SUCCESS
D2.3 v1.0
Security by Design Concept

The research leading to these results has received funding from the European Union's Horizon 2020 Research and Innovation Programme, under Grant Agreement no 700416.

Project Name	SUCCESS
Contractual Delivery Date:	April 30, 2018
Actual Delivery Date:	April 30, 2018
Contributors:	KTH, RWTH
Workpackage:	WP2 – Security, Resilience and Survivability by Design
Security:	PU
Nature:	R
Version:	1.0
Total number of pages:	19

Abstract:

This document describes the SUCCESS Security by design approach, which combines the approaches to security that have been developed over decades in the power systems and in the IT domain. The security by design approach is applied to the example of time synchronization, which is increasingly important in maintaining the stability of power systems in the presence of bidirectional power flows due to the penetration of distributed renewable micro generation.

Keyword list:
Security, Time synchronization, resource management

Disclaimer:
All information provided reflects the status of the SUCCESS project at the time of writing and may be subject to change.

Executive Summary

This second version of the deliverable illustrates the security by design concept of SUCCESS through the example of Phasor Measurement Unit (PMU) time synchronization. The choice of focusing on PMUs is motivated by their increasing importance. PMUs are already widely used in power transmission systems for direct measurement of the system state, with the purpose of improved real time situational awareness compared to legacy SCADA measurements. Furthermore, it is expected that in the future low cost PMUs will be widely used in power distribution systems as well, in order to provide real time situational awareness needed for maintaining grid stability in the presence of a high share of intermittent distributed energy generation.

PMUs rely on precise time synchronization for taking direct measurements of the system state and according to the C37.118 standard the time synchronization accuracy needs to be within 30μs. Such an accuracy can be achieved using space based (GPS, Galileo, etc.) or on network based time synchronization (IEEE 1588-2008, PTPv2). It is widely known that both of these sources of time synchronization are vulnerable to spoofing attacks, which makes PMU measurements and control algorithms based on PMU data vulnerable. The security of time synchronization is thus fundamental for the security of future power systems.

In this deliverable we show how the power system can be made secure against time synchronization attacks through appropriate system design, combining state estimation widely used in power systems for combatting measurement noise with strategic deployment of PMU measurements and IT security controls for time synchronization.

The deliverable gives a background on time synchronization, security issues with existing time synchronization methods, and demonstrates that undetectable time synchronization attacks could be performed by a powerful attacker in absence of appropriate countermeasures implemented in the system. Compared to the first version of the deliverable this version also provides a discussion of the practical aspects of implementing time synchronization attacks, it then describes two mitigation schemes for time synchronization attacks and provides algorithmic results for cost efficient mitigation following the security by design principle.

Authors

Partner	Name	e-mail
KTH Royal Institute of Technology (KTH)		
	György Dán	gyuri@kth.se
	Peiyue Zhao	peiyue@kth.se
	Ezzeldin Shereen	eshereen@kth.se

Table of Contents

1. Introduction

This document reports on the SUCCESS security by design concept and its application to the use case of time synchronization attacks against Phasor Measurement Units (PMUs). Security by design is a methodology that can capture the experience developed in the power systems and in the IT domain over the past decades, and exploits these experiences in a coordinated fashion for developing joint detection and mitigation schemes. Its final objective is a jointly reconfigurable energy and IT system that acts and reacts in a coordinated way to detect and to respond to attacks.

The document gives an overview of the role PMUs are and will be playing in power transmission and distribution systems, and a brief description of the data they measure. It gives a description of time synchronization attacks against PMUs and of their potential impact on power system operation. It then discusses existing approaches for combatting time synchronization attacks and two mitigation strategies that rely on expertise in the IT and in the power domain. Time synchronization attacks are listed as T405 in Deliverable D1.2 [1].

1.1 Security by design

The SUCCESS security by design concept combines traditional detection and mitigation techniques developed in the power systems and in the IT community for developing cost-efficient joint detection and mitigation solutions. At a high level, it leverages the cyber-physical nature of power systems, i.e., the fact that disturbances in the IT system have an effect on the physical domain (electric power), and disturbances in the physical domain can have observable effects in the IT domain.

In this document we show how to use the security by design concept for securing PMU time synchronization for two reasons. First, the problem of time synchronization security is notoriously difficult, and has not been fully understood in the context of PMUs. Second, PMU data is expected to be used for a variety of applications in future power systems, hence its security, including that of time synchronization, is essential. The security by design concept could be demonstrated based on other kinds of data, such as traditional SCADA measurement data, smart meter data, etc., but the security of those data are much better understood and have heavily been investigated in the past [12][13].

Traditionally the security of PMU time synchronization is ensured using IT solutions, including cryptographic solutions, and the detection of anomalies is based on detection algorithms that consider IT information only, e.g., monitoring the sequence of time synchronization messages. At the same time, data validity in power systems is usually verified using state estimation based on a physical model of the system. In the case of PMU data, data validity depends on time synchronization, hence the compromise of PMU time synchronization may be detectable through a state estimator. The rest of the deliverable provides novel results on the limits of power system based detection approach, and as such it shows that a joint detection and mitigation approach is necessary. It then shows how the joint detection approach could be developed to be able to ensure PMU time synchronization security.

1.2 How to read this document

This document provides a description of the security by design concept through the example of securing PMU time synchronization. As such it provides an introduction to PMUs and to time synchronization (Section 2), Security issues in time synchronization (Section 3), the theory of linear state estimation, and attacks that may bypass bad data detection (Section 4). Mitigation schemes are covered in Section 5 of the deliverable.

The deliverable is meant to be expository in nature; the related publications [9][11][16] provide all technical details, but less background information. For a description of the overall SUCCESS architecture we refer to Deliverable D4.3.

This second version of the deliverable differs from the first version in that it discusses practical aspects of implementing time synchronization attacks, it provides two mitigation schemes, and provides algorithmic results on resource efficient mitigation. This is the current version of the deliverable and as such it obsoletes version 1 of the deliverable (D2.2).

