An Ensemble Approach to Estimate the Fault-Time Instant

Cesare Alippi, Giacomo Boracchi, Vicenç Puig, and Manuel Roveri

Abstract— [1]

Since systems are prone to faults, fault detection and isolation are essential activities to be considered in safety-critical applications. In this direction, the availability of a sound estimate of the time instant the fault occurred is a precious information that a diagnosis system can fruitfully exploit, e.g., to identify information consistent with the faulty state. Unfortunately, any fault-detection system introduces a structural delay that, typically, increases in correspondence of subtle faults (e.g., those characterized by a small magnitude) with a consequence that the fault-occurrence time is overestimated.

In this paper we propose an ensemble approach to estimate the time instant a fault occurred. We focus on systems that can be described as ARMA models and faults inducing an abrupt change in the model coefficients.

I. Introduction

COMPLEX embedded systems designed to work in real-world, possibly harsh, environments e.g., networked embedded systems or large-scale sensor-actuator networks, are prone to faults or ageing effects that might affect embedded electronics, actuators and sensors. Faults and ageing must hence be detected as soon as possible to reduce the "cascade effect" on the rest of the system and activate accommodation actions (e.g., substitution of the faulty unit). A desirable feature in real-world applications [1] is the automatic fault-diagnosis, since it allows to isolate the unit (or the units) affected by faults (fault isolation) and identify the kind of fault occurred (fault identification). These mechanisms, which are activated after each fault detection, generally require an estimate of the time-instant the fault occurred.

The fault-time instant can be used to partition the data before (*healty* condition) and after the change (*faulty* condition), with the latter data particularly useful for fault isolation and identification. Typically, the fault time-instant can be computed offline, by processing a fixed-length buffer containing the most recent observations. To this purpose one may enforce Change-Point Methods (CPMs), namely, statistical techniques to assess if the data can be split in two consecutive parts that have been generated from two different models, i.e., the data contains a change point [2], [3].

In this work we address the problem of estimating T^*, namely the time instant the fault occurred, once the fault has been detected. We focus on systems whose observations can be modeled with autoregressive-moving-average (ARMA) processes, since ARMA are effective system models for autocorrelated processes. We consider a general class of faults inducing an abrupt change in the parameter vector of the ARMA model.

Very few CPMs have been developed for autocorrelated processes. For example, [4] estimates the time instant when a change in the parameters of an AR(1) process occurs, while [5], [6] address changes affecting the mean of an ARMA process. We tackle the problem of estimating the fault-time instant by looking at residuals generated by a predictive models. In principle, a viable option consists in using a CPM on the residuals computed by means of a suitable approximation model, as far as these are expected to be, before the fault, white noise. However, in the practice, residuals are not white even during healthy conditions, since predictors are typically affected by model bias that, often, introduces a temporal dependency in the residual sequence. Unfortunately, these temporal dependencies impair the performance of CPM.

Here we approach the problem of estimating the fault-time instant from a residual sequence by means of an ensemble of CPMs. Ensemble methods [7] improve the generalization ability of a single model by combining a large number of models meant to solve the same problem. These methods have been successfully applied in regression and classification [8], as well as in specific applications such as face and object recognition, intrusion detection, optical character recognition and medical diagnosis [9]. In the last years the interest of the research community on ensemble methods designed for time series prediction has exponentially grown (e.g., [10], [11]) corroborating the use of this ensemble-of-models approach in time-dependent scenarios.

To the best of our knowledge, ensemble methods have never been proposed for estimating the fault-time instant, and this paper illustrate a first attempt in this direction. A key point of the proposed solution is that the ensemble relies on random sampling of the residual sequence to reduce temporal dependencies associated with the model bias. The proposed solution is general, as it is possible to build ensembles by using different CPMs.

The paper is organized as follows: Section II introduces the problem, while the CPM and problems arising when these are used on a residuals' sequence are presented in Section III. The proposed ensemble of CPMs is introduced in Section IV and performance tested in Section V.

II. Problem Formulation

We consider systems that can be successfully modeled by an autoregressive moving-average (ARMA) process
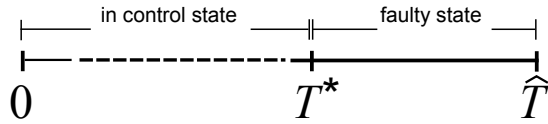
Fig. 1. The time-instants of fault detection: T^* is the time instant the fault occurred, $\widehat{T}$ is the time instant the fault is detected by the FDS. During the interval $[0, T^*)$, the system is in healthy conditions and observations are from an ARMA process having parameter vector θ_0. During the interval $[T^*, \widehat{T}]$ the system is faulty and observations are from a different ARMA process having parameter vector θ_1.

$$y(t) = \sum_{i=1}^{i=p} \phi_i y(t-i) + \sum_{i=1}^{i=q} \psi_i \epsilon(t-i) + \epsilon(t), \quad (1)$$

where $\theta = [\phi_1, \ldots, \phi_p, \psi_1, \ldots, \psi_q]$ is the parameter vector, $y(t) \in \mathbb{R}$ and $\epsilon(t) \sim \mathcal{N}(0, \sigma^2)$ denote the observation and the innovation at time t, respectively. The parameters $p > 0$ and $q > 0$ represent the orders of the autoregressive and moving-average terms, respectively. Eq. (1) describes observations when the monitored system is in healthy conditions. After the fault-time instant T^*, observations are no more described by (1).

