KES2014

Online detection of shutdown periods in chemical plants: A case study

Manuel Martin Salvador[a,*], Bogdan Gabrys[a], Indrė Žliobaitė[b,c]

[a]*Smart Technology Research Centre, Bournemouth University, United Kingdom*
[b]*Department of Information and Computer Science, Aalto University, Finland*
[c]*Helsinki Institute for Information Technology, Finland*

Abstract

In process industry, chemical processes are controlled and monitored by using readings from multiple physical sensors across the plants. Such physical sensors are also supplemented by soft sensors, i.e. adaptive predictive models, which are often used for computing hard-to-measure variables of the process. For soft sensors to work well and adapt to changing operating conditions they need to be provided with relevant data. As production plants are regularly stopped, data instances generated during shutdown periods have to be identified to avoid updating these predictive models with wrong data. We present a case study concerned with a large chemical plant operation over a 2 years period. The task is to robustly and accurately identify the shutdown periods even in case of multiple sensor failures. State-of-the-art methods were evaluated using the first half of the dataset for calibration purposes and the other half for measuring the performance. Results show that shutdowns (i.e. sudden changes) can be quickly detected in any case but the detection delay of startups (i.e. gradual changes) is directly related with the choice of a window size.
© 2014 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of KES International.

Keywords: change-point detection; online detection; shutdown periods; data streams; case study

1. Introduction and motivation

Chemical production processes are controlled both manually and automatically to achieve a desired product quality. Physical sensors around the plants provide data streams such as temperature, pressure, humidity or flow, that are essential to monitor plant operation in real time.

Historical data collected from plant sensors may be huge and can be used for different tasks like reporting (e.g. monthly productivity of a plant) or building data-driven soft sensors (i.e. predictive models to support decision making or as cover for physical sensors).

Soft sensors are usually built for predicting hard-to-measure values in real time [1]. This building process involves data cleaning such as the removal of data from shutdown periods. We are studying the automation of shutdown identification for speeding up these tasks.

---

* Corresponding author.
   E-mail address: msalvador@bournemouth.ac.uk

Adaptive soft sensors are often updated with new data for capturing well the underlying behaviour of the process that evolves over time[2]. If the data from shutdowns are not removed, the predictive model can adapt to an undesirable process state. Also, during these inactive periods of production the predicted values are not meaningful from the process point of view. Therefore, an online method is necessary to automatically detect shutdowns in order to stop model adaptation.

The decision to stop a plant is usually taken by an human operator. Despite the fact that shutdowns can be scheduled for a year ahead, they may vary depending on the operating conditions (e.g. if heat suddenly increases to a dangerous level the plant has to be stopped for safety reasons).

The problem is also challenging because there is usually no single variable that can accurately and unambiguously identify the operating state of the chemical process. The solution starts from monitoring sensor values to detect changes in the process. However, not all sensors react in the same way to a shutdown. Expert support is usually needed to select the relevant sensors to monitor.

In addition, physical sensors can fail and as a result detection may be interrupted. Therefore, a shutdown detection method that works by monitoring only one sensor is unreliable in an industrial environment. It is essential to build robust methods that are able to monitor and combine several sensors at the same time.

As part of the ongoing evaluation studies aiming at identifying robust pre-processing methods and ultimately developing of the automated data preprocessing framework (see Žliobaitė and Gabrys work[3]), in this paper we have performed a comparative analysis of the state-of-the-art change-detection methods for a challenging case study.

The data of this case study has been provided by a chemical company and it has been collected from a chemical plant over a period of 2 years of operation. Data from 11 flow sensors have been aligned by time-stamp in order to form instances. The location of sensors in the plant causes delays between sensor signals during both shutdowns and startups. These delays make the detection more challenging. Furthermore, the annotation of the shutdown periods as ground truth for evaluation purposes has not been trivial.

The paper is organized as follows: Section 2 presents the problem setting; in the Section 3 we review the related literature and we describe a new method based on control charts; Section 4 contains the dataset description and the experimentation; finally, we conclude the paper in Section 5.

2. Problem setting

A shutdown is a period of time $[t_\alpha, t_\omega]$ during which a process is inactive but its duration is not defined a priori. Process operation is monitored using a group of sensors. The visualization of the values of some sensors over time makes possible a clear distinction of groups of out-of-control values that represent shutdown periods (see Figure 1). However, other sensors are not showing any change during those periods. The selection of relevant sensors for shutdown detection is not straightforward and usually domain experts select them manually.