2. PMUs and Time Synchronization

PMUs are becoming a critical asset in power transmission systems. They are increasingly used for a variety of wide-area monitoring, protection and control (WAMPAC) applications, including phase angle monitoring and anti-islanding protection [6][4]. Low cost PMUs are also expected to be widely used in distribution networks, in order to provide real time situational awareness needed for maintaining grid stability in the presence of a high share of intermittent distributed energy generation [2]. It is thus of utmost importance to ensure the security of PMUs.

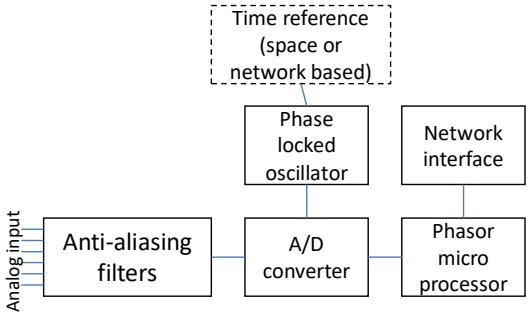

Figure 1 Block diagram of a PMU based on [14]

PMUs make the above applications possible through being able to directly measure the state of the power system, i.e., voltage phasors. A voltage phasor is a complex quantity that represents the magnitude and the phase angle of a sinusoid compared to a reference sinusoid. In order for the voltage and current phasors measured by different PMUs to be comparable, the reference sinusoid has to be the same across all PMUs in a system. Consequently, the PMUs within a synchronous power system (i.e., an interconnected AC power system) require a synchronized time reference. Figure 1 shows the high-level block diagram of a PMU.

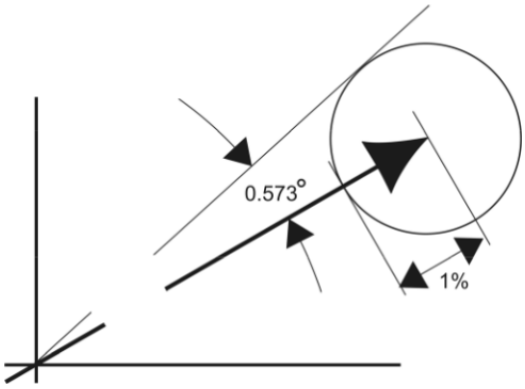

Figure 2 Example of a voltage phasor, which represents the amplitude and the relative phase of a sinusoid compared to a reference sinusoid. The circle illustrates the accuracy requirement of 1% TVE mandated by C37.118 [7]: if the phasor shown represents the true value, the measured value has to be within the circle.

PMUs have stringent accuracy constraints mandated by IEEE C37.118, as illustrated by Figure 2. According to the standard, the total vector error (TVE) of a PMU measurement, which is the ratio of magnitude of the difference between the measured voltage phasor and the actual voltage phasor and of the magnitude of the actual voltage phasor, should not exceed 1%. Hence the measured phasor has to lie within the circle shown in Figure 2. Assuming that the voltage magnitude measurement is accurate, this accuracy requirement corresponds to a phase angle error of 0.573 degree. Thus, in order for a PMU to be compliant with the standard, its time reference has to be synchronized within at most ±31.8μs for 50 Hz grids and ±26.5μs for 60 Hz grids.

Consequently, high accuracy time synchronization is an essential requirement for the wide spread use of PMUs. The time synchronization accuracy required by PMUs can today be achieved through space-based time synchronization (such as GPS, Galileo or Glonass) and through network-based time synchronization with the Precision Time Protocol (PTPv2, IEEE1588-2008).

Space based time synchronization is widely used for civilian applications, such as navigation, and is also used by the vast majority of PMUs installed today. GPS based time synchronization is based on high precision time information transmitted periodically by satellites orbiting the Earth. Precise time information is obtained by correcting the bias due to transmission and propagation delay, and can provide time synchronization accuracy in the order of 40ns.

Network based time synchronization using PTPv2 relies on one or more master clocks connected to a communication network. The master clocks periodically send messages containing their instantaneous time to clients that want to synchronize their clocks. The time information carried in the messages has to be compensated by the clients for the one way propagation delay and for the time that the message has spent in intermediate routers (referred to as transparent clocks in PTPv2). In order to make accurate time synchronization possible, PTPv2 requires that the master clocks, the transparent routers and the clients should be equipped with network interface cards capable of hardware timestamping. Time synchronization also requires an estimate of the one-way propagation time, which is obtained by performing a round trip time measurement and assuming that the propagation times are symmetric, or they have a known asymmetry, which allows computing the one-way delay. PTPv2 with appropriate timestamping hardware is able to provide time synchronization accuracy in the order of 100ns.

3. Time Synchronization Security

Although both space based and network based time synchronization can provide the accuracy required by PMUs, they both have significant vulnerabilities that an attacker could leverage for manipulating the time references of PMUs. On the one hand, GPS has been shown in the past to be vulnerable to spoofing attacks [3]. On the other hand, PTPv2 is vulnerable to the compromise of transparent routers, could be manipulated if time synchronization messages are not authenticated, and could even be attacked by changing the one way propagation delay, e.g., through a delay box [5]. Such attacks against PMU time synchronization, space based or network based, could have a detrimental effect on WAMPAC applications and on grid stability, as shown by recent theoretical and empirical work [8]. Detection and mitigation of these attacks is thus of utmost importance, and has received significant attention recently.

Detection schemes for space-based time synchronization typically rely on equipping receivers with multiple antennas, so that the direction of arrival (DoA) of the satellite signals can be verified. The underlying assumption for such mitigation solutions is that valid satellite signals come from different directions, and the difference in directions can be measured by having two or more receive antennas. On the contrary, a jammer would normally emit all spoofed satellite signals from a single antenna, and hence the direction of arrival of the spoofed satellite signals would be identical. As an example, for GPS [10] proposed a vector tracking based algorithm for detecting sudden changes of GPS time synchronization and location information based on multiple receivers.

