Security of Fully Distributed Power System State Estimation: Detection and Mitigation of Data Integrity Attacks

Ognjen Vuković and György Dán

Abstract—State estimation plays an essential role in the monitoring and supervision of power systems. In today's power systems state estimation is typically done in a centralized or in a hierarchical way, but as power systems will be increasingly interconnected in the future smart grid, distributed state estimation will become an important alternative to centralized and hierarchical solutions. As the future smart grid may rely on distributed state estimation, it is essential to understand the potential vulnerabilities that distributed state estimation may have. In this paper, we show that an attacker that compromises the communication infrastructure of a single control center in an interconnected power system can successfully perform a denial of service attack against state-of-the-art distributed state estimation, and consequently it can blind the system operators of every region. As a solution to mitigate such a denial of service attack, we propose a fully distributed algorithm for attack detection. Furthermore, we propose a fully distributed algorithm that identifies the most likely attack location based on the individual regions' beliefs about the attack location, isolates the identified region, and then reruns the distributed state estimation. We validate the proposed algorithms on the IEEE 118 bus benchmark power system.

Index Terms—Distributed Power System State Estimation, Security, Data Integrity Attacks, False Data Injection, Detection, Mitigation.

I. INTRODUCTION

Power system state estimation (SE) is an essential functionality of modern Energy Managements Systems (EMS), which allows the power system operators to get an accurate estimate of the system's state despite noisy or faulty measurement data collected by the Supervisory Control and Data Acquisition (SCADA) system at substations [1], [2]. The output of the SE, the estimated state and the resulting power flows, is the basis for various important EMS applications, such as contingency analysis used to assess how an outage would affect system stability, and optimal power flow used to compute the optimal generation profile based on some predefined criteria. Hence, an accurate state estimate is crucial both for system safety and for operating efficiency.

The importance of SE has made its security a major concern, and therefore the vulnerability of standalone SEs to so called stealth attacks has been widely studied [3], [4], [5], [6], [7], [8], [9], [10]. Stealth attacks are false data injection attacks

against the measurement data collected by the SCADA system that successfully bypass the model-based bad data detection (BDD) used in the SE [3]. To secure standalone SE, a variety of mitigation schemes were proposed recently against stealth attacks [3], [5], [6], [7], [8].

Power systems are increasingly interconnected and the trend of interconnection is expected to continue in the future smart grid. Interconnected power systems are managed by independent operators; each operator uses SE to estimate the state of the region of the interconnected system that it controls. Examples of interconnected power systems are the Western Interconnect (WECC) in the U.S., and the ENTSO-E in Europe. The safety of an interconnected power system depends on the safety of its constituent regions, as demonstrated by recent cascading failures (e.g., the U.S. North-East blackout in 2003). It is therefore very important that the operators exchange accurate information about their most recent system state in a timely manner. However, the information exchange is very limited in practice due to the sensitivity of the data, and it typically includes only the data needed for a consistent and correct estimate of power flows on the lines connecting two regions. While today the SE in interconnected power systems is mostly done hierarchically, there is an increasing interest for fully distributed SE (DSE) for future smart grids [11], [12], [13], [14], [15], as it eliminates the need for a central authority. DSE is effectively an extension of the basic SE [1], [2], and it can obtain a consistent state estimate for the entire interconnected power system.

Despite its importance, the security of DSE has not received significant attention. In an interconnected power system every region could in principle use an appropriate mitigation scheme to secure its own local SE. Nevertheless, in the case of DSE in an interconnected system, the security of one's local SE may depend on the security of other SEs, and the security of the DSE as a whole may also depend on the security of the data exchange between the regions [16]. In order to design secure and resilient DSEs for future smart grids, it is thus important to understand the potential vulnerabilities of DSE, i.e., whether or not a compromised control center or compromised data exchange between SEs could affect the DSE. If DSE is vulnerable to attacks, it is important to develop mitigation schemes for the vulnerabilities.

In this work we consider false data injection attacks on fully distributed SE. We consider an attacker that compromises a single control center so that it can manipulate the data exchanged between the control center and its neighbors. We

O. Vuković and G. Dán are with the Laboratory for Communication Networks, School of Electrical Engineering, KTH Royal Institute of Technology, Stockholm, Sweden.

E-mail: {vukovic,gyuri}@ee.kth.se

Manuscript received October 23, 2013; revised April 11, 2014.

consider one of the most recent DSE algorithms [15] and show that an attacker can effectively disable the DSE by manipulating the data exchanged by the attacked control center. We propose an algorithm to detect the attack by identifying discrepancies in the temporal evolution of the exchanged data between regions. Furthermore, we propose a distributed algorithm to mitigate the attack. The algorithm identifies the region with the compromised control center by consolidating the beliefs of the individual regions about the origin of the attack, isolates the identified region, and then restarts the DSE.

The structure of the paper is as follows. In Section II we discuss related work. In Section III we outline the DSE algorithm used for our study. In Section IV we describe the attack model and show that the false data injection attacks can disable the DSE. In Section V we propose an algorithm for attack detection, and in Section VI we propose the algorithm for mitigation. Section VII concludes the paper.

II. RELATED WORK

The vulnerability of standalone SE to false data injection attacks was first studied in [3]. There it was shown that the measurement data collected by SCADA can be corrupted so that they do not trigger the BDD system. Such attacks are often called stealth attacks. The observation was made using a linearized model of the SE, but it was shown later on a SCADA/EMS testbed that stealth attacks are also possible under a non-linear model [4]. Since then the security of standalone SE has received much attention [3], [4], [5], [6], [7], [8], [9], [10]. Various schemes were proposed to mitigate stealth attacks, through individual data protection [5], through changes to the BDD algorithm [6], and through the protection of the SCADA infrastructure [7], [8].

The vulnerability of hierarchical multi-area state estimation has been studied in [17], where the authors extended the false data injection attack presented in [3] to the case of a bi-level hierarchical state estimator, and gave some results on how the attack could impede network observability. The security of DSE against false data injection attacks on the exchanged data between neighboring operators was studied in [16] for a simple DSE [11]. It was shown that an attack can disable the DSE, i.e., can prevent it from finding a correct estimate. Furthermore, a detection scheme was proposed to detect an attack along with a simple mitigation scheme. The mitigation scheme suggested that upon detecting an attack, the regions ignore all exchanged data and perform a local SE. However, by using such a mitigation scheme, the power flows on transmission lines connecting any two regions cannot be correctly estimated. Compared to [16], in this paper we consider a state-of-the-art DSE [15], and we propose a mitigation scheme that makes it possible for the DSE to be performed between non attacked regions. Consequently, the power flows on the lines connecting the non attacked regions can be correctly estimated.

