

Malicious Data Attacks on the Smart Grid

Oliver Kosut, Liyan Jia, Robert J. Thomas, and Lang Tong

Abstract—Malicious attacks against power systems are investigated, in which an adversary controls a set of meters and is able to alter the measurements from those meters. Two regimes of attacks are considered. The strong attack regime is where the adversary attacks a sufficient number of meters so that the network state becomes unobservable by the control center. For attacks in this regime, the smallest set of attacked meters capable of causing network unobservability is characterized using a graph theoretic approach. By casting the problem as one of minimizing a supermodular graph functional, the problem of identifying the smallest set of vulnerable meters is shown to have polynomial complexity. For the weak attack regime where the adversary controls only a small number of meters, the problem is examined from a decision theoretic perspective for both the control center and the adversary. For the control center, a generalized likelihood ratio detector is proposed that incorporates historical data. For the adversary, the tradeoff between maximizing estimation error at the control center and minimizing detection probability of the launched attack is examined. An optimal attack based on minimum energy leakage is proposed.

Index Terms—Power system state estimation, false data attack, bad data detection, power network observability, smart grid security.

I. INTRODUCTION

Future smart grids will likely to be more tightly integrated with the cyber infrastructure for sensing, control, scheduling, dispatch, and billing. Already the current power grid relies on computer and communication networks to manage generation and facilitate communications between users and suppliers. While such integration is essential for a future “smart” grid, it also makes the power grid more vulnerable to cyber-attacks by adversaries around the globe. It has already been widely reported that the United State electrical grid has been penetrated by cyber spies [1].

We consider in this paper strategies of covert attack by adversaries on meters of the smart grid by injecting malicious data with the goal of biasing power system state estimation. If successful, such attacks may mislead the control center to take erroneous actions, or at the minimum, make the control center distrust the state estimate. Since some real-time markets use state estimation to determine location marginal prices (LMPs) [2], [3], malicious attacks also have impacts on real-time electricity markets.

Also considered in this paper are counter measures to malicious data attack at the control center in the form of attack

O. Kosut is with the Massachusetts Institute of Technology, Cambridge, MA {okosut}@mit.edu

L. Jia, R. J. Thomas, and L. Tong are with Cornell University, Ithaca, NY {lj92,rjt1,lt35}@cornell.edu

This work is supported in part by the National Science Foundation under CCF-0728872 and the NSF TRUST (The Team for Research in Ubiquitous Secure Technology) center under award CCF-0424422. Part of the work was presented at the IEEE SmartGridComm, Gettysburg MD, Oct. 2010, and UPEC 2010, Cardiff UK, August 2010.

detection. The problem of detecting malicious data attack can be viewed as a form of classical bad data detection. It is however important to note that, because the adversary can choose the site of attack judiciously and design attack data carefully, it is far more difficult to detect malicious data attacks than to detect random errors in the power systems. We will examine attacks with different degrees of sophistication.

The problem of malicious data attack on the power grid was first studied in [4], in which it was observed that there exist cooperative attacks on meters that all known bad data techniques will fail to detect. The authors of [4] gave a method to adjust measurements at a few meters in the grid in such a way that bad data detector will fail to perceive the corruption of the data and the estimate of network state can be perturbed arbitrarily in certain subspace.

We view the existence of these “unobservable” attacks as a fundamental limit on the detectability of malicious data attacks. Given that this fundamental limit depends on the number of meters that can be corrupted by the adversary, it is therefore natural to divide the attack into two different regimes. The *strong attack regime* is when the adversary is able to access a sufficient number of meters to launch an unobservable attack. Attacks in this regime cannot be detected by the control center (even if there is no measurement error). The *weak attack regime*, on the other hand, is when the adversary does not have access to a sufficient number of meters; their attacks can be detected, though imperfectly due to measurement errors.

