A Tight Converse to the Asymptotic Performance of Byzantine Distributed Sequential Change Detection

Yu-Chih Huang
Dept. of Commun. Eng., NTPU
New Taipei City, Taiwan
ychuang@mail.ntpu.edu.tw

Shih-Chun Lin
Dept. of ECE, NTUST
Taipei, Taiwan
sclin@mail.ntust.edu.tw

Yu-Jui Huang
Dept. of Applied Math., Univ. of Colorado
Boulder, CO 80309, USA
yujui.huang@colorado.edu

***Abstract*—The Byzantine distributed sequential change detection (BDSCD) problem is studied, where a fusion center monitors an abrupt event occurring at an unknown time through a bunch of distributed sensors. It is assume that a part of the sensors are compromised and each sensor, honest or compromised, communicates with the fusion center via a noiseless link. A new converse for this problem is presented whose first-order asymptotic delay subject to a certain false alarm rate coincides with the currently best known result achieved by the *consensus rule* proposed by Fellouris *et al*. This result characterizes the first-order asymptotic performance of BDSCD and shows that 1-bit links suffice to achieve the asymptotic optimality. The proof of the converse involves constructing an attack strategy, called the reverse attack, introducing a genie that gives the fusion center the identities of a subset of honest sensors and observations at each sensor used for generating its local report, and transforming the problem into an equivalent non-Byzantine sequential change detection but with reduced number of honest sensors.**

I. Introduction

The problem of sequential change detection (SCD) studies detecting an abnormal event as quickly as possible after its occurrence at an unknown time, subject to a certain false alarm rate. It has many applications and has been extensively researched since the early works [1], [2], [3]. A nice tutorial on SCD can be found in [4]. However, recent applications such as massive machine-type communications (mMTC) and internet of things (IoT) [5] typically involve multiple distributed sensors monitoring the event and reporting their observations to the fusion center via bandwidth-limited links. Moreover, some sensors, whose identities are unknown to the fusion center, may be compromised and try to sabotage the detection. Motivated by these applications, this paper considers the decentralized version of SCD, where a fusion center monitors the event through distributed sensors, with compromised sensors forming Byzantine attack. This problem has been studied in [6], [7] and is called Byzantine distributed SCD (BDSCD).

In [6], a special case of BDSCD, with only one compromised sensor and with infinite-bandwidth links between the center and sensors, is considered. A decision rule called second-alarm rule is proposed and analyzed. In [7], the general BDSCD problem, with either infinite-bandwidth links or 1-bit links, is investigated. Multiple rules are proposed and their corresponding asymptotic performance are analyzed. Among the rules with 1-bit links proposed in [7], the voting rule that declares the occurrence of the event after the number of received local reports exceeds a certain threshold has the best first-order asymptotic performance. When the threshold is set to be the total number of honest sensors, the asymptotic performance of the voting rule, called the consensus rule in this special case, reaches its maximum. This achieves the best asymptotic performance of the Low-Sum-CUSUM scheme in [7], which requires infinite-bandwidth links.

Despite the exciting results in [6], [7], it is thus far unclear what the best first-order asymptotic performance for BDSCD is. Although one converse can be easily obtained by assuming that a genie reveals to the fusion center the identities of all honest sensors (this simple converse will be presented in Section II-A), this converse bound and the best achievable asymptotic performance in [7] do not match. This indicates that either the best achievable scheme thus far is not optimal or the converse is not tight, or both.

In this work, we present a new converse for the first-order asymptotic performance of BDSCD. In our proof, we first construct an attack strategy referred to as the "reverse attack". Then, we assume a less powerful genie who only reveals the identities of *some* honest sensors and all the observations that each sensor uses for generating its local report. After that, inspired by the very recent work of Chen and Wang [8], we transform the genie-aided version of BDSCD into an equivalent centralized SCD, for which the cumulative sum (CUSUM) procedure is known to be optimal [9]. Evaluating the performance of CUSUM then shows the new converse. It turns out that the new converse and the consensus rule have the same first-order asymptotic performance; hence, the first-order asymptotic performance of BDSCD is characterized. An important implication of this result is that 1-bit links are good enough for BDSCD in terms of first-order asymptotic delay performance.

Notations : For a positive integer K, define $[K] := \{1, \dots, K\}$ and $[K]^+ = \{0\} \cup [K]$. Function $(x)^+$ outputs x if $x \geq 0$ and zero otherwise. For two real functions $f_1(x)$ and $f_2(x)$, as $x \to \infty$, we write $f_1(x) \sim f_2(x)$ when $f_1(x)/f_2(x) \to 1$ and $f_1(x) \gtrsim f_2(x)$ when $\liminf(f_1(x)/f_2(x)) \geq 1$.

