$. Suppose that for some i^* , $H \subseteq C_{i^*}[k]$, then

$$\text{S-CuSum}_{i^*}[k] \geq \text{S-CuSum}_H[k]. \quad (77)$$

Combining with the fact that

$$\max_i \text{S-CuSum}_i[k] \geq \text{S-CuSum}_{i^*}[k], \quad (78)$$

it then follows that

$$\begin{aligned} \left\{ \max_i \text{S-CuSum}_i[k] < b \right\} & \subseteq \left\{ \text{S-CuSum}_{i^*}[k] < b \right\} \\ & \subseteq \left\{ \text{S-CuSum}_H[k] < b \right\}. \end{aligned} \quad (79)$$

This further implies that

$$\begin{aligned} A[k] & \subseteq \left\{ \left\{ \text{S-CuSum}_H[k] < b \right\} \cap \left\{ \exists i : H \subseteq C_i[k] \right\} \right\} \\ & \subseteq \left\{ \text{S-CuSum}_H[k] < b \right\}. \end{aligned} \quad (80)$$

For any $1 \leq j \leq \ell$, it follows from the definition of CuSum statistic $W_i[k]$ that

$$\begin{aligned} & \text{S-CuSum}_H[j\alpha_b] \\ & = \min_{\substack{C' \subseteq H: \\ |C'| = |H| - \eta + 1}} \sum_{i \in C'} (W_i[j\alpha_b])^+ \\ & \geq \min_{\substack{C' \subseteq H: \\ |C'| = |H| - \eta + 1}} \sum_{i \in C'} Z_i[\max\{v_i, (j-1)\alpha_b + 1\}, j\alpha_b]. \end{aligned} \quad (81)$$

We then define

$$\begin{aligned} & A'[j\alpha_b] \\ & = \left\{ \min_{\substack{C' \subseteq H: \\ |C'| = |H| - \eta + 1}} \sum_{i \in C'} Z_i[\max\{v_i, (j-1)\alpha_b + 1\}, j\alpha_b] < b \right\}. \end{aligned} \quad (82)$$

It is clear that $\forall 1 \leq j \leq \ell$,

$$A[j\alpha_b] \subseteq A'[j\alpha_b], \quad (83)$$

and $A'[j\alpha_b]$ only depends on the samples from $(j-1)\alpha_b + 1$ to $j\alpha_b$.

We then analyze $B[k]$. By the assumption that the event propagates along the edges in the network, the sub-graph

induced on H is connected. If $\forall j \in H, W_j[k] \geq \log b$, then there must exist $C_i[k]$ that contains all nodes in H . Therefore, it follows that

$$\{\forall i : H \not\subseteq C_i[k]\} \subseteq \{\exists j \in H : W_j[k] < \log b\}. \quad (84)$$

Therefore, we have

$$\begin{aligned} B[k] &\subseteq \left\{ \left\{ \max_i \text{S-CuSum}_i[k] < b \right\} \cap \{\exists i \in H : W_i[k] < \log b\} \right\} \\ &\subseteq \{\exists i \in H : W_i[k] < \log b\}. \end{aligned} \quad (85)$$

Similarly, $\forall 1 \leq j \leq \ell$, it follows from the definition of $W_i[ja_b]$ that

$$W_i[k] \geq Z_i[\max\{v_i, (j-1)a_b + 1\}, ja_b]. \quad (86)$$

We then define

$$\begin{aligned} B'[ja_b] &= \{\exists i \in H : Z_i[\max\{v_i, (j-1)a_b + 1\}, ja_b] < \log b\}. \end{aligned} \quad (87)$$

It follows that $\forall 1 \leq j \leq \ell$,

$$B[ja_b] \subseteq B'[ja_b], \quad (88)$$

and $B'[ja_b]$ only depends on the samples from $(j-1)a_b + 1$ to ja_b .

Combining (83) and (88), equation (75) can be further bounded as follows:

$$\begin{aligned} \mathbb{P}_v(\bar{\tau}(b)/a_b > \ell) &\leq \mathbb{P}_v\left(\forall k \in \{a_b, 2a_b, \dots, \ell a_b\} : A'[k] \cup B'[k]\right) \\ &= \prod_{j=1}^{\ell} \mathbb{P}_v\left(A'[ja_b] \cup B'[ja_b]\right) \\ &\leq \prod_{j=1}^{\ell} \left(\mathbb{P}_v\left(A'[ja_b]\right) + \mathbb{P}_v\left(B'[ja_b]\right)\right). \end{aligned} \quad (89)$$

Following similar steps as from (62) to (69), we can show that for large b , $\forall 1 \leq j \leq \ell$,

$$\mathbb{P}_v\left(A'[ja_b]\right) \leq \delta/2, \quad (90)$$

where δ can be arbitrarily small.

