Stealth Attacks and Protection Schemes for State Estimators in Power Systems

György Dán and Henrik Sandberg
ACCESS Linnaeus Centre, School of Electrical Engineering
KTH, Royal Institute of Technology
Stockholm, Sweden
{gyuri,hsan}@ee.kth.se

***Abstract*—State estimators in power systems are currently used to, for example, detect faulty equipment and to route power flows. It is believed that state estimators will also play an increasingly important role in future smart power grids, as a tool to optimally and more dynamically route power flows. Therefore security of the estimator becomes an important issue. The estimators are currently located in control centers, and large numbers of measurements are sent over unencrypted communication channels to the centers. We here study stealthy false-data attacks against these estimators. We define a security measure tailored to quantify how hard attacks are to perform, and describe an efficient algorithm to compute it. Since there are so many measurement devices in these systems, it is not reasonable to assume that all devices can be made encrypted overnight in the future. Therefore we propose two algorithms to place encrypted devices in the system such as to maximize their utility in terms of increased system security. We illustrate the effectiveness of our algorithms on two IEEE benchmark power networks under two attack and protection cost models.**

I. Introduction

SCADA (Supervisory Control and Data Acquisition) systems are widely used to monitor and control the behavior of large-scale power systems. SCADA systems transmit measurement data, status information, and control signals to and from Remote Terminal Units (RTUs), which are located in substations in the grid, see for example [1], [2]. 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), see for example [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 fiber optics, satellite, and microwave connections. Data is often sent without encryption. Therefore many potential cyber security threats exist for modern power control systems, as has been pointed out in for example [3], [4]. Another motivation for this work is that future smart power grids are believed to be more dependent on accurate state estimators to fulfill their task of optimally and dynamically routing power flows. Resilience and security of smart power grids are addressed in for example [5].

In this paper, the focus is on stealth attacks (also called false-data injection attacks) against state estimators. This type of attacks was first studied in [4], to the authors' best knowledge. In [4], it was shown that an attacker can manipulate the state estimate while avoiding bad-data alarms in the control center. It was also shown that rather simple false-data attacks often can be constructed by an attacker with access to the power network model. More recently in [6], [7], further aspects of these attacks were studied. In [7], two security indices were defined that quantify how difficult it is to perform a successful stealth attack against particular measurements. In [6], it was shown how one can completely protect a state estimator from these unobservable attacks by encrypting a sufficient number of measurement devices.

Here we extend the work in [6], [7]. First of all, we propose an efficient method for computing the security index α_k introduced in [7] for sparse stealth attacks. This index is relevant to the problem because it quantifies the minimum number of measurements that need to be corrupted to perform a stealth attack with a specific goal. We also propose an extension where clusters of measurements are available at the same cost for the attacker. This is a realistic scenario if an attack is taking place from a substation, and potentially all measurements originating from the substation can be corrupted at once. Finally, we propose a protection scheme for how to allocate encryption devices to strengthen security. In [6], it is shown exactly how many measurements need to be encrypted to ensure security. It is shown that the number is equal to the number of state variables in the system. In this paper, we use the introduced security index to quantify the security when the number of encrypted measurements is insufficient to provide complete security, but one would like to maximize the

This work is supported in part by the European Commission through the FP7 project VIKING, and by the ACCESS Linnaeus Centre at KTH.

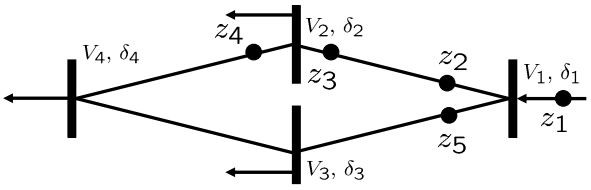

Fig. 1. A simple small 4-bus power network. Each bus has a voltage (V_i) and phase angle (δ_i) associated to it. The dots indicate available power flow measurements.

usefulness of the encrypted devices one can deploy. This is a quite realistic scenario, since the number of state variables in current systems is typically large, while measurement devices that allow for encryption seem to be rare, and their installation incurs costs. The framework can be adapted to attacks that are not unobservable, by extending the definition of α_k, as discussed in [7].

