Correlation-based Detection of PMU Time Synchronization Attacks

Ezzeldin Shereen and György Dán
Department of Network and Systems Engineering
School of Electrical Engineering and Computer Science
KTH Royal Institute of Technology, Stockholm, Sweden

***Abstract*—Real-time monitoring and control in power systems is increasingly dependent on Phasor Measurement Units (PMUs). PMUs depend on precise time synchronization, and thus it is essential to ensure the security of time synchronization. In this paper we consider the detection of low-rate time synchronization attacks against PMUs. Based on a general clock model and a PMU measurement model we provide a closed form expression for the correlation between the clock frequency adjustments and the measured PMU phase angles in the absence of an attack. Leveraging the intuition that an attack would affect the correlation between these two quantities, we propose a model-based and a non-parametric correlation-based detector for time synchronization attacks. We evaluate the proposed detectors using extensive simulations. Our results show that they outperform traditional change detection techniques for clocks with low accuracy, for which attack detection is most challenging.**

I. Introduction

Many emerging real-time smart grid control applications depend on synchrophasor measurements taken by Phasor Measurement Units (PMUs). Examples include phase angle separation monitoring and power oscillation damping. For all these applications it is fundamental to ensure precise time synchronization for the PMUs, either using space-based (e.g., GPS) or network-based (e.g., PTPv2) time synchronization. Unfortunately, both space-based and network-based solutions have been shown to be vulnerable to attacks. GPS is vulnerable to spoofing attacks [1], while PTPv2 is vulnerable both to spoofing attacks and to the manipulation of the one-way delays [2]. Time synchronization attacks against PMUs could have serious impact [3], thus their detection and mitigation is a fundamental requirement.

A variety of recent works considered the detection of GPS spoofing attacks. Authors in [4] proposed two approaches for detecting anomalies for GPS-based time synchronization. The first approach leverages that PMU locations are fixed, hence each PMU should receive time information from a certain set of GPS satellites, and lets a PMU discard information received from a satellite that should not be in line of sight at a particular time. The second approach assumes that spoofed GPS signals will unlikely be consistent with the expected behavior of the clock. Authors in [5] proposed a method for detecting GPS spoofing based on the direction of arrival (DOA) of the GPS signal compared to the predicted DOA using an extended Kalman filter.

For network-based time synchronization, [6] proposed a game theoretical framework to model the interaction between an attacker and a defender in a PTP network. Links suspected of malicious activity are put in quarantine mode in order to investigate the cause of the anomaly, but the paper does not discuss how to find anomalies. Other works focused on the more general problem of detecting anomalies in the clock frequency [7], [8]. Authors in [7] proposed a technique for detecting anomalies in atomic clocks using the notion of Dynamic Allan Variance (DAVAR), which is a measure of the stability of the clock. Authors in [8] proposed a clock frequency jump detector, which operates by estimating the trend of the frequency evolution, de-trending the frequency, and then using the de-trended frequency to detect the jumps. Nevertheless, those techniques are best suited to detect clock faults, and not for detecting sophisticated low-rate attacks that are designed to bypass detection algorithms.

An alternative approach for detecting time synchronization attacks could be based on Linear State Estimation (LSE) if PMU measurements allow observability of the power system [9]. Given the system admittance matrix, the PMU measurements of voltage and current phasors can be used to estimate the system state, e.g., the voltage phasor on each bus of the network, and to identify bad data using some well known Bad Data Detection (BDD) technique, such as the largest normalized residual test. Intuitively, a compromised time reference should result in a BDD alarm, but recent work has shown that PMU time synchronization attacks can in fact remain undetected by BDD [9], [10].

A significant shortcoming of existing detection schemes is that they are either based solely on information from the time synchronization system, or on information from the power system. In this paper we show that the detection of time synchronization attacks can be improved by combining information from the power system and information from the PMU clock. Our approach is based on a power system process model on a clock frequency model, and relies on the correlation between the measured phase angle by the PMU and the periodical frequency adjustments made to the PMU clock. We provide an expression for the correlation without attack and propose a threshold based correlation test for detecting attacks. We use extensive simulations to validate our model and to show that the proposed correlation based detection method outperforms state of the art CUSUM based detection.

The rest of this paper is organized as follows. We introduce the measurement and clock models in Section II. We present analytical results for the correlation between the phasor measurement and the clock frequency adjustment in Section III. In Section, IV we show the accuracy of our analytical results and evaluate the performance of the proposed detection approach. Finally, we conclude the paper in Section V.

