Optimal Quickest Change Detection in Sensor Networks Using Ordered Transmissions

Yicheng Chen
Lehigh University
Bethlehem, PA, USA
yic917@lehigh.edu

Rick S. Blum
Lehigh University
Bethlehem, PA, USA
rblum@lehigh.edu

Brian M. Sadler
Army Research Laboratory
Adelphi, MD, USA
brian.m.sadler6.civ@mail.mil

***Abstract*—Quickest change detection in a sensor network is considered where each sensor observes a sequence of random variables and transmits its local information on the observations to a fusion center. At an unknown point in time, the distribution of the observations at all sensors changes. The objective is to detect the change in distribution as soon as possible, subject to a false alarm constraint. We consider minimax formulations for this problem and propose a new approach where transmissions are ordered and halted when sufficient information is accumulated at the fusion center. We show that the proposed approach can achieve the optimal performance equivalent to the centralized cumulative sum (CUSUM) algorithm while requiring fewer sensor transmissions. Numerical results for a shift in mean of independent and identically distributed Gaussian observations show significant communication savings for the case where the change seldom occurs which is frequently true in many important applications.**

***Index Terms*—Communication efficient, CUSUM, minimax, ordered transmissions, quickest change detection.**

I. Introduction

Significant attention has been devoted to distributed signal processing in sensor networks for both military and civilian applications, such as intrusion detection, secure surveillance, disaster prediction, internet of things, and health monitoring [1]. A problem of particular importance is the dynamic decision problem in sensor networks to detect the occurrence of a change. Such a detection problem can be modeled as a quickest change detection (QCD) problem, see [2]–[4] and references therein.

The centralized (single sensor) formulation of QCD is well studied [2] [5]–[7]. In the classical formulation of QCD, a decision maker monitoring the environment takes a sequence of observations whose distribution changes at an unknown point in time. The objective is to detect the change as quickly as possible subject to false alarm constraints. Based on the knowledge of the distribution of the change time, two formulations of QCD, Bayesian and minimax, are proposed, and the corresponding optimal solutions are described in [2]. In this paper, we focus on the minimax formulation where we model the change time as a deterministic but unknown positive integer and minimize the worst case detection delay subject to false alarm constraints.

Compared to the centralized QCD model, in the decentralized QCD setting there are multiple distributed sensors and a common decision maker at a fusion center. Further, each sensor takes some observations, processes them and then transmits a summary to the fusion center. At an unknown change time, the distributions of the observations at all sensors change simultaneously. Based on the information received, the fusion center would like to detect the change as soon as possible subject to false alarm constraints. This decentralized QCD problem has been well investigated [8] [9]. Typically, each sensor in the network carries its own limited energy sources. Hence, communication efficiency is an important topic in the decentralized QCD problem. Methods for improving communication efficiency include quantization and reducing the number of communications, see [4] [10] [11] and references therein.

One popular approach called censoring has been shown to be an effective method to improve communication efficiency where only highly informative data is transmitted in QCD [4] [11]. Censoring yields transmission savings but always increases detection delay, unlike the ordered transmission approach we introduce next.

The work in [12]–[14] employs the idea of ordered transmissions for a nonchanging hypothesis testing problem where sensors with the most informative observations transmit first. Transmissions can be halted when sufficient information is accumulated for the fusion center to decide which hypothesis is true. In this paper, we employ ordering in QCD to save communications without any impact on detection delay, and show that the average number of communications saved by our approach increases at least proportional to the number of sensors while detection delay is not affected provided a well-behaved distance measure between the distribution before the change and the distribution after the change is sufficiently large.

The remainder of the paper is organized as follows. In Section II, we present the mathematical formulation of the quickest change detection problem. Section III reviews the classical optimal CUSUM algorithm in the context of sensor networks. We propose and analyze the ordered transmission

The work is supported by the U. S. Army Research Laboratory and the U. S. Army Research Office under grant number W911NF-17-1-0331, the National Science Foundation under Grant ECCS-1744129, and a grant from the Commonwealth of Pennsylvania, Department of Community and Economic Development, through the Pennsylvania Infrastructure Technology Alliance (PITA).