The organization of the paper is as follows. In Section II, basic power systems modeling and state estimation are introduced. In Section III, stealth attacks are defined, and efficient methods for computing the introduced security index are presented. In Section IV, protection schemes based on the security index are described, and in Section V numerical examples are considered.

II. Power Network Modeling and State Estimation

In this section, we review basic steady-state power network modeling and state-estimation techniques. Much more complete presentations are found in [1], [2], for example.

Here it is assumed that the power system has $n+1$ buses. We will only consider models of the active power flows P_{ij}, active power injections P_i, and bus phase angles δ_i, where $i,j = 1,\ldots,n+1$. 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. 1. 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 current state estimators. Assuming that the resistance in the transmission line connecting buses i and j is small compared to its reactance X_{ij}, we have that the active power flow from bus i to bus j is [1], $P_{ij} = \frac{V_i V_j}{X_{ij}} \sin(\delta_i - \delta_j)$. 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 $P_i = \sum_{k\in\mathcal{N}_i} P_{ik}$, where $\mathcal{N}_i$ is the set of all buses connected to bus i. The models we use below are based on application of these equations.

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. Thus $e = \begin{pmatrix} e_1 & \ldots & e_m \end{pmatrix}^T \in \mathcal{N}(0,R)$, where $R := \mathbf{E}ee^T$ is the diagonal measurement covariance matrix. For the example in Fig. 1 using the measurements z_1 and z_2, we obtain

$$\begin{aligned}\begin{pmatrix} z_1 \\ z_2 \end{pmatrix} &= \begin{pmatrix} P_1 \\ P_{12} \end{pmatrix} + \begin{pmatrix} e_1 \\ e_2 \end{pmatrix} \\ &= \begin{pmatrix} \frac{V_1V_2}{X_{12}} \sin(\delta_1-\delta_2) + \frac{V_1V_3}{X_{13}} \sin(\delta_1-\delta_3) \\ \frac{V_1V_2}{X_{12}} \sin(\delta_1-\delta_2) \end{pmatrix} + \begin{pmatrix} e_1 \\ e_2 \end{pmatrix}.\end{aligned}$$

In general, we denote such models by

$$z = P + e = h(x) + e \in \mathbb{R}^m, \tag{1}$$

where $h(x)$ is the power-flow model derived using the above power flow equations, and $x \in \mathbb{R}^{n+1}$ is a vector of $n+1$ unknown bus phase angles.

The state-estimation problem we consider consists of estimating n phase angles δ_i given a set of m 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 V_i of each bus is assumed to be known, as well as the reactance of each transmission line. Note that here we only analyze the dependence on the phase angles δ_i, and everything else is assumed fixed and known to the state estimator. This decoupling assumption is common in the literature, see [1], but can be relaxed to include reactive power-flow measurements and bus voltage estimates by including more unknown states. The Gauss-Newton method is often used [1] to estimate the unknown bus phase angles from power flows measurements z,

$$\hat{x}^{k+1} = \hat{x}^k + (\bar{H}_k^T R^{-1} \bar{H}_k)^{-1} \bar{H}_k^T R^{-1} (z - h(0,\hat{x}^k)), \tag{2}$$

where $\hat{x}^k \in \mathbb{R}^n$ is the estimate of the n unknown phase angles, k denotes iteration number, and H_k is the Jacobian evaluated at $\hat{x}^k$ and reference angle equal to zero, $H_k := \frac{\partial h}{\partial x}(0,\hat{x}_k) \in \mathbb{R}^{m\times(n+1)}$. With $\bar{H}_k \in \mathbb{R}^{m\times n}$ we mean H_k with the column corresponding to the reference angle being removed. We will assume the phase differences $\delta_i - \delta_j$ in the power network are all small. Then a linear approximation of (1) is accurate, and we obtain