The strong and weak regimes are divided by the smallest number χ^* of meters to which the adversary must have access in order to launch an unobservable attack. The quantity χ^* can be defined as a *security index* for the power network. As we will see, in a network with a more connected topology, or more redundant meters, the adversary must compromise a larger number of meters to perform this attack, and therefore the network is more secure. Thus quantifying χ^* is of theoretical and practical importance.

A. Summary of Results and Contributions

The main results in the study of malicious data attack in the strong attack regime are the characterization of the smallest number χ^* of adversarially-controlled meters such that an unobservable attack exists and obtaining this smallest size attack. In Sec. III, we show a connection between the classical notion of unobservability and the attack discovered in [4]; this justifies our use of the term *unobservable* to describe these attacks. This insight transforms the problem of finding the smallest unobservable attack into the problem of finding the smallest set of meters that, if removed, renders the network unobservable. We tackle this problem in a strictly graph theoretic manner, relying only on the network topology (line

diagram) without making use of specific network parameters. By exploiting the submodularity of certain graph functionals, we show that smallest size unobservable attack can be identified in polynomial time.

In the weak attack regime, the adversary has access to too few meters to perform an unobservable attack; therefore, it is possible to detect its presence. In Sec. IV, we investigate the problem in this regime, developing strategies to detect and localize malicious attacks. Under a classical decision theoretic detection formulation, we investigate the tradeoff between detection probability and false alarm. Because the adversary can choose where to attack the network and design the injected data, the problem of detecting malicious data cannot be formulated as a simple hypothesis test, and the uniformly most power test does not exist in general. We propose a detector based on the generalized likelihood ratio test (GLRT). The GLRT is not optimal in general, but it is known to perform well in practice and it has well established asymptotic optimality [5], [6], [7]. In other words, if the detector has many data samples, the detection performance of GLRT is close to optimal. The GLRT itself requires solving a combinatorial optimization problem. This makes it infeasible to use to detect a large number of corrupted meters. We therefore also study a detector using a convex regularization of the optimization problem, based on L_1 norm minimization. The convexity makes the optimization much easier to compute, but with potential sacrifices in performance. We provide numerical results for the true GLRT itself when it is feasible to use it, and the convex relaxation for larger scale problems.

We note that the proposed detector has a different structure from those used in conventional bad data detectors which usually employ a test on the state estimator residues errors [8], [9], [10]. The proposed the GLRT detector does not compute explicitly the residue error. We show, however, that when there is at most one attacked meter (a single attacked data), the GLRT is identical to the classical largest normalized residue (LNR) test using the residue error from the minimum mean square error (MMSE) state estimator.

Next we investigate malicious data attack from the perspective of an adversary who must make a tradeoff between inflicting the maximum damage and being detected by the control center. We define in Sec. V the notion of *Attacker Operating Characteristic* (AOC) that characterizes the tradeoff between the probability of being detected vs. resulting (extra) mean-square error at the state estimator. We consider the AOC to be dual to the classical *Receiver Operating Characteristic* (ROC), which characterizes the control center's tradeoff. We formulate the problem of optimal attack as minimizing the probability of being detected subject to causing the mean square error (MSE) to increase beyond a predetermined level. Unlike the strong attack regime, in which an unobservable attack is always the most damaging action for the adversary, in the weak attack regime, it is much less clear what the adversary should do. In particular, finding the attack with the optimal AOC is intractable. We present a heuristic that allows us to obtain attacks with minimum attack power leakage to the detector while increasing the mean square error at the state estimator beyond a predetermined objective. This heuristic

reduces to an eigenvalue problem that can be solved off line.

We also present a proof-of-concept analysis of the effect of a malicious data attack on the electricity market. In Sec. VI we describe how the locational marginal price (LMP) is calculated in the day-ahead and real-time power markets. In particular, the real-time price is determined based on the state estimator output, therefore it is vulnerable to malicious data attacks.