II. System Model

We consider a PMU that periodically measures a voltage or current phasor. We denote by $\alpha_p(t_i)$ the true value of the phase angle of the phasor at time instant t_i, and we denote by $O_p(t_i)$ the zero crossing time lag that this phase angle corresponds to compared to a reference signal. The correspondence is given by $\alpha_p(t_i) = 2\pi f O_p(t_i)$, where f is the mains frequency, e.g., a phase angle of $\alpha_p(t_i) = 0.573$ degree corresponds to $O_p(t_i) = 31.8\mu s$ of time in a 50Hz system. For convenience we will express all quantities in terms of time in the following.

A. Process and Measurement Model

We adopt a process and measurement model widely used in power systems [11]. The model assumes that the system state, and hence $O_p(t_i)$, follows a random walk

$$O_p(t_i) = O_p(t_{i-1}) + \omega_p(t_i), \tag{1}$$

where $\omega_p(t_i) \sim \mathcal{N}(0, \sigma_p)$ is zero mean normally distributed process noise. Hence, the instantaneous system state is the best predictor of the next state. The phase angle measurement is subject to additive measurement noise, which can be assumed to be normally distributed [11], thus if we assume that the PMU clock is perfectly synchronized then the phase angle measured by the PMU (expressed in terms of time) can be expressed as

$$\begin{aligned} O_z(t_i) &= O_p(t_i) + \omega_n(t_i) \\ &= O_p(t_{i-1}) + \omega_p(t_i) + \omega_n(t) \\ &= O_z(t_{i-1}) + \omega_p(t_i) + \omega_n(t_i) - \omega_n(t_{i-1}) \\ &= O_z(t_{i-1}) + \omega_z(t_i), \end{aligned} \tag{2}$$

where $\omega_n(t_i) \sim \mathcal{N}(0, \sigma_n)$ is the measurement noise, $\omega_z(t_i) = \omega_p(t_i) + \omega_n(t_i) - \omega_n(t_{i-1}) \sim \mathcal{N}(0, \sigma_z)$, and $\sigma_z = \sqrt{\sigma_p^2 + 2\sigma_n^2}$. That is, the sequence $\omega_z(t_i)$ consists of identically distributed (not independent) Gaussian random variables.

Our focus is on the practical case when the PMU clock is not perfectly synchronized, and we denote by $O_c(t_i)$ the time offset between the PMU clock and the reference time. Clearly, the time offset affects the measured phase angle, and thus the actual measurement becomes

$$O_m(t_i) = O_z(t_i) + O_c(t_i). \tag{3}$$

In what follows we consider the time offset $O_c(t_i)$, without and under a time synchronization attack.

B. Clock Drift Model

We model the evolution of the PMU clock frequency deviation using the Ornstein-Uhlenbeck (OU) process, which was shown to be a suitable model of the clock frequency evolution [12]. The OU process is a stationary Gauss-Markov process, and is defined as the solution to the stochastic differential equation

$$d\gamma(t) = \theta(\mu - \gamma(t))\, dt + \sigma_\gamma\, dW(t), \tag{4}$$

where $\gamma(t)$ is the clock frequency deviation at time t with respect to the nominal clock frequency, μ is the long term mean of the process, and $\theta > 0$ is the speed of reversion, which is a constant that determines how fast $\gamma(t)$ drifts to μ. Finally $W(t)$ is the Wiener process, and σ_γ is a measure of the fluctuation of $\gamma(t)$. By the Euler-Marugama method [13], (4) can be discretized and approximated as

$$\gamma(t_i) = \gamma(t_{i-1}) + \theta(\mu - \gamma(t_{i-1}))(t_i - t_{i-1}) + \omega_\gamma(t_i)\sqrt{t_i - t_{i-1}}, \tag{5}$$

where $\omega_\gamma(t_i) \sim \mathcal{N}(0, \sigma_\gamma)$ is the Wiener increment, and $t_i - t_{i-1}$ is the time duration between samples. As an example, in PTP it corresponds to the time between Synch messages, and is usally $t_i - t_{i-1} = 1$ second. Therefore, in the rest of the paper we use the indexes $t-1$ and t instead of t_{i-1} and t_i. In this case, (5) becomes

$$\gamma(t) = \gamma(t) + \theta(\mu - \gamma(t-1)) + \omega_\gamma(t). \tag{6}$$

Note that $\theta = 1$ corresponds to a Gaussian process $\gamma(t) = \mu + \omega_\gamma(t)$, and $\theta = 0$ corresponds to a pure random walk $\gamma(t) = \gamma(t-1) + \omega_\gamma(t)$ (which is not OU per definition). The OU process is a stationary process because it is always drifting to the long-term mean μ. For simplicity we assume that $\gamma(0) = \mu$. We can then model the evolution of the clock frequency as

$$\begin{aligned} \gamma(t) &= \gamma(t-1) + \theta(\gamma(0) - \gamma(t-1)) + \omega_\gamma(t) \\ &= \theta\gamma(0) + (1-\theta)\gamma(t-1) + \omega_\gamma(t), \end{aligned} \tag{7}$$

i.e., a weighted average of the long term mean $\gamma(0) = \mu$ and the previous value $\gamma(t-1)$ plus random noise $\omega_\gamma(t)$.