978-1-7281-5478-7/20/$31.00 ©2020 IEEE

approach in Section IV. Numerical results are provided to illustrate the gains of our approach in Section V. We draw our conclusions in Section VI.

II. Problem Formulation

We consider a sensor network with K sensors and a fusion center. Sensor k for $k = 1, 2, .., K$ observes the sequence $\{X_{n,k}\}_{n\geq 1}$ with n being the time slot index. At an unknown time slot τ, the distribution of $\{X_{n,k}\}_{n\geq 1}$ for all the sensors changes from f_0 to f_1 where f_0 and f_1 are the known probability density functions (pdfs) before and after the change time, respectively. Throughout this paper we make the following assumptions.

Assumption 1: The random variable $X_{n,k}$ is independent across the time slot index n and sensor index k conditioned on the change time τ.

Assumption 2: The distributions of the observations at all sensors change simultaneously at the change time τ.

The objective is to detect a change as quickly as possible after the change occurs which implies the goal is to minimize detection delay if the change occurs. As long as no change is declared, the sensors will continue observing data. Without a prior on the distribution of the change time, we employ the constraint

$$\mathbb{E}_\infty(n) \geq \gamma \tag{1}$$

where $\mathbb{E}_\infty(n)$ is the average delay when the change does not occur, and γ is a pre-specified constant. If the change occurs, then there are two formulations to evaluate the detection delay. The first formulation employs the worst case average detection delay (WADD) defined in [5] as

$$\text{WADD}(n) = \sup_{\tau\geq 1} \text{ ess } \sup \mathbb{E}_\tau \left[(n-\tau)^+ | \mathcal{I}_{\tau-1}\right] \tag{2}$$

where $\text{ess}\sup X$ denotes essential supremum of X, $\mathbb{E}_\tau$ is the expectation when the change occurs at time τ, $(x)^+ \triangleq \max\{x, 0\}$, $\mathcal{I}_{\tau-1} \triangleq (X_{[1,\tau-1],1}, ..., X_{[1,\tau-1],K})$ denotes past global information at time slot τ, and $X_{[1,\tau-1],k} \triangleq (X_{1,k}, ..., X_{\tau-1,k})$ denotes past local information at sensor k. The other formulation employs the conditional average detection delay (CADD) which is defined as [7]

$$\text{CADD}(n) = \sup_{\tau\geq 1} \mathbb{E}_\tau[n-\tau | n \geq \tau]. \tag{3}$$

Thus, the quickest change detection problem in a minimax setting can be formulated as a constrained optimization problem

$$\begin{aligned} &\min_n \text{WADD}(n) \text{ or } \min_n \text{CADD}(n) \\ &s.t.\ \mathbb{E}_\infty(n) \geq \gamma. \end{aligned} \tag{4}$$

In the following sections, we will first review the likelihood ratio procedures to solve (4) and then propose a communication-efficient approach to achieve the same detection performance while saving communications.

III. Likelihood Ratio Method

In this section, we review the likelihood ratio procedure in the quickest change detection problem [3]. The QCD problem can be modeled as a hypothesis testing problem, given by [3]

$$\begin{aligned} &H_0 : \text{no change occurs} \\ &H_1 : \text{change occurs at some finite time slot } \tau. \end{aligned} \tag{5}$$

Note that when the change occurs, all sensors are assumed to be affected simultaneously as mentioned in *Assumption 2*.

The log-likelihood ratio (LLR) up to time n for (5) is [3]

$$\begin{aligned} &LLR_n \\ &= \max_{1\leq m\leq n} \log \frac{\prod_{i=1}^{m-1} f_0\left(X_{i,[1,K]}\right) \prod_{i=m}^{n} f_1\left(X_{i,[1,K]}\right)}{\prod_{i=1}^{n} f_0\left(X_{i,[1,K]}\right)} \quad (6) \\ &= \max_{1\leq m\leq n} \sum_{i=m}^{n}\sum_{k=1}^{K} \log \frac{f_1(X_{i,k})}{f_0(X_{i,k})} \quad (7) \end{aligned}$$

where $X_{i,[1,K]} \triangleq (X_{i,[1,1]}, ..., X_{i,[1,K]})$. The result in (7) is obtained since the observations between different sensors are independent.