As a general class of faults, we consider abrupt changes in the parameter vector of the ARMA model describing the healthy conditions. In practice, the parameter vector characterizing (1) shifts permanently, at an unknown time instant T^*, from θ_0 to $\theta_1 \neq \theta_0$. In what follows we refer to θ_0 as the error-free parameters describing the healthy conditions, and with θ_1 as the fault-affected parameters.

Targeting real-world scenarios, we assume that the parameter vector of the faulty model θ_1 is unknown and the that θ_0 has to be estimated from a suitable training sequence [12]. Such an estimate is denoted by $\widehat{\theta}_0$.

From now on, we assume that a suitable fault-detection system (FDS) is inspecting the system online, possibly not solely relying on observations $y(t)$ but also exploiting additional information, including that coming from an expert.

We draw our attention on the post-detection operations and we address the issue of computing F, an estimate of the fault-time instant T^*, every time the FDS detects a fault. The time instant when the FDS detects the fault is instead denoted by $\widehat{T}$.

III. CHANGE POINT METHODS

We consider CPMs [2] consisting in hypothesis tests to analyze, in an offline manner, a given data sequence, which – under the null hypothesis – contains i.i.d. realizations of a random variable. There, each point of the sequence is considered as a potential change point to be inspected. For each candidate point, the data sequence is split in two parts (the samples before and after the potential change point, respectively) and a specific test statistic is computed to compare the probability distribution of the data in the two parts. Then, when all the points have been tested, the CPM draws the conclusion that there is enough statistical evidence to reject the null hypothesis if, in at least one point, the test statistic exceeds a specific threshold (which depends on the test statistic as well as the sequence length and a confidence level). In these cases, the sequence is considered to contain a change point, and the point yielding the maximum of the statistic is selected. On the contrary, when the test statistic never exceeds the threshold, the CPM does not reveal a change in the data distribution.

Examples of test statistics are the sample mean or sample variance, which lead to the Student or Bartlett tests [13], for Gaussian distributed data. Sometimes it is advisable not to assume any specific distribution of the data, and rely on nonparametric test statistics, e.g, the Mann-Whitney, the Mood [14] or the Lepage [15], which assess changes in the sample mean, sample variance, or both simultaneously, respectively. Also the Kolmogorov-Smirnov statistic can be used to detect more general changes in the data distribution.

One of the most critical issues in CPM consists in computing the thresholds for rejecting the null hypothesis, since often they do not admit an analytical formulation and one has to resort to Montecarlo simulations [2], [3].

We recall that the CPM framework, which has been designed to estimate the change point within a fixed length sequence [2] has been also extended to address the online (sequential) detection of changes to data streams [16], [3].

A. CPM for Abrupt Changes in the Residuals

The CPMs cannot be applied straightforward to the observation sequence $\{y(t), t = 1, \ldots, \widehat{T}\}$, as these are not i.i.d. realizations of a random variable. Instead, CMPs can be run on the residual sequence

$$\mathcal{R} = \{r(t), t = 1, \ldots, \widehat{T}\} \quad (2)$$

that are computed w.r.t. the estimated model $\widehat{\theta}_0$, i.e.,

$$r(t) = y(t) - f_{\widehat{\theta}_0}(t), \quad (3)$$

where $f_{\widehat{\theta}_0}$ is the predictor associated to the ARMA model of parameter vector $\widehat{\theta}_0$. In fact, provided that the initial estimate $\widehat{\theta}_0$ is accurate, residuals are white, thus (2) contains i.i.d. realizations of a Gaussian random variable, when $t < T^*$. Therefore, CPMs can be applied in principle to the residuals to estimate the fault-time instant.

Without loss of generality and for the sake of notation, in what follows we assume that $\mathcal{R}$ stores all the residuals (2), even though it is possible to store and process only the most recent ones in a FIFO buffer.