$$z = Hx + e, \tag{3}$$

where $H \in \mathbb{R}^{m\times(n+1)}$ is a constant Jacobian matrix. The estimation problem (2) can then be solved in one step,

$$\hat{x} = (\bar{H}^T R^{-1} \bar{H})^{-1} \bar{H}^T R^{-1} z, \tag{4}$$

where again $\bar{H}$ denotes the Jacobian with the column corresponding to the reference angle being removed. The phase-angle estimate $\hat{x}$ can be used to estimate the active power flows by $\hat{z} = \bar{H}\hat{x} = \bar{H}(\bar{H}^T R^{-1} \bar{H})^{-1} \bar{H}^T R^{-1} z$. The Bad Data Detection (BDD) system in the control center calculates the measurement residual r,

$$r := z - \hat{z} = P + e - \bar{H}\hat{x} = (I - K)z, \tag{5}$$

where the phase angle estimate $\hat{x}$ is given by (4). 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 [1], [8], [9]. A necessary condition for a successful stealth attack is to avoid such alarms.

III. Stealth Attacks

We consider a power system that consists of $n+1$ buses, and there are m meters that provide power flow measurements z to the state estimator via some, possibly shared, communication channels. An attacker is able to change some, or all, of the measurements from z into $z_a := z + a$, either by physically tampering with the individual meters or by getting access to some communication channels. 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 stealth 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.

A key observation in [4] is that an attacker that manipulates the measurements from z into $z_a := z + a$, where $a = Hc \in \mathcal{R}(H)$ and $c \in \mathbb{R}^{n+1}$ is an arbitrary vector, is *undetectable* since the residual r is not affected. That certain errors are undetectable by residual analysis has been known for a long time in the power systems community, see for example [8], [9]. It is easy to show that such a lies in the nullspace of $I - K$ in (5). 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.

A. Attack and protection cost model

To capture the cost of the attacker and the system operator we introduce a partition $\mathcal{M} = \{M_1, \ldots, M_{|\mathcal{M}|}\}$ of the set of measurements $\{1, \ldots, m\}$. The attacker can attack any number of measurements in the same block M_j of the partition at unit cost. For the operator, all measurements belonging to the same block M_j can be protected at unit cost. We denote by S the $|\mathcal{M}| \times m$ matrix whose element $S_{jk} = 1$ if $k \in M_j$, and $S_{jk} = 0$ otherwise. The cost of an attack a for the attacker is then $\|S|a|\|_0$, the number of non-zero elements in the vector $S|a|$. By $|a|$ we denote the vector of the magnitudes of the elements in a. We denote the subset of the partition protected by the operator by $\mathcal{P} \subseteq 2^{\mathcal{M}}$.

This notation allows to consider various attack and protection cost models. We consider two cost models throughout the paper.

Stealth meter attacks: This scenario corresponds to a partition $\mathcal{M} = \{\{1\}, \ldots, \{m\}\}$, i.e., every measurement is a partition block. The attacker has to gain access to each individual meter it needs to compromise in order to achieve its attack goal. The cost of the attacker is the number of meters that have to be compromised. Similarly, the protection cost of the operator is the number of meters that are protected. This scenario corresponds to physically tampering with the individual meters.

Stealth RTU attacks: This scenario corresponds to a partition of size $|\mathcal{M}| = n+1$ in which the measurements in a bus form a partition block, and there is an RTU associated to every bus. An attacker that gains access to an RTU or its communication channel can compromise any number of measurements associated with the RTU. The cost of the attacker is the number of compromised RTUs. Similarly, the protection cost of the operator is the number of RTUs that are protected. This scenario corresponds to attacks on the communication channels that carry the measurement data from individual RTUs, typically the load and branch power flows into the corresponding bus.