We focus our work in an online scenario where data from sensors are continuously arriving to the system at fixed time intervals (e.g. every second). Let $X_t = (x_{1,t}, ..., x_{N,t})$ be the vector of N sensor values at time t. The distribution of a relevant sensor (i.e. sensitive to shutdowns) is given by a Gaussian mixture model such as

$$x_{n,t} \sim \begin{cases} \mathcal{N}(\mu_{n,0}, {\sigma_{n,0}}^2), & \text{if } t \notin [t_\alpha, t_\omega] \\ \mathcal{N}(\mu_{n,1}, {\sigma_{n,1}}^2), & \text{if } t \in [t_\alpha, t_\omega] \end{cases} \tag{1}$$

The formulation of a change-point in a data stream is usually given by the stopping rule

$$T = \inf\{t : s_t(X_t) \geq \tau\} \tag{2}$$

where $s_t(X_t)$ is the statistic computed over the input data and τ is the detection threshold.

Since sensors are physically located in different places of a plant, they will perceive the change of the process state at different moments. As a consequence, when a shutdown takes place there is a time interval $[t_\alpha, t_\beta]$ in which some parts of the plant are still working while others are stopped. The same situation happens during a startup. In this case, the time interval is $[t_\psi, t_\omega]$. Figure 2 shows these time intervals in three flow sensor signals during both events. Usually, $t_\beta - t_\alpha \ll t_\omega - t_\psi$. That is, shutdowns are characterized by sudden changes while startups present gradual changes in sensor values.
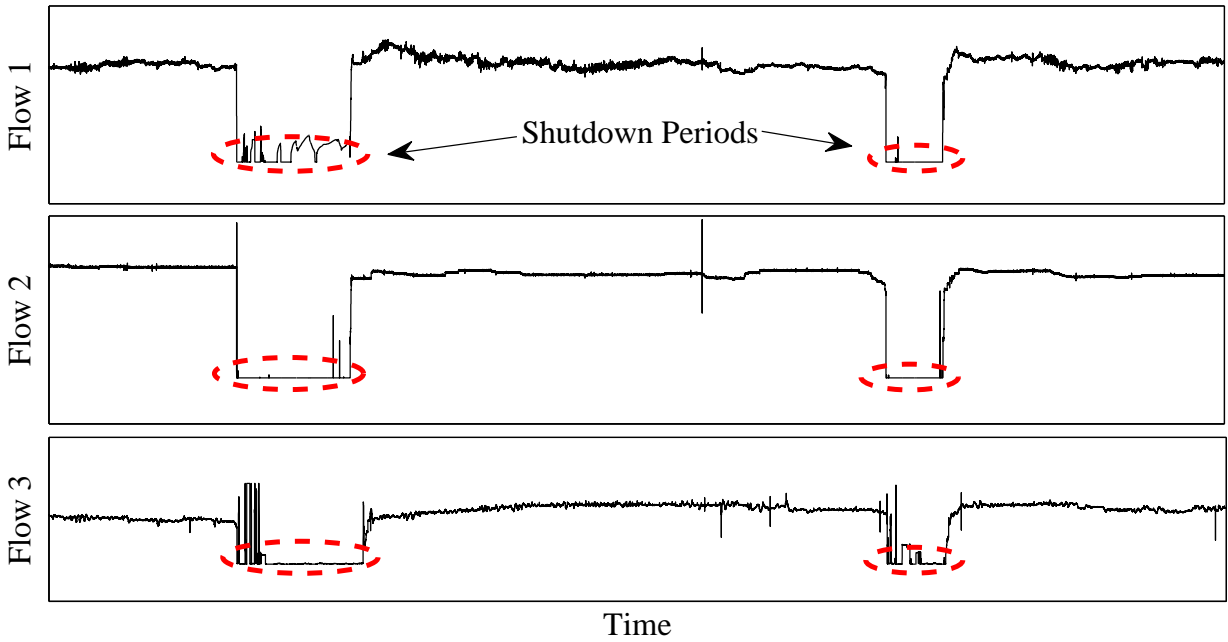
Fig. 1. This plot shows how sensors values suddenly drop when a shutdown starts. After a period of inactivity, the values suddenly increase again when shutdown ends.
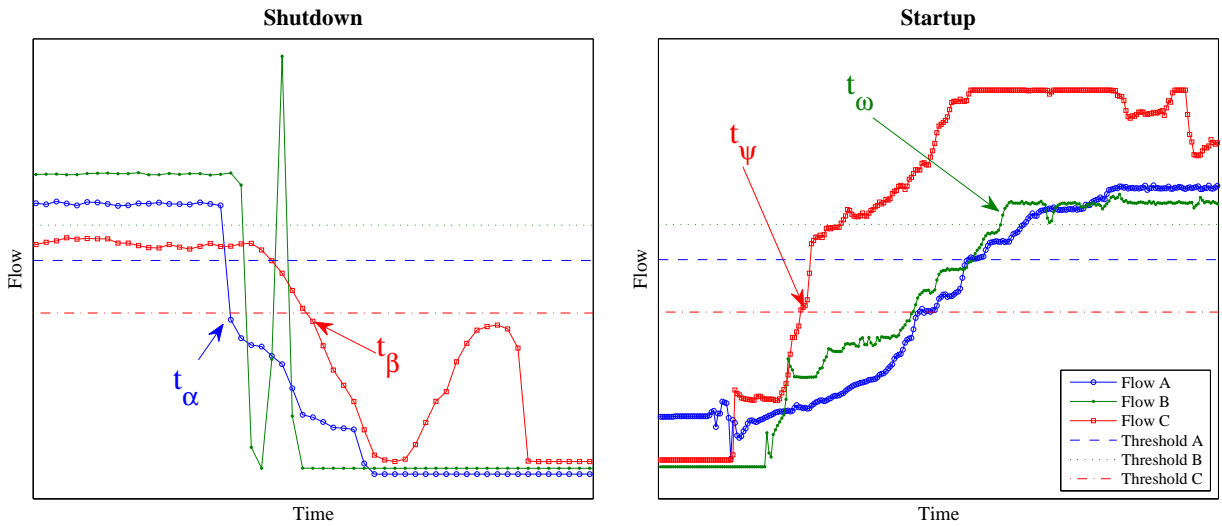