The CPM on the residual sequence $\mathcal{R}$ can be formulated as follows: each time instant $S \in \{1, \ldots, \widehat{T}\}$ is considered to be a possible change-point, and induces a partitioning of $\mathcal{R}$ into two parts

$$\begin{aligned} \mathcal{A}_S &= \{r(t), \ t = 1, \ldots, S\}, \quad (4) \\ \mathcal{B}_S &= \{r(t), \ t = S+1, \ldots, \widehat{T}\}. \end{aligned}$$

Then, the value of the test statistic $\mathcal{T}$ is computed at time instant S as

$$\mathcal{T}_S = \mathcal{T}(\mathcal{A}_S, \mathcal{B}_S), \quad (5)$$

to assess the difference between $\mathcal{A}_S$ and $\mathcal{B}_S$.

The estimate of the fault-time instant $F_\mathcal{R}$, obtained from the whole sequence $\mathcal{R}$, is the point yielding the partition that maximizes the test statistic:

$$F_\mathcal{R} = \operatorname*{argmax}_{S=1,\ldots,\widehat{T}} (\mathcal{T}_S), \qquad (6)$$

provided that the value of the test statistic at $F_\mathcal{R}$, i.e., $\mathcal{T}_{F_\mathcal{R}} = \mathcal{T}(\mathcal{A}_{F_\mathcal{R}}, \mathcal{B}_{F_\mathcal{R}})$, exceeds the corresponding threshold $h_{l,\alpha}$. This threshold depends on the statistic $\mathcal{T}$, on the length of the sequence $l = \#\mathcal{R}$, and on a defined confidence level α, which sets the percentage of type I errors (i.e., false positives) of the hypothesis test. Therefore, the CPM outcome is

$$\begin{cases} \text{The fault-time instant is } F_\mathcal{R} & \text{if } \mathcal{T}_{F_\mathcal{R}} \geq h_{l,\alpha} \\ \text{No fault-time instant identified,} & \text{if } \mathcal{T}_{F_\mathcal{R}} < h_{l,\alpha} \end{cases}. \qquad (7)$$

Any occurrence of the latter case corresponds to situations where the CPM is not able to estimate the fault-time instant. Therefore, the CPM outcomes can be interpreted as a validation of the FDS detections, as in [17].

B. Issues Related to the Use of CPM on Residuals

Unfortunately, in real applications the residuals $\mathcal{R}$ are typically far from being i.i.d. and Gaussians even before T^*, because of model bias (3). Moreover, the residuals computed after the unknown fault-time instant i.e., $\{r(t), t = T^*, \ldots, \widehat{T}\}$ are definitively not i.i.d. and Gaussian distributed and we expect a large degree of temporal dependence among these residuals. These circumstances violate the hypothesis required by the CPM and explain why, more often than in the i.i.d. case, the statistic $\mathcal{T}_{F_\mathcal{R}}$ does not reach the corresponding threshold and the CPM is not able to estimate of the fault-time instant in $\mathcal{R}$.

In light of on the above considerations, it is convenient:

1) to exploit a nonparametric CPM rather than a parametric one assuming a Gaussian distribution in $\mathcal{R}$,
2) to introduce some sort of pre-processing to remove the time-dependency from data to improve the effectiveness of CPMs.

While there are several nonparametric CPMs available in the literature [18], [3], the latter issue has been less investigated and we address this problem by means of an ensemble of CPMs that is described in the sequel.

IV. THE ENSEMBLE OF CPMS

We pursue an ensemble approach for estimating the fault-time instant by aggregating d estimates $\{F_i, i = 1, \ldots, d\}$ of T^* provided by CPMs executed on different subsequences of residuals. Peculiarity of the proposed solution is that each of such subsequence is obtained by randomly sampling $\mathcal{R}$.

A. Computing the Individual Estimates

Let us introduce $\mathrm{D_n}(\cdot)$, the operator performing random sampling:

$$\mathcal{I}_n = \mathrm{D_n}(\{1, \ldots, m\}) \quad \text{where } \#\mathcal{I}_n = n < m, \qquad (8)$$

Algorithm 1: Ensemble of CPMs $\mathcal{E}_d(\mathcal{R})$

Input: $\mathcal{R} = \{r(t), t = 0, \ldots, \widehat{T}\}$ (the residuals), α (the confidence level of the CPM), d (the number of estimates from random sampling $\mathcal{R}$), n (the random sampling parameter).

Output: F (the estimate of the fault-time instant).