Distributed state estimation can be considered a form of consensus. A widely studied model of consensus under attack is the byzantine consensus problem [18], [19], [20], in which a number of processors have to agree on a value even if some

processors may report a false value to influence the consensus. In our work the processors are the regions, but the attack is fundamentally different; its goal is to impede the convergence of the distributed state estimation, and the mitigation scheme we propose not only provides convergence but it also allows to localize the attack.

III. SYSTEM MODEL AND STATE ESTIMATION

We consider an inter-connected power system that consists of several control areas, which we call regions. We denote the set of regions by $\mathcal{R}$, and use $|\mathcal{R}| = R$. A region $r \in \mathcal{R}$ includes a subset of all buses, and a subset of the transmission lines. Regions have no common buses, but there are shared transmission lines, which connect two regions. We refer to the shared transmission lines as *tie lines*, and to the buses connected by these lines as *border buses*.

We consider models of the active power injections at every bus, and active power flows on transmission lines [1], [2]. The active power injection and flow measurements taken in region r are denoted by the vector $z_r \in \mathbb{R}^{M_r}$, where M_r is the number of measurements in region r . The measurements equal to the actual power injections/flows plus independent random measurement noise, $z_r = f_r(x_r) + e$, where x_r is the vector of phase angles used to compute the power flows in region r . The noise e is usually assumed to have a Gaussian distribution of zero mean. We denote by W_r the diagonal measurement covariance matrix.

We refer to the vector of phase angles x_r as the state vector in region r , and we refer to a component of the vector x_r as a *state variable*. The state variables of the vector x_r correspond to the phase angles on buses that belong to region r , and to the phase angles on border buses in other regions that are needed to describe the measurements on the tie lines and to describe power injection measurements at border buses in region r . Consequently, the state variables included in vectors x_r , $\forall r \in \mathcal{R}$ are overlapping. We denote by $x_{r,r'}$ the vector of state variables of region r that correspond to state variables shared between regions r and r' . Observe that all components in the vector $x_{r,r'}$ are also contained in the vector x_r . We say that region r and region r' are neighbors if the vector $x_{r,r'}$ has at least one component, and we denote the set of all neighbors of region r by $\mathcal{N}(r)$ ($|\mathcal{N}(r)| = N(r)$). For convenience, we introduce the vector $x_{r,b}$ for all state variables that region r shares with its neighboring regions $\mathcal{N}(r)$, i.e., the components in the vectors $x_{r,r'}, \forall r' \in \mathcal{N}(r)$ form the vector $x_{r,b}$. The vectors $x_{r',r}$ and $x_{b,r}$ can be defined in a similar way.

A. Distributed State Estimation (DSE)

The state-estimation problem consists of estimating the voltage phase angles x at all buses given the power flow and injection measurement vector [2]. In the case of DSE each control center needs to estimate those phase angles that are related to its measurements, but it has to cooperate with neighboring control centers, typically by exchanging the state variables of the border buses, to ensure that the power flows on the tie lines are correctly estimated. In most of the recently proposed DSE algorithms, e.g., [11], [12], [13], [15], state

variables are exchanged at the beginning or at the end of every iteration, and are used as an input when calculating the next state vector update. For the purpose of our study, we consider a state-of-the-art algorithm proposed in [15], which is highly robust and obtains accurate estimates of the power flows on the tie lines. The algorithm works as follows.

The goal of the DSE is to estimate x_r in every region under the condition that the estimates of shared state variables match between neighboring regions. One (arbitrary) phase angle in the entire interconnected system is selected as the reference angle, and its value is fixed to zero. Each region estimates x_r by minimizing the squares of the weighted deviations of the estimated active power flows and injections (which are functions of x_r) from the measured values (comprehended in z_r). Therefore, the distributed state estimation problem can be formulated as

$$\begin{aligned} \min_{x_r, r \in \mathcal{R}} \sum_{r \in \mathcal{R}} [z_r - f_r(x_r)]^T [W_r^{-1}] [z_r - f_r(x_r)] \\ \text{s.t. } x_{r,r'} = x_{r',r} \quad \forall r \in \mathcal{R} \text{ and } \forall r' \in \mathcal{N}(r), \end{aligned} \quad (1)$$

where $f_r(x)$ is the vector of non-linear functions describing the active power flows and power injections in region r as a function of the state vector x_r .

The constraints in (1) couple the estimation across regions. In order to have a fully distributed algorithm, auxiliary variables can be introduced so that the problem can be solved using the alternating direction method of multipliers (ADMM) [15]. The resulting iterative solution scheme is

$$\begin{aligned} x_r^{(k+1)} &= (H_r^{(k)T} W^{-1} H_r^{(k)} + cD_r)^{-1} (H_r^{(k)T} z_r + cD_r p_r^{(k)}) \\ s_r^{(k+1)} &= U_{x_r} \cdot \sum_{\forall r' \in \mathcal{N}(r)} Y_{r,r'} \cdot x_{r',r}^{(k+1)} \\ p_r^{(k+1)} &= p_r^{(k)} + s_r^{(k+1)} - \frac{1}{2} (Y_{r,b} \cdot Y_{r,b}^T \cdot x_r^{(k)} - s_r^{(k)}), \end{aligned}$$

where $c > 0$ is a predefined constant, the matrix $H_r^{(k)}$ is the Jacobian of vector $f_r(x^{(k)})$, and matrices D_r , U_{x_r} , $Y_{r,r'}$ are defined as follows. D_r is a diagonal matrix whose element $d_{i,i}$ equals the number of regions sharing the i th component (state variable) of the vector x_r . U_{x_r} is a diagonal matrix whose elements are defined as: $u_{i,i}$ equals to the inverse of the number of regions (if greater than 0) sharing the i th component (state variable) of the vector x_r , and zero otherwise. Finally, $Y_{r,r'}$ is a matrix that determines the connection between vector x_r and vector $x_{r,r'}$, and its elements are: $y_{i,j} = 1$ if the i th element (state variable) in x_r corresponds to the j th element (state variable) in $x_{r,r'}$, and $y_{i,j} = 0$ otherwise. Consequently, we have

$$x_{r,r'} = Y_{r,r'}^T \cdot x_r. \quad (2)$$

Similar to (2), we introduce the matrix $Y_{r,b}$, which has a similar structure as $Y_{r,r'}$ so that we have

$$x_{r,b} = Y_{r,b}^T \cdot x_r \quad (3)$$

The matrix $Y_{b,r}$ can be defined in a similar way.