C. Clock Servo and Adjustments

The key component that synchronizes a PMU clock to its external time reference is the clock servo. Its purpose is to adjust the clock frequency so as to reduce the offset between the PMU clock and the time reference, while ensuring a smooth flow of time. Clock servos typically use a P-controller or a PI-controller for this purpose, such as the open source PTP implementation PTPd [14]. In what follows we consider a PI-controller, as a P-controller is a special case of that.

Every time a time reference is received from the time synchronization source, the clock servo computes the raw offset $\hat{O}_c(t) = \tau_{pmu}(t) - \tau_r(t)$ between the local time $\tau_{pmu}(t)$ at the PMU and the reference time $\tau_r(t)$. It then uses the raw offset for computing the clock frequency adjustment $A(t)$ as the output of a PI-Controller defined by

$$A(t) = D(t) + K_p\hat{O}_c(t), \tag{8}$$

where K_p is the proportional gain of the PI controller, and D is called the observed drift, which is the accumulated integrator error of the PI-controller given by

$$D(t) = D(t-1) + K_i \hat{O}_c(t), \tag{9}$$

where K_i is the integrator gain of the controller. Typical values for the controller gains are $K_p = 0.1$ and $K_i = 0.001$ [14]. The computed adjustment $A(t)$ determines the true offset $O_c(t)$ between the PMU clock and the correct time according to

$$O_c(t) = O_c(t-1) + \int_{t-1}^{t} \gamma(t)\, dt - A(t-1). \tag{10}$$

Note that the raw offset $\hat{O}_c(t)$ and the true offset $O_c(t)$ are not necessarily equal (e.g. in case of a time synchronization attack)

D. Attack Model

We consider an attacker that is capable of spoofing a time synchronization source (e.g. a GPS satellite signal, or a PTP synchronization message). The attacker is thus able to manipulate a series of timestamps $\tau_r^a(t), t \in \{t_1, t_2, \cdots, t_n\}$, where t_1 is the start time of the attack, and t_n is the end time of the attack. The attacked timestamps would consequently create a series of attacked raw offsets $\hat{O}_c^a(t)$.

Furthermore, we assume that the attacker is aware of the clock servo parameters $(K_p, K_i, D(t))$. This is a reasonable assumption since most clock servo implementations are open-source or can be reverse engineered. Therefore, the attacker is able to compute the effect of $\hat{O}_c^a(t)$ on the output adjustments $A^a(t)$ through (8), with high accuracy. Thus the attacker can have control over $A^a(t)$ implemented by the clock servo. At time instant $t_1 \leq t_k \leq t_n$, the attack will create a total offset of $O_c^a(t) - O_c(t) = \sum_{t=t_1}^{t_k} A(t) - A^a(t)$ between the PMU clock and the correct time, where $A(t)$ refers here to the adjustments that would have been implemented without the attack, and $O_c(t)$ refers to the true offset in case of no attack, which should normally be close to zero.

III. Attack detection using Correlation

In what follows we propose a detection approach for time synchronization attacks based on the correlation between the adjustment $A(t-1)$ introduced by the PMU clock servo at time $t-1$, and the change in the phase angle of the measured phasor at the subsequent time step, denoted by $\Delta O_m(t) = O_m(t) - O_m(t-1) = \Delta O_z(t) + \Delta O_c(t)$. Our approach is inspired by the observation that the correlation between these two quantities would be stable without an attack, but it would change due to an attack through the change in the sequence of adjustments $A^a(t)$, as we will show in the following.

A. Correlation Analysis

In what follows we present an approximate analysis of the correlation between $A(t-1)$ and $\Delta O_m(t)$ without an attack. To make the analysis tractable, we consider that the speed of reversion $\theta = 1$, which corresponds to that $\gamma(t)$ follows a Gaussian process. Furthermore, we consider that the clock servo uses a P-controller, thus $K_i = 0$ and $D(t) = 0, \forall t$. In Section IV we will quantify the effect of these assumptions on the accuracy of the analysis.