B. Minimum cost stealth attacks

In general the attacker can use any undetectable attack vector $a, a_k \neq 0$ to attack measurement k. A rational attacker would, however, be interested in finding an attack vector $a, a_k \neq 0$ with minimum cost, i.e., the number of partition blocks to which the compromised meters belong should be minimal, with the constraint that the attacker cannot compromise any protected measurement $k \in \mathcal{P}$. In order to find a minimal stealth attack on measurement k the attacker has to solve the problem

$$\begin{aligned} &\alpha_k := \min_c \|S|Hc|\|_0 \\ &\text{such that } 1 = \sum_i H_{ki} c_i, \\ &(Hc)_j = 0 \quad \forall j \in \mathcal{P}, \end{aligned} \tag{6}$$

where $\|.\|_0$ denotes the number of non-zero elements in a vector, and H_{ki} is the element (k,i) of H. In (6), we optimize over all corruptions $a = Hc \in \mathcal{R}(H)$ that do not trigger bad-data alarms and do not involve compromising protected measurements. A solution c^* to (6) can be re-scaled to obtain $a^* = a_k H c^*$ 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 blocks of measurements as possible. In total, $\alpha_k = \|S|a^*|\|_0$ blocks of measurements have to be corrupted to manipulate the measurement z_k. Unfortunately, the problem (6) is non-convex and is generally hard to solve for large problems. In the following we first describe an upper bound on α_k, then we describe an algorithm that calculates the optimal solution by exploiting the topology of the power network graph.

1) Upper bound on the minimum cost: A simple upper bound on α_k can be obtained by looking at the k-th row of H. Any column i 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, if $H_{ji} = 0 \ \forall j \in \mathcal{P}$. 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 $S|a_k^i|$, we obtain an upper bound $\hat{\alpha}_k$ on α_k. Formally we have,

$$\hat{\alpha}_k := \min_{i:H_{ki} \neq 0} \|S|H_{\cdot,i}|\|_0. \tag{7}$$

Since H is typically sparse for power networks, this bound is very fast to compute, and exists whenever $\mathcal{P} = \emptyset$.

TABLE I
THE ITERATIVE PATH AUGMENTATION ALGORITHM USED TO CALCULATE THE ATTACKS WITH MINIMAL COST FOR MEASUREMENT k.

$\mathcal{A}^{(1)} = \{M_j\},\ k \in M_j,\ \mathcal{A}^* = \emptyset$
for i = 1 to $|\mathcal{M}| - |\mathcal{P}|$
  for $A \in \mathcal{A}^{(i)}$
    $A' = \{l | l \in A, \nexists j \notin A\ s.t.\ j \sim l\}$
    if $rank(H_{\{1,\dots,m\}\setminus A'}) = n-1$ and $rank(H_{(\{1,\dots,m\}\setminus A')\cup\{k\}}) = n$ then
      $\mathcal{A}^* = \mathcal{A}^* \cup A$
    end if
  end for
  if $\mathcal{A}^* \neq \emptyset$ then return $\mathcal{A}^*$
  for $A \in \mathcal{A}^{(i)}$
    for $M_j \subseteq A$
      for $M_k \in \mathcal{N}(M_j), M_k \cap \mathcal{P} = \emptyset, M_k \cap A = \emptyset$
        $\mathcal{A}^{(i+1)} = \mathcal{A}^{(i+1)} \cup (A \cup M_k)$
      end for
    end for
  end for
end for

C. Finding the minimum cost attack

Finding α_k is equivalent to finding a set of rows $N \subseteq \{1,\dots,m\} \setminus \{k\}$ that is maximal in terms of the number of partition blocks M_j it covers, and for which the following two conditions hold

$$rank(H_N) = n-1, \quad (8)$$

$$rank(H_{N\cup\{k\}}) = n, \quad (9)$$

where H_N is the submatrix of H formed by the rows in N. Given N the attack can be constructed by calculating the nullspace of the submatrix H_N, which is 1 dimensional due to the rank-nullity theorem. Since $\forall c \in null(H_N)$ we have $(Hc)_k = 0\ \forall k \in N$, and N is maximal, it follows that $\alpha_k = \|S|Hc|\|_0$.