The DSE is said to converge when for some k^* the maximum change of the state variables in every region is smaller than the *convergence threshold* $\epsilon > 0$, i.e., $\forall r \in \mathcal{R}$, $\|x_r^{(k^*+1)} - x_r^{k^*}\|_\infty < \epsilon$, where $\|\cdot\|_\infty$ denotes the

Fig. 1. Interconnected power system with three regions. The attacker corrupts the control center of Region 1, and tampers with the state variables $x_{1,2}^{(k)}$ and $x_{1,3}^{(k)}$ sent from Region 1, and the state variables $x_{2,1}^{(k)}$ and $x_{3,1}^{(k)}$ received by Region 1. The symbol (+) indicates that the components of the attack vector are added to the corresponding components (phase angles) of the vector of exchanged state variables. The attacker cannot tamper with the state variables exchanged between Regions 2 and 3.

maximum norm of a vector. We refer to the number of iterations k^* required for convergence as the *convergence time*.

IV. A DoS ATTACK ON DSE

We consider an attacker whose goal is to perform a Denial-of-Service (DoS) attack against the DSE, i.e., to disable the DSE by preventing it from converging. The attacker compromises the communication infrastructure of a region $r^a \in \mathcal{R}$ used for data exchange between r^a and its neighbors $\mathcal{N}(r^a)$, so it can manipulate the exchanged data used as an input to the DSE. The exchanged data are the state variables defined by the vectors $x_{r,r^a}^{(k)}$, $\forall r \in \mathcal{N}(r^a)$, and the vectors $x_{r^a,r}^{(k)}$, $\forall r \in \mathcal{N}(r^a)$. We describe the attack against the state variables sent from regions $r \in \mathcal{N}(r^a)$ to region r^a (from r^a to r) at the end of iteration k by the *attack vector* $a_{r,r^a}^{(k)}$ ($a_{r^a,r}^{(k)}$). We define the attack vector $a_{r,r^a}^{(k)}$ as the vector of phase angles whose elements correspond to the value that the attacker adds to that phase angle, that is,

$$\tilde{x}_{r,r^a}^{(k)} = x_{r,r^a}^{(k)} + a_{r,r^a}^{(k)}, \quad (4)$$

where $\tilde{x}_{r,r^a}^{(k)}$ is the resulting corrupted vector of state variables. The vector $\tilde{x}_{r,r^a}^{(k)}$ is used as input to the next iteration of DSE in region r^a , instead of the originally exchanged vector $x_{r,r^a}^{(k)}$. The attack vector $a_{r,r^a}^{(k)}$ can be defined in a similar way.

In the rest of this Section, we describe the attack against the state variables sent to region r^a from its neighbors $r \in \mathcal{N}(r^a)$. The attack against the state variables sent from region r^a to its neighbors can be described in a similar way, but we omit it for brevity. For convenience, we introduce the attack vector $a_{b,r^a}^{(k)}$ for the state variables that region r^a receives from all its neighboring regions

$$a_{b,r^a}^{(k)} = [a_{r_{i_1},r^a}^{(k)} \ a_{r_{i_2},r^a}^{(k)} \ \dots]^T \quad \forall r_{i_j} \in \mathcal{N}(r^a), \quad (5)$$

and the corresponding corrupted vector of state variables

$$\hat{x}_{b,ra}^{(k)} = x_{b,ra}^{(k)} + a_{b,ra}^{(k)}, \quad (6)$$

Fig. 1 illustrates an attack on a power system with three regions. Observe that $\hat{x}_{b,ra}^{(k)}$ is the input to iteration $k + 1$ of DSE, and thus, the attack $a_{b,ra}^{(k)}$ leads to a *corrupted* state vector $\hat{x}_{ra}^{(k+1)}$.

We define the *size of the attack* as the Euclidean norm of the attack vector, i.e., $\|a_{b,ra}^{(k)}\|_2$. Intuitively, a smaller attack size implies smaller corruption added to the exchanged values, which could make the detection and the localization of the attack harder; as our results will show later, this is indeed the case. Thus, it would be natural for the attacker to look for the smallest attack vector that prevents the DSE from converging ($k^* = \infty$), or formally

$$\min_{a_{b,ra}^{(k)}, k=1, \dots} \beta \quad \text{s.t. } k^* = \infty \quad \text{and} \quad \beta = \|a_{b,ra}^{(k)}\|_2; \forall k. \quad (7)$$

Since the distributed state estimation problem is non-linear, solving (7) is non-trivial. In the following we propose an approximation of the above objective.

A. First Singular Vector Attack (FSV)

The First Singular Vector (FSV) attack is an approximation of (7) done by maximizing the introduced disturbances for a given attack size. Note that the attack vector $a_{b,ra}^{(k)}$ results in corrupted vectors

$$\begin{aligned} \hat{s}_{ra}^{(k+1)} &= s_{ra}^{(k+1)} + U_{x_r} \cdot Y_{b,ra} \cdot a_{b,ra}^{(k)} \\ \hat{p}_{ra}^{(k+1)} &= p_{ra}^{(k+1)} + U_{x_r} \cdot Y_{b,ra} \cdot a_{b,ra}^{(k)}, \end{aligned} \quad (8)$$

which yield a corrupted state vector

$$\hat{x}_{ra}^{(k+1)} = x_{ra}^{(k+1)} + K \cdot a_{b,ra}^{(k)}, \quad (9)$$

where $K = (H_r^{(k)T} W^{-1} H_r^{(k)} + cD_r)^{-1} \cdot cD_r U_{x_r} Y_{b,ra}$. Note that the addend in (9) is a vector with the same number of elements as the vector $x_{ra}^{(k+1)}$, and we refer to it as the *addend vector*. The Euclidean norm of the addend vector is maximized if the attack vector $a_{b,ra}^{(k)}$ is aligned with the first right singular vector of the matrix K , that is, with the singular vector with highest singular value. The complexity of singular vector decomposition is $O(mn^2)$ [21], low enough for the computation to be done on-line. Nevertheless, the computation of the Jacobian $H_r^{(k)}$ requires knowledge of the current system state $x_{ra}^{(k)}$ for the attacked region ra . Since the entire current system state is not exchanged between the regions, and consequently the attacker does not have access to all entries in $x_{ra}^{(k)}$, we approximate $H_r^{(k)}$ with the Jacobian calculated at the initial state $H_r^{(0)}$. Such an approximation can be easily used by a sophisticated attacker that knows the system model, which is also sufficient to obtain the matrices U_{x_r} and $Y_{b,ra}$.

