Feasibility of Time Synchronization Attacks against PMU-based State-Estimation

Ezzeldin Shereen, *Student Member, IEEE*, Marguerite Delcourt, *Student Member, IEEE*, Sergio Barreto, *Member, IEEE*, György Dán, *Senior Member, IEEE*, Jean-Yves Le Boudec, *Fellow, IEEE*, Mario Paolone, *Senior Member, IEEE*

***Abstract*—The emerging measurement technology of phasor measurement units (PMUs) makes it possible to estimate the state of electrical grids in real-time, thus opening the way to new protection and control applications. PMUs rely on precise time synchronization, therefore they are vulnerable to time-synchronization attacks which alter the measured voltage and current phases. In particular, undetectable time synchronization attacks pose a significant threat as they lead to an incorrect but credible estimate of the system state. Prior work has shown that such attacks exist against pairs of PMUs, but they do not take into consideration the clock adjustment performed by the clock-servo, which can modify the attack angles and make the attacks detectable. This cannot easily be addressed with the existing attacks, as the undetectable angle values form a discrete set and cannot be continuously adjusted as would be required to address the problems posed to the attacker by the clock servo. Going beyond prior work, this paper first shows how to perform undetectable attacks against more than two PMUs, so that the set of undetectable attacks forms a continuum and supports small adjustments. Second, it shows how an attacker can anticipate the operation of the clock servo while achieving her attack goal and remaining undetectable. Third, the paper shows how to identify vulnerable sets of PMUs. Numerical results on the 39-bus IEEE benchmark system illustrate the feasibility of the proposed attack strategies.**

I. Introduction

The control and operation of interconnected power grids often require the timely knowledge of the system state. Accurate information on the state enables or improves the performance of fundamental functions, such as security assessment, voltage control and stability analysis. Legacy measurement technologies have low measurement and streaming rates which induces a relatively low refresh rate of the system state-estimation. Nonetheless, the emerging measurement technology of PMUs makes it possible to acquire phasors that are accurate, time and phase-aligned (i.e., synchrophasors) with streaming rates of the order of tens of measurements per second [1]–[3]. PMUs require, however, precise time synchronization [4], which can be achieved by GPS [5] or by network-based time synchronization protocols, such as PTP [6]. The need for time synchronization is a weakness as existing time-synchronization techniques are known to be vulnerable to cyber-attacks [7], GPS-based synchronization is vulnerable to spoofing attacks [8] while packet-based protocols are vulnerable to delay-box insertion on transmission-lines [9]. A delay-box is a repeater-box that is inserted on the fiber, replacing the bidirectional fiber by a short fiber in one direction and a long fiber coil in the reverse direction, thus adding an asymmetric delay. Such an attack bypasses any authentication scheme used by the synchronization protocol. Importantly, none of these attacks require physical access to PMUs. Therefore, in order to assess the vulnerability of synchrophasor-based applications, in particular the state estimation of a system, it is essential to explore the feasibility and detectability of PMU time-synchronization attacks (TSAs). TSAs are of interest to other time-sensitive networks as well. Sensor networks used for source-localization, require an accurate synchronization [10], attacking the time-reference of the sensors alters the post-processing of measurements, yielding a false estimation of the location. As a consequence, an attacked boat or plane could enter on the wrong territory and an attacked sensor network could be unable to track a source. Collaborative robots in automation systems also rely on synchronization [11] and are increasingly deployed in the industry.

A significant advantage of measuring phasors with PMUs is that the state estimation problem becomes linear. This is because PMUs are capable of directly measuring voltage and current phasors, compared to traditional supervisory control and data acquisition (SCADA) measurements (e.g., the power-flow and power injections, which are nonlinear functions of the system state). A widespread technique for making the linear state estimation (LSE) more robust against attacks is to couple it with a bad-data detection (BDD) scheme (e.g., χ^2 and the largest normalized residual tests (LNR)). Nonetheless, the seminal work in [12] shows that LSE is vulnerable to bad-data injections that bypass the different BDD algorithms (both χ^2 and LNR). Subsequent articles focus on the characterization and mitigation of undetectable attacks, through either prevention or detection. The authors of [13], [14] focus on establishing an index of security that quantifies the vulnerability of sets of PMUs and on smart PMU allocation to mitigate it. Malicious undetectable attacks are proposed in [15] and techniques to identify vulnerable meters are given

E. Shereen and G. Dán are with the School of of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden. e-mail: {eshereen| gyuri}@kth.se.

M. Delcourt and J-Y. Le Boudec are with the School of Computer and Communication Sciences of the Swiss Federal Institute of Technology Lausanne, EPFL, Switzerland. e-mail: {marguerite.delcourt| jean-yves.leboudec}@epfl.ch.

S. Barreto is with Cisco Systems, Core Software Group, Lausanne, Switzerland. During the work on this paper, he was with the School of Computer and Communication Sciences of the Swiss Federal Institute of Technology Lausanne, EPFL, Switzerland. e-mail: sbarreto@cisco.com.

M. Paolone is with the School of Engineering of the Swiss Federal Institute of Technology Lausanne, EPFL, Switzerland. email: mario.paolone@epfl.ch

in [16]. Globally, false-data injection attacks require physical access to monitoring devices that are often located in secure facilities. In contrast, the authors of [17] use GPS spoofing to exclusively manipulate the time reference used by PMUs, thus shifting the phase of their measured synchrophasors, without requiring any physical access to PMUs. The attack aims at altering various smart-grid applications but does not tackle the issue of being undetected by BDD algorithms to impact durably the LSE.

The authors of [18] address this important question by showing how to compute an undetectable TSA against pairs of PMUs. They do so by solving a specific set of non-linear equations, yielding a discrete and finite set of attacks. In the case where these equations lead to no attack solution, they propose an approximated set of equations that yields an attack solution. Importantly, the attacks work on the actual LSE with complex values, not on a linear approximation of a non-linear state-estimation problem. The authors also discuss how attacks against pairs of PMUs can be combined to maximize the attacker's objective, which could be to maximize the error on a particular line power-flow, for instance. However, the clock adjustment rate of a PMU is controlled in practice by a controller typically called the clock servo [19]. The latter ensures that clock adjustments always stay below a well-defined threshold. Hence, a delay above this threshold will not be implemented by the clock servo. Instead, it will transform the intended delay into a smaller one which may not be in the discrete and finite set of undetectable attacks. Therefore, if an attacker blindly performs an attack on pairs of PMUs according to the results of [18], without anticipating the actions of the clock servo, it is likely that the intended attack-angles will be modified into smaller detectable ones. Consequently, the impact on the measurements might not be the one intended by the attacker and the attack could become detectable by the BDD algorithms. In order to overcome the limitations of the results of [18], three major contributions are made in this paper.

First, it is shown that the set of undetectable TSAs against three or more PMUs forms a connected compact set and that the number of valid attacks is uncountably infinite. This allows the attacker to anticipate the actions of the clock servo and to remain undetected by injecting small incremental delays over a period of time until the objective is maximized. It is proven that attacks can be tailored to remain undetected by reaching an optimal attack-angle in a continuous manner. It is also shown how to compute this set of valid attacks.

The second contribution of this paper is to address the practical feasibility of performing TSAs by taking PMU clock constraints in consideration. Algorithms that can achieve practically undetectable attacks under realistic conditions are proposed. In order to find attack-angles against a specific set of PMUs, it is required to build an attack angle matrix using the system topology, the measurements and the choice of PMUs to attack. When this matrix is of rank approximately or exactly equal to 1, the results of this paper show how to perform an attack. The vulnerability of a set of PMUs depends on how close the rank of the corresponding matrix is to 1. This measure is captured by the index of separation (IoS) introduced in [18], which only considers pairs of PMUs.

As a third contribution, it is shown how arbitrary sized sets of vulnerable PMUs can be found. The theory is extended in order to show that synchrophasors can be grouped in equivalence classes and that members of a class form a vulnerable set. This sufficient condition allows to efficiently find sets of PMUs to attack simultaneously. By analysing the infimum of the IoS over all measurements, it is shown that vulnerable sets can be identified based on the system topology only, without having to read measurement values. Furthermore, the paper provides evidence that under some topological conditions it is possible to attack PMUs that measure both voltage and current phasors (both sharing the same time reference). Finally, it also provides numerical evidence that attacks can be feasible against the time synchronization of PMUs whose measurements are not exactly critical.

This paper is an extension of a preliminary conference version [20], where the practical feasibility of TSAs was not considered. Contrary to [20], this paper shows that TSAs can be mounted undetectably by satisfying the constraints imposed by the PMU clock servo. In addition, it provides a sufficient condition for finding critical groups of measurements of arbitrary size, and it also shows that vulnerable sets without this condition exist. Finally, going beyond results in [20] it shows that it is possible to attack PMUs that measure two distinct synchrophasors simultaneously, and shows that the attacks bypass robust state estimation as well.

The rest of the paper is organised as follows. Section II describes the same system model as in [18] and a stronger attack model. The new contributions start in Section III which gives expressions to compute the set of possible attack-angles. In Section IV results on how to efficiently find sets of vulnerable PMUs are explained. In Section V, the practical feasibility of deploying an undetectable attack is discussed by considering multiple strategies. Numerical results are presented in Section VI to illustrate the effectiveness of the attack. Countermeasures against TSAs are discussed in Section VII. Finally, Section VIII concludes the paper.

II. System Model

The system model used throughout this paper is the same as the one used in [18], it is repeated here for completeness. A balanced transmission system that consists of N buses equipped with a PMU based measurement system is considered. Let $\mathcal{M}^V$ and $\mathcal{M}^I$ be such that $\mathcal{M} = \mathcal{M}^V \cup \mathcal{M}^I$ is the set of all voltage and nodal-current measured synchrophasors, and $M = |\mathcal{M}|$. Assume that $\mathcal{M}^V \cap \mathcal{M}^I = 0$, meaning that a PMU is dedicated to the measurement of a specific synchrophasor through time (voltage or current phasor, but not both). It is shown numerically in Section VI that it is still possible to perform undetectable attacks without this assumption. Given the $M \times N$ measurement matrix H, the measurement model is $z = Hx + e$, where $x \in \mathbb{C}^N$ is the system state (voltage phasors on all buses), $z \in \mathbb{C}^M$ the measurement vector (measured voltage and current phasors by PMUs), and $e \in \mathbb{C}^M$ the complex measurement-error. In this paper, the weighted least squares (WLS) estimator is used for the LSE. The verification matrix is defined as follows

$$F \triangleq H(H^\dagger H)^{-1}H^\dagger - I, \qquad (1)$$

where $H^\dagger$ is the conjugate transpose of H. Then, the residuals are defined by $r = Fz$ as they correspond to the difference between the observed measurements and the ones that should be observed if the state estimation is exactly correct. Clearly, $Fz = 0$ occurs if and only if there exists some state x such that $z = Hx$.

A. Attack Model

The considered attacker is able to manipulate the time synchronization of $p \geq 2$ PMUs, via GPS spoofing or delay-box insertion on transmission-lines, such that the time reference of an attacked PMU is delayed or advanced. This is equivalent to introducing p attacking angles $\alpha_i, i = 1 : p$, which correspond to the phase angle shifts of the synchrophasors measured by the attacked PMUs. The equivalence is given by $\Delta t_i = \frac{\alpha_i}{2\pi * f * 10^{-6}} = g_\alpha(\alpha_i)$, where Δt_i is the offset caused by the attack given in μs, and $f \approx 50Hz$ is the instantaneous voltage signal frequency. It is assumed that a time reference affects only one synchrophasor location. Thus, for every $i \in \{1, \cdots, p\}$, an attack changes the measured phasor z_i to $z_i' = z_i\, u_i$, where $u_i = e^{j\alpha_i} \in \mathbb{T}$, and $\mathbb{T}$ is the set of complex numbers of modulus 1, therefore the phasor magnitude is unchanged. Note that TSAs are multiplicative whereas traditional false-data injection attacks [12] are additive in nature. To identify targeted measurements, let Ψ be the $M \times p$ attack-measurement indicator matrix, defined by

$$\Psi_{m,i} = \begin{cases} 1 & \text{if } \alpha_i \text{ targets } z_m, \\ 0 & \text{otherwise.} \end{cases}$$

It is supposed that the attacker knows H and can observe the synchrophasors z. The only feature added to the attack model of [18] is that it is further supposed that the attacker is able to anticipate the actions of a regular PMU clock servo. This attack model is strong for two reasons. First, recent attacks on critical infrastructures, e.g., Stuxnet [21], had access to detailed system information. Second, such a strong model enables engineers to identify vulnerable data that require protection.