Finally, in Sec. VII we conduct numerical simulations on a small scale example using the IEEE 14-bus network. For the control center, we present simulation results that compare different detection schemes based on the *Receiver Operating Characteristics* (ROC) that characterize the tradeoff between the probability of attack detection vs. the probability of false alarm. We show that there is a substantial difference between the problem of detecting randomly appearing bad data from detecting malicious data injected by an adversary. Next we compare the GLRT detector with two classical detection schemes: the $J(\hat{\mathbf{x}})$ detector and the (Bayesian) largest normalized residue (LNR) detector [8], [9]. Our test shows improvement over the two well established detection schemes. From the adversary perspective, we compare the *Attacker Operating Characteristics* (AOC). Our result shows again that the GLRT detector gives higher probability of detection than that those of conventional detectors for the same amount MSE increase at the state estimator. We also provide simulation results on the electricity market, illustrating that even an attack in the weak attack regime can affect prices with a low probability of being detected.

B. Related Work

The study of malicious data attack is fairly recent. Liu, Ning, and Reiter was the first to address cyber-attack on power system state estimation in [4] where the authors obtained an algebraic condition for the existence of unobservable attacks. They also found that, for many standard networks, unobservable attack can be launched using only a limited number of meters.

In the strong attack regime, a fundamental problem is to characterize χ^* —the smallest number of meters required for unobservable attack [11], [12], [13]. Sandberg, Teixeira, and Johansson are the first to introduce the measure of the vulnerability of a network to malicious data attack [14] by defining a security index as the minimum number of meters to perform an unobservable attack including a given meter. Such an index is a function of the included meter, and finding such an index is difficult in general though [12] provided a lower bound; see [13] for a specific algorithm. The security index χ^* considered in [11] and in this paper does not impose restrictions on which meter to include in the set of attacking meters and is a function of the network topology only. The graph theoretic approach to finding security index has its root in the classical work of Clements, Krumpholz and Davis [15] who established the relation of network observability and the graph theoretic notion of spanning tree. It is this relation that allows us to formulate an optimization involving a submodular function.

Introducing redundant and, more importantly, trustworthy measurements is the key to defending malicious data attack.

The use of PMUs, for example, will in general reduce the vulnerability of the network, provided that PMU measurements themselves are secure. The authors of [16] found that, in order to protect a network against all unobservable attacks, a minimum size set of measurements that by themselves ensure observability need to be included. This result is corroborated and enhanced by the graph theoretic techniques presented in this paper. We note that the use of PMUs does not fundamentally change the formulation of the problem.

In the weak attack regime, the problem of detecting malicious attack was first considered in the precursors of the current paper in [17], [18], [11]. There is a natural connection between the problem studied here and the classical bad data detection as part of the original formulation of state estimation [8]. See [19] for an earlier comparison study. Malicious data attack can be viewed as the *worst interacting bad data* injected by an adversary. To this end, very little is known about the worst case scenario although the detection of interacting bad data has been considered [9], [20], [21], [22].

Finally, the impact of attack on real-time market is studied in [23], [24]. The influence of false data attacks on electricity markets is studied in [24]. A method is given to find attacks that influence LMPs at certain buses, which could be employed by a malicious intruder to turn a profit. In this paper, we focus on the effect on the LMP prices through the attacks on state estimation. See Sec. VI.

II. THE NETWORK AND ATTACK MODELS

We adopt a graph-theoretic model for the power system with an undirected graph (V, E) , where V represents the set of buses, and E is the set of transmission lines. Each line connects two meters, so each element $e \in E$ is an unordered pair of buses in V . Fig 1 shows the graph structure of the IEEE 14-bus test system, which we use in our simulations. The control center receives measurements from various meters deployed throughout the system, from which it performs state estimation. Meters come in two varieties: transmission line flow meters, which measure the power flow through a single transmission line, and bus injection meters, which measure the total outgoing flow on all transmission lines connected to a single bus. Therefore each meter is associated with either a bus in V or a line in E . We allow for the possibility of multiple meters on the same bus or line. Indeed, in our simulations, we assume that a meter is placed in every bus, and two meters on every line, one in each direction.