The corresponding likelihood ratio procedure will raise an alarm at time

$$T_{LR}(b) = \inf\{n \geq 1 : LLR_n \geq b\} \tag{8}$$

where the constant b needs to be chosen properly to satisfy the false alarm constraint in (1). It turns out that the procedure in (8) is also called the CUSUM algorithm where the fusion center declares a change at

$$T_{CS}(b) = \inf\{n \geq 1 : W_n \geq b\} \tag{9}$$

where the CUSUM statistic W_n is defined as

$$W_n \triangleq \max_{1\leq m\leq n+1} \sum_{i=m}^{n}\sum_{k=1}^{K} \log \frac{f_1(X_{i,k})}{f_0(X_{i,k})}. \tag{10}$$

In this paper, if the upper index of any summation is smaller than the lower index, then we define the summation to be zero. A nice property of the non-negative CUSUM statistic W_n is that it can be computed recursively as

$$W_n = \max\left\{0, \quad W_{n-1} + \sum_{k=1}^{K} \log \frac{f_1\left(X_{n,k}\right)}{f_0\left(X_{n,k}\right)}\right\} \tag{11}$$

with $W_0 = 0$. The above recursion is powerful in online detection since it requires little memory. At each time slot n of the CUSUM algorithm, all sensors are required to transmit their LLRs and the fusion center computes W_n according to (11). We note that when we employ the CUSUM algorithm with $W_0 = 0$, WADD(n) in (2) and CADD(n) in (3) are equal to [2]

$$\text{WADD}(n) = \text{CADD}(n) = \mathbb{E}_1[n-1] \tag{12}$$

which means the worst case detection delay occurs at $\tau = 1$. The result in (12) makes the computation of CADD and WADD in simulations straightforward.

The power of the traditional CUSUM test is enhanced with a larger number of sensors, but this also requires more communications. Here, we employ ordering to reduce the number of transmissions while maintaining optimality.

IV. Ordered Transmissions

In this section, we propose a communication-efficient approach to detect a change in distribution via ordered transmissions. The idea is to order and halt the transmissions of the LLR. Employing the approach described in Algorithm 1, at each time slot the sensor with the largest LLR magnitude transmits first and the sensors with smaller LLR magnitudes possibly transmit later. By sometimes halting transmissions before all sensors have communicated their LLRs, communications can be saved while achieving the same detection delay as the optimal CUSUM algorithm where all sensors communicate their LLRs to the fusion center.

At each time slot n for $n = 1, 2, ...$, we order transmissions according to the magnitudes of the LLRs. We denote $\log \hat{L}_{n,1}$, ..., $\log \hat{L}_{n,K}$ as the nonincreasing rearrangement of the LLRs $\{\log(f_1(X_{n,k})/f_0(X_{n,k}))\}_{k=1}^{K}$ for which

$$\left|\log \hat{L}_{n,1}\right| \geq \left|\log \hat{L}_{n,2}\right| \geq ... \geq \left|\log \hat{L}_{n,K}\right|. \tag{13}$$

Our ordered-CUSUM transmission approach is summarized in Algorithm 1. If all transmission propagation delays are known and timing is synchronized, one can schedule all transmissions back to the fusion center so they arrive in the correct order. However, even with inaccurate estimates of propagation delays or imperfect synchronization, since the fusion center receives the values to be ordered, the fusion center can put them back in order correctly as long as the fusion center waits a short period related to the uncertainty, see [15].

Algorithm 1 ordered-CUSUM.

Input: a positive constant b.
Initialize: $n = 0$, $W_0 = 0$ and a positive number η^1.
1: **while** $W_n < b$ **do**
2: $n = n + 1$ (update time slot).
3: Sensor k computes its LLR for all $k = 1, ...K$.
4: **for** $k' = 1, 2, ..., K$ **do**
5: $\log \hat{L}_{n,k'}$ is transmitted to the fusion center after time equal to $\eta/|\log \hat{L}_{n,k'}|$.
6: The fusion center computes $W_{n,k'} = W_{n-1} + \sum_{k=1}^{k'} \log \hat{L}_{n,k}$ and $t_{n,L} = -(K-k')|\log \hat{L}_{n,k'}|$.
7: **if** $W_{n,k'} \leq t_{n,L}$ **then**
8: The fusion center decides $W_n = 0$.
9: break **for** loop and go to line 2.
10: **else**
11: Continue **for** loop.
12: **end if**
13: **end for**
14: The fusion center computes W_n according to (11).
15: **end while**
16: Declare the change occurs at time slot n.