Observe that in (9) the size of the corrupted vector $\hat{x}_{ra}^{(k+1)}$ depends on the direction of the addend vector, and consequently, on the direction of the first singular vector. Since the attacker does not know the state vector $x_{ra}^{(k)}$, finding the correct

Fig. 2. IEEE 118 bus system divided into six regions. Neighboring regions are connected by a line and the number next to the line represents the number of shared state variables. Note that the reference bus (69) is not a state variable.

direction is not trivial. In order to estimate the direction, the attacker can assume that the estimates of the power flows on a tie line are closer to their actual values when using the most recent exchanged state variables. Then, the attacker applies the attack so that the introduced estimation errors take the estimates in the direction towards the previous iteration estimates.

B. Impact of FSV Attack on DSE

We show the impact of the FSV attack on the IEEE 118 bus power system, divided into six regions as shown in Fig. 2. We consider that the attacker corrupts the control center of one of the regions, and performs the attacks against the state variables sent from and to that region. Bus 69, located in region r_6 , is used as the reference bus, as specified in the IEEE 118 bus power system. Measurements are taken at every power injection and power flow, and the convergence threshold is $\epsilon = 10^{-3}$. The phase angles, thereby the state variables and the attack vector, are in radians.

As a baseline for comparison we use a simple attack, the Uniform Rotation (UR) attack, which adds a constant ϕ to every compromised state variable. The attack vector of the UR attack is thus $a_{b,ra}^{(k)} = \phi \cdot \mathbf{1}$, where $\mathbf{1}$ is the column vector of all ones with the same dimension as the vector $a_{b,ra}^{(k)}$. The size of the attack is $\|a_{b,ra}^{(k)}\|_2 = \phi \cdot \sqrt{|a_{b,ra}^{(k)}|}$, where $|a_{b,ra}^{(k)}|$ denotes the number of elements in the vector $a_{b,ra}^{(k)}$.

Fig. 3 shows the convergence time k^* (when the DSE converges) as a function of the attack size for the FSV attack and for the UR attack considering regions r_3 , r_5 , and r_6 individually as the attacked region. For all considered cases, both the FSV attack and the UR attack can prevent the DSE from converging, i.e., lead to denial of service. The FSV attack is more powerful than the UR attack: FSV requires a much

Fig. 3. Convergence time for cases when the DSE converges as a function of the attack size.

smaller attack size for a successful denial of service attack than UR. One might expect that the DSE is more sensitive when the region containing the reference bus is attacked, since it may be harder for other regions to synchronize with the reference bus. However, the results show that this is not the case: there is no significant difference when the region containing the reference bus is attacked (region r_6), and when some other region is attacked using either the FSV attack or the UR attack.

Observe that in Fig. 3 it does not take a big FSV attack to prevent the DSE from converging. For example, the FSV attack with size $\|a_{b,r^a}^{(k)}\|_2 = 0.07$ prevents the DSE from converging regardless of which region is attacked. This size corresponds to an average value of the attack vector elements of 0.0265 radians (1.51 degrees) if region r_1 is attacked, or 0.019 (1.07 degrees) if region r_6 is attacked.

Although for small attacks the DSE converges, the estimated state and thus the estimated power flows could be erroneous. Fig. 4 shows the maximum of the relative estimation error for the highest 10% of the power flows and injections as a function of convergence time (and thus the attack size). The relative estimation error increases monotonically with the convergence time, and thereby the attack size, and can exceed 15% for some power flows.

Given the potential of the FSV attack and the UR attack to prevent the DSE from converging, a natural question is whether the attacks can be detected and mitigated. In the following, we show that this is possible.

V. DETECTION OF ATTACKS

Let us start by elaborating on the convergence of the DSE. Recall that in order to solve (1) in a fully distributed fashion, the right-hand side of the condition $x_{r,r'} = x_{r',r}$ is replaced with an auxiliary variable for each $r \in \mathcal{R}$ and $\forall r' \in \mathcal{N}(r)$. In iteration k and for regions r and r' , the auxiliary variable equals to the average of the shared state variables between the regions, i.e., $(x_{r,r'}^{(k)} + x_{r',r}^{(k)})/2$ [15]. Consequently, the condition in (1) can be expressed as $x_{r,r'} = (x_{r,r'}^{(k)} + x_{r',r}^{(k)})/2$, or $(x_{r,r'}^{(k)} - x_{r',r}^{(k)})/2 = 0$. The resulting decomposed problem is solved with the ADMM, which guarantees convergence if the following criteria are satisfied (based on [22]).

Fig. 4. Relative estimation error (maximum) for the upper 10% utilized power flows and injections vs. convergence time.

Proposition 1. *If for $\forall r \in \mathcal{R}$ the function $J_r(x_r) = [z_r - f_r(x_r)]^T [W_r^{-1}] [z_r - f_r(x_r)]$ that region r minimizes (the summand in (1)), is closed, proper, and convex, and the augmented Lagrangian*

$$\mathcal{L} = \sum_{\forall r \in \mathcal{R}} J_r(x_r) + y^T \frac{x_{r,r'}^{(k)} - x_{r',r}^{(k)}}{2} + c \left\| \frac{x_{r,r'}^{(k)} - x_{r',r}^{(k)}}{2} \right\|_2^2 \quad (10)$$

(y is Lagrange multiplier) has a saddle point, then the ADMM converges and $\|(x_{r,r'}^{(k)} - x_{r',r}^{(k)})/2\|_2^2 \rightarrow 0$ as $k \rightarrow \infty$ [22, Appendix A, p. 106–110].

