On Security Indices for State Estimators in Power Networks

Henrik Sandberg, André Teixeira, and Karl H. Johansson

Abstract— **In this paper, we study stealthy false-data attacks against state estimators in power networks. The focus is on applications in SCADA (Supervisory Control and Data Acquisition) systems where measurement data is corrupted by a malicious attacker. We introduce two security indices for the state estimators. The indices quantify the least effort needed to achieve attack goals while avoiding bad-data alarms in the power network control center (stealthy attacks). The indices depend on the physical topology of the power network and the available measurements, and can help the system operator to identify sparse data manipulation patterns. This information can be used to strengthen the security by allocating encryption devices, for example. The analysis is also complemented with a convex optimization framework that can be used to evaluate more complex attacks taking model deviations and multiple attack goals into account. The security indices are finally computed in an example. It is seen that a large measurement redundancy forces the attacker to use large magnitudes in the data manipulation pattern, but that the pattern still can be relatively sparse.**

I. Introduction

In Fig. 1, a schematic block diagram of a modern power network control sytstem is shown. The power network models we consider are on the transmission level. They should be thought of as large and consisting of up to hundreds of buses that are spread out over a large geographic area (a region in a country, for example). To monitor and control the behavior of such large-scale systems, SCADA (Supervisory Control and Data Acquisition) systems are used to transmit measurements, status information, and circuit-breaker signals to and from Remote Terminal Units (RTUs) that are connected to substations, see [1]–[3]. For such large-scale systems, lost data and failing sensors are common. The incoming data is therefore often fed to a so-called *state estimator* which provides Energy Management Systems (EMS) and the human operator in the control center with hopefully accurate information at all times.

The technology and the use of the SCADA systems have evolved quite a lot since the 1970s when they were introduced. The early systems were mainly used for logging data from the power network. Today a modern system is supported by EMS such as automatic generation control (AGC), optimal power flow analysis, and contingency analysis (CA), as is indicated in Fig 1. With the advent of new sensors such as PMUs (Phasor Measurement Units), so-called Wide-Area Monitoring and Control Systems (WAMS/WAMC) will also be introduced. This provides yet another layer of control in the modern power network control systems. One motivation for this paper is that SCADA/EMS systems are increasingly more connected to office LANs in the control center. Thus these critical infrastructure systems are potentially accessible from the internet. The SCADA communication network is also heterogeneous and consists of fibre optics, satellite, and microwave connections. Data is often sent without encryption. Therefore many potential security threats exist for modern power control systems, as has been pointed out in for example [4].

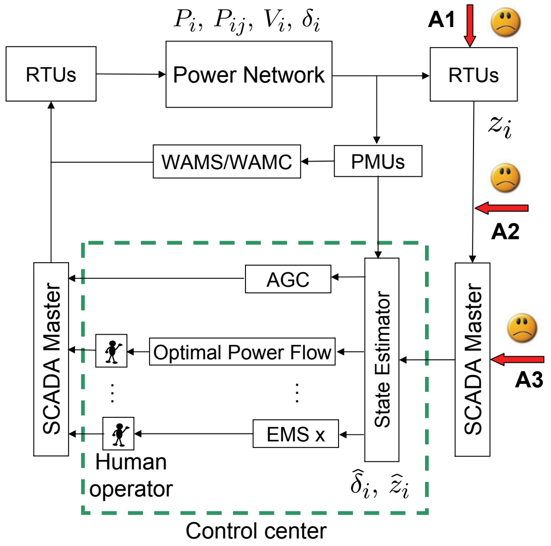

Fig. 1. A schematic block diagram of a power network, a SCADA system, and a control center. Noisy measurements (z_i) of power flows (P_i, P_{ij}) are sent over the SCADA system to the state estimator where estimates of for example the bus phase angels ($\hat{\delta}_i$) are computed. The effect of manipulations on the measurement data z_i are considered in this paper. The manipulations can arise from attacks at various levels A1–A3 in the system. Figure adapted from [4].