Fig. 2. Plot on the left shows the points t_α and t_β which are chosen in annotation as the time interval of the beginning of the shutdown period. Plot on the right shows the points t_ψ and t_ω which are annotated as the time interval of the end of the shutdown.

The statistic s_t will increase during the interval $[t_\alpha, t_\beta]$ and it will decrease during $[t_\psi, t_\omega]$. As a consequence, different stopping rules have to be used according to the type of change that we would like to detect. That is, Equation 2 is suitable for detecting shutdowns, while the following stopping rule is more suitable for detecting startups:

$$T = \inf\{t : s_t(X_t) < \tau\} \tag{3}$$

The right use of either stopping rules is associated with the process state since they are contradictory.

In the case of a shutdown we would like to detect quickly its beginning, that is, in a time point $t_a \geq t_\alpha$ as close as possible to t_α. On the other hand, in case of a startup we would like to detect the very end of the shutdown period, that is, a time point $t_z \geq t_\omega$ as close as possible to t_ω.

Deployment of a method for detecting these periods can result in three different outcomes: a) correct detection; b) false detection; and c) no detection. A good method should be able to maximize the number of correct detections and to minimize the number of false detection/no detection cases. This requirement is directly related to the reduction of the detection delay and the rate of false detections, which are the two most common metrics in the change-point detection literature.

3. Multi-sensor change-point detection methods

The detection of abrupt changes in single-sensory data (i.e. one-dimensional) has been well studied and solved. For example, the book of Basseville and Nikiforov[4] is one of the main references for this problem. Lai[5] surveys the sequential change-point detection methods in quality control and dynamical systems. A more recent state-of-the-art in single-sensor sequential change-point detection is presented by Polunchenko and Tartakovsky[6] where methods of main formulations are reviewed.

The extension of this problem to multiple-sensory data has been also addressed by several authors. Viswanathan et al.[7,8] present a two-part review of methods using different topologies and approaches. Classical methods use all the data collected until the current time t. However, those approaches are not feasible for practical purposes where data streams are continuously arriving to be processed. For this case study, we have selected and implemented a number of window limited versions of the state-of-the-art methods in order to carry out a comparative performance study.

The first chosen method was proposed by Tartakovsky and Veeravalli[9] where a likelihood ratio test is carried out for each sensor and individual results are aggregated. The statistic to define the stopping rules of this method is referred to TV in Table 1.

