Space-Time CUSUM for Distributed Quickest Detection Using Randomly Spaced Sensors Along a Path

Daniel Egea-Roca*, Gonzalo Seco-Granados*, José A. López-Salcedo*, Sunwoo Kim†
* Dpt. of Telecommunications and Systems Engineering, Universitat Autònoma de Barcelona (UAB), Spain
† Department of Electronics Engineering, Hanyang University, Korea

***Abstract*—This work investigates the distributed quickest detection problem, where a set of sensors receive independent observations and send messages to a fusion center, which makes a final decision. We are interested in detecting an event as soon as possible even though the set of affected sensors is unknown. We consider a scenario where the sensors are randomly spaced along a path, and then the affected sensors are assumed to be consecutive. Based on the assumption that the affected sensors are consecutive, we propose a solution based on the detection of a transient change in the spatial domain (i.e. from different sensors). This is done by applying a double CUSUM to detect both the appearance and disappearance of the change in the space samples. Numerical results are presented showing the superior performance of our proposed solution, for different scenarios, with respect to other methods in the literature.**

I. Introduction

The problem of detecting a change in distribution of a stochastic process appears in a wide range of fields such as quality process control [1], signal detection in wireless communications [2], and spectrum sensing for cognitive radio [3], just to mention a few. Generally, these problems deal with the case where all information about the change is available at a single sensor, only. Specifically, this problem is known as sequential change detection or quickest detection. It has been solved under different criteria since the pioneering work by Page [4], which proposes the use of the so-called CUSUM algorithm to detect a sudden change in distribution as soon as possible. Nevertheless, there are situations where the information available for the decision process is decentralized (i.e. available at different sensors distributed in space).

Conventional decentralized or distributed detection approaches often consider that the change affects either a single sensor at a given time [5], or all existing sensors at the same time [6], [7]. There are other approaches, though, that consider that the sensors are gradually affected after some propagation time. In this latter case it is often assumed that sensors are placed across a sensor array, and therefore the change will always appear first at some predetermined sensor [8] or at any possible sensor [9]. Nevertheless, these contributions assume that all sensors will ultimately be affected after some large enough time, and they adopt a Bayesian approach with *a-priori*

This work has been partly supported by the National Research Foundation of Korea (NRF-2013R1A1A2062728) and the Spanish Grant TEC2014-53656-R.

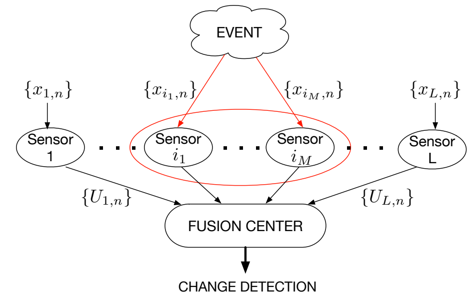

Fig. 1. Considered scheme for decentralized change detection

knowledge on the time at which the change may appear, as well as for the propagation time. In this paper, we consider a non-Bayesian approach (i.e. no prior knowledge) and we address a more general case with just an unknown subset of sensors may experience the change. This scenario was first considered in [10] and [11], and it is often the case in practice when sensors are randomly distributed over a geographical path, and just a few of these sensors become exposed to the event.

Specifically, the scenario considered herein is shown in Fig. 1, which considers a system with L sensors. At time n, an observation $x_{i,n}$ is made at sensor i and the message $U_{i,n}$, created from the available information, is sent to a central location called the fusion center. Then, based on the sensor messages, a decision about the change (event) is made at the fusion center. When the change to be detected happens, we assume that only $M < L$ consecutive sensors are affected, usually those closest to the event (i.e. $i_1, \dots, i_M$). We will take advantage here of the fact that affected sensors are consecutive, so that the problem can be formulated as a transient change detection problem in the spatial domain. Furthermore, apart from detecting the change, we will be also able to identify the unknown set of affected sensors. This scenario appears in many real-world applications in which the sensors are randomly spaced along a path, as when deployed from an aircraft, so that the affected sensors by the change can fairly be assumed to be consecutive.