The focus of this work is on the state estimator and its so-called Bad Data Detection (BDD) system that is used to remove faulty data, see [2], [3], [5]. The BDD system works by checking that the received data (z_i in Fig. 1) reasonably well matches a physical model of the power network. In the recent paper [6], it was shown how an attacker can avoid triggering the BDD system by coordinated attacks on the measurement data z_i. The attacker can corrupt these data by attacking the RTUs (A1), by tampering with the heterogeneous communication network (A2), or by breaking into the SCADA system through the control center office LAN (A3). In this paper, we further analyze this problem and quantify how sensitive the state estimator is to these

This work is supported in part by the European Commission through the VIKING project, and the ACCESS Linnaeus Center at KTH.

H. Sandberg, A. Teixeira, and K. H. Johansson are with the Automatic Control Lab, School of Electrical Engineering, Royal Institute of Technology, 100 44 Stockholm, Sweden. {hsan,andretei,kallej}@ee.kth.se

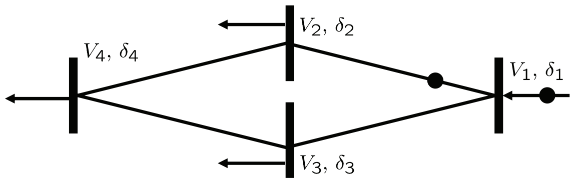

Fig. 2. A simple 4-bus power network. Each bus has a voltage (V_i) and phase angle (δ_i) associated to it. The dots indicate available active power flow measurements.

attacks.

A. Related Work and Contribution of This Paper

False-data injection attacks in power networks were first studied in [6], to the authors' best knowledge. In [6], it was shown that an attacker can manipulate the state estimate while avoiding bad-data alarms. It was also shown that rather simple false-data attacks often can be constructed by an attacker with access to the power network model. The attacker's goal in [6] was either random or targeted false-data attacks. In the targeted attacks, the goal was to change the state estimate into a specific target value.

In this paper, we study a different targeted attack scenario. Here the goal is to manipulate one power flow measurement and to change related measurements in a consistent manner so that no alarms are triggered. Or more accurately: so that the risk of alarms is not increased. At the same time, this shall be done using as small effort as possible. These targeted attacks require less knowledge about the system than the targeted attacks in [6], since the state vector is not necessarily involved. By "small-effort attacks" we here mean either to corrupt as few measurements as possible, or to corrupt the magnitude of the measurement vector as little as possible. The least efforts are then used to define security indices for each targeted measurement. The indices are bounded or computed using simple matrix search techniques or convex optimization. Our study shows that large measurement redundancy gives large magnitude attacks, but that they can still be sparse. Finally, we develop a convex optimization framework that can be used to evaluate false-data attacks which deviate from the model in order to decrease the attack effort and still only marginally increase the risk of a bad-data alarm. Multiple attack goals can also be included in this framework.

II. Power Network Modeling and State Estimation

In this section, we review basic steady-state power network modeling and state-estimation techniques.

A. Active Power Flow Models

It is assumed the power system has $n+1$ buses. Here we will only consider models of the active power flows P_{ij}, active power injections P_i, and bus phase angles δ_i, where $i, j = 1, \dots, n+1$. It is also of interest to study reactive power flows and the voltage levels, but we leave this for future work.

Consider the simple 4-bus power network in Fig. 2. We assume throughout that the power network has reached a steady state. Since measurements are only sent at a low frequency in the SCADA systems, transients cannot be seen in the state estimator. Assuming that the resistance in the transmission line connecting buses i and j is small compared to its reactance, we have that the active power flow from bus i to bus j is [2],

$$P_{ij} = \frac{V_i V_j}{X_{ij}} \sin(\delta_i - \delta_j). \tag{1}$$

At each bus i, active power can also be injected through a generator. Denote this quantity with P_i. A negative P_i indicates a power load. Assuming that there are no losses, conservation of energy yields that for all buses it holds that

$$P_i = \sum_{k \in \mathcal{N}_i} P_{ik}, \tag{2}$$

where $\mathcal{N}_i$ is the set of all buses connected to bus i. The models we use are based on application of (1) and (2) on each bus in the network.

Remark 1: It is possible to include resistive losses in (1) and shunt loads in (2), see [2], but to simplify notation we leave this out.

B. State Estimation

The state-estimation problem we consider consists of estimating n phase angles δ_i given a set of active power flow measurements. One has to fix one (arbitrary) bus phase angle as reference angle, for example $\delta_1 := 0$, and therefore only n angles have to be estimated. The voltage level of each bus is assumed to be known, as well as the reactance of each transmission line.