We assume a standard DC power flow model from a linearized version of the AC power flow model. In the absence of attack, the control center obtains meter measurements

$$\mathbf{z} = \mathbf{H}\mathbf{x} + \mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_e), \quad (1)$$

where $\mathbf{z} \in \mathbb{R}^m$ is the vector of power flow measurements, $\mathbf{x} \in \mathbb{R}^n$ is the system state, and \mathbf{e} is the Gaussian measurement noise with zero mean and diagonal covariance matrix $\mathbf{\Sigma}_e$. Note we are assuming the measurement noise to be zero-mean. If this were not the case, as long as the mean is known, it can simply be subtracted off with no impact on our results, but we make this assumption for convenience.

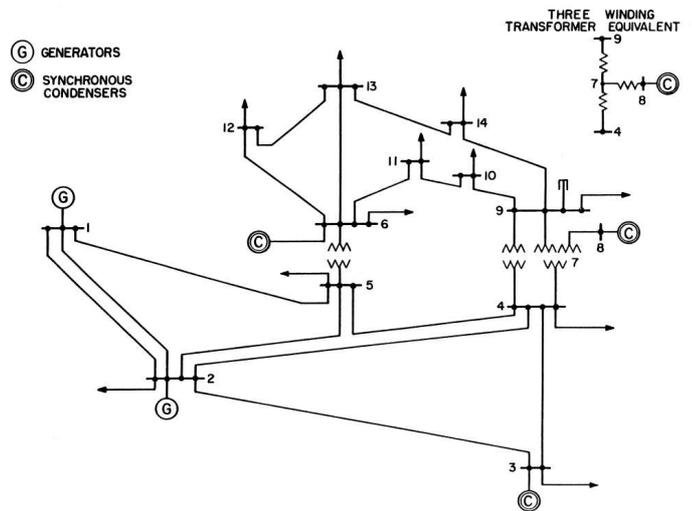


Fig. 1. IEEE 14 bus test system.

The measurement matrix \mathbf{H} depends on the topology of the network, the susceptance of each transmission line, and the placement of the meters. The matrix \mathbf{H} is generated from this information as follows. Suppose $(i, j) \in E$; that is, buses i and j are connected by a transmission line. The DC power flow through this line from bus i to j is $B_{ij}(x_i - x_j)$, where B_{ij} is the susceptance of the line (i, j) . We may also write this power flow as $h_{ij}\mathbf{x}$, where

$$h_{ij} = \begin{bmatrix} 0 & \cdots & 0 & \underbrace{B_{ij}}_{i\text{th element}} & 0 & \cdots & 0 & \underbrace{-B_{ij}}_{j\text{th element}} & 0 & \cdots & 0 \end{bmatrix}. \quad (2)$$

Therefore, if a meter measures the flow through line (i, j) , the associated row of \mathbf{H} is given by $h_{(i,j)}$. A bus injection meter measures the total power flow on all lines incident to a particular node. Therefore the row of \mathbf{H} associated with a meter on bus i is given by

$$\sum_{j:(i,j) \in E} h_{ij}. \quad (3)$$

In the presence of malicious data attack, the data collected at the control center satisfy

$$\mathbf{z} = \mathbf{H}\mathbf{x} + \mathbf{a} + \mathbf{e} \quad (4)$$

where vector \mathbf{a} is malicious data injected by an adversary. The injected vector \mathbf{a} is sparse with respect to the meters controlled by the adversary. That is, $a_i \neq 0$ only if the i th meter is controlled by the adversary. In general, we assume that the adversary may control any set of up to k meters. Therefore we impose the constraint $\mathbf{a} \in \mathcal{A}_k = \{\mathbf{a} : \|\mathbf{a}\|_0 = k\}$, which is the set of sparse vectors with at most k nonzero entries.

We assume that the adversary has access to network parameters \mathbf{H} and is able to coordinate attacks from different meters. These assumptions, and that the adversary may choose any set of k meters it likes, give the adversary more power than perhaps possible in practice, which is a well adopted practice when analyzing security. Thus the results we obtain are in general conservative.