The above considered scenario is of interest for applications such as environmental monitoring [12], where for instance a fire will affect the closest sensors to the fire, interference

ISBN 978-0-9928626-7-1 © EURASIP 2017

monitoring [13], where the affected sensors will be those close to the interference source or not obstructed by any obstacle, or intrusion detection in security monitoring systems [14]. Based on this observation our contribution is threefold. Firstly, we propose a CUSUM-based solution to the distributed quickest detection problem with unknown number of affected sensors. To do so we transform the problem to that of detecting a change within the spatial domain (i.e. from different sensors). This is in contrast to traditional solutions, which combine spatial samples to detect a change in the temporal domain. Secondly, based on the assumption that the affected sensors are consecutive, we formulate the spatial change detection problem as a transient change detection problem within the spatial samples. Thirdly, with the previous formulation, beyond detecting the temporal change we are able to localize those sensors affected by the change without any additional cost.

It is worth pointing out that the proposed solution is developed by assuming consecutive affected sensors, but as we will show with numerical results, it is also a good option when the sensors are not consecutive. The rest of this paper is organized as follows. Section II presents the signal model and introduces some preliminaries for the distributed quickest detection problem. Section III presents our proposed solution. Finally, Section IV shows numerical performance results, while Section V concludes the paper.

II. Signal Model and Preliminaries

Suppose we have L randomly spaced sensors along a path communicating with a fusion center (see Fig. 1), and each local sensor is taking measurements over time, namely $\{x_{i,n}\}_{n=1}^{\infty}$ for $i = 1, \ldots, L$. Initially, the measurements are under nominal conditions ($\mathcal{H}_0$) and $x_{i,n}$ is distributed according to the density $f_{i,0}$ at the i-th sensor. At some unknown time v, an event (change) occurs driving the measurement of an unknown subset of M sensors to different conditions ($\mathcal{H}_1$), so that if the i-th sensor is affected, the density function of its local observation $x_{i,n}$ changes from $f_{i,0}$ to $f_{i,1}$ at time v. It is important to note that this notation allows each sensor to have different distributions. The subset of M affected sensors will be denoted by $\mathcal{I}_M = \{i_1, \ldots, i_M\} \in \mathcal{I}$, with $\mathcal{I}$ the set of all possible subsets of affected sensors, with a total of $|\mathcal{I}|$ elements. Both M and $\mathcal{I}_M$ are assumed to be unknown. The problem is to detect the presence of a change affecting M of the L available sensors by taking advantage of the diversity given by the information provided by the $L > 1$ sensors, in contrast to the traditional problem where we only have one sensor (i.e. $L = 1$).

In change detection problems, the time at which the change is detected is determined by means of a stopping time T, which is given by some detection rule based on the measurements $\{x_{i,n}\}_{n\geq 1}$, for $i = 1, \ldots, L$. We assume these measurements to be independent across time and space. In particular, we want to find a stopping time T, which will determine a change as quickly as an event occurs, with some restriction in false alarms. Let us denote by $\mathbb{E}_v^{(\mathcal{I}_M)}$ the expectation when a change appears at time v and sensors $i \in \mathcal{I}_M$, and denote by $\mathbb{E}_\infty$ the same when there are no changes. Then, the previous criterion can be formulated as the minimization of [15]

$$D(T) \doteq \sup_{v\geq 1} \operatorname{essup} \mathbb{E}_v^{(\mathcal{I}_M)} \left[(T - v + 1)^+ \,|\boldsymbol{X}_{v-1} \right], \qquad (1)$$

subject to the global false alarm constraint

$$N(T) \doteq \mathbb{E}_\infty(T) \geq \gamma, \qquad (2)$$

where $\boldsymbol{X}_{v-1} \doteq [\boldsymbol{x}_{1,[1,v-1]} \ldots \boldsymbol{x}_{L,[1,v-1]}]$ denotes past global information (i.e. at the fusion center) at time v, with $\boldsymbol{x}_{i,[1,v-1]} = [x_{i,1}, \ldots, x_{i,v-1}]^\top$ the past information at the i-th sensor; and $\gamma > 0$ is a desired finite constant for the global false alarm rate. Moreover, $(x)^+ = \max(0, x)$ and essup denotes the essential supremum.