An advantage of the ordered transmission approach can be summarized in the following theorem.

Theorem 1: Under *Assumptions 1 and 2*, if there exists a time slot n such that $W_{n,k'} \leq t_{n,L}$ with $k' < K$, then Algorithm 1 gives the same detection performance as the optimal full communication CUSUM algorithm (all sensors transmit their LLRs to the fusion center at each time slot), while using a smaller number of transmissions.

Proof: During time slot n (see Algorithm 1), when the fusion center receives a new LLR, it updates $t_{n,L}$ (a threshold) according to

$$t_{n,L} \triangleq -(K-k')\left|\log \hat{L}_{n,k'}\right| \tag{14}$$

and compares it with

$$W_{n,k'} \triangleq W_{n-1} + \sum_{k=1}^{k'} \log \hat{L}_{n,k}. \tag{15}$$

Note that the largest possible positive contribution from the sum of the sensor LLRs that have not yet transmitted is $-t_{n,L}$. If $W_{n,k'} \leq t_{n,L}$, then W_n has to be zero based on (11), regardless of the LLRs that have not yet been transmitted. Hence, even without receiving further transmissions, the fusion center can implement the optimum CUSUM algorithm at this time slot n. In fact, if $k' < K$ is true, then Algorithm 1 can at least save one transmission, thus Algorithm 1 will have a smaller average number of transmissions than the optimal CUSUM algorithm. Otherwise they employ the same number of transmissions. ■

One interesting question is whether the communication saving gains of the ordered transmission approach described in Algorithm 1 are large or not. To study this, we assume that there exists a distance measure denoted as s between f_0 and f_1. This distance measure s satisfies the following fairly mild assumption.

Assumption 3: For the hypothesis testing problem considered in (5), we assume that the probability $\Pr(\log(f_1(X_{n,k})/f_0(X_{n,k})) < 0 | n < \tau) \to 1$ as $s \to \infty$ and $\Pr(\log(f_1(X_{n,k})/f_0(X_{n,k})) > 0 | n \geq \tau) \to 1$ as $s \to \infty$ for all $k = 1, ..., K$.

Next we provide the following theorem on the average number of transmissions saved by Algorithm 1.

Theorem 2: Under *Assumptions 1–3*, consider the approach in Algorithm 1 for the quickest change detection problem in (5). With a sufficiently large s (distance measure), the average number of transmissions saved over the optimal CUSUM algorithm increases at least as fast as proportional to K while the detection delay is not affected.

Proof: The proof of *Theorem 2* is omitted due to space constraints. ■

Although *Theorem 2* requires a sufficiently large distance measure s, numerical results indicate that s does not have to be very large in some cases of interest to observe the event in *Theorem 2*.

[1] η is a global fixed constant. The value of η can be made as small as the system will allow.

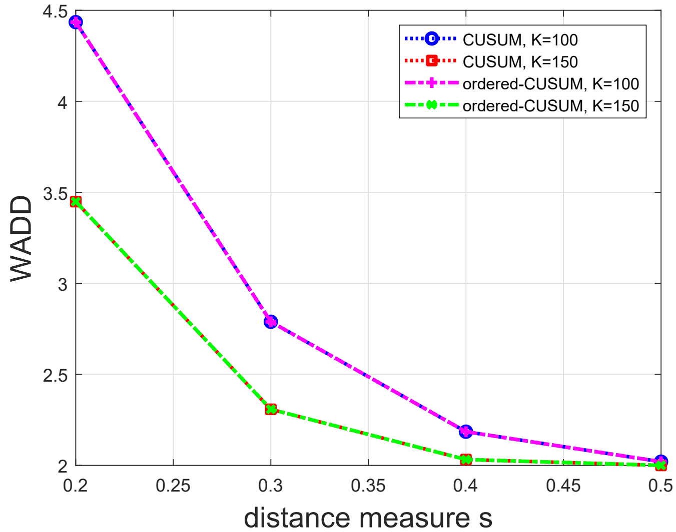

Fig. 1: The worst case average detection delay versus the distance measure for $K = 100$ and $K = 150$.