The m active power flow measurements are denoted by z_i, and are equal to the actual power flow plus independent random measurement noise e_i, which we assume has a Gaussian distribution of zero mean,

$$e = \begin{pmatrix} e_1 \\ \vdots \\ e_m \end{pmatrix} \in \mathcal{N}(0, R),$$

where $R := \mathbf{E} ee^T$ is the diagonal measurement covariance matrix. For the example in Fig. 2 using the indicated measurements of P_1 and P_{12}, we obtain

$$\begin{aligned} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} &= \begin{pmatrix} P_1 \\ P_{12} \end{pmatrix} + \begin{pmatrix} e_1 \\ e_2 \end{pmatrix} \\ &= \begin{pmatrix} \frac{V_1 V_2}{X_{12}} \sin(\delta_1 - \delta_2) + \frac{V_1 V_3}{X_{13}} \sin(\delta_1 - \delta_3) \\ \frac{V_1 V_2}{X_{12}} \sin(\delta_1 - \delta_2) \end{pmatrix} + \begin{pmatrix} e_1 \\ e_2 \end{pmatrix}. \end{aligned}$$

In general, we denote such models by

$$z = P + e = h(x) + e \in \mathbb{R}^m, \tag{3}$$

where $h(x)$ is the power-flow model derived using (1)–(2), and $x \in \mathbb{R}^n$ is a vector of n bus phase angles. Note that here we only analyze the dependence on the phase angles δ_i, and everything else is assumed fixed and known to the

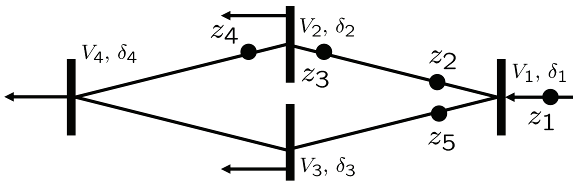

Fig. 3. Same example as in Fig. 2, but with five measurements $z_1 - z_5$ (indicated by dots). This system is observable.

state estimator. This decoupling assumption is common in the literature, see [2], but can be relaxed to include reactive power-flow measurements and bus voltage estimates.

The Gauss-Newton method is often used [2] to estimate the n unknown bus phase angles from power flows measurements z,

$$\hat{x}^{k+1} = \hat{x}^k + (H_k^T R^{-1} H_k)^{-1} H_k^T R^{-1}(z - h(\hat{x}^k)), \qquad (4)$$

where $\hat{x}^k \in \mathbb{R}^n$, k denotes iteration number, and H_k is the Jacobian evaluated at $\hat{x}^k$,

$$H_k := \frac{\partial h}{\partial x}(\hat{x}_k) \in \mathbb{R}^{m \times n}.$$

We will assume the phase differences $\delta_i - \delta_j$ in the power network are all small. Then a linear approximation of (3) is accurate, and we obtain

$$z = Hx + e, \qquad (5)$$

where $H \in \mathbb{R}^{m \times n}$ is a constant Jacobian matrix. The estimation problem (4) can then be solved in one step,

$$\hat{x} = (H^T R^{-1} H)^{-1} H^T R^{-1} z. \qquad (6)$$

The phase-angle estimate $\hat{x}$ can be used to estimate the active power flows by

$$\hat{z} = H\hat{x} = H(H^T R^{-1} H)^{-1} H^T R^{-1} z =: Kz, \qquad (7)$$

where K is the so-called "hat matrix" [2]. The BDD system uses such estimates to identify faulty sensors and bad data by comparing the estimate $\hat{z}$ with z, see below.