II. Problem Formulation and Known Results

The BDSCD problem consists of a fusion center and K sensors indexed by $[K]$. Among these sensors, there is an *unknown* subset $\mathcal{N} \subset [K]$ of honest sensors, with the remaining $M := K - |\mathcal{N}|$ sensors being potentially compromised. The

goal of the honest sensors is to monitor an event and help the fusion center decide whether the event has occurred, while the goal of compromised sensors is to collaboratively confuse the fusion center. Although the exact information about which sensors are honest and which sensors are compromised is unknown, we assume that M, the maximum number of sensors the attacker can compromise, is known by the fusion center. Moreover, it is assumed that there are more honest sensors than compromised sensors, i.e., $|\mathcal{N}| > M$. The observations of all K sensors are sequences of independent random variables with known distributions, subject to the same distribution change at a unknown but deterministic time ν. Before the change time ν, sensor k's observations $X_1^k, X_2^k, \ldots, X_\nu^k$ are independent and identically distributed (i.i.d.) with the density P_0, while $X_{\nu+1}^k, X_{\nu+2}^k, \ldots$ are i.i.d. with the density P_1. If the change never happens, i.e., $\nu = \infty$, X_t^k are i.i.d. with P_0 for all t. We denote by $\mathbf{X}_t = [X_t^1, X_t^2, \ldots, X_t^K]$ the collection of observations at time t and we use the notation $\mathbf{X}_{t_1}^{t_2}$ for $t_1 < t_2$ to denote the collection $[\mathbf{X}_{t_1}, \mathbf{X}_{t_1+1}, \ldots, \mathbf{X}_{t_2}]$. Also, we define the Kullback-Leibler information from P_0 to P_1 as [10], $I := \int \log\left(\frac{P_1(x)}{P_0(x)}\right) P_1(x)dx$. Throughout the paper, we assume that I is finite and strictly positive and $\int \log\left(P_1(x)/P_0(x)\right)^2 P_1(x)dx \leq \infty$.

All the local reports from honest or compromised sensors belong to the set $\mathcal{X}$, which satisfies the underlying bandwidth constraint on the noiseless link between each sensor and the fusion center. It is worth emphasizing that this setting encompasses many scenarios discussed in existing works including $\mathcal{X} = \{0, 1\}$ and $\mathcal{X} = \mathbb{R}$ in [7] and $\mathcal{X}$ being a set of finite alphabets in [9]. At each time index t, the honest sensor k individually makes a local decision by mapping its own observations up to time t to an element in $\mathcal{X}$, and then chooses to report it or not according to the adopted reporting mechanism. Based on the received local reports from all sensors, the fusion center adopts a stopping rule to determine when to declare the event has occurred. A change detection rule T includes such a stopping rule and local rules at honest sensors. The M compromised sensors, on the other hand, try to disrupt/confuse the fusion center by sending attack signals. We assume a very powerful attacker that knows the exact change-time ν and have the access to the current and past observations of all nodes. The symbols sent by the compromised sensors at time t are then produced by g, a function (called an attack strategy) with inputs ν, $\mathbf{X}_1^t$, and T. We denote by $\mathcal{G}$ the set of all attack strategies including all possible g with no more than M compromised sensors. Following [7], we analyze the performance of rule T by its worst-case expected detection delay and mean time to false alarm in the sense of Lorden [2], under the worst attack strategy among $\mathcal{G}$. Specifically, we define the performance metrics as follows.

- **Detection Delay:** The worst-case mean detection delay

$$\mathcal{D}[T] := \sup_{g\in\mathcal{G},\nu} \operatorname{ess\,sup} \mathbb{E}_\nu^g[(T-\nu)^+ | \mathbf{X}_1^\nu], \tag{1}$$

  where $\mathbb{E}_\nu^g[.]$ means the expectation is taken w.r.t. P_0 when $t \leq \nu$ and w.r.t. P_1 when $t > \nu$ under the attack strategy $g \in \mathcal{G}$.

- **False Alarm:** Without any abnormal changes (i.e. $\nu = \infty$), the worst-case mean time to false alarm is

$$\mathcal{A}[T] := \inf_{g\in\mathcal{G}} \mathbb{E}_\infty^g[T], \tag{2}$$

  where $\mathbb{E}_\infty^g[.]$ means the expectation is w.r.t. P_0 for all t (i.e., $\nu = \infty$) under the attack strategy $g \in \mathcal{G}$.

The main theme of this paper is to investigate the optimal asymptotic behavior of how the mean detection delay scales with the mean time to false alarm in the worst case. Specifically, for an optimal detection rule T that has $\mathcal{A}[T] = \gamma$, we want to characterize how $\mathcal{D}[T]$ grows with $\log(\gamma)$ as $\gamma \to \infty$.