1- $i = 0$, $l = \#\mathcal{R}$,
2- **while** $(i \leq d)$ **do**
3- $\mathcal{I}_n^{(i)} = \mathrm{D_n}(\{1, \ldots, l\})$
4-   **foreach** $S \in \mathcal{I}_n^{(i)}$ **do**
5- $\mathcal{A}_S = \{r(t), \ t \in \mathcal{I}_n^{(i)}, t \leq S\}$,
6- $\mathcal{B}_S = \{r(t), \ t \in \mathcal{I}_n^{(i)}, t > S\}$,
7- $\mathcal{T}_S = \mathcal{T}(\mathcal{A}_S, \mathcal{B}_S)$
     end
8- $F_i = \operatorname*{argmax}_{S \in \mathcal{I}_n} (\mathcal{T}_S)$, $\mathcal{T}_{\mathcal{F}_i} = \max_{S \in \mathcal{I}_n} (\mathcal{T}_S)$,
9-   **if** $(\mathcal{T}_{\mathcal{F}_i} \geq h_{n,\alpha})$ **then**
10- $\omega_i = 1$,
     else
11- $\omega_i = 0$,
     end
  end
12- Compute $\mathcal{F}_\mathcal{R}$ as in (11) and the corresponding value of the statistic $\mathcal{T}_{\mathcal{F}_\mathcal{R}}$
13- **if** $(\mathcal{T}_{\mathcal{F}_\mathcal{R}} \geq h_{l,\alpha})$ **then**
14- $\omega_{d+1} = 1$,
  else
15- $\omega_{d+1} = 0$,
  end
16- $F = \frac{\sum_{i=1}^{d} \omega_i F_i + \omega_{d+1} F_\mathcal{R}}{\sum_{i=1}^{d+1} \omega_i}$

---

which generates a subsequence $\mathcal{I}_n$ containing n time instant randomly chosen – without repetition – from the first m integers. It is important to remark that the elements in $\mathcal{I}_n$ are sorted, thus *preserving the temporal order*. We denote the set of the residuals at time instants in $\mathcal{I}_n$ as

$$\mathcal{R}_{\mathcal{I}_n} = \{r(t), \ t \in \mathcal{I}_n\} \qquad (9)$$

Let us denote by $\mathrm{CPM}_\mathcal{T}$ the operator associated to a CPM using test statistic $\mathcal{T}$ which, given a sequence $\mathcal{X}$, provides

$$F_\mathcal{X} = \mathrm{CPM}_\mathcal{T}(\mathcal{X}), \qquad (10)$$

the index maximizing $\mathcal{T}$ when determining a change point in $\mathcal{X}$, as described in Section III-A.

Thereby, the estimate $\mathcal{F}_\mathcal{R}$ provided by (4) - (6) becomes:

$$F_\mathcal{R} = \mathrm{CPM}_\mathcal{T}(\mathcal{R}). \qquad (11)$$

Following the notation in (10), we denote by F_i the estimate obtained from the i-th random sampling of $\mathcal{R}$, namely

$$\mathcal{I}_n^{(i)} = \mathrm{D_n}(\{1, \ldots, l\}) \qquad (12)$$

$$j = \mathrm{CPM}(\mathcal{R}_{\mathcal{I}_n^{(i)}}) \qquad (13)$$

$$F_i = \mathcal{I}_n^{(i)}[j], \qquad (14)$$

where $l = \#\mathcal{R}$, and $\mathcal{I}_n^{(i)}[j]$ denotes the j-th element of the sequence $\mathcal{I}_n^{(i)}$. Note that $\mathcal{I}_n^{(i)}[j]$ maps the estimate of the change-point from the subsequence in $\mathcal{R}_{\mathcal{I}_n^{(i)}}$ back on $\mathcal{R}$ indexes, thus in the temporal domain.

B. Aggregation

The ensemble $\mathcal{E}_d(\mathcal{R})$ provides an estimate F of T^* by aggregating d individual estimates $\{F_i, i = 1, \dots, d\}$ from (12) together with $F_{\mathcal{R}} = \mathrm{CPM}_\mathcal{T}(\mathcal{R})$. Aggregation of these estimates is performed by averaging:

$$F = \frac{\sum_{i=1}^{d} \omega_i F_i + \omega_{d+1} F_\mathcal{R}}{\sum_{i=1}^{d+1} \omega_i} \qquad (15)$$

where weights ω_i, $i = 1, \dots, d+1$ are binary and given by

$$\omega_i = \begin{cases} 0, & \text{if } \mathcal{T}_{F_i} < h_{n,\alpha} & i = 1, \dots, d \\ 1, & \text{if } \mathcal{T}_{F_i} \geq h_{n,\alpha} & i = 1, \dots, d \\ 0, & \text{if } \mathcal{T}_{F_\mathcal{R}} < h_{l,\alpha} & i = d+1 \\ 1, & \text{if } \mathcal{T}_{F_\mathcal{R}} \geq h_{l,\alpha} & i = d+1 \end{cases} . \qquad (16)$$

where $h_{n,\alpha}$ and $h_{l,\alpha}$ are the test-statistic thresholds. Equations (15) and (16) ensure that we aggregate only estimates provided by CPMs that are able to assess the change in the residual distribution.