Then by the Weak Law of Large Numbers, for large b , we obtain

$$\begin{aligned} \mathbb{P}_v(B'[ja_b]) &\leq \sum_{i \in H} \mathbb{P}_v\left(Z_i[\max\{v_i, (j-1)a_b + 1\}, ja_b] < \log b\right) \\ &\leq \sum_{i \in H} \delta' \\ &\triangleq \delta/2, \end{aligned} \quad (91)$$

where δ' can be made arbitrarily small for large b , and hence so can δ .

Therefore,

$$\mathbb{P}_v(\bar{\tau}(b)/a_b > \ell) \leq \delta^\ell, \quad (92)$$

and following steps similar to those in (70) to (72), we conclude the proof.

ACKNOWLEDGMENT

The authors would like to thank Dr. Ananthram Swami and Dr. Georgios Fellouris for valuable discussions and comments.

REFERENCES

- [1] S. Zou and V. V. Veeravalli, "Quickest detection of dynamic events in sensor networks," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, Calgary, AB, Canada, Apr. 2018, pp. 6907–6911.
- [2] S. Zou, V. V. Veeravalli, J. Li, and D. Towsley, "Quickest detection of significant events in structured networks," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Oct. 2018, pp. 1307–1311.
- [3] V. V. Veeravalli and T. Banerjee, "Quickest change detection," in *Academic Press Library in Signal Processing: Array and Statistical Signal Processing*, vol. 3. New York, NY, USA: Elsevier, 2013, pp. 209–256.
- [4] H. V. Poor and O. Hadjiladis, *Quickest Detection*. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [5] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*, vol. 104. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.
- [6] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Change-point Detection*. Boca Raton, FL, USA: CRC Press, 2014.
- [7] N. A. Christakis and J. H. Fowler, "Social network sensors for early detection of contagious outbreaks," *PLoS ONE*, vol. 5, no. 9, p. e12948, 2010.
- [8] F. Pervaiz, M. Pervaiz, N. A. Rehman, and U. Saif, "Flubreaks: Early epidemic detection from google flu trends," *J. Med. Internet Res.*, vol. 14, no. 5, p. e125, 2012.
- [9] A. Guille, H. Hacid, C. Favre, and D. A. Zighed, "Information diffusion in online social networks: A survey," *ACM SIGMOD Rec.*, vol. 42, no. 2, pp. 17–28, 2013.
- [10] S. Rajasegarar, C. Leckie, and M. Palaniswami, "Anomaly detection in wireless sensor networks," *IEEE Wireless Commun.*, vol. 15, no. 4, pp. 34–40, Aug. 2008.
- [11] A. Shabtai, R. Moskovitch, Y. Elovici, and C. Glezer, "Detection of malicious code by applying machine learning classifiers on static features: A state-of-the-art survey," *Inf. Secur. Tech. Rep.*, vol. 14, no. 1, pp. 16–29, 2009.
- [12] A. G. Tartakovsky and V. V. Veeravalli, "Change-point detection in multichannel and distributed systems," in *Applied Sequential Methodologies: Real-World Examples With Data Analysis*, vol. 173. Boca Raton, FL, USA: CRC Press, 2004, pp. 339–370.
- [13] 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 Trans. Signal Process.*, vol. 54, no. 9, pp. 3372–3382, Sep. 2006.
- [14] Y. Mei, "Efficient scalable schemes for monitoring a large number of data streams," *Biometrika*, vol. 97, no. 2, pp. 419–433, Apr. 2010.
- [15] Y. Xie and D. Siegmund, "Sequential multi-sensor change-point detection," *Ann. Statist.*, vol. 41, no. 2, pp. 670–692, 2013.
- [16] G. Fellouris and G. Sokolov, "Second-order asymptotic optimality in multisensor sequential change detection," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3662–3675, Jun. 2016.
- [17] V. Raghavan and V. V. Veeravalli, "Quickest change detection of a Markov process across a sensor array," *IEEE Trans. Inf. Theory*, vol. 56, no. 4, pp. 1961–1981, Apr. 2010.
- [18] M. Ludkovski, "Bayesian quickest detection in sensor arrays," *Sequential Anal.*, vol. 31, no. 4, pp. 481–504, 2012.
- [19] O. Hadjiladis, H. Zhang, and H. V. Poor, "One shot schemes for decentralized quickest change detection," *IEEE Trans. Inf. Theory*, vol. 55, no. 7, pp. 3346–3359, Jul. 2009.
- [20] S. Zou, G. Fellouris, and V. V. Veeravalli, "Asymptotic optimality of D-CuSum for quickest change detection under transient dynamics," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Aachen, Germany, Jun. 2017, pp. 156–160.