Mei[10] proposes a family of scalable schemes for global online monitoring of data streams based on CUSUM statistics from each individual data stream. The same author extends this work where the fact that the change point may be different in each data stream is taking into account[11]. The statistic monitored by Mei's method is referred to MEI in Table 1.

Xie and Siegmund[12] propose a mixture procedure based on the aggregation of the local generalized likelihood ratio (GLR) statistic of each sensor. This method assumes that the pre- and post-change distributions are Gaussian with pre-change mean being zero. This mixture includes a fraction of affected sensors by the change that has to be fixed a priori. The statistics of two different stopping rules proposed by the authors can be found in Table 1 (as XS1 and XS2).

Finally, we have implemented a method based on Shewhart's control charts[13]. These charts are widely used in the industry to distinguish between two states of a process (i.e. in-control and out-of-control). Section 3.1 provides more detailed explanation of this method. The statistic monitored by this method is referred to SGZ in Table 1.

In the recent years, multi-sensor change-point detection methods have been applied for example to fault detection[14] and intrusion detection[15], but no works in shutdown periods detection are available so far to the best of our knowledge.

A common assumption in the literature is that a change has to be detected as soon as possible. This is also true in the case of detecting the beginning of a shutdown. On the other hand, the pipeline structure of a big chemical plant means that the re-initialization of the sensors after a startup is delayed according to their spatial location. Thus, the detection of a startup should be deferred until all the parts of the chemical plant are working in a steady state which makes the startup detection a much more challenging problem.

3.1. Multi-sensor change-point detection method based on control charts

In order to control a quality measure, an upper and lower thresholds are computed with historical data. It is common to get these thresholds using the 3σ-rule, which state that for a normal distribution the 99.7% of values lies in the interval $(\mu - 3\sigma, \mu + 3\sigma)$ where μ is the mean and σ is the standard deviation of the sample.

To extend this method to our case study, the limits for each sensor are computed as $L_n = \tilde{\mu}_n - 3\hat{\sigma}_n$ and $U_n = \tilde{\mu}_n + 3\hat{\sigma}_n$, where $\tilde{\mu}_n$ is the median and $\hat{\sigma}_n \approx 1.4826 \text{MAD}$ is the estimation of the standard deviation using the median absolute deviation computed over the historical data.

The presence of outliers in the data might lead to false detections. Therefore, to increase robustness of the method we have introduced a window of out-of-control values for each sensor. Thus,

$$C_{n,t} = (\mathcal{B}(x_{n,t-r}), \ldots, \mathcal{B}(x_{n,t})) \tag{4}$$

is the window of r binary values for the sensor n where

$$\mathcal{B}(x_{n,t}) = \begin{cases} 1, & \text{if } x_{n,t} \notin [L_n, U_n] \\ 0, & \text{otherwise} \end{cases} \tag{5}$$

A window for each sensor is monitored and weighted according to its reliability. Binary weights are used to discard those sensors which might be failing. A similar approach to identify faulty sensors is presented by Seron et al.[16]. Let

$$\gamma_{n,t} = \sum c_i, \quad \forall c_i \in C_{n,t} \tag{6}$$

be the number of out-of-control values in the window $C_{n,t}$. Each weight is then updated in each time t as

$$w_n = \begin{cases} 1, & \text{if } \gamma_{n,t} \in [\mathcal{L}, \mathcal{U}] \\ 0, & \text{otherwise} \end{cases} \tag{7}$$

where $\mathcal{L}$ and $\mathcal{U}$ are the thresholds of $C_{1..N,t}$ computed using the Hampel identifier[17].

Final decision for detecting a change is taken by aggregating the weighted counters of all sensors. Therefore, using

$$s_t(X_t) = \max_{1 \le n \le N} (w_n \gamma_{n,t}) \tag{8}$$

as a statistic (SGZ in Table 1) we ensure both quick detection during a shutdown and deferred detection during a startup. This aggregation can therefore deal with the delays between sensors due to their spatial location. Pseudo-code for this method is presented in Algorithm 1.

4. Experimental evaluation

The goal of this experimental evaluation is to compare the performance and reliability of different multi-sensor change-point detection methods in our case study and to select the most suitable for a production environment. For that purpose, we first have defined the evaluation measures and then we have established an experimental protocol. Finally, we discuss the results of the conducted experiments.