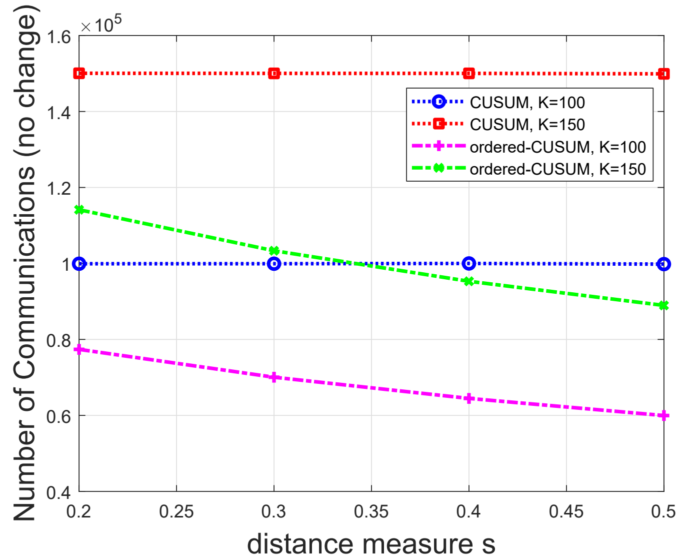

Fig. 2: Number of communications when the change does not occur versus the distance measure for $K = 100$ and $K = 150$.

V. Numerical Results

We first compare the performance of our approach with CUSUM as a function of the distance measure s between f_0 and f_1. Recall that CUSUM requires all sensors to transmit during each time slot. We plot WADD versus the distance measure in Fig. 1 for the parameters $\tau = 1$, $\gamma = 10^3$, $f_0 = \mathcal{N}(0, 1)$ and $f_1 = \mathcal{N}(s, 1)$. Fig. 1 illustrates that our approach has the same worst average detection delay as CUSUM. It also illustrates that WADD decreases as s is increased because the change time becomes easier to detect as s becomes larger. In Fig. 2, we plot the number of transmissions versus the distance measure s when the change does not occur. This indicates that our approach can save a significantly large number of communications compared to CUSUM. Fig. 1 and Fig. 2 together show that more sensors can lead to a smaller WADD at the cost of more transmissions which implies a basic tradeoff between the number of communications and the detection delay when the false alarm constraint is fixed.

Fig. 3 shows WADD as a function of the number of sensors K for our approach and CUSUM for $\gamma = 10^3$, $f_0 = \mathcal{N}(0, 1)$ and $f_1 = \mathcal{N}(s, 1)$ with $s = 0.2$ and $s = 0.3$. It indicates that our ordered-CUSUM algorithm provides the same detection performance as CUSUM. It also illustrates that WADD decreases as the number of sensors K is increased, which is reasonable since larger K implies more observations can be obtained per observation time slot which can help us detect the change more quickly. The result in Fig. 3 is consistent with our intuition that a change with a larger distance measure s is easier to detect which results in a smaller WADD. With the same parameter setting, Fig. 4 shows that the average number of communications saved by our approach increases approximately linearly with K for every value of s. In addition, Fig. 4 indicates that the rate of increase with K becomes faster when the distance measure s is increased.