With such capabilities, the attacker's goal is to compromise the LSE provoking wrong power attribution. For instance the objective could be to provoke a black out by making the system over or under-estimate the power in a region of the grid.

B. Undetectability Condition

The condition for undetectability is that the attack must not modify the residuals. In other words, the residuals obtained from the state-estimation with the attacked and unattacked measurements must be identical. Hence, the resulting state-estimation while being false, remains plausible in the sense that it could be the result of an estimation after a natural trajectory of the grid. Therefore, it is expected that no BDD scheme based on residual analysis, is able to detect the proposed attacks. The vast majority of state-of-the-art BDD algorithms are variants of the two most widespread techniques, namely the LNR test and the χ^2-test [22]–[24]. Both techniques make use of the fact that the residuals are typically distributed as a Gaussian distribution with zero mean and easily-computable standard deviation [23]. In the LNR test, if the largest normalized residual is above a certain threshold η_{BDD}, then its associated measurement is marked as potential Bad Data (BD). The χ^2-test (Chi-squared test) is an alternative bad-data detection method that exploits the property that the sum of normally-distributed random variables is a variable with a χ^2 distribution and a certain number of degrees of freedom. If the sum of the residuals does not follow this distribution with a certain confidence level, one or more of the measurements are suspected to be corrupt. Since an undetectable attack does not modify residuals, it does not impact the distribution of residuals and thus these tests are expected to fail.

The condition for attack undetectability translates to: $Fz = Fz'$, which is made more tractable by the authors of [18] by introducing the $p \times p$ attack-angle matrix W, as the Hermitian complex matrix given by

$$W \triangleq \Psi^T \text{ diag}(z)^\dagger F^\dagger F \text{ diag}(z)\Psi. \qquad (2)$$

As shown in (Theorem 1, [18]), an attack $\alpha = (\alpha_1, \ldots, \alpha_p)$ is undetectable if and only if

$$W(\vec{u} - \vec{1}) = 0, \qquad (3)$$

where $\vec{u} = (u_1, ..., u_p)^T$, and $\vec{1} = (1, ..., 1)^T$. Equation (3) is called the undetectability condition for an attacking vector $\vec{u}$. Note that (3) is independent of the noise model. Hence, it is valid whether PMU errors can be modelled by a Gaussian distribution or not [25]. Finding sets of valid attack-angles requires solving non-linear equations derived from (3), as discussed next.

III. Computing undetectable attack-angles

This section presents a closed form expression to compute an undetectable TSA involving $p = 3$ time references. Then it shows how to extend this result for any $p \geq 2$. This contribution represents a great improvement compared to the case of $p = 2$ considered in [18], as it allows a continuum of non-trivial feasible attacks, which is needed to address the constraints required by the clock servo.

A. Computing Attack Angles for $p = 3$

The measurements taken by the three PMUs to be attacked are denoted by $[z_1, z_2, z_3]$, and the corresponding attack-angles by $\alpha_1, \alpha_2, \alpha_3$. It has been shown in [18] that attacks are feasible when the effective rank of W is 1. In this case, (3) can be rewritten as

$$w_1(u_1 - 1) + w_2(u_2 - 1) = -w_3(u_3 - 1), \qquad (4)$$

where $w = [w_1 w_2 w_3]$ is the row of largest norm of the attack-angle matrix W.

In what follows $C = (c, r)$ denotes a circle in the complex plane with center c and radius r. An algebraic approach is used to solve (4) and to provide a closed-form solution. Namely, the equation is interpreted as the intersection of the right-hand side with the left-hand side, which represent in the complex plane a circle $C_3 = (w_3, |w_3|)$, and an annular region defined by an inner circle $C_i = (-(w_1 + w_2), ||w_1| - |w_2||)$ and an

Algorithm 1 Compute-Feasible-Angles(w)

```
Input: $w$ (row of largest norm of $W$)
  $C_3 \leftarrow (w_3, |w_3|)$
  $C_i \leftarrow (-(w_1 + w_2), ||w_1| - |w_2||)$
  $C_o \leftarrow (-(w_1 + w_2), |w_1| + |w_2|)$
  Compute $\mathcal{I}_i = C_3 \cap C_i = \{I_1, I_2\}$.
  Compute $\mathcal{I}_o = C_3 \cap C_o = \{I_1, I_2\}$.
  Compute $\Theta_3$ using Proposition 1
Output: $\Theta_3$
```

outer circle $C_o = (-(w_1 + w_2), |w_1| + |w_2|)$. The following result characterises the set Θ_3 of feasible values for α_3.

Proposition 1. *For $p = 3$ and rank$(W) = 1$, the set U_3 of feasible values of u_3, i.e., $U_3 = \{u_3 : u_3 = e^{i\alpha_3}\ \forall \alpha_3 \in \Theta_3\}$ is either a non-empty connected compact subset of $\mathbb{T}$ or the union of two non-empty connected compact subsets of $\mathbb{T}$. Furthermore, $u_3 = 1 \in U_3$.*

Proof. Let $\mathcal{I}_o$ and $\mathcal{I}_i$ be the set of intersection points of the circle C_3 with the outer and the inner circle, respectively. Four cases are distinguished.

1) $|\mathcal{I}_o| + |\mathcal{I}_i| = 1$, i.e., C_3 is tangent to one of the circles. This intersection point must be the one corresponding to $\alpha_3 = 0$, because $\alpha_1 = \alpha_2 = \alpha_3 = 0$ (no attack) is a solution to (4). Thus $\Theta_3 = \{0\}$.
2) $2 \leq |\mathcal{I}_o| + |\mathcal{I}_i| < 4$, i.e., C_3 intersects with one of the circles at two points and could be tangent to the other circle. Let the two intersection points (not the tangent) correspond to angles α_3^1 and α_3^2. If $\{\alpha_3^1, \alpha_3^2\} \in [0, 2\pi]$ and $\alpha_3^1 < \alpha_3^2$ then we have two intervals, $[\alpha_3^1, \alpha_3^2]$ and $[\alpha_3^2, \alpha_3^1 + 2\pi]$, and the set of feasible values is the one including 0, since $\alpha_3 = 0$ is a feasible solution. Hence, $\Theta_3 = [\alpha_3^2, \alpha_3^1 + 2\pi]$.
3) $(|\mathcal{I}_o| = |\mathcal{I}_i| = 2)$, i.e., four intersection points. Let the corresponding angles in increasing order be $\{\alpha_3^1, \alpha_3^2, \alpha_3^3, \alpha_3^4\}$. Observe that due to the ordering, angles 1 and 2 correspond to intersection points with the same circle. The feasible set consists of the intervals between angles that correspond to intersection points with different circles. Thus, $\Theta_3 = [\alpha_3^2, \alpha_3^3] \cup [\alpha_3^4, \alpha_3^1 + 2\pi]$. Notice that the second interval includes $\alpha_3 = 0$.
4) $|\mathcal{I}_o| + |\mathcal{I}_i| = 0$ or $|\mathcal{I}_o| + |\mathcal{I}_i| = \infty$. Since $\alpha_3 = 0$ is a feasible solution, it is clear that $|\mathcal{I}_o| + |\mathcal{I}_i| = 0$ implies that C_3 is inside the annular region, while $|\mathcal{I}_o| + |\mathcal{I}_i| = \infty$ implies C_3 coincides with one of the circles. Thus, in both cases $\Theta_3 = [0, 2\pi[$

Note that Θ_3 always includes the intersection angles because they correspond to feasible solutions, hence the set of feasible solutions is closed. Furthermore, due to the structure of the circle group $\mathbb{T}$, an interval of feasible angles maps into a connected set. Moreover, in all four cases, $0 \in \Theta_3$. In other words, $1 \in U_3$. □

Algorithm 1 shows the pseudo-code for computing the set Θ_3 of feasible values for the attack angle α_3. Let $C_x \cap C_y$ denote the intersection between two circles C_x and C_y. This can be efficiently computed by the following Lemma.

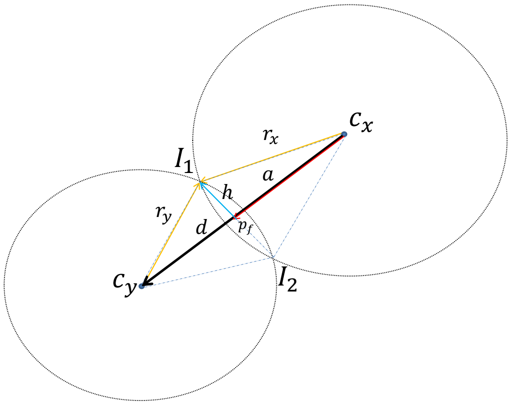

Fig. 1: Example illustrating the intersection between two circles

Lemma 1. *Consider two circles, $C_x = (c_x, r_x)$ and $C_y = (c_y, r_y)$, in the complex plane. Assume that $r_x > r_y$ and that the two circles intersect. Let $\mathcal{I} = \{I_1, I_2\}$ be the set of intersection points. I_1 and I_2 are given by*

$$I_1 = c_x + a + h, \quad I_2 = c_x + a - h$$

where

$$a = d \cdot \frac{d\bar{d} + r_x^2 - r_y^2}{2d\bar{d}}, \quad h = d \cdot i \cdot \sqrt{\frac{r_x^2 - a\bar{a}}{d\bar{d}}}, \quad d = c_y - c_x.$$

Proof. Figure 1 illustrates the problem of finding the intersection of the circles. Let p_f be the point of intersection of the line connecting c_x to c_y and the radical axis of the two circles. Let d be the vector directed from c_x to c_y, that is, $d = c_y - c_x$. Furthermore, let vector a be the vector directed from c_x towards p_f, and h be the vector directed from p_f to I_1. Note that a points in the same direction as d and h is perpendicular to both vectors. By inspecting the two triangles (c_x, p_f, I_1) and (c_y, p_f, I_1) the following two equalities hold

$$|a|^2 + |h|^2 = r_x^2, \quad (|d| - |a|)^2 + |h|^2 = r_y^2$$

solving the two equations for $|a|$ and using $|d|^2 = d\bar{d}$ leads to

$$|a| = \frac{d\bar{d} + r_x^2 - r_y^2}{2|d|}.$$

Since a is parallel to d, it is the case that $a = d \cdot \frac{|a|}{|d|}$, which yields the expression for a in the lemma. Because $|h| = \sqrt{r_x^2 - a\bar{a}}$, and since h is perpendicular to d, it must be that $h = d \cdot i \cdot \frac{|h|}{|d|}$, which in turn yields the expression for h in the lemma. The intersection points can be computed as

$$I_1 = c_x + a + h, \quad I_2 = c_x + a - h.$$

Note that if the two circles intersect only at one point, then $h = 0$, leading to $I_1 = I_2$. □

For each possible attack-angle α_3, valid attack-angles α_1 and α_2 can be efficiently computed as follows. Substituting $s = -w_3(u_3 - 1)$ into (4) leads to

$$w_1(u_1 - 1) = s - w_2(u_2 - 1). \tag{5}$$

Algorithm 2 Compute-Angle-Pairs(α_3, w)

```
Input: $\alpha_3 \in \Theta_3$, and $w$ (row of largest norm of $W$).
  $S \leftarrow \emptyset$
  $s \leftarrow w_3(e^{i\alpha_3} - 1)$
  $C_1 \leftarrow (-w_1, |w_1|)$
  $C_2 \leftarrow (w_2 + s, |w_2|)$
  Compute $\mathcal{I}_{12} = C_1 \cap C_2 = \{I_1, I_2\}$.
  for all $I \in \mathcal{I}_{12}$ do
    Compute $(\alpha_1, \alpha_2)$ by using $I$ in equations (6) , (7)
    $S \leftarrow S \cup (\alpha_1, \alpha_2)$
  end for
Output: $S$
```

Proposition 2. *For each $\alpha_3 \in \Theta_3$ there exist either one or two pairs of (α_1, α_2). A closed form expression of such a pair is given by*

$$u_1 = \frac{I}{w_1} + 1, \quad u_2 = \frac{-w_3(u_3 - 1) - I}{w_2} + 1 \tag{6}$$

$$\alpha_1 = arg(u_1), \quad \alpha_2 = arg(u_2), \tag{7}$$

where I corresponds to an intersection point between the left-hand side of (5) denoted by $C_1 = (-w_1, |w_1|)$ and the right-hand side denoted by $C_2 = (w_2 - w_3(u_3 - 1), |w_2|)$.