4.1. Dataset of a chemical plant

The dataset is composed of 109,627 records, sampled every 10 minutes from the archive of readings of the chemical plant. Each record contains 11 numerical values from flow sensors. The structure of the sensors is unknown. The records have been aligned by time-stamp in order to form instances. Missing values have been interpolated to simplify the procedures, but apart from that the dataset has not been preprocessed, so it contains outliers and noise as we would expect from the online process.

The dataset is not publicly available but can be requested to the corresponding authors by email.

4.2. Evaluation measures

In change-point detection literature there is a trade-off between detection delay and false alarm rate. While the objective is to minimize both measures, a threshold for a quick detection delay can increase the number of false

---

Algorithm 1 Function processSample(X_t)

1: **for** $n = 1 \rightarrow N$ **do**
2:      $\mathcal{B}$ = isOutlier($x_{n,t}$)   ▷ Eq. 5 equivalent
3:      C_n.removeOldest()   ▷ C_n is a list of size r initialized as a global variable
4:      C_n.append($\mathcal{B}$)   ▷ Eq. 4 equivalent
5: **end for**
6: $[\mathcal{L}, \mathcal{U}]$ = getReliabilityThresholds($C_{1,...,N}$)
7: $s_t = 0$
8: **for** $n = 1 \rightarrow N$ **do**
9:      γ_n = sum(C_n)   ▷ Eq. 6 equivalent
10:     w_n = inLimits($\gamma_n, \mathcal{L}, \mathcal{U}$)   ▷ Eq. 7 equivalent
11:     s_t = max($s_t, w_n \cdot \gamma_n$)   ▷ Eq. 8 equivalent
12: **end for**
13: **if** processActive **and** $s_t \geq \tau$ **then**   ▷ Eq. 2 equivalent
14:     shutdownDetected()
15:     processActive = false   ▷ processActive is defined as a global variable
16: **else if** ¬processActive **and** $s_t < \tau$ **then**   ▷ Eq. 3 equivalent
17:     startupDetected()
18:     processActive = true
19: **else**
20:     continueProcessSample(X_t)
21: **end if**

---

alarms (i.e. incorrect detections). In our case study we want to avoid false detection at all cost but at the same time we would not like to get long delays.

Although shutdowns and startups are both changes from the theoretical point of view, we distinguish between them during the experiments because they are different in practice as we explained in Section 2. Therefore, we measure the shutdown's detection delay as $\Delta_\alpha = t_a - t_\alpha$ and the startup's detection delay as $\Delta_\omega = t_z - t_\omega$ where t_a and t_z are the times of the detection for each case.

4.3. Experimental setting

The dataset has been split in two equally-sized parts. The first half is used for the calibration of parameters and the second half for evaluating the methods. Each half contains 22 change-points (11 shutdowns and 11 startups) that have been manually annotated.

Table 1 contains the distinctive formulas for calculating the time changing statistical values used in the stopping rule of each method. In this table, t is the current time, r is the size of a temporal window,

$$\ell_n(t, k, \mu_n) = \sum_{i=k+1}^{t} (\mu_n x_{n,i} - \mu_n^2/2) \tag{9}$$

is the log-likelihood of observations accumulated by time t, μ_n is the mean of the data during a shutdown,

$$\hat{\mu}_{n,k,t} = \frac{\sum_{i=k+1}^{t} x_{n,i}}{t - k} \tag{10}$$

is the maximum likelihood estimator of the mean, and $\Pr(x_t)$ is the posterior probability of x_t in the distribution $\mathcal{D}_0$ or $\mathcal{D}_1$, according to the process state.

The limits of the stopping rules for all the methods have to be chosen to minimize both the detection delay and the rate of false alarms. In the case of XS1, XS2, MEI and TV this limit lies in $0 < \tau \lesssim \sum_{n=1}^{N} r \cdot \mu_{n,1}^2/2$, and for SGZ

Table 1. Formulas used for $s_t(\boldsymbol{X}_t)$ in Equations 2 and 3

Method	$s_t(\boldsymbol{X}_t)$
TV	$\max\limits_{t-r \leq k \leq t} \sum\limits_{n=1}^{N} \ell_n(t, k, \mu_n)$
MEI	$\sum\limits_{n=1}^{N} \max\limits_{t-r \leq k \leq t} \ell_n(t, k, \mu_n)$
XS1	$\max\limits_{t-r \leq k \leq t} \sum\limits_{n=1}^{N} \log(1 - p_0 + p_0 \exp[\ell_n^{+}(t, k, \mu_n)])$
XS2	$\max\limits_{t-r \leq k \leq t} \sum\limits_{n=1}^{N} \log(1 - p_0 + p_0 \exp[(\hat{\mu}_{n,k,t}^{+})^2/2])$
SGZ	$\max\limits_{1 \leq n \leq N} (w_n \gamma_{n,t})$