In general, finding the maximal set N is a combinatorial optimization problem. For sparse power network graphs, however, it is possible to calculate the optimal solutions even for systems with hundreds of state variables and measurement points using the iterative path augmentation algorithm described in the following.

The iteration starts with an attack that consists of the partition block to which measurement k belongs, $\mathcal{A}^{(1)} = \{M_j\}, k \in M_j$. In iteration i the algorithm first considers all attacks of cost i. For every attack $A \in \mathcal{A}^i$ it creates the corresponding attack A' by only keeping the rows l of H for which there is no row j not in attack A that is linearly dependent on row l ($l \sim j$). It then verifies if the set $N = \{1,\dots,m\} \setminus A'$ satisfies the rank conditions (8) and (9). If no such attack is found, the algorithm augments every attack $A \in \mathcal{A}^i$ of cost i with one additional partition block M_k that is unprotected ($M_k \cap \mathcal{P} = \emptyset$) and is neighboring to a partition block already in the attack ($M_k \in \mathcal{N}(M_j)$ for some $M_j \subseteq A$). The pseudo-code of the algorithm is shown in Table I.

IV. PROTECTION AGAINST STEALTH ATTACKS

In this section we consider that the operator has a budget π in terms of the number of protected measurement partition blocks that it can spend. Thus, $C_{\mathcal{M}}(\mathcal{P}) \leq \pi$, where $\mathcal{P}$ denotes the set of chosen protected measurements, and $C_{\mathcal{M}}(\mathcal{P})$ denotes the cost of protecting $\mathcal{P}$ considering the partition $\mathcal{M}$, and can be calculated as the number of partition blocks M_j s.t. $M_j \cap \mathcal{P} \neq \emptyset$. The goal of the operator is to achieve the best possible protection of the state estimator against stealth attacks given its budget. Instead of picking the set of protected measurements at random, as done in [6], we propose three deterministic algorithms to choose the set of protected measurements. We evaluate the performance of the algorithms in Section V.

A. Perfect protection

Ideally, the set of protected measurements $\mathcal{P}$ should be such that no stealth attacks are possible, i.e., $\alpha_k = \infty, \forall k \in \{1,\dots,m\}$. We refer to such a protection as *perfect protection*.

1) Stealth meter attacks: In the case of meter attacks in order to achieve perfect protection it is necessary and sufficient for the operator to protect $|\mathcal{P}| = n$ measurements chosen such that $rank(H_{\mathcal{P}}) = n$ [6]. The budget required to achieve perfect protection is thus $\pi = n$.

2) Stealth RTU attacks: In the case of RTU attacks the condition $\pi = n$ is not necessary, since we now count the number of protected blocks, which can contain more than one measurement each. In particular the following result holds.

Definition 1: Let us call the *RTU level power network graph* the graph where each vertex is an RTU in the power system, and every edge is a transmission link between the RTUs. A dominating set $\mathcal{P}$ of the RTU level power network graph is a subset of vertices such that each vertex not in $\mathcal{P}$ is adjacent to at least one member in $\mathcal{P}$.

Proposition 1: Consider a perfect RTU protection $\mathcal{P}$. Then $\mathcal{P}$ is a dominating set of the RTU level power network graph.

Proof: Assume that $\mathcal{P}$ is not a dominating set. Then there is an RTU k for which no neighboring RTU is in the protected set, $j \notin \mathcal{P}\ \forall j \in \mathcal{N}(k)$. A stealth attack can then be constructed based on the column corresponding to the state variable in bus k as done to obtain the upper bound on α_k. ■

A dominating set of the RTU level power network graph is not necessarily a perfect protection as we show it in the next section. Nevertheless, we can use the proposition to define an effective algorithm to find a perfect protection $\mathcal{P}$.

Dominating Set Augmentation Algorithm (DSA):
Initialize the set of protected measurements $\mathcal{P}$ with a minimal dominating set of the RTU level power network graph. Iterate over $k = 1,\dots,m$ and set $\mathcal{P} = \mathcal{P} \cup \{k\}$ if $\alpha_k < \infty$ for some k.