The problem of monitoring a single sensor (i.e. $L = 1$) has been widely addressed in the quickest detection literature [16], [17]. For this problem, the optimal procedure for the non-Bayesian framework is the CUSUM algorithm, which can recursively be computed at the i-th sensor, at time n, as

$$g_{i,n} = (g_{i,n-1} + \mathrm{LLR}_{i,n})^+, \qquad (3)$$

with $g_{i,0} = 0$ and $\mathrm{LLR}_{i,n} \doteq \ln(f_{i,1}(x_{i,n})/f_{i,0}(x_{i,n}))$ the log-likelihood ratio (LLR) of the observation $x_{i,n}$. The global problem (i.e. $L > 1$), though, in addition includes the nuisance parameter $\mathcal{I}_M \in \mathcal{I}$, and then a generalized LLR should be used instead, leading to the following stopping time [10]:

$$T_{\mathrm{LR}} = \inf\left\{ n \geq 1 : \max_{\mathcal{I}_M \in \mathcal{I}} g_n^{(\mathcal{I}_M)} \geq h \right\}, \qquad (4)$$

with $h > 0$ the detection threshold chosen to satisfy the false alarm constraint in (2) and $g_n^{(\mathcal{I}_M)}$ the CUSUM algorithm applied to the subset of sensors $i \in \mathcal{I}_M$, defined as

$$g_n^{(\mathcal{I}_M)} \doteq \left(g_{n-1}^{(\mathcal{I}_M)} + \sum_{j=1}^{M} \mathrm{LLR}_{i_j,l} \right)^+. \qquad (5)$$

Based on this recursive form, T_{LR} can be implemented as simultaneously evaluating a total of $|\mathcal{I}|$ CUSUMs with each CUSUM considering a specific subset of affected sensors $\mathcal{I}_M$.

The implementation in (4) is very suitable when some prior information about the number of affected sensors is available. This is the case, for instance, when we know that only one sensor is affected when the event appears (i.e. $M = 1$). In this case, from (4), we have the so-called Max-CUSUM

$$T_{\mathrm{mx}} \doteq \inf\left\{ n \geq 1 : \max_{1\leq i\leq L} g_{i,n} \geq h \right\}. \qquad (6)$$

That is, we raise an alarm at the global level as soon as the maximum of the local CUSUMs is above the detection threshold h [5]. Unfortunately, when the number M of affected sensors is completely unknown, the implementation in (4) is infeasible, as it requires to evaluate too many subsets. Alternative efficient approaches were presented in [11]. One of these is the henceforth referred to as Hard-CUSUM, which proposes to raise a global alarm at time

$$T_{\mathrm{hd}} \doteq \inf\left\{ n \geq 1 : \sum_{i=1}^{L} g_{i,n} \mathbb{1}_{\{g_{i,n}\geq b\}} \geq h \right\}, \qquad (7)$$

ISBN 978-0-9928626-7-1

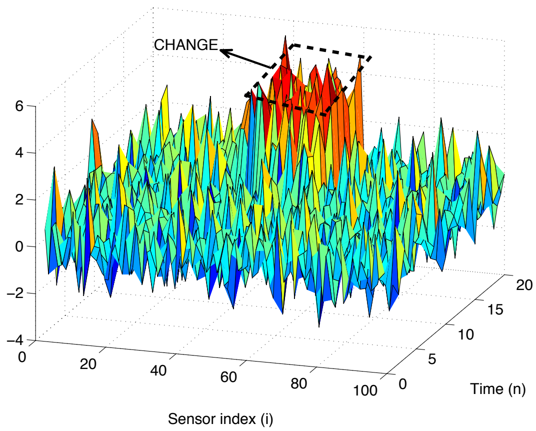

Fig. 2. Space-Time plane with $L = 100$ sensors and a Gaussian mean change at time $v = 10$ of $M = 20$ appearing at sensors $i \in \{40, 41, \ldots, 59\}$.

with $b \geq 0$ the local threshold and $\mathbb{1}_{\{A\}}$ the indicator function of the event A, which is 1 if the event is true and 0 otherwise. In other words, the fusion center raises an alarm as soon as the sum of those local CUSUMs above the local threshold b is too large. Results in [11] show an improvement of the detection delay for a fixed false alarm rate for specific values of thresholds (i.e. b and h) of the Hard-CUSUM with respect to other approaches available in the literature.