III. THE STRONG ATTACK REGIME

We consider in this section the case when the adversary is able to launch an attack from a sufficiently large number of well chosen meters. Attacks in the so-called strong attack regime are defined by the attack given in [4]. In particular, an attack vector \mathbf{a} belongs to the strong attack regime if $\mathbf{a} = \mathbf{H}\mathbf{c}$ for some \mathbf{c} . For such an attack, we have

$$\mathbf{z} = \mathbf{H}\mathbf{x} + \mathbf{a} + \mathbf{e} = \mathbf{H}(\mathbf{x} + \mathbf{c}) + \mathbf{e}.$$

Therefore \mathbf{x} is indistinguishable from $\mathbf{x} + \mathbf{c}$. If both \mathbf{x} and $\mathbf{x} + \mathbf{c}$ are valid network states, the adversary's injection of data \mathbf{a} when the true state is \mathbf{x} will lead the control center to believe that the true network state is $\mathbf{x} + \mathbf{c}$. Because vector \mathbf{c} can be scaled arbitrarily, the adversary can perturb the network state (in the view of control center) arbitrarily.

Since no detector can distinguish \mathbf{x} from $\mathbf{x} + \mathbf{c}$, we call hereafter an attack vector \mathbf{a} *unobservable* if it has the form $\mathbf{a} = \mathbf{H}\mathbf{c}$. Note that it is unlikely that random bad data \mathbf{a} will satisfy $\mathbf{a} = \mathbf{H}\mathbf{c}$. An adversary, on the other hand, can synthesize its attack vector to satisfy the unobservable condition.

The attacks in the strong attack regime is defined by the algebraic condition that a k -sparse vector satisfying $\mathbf{a} = \mathbf{H}\mathbf{c}$ for some vector \mathbf{c} . Given a network and the corresponding factor matrix \mathbf{H} , we define χ^* as the smallest k such that $\mathbf{a} = \mathbf{H}\mathbf{c} \in \mathcal{A}_k$ for some \mathbf{c} . Equivalent, χ^* is given by

$$\chi^* = \min_{\substack{\mathbf{a}: \mathbf{a} = \mathbf{H}\mathbf{c} \\ \text{for some } \mathbf{c}}} \|\mathbf{a}\|_0 \quad (5)$$

where $\|\mathbf{a}\|_0$ is the number of non-zero entries in \mathbf{a} .

In this section, we characterize χ^* and find the corresponding attack \mathbf{a} . It is interesting to note that, although the definition given in (5) seems to suggest that χ^* depends on \mathbf{H} , our result shows that χ^* depends only on the topology of the network, not on specific values of the matrix \mathbf{H} .

A. Unobservable Attacks and Network Observability

We establish first a connection between the unobservable attack and the classical notion of network observability [25]. The following theorem shows that the algebraic condition that defines the unobservable attack is equivalent to the classical network unobservability condition [25].

Theorem 1: A k -sparse attack vector \mathbf{a} comprises an unobservable attack if and only if the network becomes unobservable when the k meters associated with the nonzero entries of \mathbf{a} are removed from the network; that is, the $(m - k) \times n$ submatrix of \mathbf{H} taken from the rows of \mathbf{H} corresponding to the zero entries of \mathbf{a} does not have full column rank.

Proof: Without loss of generality, let \mathbf{H} be partitioned into $\mathbf{H}^T = [\mathbf{H}_1^T \mid \mathbf{H}_2^T]$, and submatrix \mathbf{H}_1 does not have full column rank, i.e. there exists a vector $\mathbf{c} \neq \mathbf{0}$ such that $\mathbf{H}_1\mathbf{c} = \mathbf{0}$. We now have $\mathbf{a} = \mathbf{H}\mathbf{c} \in \mathcal{A}_k$, which is unobservable by definition. Conversely, consider an unobservable $\mathbf{a} = \mathbf{H}\mathbf{c} \in \mathcal{A}_k$. Without loss of generality, we can assume that the first $m - k$ entries of \mathbf{a} are zero. We therefore have $\mathbf{H}_1\mathbf{c} = \mathbf{0}$ where \mathbf{H}_1 is the submatrix made of the first $m - k$ rows of \mathbf{H} . ■