Proof. Both the left- and right-hand sides in (5) represent circles in the complex plane. They will be referred to as C_1 and C_2, respectively. C_1 is centered at $c_1 = -w_1$ with radius $r_1 = |w_1|$, and C_2 is centered at $c_2 = w_2 + s$ with radius $r_2 = |w_2|$. An intersection point of these circles corresponds to a solution to (5). The two circles intersect as $\alpha_3 \in \Theta_3$. Again, Lemma 1 can be used to find the set of intersection point(s) I_{12}, and each intersection point corresponds to a pair (α_1, α_2). For each intersection point $I \in I_{12}$ the corresponding (u_1, u_2) and (α_1, α_2) can be computed by equating the left and right hand sides of (5) to I, yielding the expression in the proposition. □

The procedure of computing (α_1, α_2) is illustrated in Algorithm 2. To summarize, the set $\Theta(z) \subset \mathbb{R}^3$ of undetectable attacks for $p = 3$ is a two dimensional manifold in $\mathbb{R}^3$, characterized by one degree of freedom.

B. Computing Attack Angles for any $p \geq 2$

The following describes an algorithm for computing undetectable attacks for the general case of $p \geq 2$. In this case, (4) becomes

$$\sum_{i=1}^{p-1} w_i(u_i - 1) = -w_p(u_p - 1) \tag{8}$$

where w_i is the entry in the row of the largest norm and the i^{th} column of W. In (8), the right hand side represents a circle $C_p = (w_p, |w_p|)$ in the complex plane, while the left hand side represents an annular region that is defined by an inner circle $C_i = (-\sum_{i=1}^{p-1} w_i, \max\{0, 2|w_{i^*}| - \sum_{i=1}^{p-1} |w_i|\})$ and an outer circle $C_o = (-\sum_{i=1}^{p-1} w_i, \sum_{i=1}^{p-1} |w_i|)$, where $i^* = \arg\max_{i \in \{1..p-1\}} |w_i|$. Similar to the procedure of the case when $p = 3$, the feasible set Θ_p of α_p can be computed by Algorithm 1, given the parameters of the circles. For any choice of $\alpha_p^* \in \Theta_p$ (and corresponding u_p^*) equation (8) can be rewritten as

$$\sum_{i=1}^{p-2} w_i(u_i - 1) = -w_{p-1}(u_{p-1} - 1) + s_p$$

where $s_p = -w_p(u_p^* - 1)$. Again, Algorithm 1, with the appropriate parameters of the circles, can be used to compute the feasible range Θ_{p-1} of α_{p-1}. Computing the feasible regions for $p - 2$ iterations results in

$$w_1(u_1 - 1) = \sum_{i=3}^{p} s_i - w_2(u_2 - 1) \tag{9}$$

Notice that (9) has the same form as (5). Therefore, α_1 and α_2 can be computed using Algorithm 2. Hence, it is expected that the set $\Theta(z) \subset \mathbb{R}^p$ of undetectable attacks is a $p - 1$ dimensional manifold in $\mathbb{R}^p$, characterized by $p - 2$ degrees of freedom.

IV. Finding Sets of p Vulnerable Measurements

This section establishes a sufficient condition for finding vulnerable sets of PMUs of size p.

Recall that the attacks target sets of measurements whose corresponding W matrix is of rank approximately or exactly equal to 1. Therefore, the vulnerability of a set is given by a measure of how close to 1 the rank of W is. This quantity is introduced in [18] for $p = 2$ measurements as the index of separation (IoS), the authors also showed that its infimum (IoS*) can be computed from the topology only without access to the measurement values. For a pair of measurements (z_i, z_j), if the value of the IoS of the corresponding W matrix, denoted by $IoS_{(i,j)}(z_i, z_j)$, is close to 1, then W is well approximated by a matrix of rank 1 and the attack can be performed on the pair. If the $IoS^*_{(i,j)} = \min_{z_i, z_j} IoS_{(i,j)}(z_i, z_j)$ is close to 1, then whatever the actual measurement value, the $IoS_{(i,j)}$ will also be close to 1 hence the pair of PMUs is vulnerable. Note that a pair for which the IoS^* is far from 1 might still be vulnerable at a particular time instant due to measurement values for which the IoS gets close to 1. Both can be computed in closed form using Eqs(16-17) in [18].

The next contribution of this paper is to extend the theory by showing that measurements can be grouped in classes such that any combination of measurements within a class produces a W matrix of rank equal to 1.

Theorem 1. *For a given value of the measurement vector z and for two measurements (i, j) we say that $i\mathcal{R}j$ if and only if $i = j$ or $IoS_{i,j}(z_i, z_j) = 1$.*

1) *The relation $\mathcal{R}$ defined in this way is an equivalence relation over the set of all possible measurements.*
2) *For any set $\mathcal{P}$ of $p \geq 2$ measurements, the corresponding W matrix has rank 1 if and only if all measurements in $\mathcal{P}$ are equivalent under $\mathcal{R}$.*

The proof is given in the appendix. The first item infers that the IoS has the transitivity property. In practice, Theorem 1 gives a sufficient condition for finding vulnerable sets of size p. An attacker will look for a set $\mathcal{P}$ of at least p measurements that mutually have $IoS = 1$. Any combination of at least 2 measurements within $\mathcal{P}$ has a W matrix of rank 1. Hence all or a subset of the measurements of this set can be the target of a powerful attack. This method is efficient, compared to a brute-force approach, which includes computing the rank of W for each combination of p PMUs, and thus has exponential complexity, which makes it intractable even for small size grids.

Theorem 2 establishes a similar result that holds for IoS^*. The proof is not given as it is similar to the proof of Theorem 1.

Theorem 2. *For two measurements (i, j) we say that $i\mathcal{R}^*j$ if and only if $i = j$ or $IoS^*_{i,j} = 1$.*

1) *The relation $\mathcal{R}^*$ defined in this way is an equivalence relation over the set of all possible measurements.*
2) *For any set $\mathcal{P}$ of $p \geq 2$ measurements, the corresponding W matrix has rank 1 if all measurements in $\mathcal{P}$ are equivalent under $\mathcal{R}^*$.*

Note that for any value of the measurement vector z, the relation $\mathcal{R}^*$ is a subset of $\mathcal{R}$.

Theorem 2 is thus more restrictive and enables to find fewer attackable measurements than Theorem 1 does. Nonetheless, the relation $\mathcal{R}^*$ can be computed a-priori without knowledge of the actual values of the measurements.

The next result shows that there is a link between the criticality of a set of measurements and its vulnerability to TSAs based on rank-1 approximations. First criticality and independence of measurements are defined.

Definition 1. *A set of measurements is said to be critical if removing this set makes the system non-observable, i.e., the matrix obtained by deleting the corresponding rows in H does not have full rank.*

Note that the matrix H is complex, hence "deleting one row" means removing the real and imaginary parts of one complex PMU measurement simultaneously.

Definition 2. *Two measurements are said to be independent if the corresponding rows of the H matrix are linearly independent over $\mathbb{C}$. When two measurements are not independent, one is a known complex multiple of the other, i.e., they are essentially measuring the same complex quantity.*

Using the two definitions the following result can be formulated.

Theorem 3. *Assume that the number of measurements M and the number of states N satisfy $N \geq 3$ and $M \geq N+2$. Assume that every single measurement is non-critical. Then, a pair of measurements $\{i, j\}$ (with $i \neq j$) is a critical set if and only if i, j are independent and $IoS^*_{i,j} = 1$.*

The proof is given in the Appendix. In [18], large attacks targeted groups of pairs of PMUs such that their IoS^* values were equal to 1. In this paper, Theorem 2 further establishes that such groups are in fact equivalence classes. Theorem 3 gives a novel condition for identifying vulnerable sets by finding critical pairs. This new technique is linked to the analysis of the rank of matrix H while studying the IoS and IoS^* values correspond to analysing the rank of matrix W and values of matrix F respectively. In order to attack $p \geq 2$ PMUs, a target set can thus be found by identifying sets of measurements of cardinality p such that all combinations of two measurements are critical or have a corresponding IoS^* value equal to 1. Incidentally, Theorems 2 and 3 also establish that the criticality of a pair defines an equivalence relation on the set of non-critical measurements, a result of independent interest that can be used in other contexts than TSAs. Also note that vulnerability to TSAs with rank-1 approximations is not strictly equivalent to pairwise criticality: it is possible to find non-critical pairs of measurements that have $IoS \approx 1$ for some values of the measurement vector z (see Section VI).

Algorithm 3 Optimize-Attack-Angles(z, w, L, φ)

```
Input: $z$ (non-attacked measurement), $w$ (row of largest norm of
  $W$), $L$ (number of grid search points), and $\varphi$ (attacker objective
  function)
  $\mathcal{A} \leftarrow \emptyset$
  $\Theta_3 \leftarrow$ Compute-Feasible-Angles($w$)
  $\eta^* \leftarrow \min\{\eta > 0 : |\{0, \eta, 2\eta, \cdots, 2\pi\} \cap \Theta_3| = L\}$
  $A_3 \leftarrow \{0, \eta^*, 2\eta^*, \cdots, 2\pi\} \cap \Theta_3$
  for $\alpha_3 \in A_3$ do
      $S \leftarrow$ Compute-Angle-Pairs($\alpha_3, w$)
      for $(\alpha_1, \alpha_2) \in S$ do
          $\mathcal{A} \leftarrow \mathcal{A} \cup (\alpha_1, \alpha_2, \alpha_3)$
      end for
  end for
  $(\alpha_1^*, \alpha_2^*, \alpha_3^*) \leftarrow \arg\max_{(\alpha_1,\alpha_2,\alpha_3)\in\mathcal{A}} \varphi(z, \alpha_1, \alpha_2, \alpha_3)$
Output: $(\alpha_1^*, \alpha_2^*, \alpha_3^*)$
```

V. Strategies for Implementing Undetectable Attacks

This section discusses how an attacker could use the presented methods to perform a TSA.