Proposition 1. *Consider that $\theta = 1, K_i = 0$. Then the correlation $\rho_{(\Delta O_m(t), A(t-1))}$ between $\Delta O_m(t)$ and $A(t-1)$ can be approximated as*

$$\begin{aligned} \tilde{\rho}_{(\Delta O_m(t), A(t-1))} &= \frac{cov((\Delta O_m(t), A(t-1)))}{\sigma_{\Delta O_m} \sigma_A} \\ &= \frac{E[(\Delta O_m(t) - \mu_{\Delta O_m})(A(t-1) - \mu_A)]}{\sigma_{\Delta O_m} \sigma_A}, \end{aligned} \tag{11}$$

where $cov((\Delta O_m(t), A(t-1))), \sigma_{\Delta O_m}$, and σ_A are given in (12)

Proof. Observe that for $\theta = 1$ eqn. (7) can be simplified as $\gamma(t) = \gamma(0) + \omega_\gamma(t)$, and for $K_i = 0$ the frequency adjustments can be computed as $A(t) = K_p O_c(t)$. Furthermore, we can approximate (10) using linear interpolation between $t-1$ and t to obtain

$$O_c(t) = O_c(t-1) + \frac{\gamma(t) + \gamma(t-1)}{2} - A(t-1). \tag{13}$$

This approximation is reasonable since γ is not expected to change rapidly during one second. Assuming $A(0) = O_c(0) = 0$, (13) becomes

$$\begin{aligned} O_c(t) &= \Big(\sum_{k=1}^{t-1} \frac{(1-K_p)^{t-k} + (1-K_p)^{t-k-1}}{2} * \omega_\gamma(k) \Big) \\ &+ \frac{\omega_\gamma(t)}{2} + \sum_{k=0}^{t} (1-K_p)^k \gamma(0). \end{aligned}$$

For very large t the last term converges to $\frac{\gamma(0)}{K_p}$. Therefore,

$$\begin{aligned} O_c(t) &= \Big(\sum_{k=1}^{t-1} \frac{(1-K_p)^{t-k} + (1-K_p)^{t-k-1}}{2} * \omega_\gamma(k) \Big) \\ &+ \frac{\omega_\gamma(t)}{2} + \frac{\gamma(0)}{K_p}. \end{aligned}$$

and hence,

$$\begin{aligned} \Delta O_c(t) &= \Big(\sum_{k=1}^{t-2} \frac{(1-K_p)^{t-k} - (1-K_p)^{t-k-2}}{2} * \omega_\gamma(k) \Big) \\ &+ \frac{(1-K_p)\omega_\gamma(t-1)}{2} + \frac{\omega_\gamma(t)}{2}. \end{aligned} \tag{14}$$

Using $\Delta O_m(t) = \Delta O_c(t) + \omega_z(t)$ and $A(t-1) = K_p O_c(t-1)$ in (11), and using the fact that $E[\omega_\gamma(i)\omega_\gamma(j)] = 0$ for $i \neq j$, and that $\mu_{\Delta O_m} = 0, \mu_A = \gamma(0)$, after some algebraic manipulation we obtain (12), which proves the result. □

Observe that the expression for $\tilde{\rho}(\Delta O_m(t), A(t-1))$ does not depend on σ_z and σ_γ, but on their ratio $\sigma^* = \frac{\sigma_z}{\sigma_\gamma}$. Therefore, we can write $\tilde{\rho}_{(\Delta O_m(t), A(t-1))} = f_\rho(t, \sigma^*, K_p)$.

$$
\begin{aligned}
cov((\Delta O_m(t), A(t-1))) &= K_p\Big(\frac{\sigma_\gamma}{2}\Big)^2 * \left[\frac{(1-(1-K_p)^{2t-4})((1-K_p)^3+(1-K_p)^2-(2-K_p))}{K_p(2-K_P)}+1-K_p\right], \\
\sigma^2_{\Delta O_m} &= \Big(\frac{\sigma_\gamma}{2}\Big)^2 * \left[\frac{(1-(1-K_p)^{2t-4})(1-(1-K_p)^2)^2}{K_p(2-K_p)}+(1-K_p)^2+1\right]+\sigma_z^2, \\
\sigma^2_{A} &= \Big(\frac{K_p\sigma_\gamma}{2}\Big)^2 * \left[\frac{(1-(1-K_p)^{2t-4})(2-K_p+(1-K_p)^2))}{K_p(2-K_p)}+1\right]. \qquad (12)
\end{aligned}
$$

B. Correlation-based Time Synchronization Attack Detection

In what follows we propose two correlation-based approaches for the detection of time synchronization attacks against PMUs: a model based (parametric) approach, and a non-parametric approach. Both approaches are based on computing the difference between a predicted correlation $\hat{\rho}$ and the measured one.