III. Space-Time Double-CUSUM (STD-CUSUM)

Our proposed scheme for distributed quickest detection is based on the observation that physical phenomena typically affect a group of neighboring sensors, where the $M \leq L$ affected sensors can fairly be assumed to be consecutive. This can equivalently be understood as a transient change detection problem in the spatial domain. That is, for a given time instant, the spatial samples coming from each sensor can be processed sequentially as a conventional sequential detector does in the time domain. Under $\mathcal{H}_0$ all spatial measurements follow certain distribution $f_{i,0}$, whereas after the change, M of these measurements in the space domain experience a change in their distribution (i.e. from $f_{i,0}$ to $f_{i,1}$). This behavior is depicted in Fig. 2, where the so-called active space-time plane is highlighted. The illustration assumes a change in the mean of a Gaussian distribution at time $v = 10$ of $M = 20$ sensors (i.e. $i \in \{40, 41, \ldots, 59\}$) out of a total of $L = 100$ sensors.

If we take a look at one of the spatial realizations under $\mathcal{H}_1$ in Fig. 2 (i.e. for time instants $n \geq 10$) we see how the mean of some sensors measurements abruptly changes. Thus, for a given time instant, the problem can be thought as that of change detection using spatial samples. With this behavior in mind, we can think on applying the CUSUM algorithm on the spatial domain, and then use only the information of those sensors declared to be affected for making a global decision. Indeed, this is similar to the behavior of the Hard-CUSUM presented in (7), which only uses the information of those sensors whose local CUSUM has detected a change.

Nevertheless, the local CUSUMs used for the Hard-CUSUM are not exploiting the space dimension, in the sense that they only consider the corresponding local sensor, thus disregarding spatial samples from neighboring sensors. As a result, it may take some time until all the local CUSUMs of the affected sensors are above the local threshold, thus incurring in a penalty in terms of detection delay. This is so because some measurements may seem not to be affected, even if they actually are. Our proposed approach is based on the idea of exploiting the space dimension in such a way that we have some statistic that increases along the space dimension as long as the sensor measurement seems to be affected and it is restarted just when the change in the space dimension disappears. Thus, the statistic will be above the local threshold even if some sensor measurements seem not to be affected.

A. STD-CUSUM detection rule

Assuming that the $M < L$ affected sensors are consecutive and unknown, we can model the problem of detecting the change in the space dimension as a transient change detection problem, with transient length equal to M samples. In fact, L is fixed, so this detection problem can be seen as an off-line problem. In other words, at each time we have L samples and we want to decide if there is some change within these samples. This problem is analyzed in [18], but it leads to a complex solution due to the large number of trials we need to evaluate, similarly to what happened in (4). Our contribution is intended to circumvent this limitation by using a simple and computationally efficient space-time CUSUM detector.

Actually, in order to have the desired behavior (i.e. detecting the transient change in the space dimension) we need to use two different CUSUMs, one for detecting the change from $\mathcal{H}_0$ to $\mathcal{H}_1$ and another for detecting the change from $\mathcal{H}_1$ to $\mathcal{H}_0$. Let $g_{i,n}$ denote the local CUSUM for detecting the former change, defined as in (3), and let $\tilde{g}_{i,n}$ denote the local CUSUM of the i-th sensor, but for detecting the disappearance of the change, defined as

$$\tilde{g}_{i,n} \doteq \left(\tilde{g}_{i,n-1} - \mathrm{LLR}_{i,n}\right)^+ . \tag{8}$$

Furthermore, let us define the CUSUM for detecting the appearance and disappearance of the change in the space dimension for a given time instant n, respectively, as

$$\begin{aligned} G_{i,n} &\doteq \left(G_{i-1,n} + g_{i,n-1} + \mathrm{LLR}_{i,n}\right)^+ \mathbb{1}_{\{\widetilde{G}_{i,n-1} < b\}} \\ \widetilde{G}_{i,n} &\doteq \left(\widetilde{G}_{i-1,n} + \tilde{g}_{i,n-1} - \mathrm{LLR}_{i,n}\right)^+ \mathbb{1}_{\{G_{i,n} \geq b\}} \end{aligned}, \tag{9}$$

with $G_{i,0} = 0$ and $\widetilde{G}_{i,0} = 0$ for $i = 1, \ldots, L$. It is worth pointing out that this recursive form includes both time (with $g_{i,n-1}$) and space (with $G_{i-1,n}$) integration, and it is here where our contribution exploit both time and space dimensions. With this framework, we define the Space-Time Double-CUSUM (STD-CUSUM) as

$$g_{i,n}^{(\mathrm{STD})} \doteq G_{i,n} \mathbb{1}_{\{\widetilde{G}_{i,n} < b\}} . \tag{10}$$