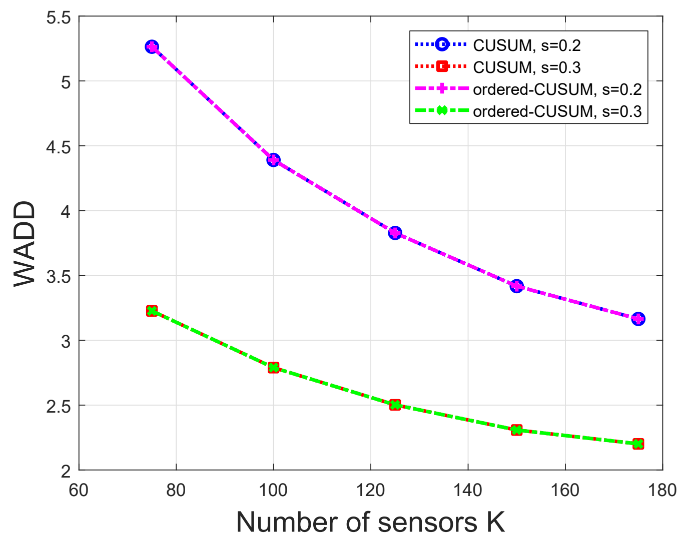

Fig. 3: The worst case average detection delay versus the number of sensors K for $s = 0.2$ and $s = 0.3$.

In Fig. 5, we compare WADD of our approach and CUSUM as a function of γ for $\tau = 1$, $f_0 = \mathcal{N}(0, 1)$ and $f_1 = \mathcal{N}(0.5, 1)$. This illustrates the basic tradeoff between WADD and γ. Fig. 6 also indicates that our approach can reduce the number of communications needed as compared to the number needed in the optimal CUSUM algorithm.

VI. Conclusion

In this paper, a new communication-efficient QCD approach has been proposed in sensor networks that reduces the number of transmissions without any impact on detection delay when compared to the classical CUSUM algorithm. In our approach, the sensors with more informative observations transmit their data to the fusion center earlier during each time slot. We have

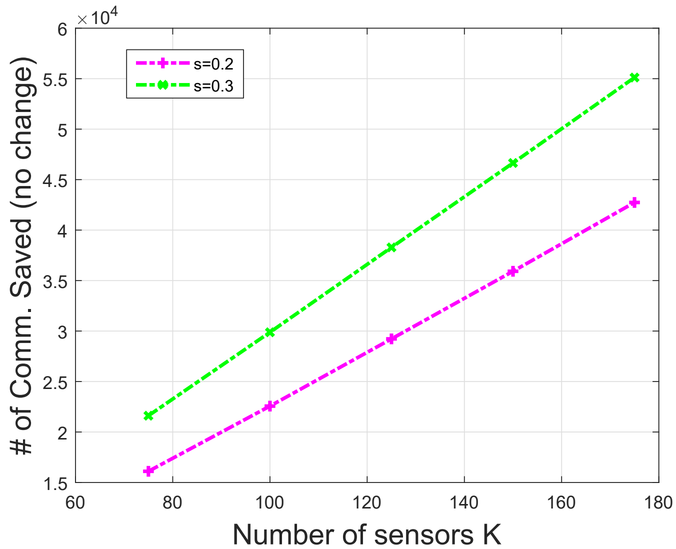

Fig. 4: Number of communication saved when the change does not occur versus the number of sensors K for $s = 0.2$ and $s = 0.3$.

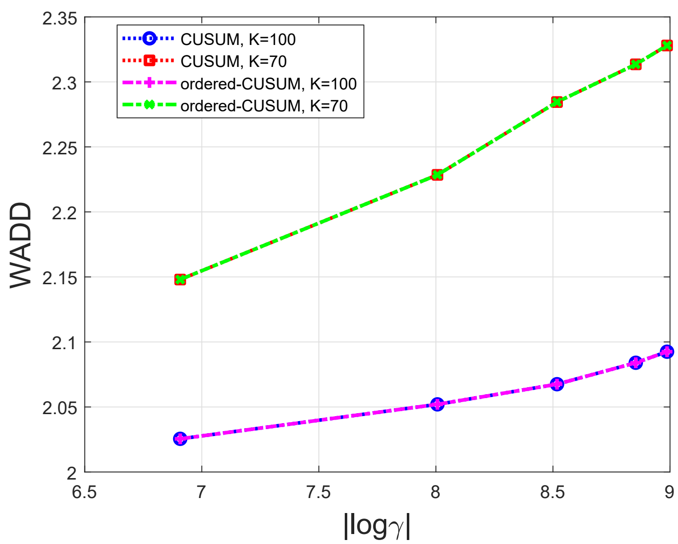

Fig. 5: The worst case average detection delay versus the false alarm constraint $|\log\gamma|$ for $K = 100$ and $K = 70$.