In the case of PTPv2, security was originally planned to be provided by protecting the integrity of messages through a hashed message authentication code (HMAC), proposed in Annex K of the standard, but Annex K has never been used due to the computational burden of HMACs. A security extension for PTPv2, called SecurityTLV, is currently under standardization in IEEE. SecurityTLV provides message integrity in one of two ways. One solution adopted by the standard relies on a symmetric group key managed using the GDOI protocol, and another solution relies on a hash chain used for authentication managed using the TESLA protocol. Yet, even if the standard will be extended with SecurityTLV, PTPv2 will remain vulnerable to the compromise of transparent routers and to delay boxes. Providing a system design for power systems that is secure against time synchronization attacks is thus fundamental and is the subject of the following sections.

4. State estimation for mitigation of PMU time synchronization attacks

Beyond technology-specific algorithms (DoA detection, SecurityTLV), in the context of power systems state estimation and related bad data detection (BDD) could be a promising candidate for detecting time synchronization attacks. State estimation is already widely used in power system energy management systems (EMS), in conjunction with traditional SCADA measurements, for combatting measurement noise and for detecting potential bad data due to malfunctioning measurement devices. In what follows we provide a brief overview of linear state estimation (LSE) that can be used in conjunction with voltage and current phasors measured by PMUs.

4.1 System model and linear state estimation

We consider a one-phase direct-sequence equivalent of a three phase transmission network with *N* buses, and M PMUs measuring voltage phasors, current phasors or both. We denote the measurement vector by $z \in C^M$, the system state by $x \in C^M$, and by *H* the *MxN* complex measurement matrix. The measurement model is $z = Hx + e$, where $e \in C^M$ is the complex measurement error. Given a measurement $z \in C^M$, the least squares system state estimate is given by $\hat{x} = (H^* H)^{-1} H^* z$, where H^* is the conjugate transpose of *H*. Based on the state estimate and using the linear measurement model, we can express the estimated measurement vector as $\hat{z} = H\hat{x} + e$. The measurement residual, the difference between the estimated measurement vector and the actual measurement vector, is then defined as $r = \hat{z} - z$. The residual vector *r* can be used for bad data detection using one of many statistical tests. The basis for the statistical tests is the assumption that measurement residuals follow a normal distribution $r \sim N(0, \Omega)$, where

$$\Omega = SR' = (I - K)R' = (I - H'G^{-1}H'^{T}R'^{-1})R' \\ = R' - H'G^{-1}H'^{T}.$$

In the above, $G = (H^* H)^{-1}$, and the normalized residuals can be defined as

$$r_m^N = \frac{r_m}{\sqrt{\Omega_{m,m}}} \sim N(0,1)$$

A commonly used statistical test is the largest normalized residual test, which compares the largest residual to a predefined threshold, typically set to 3 standard deviations, and if it exceeds the threshold, the data is marked as bad data. Figure 3 shows an illustration of state estimation for a system with a one-dimensional state and two measurements.

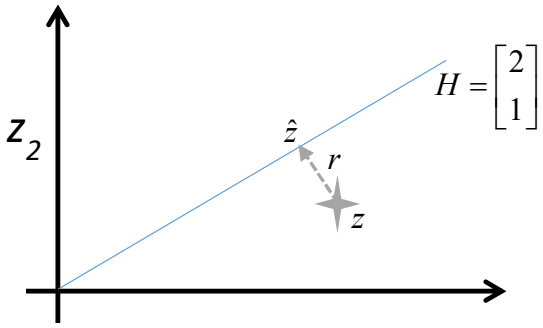

Figure 3 Illustration of linear state estimation for a system with scalar state ($x \in R$) and two measurements $z \in R^2$, with measurement matrix $H=[2\ 1]^T$. Without noise the measurement vector should lie on the blue line. For a measurement vector z, LSE finds a state estimate $\hat{x}$ for which the resulting $\hat{z}$ is closest to z in (weighted) Euclidean distance. The residual vector is the vector between the measurement z and the estimated value $\hat{z}$.

State estimation and bad data detection are thus designed for detecting outliers in measurement data based on a physical model of the power system and based on an assumption about the distribution of measurement noise. From the perspective of time synchronization, the fundamental question is whether LSE combined with BDD would be able to detect time synchronization attacks.

4.2 Attacker model

In order to be able to answer this question we need to define an attacker model. We consider that the attacker is able to manipulate the time reference of *p* PMUs, and for simplicity we assume that it can introduce a different time synchronization bias at each of the *p* PMUs. We denote by α_i the phase angle offset caused by the time synchronization bias at the PMU (*i*=1,…,*p*). As an example, recall that a time synchronization bias of 31μs would result in a phase angle offset of approximately 0.5 degrees in the measured voltage or current phasor. The attack that results in a phase angle offset of α_i at PMU measurement *m* thus corresponds to a rotation $u_i \in C, |u_i| = 1$ of the measured quantity, and the resulting measurement is $z_m u_i$.

4.3 Undetectable attacks against PMU time synchronization

In order to answer the question whether every time synchronization attack can be detected by LSE and BDD let us define for the linear measurement model the complex *MxM* verification matrix

$$F = H(H^* H)^{-1} H^* - I .$$

The key observation is that *Fz=0* for some measurement vector *z* if and only if there exists a state *x* with *z=Hx*. Consider now an attack against PMU time synchronization that results in a change Δz of the measurement vector. Since the change is caused by a manipulated time reference, we can write $\Delta z_i = (u_i - 1) z_i$, where $u_i \in C, |u_i| = 1$ is essentially a rotation. By definition, this attack would be undetectable by the BDD if and only if $F\Delta z = 0$, and the question is whether an attacker would be able to find a set of u_i values such that $F\Delta z = 0$.