A. Known results

For the considered problem with $|\mathcal{N}| = 1$ and $M = 0$, the problem reduces to the standard SCD problem for which it was shown in [2], [3] that Page's CUSUM procedure T_{single} [1] achieves the optimal scaling that for $\mathcal{A}[T_{\text{single}}] = \gamma$, the expected detection delay scales like $\mathcal{D}[T_{\text{single}}] \sim \log\gamma/I$ as $\gamma \to \infty$. For $M = 0$ but general $|\mathcal{N}|$, Mei in [9] developed a scheme T_{multiple}, where each sensor performs CUSUM according to its local observations and sends a binary report to the fusion center, which declares the occurrence of the event when all the $|\mathcal{N}|$ sensors say so. It was then shown that this scheme is asymptotically optimal that for $\mathcal{A}[T_{\text{multiple}}] = \gamma$, the expected detection delay scales like $\mathcal{D}[T_{\text{multiple}}] \sim \log\gamma/|\mathcal{N}|I$ as $\gamma \to \infty$.

Very recently, in [7], the BDSCD for general $|\mathcal{N}|$ and M was discussed and multiple schemes were analyzed. Among these schemes, the L-voting rule T_L achieves the best scaling. In this scheme, each sensor adopts the same procedure as in T_{multiple}, and the fusion center declares the occurrence of the event at the first time at least L sensors raises alarms simultaneously. The scaling achieved by this scheme is summarized as follows.

Theorem 1 ([7, Theorem 26])**.** *Let L be an integer satisfying $M < L \leq |\mathcal{N}|$. For $\mathcal{A}[T_L] = \gamma$, as $\gamma \to \infty$, the worst-case mean detection delay of the voting rule scales like*

$$\mathcal{D}[T_L] \sim \frac{\log\gamma}{(L-M)I}, \tag{3}$$

The best asymptotic performance reported in [7] is the above one with $L = |\mathcal{N}|$, which also coincides with another scheme in [7], Low-Sum-CUSUM, that requires infinite bandwidth. This leads us to conjecture that (3), with $L = |\mathcal{N}|$, is the optimal first-order behavior. A tight converse is then necessary to verify this conjecture.

We would like to point out that a converse can be obtained by revealing the identities of all $|\mathcal{N}|$ honest sensors and using the asymptotic optimality in [9] with the honest sensors $\mathcal{N}$ as

Theorem 2 (Simple converse)**.** *For any Byzantine change detection rule T with $\mathcal{A}[T] = \gamma$, as $\gamma \to \infty$, the worst-case mean detection delay meets $\mathcal{D}[T] \gtrsim \frac{\log\gamma}{|\mathcal{N}|I}$.*

Unfortunately, this converse is not tight compared to (3).

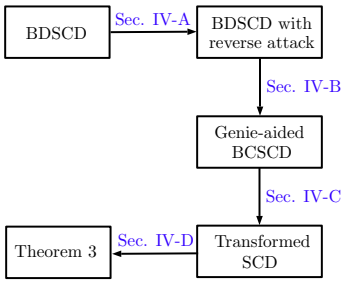

Fig. 1. Diagram of proof steps.

III. Main Results

In this section, we present the main result of this paper, which is a new converse of the first-order asymptotic performance of BDSCD. This main result, together with achievability of the consensus rule in Theorem 1 which requires only 1-bit link, characterizes the optimal first-order behavior of BDSCD.

Theorem 3 (Tight converse)**.** *For any Byzantine change detection rule T with $\mathcal{A}[T] = \gamma$, as $\gamma \to \infty$, the worst-case mean detection delay is lower bounded as*

$$\mathcal{D}[T] \gtrsim \frac{\log \gamma}{(|\mathcal{N}| - M)I}. \tag{4}$$

To prove this theorem, we first note that if the optimal asymptotic scaling is lower bounded by $\eta(\gamma)$ under an attack strategy, then it is also lower bounded by $\eta(\gamma)$ under the *worst* attack. We thus proceed by constructing an attack strategy, called the *reverse attack*, which yields the lower bound in Theorem 3. This is done by assuming that a genie provides the identities of $|\mathcal{N}| - M$ out of $|\mathcal{N}|$ honest sensors and the local observations used for generating the local report at every sensor. After that, by absorbing the impact of the reverse attack into pre/post-change distributions, the problem is transformed into an equivalent centralized SCD problem for which CUSUM is known to be optimal. Finally, evaluating the CUSUM procedure for the transformed problem reveals the connection to another non-Byzantine SCD with only $|\mathcal{N}| - M$ honest sensors. A sketch of the proof of the new converse is outlined in Fig. 1 and the details are given in the next section.

IV. The Proof

The proof presented in this section follows closely the steps shown in Fig. 1.

A. The reverse attack

For the ease of presentation, we let $P_{0,1} = P_0$ and $P_{1,1} = P_1$. Recall that each honest sensor k's observation sequence X_t^k is drawn i.i.d. according to $P_{0,1}$ before the change time ν and i.i.d. according to $P_{1,1}$ after ν. We construct an attack strategy as follows. For each compromised sensors k', it generates a fake observation sequence $X_t^{k'}$, which is then input to the assigned local decision function for forming the fake report. The fake observation sequence is generated i.i.d. according to $P_{0,2}$ and $P_{1,2}$ before and after the change time ν, respectively. i.e., the compromised sensors form fake reports according to observations based on wrong distributions. To establish the tight converse, we will set $P_{0,2} = P_{1,1} = P_1$ and $P_{1,2} = P_{0,1} = P_0$ in the very end of the proof; therefore, we call this attack strategy the "reverse attack". However, most of the steps in the proof stay valid for general densities $P_{0,2}$ and $P_{1,2}$.