The implication from the above theorem is that the attack discovered in [4] is equivalent to removing k meters from the

network, thus making the network not observable. With this perspective, it is in some sense not surprising that such an attack exists: if an adversary controls a set of meters whose absence makes the system unobservable, then there must be some aspect of the system state that can only be learned through those meters. Thus, the adversary has complete control over what the control center learns about this aspect of the state.

We further note that even though an unobservable attack is equivalent to the network being made unobservable, the adversarial attack is still much more destructive. When the network is unobservable because there are insufficient meters, the control center knows; it knows exactly what aspects about the system state it can gather information about, and which it cannot. However, in the case of an unobservable adversarial attack, the control center does not know it is under attack, nor which of several possible attacks is being executed. Therefore the situation is much more precarious, because the control center does not even know what it does not know.

B. Unobservable Attacks on AC Power Flow

Our primary focus on this paper is on the DC power flow model, but we briefly note here that because of the connection made by Theorem 1 between unobservable attacks and classical unobservability, unobservable attacks can be constructed even under the more realistic, and nonlinear, AC power flow model. In particular, observability is the same for both models, so if the adversary controls enough meters the absence of which makes the system unobservable, it can similarly manipulate the control center's estimate of the unobservable part of the state, even under the nonlinear model.

To be more precise, we may write the AC power flow model as follows. The state of each bus i is given by its voltage magnitude $|V_i|$ and its phase θ_i . The state \mathbf{x} is the vector composed of all voltages and all phases. A meter sitting at bus i and measuring the flow through line $(i, j) \in E$ measures the real and reactive power flows through this line. We denote this $h_{ij}(\mathbf{x})$, which can be written

$$h_{ij}(\mathbf{x}) = \begin{bmatrix} P_{ij} \\ Q_{ij} \end{bmatrix} = \begin{bmatrix} |V_i|(|V_j|G_{ij} \cos \theta_{ij} + |V_j|B_{ij} \sin \theta_{ij} - |V_i|G_{ij}) \\ |V_i|(|V_j|G_{ij} \sin \theta_{ij} - |V_j|B_{ij} \cos \theta_{ij} + |V_i|B_{ij}) \end{bmatrix} \quad (6)$$

where $\theta_{ij} = \theta_i - \theta_j$, and B_{ij} and G_{ij} are the susceptance and conductance of line (i, j) respectively. Note that (6) is the AC equivalent of (2). A meter measuring the total power flow at bus i measures

$$\sum_{j:(i,j) \in E} h_{ij}(\mathbf{x}). \quad (7)$$

Again, (7) is the AC equivalent of (3). Let $\mathbf{h}(\mathbf{x})$ be the complete vector of measurements given the state \mathbf{x} . For a set of meters \mathcal{A} , let $\mathbf{h}_{\mathcal{A}}(\mathbf{x})$ be the subvector of measurements taken by just these meters.

Even though the measurement equations for AC power flow are nonlinear, it is clear from (6) and (7) that the phases only enter as $\theta_i - \theta_j$ for $(i, j) \in E$. Suppose the true state were