A. Computing an Optimal Undetectable Attack

For attacking $p = 3$ measurements, consider that the attacker has an objective function $\varphi(z, \alpha_1, \alpha_2, \alpha_3)$ he wants to maximize, e.g., the difference between the estimated and the actual power-flow on a transmission line. The attacker can observe the measurements taken at time instants $\{t^0, t^1, \cdots, t^k, \cdots\}$ and knows the instantaneous attack-angles $\alpha^k = \{\alpha_i^k, i \in \{1, \cdots, p\}\}$ that (s)he already implemented. Therefore, given an observed measurement $z^{k'}$ taken at time t^k (possibly already attacked), the attacker can compute the non-attacked measurement $z^k = \{z_i^k : z_i^k = z_i^{k'} e^{-j\alpha_i^k}, i \in \{1, \cdots, p\}\}$, and the angles $(\alpha_1^{k*}, \alpha_2^{k*}, \alpha_3^{k*}) = \arg\max_{(\alpha_1,\alpha_2,\alpha_3)\in\Theta(z^k)} \varphi(z^k, \alpha_1, \alpha_2, \alpha_3)$ that would maximize the attack objective. Finding an approximately optimal solution is feasible, even if φ is non-convex, e.g., using a simple grid search over $\Theta_3(z^k)$ as shown in Algorithm 3. If $(\alpha_1^{k*}, \alpha_2^{k*}, \alpha_3^{k*}) \neq (\alpha_1^k, \alpha_2^k, \alpha_3^k)$, the attacker has to adjust the time references of the PMUs. Note that Algorithm 3 can be easily extended for any $p > 3$ by nesting an additional for loop for each additional attacked measurement $\{\alpha_4, \alpha_5, etc., \cdots\}$, and by updating the implementations of the functions *Compute-Feasible-Angles()* and *Compute-Angle-Pairs()* with the appropriate circle definitions mentioned in

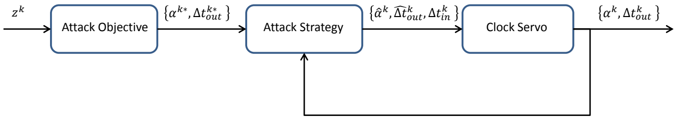
z^k
Attack Objective
$\{\alpha^{k*}, \Delta t_{out}^{k*}\}$
Attack Strategy
$\{\hat{\alpha}^k, \widehat{\Delta t}_{out}^k, \Delta t_{in}^k\}$
Clock Servo
$\{\alpha^k, \Delta t_{out}^k\}$

Fig. 2: Block diagram of clock servo aware timing attacks.

Section III-B. Alternatively a recursive procedure could be used to search through the grid.

Next, methods to implement the optimal angles in an undetectable manner are discussed.

B. Clock Servo and Brute-Force Attack

Regardless of the employed time-synchronization mechanism, PMUs adjust their internal clock smoothly based on the external time reference. The component used to regulate this adjustment is typically called a clock servo. The clock servo can either be a hardware or a software component. Given a sequence of time instants $\{t^0, t^1, \cdots, t^k, \cdots\}$, the clock servo takes as input the observed clock offset Δt_{in}^k at time t^k, and outputs the target offset Δt_{out}^k that will be implemented by time t^{k+1}. In a TSA scenario, recall that the implemented target offset Δt_{out}^k in micro-seconds is related to the implemented attack-angle α^k in radians by the relation $\Delta t_{out}^k = g_\alpha(\alpha^k)$. For now, assume that initially $\Delta t_{out}^0 = 0$ and $\alpha^0 = 0$, i.e., there is no attack.

An attacker that compromises the time-synchronization mechanism can control Δt_{in}^k, and would want to set $\Delta t_{out}^k = \Delta t_{out}^{k*} = g_\alpha(\alpha^{k*})$, where α^{k*} is the attack-angle computed for the PMU that maximizes the attack impact at time t^k computed as discussed above. The intended output offset by the attacker might not always be implemented by the clock servo. Therefore, $\Delta \hat{t_{out}}^k$ and $\hat{\alpha}^k$ are respectively defined as the offset and the attack-angle intended by the attacker at time t^k; and Δt_{out}^k and α^k denote the actual values implemented by the clock servo. Figure 2 summarises the procedure of computing and implementing a TSA.

Brute-Force Attack (BF): A naive attacker that is unaware of the clock servo would provide the servo with $\Delta t_{in}^k = \Delta t_{out}^{k*}$ at every time instant t^k. Nonetheless, due to the clock servo, the implemented Δt_{out}^k could be different from Δt_{in}^k and hence the attack might become detectable, as the adjustments of the clocks of the individual PMUs could result in a trajectory far from the set $\Theta(z^k)$ of undetectable TSAs. This attack corresponds to TSAs considered in previous work [18].

In what follows, it is shown how to implement attacks against two clock servo implementations: output constrained proportional-integral-controller (PI-controller), and output constrained P-controller. Since the latter is a special case of the former, the general case (the PI-controller) is considered first.

C. Output Constrained PI-Controller Clock Servo

The basic equation of a PI-controller is $y[t+1] = y[t] + K_p * e[t] + K_i \int_0^t e[\tau] d\tau$, where K_p and K_i are called the proportional and the integral gains of the controller, respectively. The measurement error at time t is $e[t] = y_{desired} - y[t]$. An output constrained PI-controller (OCPI(ρ)) clock servo adjusts the clock offset depending on the output of a PI-controller, and limits the change in the output during one time step by some threshold ρ. An example of such a clock servo is the one usually used in Precision Time Protocol version 2 (PTPv2). Therefore, the following provides a description of PTPv2 and its widely-used implementation, PTPd [26].

PTPv2 (IEEE1588-2008) is the latest standard protocol for network-based time synchronization [27]. It synchronizes the clock of one or more slave devices to that of a master clock by exchanging timestamps over a network. PTPd and LinuxPTP are widely used open-source implementations of PTPv2 for Unix based systems. The PTPd clock servo is used for adjusting the tick rate of the clock (the number of system clock ticks per second) as follows. First, the PMU clock estimates the master-to-slave d_{m2s} and the slave-to-master d_{s2m} delays from the exchanged timestamps, and uses them to estimate the one way propagation delay d_{prop} between the slave and the master (in μs) according to an infinite impulse response (IIR) low pass filter with the equation

$$s\, d_{prop}^k - (s-1) d_{prop}^{k-1} = (x^k + x^{k-1})/2,$$

where x^k is the filter input at time step k and s is the filter stiffness that controls the cut-off and the phase of the filter. The filter input is computed as $x^k = (d_{m2s} + d_{s2m})/2$. Furthermore, the estimated offset (clock error) from the master $\hat{o}^k$ is computed from d_{prop}^k using a finite impulse response (FIR) low pass filter (two sample average) according to

$$\hat{o}^k = ({\Delta t_{in}}^k + {\Delta t_{in}}^{k-1})/2, \tag{10}$$

where, ${\Delta t_{in}}^k$ is the observed offset computed as ${\Delta t_{in}}^k = d_{prop}^k - d_{m2s}$. Next, $\hat{o}^k$ is fed as the error signal to a discretized PI-controller that is used for computing the tick-rate adjustment:

$$\Delta t_{out}^{k-1} - \Delta t_{out}^k = d_{i,b}^k + K_p\, \hat{o}^k. \tag{11}$$

To interpret (11), observe that $\Delta t_{out}^{k-1} - \Delta t_{out}^k$ represents the tick-rate adjustment that needs to be applied to the clock. The current controller error is $\hat{o}^k$, which is the estimated offset at time step k, and the bounded accumulated controller error (drift) at time step k is

$$d_{i,b}^k = \begin{cases} -\tau_d, & d_i^k < -\tau_d \\ d_i^k, & -\tau_d \le d_i^k \le \tau_d \\ \tau_d, & d_i^k > \tau_d \end{cases}, \tag{12}$$

where

$$d_i^k = d_{i,b}^{k-1} + K_i\, \hat{o}^k, \tag{13}$$

τ_d is a limit on the accumulated error, and $d_i^0 = 0$. In real systems, typical values of the controller gains are $K_p = 0.1$ and $K_i = 0.001$. Furthermore, the servo makes sure that the adjustment magnitude is bounded, i.e., $|\Delta t_{out}^{k-1} - \Delta t_{out}^k| < \tau_d$ similar to (12). Finally, $\Delta t_{out}^{k-1} - \Delta t_{out}^k$ is passed to the Unix kernel function *adjtimex()* to implement the adjustment. Thus, the PTPd clock servo is an OCPI(τ_d) clock servo.

In the case of PTP, a TSA on a PMU can be implemented by changing the propagation times between the PTP master and the PMU (the slave), which causes a change in both the master-to-slave and the slave-to-master delays. In what follows it is shown that the attacker can manipulate Δt_{in}^k such that the attack angle follows a desired sequence, which it can use for performing an undetectable attack.

OCPI(ρ) Clock Servo Aware Attack (OCPI): If the attacker is aware that a PMU uses an OCPI(ρ) clock servo, it provides the clock servo at every second with a computed Δt_{in}^k that results in a desired $\Delta \hat{t}_{out}^{k} = g_\alpha(\hat{\alpha}^k)$, with the constraint $|\Delta t_{out}^{k-1} - \Delta \hat{t}_{out}^{k}| < \varrho$, for some constant $\varrho > 0$. Note that, the notations $\Delta \hat{t}_{out}^{k}$ and $\hat{\alpha}^k$ are used because the values might not be implemented by the servo due to the constraint on d_i^k. The supplied Δt_{in}^k can be calculated by solving

$$\Delta t_{in}^k = \frac{d_{i,b}^{k-1} + \Delta \hat{t}_{out}^{k} - \Delta t_{out}^{k-1}}{\frac{K_i}{2} + \frac{K_p}{2}} - \Delta t_{in}^{k-1}, \qquad (14)$$

which is obtained by substituting (13) and (10) in (11). Note that at time t^k the values of $d_{i,b}^{k-1}$, Δt_{out}^{k-1} and Δt_{in}^{k-1} are already known, hence Δt_{in}^k is only a function of the desired $\Delta \hat{t}_{out}^{k}$.

A simplified version of the OCPI-controller where $K_i = 0$ is also considered. It is referred to as an output constrained P-controller (OCP(ρ)) and the corresponding attack is referred to as the OCP attack. To the best of our knowledge, there are no PTP implementations where the servo is a P-controller. However, considering the OCP attack allows to evaluate the importance of the knowledge of K_i in performing an undetectable attack against an OCPI servo. The results for such a scenario are presented in the next section.

VI. Numerical Results

In this section, simulation results based on the IEEE 39-bus system are provided. Section VI-A describes the electrical model and the methodology considered to evaluate the proposed attacks. Section VI-B considers an attack against $p = 5$ PMUs each measuring a single distinct synchrophasor. For this scenario, it is shown that using the Brute Force strategy as in [18], without taking the servo constraints into consideration, makes the attack detectable by BDD, and that using the attack strategies presented in Section V enables the attack to remain undetected. It is also shown that when attacking $p = 2$ PMUs among this set of 5, the discreteness property of the set of undetectable attack-angles prevents the implementation of servo-aware attacks, thus leading to attack detection. Furthermore, it is shown that implementing robust state-estimation techniques does not counter the presented undetectable attacks. Section VI-C gives numerical evidence that attacks are also possible against PMUs that measure both voltage and current synchrophasors simultaneously, namely when both measurements at a bus share the same time reference. Lastly, Section VI-D shows that pairwise criticality or the shared equivalence class property of attacked measurements are not necessary conditions for vulnerability to TSAs, by demonstrating a practically undetectable attack, using rank-1 approximation of W, on a set of $p = 3$ PMUs where the pairwise IoS^* values of the measurements are strictly less than 1.