- [21] S. Zou, G. Fellouris, and V. V. Veeravalli, "Quickest change detection under transient dynamics: Theory and asymptotic analysis," *IEEE Trans. Inf. Theory*, vol. 65, no. 3, pp. 1397–1412, Mar. 2018.
- [22] G. Rovatsos, S. Zou, and V. V. Veeravalli, "Quickest change detection under transient dynamics," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, New Orleans, LA, USA, Mar. 2017, pp. 4785–4789.
- [23] G. Rovatsos, J. Jiang, A. D. Domínguez-García, and V. V. Veeravalli, "Statistical power system line outage detection under transient dynamics," *IEEE Trans. Signal Process.*, vol. 65, no. 11, pp. 2787–2797, Jun. 2017.
- [24] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Trans. Inf. Theory*, vol. 44, no. 7, pp. 2917–2929, Nov. 1998.
- [25] T. L. Lai, "Sequential changepoint detection in quality control and dynamical systems," *J. Roy. Stat. Soc. Ser. B, Methodol.*, vol. 57, no. 4, pp. 613–658, 1995.
- [26] T. L. Lai and H. Xing, "Sequential change-point detection when the pre- and post-change parameters are unknown," *Sequential Anal.*, vol. 29, no. 2, pp. 162–175, 2010.
- [27] D. Siegmund and E. S. Venkatraman, "Using the generalized likelihood ratio statistic for sequential detection of a change-point," *Ann. Statist.*, vol. 23, no. 1, pp. 255–271, 1995.
- [28] E. Arias-Castro, D. L. Donoho, and X. Huo, "Near-optimal detection of geometric objects by fast multiscale methods," *IEEE Trans. Inf. Theory*, vol. 51, no. 7, pp. 2402–2425, Jul. 2005.
- [29] L. Addario-Berry, N. Broutin, L. Devroye, and G. Lugosi, "On combinatorial testing problems," *Ann. Statist.*, vol. 38, no. 5, pp. 3063–3092, 2010.
- [30] E. Arias-Castro, E. J. Candès, H. Helgason, and O. Zeitouni, "Searching for a trail of evidence in a maze," *Ann. Statist.*, vol. 36, no. 4, pp. 1726–1757, 2008.
- [31] E. Arias-Castro, E. J. Candès, and A. Durand, "Detection of an anomalous cluster in a network," *Ann. Statist.*, vol. 39, no. 1, pp. 278–304, 2011.
- [32] S. Zou, Y. Liang, and H. V. Poor, "Nonparametric detection of geometric structures over networks," *IEEE Trans. Signal Process.*, vol. 65, no. 19, pp. 5034–5046, Oct. 2017.
- [33] S. Chen, Y. Yang, S. Zong, A. Singh, and J. Kovačević, "Detecting localized categorical attributes on graphs," *IEEE Trans. Signal Process.*, vol. 65, no. 10, pp. 2725–2740, May 2017.
- [34] J. Qian, V. Saligrama, and Y. Chen, "Connected sub-graph detection," in *Proc. Int. Conf. Artif. Intell. Statist. (AISTATS)*, Apr. 2014, pp. 796–804.
- [35] J. Sharpnack, A. Rinaldo, and A. Singh, "Changepoint detection over graphs with the spectral scan statistic," in *Proc. Int. Conf. Artif. Intell. Statist. (AISTATS)*, vol. 13, Scottsdale, AZ, USA, May 2013, pp. 545–553.
- [36] C. Aksoylar, L. Orecchia, and V. Saligrama, "Connected subgraph detection with mirror descent on SDPs," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Sydney, NSW, Australia, Aug. 2017, pp. 51–59.
- [37] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, nos. 1–2, pp. 100–115, 1954.
- [38] A. G. Tartakovsky, "Asymptotic performance of a multichart CUSUM test under false alarm probability constraint," in *Proc. IEEE Conf. Decis. Control*, Dec. 2005, pp. 320–325.
- [39] M. Pollak, "Optimal detection of a change in distribution," *Ann. Statist.*, vol. 13, no. 1, pp. 206–227, 1985.
- [40] M. Blum, R. W. Floyd, V. Pratt, R. L. Rivest, and R. E. Tarjan, "Time bounds for selection," *J. Comput. Syst. Sci.*, vol. 7, no. 4, pp. 448–461, Aug. 1973.
- [41] S. Li, Y. Xie, H. Dai, and L. Song, "M-statistic for kernel change-point detection," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2015, pp. 3366–3374.
- [42] T. S. Lau, W. P. Tay, and V. V. Veeravalli, "A binning approach to quickest change detection with unknown post-change distribution," 2018, *arXiv:1801.04756*. [Online]. Available: <https://arxiv.org/abs/1801.04756>
- [43] S. Zou, V. V. Veeravalli, J. Li, D. Towsley, and A. Swami, "Distributed quickest detection of significant events in networks," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, Brighton, U.K., May 2019, pp. 8454–8458.
- [44] G. Fellouris and A. G. Tartakovsky, "Multichannel sequential detection—Part I: Non-i.i.d. data," *IEEE Trans. Inf. Theory*, vol. 63, no. 7, pp. 4551–4571, Jul. 2017.