As an example, assuming the voltages $V_i = 1$ and reactances $X_{ij} = 1$ for the network in Fig. 2, we obtain the model

$$H = \begin{pmatrix} -1 & -1 & 0 \\ -1 & 0 & 0 \end{pmatrix},$$

where $x = \begin{pmatrix} \delta_2 & \delta_3 & \delta_4 \end{pmatrix}^T$, and $\delta_1 = 0$ is the reference bus. However, $H^T H$ is not invertible and it is not possible to use (6) to obtain a unique estimate $\hat{x}$. This network is therefore called *unobservable* [2]. If we add more measurements, such as in the network in Fig. 3, the model becomes

$$H = \begin{pmatrix} -1 & -1 & 0 \\ -1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & 0 \end{pmatrix}, \qquad (8)$$

where $P = \begin{pmatrix} P_1 & P_{12} & P_{21} & P_{24} & P_{13} \end{pmatrix}^T$. Here $H^T H$ is invertible and it is possible to estimate the phase angles in the system. Assuming the measurement error covariance $R = I$, the hat matrix becomes

$$K = \begin{pmatrix} 0.60 & 0.20 & -0.20 & 0 & 0.40 \\ 0.20 & 0.40 & -0.40 & 0 & -0.20 \\ -0.20 & -0.40 & 0.40 & 0 & 0.20 \\ 0 & 0 & 0 & 1.00 & 0 \\ 0.40 & -0.20 & 0.20 & 0 & 0.60 \end{pmatrix}. \qquad (9)$$

The hat matrix shows how the power flow measurements z are weighted together to form a power flow estimate $\hat{z}$. The rows of the hat matrix can be used to study the measurement redundancy in the system [2]. Typically a large degree of redundancy (many non-zero entries in each row) is desirable to compensate for noisy or missing measurements. In (9), it is seen that all measurements are redundant except the measurement of P_{24} which is called a *critical measurement*. Without the critical measurement observability is lost. From the hat matrix one is lead to believe that the critical measurement is sensitive to attacks. This is indeed the case as we shall see, but also some of the other measurements are sensitive to attacks. This is however not as easy to see from the hat matrix and we therefore take a different approach to quantify the security here.

III. Problem Formulation

The scenario we consider is that an attacker gains access to the measurements through attacks A1–A3, and is able to change some, or all, of the measurements from z into $z_a := z + a$. The *attack vector* a is the corruption added to the real measurement z. The attacker's goal is to fool the EMS and the human operator that a particular power flow measurement is $z_{k,a} := z_k + a_k$ and not z_k, for some k and fixed scalar a_k. A necessary condition for a stealthy attack is that the BDD system is not triggered (or more accurately, that the alarm risk is not increased). To just corrupt the corresponding measurement z_k into $z_k + a_k$ will typically trigger a bad-data alarm, as seen in the next section. We will consider how many, and by how much, other measurements z_i, $i \neq k$, need to be corrupted in coordination with z_k to avoid triggering alarms. A power flow measurement z_k that requires more and larger corruptions to be altered in stealth is here considered more secure, and will obtain larger security indices, as defined below.

Remark 2: An optimal solution to the above problem in terms of the 2-norm of the attack vector a has recently been presented in [7]. The stealthy attack vector a of minimal 2-norm, $\|a\|_2 = \sqrt{a^T a}$, that achieves $z_{k,a} := z_k + a_k$ is given by $a = \frac{a_k}{K_{kk}} K_{\cdot,k}$, where $K_{\cdot,k}$ is the k-th column of the hat matrix (7) using $R = I$. Generally these attack vectors are not *sparse* (except for critical measurements), however. This can be seen in the example (9). The present study is motivated by the fact that an attacker most likely would use sparse attack vectors, and corrupt as few measurement devices as possible.

IV. Sparse Attacks and the Security Index α_k

In the control center, the measurement residual r,

$$r := z - \hat{z} = P + e - H\hat{x} = (I-K)z, \qquad (10)$$

is computed and analyzed in the BDD system. The phase angle estimate $\hat{x}$ is given by (6). If the residual r is larger than expected (measurement errors e will typically make $r \neq 0$), then an alarm is triggered and bad measurements z_i are identified and removed [2], [5], [8]. A key observation in [6] is that an attacker that manipulates the measurements from z into $z_a := z + a$, where $a = Hc \in \mathcal{R}(H)$ and c is an arbitrary vector, is *undetectable* since the residual r is not affected. That certain errors are undetectable by residual analysis has been know for a long time in the power systems community, see for example [5], [8]. It is easy to show that such a lies in the nullspace of $I - K$ in (10). Intuitively this is clear since z_a corresponds to an actual physical state in the power network (minus the measurement error e). The BDD system only triggers when the measurements deviate too much from a possible physical state, at least as long as the linear model is valid.