A. Electrical Model and Evaluation Methodology

In this simulation, it is assumed that the IEEE 39-bus system network has 13 PMUs that measure voltage phasors, 22 PMUs that measure injected-current phasors, and 12 zero-injection buses as illustrated in Figure 3. It is also assumed that Bus 31 is the connection point to the external grid. Note that a different topology would result in a different verification matrix and thus different attack-locations and attack-angles than the ones presented here. Also note that PMUs are considered to have an OCPI clock servo with threshold $\rho = 500\mu$s and that the OCP and OCPI attack strategies performed by an attacker are done with a threshold of $\rho = 20\mu$s. The used load profiles were obtained from real measurements taken at 50 frames-per-second by real PMUs installed in the 125-kV sub-transmission network of Lausanne, Switzerland. For this reason, the load profiles present time-domain behaviour typical of transmission networks. For a set of p attacked PMUs and for a target transmission line chosen by the attacker, a TSA is simulated using the following procedure:

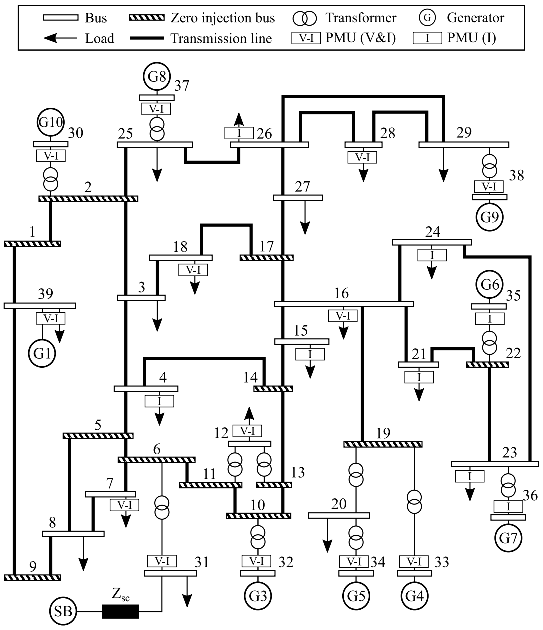

Fig. 3: Benchmark IEEE 39-bus transmission system and PMU locations.

- Creation of the non-attacked measurements:
  - At each time step k, a load flow is computed on the system to determine the true state of the power-grid, based on the load profiles.
  - The true state of the system is perturbed with randomly-generated Gaussian noise, depending on the accuracy of each PMU, to generate the measurement vector z^k (assuming class 0.1 voltage and current sensors).
- Attack computations:
  - Previous measurements are used to compute an estimate $\tilde{z}^k$ of z^k.
  - The estimate $\tilde{z}^k$ is used to compute the optimal intended attack-angles α^{k*} by using grid search as in Algorithm 3.
  - The implemented attack-angle $\hat{\alpha}^k$ is computed according to the chosen attack strategy: BF, OCP or OCPI; and applied to the PMU clock servo, resulting in the attacked measurement vector $z^{k'}$.

- State estimation and attack detection:
  - The WLS estimation is performed, both with the unattacked and attacked measurement vectors z^k and $z^{k'}$, and the measurement residuals $r^k = Fz^k$ and $r^{k'} = Fz^{k'}$ are computed.
  - The LNR and/or χ^2 tests [23] are performed for the residuals r^k and $r^{k'}$.
  - The estimated power-flow is computed on the target line, with and without the attack.

At every new time-step, it is assumed that the attacker can use the previous measurement values to estimate the current measurement value and compute valid attack-angles with respect to this estimate. If the estimate is significantly inaccurate, the solution set of possible attack-angles will in fact be detectable. In order to determine the effect of sudden changes in the system state on the attack detectability, a sudden increase of factor 2 in the active power (referred to as an "inrush") is introduced in one of the buses after $t = 300$ seconds from the start of the simulation. Note that unless the clock-servo implements attack-angles that are different from the ones intended by the attacker, the presented strategies don't impact the residuals, hence detection methods based on the normality of the residuals are expected to fail in identifying the attack.

B. Practical Feasibility of Attacks

For the PMU allocation shown in Figure 3, the analysis presented in section IV was applied to find equivalence classes under $\mathcal{R}^*$ by computing the pairwise IoS^* between measurements. The following equivalence classes were found:

- Class 1 contained 5 measurements: the voltage measurements at buses $\{28, 38\}$, and the current measurements at buses $\{26, 28, 38\}$.
- Class 2 contained 2 measurements: the voltage and current measurements at bus 34.
- Class 3 contained 2 measurements: the voltage and current measurements at bus 37.
- Class 4 contained 6 measurements: the current measurements at buses $\{16, 21, 23, 24, 35, 36\}$.

For the remaining measurements, each measurement constitutes a separate equivalence class.

This PMU allocation contained a total of 34 measurements, and thus the IoS^* of $\binom{34}{2} = 561$ pairs of PMUs had to be checked in order to construct the equivalence classes of attackable measurements. Note that without the knowledge on the equivalence classes, an attacker would have to compute the rank of the W matrix corresponding to $\binom{34}{3} = 5984$ combinations of measurements in order to find whether a $p = 3$ attack exists, or $\binom{34}{5} = 278256$ combinations in order to find whether a $p = 5$ attack exists.

A TSA was mounted on a subset of $p = 5$ PMUs from equivalence class 4, namely the current measuring PMUs at buses $\{21, 23, 24, 35, 36\}$. The goal of the attack was to minimize the apparent power-flow on the line between buses 16 and 24. In this scenario, the simulated inrush was located at bus 21, which is one of the attacked buses. To implement the attack strategies for $p = 5$, the $p = 5$ extension of Algorithm 3 was used to compute the optimal angles α^{k*} at each time-step. First, the impact and detectability of this attack applied according to the different strategies described in Section V are discussed. Then, it is shown that a Robust State Estimation technique is just as vulnerable as a non-robust LSE. Finally, the number of PMUs that an attacker should target within this set of five is investigated.

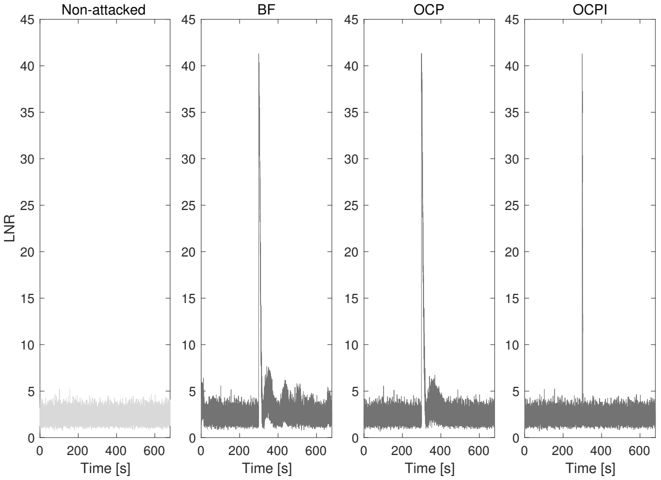

Fig. 4: LNR Test results for attacking the $p = 5$ equivalence class PMUs: the OCPI strategy is closest to the non-attacked scenario, except for the spike caused by the inrush, where all attacks are detectable.

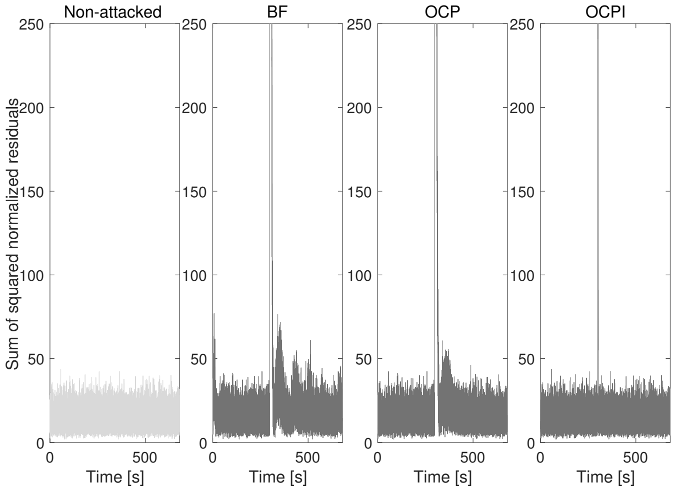

Fig. 5: χ^2Test results for attacking the $p = 5$ equivalence class PMUs: the OCPI strategy is closest to the non-attacked scenario, except for the spike caused by the inrush, where all attacks are detectable. The inrush spike goes above 2000, the figure was zoomed closer to zero for better comparison.

1) Impact and Detectability of the Different Strategies:

The results of applying the LNR and χ^2 tests on the obtained residual vectors are shown in Figures 4 and 5 respectively. The x-axis shows the simulation time, while the y-axis shows the value of the largest normalized residual and of the sum of squared normalized residuals at each time-instant, respectively. The figure shows that the two tests are

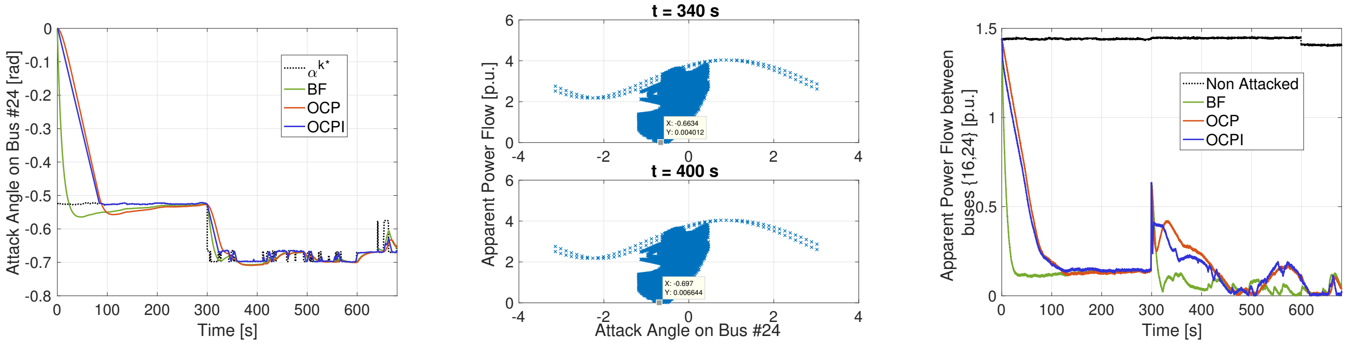

(a) (b) (c)

Fig. 6: Attack strategies and their impact: (a) Attack-angles for the optimal intended case α^{k*} and the ones implemented by the different strategies on the current measuring PMU at bus 24. Notice the spikes in α^{k*} after the inrush. (b) The power-flow on the transmission line between buses 16 and 24 as a function of the attack angle at bus 24 at two different time instants after the inrush. The optimal (minimal) angles changes significantly between the two instants explaining the spikes in the previous figure. (c) The estimated power-flow on the line between buses 16 and 24 is largely affected by the attack strategies.

equivalent in terms of attack detection. The lower the LNR and sum of squared normalized residuals, the stealthier the attack is. It is observed that the OCPI attack strategy yields values that are very close to the non-attacked measurements. The brute-force strategy and the OCP strategy, on the other hand, produce high values for long durations, especially after the inrush. The only time-instants when the OCPI attack produces high residuals are right after the inrush, where a spike is observed. This occurs because with the inrush the attacker mis-estimates the measurement value and does not compute the valid set of feasible attack-angles. As soon as the attacker is able to successfully estimate the new measurements, it is observed that the LNR and the sum of squared normalized residuals values decrease back to undetected values. However this decline is gradual as the clock-servo does not implement the new optimal attack-angles directly. This is the case for all strategies, nevertheless the OCPI strategy allows to regain stable undetectable conditions faster. The figure suggests that sudden changes in the system state present a natural counter-measure to TSAs.