1) Model-based Approach: In the model-based approach the predicted correlation $\hat{\rho}$ is computed using the correlation analysis presented in Section III-A. First, we estimate σ_γ and σ_z (hence σ^*) based on previous knowledge about the PMU clock accuracy and the power system stability. The predicted correlation can then be computed as $\hat{\rho} = f_\rho(t, \sigma^*, K_p)$ for a sufficiently high value of t, where K_p is a known clock servo parameter. Note that the accuracy of the estimated σ^* plays a crucial role in the accuracy of the predicted correlation.

2) Non-Parametric Approach: In the non-parametric approach the predicted correlation $\hat{\rho}$ is computed from the historical values of $A(t-1)$ and $\Delta O_m(t)$ from the actual system, when the PMU is known to be in a normal state (non-attacked). The advantage of this approach is that it does not depend on the accuracy of the system model, but it requires the system to be in a non-attacked state for the computation of $\hat{\rho}$.

Given $\hat{\rho}$, for both approaches, upon every time step t we compute the correlation $\rho_N(t)$ on-line based on $A(\tau-1)$ and $\Delta O_m(\tau)$, $\tau = t-N+1, \ldots, t$, where N is the correlation window length. An alarm is raised if $|\rho_N(t) - \hat{\rho}| > \eta_\rho$, where η_ρ is the detection threshold.

IV. Numerical Results

In this section we evaluate our proposed detection approaches. First, we show that $f_\rho(t, \sigma^*, K_p)$ is an accurate estimate of the correlation $\rho_N(t)$ under the considered assumptions, and is fairly accurate even when the assumptions are relaxed. We then evaluate the proposed detection approach and compare it to traditional change detection methods used for clock anomaly detection, and show that the proposed approach performs better for low-accuracy clocks, for which anomaly detection is most challenging.

A. Correlation Accuracy for $\theta = 1, K_i = 0$

To assess the accuracy of $f_\rho(t, \sigma^*, K_p)$, we simulated a clock with $\theta = 1$ and a P-controller clock servo ($K_i = 0$). The results reported are the averages of 5000 simulations of a PMU according to the previously mentioned measurement and clock models. Figure 1 shows $f_\rho(t, \sigma^*, K_p)$ and $\rho_N(t)$ as a function of σ^* for various window sizes N. The figure shows that $f_\rho(t, \sigma^*, K_p)$ hardly changes when σ^* is either very high or very low, and thus would be insensitive to estimation errors of σ^* in these regions. The figure also shows that the accuracy of $f_\rho(t, \sigma^*, K_p)$ is increasing as N increases. The reason for the discrepancy between $f_\rho(t, \sigma^*, K_p)$ and the emprical correlation $\rho_N(t)$ is due to the difference between time averages and ensemble averages. While $\rho_N(t)$ is the correlation between a sequence of N pairs of random variables, $f_\rho(t, \sigma^*, K_p)$ is the correlation between $\Delta O_m(t)$ and $A(t-1)$. To obtain the equivalent of $f_\rho(t, \sigma^*, K_p)$ numerically one would have to compute the correlation between $\Delta O_m(t)$ and $A(t-1)$ across multiple simulations. To show that this is indeed the case, Figure 1 also shows $\rho_{runs}(t)$ computed across simulations, and shows that this ensemble average correlation is an excellent match for $f_\rho(t, \sigma^*, K_p)$. Observe that $\rho_{runs}(t)$ can not be computed in practice, as it is impractical to run multiple copies of the same time synchronization system, but ρ_N can be computed efficiently. Overall, the results show that $f_\rho(t, \sigma^*, K_p)$ is a good approximation of $\rho_N(t)$ as long as the window size N is large enough. In what follows we use $N = 200$, unless otherwise noted.

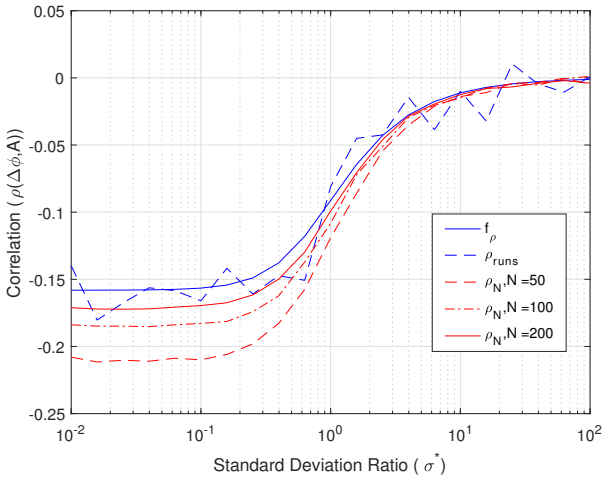

Fig. 1. Comparison of $f_\rho(t, \sigma^*, K_p)$, ρ_{runs}, and ρ_N vs. σ^*. $K_p = 0.1$, and $t = 2000$.