In light of this, and the problem introduced in Section III, it is natural to consider the following problem:

$$\begin{aligned} \alpha_k &:= \min_c \|Hc\|_0 \\ &\text{such that } 1 = \sum_i H_{ki} c_i, \end{aligned} \qquad (11)$$

where $\|Hc\|_0$ denotes the number of non-zero elements in the vector $a = Hc$, and H_{ki} is the element (k, i) of H. In (11), we optimize over all corruptions $a = Hc \in \mathcal{R}(H)$ that do not trigger bad-data alarms. A solution c^* to (11) can be re-scaled to obtain $a^* = a_k Hc^*$ such that the measurement attack $z_a = z + a^*$ achieves the attacker's goal $z_{k,a} = z_k + a_k$, and at the same time corrupts as few measurements as possible. In total, $\alpha_k = \|a^*\|_0$ measurements have to be corrupted to manipulate the measurement z_k. Unfortunately, the problem (11) is non-convex and is generally hard to solve for large problems. However, it is easy to get bounds on α_k even for large models, as shown next.

It is clear that the lower bound $\alpha_k \geq 1$ holds, since at least one measurement (z_k) is corrupted. One can also show that if measurement z_k is a critical measurement, then $\alpha_k = 1$. A simple upper bound can be achieved by looking at the k-th row of H: Every column of H with a non-zero entry in the k-th row can be used to construct a false-data attack vector a that achieves the attack goal. Assume that H_{ki} is non zero. Then the attack vector

$$a_k^i := \frac{a_k}{H_{ki}} H_{\cdot,i},$$

where $H_{\cdot,i}$ denotes the i-th column of H, achieves the attack goal. By selecting the sparsest vector among all a_k^i, we obtain an upper bound $\bar{\alpha}_k^1$ on α_k. Formally we have,

$$\bar{\alpha}_k^1 := \min_{i: H_{ki} \neq 0} \|H_{\cdot,i}\|_0.$$

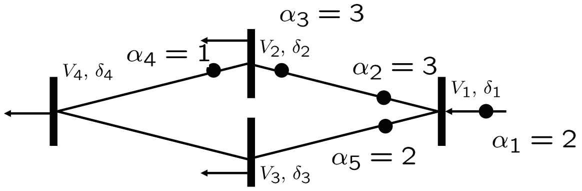

Fig. 4. A power network and its security indices α_k. The flow P_{24} with $\alpha_4 = 1$ is easiest to attack. Only one measurement has to be corrupted. The flows P_{21} and P_{12} with index $\alpha_2 = \alpha_3 = 3$ are hardest to attack, and require a coordinated attack involving three sensors.

TABLE I

The security index α_k, the bound $\bar{\alpha}_k^1$, and the sparsest attack vectors for the power network in Fig. 4

Measurement	Power flow	α_k	$\bar{\alpha}_k^1$	a^*
z_1	P_1	2	2	$\begin{pmatrix}1 & 0 & 0 & 0 & 1\end{pmatrix}^T$
z_2	P_{12}	3	4	$\begin{pmatrix}1 & 1 & -1 & 0 & 0\end{pmatrix}^T$
z_3	P_{21}	3	4	$\begin{pmatrix}-1 & -1 & 1 & 0 & 0\end{pmatrix}^T$
z_4	P_{24}	1	1	$\begin{pmatrix}0 & 0 & 0 & 1 & 0\end{pmatrix}^T$
z_5	P_{13}	2	2	$\begin{pmatrix}1 & 0 & 0 & 0 & 1\end{pmatrix}^T$

Since H is typically sparse for power networks, this bound seems many times to be pretty good and is also very fast to compute. A second upper bound, $\bar{\alpha}_k^2$, is discussed in the next section, and the best of them can be used as an upper bound of α_k

$$\bar{\alpha}_k := \min\{\bar{\alpha}_k^1, \bar{\alpha}_k^2\}. \qquad (12)$$

Obtaining better easily computed bounds, or even to characterize the exact solution of (11) is an interesting problem for future work.

Remark 3: To obtain a better bound $\bar{\alpha}_k^1$, one can include a column in H that corresponds to the reference bus ($\partial h/\partial \delta_1$).