Observe that if the conditions in Proposition 1 are satisfied, and therefore the DSE converges without an attack, the disagreement $\|(x_{r,r'}^{(k)} - x_{r',r}^{(k)})/2\|_2^2$ may not decrease monotonically. However, for large k and when the DSE approaches a solution, one may expect that

$$\|(x_{r,r'}^{(k+1)} - x_{r',r}^{(k+1)})/2\|_2^2 < \|(x_{r,r'}^{(k)} - x_{r',r}^{(k)})/2\|_2^2 \quad (11)$$

holds for all state variables exchanged between regions. In what follows we investigate if a normalized version of (11) can be used to detect convergence problems due to an attack.

Definition. *The mean squared disagreement (MSD) between regions r and r' at iteration k is*

$$d_{r,r'}^{(k)} = \frac{\|(x_{r,r'}^{(k)} - x_{r',r}^{(k)})/2\|_2^2}{|x_{r,r'}^{(k)}|}, \quad (12)$$

where $|x_{r,r'}^{(k)}|$ denotes the number of elements in vector $x_{r,r'}^{(k)}$. Observe that by definition $d_{r,r'}^{(k)} = d_{r',r}^{(k)}$.

Fig. 5 shows the evolution of the MSD $d_{r_6,r'}^{(k)}$ between region r_6 and its neighbors $r' \in \mathcal{N}(r_6)$ without an attack: it decreases for all $r' \in \mathcal{N}(r_6)$. Fig. 6 and Fig. 7 show the evolution of the MSDs of regions r_6 and r_5 , which are neighbors of the attacked regions $r^a = r_2 \in \mathcal{N}(r_6)$ and $r^a = r_6 \in \mathcal{N}(r_5)$, for the FSV attack and for the UR attack, respectively. Observe that not all MSDs decrease with the iterations, which is in contrast to Proposition 1. This is the phenomenon we use to detect convergence problems as described in the following.

Proposition 2. *Let $\sup\{\cdot\}$ be the supremum of a set. If the conditions in Proposition 1 are satisfied, but for large k there*

Fig. 5. Evolution of the MSDs $d_{r_6, r'}^{(k)}$ observed in region r_6 for $r' \in \mathcal{N}(r_6)$. No attack

Fig. 6. Evolution of the MSDs $d_{r_6, r'}^{(k)}$ observed in region r_6 for $r' \in \mathcal{N}(r_6)$ in presence of FSV attack in region $r^a = r_2 \in \mathcal{N}(r_6)$ for attack size 0.1.

Fig. 7. Evolution of the MSDs $d_{r_5, r'}^{(k)}$ observed in region r_5 for $r' \in \mathcal{N}(r_5)$ in presence of UR attack in region $r^a = r_6 \in \mathcal{N}(r_5)$ for attack size 0.7.

Fig. 8. Markov chain based on BALs used for the attack localization.

are some r and $r' \in \mathcal{N}(r)$ such that $\sup\{d_{r, r'}^{(k')} : k' > k\} > 0$, $\|x_r^{(k+1)} - x_r^{(k)}\|_\infty > \epsilon$, and $\exists t \in \mathbb{N}$ so that

$$\sup\{d_{r, r'}^{(k')} : k' > k\} > \sup\{d_{r, r'}^{(k')} : k' > k + t\} \quad (13)$$

then there is a convergence problem (an attack).

Proof. The proof follows from Proposition 1. If the conditions of Proposition 1 hold, then $\|(x_{r, r'}^{(k)} - x_{r', r}^{(k)})/2\|_2^2 \rightarrow 0$ and $d_{r, r'}^{(k)} \rightarrow 0$ as $k \rightarrow \infty$. Consequently, $\sup\{d_{r, r'}^{(k')} : k' > k\} \rightarrow 0$. $\square$

The regions can thus use Proposition 2 to detect an attack.

VI. MITIGATION OF ATTACKS

Given that we can detect an ongoing attack, the next important question is whether it is possible to mitigate the attack. In the following we propose a mitigation algorithm that first aims at localizing the region where a detected attack originates from, and then isolates the region so that the DSE can converge.

A. Distributed localization and mitigation algorithm

We start with the definition of the beliefs of the individual regions, which is the basis for the localization algorithm.

Definition. Let $\tilde{d}_{r, r'}^{(k)} = \alpha^{(k)} d_{r, r'}^{(k)} + (1 - \alpha^{(k)}) \tilde{d}_{r, r'}^{(k-1)}$ be the weighted moving average (WMA) of the MSD $d_{r, r'}^{(k)}$. The smoothing factor $\alpha^{(k)} \in (0, 1)$ and satisfies $\sum_{k=0}^{\infty} \alpha^{(k)} = \infty$. The belief of attack direction of region r that its neighbor $r' \in \mathcal{N}(r)$ is the attacked region at iteration k is defined as

$$B_{r, r'}^{(k)} = \frac{\tilde{d}_{r, r'}^{(k)}}{\sum_{r' \in \mathcal{N}(r)} \tilde{d}_{r, r'}^{(k)}}. \quad (14)$$

Observe that regions have beliefs only about their neighbors. i.e., $B_{r, r'}^{(k)} = 0$, $\forall r' \notin \mathcal{N}(r)$. Furthermore, the beliefs are not necessarily symmetric, i.e., $B_{r, r'}^{(k)} \neq B_{r', r}^{(k)}$ is possible.

Given the beliefs $B_{r, r'}^{(k)}$ of the regions, our goal is to find the region that is most likely to be compromised consistent with all beliefs. Before we introduce the distributed localization algorithm we describe a hypothetical localization scheme based on a global observer, which motivates the proposed algorithm.

Motivation: Assume there exists a token that the regions use to express their beliefs about the attack location: when region r receives the token, it will pass the token to region r' with probability $B_{r, r'}$. Moreover, assume that there exists a global observer that observes every passing of the token and that keeps count of how many times the token visits each region. The observer uses the counts to calculate for every region the empirical frequency of token visits: the number of

token visits to the region divided by the number of token visits to all regions. It then identifies the region $\hat{r}^a$ with the highest empirical frequency as the most likely compromised region.