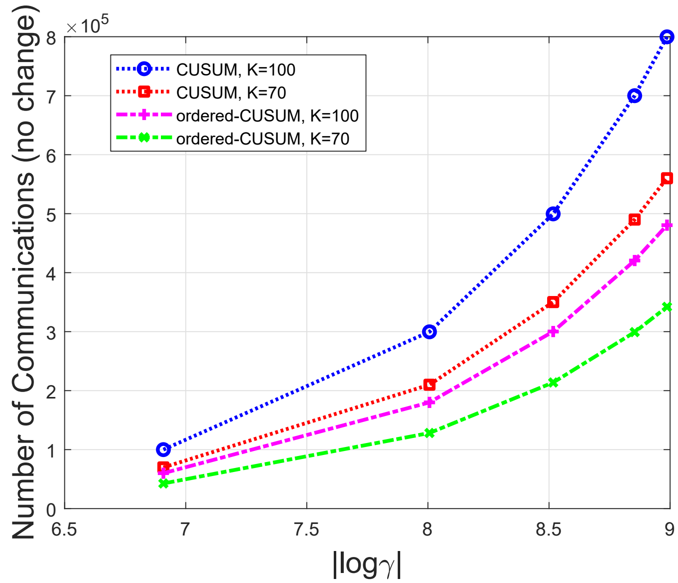

Fig. 6: Number of communications when the change does not occur versus the false alarm constraint $|\log\gamma|$ for $K = 100$ and $K = 70$.

shown that the average number of transmissions saved by our approach can increase at least as fast as proportional to the number of sensors when the distance measure between f_0 and f_1 is sufficiently large.

References

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: a survey," *Computer networks*, vol. 38, no. 4, pp. 393–422, 2002.

[2] V. V. Veeravalli and T. Banerjee, "Quickest change detection," in *Academic Press Library in Signal Processing*. Elsevier, 2014, vol. 3, pp. 209–255.

[3] Y. Mei, "Efficient scalable schemes for monitoring a large number of data streams," *Biometrika*, vol. 97, no. 2, pp. 419–433, 2010.

[4] T. Banerjee and V. V. Veeravalli, "Data-efficient quickest change detection in sensor networks," *IEEE Transactions on Signal Processing*, vol. 63, no. 14, pp. 3727–3735, 2015.

[5] G. Lorden *et al.*, "Procedures for reacting to a change in distribution," *The Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1897–1908, 1971.

[6] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Transactions on Information Theory*, vol. 44, no. 7, pp. 2917–2929, 1998.

[7] M. Pollak, "Optimal detection of a change in distribution," *The Annals of Statistics*, pp. 206–227, 1985.

[8] V. V. Veeravalli, "Decentralized quickest change detection," *IEEE Transactions on Information theory*, vol. 47, no. 4, pp. 1657–1665, 2001.

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

[10] A. G. Tartakovsky and V. V. Veeravalli, "Asymptotically optimal quickest change detection in distributed sensor systems," *Sequential Analysis*, vol. 27, no. 4, pp. 441–475, 2008.

[11] Y. Mei, "Quickest detection in censoring sensor networks," in *2011 IEEE International Symposium on Information Theory Proceedings*. IEEE, 2011, pp. 2148–2152.

[12] R. S. Blum and B. M. Sadler, "Energy efficient signal detection in sensor networks using ordered transmissions," *IEEE Transactions on Signal Processing*, vol. 56, no. 7, pp. 3229–3235, 2008.

[13] Y. Chen, B. M. Sadler, and R. S. Blum, "Ordered transmission for efficient wireless autonomy," in *2018 52nd Asilomar Conference on Signals, Systems, and Computers*. IEEE, 2018, pp. 1299–1303.

[14] Y. Chen, R. S. Blum, B. M. Sadler, and J. Zhang, "Testing the structure of a gaussian graphical model with reduced transmissions in a distributed setting," *IEEE Transactions on Signal Processing*, vol. 67, no. 20, pp. 5391–5401, 2019.

[15] R. S. Blum, "Ordering for estimation and optimization in energy efficient sensor networks," *IEEE Transactions on Signal Processing*, vol. 59, no. 6, pp. 2847–2856, 2011.