B. Genie-aided Byzantine centralized SCD

First note that the worst case happens when there are M compromised sensors. Also since the identities of the sensors are unknown, the fusion center cannot enhance the worst-case performance by selectively accepting reports. If the fusion center accepts reports from $K - K'$, $K' \leq |\mathcal{N}|$, sensors only, in the worst case, the problem reduces to the BDSCD with M compromised sensors and $|\mathcal{N}| - K'$ honest sensors, which results in a worse performance. Moreover, when $K' > |\mathcal{N}|$, we are left with only compromised sensors in the worst case, which is obviously worse than accepting all reports. We therefore only have to consider that the fusion center uses reports from all K sensors for detection in what follows.

The K sensors are divided into three groups where the first two groups consist of M sensors each, while the last one has $|\mathcal{N}| - M$ sensors. All sensors in the first or third groups are honest while those in the second group are compromised. Assume that there is a genie giving away the identities of $|\mathcal{N}| - M$ honest sensors to the fusion center. For the rest M honest sensors and M compromised sensors, the identities are unknown to the fusion center. Without loss of generality, we assume that sensors in the first two groups have indices $[2M]$. We also give the observations used at each sensor (fake observations if the sensor is compromised) for generating its local report and the densities $P_{0,2}$ and $P_{1,2}$ to the fusion center. Let $s : [K] \to [2]$ be a function that assigns each sensor to group 1 or 2 in such a way that exactly M out of the first $2M$ sensors are assigned to compromised group 2 and the last $|\mathcal{N}| - M$ sensors are all assigned to honest group 1. Let $\mathcal{S}$ be the collection of all possible assignments s. Clearly, there are total $|\mathcal{S}| = \binom{2M}{M}$ such assignments. For $\theta \in \{0, 1\}$, the product density under the compromised group assignment s is

$$P_{\theta,s}(\mathbf{X}_t) = \prod_{k'=1}^{2M} P_{\theta,s(k')}(X_t^{k'}) \prod_{k=2M+1}^{K} P_{\theta,1}(X_t^k). \tag{5}$$

Now, we are facing a composite change detection problem, which we refer to as genie-aided Byzantine centralized SCD (BCSCD). Before the change time ν, the random vectors $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_{\nu-1}$ are i.i.d. over time with density $P_{0,s}$ while $\mathbf{X}_\nu, \mathbf{X}_{\nu+1}, \ldots$ are generated with density $P_{1,s}$, for some $s \in \mathcal{S}$. In this genie-aided version, the fusion center knows everything about the compromised sensors except for their exact locations. With slight abuse of notations, as (1), the mean detection delay of this problem is given by

$$\mathcal{D}[T] := \sup_{s \in \mathcal{S}, \nu} \operatorname{ess\,sup} \mathbb{E}_\nu^s[(T - \nu)^+ | \mathbf{X}_1^\nu]; \tag{6}$$

also as (2), the mean time to false alarm is

$$\mathcal{A}[T] := \inf_{s\in\mathcal{S}} \mathbb{E}^s_\infty[T]. \tag{7}$$

C. Transformed Centralized SCD

We transform the genie-aided BCSCD problem into an equivalent centralized SCD problem for which CUSUM is known to be optimal. Let $\tau_{2M}(\mathbf{X}_t)$ be the masked ordering map that puts the first $2M$ elements of its input $\mathbf{X}_t$ in descending order while keeps the other $|\mathcal{N}|-M$ positions unchanged. A decision rule $T(.)$ of the genie-aided BCSCD problem is said to be masked symmetric if it can be represented as $T(\{\mathbf{X}_t\}_{t\geq 1}) = \tilde{T}(\{\tau_{2M}(\mathbf{X}_t)\}_{t\geq 1})$ for some decision rule $\tilde{T}$.

In the transformed centralized SCD problem, the fusion center observes $\tilde{\mathbf{X}}_t = \tau_{2M}(\mathbf{X}_t)$ at time t. Let $\tilde{P}_\theta(\tilde{\mathbf{X}}_t)$ be the density of $\tau_{2M}(\mathbf{X}_t)$, where $\mathbf{X}_t$ is generated according to density $P_{\theta,s}, \theta\in\{0,1\}$. Before the change, the observations $\{\tilde{\mathbf{X}}_t\}$ follow $\tilde{P}_0$ while after the change, they follow $\tilde{P}_1$. Also,

$$\tilde{P}_\theta(\tilde{\mathbf{X}}_t) = \sum_{\mathbf{X}_t:\tau_{2M}(\mathbf{X}_t)=\tilde{\mathbf{X}}_t} P_{\theta,s}\left(\mathbf{X}_t\right), \tag{8}$$