In order to model the attack, we can define the attack-measurement matrix Ψ as

$$\Psi_{m,i} = 1 \text{ if } m \in \mathcal{A}_i \text{ and } \Psi_{m,i} = 0 \text{ otherwise,}$$

and we define the complex Hermitian *p*×*p* attack angle matrix

$$W \triangleq \Psi^T \operatorname{diag}(z)^\dagger F^\dagger F \operatorname{diag}(z) \Psi$$

The important observation is that an attack $\alpha=(\alpha_1,\ldots,\alpha_p)$ is undetectable if and only if

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

We refer to [9] for a proof and for more details concerning this important observation.

4.3.1 Attack against a single time reference

If an attacker can only manipulate a single time reference, i.e., *p*=1, then W is a single complex number, and equation (1) has only one solution, $u_i=1$, which corresponds to no attack (excluding the unlikely event that $W_{1,1}=0$). An undetectable attack is thus not possible, and the BDD can be used for detecting an attack.

4.3.2 Attack against two time references

If an attacker can manipulate two time references, i.e., p=2, then W is a 2x2 complex Hermitian matrix and $\alpha=(\alpha_1,\alpha_2)$. The observation in this case is that if W is rank deficient, i.e., *r(W)=1*, then there is a single undetectable attack, and the attack angles can be computed as

$$\alpha_1 = 2\arg(W_{1,1} + W_{1,2})(mod\ 2\pi)$$
$$\alpha_2 = -2\arg(W_{1,2}) + 2\arg(W_{1,1} + W_{1,2})(mod\ 2\pi)$$

We refer to [9] for a proof. The importance of this result is that it shows the possibility of PMU time synchronization attacks that are undetectable by BDD. At the same time it shows that the attack angles are a function of the instantaneous measurement values z_m (since *W* is a function of *z*) and thus the attacker needs to have access to the measurement values and has to compute the attack angles in real time.

On the contrary, if W is full rank, i.e., r(W)=2, then an undetectable attack is not possible, i.e., the BDD would detect any attack.

4.3.3 Attack against three or more time references

If an attacker can manipulate three time references, i.e., p=3, then W is a 3x3 complex Hermitian matrix and $\alpha=(\alpha_1,\alpha_2,\alpha_3)$. The important result in this case is that if W is rank deficient, i.e., *r(W)=1*, then there is a continuum of attacks possible, and each attack corresponds to an intersection point between an annular region and a circle in the complex plane. Similar to the case of p=3, the attack angles can be computed in closed form. We refer to [11] for the computation of the attack angles.

The above result can be generalized to *p>3*. Using the same methodology as used for computing an attack for *p=3*, a recursive algorithm can be used for computing an attack angle vector for *p>3*.

On the contrary, if $r(W)\geq 2$ then an undetectable attack cannot be performed against PMU time synchronization for *p=3*, and the BDD would detect any attack.

4.3.4 Identifying attack locations

It is apparent from the above results that the possibility of performing an undetectable PMU time synchronization attack depends on the rank of the W matrix, which in turn depends on the attack-measurement matrix Ψ (i.e., the locations of the PMUs whose time reference is under attack) and on the measurement vector *z*.

The next important result provides a sufficient condition for identifying pairs of PMUs for which *r(W)=1* independent of the measurement vector *z*. To formulate the sufficient condition let us define the index of separation for matrix *W* as the ratio of its largest eigenvalue and the sum of its eigenvalues, i.e.,

$$\mathrm{IoS} = \frac{\lambda_{\max}}{\sum_i \lambda_i}$$

Clearly, for *p=2* if IoS=1 then the smaller eigenvalue $\lambda_2=0$, and thus *r(W)=1*. As shown in [9], for *p=2* it is possible to provide a lower bound on the IoS, referred to as the minimum index of separation, as

$$IoS^* = \frac{1}{2} + \frac{|f_{12}|}{2\,(f_{11}f_{22})^{\frac{1}{2}}}$$

where

$$f_{i,j} = \sum_{l,m}\sum_{n} \Psi_{l,i}\Psi_{m,j}\bar{F}_{n,l}F_{n,m}$$

Observe that the minimum index of separation does not depend on the measurement vector *z*. Furthermore, $1=\mathrm{IoS}^*\leq\mathrm{IoS}$ implies that *r(W)=1*. Since *IoS** only depends on the attack measurement matrix Ψ and the verification matrix *F*, it can be computed offline based on the measurement model *H*, and can be used for identifying pairs of PMUs for which an undetectable attack exists. We refer to [9] for details.

The above results can be extended to p>2, as shown in [11], as the pairs of undetectable PMUs form equivalence classes. This observation makes it possible to identify sets of PMU

measurements for p>2 for which r(W)=1. Furthermore, for p≥3 the set of undetectable attacks forms a continuum, which makes it possible for the attacker to adapt one of the attack angles in order to maximize the attack impact. As an illustration, Figure 4 shows the apparent power flow as a function of the attack angle α_3 for two different system states. We refer to [11] for details.