C. The Algorithm

Algorithm 1 details the proposed ensemble of CPMs. In particular, the loop at lines 2 - 11 perform d-times the random sampling of $\mathcal{R}$, computes the corresponding estimates $\{F_i, i = 1, \dots, d\}$ and also the values of the corresponding statistic $\mathcal{T}_{\mathcal{F}_i}$. The values of the test statistic define the aggregation weights $\{\omega_i, i = 1, \dots, d\}$, as in (16). Then, the estimate and the aggregation weight from the whole residual sequence $\mathcal{R}$ is computed at lines 12 - 15. Finally, the $d+1$ estimates are aggregated at line 16 to provide the ensemble output F.

As a guideline for building the CPM ensemble, we remark that, in particular when l is small, it is not worth considering large values of d, to reduce the probability of having individual estimates from the same subsequences: the upper bound for d is $\binom{l}{l/2}$, which, however, is often large.

V. Experiments

The experiments have been performed using the Lepage [15] test statistic, following the implementation presented in [3]. The Lepage CPM is a nonparametric test able to assess changes in the mean or in the variance between two populations. Unfortunately, it is not available any analytical formula to compute thresholds $h_{n,\alpha}$, and these are computed by means of Montecarlo simulations. The *CPM* package [19] implemented in R statistical software provides such thresholds.

Beside $\mathrm{CPM}_\mathcal{T}(\mathcal{R})$, i.e., the CPM executed on the whole residual sequence $\mathcal{R}$, we consider $\mathcal{E}_{10}, \mathcal{E}_{25}, \mathcal{E}_{50}, \mathcal{E}_{100}$, i.e., four ensembles of CPM aggregating different number of estimates from random sampling ($d \in \{10, 25, 50, 100\}$), and as a further comparison, we consider the $\mathrm{CPM}_\mathcal{T}(\mathcal{R})$

with threshold set to 0, thus estimating the fault-time instant irrespective of the maximum value of the test statistic (which is not negative). Obviously, when the threshold is 0, the $\mathrm{CPM}_\mathcal{T}(\mathcal{R})$ always provides an estimate of the fault-time instant without requiring any statistical confidence; otherwise, we set α to 0.05. Next, we describe how the residual sequences are generated and then we report the figures of merit used to assess the performance of the CPMs.

A. DataSet Generation

In our experiments, each sequence is composed of 700 samples, which are generated by an ARMA reference model which undergoes an abrupt change affecting the parameter vector at time instant $T^* = 500$. The parameter vectors determining the healthy (θ_0) and faulty (θ_1) conditions have been randomly generated (and only the stable ones kept). The orders (p, q) of the ARMA have also been randomly selected as $p \in \{1, \dots, 4\}$ and $q \in \{0, 1, 2\}$. The standard deviation of the innovation σ is set to 0.1. Results have been computed on a dataset containing 2000 of such sequences.

In addition, we considered three scenarios characterizing FDSs having different average detection delays. Specifically, we set $\widehat{T} = \{530, 540, 550\}$.

The first 400 observations of each sequence are provided as a training sequence and used to identify the parameter vector of the healthy process, $\widehat{\theta}_0$. We assume that the maximum orders of the process are known, and process identification is performed by estimating the best model for each order pair (p, q), considering $p \in \{1, \dots, 4\}$ and $q \in \{0, 1, 2\}$, and then selecting $\widehat{\theta}_0$ among these by using the Akaike information criterion [20]. It is important to remark that the initial training sequences (400 observations) are never considered for estimating the fault-time instant, and that all the CPMs ($\mathcal{E}_{10}, \mathcal{E}_{25}, \mathcal{E}_{50}, \mathcal{E}_{100}$, $\mathrm{CPM}_\mathcal{R}$, $\mathrm{CPM}_\mathcal{R}$ with threshold set to 0) are executed on sequences containing $130, 140, 150$ residuals.

B. Figures of Merit

For each value of $\widehat{T}$ we test $\mathcal{E}_{10}, \mathcal{E}_{25}, \mathcal{E}_{50}, \mathcal{E}_{100}, \mathrm{CPM}_\mathcal{R}$, and $\mathrm{CPM}_\mathcal{R}$ with threshold set to 0 on the whole residual dataset, and for each of them we compute the false negative rate (FNR) i.e., the percentage of sequences in the datasets where the CPMs are not able to estimate the fault-time instant (since in these cases the values of the statistic do not allow to reject the null hypothesis). FNR values are shown in the plots of Fig. 2. The false negative rate of $\mathrm{CPM}_\mathcal{R}$ with threshold set to 0 is not reported being is trivially 0.

To provide a compact overview of the distribution of the fault-time instant estimates, we report them using boxplots, in Fig. 3.