Figure 6a shows the optimal intended attack-angle α^{k*} and the implemented attack-angles for the different strategies applied to the current measuring PMU on bus 24. Note how the BF strategy is the fastest to reach the intended angle, but at the expense of detectability. The other servo-aware strategies reach the intended attack-angle in a more gradual manner thus keeping the LNR values low to a lower normal-looking range. Furthermore, several spikes are observed in the intended optimal attack-angle α^{k*} after the inrush in Figure 6a. These spikes can be explained by Figure 6b, where the x-axis represents the attack angle at Bus 24 and the y-axis represents the apparent power-flow on the line between buses 16 and 24. The sub-figures show all grid points considered by Algorithm 3 (its extension for $p = 5$) before choosing the optimal angles, at two different time-instants after the inrush ($t = 340$s and $t = 400$s). Note that multiple values of the apparent power-flow are obtained for every value of the attack-angle. These values correspond to different choices of the other attack-angles (on buses $\{21, 23, 35, 36\}$). By analysing Figure 6b it can be observed that the minimum power-flow is obtained by choosing substantially different values of the attack-angle ($\alpha^{340*} \approx -0.66$ rad, $\alpha^{400*} \approx -0.7$ rad), which is reflected by the spikes in Figure 6a. Note that the spikes of α^{k*} values impacts the angles implemented by the BF strategy to a greater extent than the angles implemented by the OCP and OCPI strategies. This is due to the fact that the rate of change of the attack-angle for the OCP and OCPI strategies is limited by the threshold parameter ρ.

The impact of the different attack strategies on the target bus is illustrated in Figure 6c. The latter shows that the attacks are able to create a mis-estimation of the power-flow by one order of magnitude and that this mis-estimation is more gradual in servo-aware strategies.

2) *Impact on Robust State-Estimation:*

In what follows, the LSE is assumed to be using the LNR test [23] to be robust against bad-data. In order to simulate a robust LSE, the measurements with the highest normalized residual above threshold η_{BDD} are iteratively removed if removing the measurement does not impact the observability of the system. At the end of the process, if all measurements have normalized residuals below η_{BDD}, then the control center believes it has successfully removed bad-data and thus it will trust the computed state estimate. This robust LSE technique was performed after the attacks on the same $p = 5$ PMUs with $\eta_{BDD} = 5.5$. The choice of the threshold depends on the particular system. It is highly dependent on the number of samples and on the auto-correlation of residuals, even if they are normalized. The value of 3 is mentioned in [23] as an example but, in reality, it is important to set it according to a non-attacked control scenario on a case-by-case basis, so as to avoid excessive false positives (false alarms). Figure 7a shows the apparent power-flow after the use of the bad-data removal scheme. The results are similar to those shown in Figure 6c, which shows that the scheme is not able to remove all of the attacked measurements from the data-set used for state-estimation and thus the impact of the attack is unaffected. This can be explained by the fact that the robust state-estimation techniques rely on the analysis of the distribution of residuals, which is unchanged by an undetectable attack. Figure 7b shows that the control center has removed measurements with

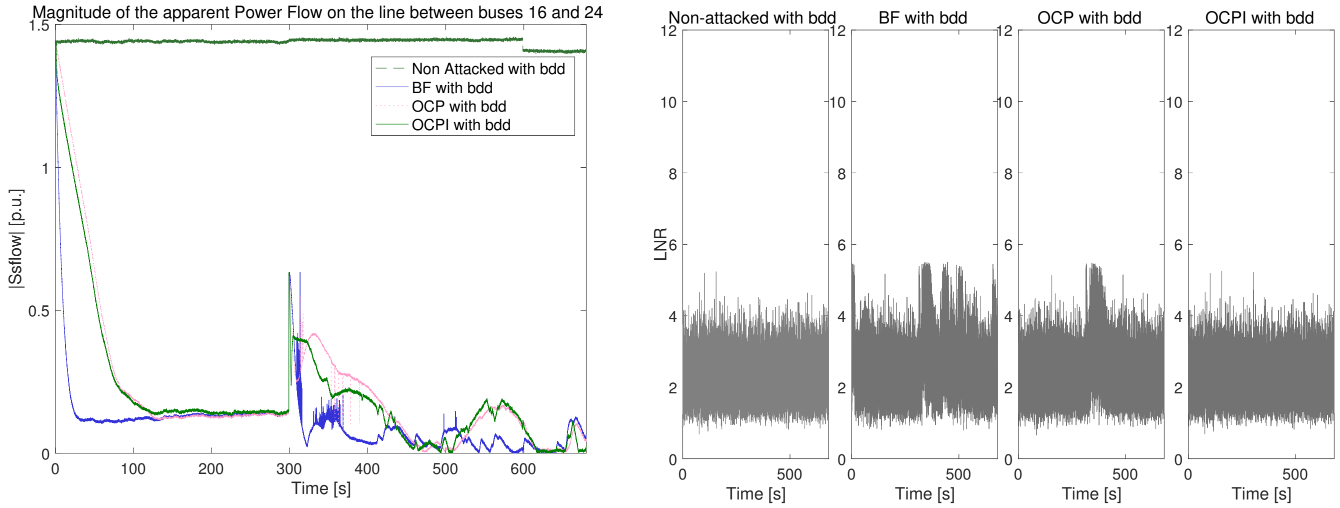

(a)

(b)

Fig. 7: Attack strategies impact with robust LSE: (a) Apparent power-flow obtained without attack and with attacks for $p = 5$ after bad-data removal shows that the mis-estimation is not countered by the robust technique employed. (b) LNR tests show that the LNR values are reduced, so the control center will be enclined to believe it managed to remove all bad-data.

high LNR values but the mis-estimation shown in Figure 7a confirms that not all attacked measurements were removed.

3) Optimal Number of PMUs to Target:

In Section III it was mentioned that when mounting an attack on $p = 2$ PMUs, the solution set of attack-angles that an attacker can choose to use while remaining undetected, is a singleton. Hence the attacker is not able to slowly change the attack angles in accordance with the PMU clock servo. However, by increasing the size of the set of targeted PMUs to $p \geq 3$, the set of undetectable attacks forms a continuum, which allows the attacker to slowly change the attack-angles so as to bypass the BDD. In order to illustrate the impact of choosing more or fewer PMUs to attack, the OCPI attacks were performed on subsets of the 5 PMUs attacked previously. A $p = 2$ case targeting PMUs 21 and 24, and a $p = 3$ case targeting PMUs 21, 23, and 24 are considered. For both cases, the attacker's objective is the same, namely to minimize the apparent power-flow on the line between buses 16 and 24.

Figure 8a shows the obtained apparent power-flow on the targeted line for the unattacked and attacked scenarios for different values of p with the best possible attack strategy. Notice that the apparent power-flow in the case of $p = 2$ drops abruptly at each change of optimal attack-angle, namely at the beginning of the attack and right after the inrush. In fact, in this case, the OCPI attack strategy corresponds to the brute-force strategy as the solution set of attack-angles is finite. Figure 8b shows the LNR test values for the different cases at the beginning of the attack. It can be noticed that the drastic change of apparent power-flow for $p = 2$ causes high, suspicious LNR values. Therefore, it is of the utmost importance that the attacker use a smart servo-aware strategy if she wishes to remain undetected, which requires targeting at least $p = 3$ PMUs. Furthermore, Figure 8a shows that in stable conditions, the apparent power-flow mis-estimation grows with p. However, one must be careful with choosing too high values of p as they also mean larger degrees of freedom which could lead to spikes in the optimal attack-angle and unstable-looking apparent power-flows as observed in Figures 6c and 8a, which could be used as a counter-measure for attack detectability. Also note that because of the increasing number of degrees of freedom, finding the optimal attack-angles takes an increasing amount of time. Therefore, an attacker would need to find a tradeoff between the cost of the attack and the level of mis-estimation she wishes to create. Experimentally, it was observed that $p = 3$ allowed for a large and stable mis-estimation of the power-flow, while remaining undetected.

C. Attacks on Voltage and Current Measuring PMUs

The previous scenario illustrated an attack against a set of $p = 5$ PMUs measuring current phasors only. One question that remains is whether it is possible to attack PMUs that measure both voltage and injected-current phasors simultaneously. Since both measurements are taken by one PMU (using one time reference), the implemented offset, and hence the attack-angles on both phasors will have to be the same, which poses a new constraint on the attacker.

Considering the same PMU allocation as before, another attackable set (pairwise $IoS^* = 1$) of 5 measurements taken by 3 PMUs was found. The attacked measurements are the injected-current phasor at bus 26, the voltage and injected-current phasors at bus 28, and the voltage and injected-current phasors at bus 38. Note that attacking these 5 measurements constitutes a $p = 3$ attack as the attack-angles applied to different measurements at the same bus will have to be the same. Since $p = 3$, there is still a continuum of undetectable attacks which enables the attacker to use the proposed attack strategies for gradually changing the attack-angles. In this scenario the attacker aims to minimize the apparent power-flow on the line between buses 28 and 29. The inrush location was also changed from bus 21 to bus 28 (closer to the attacked buses) to observe its impact on attack detectability.

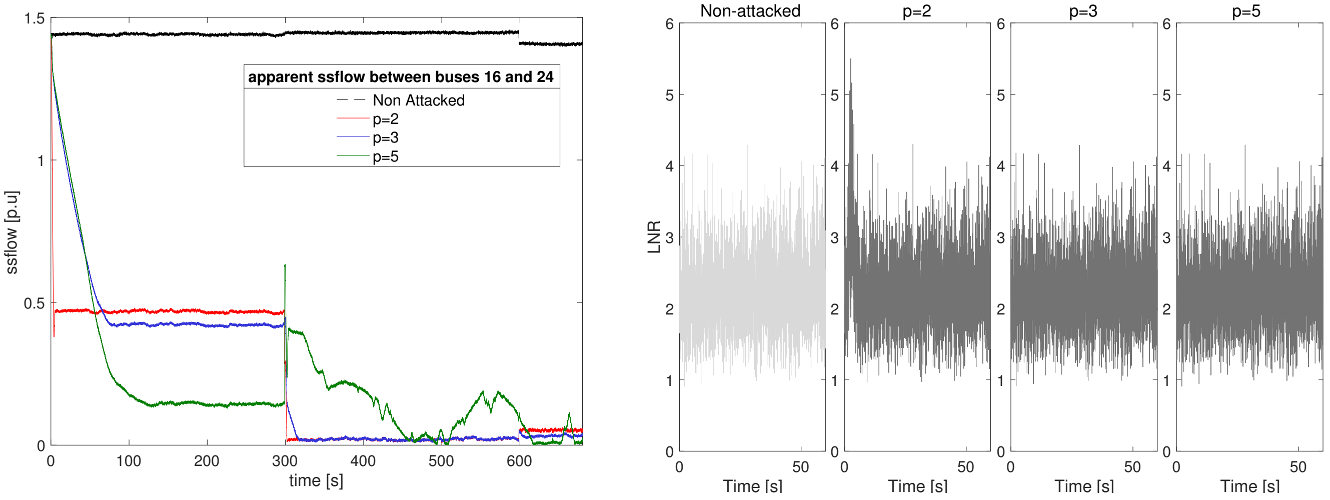

(a)

(b)

Fig. 8: Impact of the chosen number of PMUs to attack: (a) Apparent power-flow obtained without attack and with attacks with $p \in \{2, 3, 5\}$, notice that $p = 2$ is not gradual because the solution set of undetectable attacks is finite. (b) zoom on LNR tests at the beginning of the attacks shows that OCPI attacks for $p \geq 3$ are undetected but for $p = 2$ some residuals are too high which is explained by the
he

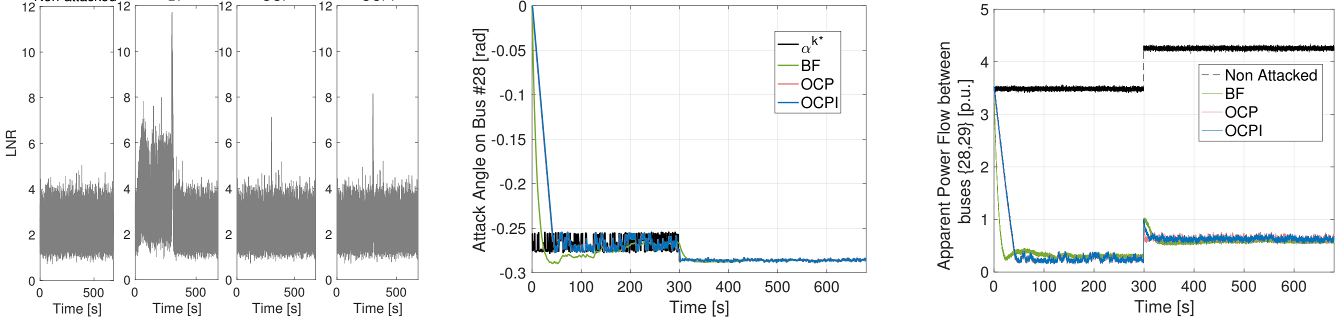

(a)

(b)

(c)

Fig. 9: Attacking PMUs measuring voltages and nodal injected-currents simultaneously: (a) The LNR test results in the non-attacked and OCPI attack scenarios cannot be differentiated, except at the inrush. (b) attack-angles implemented by the OCPI strategy follow the intended optimal angles closely. (c) The power-flow in per-unit on the transmission line between buses 28 and 29 is significantly decreased.