ISBN 978-0-9928626-7-1 © EURASIP 2017

TABLE I
DETECTION DELAYS WITH $L = 100$ AND $L = 20$ FOR DIFFERENT VALUES OF M CONSECUTIVE AFFECTED SENSORS

γ	Detection Scheme	$L = 100$					Detection Scheme	$L = 20$			
		# sensors affected M						# sensors affected M			
		80	20	10	5	1		16	7	2	1
10^4	$T_{\mathrm{hd}}(h = 106.4, b = 0.5)$	7.3	20.2	33.8	56.1	195.5	$T_{\mathrm{hd}}(h = 32.9, b = 0.5)$	12.1	24.4	69.9	122.8
	$T_{\mathrm{hd}}(h = 62.3, b = 2.3)$	9.2	19.7	31.9	53.7	191.6	$T_{\mathrm{hd}}(h = 25, b = 2.3)$	13.7	24.1	70.1	126.8
	$T_{\mathrm{hd}}(h = 29.7, b = 4.6)$	14.2	21.9	29.9	43.3	152.6	$T_{\mathrm{hd}}(h = 16.5, b = 4.6)$	19.3	27.1	59.8	112.4
	$T_{\mathrm{sc}}(h = 76, b = 0.25)$	5.5	17.2	29.8	50.5	164.4	$T_{\mathrm{sc}}(h = 24.2, b = 0.25)$	10.4	20.9	57.5	98.3
	$T_{\mathrm{sc}}(h = 63.5, b = 2.6)$	4.9	14.5	24.7	43.4	170.1	$T_{\mathrm{sc}}(h = 22.1, b = 2.6)$	8.9	18.3	55.5	102.9
	$T_{\mathrm{sc}}(h = 47, b = 5)$	**3.8**	**11.7**	**20.1**	**34.6**	142.8	$T_{\mathrm{sc}}(h = 18, b = 5)$	**7.9**	**16.2**	**50.4**	96.0
	$T_{\mathrm{mx}}(h = 11.12)$	32.7	39.9	45.2	52.3	**85.5**	$T_{\mathrm{mx}}(h = 9.75)$	34.16	41.36	57.89	**74.42**

B. Identification of Affected Sensors

The introduction of the STD-CUSUM leads to the desired behavior, i.e., we have a CUSUM-based statistic that is able to detect the appearance and disappearance of the change in the space dimension. Doing so we are exploiting the space dimension in such a way that, at time n, the STD-CUSUM increases along the space dimension as long as the sensor measurements seems to be affected by the change, so that it may still be above the local threshold even if some of the measurements seem not to be affected. With this formulation, we are able to localize those sensors affected by the change, even if they are not consecutive, at time n, as

$$\hat{\mathcal{I}}_M(n) = \left\{\hat{i}_1, \dots, \hat{i}_M\right\} = \left\{i : g_{i,n}^{(\mathrm{STD})} \geq b\right\}. \quad (11)$$

C. Global STD-CUSUM Stopping Time

Let us now focus on the global detection carried out in the fusion center. The idea is, at time n, to apply the STD-CUSUM for detecting those sensors that are affected by the change, so that we can use the expression in (4) but substituting $\mathcal{I}_M$ by their estimates at time n given by (11). Doing so, we get rid off the maximization over $\mathcal{I}_M \in \mathcal{I}$ in (4), drastically improving the computational burden. To do so, we suppose that, at time n, each sensor i sends its LLR (i.e. $U_{i,n} = \mathrm{LLR}_{i,n}$), and then the fusion center calculates the STD-CUSUM in (10) so that it can raise a global alarm at time

$$T_{\mathrm{sc}} = \inf\left\{n \geq 1 : \sum_{i=1}^{L} g_{i,n} \mathbb{1}_{\left\{g_{i,n}^{(\mathrm{STD})} \geq b\right\}} \geq h\right\}. \quad (12)$$

Both thresholds b and h should be numerically fixed according to the desired mean time between false alarms.