This hypothetical token passing scheme defines a random walk on a graph: the vertices are the regions and there is an edge between vertices r and r' if $B_{r,r'} > 0$. The random walk can then be modeled by a Markov chain. Figure 8 shows the Markov chain for the interconnected system in Figure 2. The state transition matrix $B^{(k)}$ of the Markov chain is the right stochastic $R \times R$ matrix in which every row and every column corresponds to a region, and the entries of the matrix are the beliefs of attack direction $B_{r,r'}^{(k)} \forall r' \in \mathcal{R}$. Thus, row r contains the beliefs of region r . Under appropriate conditions, which we will discuss later, the empirical frequency computed by the global observer converges to the stationary distribution $\pi^{(k)}$ of the Markov chain, which satisfies $\pi^{(k)} B^{(k)} = \pi^{(k)}$ [23]. Consequently, $\hat{r}^{a(k)} = \text{argmax}_r \pi^{(k)}$.

The Belief Consensus Localization (BCL) Algorithm: The BCL algorithm with convergence threshold ϵ^L consists of five steps executed by the regional control centers, and is shown in Fig. 9.

- 1) Flood the MSDs $d_{r,r'}^{(k)}$ so that every region obtains all MSDs in the system. A flooding protocol, such as the one used in OSPF [24] can be used for this purpose.
- 2) Every region verifies that $d_{r,r'}^{(k)} = d_{r',r}^{(k)} \forall r \in \mathcal{R}, r' \in \mathcal{N}(r)$.
- 3) Compute the beliefs $B_{r,r'}^{(k)} \forall r \in \mathcal{R}, r' \in \mathcal{N}(r)$ according to (14). Construct the state transition matrix $B^{(k)}$.
- 4) Compute the stationary distribution $\pi^{(k)}$, the solution to $\pi^{(k)} B^{(k)} = \pi^{(k)}$.
- 5) If $\|\pi^{(k)} - \pi^{(k-1)}\|_\infty < \epsilon^L$ then $k^L = k$. BCL reached convergence, $\hat{r}^{a(k^L)} = \text{argmax}_r \pi^{(k^L)}$.

Fig. 9. Pseudo-code of the BCL Algorithm

Observe that due to Step 2 the attacker cannot tamper with the MSDs sent from region r^a without being noticed, and as a consequence all regions obtain the same matrix $B^{(k)}$ in Step 3. In what follows we show that the proposed BCL algorithm is correct, i.e., all regions identify the same region $\hat{r}^{a(k^L)}$ and the algorithm leads to a solution.

Proposition 3. Consider a system with $R > 2$ regions. If (i) there exists a 3-clique in the graph $G = (\mathcal{R}, E)$ where $E = \{e_{r,r'} | r \in \mathcal{R}, r' \in \mathcal{N}(r)\}$, and (ii) for finite k the DSE does not converge, then the stationary distribution $\pi^{(k)}$ exists, it is unique and it can be computed.

Proof. For sufficiently small k the disagreements between neighboring regions $d_{r,r'}^{(k)} > 0$, because of the initial disagreements on the shared state variables and because of the lack of synchronization to the reference bus. Consequently, the moving average $\tilde{d}_{r,r'}^{(k)} > 0$ since $\alpha^{(k)} > 0$, and so are the beliefs $B_{r,r'}^{(k)} > 0, \forall r, r' \text{ s.t. } r \in \mathcal{N}(r')$. This implies that

the state transition diagram of the Markov chain described by $B^{(k)}$ is a symmetric directed graph, and thus all states of the Markov chain lie in a single communicating class, i.e., the chain is irreducible. Since the Markov chain is irreducible, it has a stationary distribution [23, Proposition 1.14] and this distribution is unique [23, Corollary 1.17]. Although $B_{r,r}^{(k)} = 0 \forall r \in \mathcal{R}$, for $R > 2$ condition (i) ensures that the Markov chain is aperiodic. Aperiodicity in turn is a sufficient condition for the (irreducible) Markov chain to converge to its stationary distribution [23, Theorem 4.9], i.e., the chain is ergodic. Since all regions obtain the same matrix $B^{(k)}$, and the stationary distribution $\pi^{(k)}$ is unique, all regions obtain the same distribution $\pi^{(k)}$. $\square$

The above proposition shows that after a particular iteration k the BCL algorithm is correct. Nonetheless, the stationary distribution $\pi^{(k)}$ is a function of the matrix $B^{(k)}$, which can change at every iteration k . The following proposition establishes the convergence of $\pi^{(k)}$, which implies that the BCL algorithm eventually terminates.

Proposition 4. If $\alpha^{(k)} \rightarrow 0$ as $k \rightarrow \infty$, then $\pi^{(k)} - \pi^{(k-1)} \rightarrow \mathbf{0}_{1 \times R}$. Furthermore, if the attacked system state follows an asymptotically periodic orbit then the stationary distributions $\pi^{(k)}$ converge in k to a stationary distribution vector π^* , and $\hat{r}^{a(k)} \rightarrow \hat{r}^{a*}$.

Proof. We start the proof of Proposition 4 by formulating the following lemma based on results in [25], which will allow us to prove the first part of the proposition ($\pi^{(k)} - \pi^{(k-1)} \rightarrow \mathbf{0}_{1 \times R}$).

Lemma 5. Let C be a right stochastic matrix that describes an irreducible Markov chain with stationary distribution vector $\pi_C = \pi_C C$, and let Π_C be the matrix with the same size as C and all columns equal π_C . Let us denote by $Z = [I - C + \Pi_C]^{-1}$ the fundamental matrix of C . Furthermore, let D be another right stochastic matrix that describes an irreducible Markov chain, and is sufficiently close to C so that all eigenvalues of the differential matrix $U = [D - C]Z$ are strictly less than unity in magnitude. Then

$$\pi_D = \pi_C + \sum_{n=1}^{\infty} \pi_C U^n, \quad (15)$$

and consequently $\pi_D - \pi_C \rightarrow \mathbf{0}_{1 \times |\pi_C|}$ as $D - C \rightarrow \mathbf{0}_{|\pi_C| \times |\pi_C|}$, where $|\pi_C|$ denotes the number of elements in the vector π_C .