The algorithm terminates after one iteration and provides a perfect protection $\mathcal{P}$. For sparse power network graphs the budget required to achieve perfect protection is $\pi \ll n$.

B. Non-perfect protection

In practice the operator's budget π in terms of the number of measurements that can be protected might be insufficient for perfect protection. Then the operator would be interested in protecting a set of measurements $\mathcal{P}$ that maximizes its protection level according to some metric. We consider two possible metrics in this paper: the *maximal minimum attack cost* and the *maximal average attack cost*.

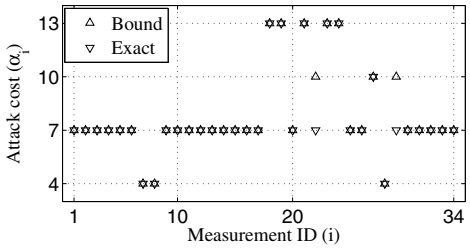

Fig. 2. The minimum attack costs α_k and their upper bounds $\hat{\alpha}_k$ for the IEEE 14 bus network. The bound is almost always tight. The shortest attacks involve the same measurements for the meter attacks and the RTU attacks.

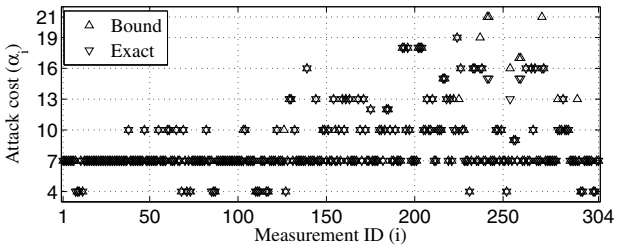

Fig. 3. The minimum attack costs α_k and their upper bounds $\hat{\alpha}_k$ for the IEEE 118 bus network. The bound is almost always tight. The shortest attacks involve the same measurements for the meter attacks and the RTU attacks.

1) Maximal minimum attack cost: According to the first criterion the goal of the operator is to maximize the minimum attack cost among all measurements that are possible to attack

$$\mathcal{P}^{MM} = \arg \max_{\mathcal{P}: C_{\mathcal{M}}(\mathcal{P}) \leq \pi} \min_k \alpha_k. \tag{10}$$

A simple greedy algorithm that aims to find an optimal set of protected measurements $\mathcal{P}$ in the sense of (10) for given budget π is the following.

Most Shortest Minimal Attacks Algorithm (MSM):
Initially set $\mathcal{P} = \emptyset$. Then in every iteration calculate $\alpha_k, \forall k \in 1, \ldots, m$ and $\min_k \alpha_k$. Pick partition block M_j that appears in most minimal attacks $A \in \mathcal{A}^*$ with least cost, i.e., $C_{\mathcal{M}}(A) = \min_k \alpha_k$. Set $\mathcal{P} = \mathcal{P} \cup M_j$. Continue until $C_{\mathcal{M}}(\mathcal{P}) = \pi$.

2) Maximal average minimum attack cost: According to the second criterion the goal of the operator is to maximize the average minimum attack cost of the measurements that are possible to attack

$$\mathcal{P}^{MA} = \arg \max_{\mathcal{P}: C_{\mathcal{M}}(\mathcal{P}) \leq \pi} \frac{1}{|\{k : \alpha_k \neq \infty\}|} \sum_{k: \alpha_k \neq \infty} \alpha_k \tag{11}$$

A simple greedy algorithm that aims to find an optimal set of protected measurements $\mathcal{P}$ in the sense of (11) for given budget π is the following.

Most Minimal Attacks Algorithm (MMA):
Initially set $\mathcal{P} = \emptyset$. Then in every iteration calculate $\alpha_k, \forall k \in 1, \ldots, m$. Pick a partition block M_j that appears in most minimal attacks $A \in \mathcal{A}^*$. Set $\mathcal{P} = \mathcal{P} \cup M_j$. Continue until $C_{\mathcal{M}}(\mathcal{P}) = \pi$.