where the equality follows from that the absolute value of the Jacobian of a permutation is always 1. Following the proof of part 1 of [8, Lemma 4.1], we can easily show that for all assignments $s\in\mathcal{S}$, the density $\tilde{P}_s(\tilde{\mathbf{X}}_t)$ does not depend on s. Suppose the change occurs at the time index ν. Under hypothesis $\tilde{H}_1$, the random vectors $\tilde{\mathbf{X}}_1, \tilde{\mathbf{X}}_2, \ldots, \tilde{\mathbf{X}}_{\nu-1}$ are drawn i.i.d. over time with density $\tilde{P}_0$ while $\tilde{\mathbf{X}}_\nu, \tilde{\mathbf{X}}_{\nu+1}, \ldots$ are generated i.i.d. with density $\tilde{P}_1$. Under hypothesis $\tilde{H}_0$, there is no change, i.e. $\nu=\infty$, and $\tilde{\mathbf{X}}_t$ are drawn i.i.d. with density $\tilde{P}_0$ for all t.

We first focus on a masked symmetric rule T, and show that the detection delay of the genie-aided BCSCD is identical to that of the transformed SCD problem, defined as $\mathcal{D}[\tilde{T}] := \sup_\nu \operatorname{ess\,sup} \mathbb{E}_\nu[(\tilde{T}-\nu)^+|\tilde{\mathbf{X}}_1^\nu]$. Specifically, we will show

$$\begin{aligned}\mathcal{D}[T] &= \sup_{s\in\mathcal{S},\nu} \operatorname{ess\,sup} \mathbb{E}^s_\nu[(\tilde{T}(\{\tau_{2M}(\mathbf{X}_t)\}_{t\geq1})-\nu)^+|\mathbf{X}_1^\nu] \\ &\overset{(a)}{=} \sup_\nu \operatorname{ess\,sup} \mathbb{E}_\nu[(\tilde{T}(\{\tilde{\mathbf{X}}_t\}_{t\geq1})-\nu)^+|\tilde{\mathbf{X}}_1^\nu] = \mathcal{D}[\tilde{T}]. \end{aligned}\tag{9}$$

The first equality is from the definition in (6) and we will devote to prove equality (9a). To do this, note that

$$\begin{aligned}&\mathbb{E}^s_\nu[(\tilde{T}(\{\tau_{2M}(\mathbf{X}_t)\}_{t\geq1})-\nu)^+|\mathbf{X}_1^\nu] \\ &=\sum_{z=0}^{\infty} 1-\mathbb{P}\left(\left(\tilde{T}(\{\tau_{2M}(\mathbf{X}_t)\}_{t\geq1})-\nu\right)^+\leq z\middle|\mathbf{X}_1^\nu\right) \\ &=\sum_{z=0}^{\infty} 1-\int 1_{\left\{\left(\tilde{T}(\{\tau_{2M}(\mathbf{X}_t)\}_{t\geq1})-\nu\right)^+\leq z\right\}}\prod_{t=\nu+1}^{\nu+z} P_{1,s}(\mathbf{X}_t)d\mathbf{X}_{\nu+1}^{\nu+z} \\ &\overset{(a)}{=}\sum_{z=0}^{\infty} 1-\int 1_{\left\{\left(\tilde{T}(\{\tilde{\mathbf{X}}_t\}_{t\geq1})-\nu\right)^+\leq z\right\}}\prod_{t=\nu+1}^{\nu+z} \tilde{P}_1(\tilde{\mathbf{X}}_t)d\tilde{\mathbf{X}}_{\nu+1}^{\nu+z} \\ &=\sum_{z=0}^{\infty} 1-\mathbb{P}\left(\left(\tilde{T}(\{\tilde{\mathbf{X}}_t\}_{t\geq1})-\nu\right)^+\leq z\middle|\tilde{\mathbf{X}}_1^\nu\right) \\ &=\mathbb{E}_\nu[(\tilde{T}(\{\tilde{\mathbf{X}}_t\}_{t\geq1})-\nu)^+|\tilde{\mathbf{X}}_1^\nu], \end{aligned}\tag{10}$$

where $\mathbb{P}()$ is the probability and $1_{\{.\}}$ is the indicator function; and (a) follows from the change of variables in integration [11] and the fact that $\tilde{P}_1(.)$ does not depend on s.