IV. NUMERICAL RESULTS

This section is intended to provide numerical simulations to illustrate the performance of our proposed scheme and compare it to other schemes in the literature. We consider two scenarios, one using $L = 100$ and another with $L = 20$ iid sensors. For the sake of clarity and illustration we suppose iid Gaussian observations at each sensor with mean 0 and variance 1 before the change (i.e. $\mathcal{H}_0$) and with mean 0.5 and variance 1 after the change (i.e. $\mathcal{H}_1$). In practice, these observations can be any detection metric related with the event of interest, like temperature, power, integrity metric, ... The results presented next under case #1 consider the affected sensors to be consecutive and case #2 considers the case when all the affected sensors are not consecutive.

In particular, the presented results compare the performance of the STD-CUSUM in terms of detection delay, with a fixed value for the time between false alarm, with the Hard- and Max-CUSUM. For our proposed detector (STD-CUSUM) we use different values of the local threshold b and we find a value h, from numerical simulations, such that $N(T) \approx \gamma = 10^4$. The selected local thresholds are chosen so that we can make a wide comparison with those selected in [11] for the Hard-CUSUM. Next, using the obtained threshold h, we obtain the detection delays of the analyzed algorithms with different number of affected sensors M. This is done for both scenarios with $L = 100$ and $L = 20$. For the other schemes we use the same values for the thresholds as in [11], and then the obtained results are the same.

A. CASE 1: Consecutive affected sensors

In this section, we consider the case when the affected sensors are consecutive, as it would be the case in a sensor array. In Table I we can see how for $L = 100$ the STD-CUSUM is in general better than the other analyzed schemes. In particular, for $M > 5$ the STD-CUSUM is better than the Hard-CUSUM for any of the used thresholds. For lower M only with the configuration using the larger local threshold ($b = 5$) in the STD-CUSUM is better than the Hard-CUSUM (for any of its configurations), giving a gain of 10 samples in terms of detection delay. Regarding the Max-CUSUM, the STD-CUSUM outperforms it for any value of M, except for $M = 1$. It is interesting to see that, in contrast to the Hard-CUSUM, the larger the local threshold b the better performance of our proposed scheme in terms of detection delay for a fixed false alarm rate for all the simulated values of M. So, for the STD-CUSUM it is of interest to fix a large local threshold b independently of the number of affected sensors M. This is a very good advantage since M is unknown and then with our proposed scheme we make sure it will work well for any value of M.

Similar results are obtained for $L = 20$ in Table I. We see how the STD-CUSUM is better than the Hard-CUSUM for any of the used configurations for all M, confirming the outperformance of our proposed solution for different configurations of the local threshold b even for low values of L. Finally, it is worth pointing out that the larger M the

ISBN 978-0-9928626-7-1

TABLE II
DETECTION DELAYS FOR NON CONSECUTIVE AFFECTED SENSORS

γ	Detection Scheme	$L = 20$			
		# sensors affected M			
		16	7	4	2
10^4	$T_{\rm sc}(h = 24.2, b = 0.25)$	10.3	20.5	32.3	57.6
	$T_{\rm sc}(h = 22.1, b = 2.6)$	9.4	18.8	30.3	56.2
	$T_{\rm sc}(h = 18, b = 5)$	**8.7**	**18.7**	**28.7**	**51.4**
	Detection Scheme	$L = 100$			
		# sensors affected M			
		80	20	10	5
	$T_{\rm sc}(h = 76, b = 0.25)$	5.6	16.5	27.9	47.6
	$T_{\rm sc}(h = 63.5, b = 2.6)$	5.2	15.1	25.3	44.1
	$T_{\rm sc}(h = 47, b = 5)$	**4.1**	**14.0**	**22.8**	**37.4**

larger is the improvement of the STD-CUSUM with respect to the Hard-CUSUM, which is intuitive since the STD-CUSUM increases as the transient change in the space dimension is larger. So, we can conclude that the STD-CUSUM is good for large L and M, but it still is a good option otherwise.

B. CASE 2: Non consecutive affected sensors

Now, we show, in Table II, the obtained results when the M affected sensors are not consecutive, as it would be the case in a randomly deployed sensor network. We only show the results for the STD-CUSUM, since the results for the Hard- and Max-CUSUM are the same as in the Table I. Firstly, we see how the detection delay in this case, for the same configuration of thresholds, is slightly degraded compared to the case of consecutive affected sensors (see Table I). This is in line with the fact that the STD-CUSUM works better when the number of consecutive affected sensor is greater.