C. Discussion

Three main comments arise about the advantages of the use of the ensemble of CPMs from experimental results depicted in Fig. 2 and 3. First, enforcing an ensemble of CPMs allows us to significantly reduce the false negatives. Second, the ensemble of CPMs guarantees a lower statistical
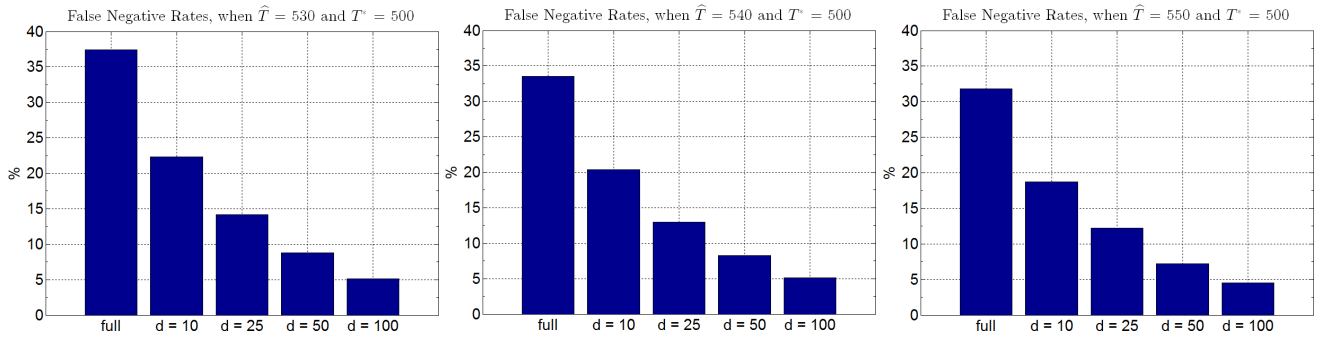
Fig. 2. The values of FNR for the three considered scenarios $\widehat{T} = \{530, 540, 550\}$.
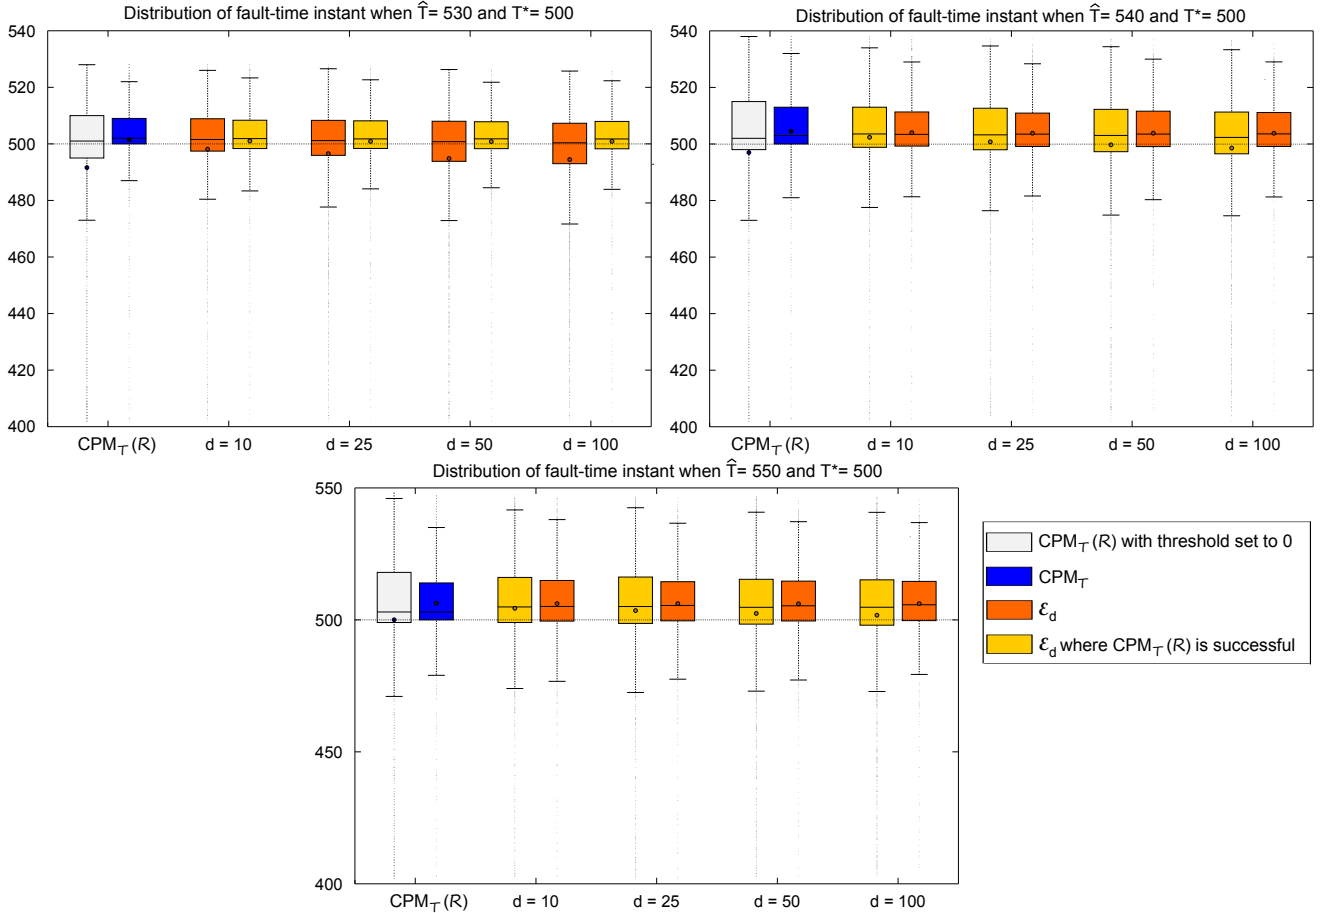