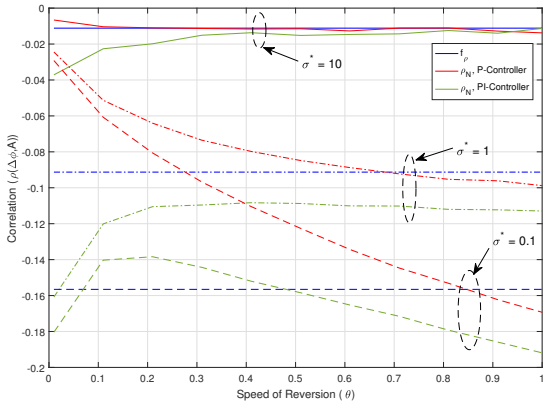

Fig. 2. Comparison of $f_\rho(t, \sigma^*, K_p)$, and ρ_N vs. θ. $\sigma^* = \{0.1, 1, 10\}$, $K_p = 0.1$, $N = 200$, and $t = 2000$.

B. Correlation accuracy for $\theta \leq 1$ and $K_i \geq 0$

To evaluate the sensitivity of $f_\rho(t, \sigma^*, K_p)$ to θ and K_i, Figure 2 shows f_ρ and ρ_N as a function of θ, for various values of σ^*, and for $K_i = 0$ (P-Controller) and $K_i = 0.001$ (PI-Controller, as used in PTPd). For $K_i = 0$ the figure shows that f_ρ is a reasonably good approximation for low θ, and its accuracy increases with θ. For $K_i = 0.001$ the results show that f_ρ is a reasonably good approximation, especially for $\sigma^* = 10$, but the trend is less obvious, because the approximation error has three sources; (1) the approximation of the ensemble average with the time average (c.f., Fig 1), (2) the error due to $\theta \neq 1$, and (3) the error due to $K_i \neq 0$. Overall, we can conclude that f_ρ is close to the empirical correlation ρ_N despite the approximation.

C. Attack Detection

Finally, we turn to the detection of time synchronization attacks using the proposed detectors. For the evaluation we consider two PMU clocks: a very accurate clock ($\gamma(0) = 100ns, \sigma_\gamma = 10ns$) referred to as clock A, and a less accurate clock ($\gamma(0) = 1\mu s, \sigma_\gamma = 100ns$) referred to as clock B. For both clocks we use a PI-controller ($K_p = 0.1, K_i = 0.001$) and a speed of reversion of $\theta = 10^{-6}$, which is a reasonable value for real clocks, which usually have weak tendencies to revert to the long-term mean frequency. For the measurement we consider $\sigma_z = 2.2\mu s$, which is realistic assuming the use of a class 0.1 sensor in the PMU. We then obtain $\sigma^* = 220$ and $\sigma^* = 22$ for clock A and clock B, respectively.

In order to simulate an attack we follow the following procedure. Every attack starts at $t_1 = 400$ seconds, and manipulates the raw offset sequence such that it follows $\hat{O}_c^a(t) \sim \mathcal{N}(\hat{O}_c(t) + \mu_a, \sigma_a)$, where $\mu_a \neq 0$, and $\hat{O}_c(t)$ represents the corresponding computed raw offset if no attack was present. This allows us to model an attacker that intends to accelerate ($\mu_a > 0$) or decelerate ($\mu_a < 0$) the clock. The attacker stops the attack at time $t_n = 3400$ seconds. Figure 3 shows the observed values of ρ_N in various attack scenarios (different μ_a and σ_a) as a function of time. We observe that when μ_a is relatively high, the correlation value changes significantly when the attack starts and when it ends.

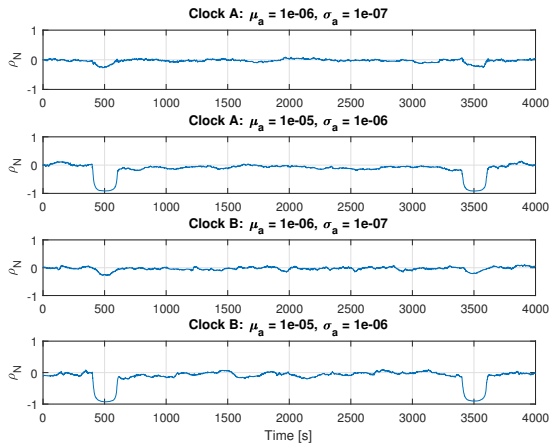

Fig. 3. Time series of ρ_N under different attack scenarios for clocks A and B. $K_p = 0.1, K_i = 0.001, \theta = 10^{-6}$, and $N = 200$.