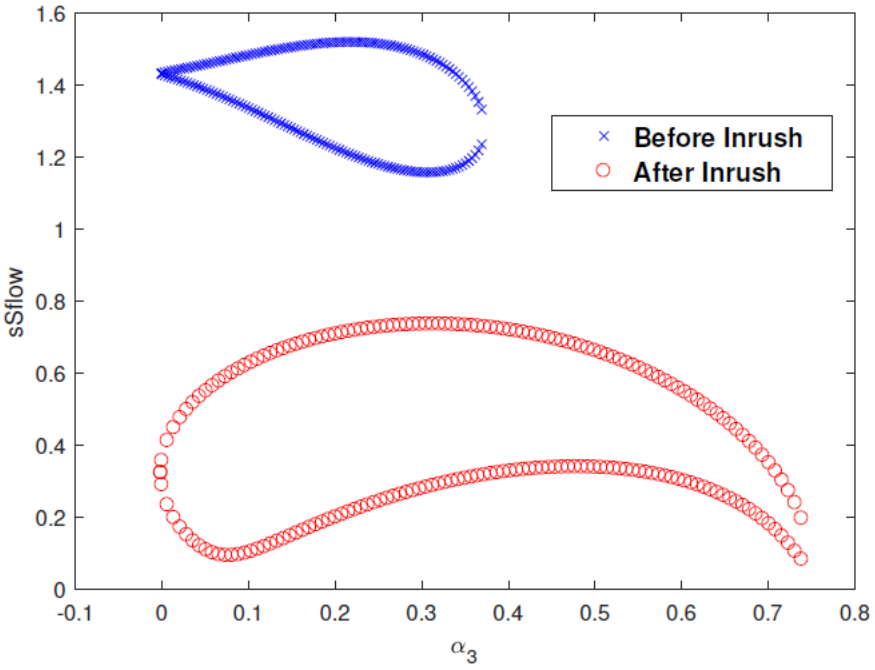

Figure 4 Illustration of attack impact in terms of apparent power flow as a function of the attack angle for p=3. The set of undetectable attacks forms a continuum, and one of the attack angles can be chosen so as to maximize the attack impact.

4.4 Practical considerations

In general, PMUs use a clock servo to adjust their internal clock smoothly based on an external time reference (obtained through GPS or PTPv2). The clock servo can either be a hardware or a software component. In order to be undetectable, an attacker would thus have to take into account the operation of the clock servo when implementing a time synchronization attack.

Let us take, as an example, an attack against p=3 measurements. We denote by $\phi(z, \alpha_1, \alpha_2, \alpha_3)$ the objective function of the attacker, which it aims at maximizing by choosing $\alpha_1, \alpha_2, \alpha_3$. The attacker can observe the measurements taken at time instants (t_0, t_1,..., t_k,...) and knows the instantaneous attack angles $\alpha^k = \left(\alpha_i^k, i \in \{1, \dots, p\}\right)$. Thus, given a measurement z^k taken at time t^k (possibly already attacked), the attacker can compute the non-attacked measurement z^k and the angles $(\alpha_1^{k*}, \alpha_2^{k*}, \alpha_3^{k*}) = argmax_\{(\alpha_1, \alpha_2, \alpha_3)\}\ \phi(z^k, \alpha_1, \alpha_2, \alpha_3)$that would maximize the attack impact given z^k. If $(\alpha_1^{k*}, \alpha_2^{k*}, \alpha_3^{k*})$ <> $(\alpha_1^k, \alpha_2^k, \alpha_3^k)$ then the attacker has to adjust the attack angles by taking into account the clock servo.

An example of a clock servo is shown in Figure 5. The clock servo consists of an infinite impulse response filter (IIR), a finite impulse response filter (FIR) and a proportional integrator (PI) controller. Thus, while adjusting the attack angles, the attacker has to ensure that the changes it makes to the master-to-slave delay will cause the output of the PI controller to adjust the clock in a way that the attack angles remain undetectable. In [16] we showed how an attacker can compute a sequence of inputs to the PTP clock servo that would ensure an attack to remain undetectable despite changes in the attack angles.

As a counterexample, if an attacker changes the attack angle according to a step function from $(\alpha_1^k, \alpha_2^k, \alpha_3^k)$ to $(\alpha_1^{k*}, \alpha_2^{k*}, \alpha_3^{k*})$ would cause the clock servo to generate a sequence of attack angles that is detectable.

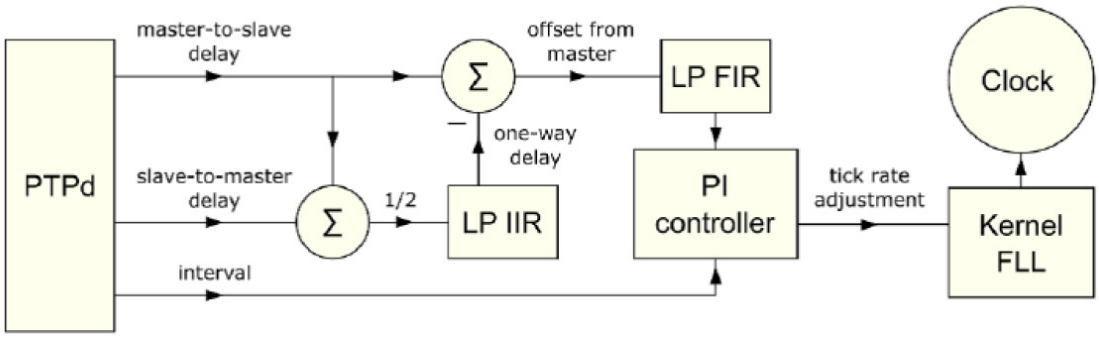

Figure 5. Illustration of the PTP Clock servo from [15]. The attack can manipulate the master-to-slave and the slave-to-master delay so as to influence the tick rate adjustment, which adjusts the clock. The manipulation has to consider the clock servo's operation so as to ensure that the attack remains undetectable while adjusting the attack angles.

4.5 Numerical examples

We validated the attack methodology on the IEEE 39 bus benchmark transmission system, shown in Figure 6 together with the considered PMU locations. A detailed description of the methodology is given in [9].