Fig. 3. The boxplot provides a compact representation of the empirical distribution of the fault-time instant estimates. The central line in each box represents the sample median, the bottom and top of each box represents the 25% and 75% quantile, respectively. The circular marker corresponds to the sample mean, while the whiskers are used to identify outliers.

dispersion than the $\mathrm{CPM}_{\mathcal{T}}(\mathcal{R})$ with threshold set to 0. Third, the ensemble estimates are comparable with those of $\mathrm{CPM}_{\mathcal{T}}(\mathcal{R})$ when this is successful. These comments are detailed and justified in the rest of the section.

Specifically, Fig. 2 shows that the ensembles $\mathcal{E}_d$ substantially reduce the FNR, in particular when d is large. Such FNR reduction can be motivated by the fact that the random sampling reduces the temporal dependency in $\mathcal{R}$ and, in practice, the CPMs on the subsequence may identify the fault-time instant when $\mathrm{CPM}_T(\mathcal{R})$ does not. In all the

considered solutions the FNR decreases with $\widehat{T}$, since the number of residuals with the fault increases.

The boxplots in Fig. 3 provide a compact overview of the empirical distributions of the fault-time instant estimated by the considered CPMs. In particular, the analysis of the sample mean and statistical dispersion of the estimates allows us to highlight the advantages of using the ensemble of CPMs.

Interestingly, all the considered CPMs suffers from a bias that increases with $\widehat{T}$, as the values of the sample mean (represented by circular markers) show in each boxplot. In

the i.i.d. case, increasing the number of faulty residuals typically ease the detection, as this consolidates the statistical evidence for rejecting the null hypothesis. In contrast, when the faulty data are not i.i.d. and are characterized by some sort of *structure*, the partitioning maximizing the test statistic may contain faulty residuals in both parts, and the estimated fault time instant may be significantly after T^*. This situation arises, for instance, when the faulty residuals follow trends. Therefore, increasing values of $\widehat{T}$ delay the fault-time instant estimates and, as a consequence, the sample means and the boxes are shifted upwards, see Fig. 3. Also the number of outliers (displayed with small gray dots out of the boxplot whiskers) in the right tails increase with $\widehat{T}$, and is particularly evident when $\widehat{T} = 150$.

More interesting considerations emerge from the statistical dispersions of the estimates of the fault-time instant. Fig. 3 shows that, in all the considered scenarios, $\mathrm{CPM}_T(\mathcal{R})$ with threshold set to 0 (light gray boxes) present the largest interquantile range (that determines the height of each box, as it collects 50% of the population about the sample median) and a substantial number of outliers. Because of these outliers, the sample mean falls below the 25% quantile when $\widehat{T} = 530$.

The dispersion of the ensemble estimates $\mathcal{E}_d$ (orange boxes) increases with d, and a similar consideration holds for the number of outliers. However, this effect mostly affects the first boxplot of Fig. 3, where only 30 faulty residuals are provided: here, the randomly sampled subsequences may not contain enough faulty residuals. Indeed, the distribution of the ensemble estimates becomes more compact and almost symmetric when $\widehat{T}$ increases, as shown in the other boxplots of Fig. 3.

The boxplots of Fig. 3 also report the distribution of the estimates provided by $\mathcal{E}_d$, on sequences where $\mathrm{CPM}_T(\mathcal{R})$ is successful (represented by the yellow boxplots). The dispersion and the location of these boxplots are very close to that of $\mathrm{CPM}_T(\mathcal{R})$, indicating that the aggregation phase successfully compensates possible inaccuracy in the individual estimates F_i. In fact, each of these estimates itself could be less accurate than the estimate from the whole dataset: just consider that the true fault-time instant T^* is not always included in the subsequences.

We also comment that the empirical distributions of all the CPMs estimates are characterized by outliers in the left tail. Such behavior is, in our opinion, due to model bias in the estimate of θ_0, yielding residuals that, during the healthy conditions, are not white.