Secondly, we see that the loss of performance of the STD-CUSUM with respect the case when all the M sensors are consecutive is greater for lower values of M. This can be seen for instance for $L = 100$ when $M = 80$ the detection delay is similar for both cases (i.e. M consecutive or not consecutive sensors), whereas when $M = 5$ we have a difference of 3 samples. This is due to the fact that for large M there will likely be some subset of affected sensors that are consecutive, and then the loss of performance is negligible, whereas for small M it may happen that no consecutive sensors are affected, thus incurring in some degradation. Nevertheless, comparing with the results in Table I for the Hard-CUSUM we see that the STD-CUSUM still gives lower detection delay, so we can conclude that the STD-CUSUM outperforms the Hard-CUSUM even when the affected sensors are not consecutive.

V. CONCLUSIONS

This work has studied the problem of distributed quickest detection when the change to be detected is only visible for some unknown number of sensors. This was motivated by the fact that many physical phenomena typically affect just a group of neighboring sensors, we have considered a scenario where the affected sensors are consecutive. Based on this assumption we have proposed a solution working with a spatial change detection problem. In contrast to previous solutions, which only exploit the temporal dimension or make use of a-priori information, our proposed solution fully takes advantage of the spatial dimension by formulating the problem as that of spatial change detection. Furthermore, no a-priori information is considered regarding neither the change point or the number of affected sensors. With the proposed solution, beyond detecting the change, we are able to localize the affected sensors. Numerical results have shown that using the information of the localized affected sensors at the fusion center improves the global detection delay with respect to previous contributions assuming an unknown number of sensors, too. This has been shown to be true even when the affected sensors are not consecutive, which makes the proposed technique applicable regardless of the sensor deployment.

REFERENCES

[1] S. P. Mertikas, “Automatic and online detection of small but persistent shifts in GPS station coordinates by statistical process control,” *GPS Solutions*, vol. 5, no. 1, pp. 39–50, 2001.

[2] T. Oskiper and H. V. Poor, “Quickest detection of a random signal in background noise using a sensor array,” *EURASIP Journal on Applied Signal Processing*, vol. 2005, pp. 13–24, 2005.

[3] E. Axell, G. Leus, E. G. Larsson, and H. V. Poor, “Spectrum sensing for cognitive radio: State-of-the-art and recent advances,” *IEEE Signal Processing Magazine*, vol. 29, no. 3, pp. 101–116, 2012.

[4] E. S. Page, “Continuous inspection schemes,” *Biometrika*, pp. 100–115, 1954.

[5] A. G. Tartakovsky, B. L. Rozovskii, R. B. Blazeka, and H. Kim, “Detection of intrusions in information systems by sequential change-point methods,” *Statistical Methodology*, vol. 3, no. 3, pp. 252–293, 2006.

[6] 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.

[7] S. Zarrin and T. J. Lim, “Cooperative quickest spectrum sensing in cognitive radios with unknown parameters,” in *Proc. IEEE Global Telecommunications Conference (GLOBECOM)*, 2009, pp. 1–6.

[8] V. Raghavan and V. V. Veeravalli, “Quickest detection of a change process across a sensor array,” in *Proc. International Conference on Information Fusion*, 2008, pp. 1–8.

[9] D. Li, L. Lai, and S. Cui, “Quickest change detection and identification across a sensor array,” in *Proc. IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2013, pp. 145–148.

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

[11] ——, “Quickest detection in censoring sensor networks,” in *Proc. IEEE Intl. Symp. on Information Theory (ISIT)*, 2011, pp. 2148–2152.

[12] S. N. Simić and S. Sastry, “Distributed environmental monitoring using random sensor networks,” in *Information Processing in Sensor Networks*. Springer, 2003, pp. 582–592.

[13] J. Lindstrom, D. M. Akos, O. Isoz, and M. Junered, “GNSS interference detection and localization using a network of low cost front-end modules,” in *Proc. ION GNSS*, 2007.

[14] S. R. Snapp *et al.*, “DIDS (distributed intrusion detection system)-motivation, architecture, and an early prototype,” in *Proc. National Computer Security Conference*. Washington, DC, 1991, pp. 167–176.

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

[16] H. V. Poor and O. Hadjiliadis, *Quickest Detection*. Cambridge University Press, 2009.

[17] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*. Englewood Cliffs: Prentice Hall, 1993.

[18] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume 2: Detection Theory*, 1998.

ISBN 978-0-9928626-7-1 © EURASIP 2017