V. Numerical results

In the following we present numerical results obtained using the proposed algorithms for the IEEE 14-bus and the IEEE 118-bus benchmark power networks. These networks were also analyzed in [4]. We added power flow measurements at each bus, and at every end of every interconnecting transmission line. For the IEEE 14 bus network there are $m = 54$ measurements, all assumed equally good ($R = I$), and the matrix H has size 54×14. For the IEEE 118 bus network there are $m = 490$ measurements, all assumed equally good ($R = I$), and the matrix H has size 490×118. These considered systems have more measurements than is normal in power systems, and should therefore have large measurement redundancy. For the computations, we used the Matlab package MatPower [10].

A. Minimum cost attack

We start the evaluation with considering the case of stealth meter attacks. Figs. 2 and 3 show the upper bound $\hat{\alpha}_k$ of the minimal attack length calculated using (7) and the exact minimal attack length α_k obtained using the iterative algorithm presented in Section III-C for the IEEE 14-bus and the IEEE 118-bus network, respectively. Except for a few meters, the bound $\hat{\alpha}_k$ is tight. Furthermore, most measurements can be attacked by modifying only 6 other measurements for both networks. A closer look at the attack vectors reveals that the meters that constitute the minimal attack belong to $\lfloor (\alpha_k - 1)/3 \rfloor$ RTUs, and hence the minimal RTU attacks are the same as the minimal meter attacks. ($\alpha_k - 1$ is not a multiple of 3 for some measurements in the IEEE 118 bus network, because there are parallel transmission lines between certain buses.)

B. Protection against stealth attacks

In order to understand the importance of the individual measurements in Fig. 4 we show the least minimum cost attack and the average minimum cost attack as a function of $\mathcal{P} = \{i\}$ for the case of meter attacks. The least minimum cost attack increases only when the protected measurements are the ones involved in the attack $A = \{7, 8, 18, 38\}$. The average minimum cost attack shows some variation depending on the protected measurement. In general, however, protecting a single meter does not provide significant improvement in terms of minimum attack costs. The same conclusion can be drawn from Fig. 5, which shows the the least minimum cost attack and the average minimum cost attack as a function of $\mathcal{P} = \{M_i\}$ for the case of RTU attacks on the IEEE 118-bus network.

Figure 6 shows the minimum attack cost and the average attack cost as a function of the protection budget π obtained using the MSM and the MMA algorithms. Using MMA the average minimum attack cost increases with the protection budget, but the least minimum attack cost is unchanged while $\pi \leq 8$. Using MSM the minimum attack cost increases faster than using MMA, but the average minimum attack cost is lower. For a budget of $\pi = n = 13$ both MMA and MSM find the set of meters that provides perfect protection. Hence, incremental protection of the meters does not lead to extra costs for the operator even if the ultimate goal is perfect protection.

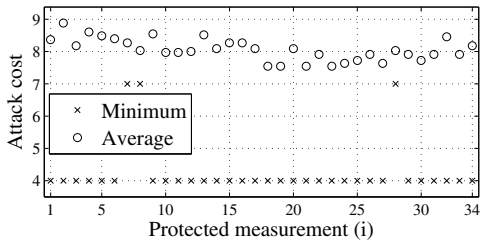

Fig. 4. The minimum attack costs $\min_{k\neq i}\alpha_k$ and the average minimum attack costs $\sum_{k\neq i}\alpha_k/(m-1)$ for the IEEE 14 bus network for the case of meter attacks.

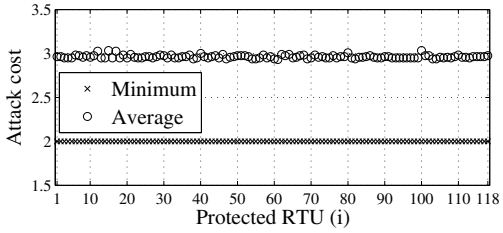

Fig. 5. The minimum attack costs $\min_{k\neq i}\alpha_k$ and the average minimum attack costs $\sum_{k\neq i}\alpha_k/n$ for the IEEE 118 bus network for the case of RTU attacks.