Now, for a fixed ν and for each $s\in\mathcal{S}$, let $\mathbb{P}^s$ and $\tilde{\mathbb{P}}$ denote the probability measures on $\mathbb{R}^{K\times\nu}$ with densities specified by (5) and (8) with $\theta=0$, respectively. To establish (9a), observe that for any $x\in\mathbb{R}^{K\times\nu}$, from (10), we have

$$\begin{aligned}&\mathbb{E}^s_\nu[(\tilde{T}(\{\tau_{2M}(\mathbf{X}_t)\}_{t\geq1})-\nu)^+|\mathbf{X}_1^\nu](x) = \\ &\mathbb{E}_\nu[(\tilde{T}(\{\tilde{\mathbf{X}}_t\}_{t\geq1})-\nu)^+|\tilde{\mathbf{X}}_1^\nu](\tau_{2M}(x)).\end{aligned}\tag{11}$$

Let D_M denote $\operatorname{ess\,sup}\mathbb{E}^s_\nu[(\tilde{T}(\{\tau_{2M}(\mathbf{X}_t)\}_{t\geq1})-\nu)^+|\mathbf{X}_1^\nu]$ where the essential supremum is taken under $\mathbb{P}^s$. By definition, there exists $\Omega\subseteq\mathbb{R}^{K\times\nu}$ with $\mathbb{P}^s(\Omega)=1$ such that $D_M \geq \mathbb{E}^s_\nu[(\tilde{T}(\{\tau_{2M}(\mathbf{X}_t)\}_{t\geq1})-\nu)^+|\mathbf{X}_1^\nu](x)$ for all $x\in\Omega$. By (11), we have

$$D_M \geq \mathbb{E}_\nu[(\tilde{T}(\{\tilde{\mathbf{X}}_t\}_{t\geq1})-\nu)^+|\tilde{\mathbf{X}}_1^\nu](\tau_{2M}(x)), \quad \forall x\in\Omega.$$

Note that

$$\tilde{\mathbb{P}}(\tau_{2M}(\Omega)) = \int_{\tau_{2M}(\Omega)}\tilde{P}_0(y)dy = \int_\Omega P_{0,s}(x)dx = \mathbb{P}^s(\Omega)=1,$$

where the densities $\tilde{P}_0$ and $P_{0,s}$ are given by (8) and (5) respectively. We therefore conclude that $D_M \geq \operatorname{ess\,sup}\mathbb{E}_\nu[(\tilde{T}(\{\tilde{\mathbf{X}}_t\}_{t\geq1})-\nu)^+|\tilde{\mathbf{X}}_1^\nu]$ where the essential supremum here is taken under $\tilde{\mathbb{P}}$. Noting that this is true for every $s\in\mathcal{S}$ shows the relation $\geq$ in (9a). By using the same argument as above, but switching the roles of the left-hand side and right-hand side of (9a), we obtain the relation "$\leq$".

We have shown that the detection delay of genie-aided BCSCD (6) is equal to that of transformed SCD under masked symmetric rules. One can similarly prove that the mean time to false alarm of the original problem (7) is equal to that of the new problem. The rest is to show that for any fusion rule $T'(.)$, there is a masked symmetric rule $T(.)$ that is not worse than $T'(.)$. This is shown in Lemma A.1 in Appendix and then the transformation of SCD is established.

D. Establishing the converse

So far, we have transformed the genie-aided BCSCD problem into an equivalent centralized SCD problem with observations following $\tilde{P}_0$ and $\tilde{P}_1$ before and after the change point ν, respectively. For such a problem, it is well known from [3], [9, Lemma 2] that an optimal strategy is Page's CUSUM procedure given by $\tilde{\sigma}(h) = \inf\{t\in\mathbb{N}: \tilde{Y}_t \geq h\}$, where $\tilde{Y}_t = (\tilde{Y}_{t-1}+\tilde{\ell}_t)^+$ with $\tilde{Y}_0=0$, and from (8)

$$\begin{aligned}\tilde{\ell}_t &= \log\frac{\sum_{\mathbf{X}_t:\tau_{2M}(\mathbf{X}_t)=\tilde{\mathbf{X}}_t} P_{1,s}(\mathbf{X}_t)}{\sum_{\mathbf{X}_t:\tau_{2M}(\mathbf{X}_t)=\tilde{\mathbf{X}}_t} P_{0,s}(\mathbf{X}_t)} \\ &= \log\frac{\sum_{\pi\in\Pi_{2M}} P_{1,s}(\pi(\mathbf{X}_t))}{\sum_{\pi\in\Pi_{2M}} P_{0,s}(\pi(\mathbf{X}_t))} = \log\frac{\sum_{\pi\in\Pi_{2M}} P_{1,s\circ\pi^{-1}}(\mathbf{X}_t)}{\sum_{\pi\in\Pi_{2M}} P_{0,s\circ\pi^{-1}}(\mathbf{X}_t)} \\ &\overset{(a)}{=} \log\frac{\sum_{s'\in\mathcal{S}} P_{1,s'}(\mathbf{X}_t)\frac{2M!}{\binom{2M}{M}}}{\sum_{s'\in\mathcal{S}} P_{0,s'}(\mathbf{X}_t)\frac{2M!}{\binom{2M}{M}}}. \end{aligned}\tag{12}$$

where $\circ$ is the function composition operator, and (a) follows from the fact that for a compromised group assignment s,