Figure 9a shows similar behaviour as for the previous scenario. Namely, the attack is not detectable with the OCPI strategy, except directly after the inrush when the optimal attack-angles change as shown in Figure 9b. Finally, Figure 9c shows a mis-estimation of the apparent power-flow on the target line by an order of magnitude. Hence, the presented attack strategies can also be generalized to target PMUs measuring two synchrophasors simultaneously.

D. Attacking Non-critical Sets of PMUs

In the previous scenarios, the considered set of PMUs formed an equivalence class, which enabled the computation of undetectable attacks. To illustrate that practical TSAs can be implemented even if there are no critical pairs (when no equivalence classes under IoS^* exist), the feasibility of attacking sets of PMUs such that their pairwise IoS^* values are not equal to 1 is now considered. For this, a denser PMU allocation on the IEEE 39-bus system that does not allow for the existence of equivalence classes under IoS^* is considered. In this allocation, assume that a current measurement is installed on every bus in the system except for buses $\{16, 21, 23, 25\}$ as well as the zero-injection buses which are the same as in the previous setting. Moreover, a voltage measurement is installed on every bus in the system except for buses $\{28, 29\}$ as well as the zero-injection buses. For this scenario, it is assumed that two synchrophasors at the same bus are measured by two distinct PMUs. An inrush at bus 28 is also considered in this scenario. There are now no pairs of PMUs with $IoS^* = 1$, i.e., no critical pairs. However, there are pairs of PMUs for which the IoS^* values are close to one, for example, the current measuring PMUs at buses $\{28, 29, 38\}$; $IoS^*(28, 29) = 0.9996$, $IoS^*(28, 38) = 0.9978$ and $IoS^*(29, 38) = 0.9993$. Hence, an attack against these $p = 3$ PMUs is considered, with the objective of minimizing the estimated power-flow on the transmission line between buses 28 and 29. At all time-steps, the corresponding W matrix has an IoS of maximum value $0.9989 < 1$ and minimum value 0.9975, hence the three measurements never form a critical set. Since the W matrix in this scenario is not of rank equal to 1, a rank-1 approximation is used to compute the attack-angles. The observed mis-estimation of the power-flow is shown in Figure 10a for the different attack strategies. The LNR test

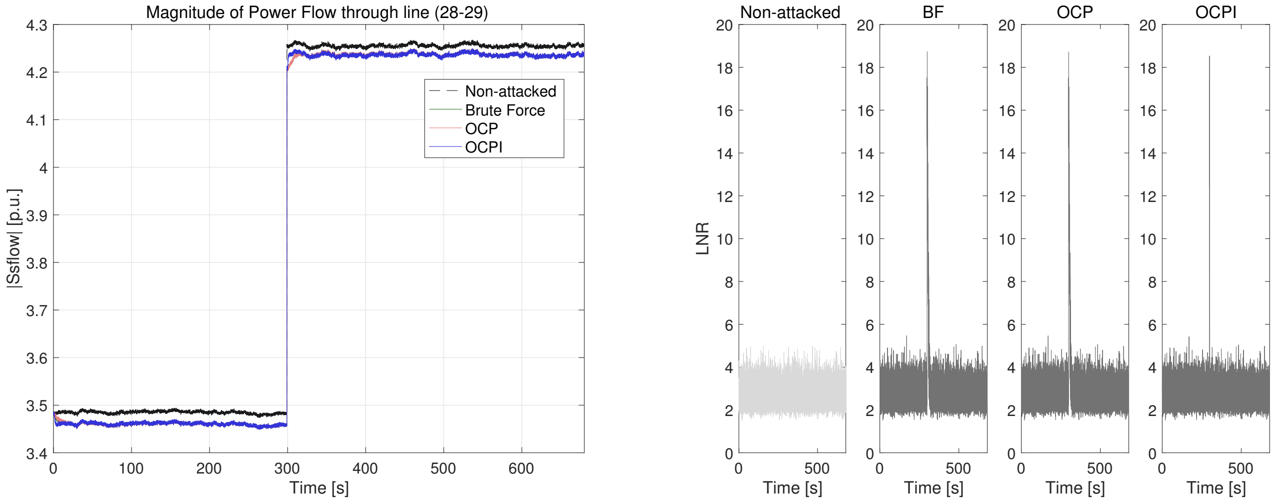

(a) (b)

Fig. 10: Results when pairs of measurements are not critical: (a) Attacks create a mis-estimation of the apparent power-flow. (b) LNR Test for $IoS^* < 1$: all strategies are undetected except for the spike caused by the inrush.

values from Figure 10b show that the attacks are undetectable by the LNR test, for all strategies. Note that the optimal attack-angles in this scenario are very small: they are, respectively, 0.023 and 0.015 rad before and after the inrush. As the optimal intended attack angles are small, they will not be significantly modified by the clock-servo and thus the residuals will not change, which explains the low LNR values for all strategies. Again, the only time instants when the residuals are high are right after the inrush because of the sudden large change of the set of feasible attack-angles.

VII. Possible Countermeasures for TSAs

As discussed before, TSAs do not change the measurement residuals and thus cannot be detected by traditional BDD algorithms that are based on power-system state-estimation. Therefore, finding approaches for mitigating and detecting TSAs is a critical issue for the security of power grids.

To mitigate TSAs, one possible approach is to authenticate time signals (e.g. GPS signals and PTP synchronization messages). For example, message authentication is expected to be a feature in the upcoming PTPv2.1 standard [28]. Although message authentication is an effective countermeasure against spoofing attacks, it is ineffective against delaying the messages using delay-box insertion on transmission lines [29]. In order to mitigate TSAs, it is recommended that the confidentiality of the phasor measurements becomes a requirement in upcoming standards for measurements in power grids, since the knowledge of the measured phasors is required for implementing undetectable TSAs.

To detect TSAs, one approach is to introduce redundancy either in the time synchronization sources or in the measurement technology. For example, the PMU could obtain timing information from different sources simultaneously (e.g., GPS and PTP) in order to verify the synchronization information as suggested in [30]. Although both GPS and PTP are vulnerable on their own, it is highly unlikely that an attacker can manipulate both synchronization systems in an undetectable manner. Similarly, the control center could verify PMU measurements with other SCADA measurements (e.g., from remote terminal units (RTUs)) [31]. The problem with this approach however, is that SCADA measurements are typically available every 4 seconds, while PMUs can provide up to 60 measurements per second. Therefore, the comparison between both measurements would not allow for reliable detection of TSAs. Another detection approach against TSAs is utilizing a denser PMU deployment in the power grid. In Section VI, it was observed that deploying more PMUs in the power system decreases the chances of finding pairs of measurements with $IoS^* = 1$. This solution can become practical, taking into consideration the recent advancements on manufacturing low-cost PMUs [32].

VIII. Conclusion

In this paper, it was shown that vulnerable sets of PMUs of arbitrary size can be found by grouping PMUs in equivalence classes with respect to the IoS. The practical feasibility of attacks was studied and different clock servo-aware strategies for implementing an attack were proposed. Numerical results illustrate the importance of using a smart attack strategy in order for the attack to remain undetected. Using the proposed attack strategy, there is no mismatch between the intended attack and the one implemented by the PMU clock. The experiments also show that attacks can be detected upon the occurrence of an inrush at a nearby bus, yet when using robust state estimation with bad-data removal, the attack was successful and completely undetectable, even during the inrush. Thus the bad-data removal scheme could potentially work in favour of the attacker if the logs of removed measurements are not closely monitored. The effects of sudden changes to the grid could be investigated further to differentiate normal grid dynamics from the ones of an attacked grid. Furthermore, numerical evidence was provided for the feasibility of undetectable attacks when attacking PMUs that measure both voltage and injected current phasors simultaneously, and when attacking sets of PMUs that are not pairwise critical.

Acknowledgement

E. Shereen and G. Dán were partly funded by The Swedish Civil Contingencies Agency (MSB) through the CERCES project.

This work is carried out within the frame of the Swiss Centre for Competence in Energy Research on the Future Swiss Electrical Infrastructure (SCCER-FURIES) with the financial support of the Swiss Innovation Agency (Innosuisse - SCCER program).

Appendix

A. Critical Measurements and W Matrix

Lemma 2. *Assume that $M \geq N+1$ (otherwise every measurement is critical). Consider the W matrix associated with the set of measurements $\mathcal{P} = \{1, ..., p\}$. Let $i \in \mathcal{P}$ and assume that $z_i \neq 0$. The measurement i is critical if and only if $W_{i,i} = 0$.*

Proof. Observe that $W_{i,i} = |z_i|^2 \sum_{j=1}^{M} |F_{j,i}|^2$ and thus $W_{i,i} = 0$ if and only if the first column of the verification matrix is identically 0, which is equivalent to $Fe^{(i)} = 0$ where $e^{(i)}$ is the column vector with 1 in the i^{th} row and 0 else. This means that any attack against the i^{th} measurement alone is undetectable; by a reasoning similar to the proof of Theorem 1 in [15], this is equivalent to measurement i being critical. □

B. Proof of Theorem 1

1. A similar result as Theorem 1 in [20] can be established, with IoS instead of IoS^*, using a similar proof, where Lemma 2 replaces Lemma 2 in [20]. This shows that $\mathcal{R}$ is an equivalence relation.

2. Assume all measurements in $\mathcal{P}$ are in the same equivalence class. The rank of W is 1 by Lemma 2 and Theorem 15 in [33]. Conversely, if the rank of W is 1, then the rank of all principal 2×2 submatrices is ≤ 1 and, again by Lemma 2, is exactly 1, which shows that $IoS_{i,j}(z_i, z_j) = 1$ for all $i, j \in \mathcal{P}$.

C. Proof of Theorem 3

Assume without loss of generality that the measurement pair is $\{1, 2\}$. Let $H^T = (H_1^T, H_2^T)$ where H_1 is a $2 \times N$ complex matrix and H_2 is an $(M-2) \times N$ complex matrix. It follows that $H^\dagger H = H_1^\dagger H_1 + H_2^\dagger H_2$. Since the system is observable, H has full rank; define $A \overset{\text{def}}{=} (H^\dagger H)^{-1}$. The complex verification matrix $F = HAH^\dagger - I_M$ that corresponds to the pair $\{1, 2\}$ can be put in the form

$$\begin{pmatrix} F_1 & F_2 \\ F_3 & F_4 \end{pmatrix} \overset{\text{def}}{=} \begin{pmatrix} H_1AH_1^\dagger - I_2 & H_1AH_2^\dagger \\ H_2AH_1^\dagger & H_2AH_2^\dagger - I_{M-2} \end{pmatrix}$$

Also define $G \overset{\text{def}}{=} F_1^\dagger F_1 + F_3^\dagger F_3$ so that the W matrix is $W = \text{diag}(\bar{z}_{1:2})\, G\, \text{diag}(z_{1:2})$.