Table 2. Limit values τ for each method and window size r

r	20	25	30	35	40	45	50	55	60	65	70	75	80	85	90	95	100
TV	16	16	16	16	16	16	16	16	16	16	**10**	10	10	10	10	10	10
MEI	18	17	17	18	19	17	18	18	19	20	19	**11**	11	11	11	11	11
XS1	80	68	**17**	80	23	58	60	25	46	57	57	57	57	57	46	57	35
XS2	16	16	16	16	16	16	16	16	16	16	**10**	10	10	10	10	10	10
SGZ	9	5	**5**	5	5	5	5	5	5	5	5	5	5	5	5	5	5

in $0 < \tau < r$. Limits for each method have been automatically set using the first half of the dataset as input. Thus, the computed limit is the minimum s_t value that avoids the false alarms and detects all the change-points. Table 2 contains these limits for each window size.

4.4. Experimental results

A sample of the s_t values computed for each method is shown in Figure 3. The statistics for TV, MEI, XS1 and XS2 are very similar because they are all based on the likelihood. The flat signal of SGZ indicates that the windows of out-of-control values is full ($r = 25$).

The left side of Figure 4 compares the median values of the detection delays during the shutdown phases. The window size almost does not affect to the these types of detections because the s_t values quickly increase during the shutdown phase and overcome the threshold τ. The MEI, TV and XS2 methods reported lower detection delays than XS1 and SGZ.

On the other hand, the window size has a significant effect on the median values of the detection delays during the startup phases as shown in the right part of Figure 4. A negative delay means that the change-point has been detected before t_ω. In this case study, we are interested in keeping a small positive detection delay. In both XS1 and SGZ a value of $r = 30$ satisfies that requirement. However, a value of $r = 75$ is needed for TV, MEI and XS2.

False alarms have only be reported for MEI with window sizes $r = \{20, 25, 30, 45, 50, 55, 60\}$ and for SGZ with window sizes $r = \{20, 65\}$. The rest of the methods have not raised any false alarms within the experimental setup.