summing over all the permuted versions $s \circ \pi^{-1}$ is equivalent to summing over all the assignments s' with each s' being involved $2M!/\binom{2M}{M}$ times. In what follows, we set $P_{0,2} = P_{1,1}$ and $P_{1,2} = P_{0,1}$ according to the reverse attack described in Sec. IV-A. We now rewrite the likelihood in (12) as

$$\tilde{\ell}_t = \log \frac{\sum_{s\in\mathcal{S}} P_{1,s}(\mathbf{X}_t)}{\sum_{s\in\mathcal{S}} P_{0,s}(\mathbf{X}_t)} \tag{13}$$

$$\overset{(a)}{=} \log \frac{\left(\sum_{s\in\mathcal{S}} \prod_{k'=1}^{2M} P_{1,s(k')}(X_t^{k'})\right) \prod_{k=2M+1}^{K} P_{1,1}(X_t^k)}{\left(\sum_{s\in\mathcal{S}} \prod_{k'=1}^{2M} P_{0,s(k')}(X_t^{k'})\right) \prod_{k=2M+1}^{K} P_{0,1}(X_t^k)}$$

$$\overset{(b)}{=} \log \frac{\prod_{k=2M+1}^{K} P_{1,1}(X_t^k)}{\prod_{k=2M+1}^{K} P_{0,1}(X_t^k)}, \tag{14}$$

where (a) follows from (5) and (b) is because of the fact that for every s, there exists a $\bar{s}$ such that whenever $s(k') = 1$, $\bar{s}(k') = 2$ and whenever $s(k') = 2$, $\bar{s}(k') = 1$; therefore,

$$\sum_{s\in\mathcal{S}} \prod_{k'=1}^{2M} P_{1,s(k')}(X_t^{k'}) = \sum_{s\in\mathcal{S}} \prod_{k'=1}^{2M} P_{0,\bar{s}(k')}(X_t^{k'})$$
$$= \sum_{\bar{s}\in\mathcal{S}} \prod_{k'=1}^{2M} P_{0,\bar{s}(k')}(X_t^{k'}), \tag{15}$$

where the first equality is from $P_{1,1} = P_{0,2}$ and $P_{1,2} = P_{0,1}$. Hence, the optimal test reduces to the standard centralized CUSUM procedure for the change detection with $|\mathcal{N}| - M$ honest sensors. Applying the results in [9] then shows (4).

V. Discussions

We discuss two byproducts obtained along the proof. Firstly, it is immediate that the "reverse attack" proposed in Section IV-A is an asymptotically worst attack for the original BDSCD problem. Secondly, in Section IV-D, we have shown that the decision rule, called mixture CUSUM, defined in the following is optimal for the genie-aided BCSCD.

Definition 1 (mixture CUSUM)**.** *For a sequence of observations $\mathbf{X}_t \triangleq (X_t^1, X_t^2, \ldots, X_t^K)$ at any time t, the mixture CUSUM algorithm computes $Y_t = (Y_{t-1} + \ell_t)^+$, where $Y_0 = 0$ and ℓ_t is the mixture likelihood ratio defined as* (13)*. The mCUSUM reports an alarm once Y_t exceeds the prescribed threshold h. The corresponding stopping time is then given by $\sigma(h) = \inf\{t \in \mathbb{N} : Y_t \geq h\}$.*

Appendix

Here we prove the following lemma.

Lemma A.1. *For any general (not necessarily masked symmetric) fusion rule $T'(\{\mathbf{X}_t\}_{t\geq 1})$, there is a masked symmetric rule $T(\{\mathbf{X}_t\}_{t\geq 1})$ that is not worse than $T'(\{\mathbf{X}_t\}_{t\geq 1})$.*

Proof: The proof is a constructive one. We first define $\pi : [K] \to [K]$ a masked permutation function that permutes the first $2M$ entries while keeps the remaining $|\mathcal{N}| - M$ entries unchanged. Let Π_{2M} be the collection of all $(2M!$ in total) such π. For $\mathbf{X}_t = [X_t^1, \ldots, X_t^K]$, we slightly abuse the notation to write $\pi(\mathbf{X}_t) = [X_t^{\pi(1)}, \ldots, X_t^{\pi(K)}]$. Let

$$T(\{\mathbf{X}_t\}_{t\geq 1}) = \frac{1}{2M!} \sum_{\pi\in\Pi_{2M}} T'(\{\pi(\mathbf{X}_t)\}_{t\geq 1}). \tag{16}$$

Following the proof of part 1 of [8, Lemma 4.2], one can show that $T(\{\mathbf{X}_t\}_{t\geq 1})$ is indeed a masked symmetric strategy. Now the detection delay $\mathcal{D}[T(\{\mathbf{X}_t\}_{t\geq 1})]$ is

$$\sup_{s\in\mathcal{S},\nu} \operatorname{ess\,sup} \mathbb{E}_\nu^s \left[\left(\frac{1}{2M!} \sum_{\pi\in\Pi_{2M}} T'(\{\pi(\mathbf{X}_t)\}_{t\geq 1}) - \nu \right)^+ \middle| \mathbf{X}_1^\nu \right].$$