Observe that by definition (14)

$$B_{r,r'}^{(k)} - B_{r,r'}^{(k-1)} = \alpha^{(k)} \frac{d_{r,r'}^{(k)} \sum_{r'' \in \mathcal{N}(r)} \tilde{d}_{r,r''}^{(k-1)} - \tilde{d}_{r,r'}^{(k-1)} \sum_{r'' \in \mathcal{N}(r)} d_{r,r''}^{(k)}}{\left(\sum_{r'' \in \mathcal{N}(r)} \tilde{d}_{r,r''}^{(k)} \right) \left(\sum_{r'' \in \mathcal{N}(r)} \tilde{d}_{r,r''}^{(k-1)} \right)}.$$

Consequently $B^{(k)} - B^{(k-1)} \rightarrow \mathbf{0}_{|\pi^{(k)}| \times |\pi^{(k)}|}$ as $\alpha^{(k)} \rightarrow 0$. Let $C = B^{(k-1)}$ and $D = B^{(k)}$. Since the Markov chains described by $B^{(k)}$ and $B^{(k-1)}$ are irreducible (c.f. Proposition 3), the conditions of Lemma 5 are satisfied for k big enough,

Fig. 10. Evolution of the largest probability in the stationary probability vector $\pi^{(k)}$ for six different scenarios. Three smoothing factors $\alpha^{(k)} = 1/k^2$, $\alpha^{(k)} = 1/k$ or $\alpha^{(k)} = 0.2$, and two attack locations $r^a = r_3$ and $r^a = r_6$. The attack size is 0.05.

and thus $\pi^{(k)} - \pi^{(k-1)} \rightarrow 0_{|\pi|}$. Consider now the orbit of the system state for large k . If the attacked system state follows an asymptotically periodic orbit then the disagreements follow an asymptotically periodic orbit too. The smoothing factor by definition satisfies $\sum_{k=0}^{\infty} \alpha^{(k)} = \infty$, and thus if it also satisfies $\alpha^{(k)} \rightarrow 0$ then the smoothed disagreements $\tilde{d}_{r,r'}^{(k)}$ converge to the mean disagreement of the limiting periodic orbit, and so do the beliefs $B_{r,r'}^{(k)} \rightarrow B_{r,r'}^*$. Consequently $\pi^{(k)} \rightarrow \pi^*$ and $\hat{r}^{a(k)} \rightarrow \hat{r}^{a*}$. $\square$

Observe that the proposition does not hold for constant $\alpha^{(k)}$, but it does hold, for example, if $\alpha^{(k)} = 1/k$.

The mitigation algorithm uses BCL to identify the region $\hat{r}^{a*}$ that is most likely to be attacked, isolates the region, and reruns the DSE for the remaining regions until the DSE eventually converges.

B. Numerical results

Fig. 10 shows the evolution of the largest element of the stationary probability vector $\pi^{(k)}$ for different values of the smoothing factor $\alpha^{(k)}$ ($1/k^2$, $1/k$, 0.2), each considered in a separate scenario. We considered two attacked regions, $r^a = r_3$ and $r^a = r_6$; the attack size is 0.05 for which the DSE does not converge. In the case of $\alpha^{(k)} = 1/k^2$, the largest element of $\pi^{(k)}$ converges relatively quickly. The fast convergence of $\pi^{(k)}$ compared to $\alpha^{(k)} = 1/k$ and $\alpha^{(k)} = 0.2$ does, however, come at a price: the identified region $\hat{r}^{a(k)} = \arg\max_r \pi^{(k)}$ is not the attacked region, for both $r^a = r_3$ and $r^a = r_6$ region $\hat{r}^{a(k)} = r_1$ is erroneously identified as attacked. Observe that $\alpha^{(k)} = 1/k^2$ does not satisfy the condition $\sum_{k=0}^{\infty} \alpha^{(k)} = \infty$ required in the definition of $\tilde{d}_{r,r'}^{(k)}$ in Section VI-A, and shows the importance of the condition. For $\alpha^{(k)} = 1/k$ and $\alpha^{(k)} = 0.2$, which do satisfy the condition, the largest element of $\pi^{(k)}$ converges slower, but the attacked region is correctly identified ($\hat{r}^{a(k)} = r^a$) eventually.

Although convergence cannot be guaranteed for a constant smoothing factor $\alpha^{(k)}$, because the condition $\alpha^{(k)} \not\rightarrow 0$ in Proposition 4 is not satisfied, a constant weighting factor is nevertheless useful for exploring the impact of smoothing on the localization time. Since in this case the stationary distribution vector $\pi^{(k)}$ may not converge, there may not exist a k^L for which $\|\pi^{(k^L)} - \pi^{(k^L-1)}\|_{\infty} < \epsilon^L$. Still, after some

Fig. 11. The first identification time (k^F) as a function of the smoothing factor ($\alpha^{(k)}$) considering the attacks in every region $\forall r \in \mathcal{R}$ individually with attack size 0.1 (left) and 0.5 (right)

number of iterations k^F the algorithm can correctly identify $\hat{r}^{a(k^F)}$ as the attacked region, that is, $\hat{r}^{a(k^F-1)} \neq r^a$, but $\hat{r}^{a(k^F)} = r^a \forall k \geq k^F$. We refer to k^F as the *first identification time*.

Fig. 11 shows the first identification time k^F as a function of $\alpha^{(k)}$ considering an attack in various regions $r \in \mathcal{R}$ for attack size 0.1 (left) and 0.5 (right). The first identification time depends on the region that is attacked as well as on the attack size: for larger attack size the localization is significantly faster (localization time is lower). For most of the regions, the optimal $\alpha^{(k)}$ is in the range $(0.2, 0.3)$ and a very high $\alpha^{(k)} > 0.7$ may even make localization fail for the smaller considered attack size (0.1). This observation supports that a small smoothing factor is in general preferable, even if it may lead to a larger localization time.

VII. CONCLUSION

We addressed the vulnerability of fully distributed state estimation to data integrity attacks. We considered an attacker that compromises the communication infrastructure of a single control center and can manipulate the state variables exchanged between the control center and its neighbors. We showed that a denial of service attack can be launched against a state of the art state estimator this way. We proposed an attack detection algorithm based on the convergence properties of the distributed state estimation algorithm and based on the evolution of the exchanged state variables. Furthermore, we proposed an attack mitigation algorithm based on the consensus of the beliefs of the individual regions about the attack location, formulated as the stationary distribution of a random walk on a graph. We established existence, uniqueness, and convergence of the stationary distribution. We showed the efficiency of the attack detection and mitigation algorithms via simulations on an IEEE benchmark power system, and we used the simulations to illustrate the trade-off between localization speed and localization accuracy. Our numerical results also show that strong attacks can often be localized and mitigated faster than weak attacks.