In Fig. 4 and in Table I, the security indices α_k and sparse attack vectors for the model (8) are shown. The index makes it easy to locate flows whose measurements are relatively easy to attack without triggering bad-data alarms. In this example, the critical measurement of P_{24} with $\alpha_4 = 1$ is easiest to attack, and P_{21} and P_{12} with index $\alpha_2 = \alpha_3 = 3$ are hardest to attack. It is also seen that the upper bound $\bar{\alpha}_k^1$ is tight in most cases.

Comparing with the hat matrix (9), it is seen that the number of non-zero elements in each row of the hat matrix is not correlated to the number of sensors that has to be involved in a stealthy attack, except in the case of critical measurements (z_4). For example, the measurement z_1 is quite redundant since the estimate $\hat{z}_1$ depends on z_1, z_2, z_3, z_5. But in fact only two measurements (z_1, z_5) have to be manipulated when z_1 is attacked. A large diagonal entry in the hat matrix K seems correlated with a smaller security index, however. Nevertheless, it is not clear from the hat matrix how many, and which, measurements that can be involved in a false-data attack. Hence it seems that measurement redundancy analysis as commonly performed in power systems is not appropriate to evaluate the system's security, and the introduction of other metrics is appropriate.

V. Small Magnitude Attack Vectors and the Security Index β_k

Next we consider a different security index which we denote by β_k. The security index α_k is appropriate to measure resistance against an attacker with limited access to the number of measurements. However, the magnitude of the elements in a sparse attack vector a can be large, and this can be an issue since the power system is nonlinear. An attack vector a with large elements may push the estimator into the nonlinear regime which may lead to bad-data alarms even if $a \in \mathcal{R}(H)$, or non-convergence of the Gauss-Newton method (4). Thus an attacker may want to construct small magnitude attack vectors while achieving his goals. It is also well known that the minimization of the 1-norm that we use below often gives rise to sparse solutions, see for example [9]. Therefore it seems that β_k is a good compromise between a sparse and a small attack vector. The method we introduce below is also based on convex optimization tools, and it is relatively easy to extend this framework to include multiple attack goals and model deviations etc.

The 1-norm of an attack vector a is $\|a\|_1 := \sum_i |a_i|$. This is a measure of the total amount of changes added to the measurement vector z. Let us next study the convex optimization problem

$$\begin{aligned} \beta_k &:= \min_c \|Hc\|_1 \\ &\text{such that } 1 = \sum_i H_{ki} c_i, \end{aligned} \tag{13}$$

which can be re-cast into a linear program. A solution c^* to (13) can be re-scaled to obtain $a^* = a_k H c^*$ such that the measurement attack $z_a = z + a^*$ achieves $z_{k,a} = z_k + a_k$, and at the same time the minimal amount of additional power, $\|a^*\|_1$, is added to the measurement vector z. We can interpret the dimensionless quantity β_k as the minimal possible amplification of the attack a_k: The attacker wants to add a_k MWs to the power-flow measurement z_k, but must in the process of doing so add a total change of $\beta_k a_k$ MWs to z in order to avoid triggering alarms.

Remark 4: Since the 1-norm optimal solutions a^* often are sparse, a natural upper bound of α_k is

$$\bar{\alpha}_k^2 := \|a^*\|_0,$$

to be used in (12). One could consider to possibly further improve the bound by using reweighted 1-norm minimization [9].

Remark 5: It is clear that the lower bound $\beta_k \geq 1$ holds. We also have the upper bound $\beta_k \leq \min_{j:H_{kj}\neq 0} \sum_i |H_{ij}/H_{kj}|$. But since β_k can be computed exactly using tools such as CVX [10], these bounds do not seem as important as the bounds on α_k.

It is possible to refine the index β_k to take more complex attack scenarios into account, as long as the constraints are convex. For example, the attacker may be willing to take risks and slightly increase the chance of bad-data alarms. By adding a bias $d \notin \mathcal{R}(H)$ to the attack vector, $a = Hc + d$, it no longer lies in the nullspace of $I - K$, and the risk of a bad-data alarm is increased. The benefit of introducing a bias (from the attacker's point of view) is that it may decrease the size of a and increase its sparsity. It would also be possible to interpret d as an error in the attacker's model.