\mathbf{x} , and the adversary constructs a vector \mathbf{x}' which changes only the phases, and moreover has the property that the phase differences $\theta_i - \theta_j$ do not differ between \mathbf{x}' and \mathbf{x} except for lines (i, j) only observed by meters controlled by the adversary. Thus, if the set of adversarial meters is \mathcal{S} , $\mathbf{h}_{\mathcal{S}}(\mathbf{x}) = \mathbf{h}_{\mathcal{S}^c}(\mathbf{x}')$, where \mathcal{S}^c is the set of honest meters. That is, the non-adversarial meters cannot distinguish \mathbf{x} from \mathbf{x}' , so if the adversary replaces $\mathbf{h}_{\mathcal{S}}(\mathbf{x})$ with $\mathbf{h}_{\mathcal{S}}(\mathbf{x}')$, the control center will mistake \mathbf{x} for \mathbf{x}' . Because the phase differences enter linearly into the measurement equations, these phase alterations can be made arbitrarily large without detection by the control center. In this sense, the attack behaves similarly to the unobservable attack for the DC model. Moreover, because the phases enter linearly in the measurement equations in exactly the same way that the states enter into the linear measurement equations for the DC model (2), (3), this attack exists exactly when an unobservable attack on the DC model exists.

Observe that the attack described to manipulate the phase estimates relies on the adversary being able to calculate $\mathbf{h}_{\mathcal{S}}(\mathbf{x}')$. Depending on the adversary's knowledge of the system state, this may be much harder to do than to calculate the bad data for an unobservable attack in the DC model. Recall that in the DC model, the added vector \mathbf{a} depends only on network characteristics, and not on the current state. To perform the attack, the adversary need only add this pre-computed vector to the true measurements. Because of the nonlinearity of the function \mathbf{h} , to perform the attack on the AC model the adversary would need to know the exact network state, even at buses where it controls no meters. However, since the DC model is a linearization of the AC model, small versions of the DC attack would work approximately on the AC model. It may be possible for the control center to employ the nonlinearity of the true system in order to better detect linear attacks, but this is beyond the scope of this paper.

C. Graph-Theoretic Characterization of Minimum Size Unobservable Attacks

Given a line diagram of a power network and the locations of meters, the vulnerability of the network can be characterized by the smallest number of meters that the network can be made unobservable by an adversary. From Theorem 1, we know that there is an unobservable k -sparse attack vector \mathbf{a} if and only if it is possible to remove k rows from \mathbf{H} and cause \mathbf{H} not to have full column rank. This algebraic condition, however, is difficult to use to determine the minimum size of an unobservable attack. We develop here a graph theoretic approach by exploiting the extra structure on \mathbf{H} imposed by the network topology.

The following theorem exactly characterizes the security index χ^* , the smallest number of meters to make the system unobservable. The proof builds on the result of [15], which gave an efficient method to determine the observability of a network based only on its topology. In the sequel, denote $|\mathcal{X}|$ as the number of elements in a set \mathcal{X} .

Theorem 2 (Characterization of security index χ^):* For a set of lines $\mathcal{A} \subseteq E$, let $g(\mathcal{A})$ be the set of meters either on lines in \mathcal{A} or on buses adjacent to lines in \mathcal{A} . Let $h(\mathcal{A})$ be

the number of connected components in the graph $(\mathcal{V}, E \setminus \mathcal{A})$; i.e., the original graph after all lines in \mathcal{A} have been removed. The security index χ^* is given by

$$\chi^* = \min_{\mathcal{A} \subseteq E} |g(\mathcal{A})| - h(\mathcal{A}) + 2. \quad (8)$$

Moreover, this quantity may be calculated in polynomial time in the size of the network.

We prove Theorem 2 with Lemma 1 and Lemma 2 below. The first shows that it is possible to remove a set of meters of size given in (8) and make the system unobservable. The second allows us to conclude that it is impossible to remove fewer meters than the quantity given in (8) to make the system unobservable.

Lemma 1: For all $\mathcal{A} \subseteq E$, removing an arbitrary subset of $g(\mathcal{A})$ of size $|g(\mathcal{A})| - h(\mathcal{A}) + 2$ makes the system unobservable.