Next, to assess the proposed correlation-based detection approaches, we created multiple simulated datasets of the time series $A(t)$ and $\Delta O_m(t)$. Each dataset corresponds to different values of μ_a and σ_a (we consider only $\sigma_a = \mu_a/10$), and consists of 2000 simulations, each of length 4000 seconds. In the first 1000 simulations, the PMU clock is running normally, without an attack. In the second 1000 simulations, we simulate an attack that starts at $t_1 = 400$ seconds and ends at $t_2 = 3400$ seconds as before. We then applied the proposed correlation detection approaches to the data sets. As a baseline we used the cumulative sum (CUSUM) method, which is a state of the art change detection algorithm, on the adjustment sequence $A(t)$. To implement CUSUM over $A(t)$, we computed the mean adjustment $\bar{A}$ and the standard deviation s_A over the first 200 seconds of each time series. An alarm is then generated if the cumulative sum of $(A(t) - \bar{A})$ exceeds $\eta_C s_A$, where η_C is the detection threshold.

To compare the detection approaches, we utilize the Receiver Operating Characteristic (ROC) curve to compute the area under the ROC curve (AUC). A ROC curve shows the tradeoff between the true positive rate (TPR) and the false positive rate (FPR) of a detector, where the TPR is defined as the ratio of attacked simulations for which at least one alarm is raised, while FPR is defined as the ratio non-attacked simulations during which at least one alarm is raised. It is worth mentioning that a ROC curve is not obtained by using only one value of the detection threshold. Instead, each point in the ROC curve corresponds to one value of either η_ρ or η_C. The AUC is computed based on the ROC curve, and is a widely accepted measure of the performance of a detector. A perfect detector would have AUC=1, i.e., there exists a threshold value for which TPR=1 and FPR=0. Figure 4 shows the AUC obtained by implementing the detection approaches on the datasets described above. For the sake of comparison we show results for three model-based correlation detectors, each using a different thresholding technique. Detector 1 (D1) raises an alarm if any $\rho_N(t)$ satisfies the detection condition $|\rho_N(t) - \hat{\rho}| > \eta_\rho$. Detector 2 (D2) raises an alarm only if there exist at least 10 successive values of $\rho_N(t)$ that satisfy

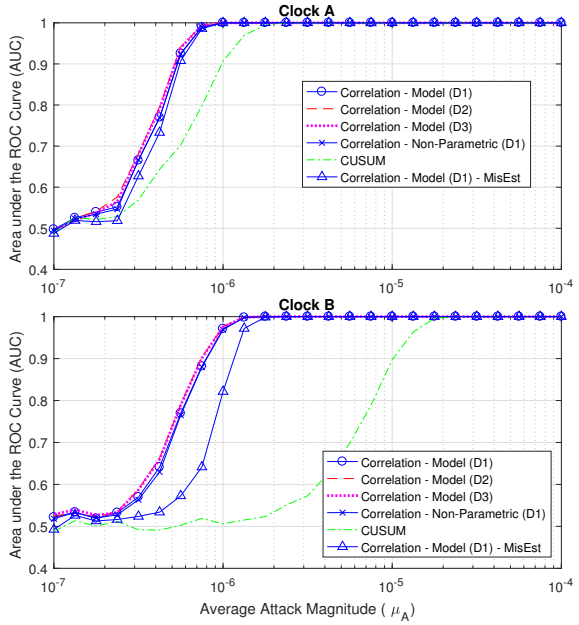

Fig. 4. Comparison of the detection performance between CUSUM and the correlation-based approaches, for clocks A and B. $K_p = 0.1, K_i = 0.001, \theta = 10^{-6}$, and $N = 200$.

the detection condition. Detector 3 (D3) divides data into non-overlapping windows, each of length 200 seconds, and raises an alarm if more than 10% of the $\rho_N(t)$ values in any window satisfy the detection condition. The figure shows that the performance of the three detectors are very similar. In what follows we only consider D1 since it is simpler to implement. The results also show that even though we simulate a PI-Controller and $\theta = 10^{-6}$, the model based approach is comparable to the non-parametric approach. In fact in most of the cases, the model based approach achieves slightly better performance. Another observation is that all detection approaches perform better for clock A than for clock B, which is expected, as clock A is a very stable clock and thus any change is easier to detect.