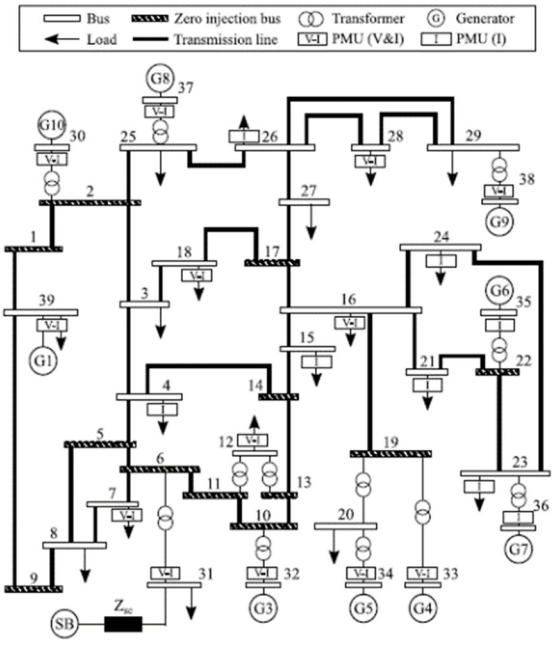

Figure 6 Benchmark IEEE 39 bus system showing PMU locations.

Figure 7 shows how undetectable PMU time synchronization attacks can mislead state estimation resulting in a significant underestimation of the power flow on a transmission line. The figure shows the estimated apparent power flow without attack, when attacking a single pair of PMUs using the undetectable attack angles computed for p=2, when attacking multiple disjoint pairs of PMUs using the undetectable attack angles computed for p=2, and when attacking overlapping pairs of PMUs using the undetectable attack angles computed for p=2 in an order identified using a greedy algorithm proposed in [9].

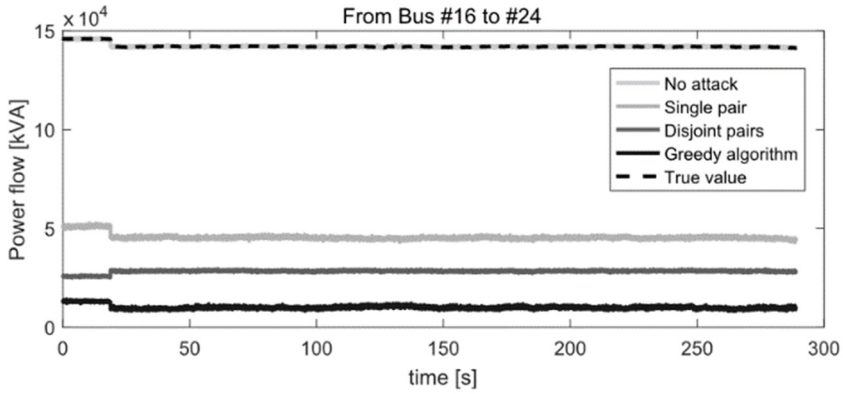

Figure 7 Underestimation of power flow due to time synchronization attack against 2 (single pair) and multiple (Disjoint) PMUs, and due to an attack that maximizes misestimation (Greedy).

To show the importance of considering the clock servo in adjusting the attack angles, Figure 8 shows the largest normalized residual (LNR) as a function of time for 5 scenarios. The results shown in the panel on the left (Non-attacked) show a scenario without a time synchronization attack, and hence the LNR values are due to measurement noise only. The panel BF shows an attack that changes the attack angles step-wise (brute force), and shows that doing so results in increased LNR values. The spike around t=300s is due an inrush, i.e., a large change in the state of the system. The panel OCP shows results corresponding to an attacker that takes into account a clock servo for computing the attack angles, but assumes that the clock servo uses a proportional controller, even though the real clock servo uses a PI controller. The panel OCPI shows results corresponding to an attacker that correctly assumes that the clock servo contains a PI controller. As the results show the LNR values are very low, close to the LNR values without attack. The only main difference is the spike due to the inrush around t=300s. The panel on the right (OCPI with BDD) shows the LNR values assuming that the operator iteratively removes measurement values that have a large LNR value and recomputes the state estimate until there is not measurement with an LNR value above the detection threshold. The results show that the LNR values are indeed close to the scenario without attack.

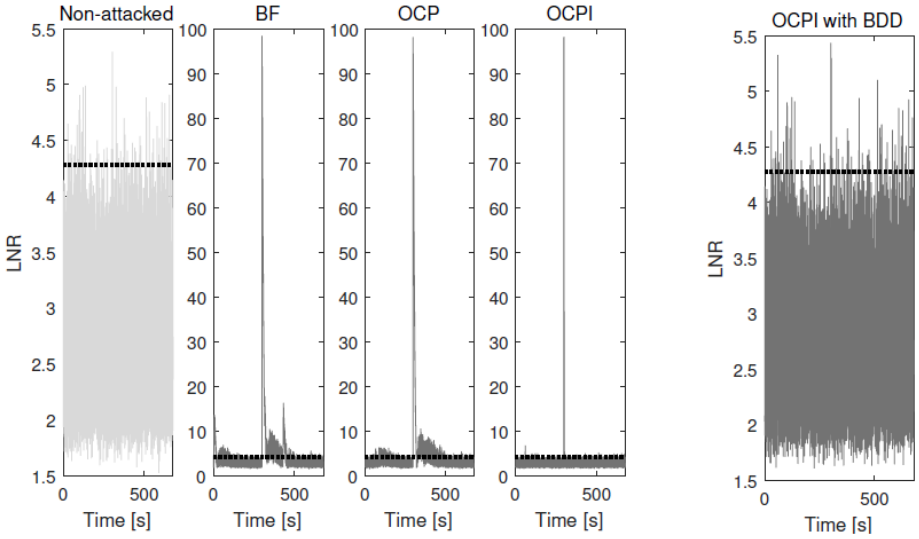

Figure 8 Largest normalized residual as a function of time (left) without attack (BF) with brute force attack, (OCP) assuming a proportional controller in the clock servo, and (OCPI) assuming a PI controller in the clock servo. The panel on the right shows the impact of bad data detection and removal on the remaining largest normalized residuals.