1. Now the if part of the theorem can be proven. Assume that $IoS^*_{1,2} = 1$ and i, j are independent. By definition of IoS^*, W and therefore G do not have full rank. Also, since 1 and 2 are independent, H_1 has full rank. The proof proceeds by contradiction: assume that $\{1, 2\}$ is non critical, it then follows that H_2 has full rank, which is a contradiction by Lemma 5.

2. Conversely, assume that $\{1, 2\}$ is a critical pair. The rank of H_2 is $\leq N-1$, thus nullity$(H_2) \geq 1$. By Lemma 4, nullity$(G) \geq 1$ and thus G does not have full rank; thus the same holds for W for any value of z, i.e. $IoS^*_{1,2} = 1$.

Furthermore, assume that measurements 1 and 2 are not independent. Since $\{1, 2\}$ is a critical pair, one of the two measurements is also critical, which is impossible by hypothesis. □

Lemma 3. *For $v \in \mathbb{C}^N$: if $H_2 v = 0$ then $F_1 H_1 v = F_3 H_1 v = 0$.*

Proof. By definition of A, $AH^\dagger H = I_N$ hence $AH_1^\dagger H_1 + AH_2^\dagger H_2 = I_N$, thus

$$\begin{aligned} F_1 H_1 &= H_1 A H_1^\dagger H_1 - H_1 && (15) \\ &= H_1 - H_1 A H_2^\dagger H_2 - H_1 = -H_1 A H_2^\dagger H_2 && (16) \\ F_3 H_1 &= H_2 A H_1^\dagger H_1 = H_2 - H_2 A H_2^\dagger H_2. && (17) \end{aligned}$$

The lemma follows immediately. □

Lemma 4. *If H has full rank then nullity$(H_2) \leq$ nullity(G).*

Proof. By Lemma 3, for all v such that $H_2 v = 0$, $H_1 v$ is in the nullspace of G.

Let $k =$ nullity(H_2) and let $(v_1, ...v_k)$ be k linearly independent vectors in the right nullspace of H_2. Let us show that $(H_1 v_1, ...H_1 v_k)$ are linearly independent. First observe that for $i = 1...k$ we have

$$Hv_i = \begin{pmatrix} H_1 v_i \\ H_2 v_i \end{pmatrix} = \begin{pmatrix} H_1 v_i \\ 0 \end{pmatrix}. \qquad (18)$$

Now assume that for some complex numbers $\lambda_1, ..., \lambda_k$:

$$\lambda_1 H_1 v_1 + ... + \lambda_k H_1 v_k = 0 \qquad (19)$$

it follows that

$$\begin{aligned} \lambda_1 H v_1 + ... + \lambda_k H v_k &= 0 && (20) \\ H(\lambda_1 v_1 + ... + \lambda_k v_k) &= 0 && (21) \end{aligned}$$

Since H has full rank and $N < M$ it follows that the nullity of H is 0. The previous equation thus implies that

$$\lambda_1 v_1 + ... + \lambda_k v_k = 0 \qquad (22)$$

which implies that $\lambda_i = 0$ for $i = 1 : k$. Thus the only null linear combination of $H_1 v_1, ...H_1 v_k$ is the trivial one, which means that $H_1 v_1, ...H_1 v_k$ are linearly independent. Since they are all in the nullspace of G, it follows that nullity$(G) \geq k =$ nullity(H_2). □

Lemma 5. *If H_1 and H_2 both have full rank then F_1, F_3 and G also have full rank.*

Proof. 1. Since H_2 has full rank and $N \leq M-2$, the rank of $H_2 H_2^\dagger$ is N and range$(H_2 H_2^\dagger) = \mathbb{C}^N$. Similarly, A is invertible and thus range$(AH_2 H_2^\dagger) = \mathbb{C}^N$. Also, H_1 has full rank and $N > 2$ thus range$(H_1) = \mathbb{C}^2$. Thus range$(H_1 A H_2^\dagger H_2) = \mathbb{C}^2$.

Now using (16), it follows that range$(F_1 H_1) = \mathbb{C}^2$. Thus

$$\mathbb{C}^2 = \text{range}(F_1 H_1) \subseteq \text{range}(F_1) \subseteq \mathbb{C}^2, \qquad (23)$$

hence range$(F_1) = \mathbb{C}^2$ and F_1 has full rank.

2. Let v be in the nullspace of F_3, i.e. $H_2 A H_1^\dagger v = 0$. Since H_2 has full rank and $N \leq M-2$, the nullspace of H_2 is reduced to 0, thus $AH_1^\dagger v = 0$. Now A is invertible thus $H_1^\dagger v = 0$. Furthermore, since H_1 (hence also $H_1^\dagger$) has full rank and $2 \leq N$, the nullspace of $H_1^\dagger$ is reduced to 0 and finally $v = 0$. Thus the nullspace of F_3 is reduced to 0. Since $2 \leq M-2$, F_3 has full rank.

3. Let v be in the nullspace of G. It follows that $v^\dagger F_1^\dagger F_1 v + v^\dagger F_3^\dagger F_3 v = 0$, i.e. $\|F_1 v\|^2 + \|F_3 v\|^2 = 0$ thus $F_3 v = 0$ and $v = 0$ by item 2. Since G is a square matrix, it follows that G has full rank. □

References

[1] "IEEE standard for synchrophasor measurements for power systems," *IEEE Std C37.118.1-2011 (Revision of IEEE Std C37.118-2005)*, pp. 1–61, Dec 2011.

[2] P. Romano and M. Paolone, "Enhanced interpolated-dft for synchrophasor estimation in fpgas: Theory, implementation, and validation of a pmu prototype," *IEEE Trans. on Instrumentation and Measurement*, 2014.

[3] Y. hua Tang ; Gerard N. Stenbakken ; Allen Goldstein, "Calibration of phasor measurement unit at nist," *IEEE Trans. on Instrumentation and Measurement*, 2013.

[4] G. Barchi, D. Fontanelli, D. Macii, and D. Petri, "On the accuracy of phasor angle measurements in power networks," *IEEE Trans. on Instrumentation and Measurement*, 2015.

[5] A. Carta, N. Locci, C. Muscas, and S. Sulis, "A flexible gps-based system for synchronized phasor measurement in electric distribution networks," *IEEE Trans. on Instrumentation and Measurement*, 2008.

[6] M. Lixia, A. Benigni, A. Flammini, C. Muscas, F. Ponci, and A. Monti, "A software-only ptp synchronization for power system state estimation with pmus," *IEEE Trans. on Instrumentation and Measurement*, 2012.

[7] Y. W. . J. P. Hespanha, "Distributed estimation of power system oscillation modes under attacks on gps clocks," *IEEE Trans. on Instrumentation and Measurement*, 2018.

[8] 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.

[9] S. Barreto, A. Suresh, and J.-Y. Le Boudec, "Cyber-attack on Packet-Based time synchronization protocols: the undetectable delay box," in *2016 IEEE Int. Instrumentation and Measurement Technology Conf. (I2MTC)*, Taipei, Taiwan, May 2016.

[10] R. Zhang, F. Hflinger, and L. Reindl, "Tdoa-based localization using interacting multiple model estimator and ultrasonic transmitter/receiver," *IEEE Trans. on Instrumentation and Measurement*, 2013.

[11] R. G. Lins, S. N. Givigi, and P. R. G. Kurka, "Vision-based measurement for localization of objects in 3-d for robotic applications," *IEEE Trans. on Instrumentation and Measurement*, 2015.

[12] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," in *Proceedings of the 16th ACM Conf. on Computer and Communications Security*, ser. CCS '09. New York, NY, USA: ACM, 2009, pp. 21–32. [Online]. Available: http://doi.acm.org/10.1145/1653662.1653666

[13] G. Dan and H. Sandberg, "Stealth attacks and protection schemes for state estimators in power systems," in *2010 First IEEE Int. Conf. on Smart Grid Communications*, Oct 2010, pp. 214–219.

[14] O. Vukovic, K. C. Sou, G. Dan, and H. Sandberg, "Network-aware mitigation of data integrity attacks on power system state estimation," *IEEE Journal on Selected Areas in Communications*, vol. 30, no. 6, pp. 1108–1118, July 2012.

[15] O. Kosut, L. Jia, R. J. Thomas, and L. Tong, "Malicious data attacks on the smart grid," *IEEE Trans. on Smart Grid*, vol. 2, no. 4, pp. 645–658, Dec 2011.

[16] A. Anwar, A. N. Mahmood, and Z. Tari, "Identification of vulnerable node clusters against false data injection attack in an ami based smart grid," *Inf. Syst.*, vol. 53, no. C, pp. 201–212, Oct. 2015. [Online]. Available: http://dx.doi.org/10.1016/j.is.2014.12.001

[17] Z. Zhang, S. Gong, A. Dimitrovski, and H. Li, "Time synchronization attack in smart grid: Impact and analysis," *IEEE Trans. on Smart Grid*, vol. 4, no. 1, pp. 87–98, March 2013.

[18] S. Barreto, M. Pignati, G. Dan, J.-Y. Le Boudec, and M. Paolone, "Undetectable pmu timing-attack on linear state-estimation by using rank-1 approximation," *IEEE Trans. on Smart Grid*, 2016.

[19] D. Fontanelli, D. Macii, S. Rinaldi, P. Ferrari, and A. Flammini, "A servo-clock model for chains of transparent clocks affected by synchronization period jitter," *IEEE Trans. on Instrumentation and Measurement*, 2014.

[20] S. Barreto, E. Shereen, M. Pignati, G. Dan, J.-Y. Le Boudec, and M. Paolone, "A continuum of undetectable timing-attacks on pmu-based linear state-estimation," in *IEEE Smart Grid Comm*, 2017.

[21] D. Kushner, "The real story of stuxnet," *IEEE Spectrum*, vol. 50, no. 3, pp. 48–53, March 2013.

[22] A. Monticelli, "Electric power system state estimation,," in *Proceedings of the IEEE*, 2000.

[23] A. Abur and A. Exposito, *Power system state estimation: theory and implementation*. CRC, 2004, vol. 24.

[24] J. Grainger and W. Stevenson, *JPower system analysis*. McGraw-Hill New York, 1994, vol. 152.

[25] S. Wang, J. Zhao, Z. Huang, and R. Diao, "Assessing gaussian assumption of pmu measurement error using field data," *IEEE Trans. on Power Delivery*, 2017.

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

[27] "IEEE std1588-2008, IEEE standard for a precision clock synchronization protocol for networked measurement and control systems," *IEEE Instrumentation and Measurement Society*, July 24, 2008.

[28] E. Shereen, F. Bitard, G. Dan, T. Sel, and S. Fries, "Next steps in security for time synchronization: Experiences from implementing IEEE 1588 v2.1," to appear in Int. IEEE Symposium on Precision Clock Synchronization for Measurement, Control, and Communication, Sep 2019.

[29] S. Barreto, A. Suresh, and J. Y. L. Boudec, "Cyber-attack on packet-based time synchronization protocols: The undetectable delay box," in *2016 IEEE Int. Instrumentation and Measurement Technology Conf. Proceedings*, May 2016, pp. 1–6.

[30] A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen, and G. Lachapelle, "Gps vulnerability to spoofing threats and a review of antispoofing techniques," *International Journal of Navigation and Observation*, 2012.

[31] J. Zhang and A. D. Domnguez-Garca, "On the failure of power system automatic generation control due to measurement noise," *IEEE PES General Meeting*, 2014.

[32] A. Angioni, G. Lipari, M. Pau, F. Ponci, and A. Monti, "A low cost pmu to monitor distribution grids," *IEEE Int. Workshop on Applied Measurements for Power Systems (AMPS)*, 2017.

[33] L. E. Dickson, *Modern algebraic theories*, 1930.