ACKNOWLEDGMENT

This work was partially funded by the EIT ICTLabs through activity ASES 13030 and by the Swedish Research Council through project 2010:5812.

REFERENCES

- [1] A. Monticelli, "Electric power system state estimation," *Proc. of the IEEE*, vol. 88, no. 2, pp. 262–282, 2000.
- [2] A. Abur and A. G. Exposito, *Power System State Estimation: Theory and Implementation*. Marcel Dekker, Inc., 2004.
- [3] Y. Liu, P. Ning, and M. Reiter, "False data injection attacks against state estimation in electric power grids," in *Proc. of the 16th ACM conference on Computer and Communications Security (CCS)*, 2009, pp. 21–32.
- [4] A. Teixeira, G. Dán, H. Sandberg, and K. H. Johansson, "A cyber security study of a SCADA energy management system: Stealthy deception attacks on the state estimator," in *Proc. IFAC World Congress*, Aug. 2011.
- [5] R. B. Bobba, K. M. Rogers, Q. Wang, H. Khurana, K. Nahrstedt, and T. J. Overbye, "Detecting false data injection attacks on dc state estimation," in *Preprints of the First Workshop on Secure Control Systems, CPSWEEK*, Stockholm, Sweden, April 2010.
- [6] O. Kosut, L. Jia, R. Thomas, and L. Tong, "Malicious data attacks on smart grid state estimation: Attack strategies and countermeasures," in *Proc. of IEEE SmartGridComm*, Oct. 2010.
- [7] G. Dán and H. Sandberg, "Stealth attacks and protection schemes for state estimators in power systems," in *Proc. of IEEE SmartGridComm*, Oct. 2010.
- [8] O. Vuković, K. C. Sou, G. Dán, and H. Sandberg, "Network-aware mitigation of data integrity attacks on power system state estimation," *IEEE JSAC: Smart Grid Communications Series*, vol. 30, no. 6, pp. 176–183, July 2012.
- [9] T. T. Kim and H. V. Poor, "Strategic protection against data injection attacks on power grids," *IEEE Trans. on Smart Grid*, vol. 2, pp. 326–333, Jun. 2011.
- [10] A. Giani, E. Bitar, M. Garcia, M. McQueen, P. Khargonekar, and K. Poolla, "Smart grid data integrity attacks: Characterizations and countermeasures," in *Proc. of IEEE SmartGridComm*, Oct. 2011.
- [11] M. Shahidepour and Y. Wang, *Communication and Control in Electric Power Systems*. John Wiley and Sons, 2003.
- [12] A. J. Conejo, S. d. Torre, and M. Canas, "An optimization approach to multiarea state estimation," *IEEE Transactions on Power Systems*, vol. 22, no. 1, pp. 213–221, February 2007.
- [13] L. Xie, D.-H. Choi, S. Kar, and H. V. Poor, "Fully distributed state estimation for wide-area monitoring systems," *IEEE Transactions on Smart Grid*, vol. 3, no. 3, pp. 1154–1169, September 2012.
- [14] X. Li and A. Scaglione, "Robust decentralized state estimation and tracking for power systems via network gossiping," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 7, pp. 1184–1194, July 2013.
- [15] V. Kekatos and G. B. Giannakis, "Distributed robust power system state estimation," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1617–1626, 2013.
- [16] O. Vuković and G. Dán, "On the security of distributed power system state estimation under targeted attacks," in *Proc. of ACM Symposium on Applied Computing (SAC)*, March 2013.
- [17] Y. Feng, C. Foglietta, A. Baiocco, S. Panzieri, and S. D. Wolthusen, "Malicious false data injection in hierarchical electric power grid state estimation systems," in *Proc. of the Fourth International Conference on Future Energy Systems*, ser. e-Energy '13, 2013, pp. 183–192.
- [18] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382–401, July 1982.
- [19] M. J. Fischer, N. A. Lynch, and M. Merritt, "Easy impossibility proofs for distributed consensus problems," in *Proc. of the ACM symposium on Principles of distributed computing*, 1985, pp. 59–70.
- [20] N. H. Vaidya and V. K. Garg, "Byzantine vector consensus in complete graphs," in *Proc. of the ACM Symposium on Principles of distributed computing*, July 2013, pp. 65–73.
- [21] R. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1990.
- [22] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2010.
- [23] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times*. American Mathematical Society, 2009.
- [24] R. Coltun, D. Ferguson, J. Moy, and A. Lindem, "OSPF for IPv6," Internet Requests for Comments, RFC Editor, RFC 5340, July 2008. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc5340.txt>
- [25] P. J. Schweitzer, "Perturbation theory and finite markov chains," *Journal on Applied Probability*, vol. 5, no. 2, pp. 401–413, August 1968.

power systems, power system communication technologies, communication security and availability, and resource management for networked systems.

of Computer Science in 2008, and a Fulbright research scholar at University of Illinois at Urbana-Champaign in 2012–2013. His research interests include the design and analysis of content management and computing systems, game theoretical models of networked systems, and cyber-physical system security in power systems.

Ognjen Vuković is a PhD student at KTH Royal Institute of Technology, Stockholm, Sweden. He received his M.Sc. degree in Telecommunications, System engineering and Radio Communications in 2010 from the Faculty of Electrical Engineering, University of Belgrade. He received his Licentiate Degree in Electrical Engineering in 2013 from KTH The Royal Institute of Technology, Stockholm, Sweden, where he is currently pursuing his PhD in the Laboratory of Communication Networks. His research interests include cyber-physical security of

György Dán is an associate professor at KTH Royal Institute of Technology, Stockholm, Sweden. He received the M.Sc. in computer engineering from the Budapest University of Technology and Economics, Hungary in 1999, the M.Sc. in business administration from the Corvinus University of Budapest, Hungary in 2003, and the Ph.D. in Telecommunications from KTH in 2006. He worked as a consultant in the field of access networks, streaming media and videoconferencing 1999–2001.

He was a visiting researcher at the Swedish Institute