5. Joint Mitigation of Time Synchronization Attacks

In this section we discuss the joint use of IT-based solutions and the LSE based solution for the detection and mitigation of time synchronization attacks against PMU measurements. Throughout the section we denote by $\mathcal{C} = \{C_1, \dots, C_A\}$ the set of time synchronization attacks that can be performed in the power system, as described above. Performing an attack C_a involves attacking the time reference of a set of PMUs, the sets of PMUs that have to be attacked for an attack C_a and an attack $C_{a'}$ may or may not overlap.

We consider two solutions for mitigation. The first solution is based on securing the time reference of individual PMUs. The second solution is based on installing additional PMUs in the system.

5.1 Securing PMUs

The time reference of an individual PMU can be secured in various ways. For instance, more than one time synchronization mechanism can be utilized (e.g. use GPS and PTP simultaneously). Another option is the use of physical layer techniques such as analyzing the angle of arrival of the incoming GPS signal [10].

Let M be the set of measurements appearing in possible attacks. M is thus the set of possible PMUs for the network operator to secure. Securing a PMU $m \in M$ would mitigate any attack that involves attacking m. Let us denote by $\mathcal{C}^m = \{a: m \in C_a\}$ be the set of time synchronization attacks that include PMU $m \in M$. Similarly, let $\mathcal{C}^M = \{a: M \cap C_a \neq \emptyset\}$ be the set of attacks that include any of the PMUs in the set M.

Observe that in order to mitigate all time synchronization attacks, the objective would be to find a set $M^*$of PMUs such that $\mathcal{C}^{M^*} = \mathcal{C}$, i.e., all attacks are mitigated. Alternatively, if there is a budget constraint *k* in terms of the number of PMUs that can be secured, i.e., $|M| \leq k$, then the objective would be to find a set $M^+ = argmin_{M,|M|\leq k}|\mathcal{C}\backslash\mathcal{C}^M|$.

5.2 Deploying additional PMUs

An alternative to securing PMU time synchronization is to deploy additional PMUs in the power system to mitigate PMU timing attacks. Let us denote by I the set of PMUs that is to be deployed, and denote by H^I the measurement matrix after deploying the set I of additional PMUs. The matrix H^I can be obtained by concatenating $|I|$ rows to the original measurement matrix H.

Before addressing mitigation, an important first question is whether deploying additional PMUs could make new time synchronization attacks possible. Our results show that if the system was observable before deploying additional PMUs then deploying additional PMUs does not make new time synchronization attacks possible. We refer to [16] for the proof.

Thus, deploying an additional PMU i would either mitigate one or more attacks $C_a \in \mathcal{C}$, or it does not mitigate any attack. Let us denote by $\mathcal{C}^i = \{a: i\ mitigates\ C_a\}$ be the set of time synchronization attacks that are mitigated by deploying an additional PMU i. Similarly, let $\mathcal{C}^I = \cup_{i\in I}\ \mathcal{C}^i$ be the set of attacks mitigated by deploying the set I of PMUs.

In order to mitigate all time synchronization attacks, the objective would be to find a set $I^*$of additional PMUs such that $\mathcal{C}^{I^*} = \mathcal{C}$, i.e., all attacks are mitigated. Alternatively, if there is a budget constraint *k* in terms of the number of PMUs that can be deployed, i.e., $|I| \leq k$, then the objective would be to find a set $I^+ = argmin_{I,|I|\leq k}|\mathcal{C}\backslash\mathcal{C}^I|$.

5.3 Mitigation algorithms

Although the above two mitigation schemes, securing existing PMUs and deploying additional PMUs, appear to be substantially different, the algorithmic problem of finding optimal sets of PMUs is the same. In what follows we first consider the problem of finding an optimal set of PMUs under a constraint on the number of PMUs to be secured or to be deployed. We then consider the problem of finding the minimum number of PMUs to be secured or to be deployed so as to mitigate all attacks. In order to unify the notation of the previous two subsections, in what follows we will denote by m a mitigation action (deploying an additional PMU, or securing

an existing PMU), by $\mathcal{C}^m$ the set of attacks that the mitigation action would mitigate, and by $\mathcal{M}$ the set of available mitigation actions. Recall that in the case of securing PMUs the set $\mathcal{M}$ corresponds to the set of PMUs installed in the system (or a subset of those, if the time reference of some PMUs cannot be secured), while in the case of deploying new PMUs the set $\mathcal{M}$ corresponds to potential measurement locations where there is no PMU deployed yet.

5.3.1 Mitigation under budget constraint

Our first problem is thus to find a set $M^* \subseteq \mathcal{M}$ of mitigation actions such that $|M^*| \leq k$ and the number $\left|\cup_{m \in M^*} \mathcal{C}^m\right|$ of mitigated attacks is maximized. It can be shown that this problem is equivalent to the maximum set cover problem [16], which is a well-known NP-hard combinatorial optimization problem. Furthermore, it can be shown that a simple greedy algorithm that upon every iteration selects the mitigation action that mitigates most attacks among the attacks not yet mitigated, provides an $\frac{e}{e-1}$ approximation of the optimal solution [17].

5.3.2 Mitigation with minimal set of PMUs

Our second problem is to find a set of mitigation actions $M^* \subseteq \mathcal{M}$ such that $\left|\cup_{m \in M^*} \mathcal{C}^m\right| = \mathcal{C}$, and M^* is minimal. It can be shown that this problem is equivalent to the minimum hitting set problem [16], which is another well-known NP-hard optimization problem. Furthermore, the same greedy algorithm as outlined above, provides a $1 + ln\gamma$ approximation for the problem, where γ is the maximum number of mitigation actions that can mitigate an attack.

We can provide a stronger result for the case time synchronization attacks against sets of PMUs for which r(W)=1, for the case of securing the time synchronization of PMUs [16]. In this case, as it was shown in [11], measurements form equivalence classes, that is, if PMU pairs (i,j) and (j,k) are attackable then PMU pairs (i,k) are attackable as well. If we denote by E the number of equivalence classes then it can be shown that it is necessary and sufficient to choose a set of mitigation actions of cardinality$|\mathcal{M}| - E$.