Shaofeng Zou (S'14–M'16) received the Ph.D. degree in Electrical and Computer Engineering from Syracuse University in 2016. He received the B.E. degree (with honors) from Shanghai Jiao Tong University, Shanghai, China, in 2011. He was a postdoctoral research associate at the Coordinated Science Lab, University of Illinois at Urbana-Champaign during 2016–2018. He joined the Department of Electrical Engineering, University at Buffalo in 2018, where he is currently an Assistant Professor. Dr. Zou's research interests include statistical signal processing, machine learning, and information theory.

Venugopal V. Veeravalli (M'92–SM'98–F'06) received the B.Tech. degree (Silver Medal Honors) from the Indian Institute of Technology, Bombay, India, in 1985, the M.S. degree from Carnegie Mellon University, Pittsburgh, PA, USA, in 1987, and the Ph.D. degree from the University of Illinois at Urbana-Champaign, Champaign, IL, USA, in 1992, all in electrical engineering.

He joined the University of Illinois at Urbana-Champaign in 2000, where he is currently the Henry Magnuski Professor in the Department of Electrical and Computer Engineering, and where he is also affiliated with the Department of Statistics, the Coordinated Science Laboratory, and the Information Trust Institute. Prior to joining the University of Illinois, he was on faculty of the ECE Department at Cornell University. He served as a Program Director for communications research at the U.S. National Science Foundation from 2003 to 2005. His research interests span the theoretical areas of statistical inference, machine learning, and information theory, with applications to data science, wireless communications, and sensor networks.

Prof. Veeravalli was a Distinguished Lecturer for the IEEE Signal Processing Society during 2010 and 2011. He has been on the Board of Governors of the IEEE Information Theory Society. He has been an Associate Editor for Detection and Estimation for the IEEE TRANSACTIONS ON INFORMATION THEORY and for the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS. Among the awards he has received for research and teaching are the IEEE Browder J. Thompson Best Paper Award, the National Science Foundation CAREER Award, and the Presidential Early Career Award for Scientists and Engineers (PECASE), and the Wald Prize in Sequential Analysis.

Jian Li (S'15–M'17) is an Assistant Professor of Computer Engineering in Department of Electrical and Computer Engineering at Binghamton University, the State University of New York. He was a postdoctoral research associate at the University of Massachusetts Amherst. He received his Ph.D. in Computer Engineering from Texas A&M University in December, 2016, and B.E. in Electrical Engineering from Shanghai Jiao Tong University in June, 2012. His current research interests lie broadly in the interplay of large scale networked systems and big data analytics focusing on applying machine learning, online algorithms, game theory, and signal processing technologies to big data applications.

Don Towsley (M'78–SM'93–F'95–LF'15) holds a B.A. in Physics (1971) and a Ph.D. in Computer Science (1975) from University of Texas. He is currently a Distinguished Professor at the University of Massachusetts in the College of Information & Computer Sciences. He has held visiting positions at numerous universities and research labs. His research interests include network science, performance evaluation, and quantum networking.

He is co-founder and Co-EiC of the new *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* (TOMPECS), and has served as Editor-in-Chief of IEEE/ACM TRANSACTIONS ON NETWORKING and on numerous editorial boards. He has served as Program Co-chair of several conferences including INFOCOM 2009.

He is a corresponding member of the Brazilian Academy of Sciences and has received several achievement awards including the 2007 IEEE Koji Kobayashi Award and the 2011 INFOCOM Achievement Award. He has received numerous paper awards including the 2012 ACM SIGMETRICS Test-of-Time Award, a 2008 SIGCOMM Test-of-Time Paper Award, and a 2018 SIGMOBILE Test-of-time Award. He also received the 1998 IEEE Communications Society William Bennett Best Paper Award. Last, he has been elected Fellow of both the ACM and IEEE.