VI. CONCLUSIONS

We introduce an ensemble method to estimate the fault-time instant in systems whose observations can be modeled by ARMA processes and that are subject to fault inducing an abrupt change in the parameter vector. We describe an ensemble of CPMs that relies on random sampling to reduce the temporal dependency in the analyzed residual sequence, since it may impair the performance of CPMs. The ensemble

of CPMs provides reduced FNR values w.r.t. to a single CPM executed on the whole residual sequence.

Though we are aware that such ensemble of CPMs is not tuned to operate with a predefined probability of type-I errors, we believe that the ensemble approach deserves further investigation, as it ease the estimation of the fault-time instant when residuals are not white.

The ensemble of CPMs can be straightforwardly adopted to estimate the fault-time instant when observations follow more general models, not only ARMA processes or dynamical systems. Ongoing works investigate different aggregation strategies to form the ensemble, techniques other than random sampling to compute the individual estimates, as well as integrating in the ensemble, CPMs relying on different test statistics.

REFERENCES

[1] R. Isermann, *Fault-diagnosis systems: an introduction from fault detection to fault tolerance*.   Springer Verlag, 2006.

[2] D. M. Hawkins and P. Qiu, "The changepoint model for statistical process control," *JOURNAL OF QUALITY TECHNOLOGY*, vol. Vol. 35, No. 4, pp. 355–366, 2003.

[3] G. J. Ross, D. K. Tasoulis, and N. M. Adams, "Nonparametric monitoring of data streams for changes in location and scale," *Technometrics*, vol. 53, no. 4, pp. 379–389, 2011.

[4] D. H. Timmer and J. J. Pignatiello, "Change point estimates for the parameters of an ar(1) process," *Quality and Reliability Engineering International*, vol. 19, no. 4, pp. 355–369, 2003.

[5] M. B. Perry and J. J. Pignatiello, "Identifying the time of step change in the mean of autocorrelated processes," *Journal of Applied Statistics*, vol. 37, no. 1, pp. 119–136, 2010.

[6] M. B. Perry, G. R. Mercado, and J. J. Pignatiello, "Phase II monitoring of covariance stationary autocorrelated processes," *Quality and Reliability Engineering International*, vol. 27, no. 1, pp. 35–45, 2011.

[7] A. Krogh and P. Sollich, "Statistical mechanics of ensemble learning," *Physical Review E*, vol. 55, no. 1, p. 811, 1997.

[8] T. Dietterich, "Ensemble methods in machine learning," *Multiple classifier systems*, pp. 1–15, 2000.

[9] Z.-H. Zhou, *Ensemble Methods: Foundations and Algorithms*.   Chapman & Hall, 2012.

[10] J. Wichard and M. Ogorzalek, "Time series prediction with ensemble models," in *Neural Networks, 2004. Proceedings. 2004 IEEE International Joint Conference on*, vol. 2.   IEEE, 2004, pp. 1625–1630.

[11] E. Mangalova and E. Agafonov, "Time series forecasting using ensemble of ar models with time-varying structure," in *Evolving and Adaptive Intelligent Systems (EAIS) 2012, IEEE Conference on*.

[12] L. Ljung, *System identification*.   Wiley Online Library, 1999.

[13] D. M. Hawkins and K. D. Zamba, "Statistical process control for shifts in mean or variance using a changepoint formulation," *Technometrics*, vol. Vol. 47, No. 2, pp. 164–173, 2005.

[14] A. M. Mood, "On the asymptotic efficiency of certain nonparametric two-sample tests," *The Annals of Mathematical Statistics*, vol. Vol. 25, No. 3, pp. 514–522, September 1954.

[15] Y. Lepage, "A combination of Wilcoxon's and Ansari-Bradley's statistics," *Biometrika*, vol. Vol. 58, No. 1, pp. 213–217, April 1974.

[16] D. M. Hawkins and K. D. Zamba, "A change-point model for a shift in variance," *Journal of Quality Technology*, vol. 37, no. 1, pp. 21–31, 2005.

[17] C. Alippi, G. Boracchi, and M. Roveri, "A hierarchical, nonparametric, sequential change-detection test," in *International Joint Conference on Neural Networks (IJCNN 2011)*, 31 2011-aug. 5 2011, pp. 2889 –2896.

[18] G. Ross and N. M. Adams, "Two nonparametric control charts for detecting arbitrary distribution changes," *Journal of Quality Technology*, vol. Vol 44, No. 22, pp. 102–116, 2012.

[19] G. J. Ross, "Parametric and nonparametric sequential change detection in R: The cpm package," *Journal of Statistical Software*, Forthcoming.

[20] H. Akaike, "A new look at the statistical model identification," *Automatic Control, IEEE Transactions on*, vol. 19, no. 6, pp. 716 – 723, dec 1974.