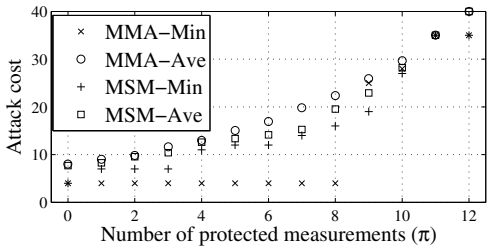

Fig. 6. The minimum attack cost $\min_k \alpha_k$ and the average minimum attack cost $\sum_{k:\alpha_k<\infty}\alpha_k/|\{k:\alpha_k<\infty\}|)$ for the IEEE 14 bus network for the case of meter attacks.

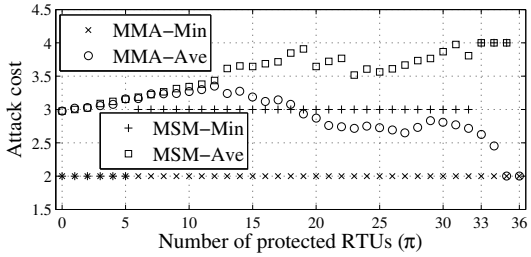

Fig. 7. The minimum attack cost $\min_k \alpha_k$ and the average minimum attack cost $\sum_{k:\alpha_k<\infty}\alpha_k/|\{k:\alpha_k<\infty\}|)$ for the IEEE 118 bus network for the case of RTU attacks.

For the case of RTU attacks perfect protection for the IEEE 14-bus network can be achieved by protecting 4 RTUs, $\mathcal{P}=\{2,6,7,9\}$. Both MMA and MSM find this set for a budget of $\pi=4$. This is a minimal dominating set of the power network graph. $\mathcal{P}=\{2,8,10,13\}$, which is also a minimal dominating set of the power network graph, does not provide perfect protection, as RTU 4 can be attacked by tampering with 3 RTUs $(4,7,9)$ and 8 measurements $(A=\{4,7,9,22,23,29,42,43\})$.

For the IEEE 118-bus network the minimal dominating set of the power network graph contains 32 RTUs, but it does not provide perfect protection. We used the DSA algorithm to find the RTUs $(5,23,69,77)$ that need to be added to the dominating set in order to achieve perfect protection. Figure 7 shows the results obtained using the MSM and MMA algorithms. MSM and MMA achieve perfect protection by protecting 36 and 37 RTUs respectively. We note two important differences compared to the case of measurement attacks (Fig. 6). First, the minimal attacks and the average attack length are rather small even close to perfect protection. Second, MSM outperforms MMA both in terms of minimal and average attack cost. This suggests that under the RTU attack cost model perfect protection is desirable if all measurements are equally important.

VI. Conclusion

In this paper, we considered the problem of finding and mitigating stealth attacks against the state estimator used in power networks. We described a security index that helps to locate power flows whose measurements are potentially easy to manipulate. We proposed an algorithm that can be used to find the least cost stealth attacks under a general model of attack and protection cost. We proposed three greedy algorithms to obtain perfect protection and partial protection against stealth attacks given a limited budget for protection. We used the proposed algorithms to evaluate the cost of attacking measurements on the IEEE 14-bus and the IEEE 118-bus power systems and showed how the incremental deployment of protected measurements can be best used to increase system security under two specific attack cost models.

References

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

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

[3] A. Giani, S. 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.

[4] 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.

[5] National Energy Technology Laboratoty, "Smart grid principal characteristics: Operates resiliently against attack and natural disasters," U.S. Department of Energy, Tech. Rep., September 2009.

[6] 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 2010*, 2010.

[7] H. Sandberg, A. Teixeira, and K. H. Johansson, "On security indices for state estimators in power networks," in *Preprints of the First Workshop on Secure Control Systems, CPSWEEK 2010*, 2010.

[8] 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.

[9] 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.

[10] 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.