The measurement residual r (10) in the BDD system is distributed according to

$$r \in \mathcal{N}(Sd, \Omega), \quad \Omega := SR,$$

where $\mathcal{N}$ is the Gaussian distribution, Ω the covariance, and Sd the expected value of the residual. $S := I - K$ is the so-called *residual sensitivity matrix* [2] (remember that K is the hat matrix (7)). Hence $d \neq 0$ changes the expected value of the residual. But it should be clear that if the normalized residual $\|\text{diag}(\Omega)^{-1/2} Sd\|_p$ is small, the risk of a bad-data alarm is still small. Hence, one can introduce a security index β_k^ϵ by

$$\begin{aligned} &\beta_k^\epsilon := \min_a \|a\|_1 \\ &\text{such that } 1 = a_k, \quad \|\text{diag}(\Omega)^{-1/2} Sa\|_p \leq \epsilon, \end{aligned} \tag{14}$$

where we have used that $Sa = S(Hc+d) = Sd$. Depending on the exact BDD system that is being used by the SCADA-system operator and the choice of integer p, the size of ϵ can be related to an increase in probability of a bad-data alarm, see [7]. Common BDD-methods include chi-squares tests and normalized residual tests [2]. Note that the attacker needs to be more informed to solve (14) than to solve (13) since R is needed.

It is also clear that the above framework can be generalized to study attacks with coordinated goals. The optimization problem

$$\begin{aligned} &\min_a \|a\|_1 \\ &\text{such that } a \in \mathcal{G}, \quad \|\text{diag}(\Omega)^{-1/2} Sa\|_p \leq \epsilon, \end{aligned} \tag{15}$$

where $\mathcal{G}$ is a convex set of attack goals, possibly involving more than one measurement, is one such generalization. For example, $\mathcal{G}$ could be intervals such as $\mathcal{G} = \{0.9 \leq a_1 \leq 1.1, -1.1 \leq a_2 \leq -0.9\}$. By solving (15) for various scenarios it is possible for the SCADA-system operator to test the security of the state estimator.

VI. Example: The IEEE 14-bus Power Network

Here we consider the IEEE 14-bus benchmark power network that was also analyzed in [6]. A different perspective is taken here and we compute its security indices and compare with two heuristic redundancy measures. For the computations, the MATLAB package MATPOWER [11] and the optimization toolbox CVX [10] are used. Power flow measurements are added at each bus, and at every end of every interconnecting transmission line. In total there are $m = 54$ measurements, all assumed equally good $R = I$, and the matrix H has size 54×13. This considered system has more measurements than is normal in a power system, and should therefore have large measurement redundancy. The question is: Does this imply security against false-data attacks?

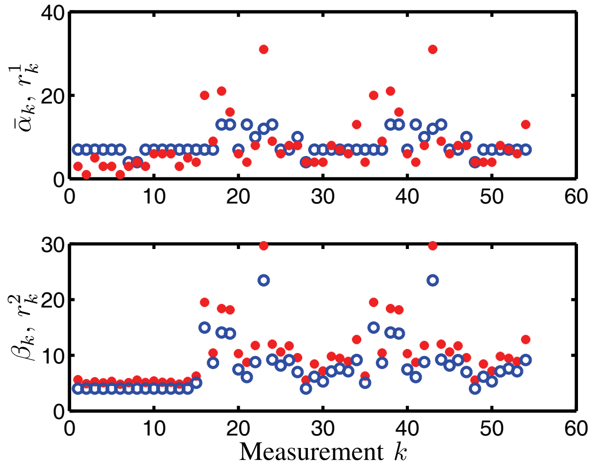

Fig. 5. In the upper plot, the security index bound $\bar{\alpha}_k$ (blue rings) and the redundancy measure r_k^1 (red full circles) are plotted versus measurement number. In the lower plot, the security index β_k (blue rings) and the redundancy measure r_k^2 (red full circles) are plotted. There is no simple connection between $\bar{\alpha}_k$ and r_k^1, whereas the variations in β_k and r_k^2 correlate very well.