Importantly, the proposed correlation-based approaches perform significantly better than CUSUM for both clocks, especially for clock B. Recall that clock B is not very stable, which makes it hard for CUSUM to avoid false positives. However, this is not an issue for the correlation-based approaches, since they do not only rely on the clock stability, but also on the phasor measurements. Finally, we evaluate the sensitivity of the model-based approach to a potential mis-estimation of σ^*, and hence $f_\rho(t, \sigma^*, K_p)$. Clock A and clock B are are mis-estimated by one order of magnitude ($\sigma^* = 22$ and 2.2, respectively). The results ("MisEst" curve) show that the detector is less sensitive to mis-estimation for clock A than for clock B. This is due to the fact that the change in $f_\rho(t, \sigma^*, K_p)$ is very limited when σ^* is high (c.f., Fig 1). However, in both cases, we observe that the non-parametric approach provides a better alternative when it is hard to estimate σ^*. In general, based on the observed results we can conclude that the proposed correlation-based detection approaches work well, especially for clocks with low accuracy, for which traditional change detection techniques, such as CUSUM, do not.

V. Conclusions

In this paper we have considered the problem of detecting time synchronization attacks on PMUs. We provided a cyber-physical model of PMU phase angle measurements, which allowed us to express the correlation between the clock frequency adjustments implemented by the PMU clock and the measured phase angles, in the absence of time synchronization attacks. Based on the intuition that an attack would affect this correlation, we proposed two correlation-based attack detectors. We used extensive simulations to evaluate the proposed detectors, and showed that they are effective compared to traditional change detection schemes especially for low accuracy clocks. Our approach shows that attack detection based on cyber-physical models can be a useful complement to traditional anomaly detection methods for increasing the security of future power systems.

References

[1] X. Jiang, J. Zhang, B. Harding, J. Makela, and A. Dominguez-Garcia, "Spoofing GPS receiver clock offset of phasor measurement units," *IEEE Trans. on Power Systems*, vol. 28, no. 3, pp. 3253–3262, Aug 2013.

[2] S. Barreto, A. Suresh, and J. LeBoudec, "Cyber-attack on Packet-Based time synchronization protocols: the undetectable delay box," in *IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, Taipei, Taiwan, May 2016.

[3] M. S. Almas, L. Vanfretti, R. S. Singh, and G. M. Jonsdottir, "Vulnerability of synchrophasor-based wampac applications to time synchronization spoofing," *IEEE Trans. on Smart Grid*, pp. 1–1, 2017.

[4] F. Zhu, A. Youssef, and W. Hamouda, "Detection techniques for data-level spoofing in GPS-based phasor measurement units," in *2016 International Conference on Selected Topics in Mobile Wireless Networking (MoWNeT)*, Apr 2016, pp. 1–8.

[5] C. H. Kang, S. Y. Kim, and C. G. Park, "Adaptive complex-EKF-based DOA estimation for GPS spoofing detection," *IET Signal Processing*, vol. 12, no. 2, pp. 174–181, 2018.

[6] E. Lisova, E. Uhlemann, W. Steiner, J. Åkerberg, and M. Björkman, "Game theory applied to secure clock synchronization with IEEE 1588," in *2016 IEEE Intl. Symp. on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS)*, Sep 2016, pp. 1–6.

[7] E. Nunzi, L. Galleani, P. Tavella, and P. Carbone, "Detection of anomalies in the behavior of atomic clocks," *IEEE Trans. on Instrumentation and Measurement*, vol. 56, no. 2, pp. 523–528, Apr 2007.

[8] L. Galleani and P. Tavella, "Robust detection of fast and slow frequency jumps of atomic clocks," *IEEE Trans. on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 64, no. 2, pp. 475–485, Feb 2017.

[9] S. Barreto, M. Pignati, G. Dán, J. L. Boudec, and M. Paolone, "Undetectable timing-attack on linear state-estimation by using rank-1 approximation," *IEEE Trans. on Smart Grid*, vol. 9, no. 4, pp. 3530–3542, July 2018.

[10] S. B. Andrade, J. L. Boudec, E. Shereen, G. Dán, M. Pignati, and M. Paolone, "A continuum of undetectable timing-attacks on pmu-based linear state-estimation," in *IEEE Intl. Conf. on Smart Grid Communications (SmartGridComm)*, Oct 2017, pp. 473–479.

[11] L. Zanni, S. Sarri, M. Pignati, R. Cherkaoui, and M. Paolone, "Probabilistic assessment of the process-noise covariance matrix of discrete Kalman filter state estimation of active distribution networks," in *2014 Intl. Conf. on Probabilistic Methods Applied to Power Systems (PMAPS)*, Jul 2014, pp. 1–6.

[12] E. Bibbona, G. Panfilo, and P. Tavella, "The Ornstein-Uhlenbeck process as a model of a low pass filtered white noise," *Metrologia*, vol. 45, p. S117, Dec 2008.

[13] G. Maruyama, "Continuous markov processes and stochastic equations," *Rend Circ Math Palermo*, 1955.

[14] K. Correll, N. Barendt, and M. Branicky, "Design considerations for software only implementations of the IEEE 1588 precision time protocol," in *Conf. on IEEE 1588 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, 2006.