After the experimentation we conclude that the methods behave similarly although selection of window size makes a difference. The MEI method detects the changes quicker than the other methods but at the same time it raises false alarms in some of the cases. Any of the following configurations would be suitable to be implemented in a production environment: (TV, $r = 70$), (MEI, $r = 75$), (XS1, $r = 30$), (XS2, $r = 70$) and (SGZ, $r = 30$). If memory requirements are costly, we would choose the XS1 method because its performance is better than the others with lower window size.
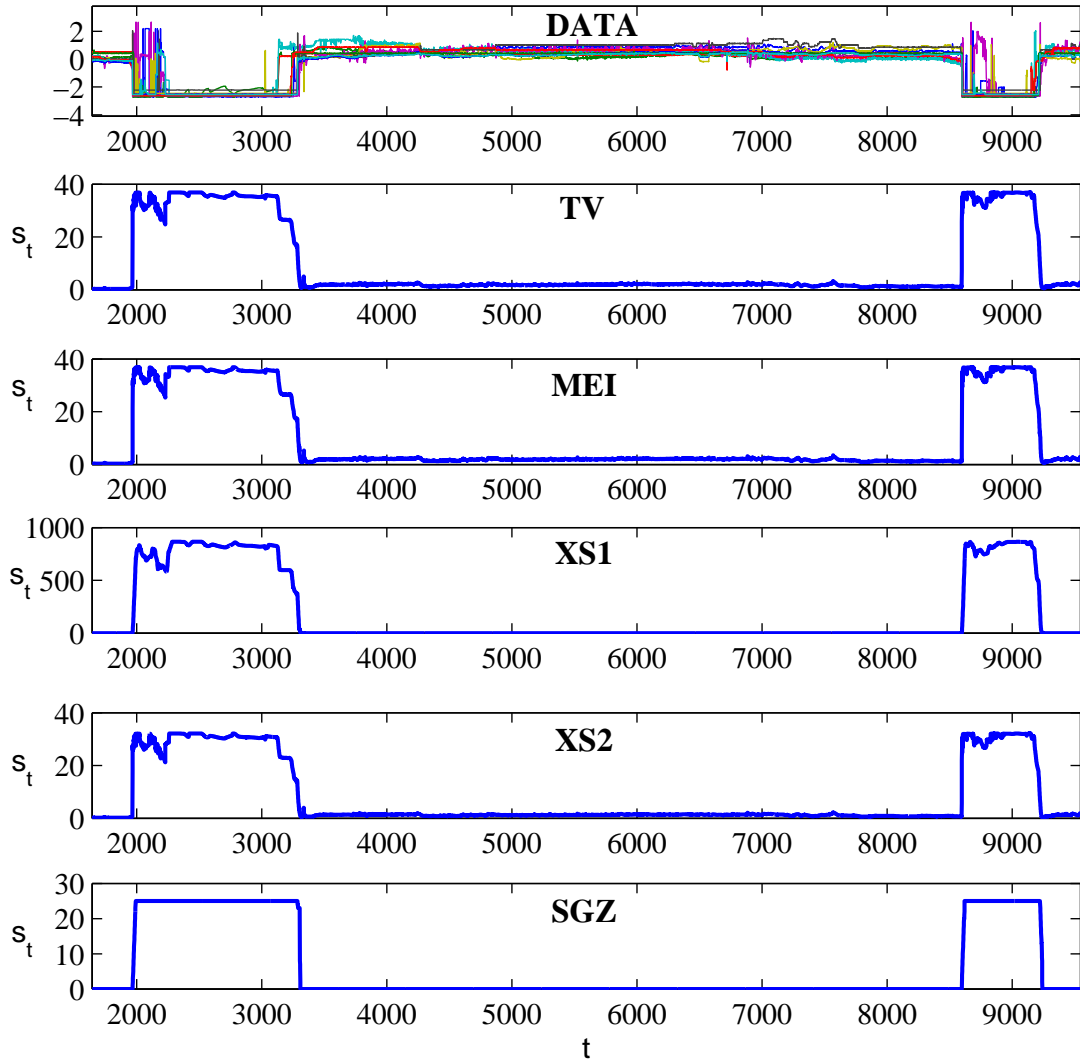
Fig. 3. Subset of the observed data and s_t values for all the methods for $r = 25$.
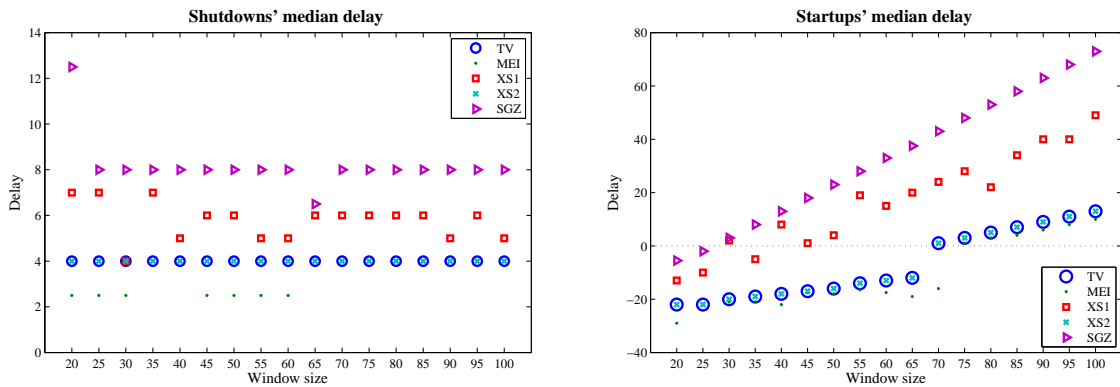

Fig. 4. Median of the detection delays of the shutdowns (left) and the startups (right).

5. Conclusions

We have presented a case study obtained from a chemical production company. The task has been to identify the shutdown periods of a chemical plant using the information provided by 11 flow sensors. For that, we have firstly stated the problem setting and then selected a number of state-of-the-art methods in multi-sensor change-point detection. In addition, we have developed a robust method based on control charts which are very popular in the industry.