In Fig. 5, the security indices $\bar{\alpha}_k$ (bound) and β_k are plotted versus measurement number. For comparison, two heuristic measurement redundancy quantities are also plotted. These are defined by

$$\begin{aligned} r_k^1 &:= \#\{|K_{ik}/K_{kk}| \geq 0.33;\ i = 1, \ldots, m\} \geq 1, \\ r_k^2 &:= \sum_i |K_{ik}/K_{kk}| \geq 1, \end{aligned}$$

where K is the hat matrix (7). The scaled columns of K are minimal stealthy 2-norm attacks, see Remark 2. Hence these are valid attack vectors, and $\beta_k \leq r_k^2$ with equality for critical measurements. The quantity r_k^1 counts the number of elements in such an attack vector whose magnitude is at least 33% of the attacked measurement. One could expect that those large elements are involved in a sparse attack, and would give a good estimate of α_k. The number 33% is chosen somewhat arbitrarily. However, in these numerical experiments r_k^1 always failed to give accurate predictions of α_k no matter this choice.

As seen in the upper plot of Fig. 5, there is no simple connection between the sparsity of possible attacks (or at least with the bound $\bar{\alpha}_k$) and the quantity r_k^1. Sometimes r_k^1 is too large, and sometimes too small, and it is hard to conclude anything other than that this heuristic must be considered as bad. The number of sensors needed for an attack seemingly has little to do with it.

In the lower plot, the index β_k is plotted together with r_k^2. There is clearly strong correlation between variations in β_k and r_k^2. Maybe this is not so surprising given Remark 2. But note that the optimal 1-norm attacks often are much sparser. To summarize: Large measurement redundancy in terms of r_k^2 seems to give larger security with respect to the security measure β_k (attack vector magnitude), but the quantity r_k^1 has little to do with the security measure α_k (attack vector sparsity).

VII. Summary and Future Work

In this paper, we have introduced two security indices for state estimators in power networks. The indices help to locate power flows whose measurements are potentially easy to manipulate. Large indices indicate that a large coordinated attack is needed in order to not trigger an alarm in the control center. We also showed how convex optimization tools can be used to evaluate attacks, taking deviations from the exact power system model and multiple attack goals into account. We have also seen that simple measurement redundancy quantities seem to give security in terms of attack vector magnitude, but not in terms of attack vector sparsity. This was demonstrated on an IEEE 14-bus network with large measurement redundancy.

For future work, we intend to study how one can use these indices and tools to increase the security. It is also interesting to study the influence of model errors in H.

Acknowledgments

The authors would like to thank Mr. Zhu Kun and Dr. György Dán for helpful and stimulating discussions.

References

[1] M. Shahidehpour, W. F. Tinney, and Y. Fu, "Impact of security on power systems operation," *Proceedings of the IEEE*, vol. 93, no. 11, pp. 2013–2025, 2005.

[2] A. Abur and A. G. Exposito, *Power System State Estimation: Theory and Implementation*. Marcel Dekker, Inc., 2004.

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

[4] A. Giani, S. Sastry, K. H. Johansson, and H. Sandberg, "The VIKING project: An initiative on resilient control of power networks," in *Proceedings of the 2nd International Symposium on Resilient Control Systems*, Idaho Falls, Idaho, 2009.

[5] L. Mili, T. V. Cutsem, and M. Ribbens-Pavella, "Bad data identification methods in power system state estimation - a comparative study," *IEEE Transactions on Power Apparatus and Systems*, vol. 104, no. 11, pp. 3037–3049, Nov. 1985.

[6] Y. Liu, P. Ning, and M. Reiter, "False data injection attacks against state estimation in electric power grids," in *Proceedings of the 16th ACM conference on Computer and communications security*, Chicago, Illinois, 2009, pp. 21–32.

[7] A. Teixeira, S. Amin, H. Sandberg, K. H. Johansson, and S. S. Sastry, "Cyber-security analysis of state estimators in electric power systems," Submitted to IEEE Conference on Decision and Control, March 2010.

[8] F. F. Wu and W.-H. E. Liu, "Detection of topology errors by state estimation," *IEEE Transactions on Power Systems*, vol. 4, no. 1, pp. 176–183, Feb. 1989.

[9] E. Candès, M. Wakin, and S. Boyd, "Enhancing sparsity by reweighted l_1 minimization," *J. Fourier Anal. Appl.*, vol. 14, pp. 877–905, 2008.

[10] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming (web page and software)," http://stanford.edu/ boyd/cvx, June 2009.

[11] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, "MATPOWER's extensible optimal power flow architecture," in *Power and Energy Society General Meeting*. IEEE, July 2009, pp. 1–7.