6. Conclusion

This deliverable documents the SUCCESS security by design concept on the example of time synchronization security for PMUs, a topic of increasing importance for the power systems industry. The security by design concept combines approaches developed in the power systems and in the IT community over the past decades into a joint detection and mitigation scheme. The work undertaken in SUCCESS has made three important methodological contributions. First, it has shown that reliance on state estimation only may be insufficient for detecting time synchronization attacks against PMUs. Second, it has provided a methodology for identifying sets of PMUs that are vulnerable to undetectable time synchronization attacks. Third, it has provided a methodology for mitigating undetectable time synchronization attacks through providing algorithms for optimal security investments in secure time synchronization and for optimal deployment of additional PMUs.

7. References

[1] SUCCESS D1.2 V2.0, “Identification of Existing Threats, V2”, April 2017

[2] Pau, Marco, et al. "Low voltage system state estimation based on smart metering infrastructure”, Applied Measurements for Power Systems (AMPS), 2016 IEEE International Workshop on. IEEE, 2016.

[3] Tippenhauer, Nils Ole, Pöpper, Christina, Rasmussen, Kasper Bonne, Capkun, Srdjan, ” On the requirements for successful GPS spoofing attacks,” in Proc. of ACM CCS, 2011

[4] Teixeira, André, Dán, György, Sandberg, Henrik, Berthier, Robin, Bobba, Rakesh B. and Valdes, Alfonso, ``Security of Smart Distribution Grids: Data Integrity Attacks on Integrated Volt/VAR Control and Countermeasures," *in Proc. of American Control Conference (ACC), Jun. 2014.*

[5] Barreto, Sergio, Suresh, Aswin, Le Boudec, Jean-Yves, “Cyber-attack on Packet-Based Time Synchronization Protocols: the Undetectable Delay Box,” in Proc. of IEEE International Instrumentation and Measurement Technology Conference, 2016

[6] V. Terzija, G. Valverde, D. Cai, P. Regulski, V. Madani, J. Fitch, S. Skok, M. M. Begovic, and A. Phadke, “Wide-area monitoring, protection, and control of future electric power networks,” Proc. of the IEEE, vol. 99, no. 1, pp. 80–93, Jan 2011.

[7] C37.118.1-2011, “IEEE Standard for Synchrophasor Measurements for Power Systems,”, IEEE, 2011

[8] Almas, M.S, Vanfretti L., Singh R. S., Jonsdottir G. “Vulnerability of Synchrophasor-based WAMPAC Applications to Time Synchronization Spoofing”, IEEE Trans. on Smart Grid: DOI: 10.1109/TSG.2017.2665461

[9] Barreto S, Pignati M, Dán G, Le Boudec J-Y, Paolone M, (2017) Undetectable Timing-Attack on Linear State-Estimation by Using Rank-1 Approximation. IEEE Trans Smart Grid: DOI: 10.1109/TSG.2016.2634124

[10] Ng, Y., Gao, G.X , “Advanced Multi-Receiver Position-Information-Aided Vector Tracking for Robust GPS Time Transfer to PMUs”, GNSS 2015

[11] Shereen, E., Barreto, S., Pignati, M., Dán G, Le Boudec J-Y, Paolone M, “Exact Characterization of Undetectable PMU Time Synchronization Attacks against Linear State Estimation”, in Proc. of IEEE SmartGridComm, Oct. 2017,

[12] Vukovic, O., Sou, K-C., Dán, G., Sandberg, H. “Network-aware Mitigation of Data Integrity Attacks on Power System State Estimation," IEEE Journal on Selected Areas in Communications (JSAC), vol. 30, no. 6, July 2012, pp. 1108-1118

[13] Teixeira, A., Dán, G. Sandberg, H., Berthier, R., Bobba, R.B., Valdes, A., ``Security of Smart Distribution Grids: Data Integrity Attacks on Integrated Volt/VAR Control and Countermeasures," *in Proc. of American Control Conference (ACC), Jun. 2014.*

[14] Hart, D.G., Uy, D., Gharpure, V., Novosel, D., Karlsson, D., Kaba, M., ” PMUs – A new approach to power network monitoring,”, ABB Review, vol 1., 2001.

[15] PTPv2 standard, “1588-2008 - IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems”, IEEE, 2008

[16] E. Shereen, G. Dán, “Algorithms for Strategic Mitigation of PMU Time Synchronization Attacks,“ Swedish National Computer Networking Workshop (SNCNW), under submission, 2018

[17] U. Feige, “A threshold of ln n for approximating set cover” , Journal of the ACM, vol. 45, no. 4, Jul 1998

8. List of Abbreviations

BDD	Bad Data Detection
CENELEC	European Committee for Electro technical Standardization
COTS	Commercial off-the-shelf
DEMS	Decentralised energy management system
DER	Distributed Energy Resources
DMS	Distribution Management System
EMS	Energy Management System
ERP	Enterprise Resource Planning
ESB	Electricity Supply Board
ESCO	Energy Service Companies
ESO	European Standardisation Organisations
ETP	European Technology Platform
ETSI	European Telecommunications Standards Institute
HMAC	Hashed Message Authentication Code
HV	High Voltage
ICT	Information and Communication Technology
IEC	International Electro-technical Commission
LV	Low Voltage
LSE	Linear State Estimator
MPLS	Multiprotocol Label Switching
MV	Medium Voltage
NIST	National Institute of Standards and Technology
OPEX	OPerational EXpenditure
PMU	Phasor Measurement Unit
SE	State Estimator
SCADA	Supervisory Control and Data Acquisition
SDH	Synchronous Digital Hierarchy
SDN	Software defined Networks
SoA	State of the Art
SS	Secondary Substation
TVE	Total Vector Error
WAMPAC	Wide Area Monitoring Protection and Control
WP	Work Package