We have conducted a series of experiments using a common framework in order to evaluate the performance of all these methods for our case study. The results point to the XS1 method as the most suitable approach for our purposes since it has a small detection delay both in shutdown and startup periods while keeping the memory requirements lower than the other methods.

As indicated in the introduction of this paper the presented automatic shutdown/startup detection study, as well as the evaluation and adaptation of existing methods, form a small part of the ongoing and future development of an automated preprocessing framework for adaptive soft sensors.

In this context the most immediate future work includes a study of how the use of the evaluated methods in an online preprocessing environment can affect the performance of adaptive the soft-sensors and their predictive accuracy.

References

1. Kadlec, P., Gabrys, B., Strandt, S.. Data-driven Soft Sensors in the process industry. *Computers & Chemical Engineering* 2009;**33**(4):795–814. doi:10.1016/j.compchemeng.2008.12.012.
2. Kadlec, P., Grbić, R., Gabrys, B.. Review of adaptation mechanisms for data-driven soft sensors. *Computers & Chemical Engineering* 2011;**35**(1):1–24. doi:10.1016/j.compchemeng.2010.07.034.
3. Zliobaite, I., Gabrys, B.. Adaptive Preprocessing for Streaming Data. *IEEE Transactions on Knowledge and Data Engineering* 2014; **26**(2):309–321. doi:10.1109/TKDE.2012.147.
4. Basseville, M., Nikiforov, I.. *Detection of abrupt changes: theory and application*. Prentice-Hall, Inc; 1993. ISBN 0-13-126780-9.
5. Lai, T.. Sequential changepoint detection in quality control and dynamical systems. *Journal of the Royal Statistical Society Series B (Methodological)* 1995;**57**(4):613–658.
6. Polunchenko, A., Tartakovsky, A.. State-of-the-Art in Sequential Change-Point Detection. *Methodology and Computing in Applied Probability* 2011;**14**(3):649–684. doi:10.1007/s11009-011-9256-5.
7. Viswanathan, R., Varshney, P.. Distributed detection with multiple sensors I. Fundamentals. *Proceedings of the IEEE* 1997;**85**(1):54–63. doi:10.1109/5.554208.
8. Blum, R., Kassam, S., Poor, H.. Distributed detection with multiple sensors I. Advanced topics. *Proceedings of the IEEE* 1997;**85**(1):64–79. doi:10.1109/5.554209.
9. Tartakovsky, A., Veeravalli, V.. Asymptotically Optimal Quickest Change Detection in Distributed Sensor Systems. *Sequential Analysis* 2008;**27**(4):441–475. doi:10.1080/07474940802446236.
10. Mei, Y.. Efficient scalable schemes for monitoring a large number of data streams. *Biometrika* 2010;**97**(2):419–433. doi:10.1093/biomet/asq010.
11. Mei, Y.. Final Report for the Project Robust Rapid Change-Point Detection in Multi-Sensor Data Fusion and Behavior Research. Tech. Rep.; Georgia Institute of Technology; 2011.
12. Xie, Y., Siegmund, D.. Sequential multi-sensor change-point detection. *The Annals of Statistics* 2013;**41**(2):670–692. doi:10.1214/13-AOS1094.
13. Shewhart, W.. *Economic Control of Quality of Manufactures Product*. Asq Press; 1931. ISBN 0-87389-076-0.
14. Rajagopal, R., Nguyen, X., Ergen, S., Varaiya, P.. Distributed Online Simultaneous Fault Detection for Multiple Sensors. In: *2008 International Conference on Information Processing in Sensor Networks (ipsn 2008)*. IEEE. ISBN 978-0-7695-3157-1; 2008, p. 133–144. doi:10.1109/IPSN.2008.41.
15. Tartakovsky, A., Rozovskii, B., Blazek, R.. A novel approach to detection of intrusions in computer networks via adaptive sequential and batch-sequential change-point detection methods. *IEEE Transactions on Signal Processing* 2006;**54**(9):3372–3382. doi:10.1109/TSP.2006.879308.
16. Seron, M., De Dona, J., Olaru, S.. Fault Tolerant Control Allowing Sensor Healthy-to-Faulty and Faulty-to-Healthy Transitions. *IEEE Transactions on Automatic Control* 2012;**57**(7):1657–1669. doi:10.1109/TAC.2011.2178716.
17. Davies, L., Gather, U.. The Identification of Multiple Outliers. *Journal of the American Statistical Association* 1993;**88**(423):782. doi:10.2307/2290763.