Thus $\mathcal{D}[T(\{\mathbf{X}_t\}_{t\geq 1})]$ is no longer than

$$\frac{1}{2M!} \sum_{\pi\in\Pi_{2M}} \sup_{s\in\mathcal{S},\nu} \operatorname{ess\,sup} \mathbb{E}_\nu^s \left[(T'(\{\pi(\mathbf{X}_t)\}_{t\geq 1}) - \nu)^+ | \mathbf{X}_1^\nu \right]$$

$$\overset{(a)}{=} \frac{1}{2M!} \sum_{\pi\in\Pi_{2M}} \sup_{s\in\mathcal{S},\nu} \operatorname{ess\,sup} \mathbb{E}_\nu^{s\circ\pi^{-1}} \left[(T'(\{\mathbf{X}_t\}_{t\geq 1}) - \nu)^+ | \mathbf{X}_1^\nu \right]$$

$$\overset{(b)}{=} \frac{1}{2M!} \sum_{\pi\in\Pi_{2M}} \sup_{s'\in\mathcal{S},\nu} \operatorname{ess\,sup} \mathbb{E}_\nu^{s'} \left[(T'(\{\mathbf{X}_t\}_{t\geq 1}) - \nu)^+ | \mathbf{X}_1^\nu \right]$$

$$= \frac{1}{2M!} \sum_{\pi\in\Pi_{2M}} \mathcal{D}[T'(\{\mathbf{X}_t\}_{t\geq 1})] = \mathcal{D}[T'(\{\mathbf{X}_t\}_{t\geq 1})]. \tag{17}$$

Note that essential supremum of the right-hand side of (a) is taken under the probability measure whose density is specified by (5) under $\theta = 0$ and compromised group assignment $s \circ \pi^{-1}$. Then (a) can be proved similar to (9a) by the fact

$$\mathbb{E}_\nu^s \left[(T'(\{\pi(\mathbf{X}_t)\}_{t\geq 1}) - \nu)^+ \;\middle|\; \mathbf{X}_1^\nu \right](x) =$$
$$\mathbb{E}_\nu^{s\circ\pi^{-1}} \left[(T'(\{\mathbf{X}_t\}_{t\geq 1}) - \nu)^+ \;\middle|\; \mathbf{X}_1^\nu \right](\pi(x)), \forall x \in \mathbb{R}^{K\times\nu}$$

since the permutation $\pi(.)$ is one-to-one; and (b) is due to the fact that $\{s \circ \pi^{-1} | s \in \mathcal{S}\} = \mathcal{S}$. We can similarly show that for the mean time to false alarm of the new rule, $\mathcal{A}[T(\{\mathbf{X}_t\}_{t\geq 1})] \geq \mathcal{A}[T'(\{\mathbf{X}_t\}_{t\geq 1})]$. Then we conclude that the masked symmetric strategy $T(\{\mathbf{X}_t\}_{t\geq 1})$ is at least as good as $T'(\{\mathbf{X}_t\}_{t\geq 1})$. ∎

References

[1] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, pp. 100—-115, Jun. 1954.

[2] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, no. 6, pp. 1897—-1908, Dec. 1971.

[3] G. V. Moustakides, "Optimal detection of a change in distribution," *Ann. Statist.*, vol. 14, pp. 1379—-1387, Dec. 1986.

[4] V. V. Veeravalli and T. Banerjee, "Quickest change detection," *arXiv:1210.5552v1 [math.ST]*, Oct. 2012.

[5] Y.-P. E. Wang, X. Lin, A. Adhikary, A. Grovlen, Y. Sui, Y. Blankenship, J. Bergman, and H. S. Razaghi, "A primer on 3GPP narrowband internet of things," *IEEE Comm. Mag.*, vol. 55, no. 3, pp. 117–123, March 2017.

[6] E. Bayraktar and L. Lai, "Byzantine fault tolerant distributed quickest change detection," *SIAM J. Control Optim.*, vol. 53, no. 2, pp. 575–591, 2015.

[7] G. Fellouris, E. Bayraktar, and L. Lai, "Efficient Byzantine sequential change detection," *IEEE Trans. Inf. Theory*, vol. 64, no. 5, pp. 3346–3360, May 2018.

[8] W.-N. Chen and I.-H. Wang, "Anonymohs heterogeneous distributed detection: Optimal decision rules, error exponents, and the price of anonymity," *arXiv:1805.03554v2 [cs.IT]*, May 2018.

[9] Y. Mei, "Information bounds and quickest change detection in decentralized decision systems," *IEEE Trans. Inf. Theory*, vol. 51, no. 7, pp. 2669–2681, Jul. 2005.

[10] T. Cover and J. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.

[11] P. Billingsley, *Probability and measure*, 3rd ed. Wiley-Interscience., 2008.