Proof: Let $\bar{\mathcal{V}}$ and \bar{E} be the sets of buses and lines respectively with a meter placed on them. Theorem 5 in [15] states that the power system given by $(\mathcal{V}, E, \bar{\mathcal{V}}, \bar{E})$ is observable if and only if there exists a $\mathcal{F} \subseteq E$ comprising a spanning tree of \mathcal{V} and an assignment function

$$\phi : \mathcal{F} \rightarrow \bar{\mathcal{V}} \cup \bar{E} \quad (9)$$

satisfying:

- 1) If $l \in \bar{E}$, then $\phi(l) = l$.
- 2) If $\phi(l) \in \bar{\mathcal{V}}$, then line l is incident to the bus $\phi(l)$.
- 3) If $l_1, l_2 \in \mathcal{F}$ are distinct, then $\phi(l_1) \neq \phi(l_2)$.

The principle behind this theorem is that a bus injection meter may “impersonate” a single line meter on a line incident to the bus. If a bus $b = \phi(l)$ for some line l , this represents the meter at b impersonating a meter on line l . The system is observable if and only if a spanning tree \mathcal{F} exists made up of transmission lines with either real meters or impersonated meters by bus meters.

Not including the lines in \mathcal{A} , the network splits into $h(\mathcal{A})$ separate pieces. Therefore, any spanning tree \mathcal{F} must include at least $h(\mathcal{A}) - 1$ lines in \mathcal{A} . Any assignment ϕ satisfying the conditions above must therefore employ at least $h(\mathcal{A}) - 1$ meters in $g(\mathcal{A})$. Hence, if any $|g(\mathcal{A})| - h(\mathcal{A}) + 2$ of these meters are removed from the network, only $h(\mathcal{A}) - 2$ remain, which is not enough to create a full spanning tree, so the network becomes unobservable. \square

Example 1: Consider the IEEE 14-bus test system, shown in Fig. 1. Take $\mathcal{A} = \{(7, 8)\}$. Since bus 8 is only connected to the system through bus 7, removing this line from the network cuts it into two pieces. Therefore $h(\mathcal{A}) = 2$. The set of meters $g(\mathcal{A})$ consists of meters on the line (7, 8), and bus injection meters at bus 7 and 8. Theorem 1 states that if we remove $|g(\mathcal{A})|$ meters from this set—that is, all the meters in $g(\mathcal{A})$ —the system becomes unobservable. In our simulation examples, we assume there are two meters on each line, therefore it takes 4 meters to execute an unobservable attack. Furthermore, it is not hard to employ Theorem 1 to find similar 4-sparse unobservable attacks on the 30-bus, 118-bus, and 300-bus test systems.

Lemma 1 shows that the quantity on the right hand side of (8) is an upper bound on χ^* . We now show that it is a lower

Input: A line l^* for which adding a meter makes the system observable, as well as a spanning tree \mathcal{F} and function ϕ certifying this.

Output: A set $\mathcal{A} \subseteq E$ for which $|g(\mathcal{A})| - h(\mathcal{A}) + 2 \leq 0$.

```

1:  $\mathcal{B} := \{l^*\};$ 
2:  $m := 1;$ 
3:  $\mathcal{W} := \emptyset;$ 
4: Let  $\mathcal{V}_1, \mathcal{V}_2$  be split on  $\mathcal{B}$  in  $\mathcal{F}$ ;
5: Let  $\mathcal{A}$  be the set of lines joining  $\mathcal{V}_1$  to  $\mathcal{V}_2$ ;
6: while  $g(\mathcal{A}) \setminus \mathcal{W} \neq \emptyset$  do
7:   Let  $v$  be an element of  $g(\mathcal{A}) \setminus \mathcal{W}$ ;
8:    $\mathcal{W} := \mathcal{W} \cup \{v\};$ 
9:   Let  $l'$  be the line in  $\mathcal{F}$  for which  $\phi(l') = v$ ; If there is
   no such line, terminate the loop;
10:   $\mathcal{B} := \mathcal{B} \cup \{l'\};$ 
11:   $m := m + 1;$ 
12:  Let  $\mathcal{V}_1, \dots, \mathcal{V}_{m+1}$  be a partition of  $\mathcal{V}$  given by the
   connected components after  $\mathcal{B}$  is removed from  $\mathcal{F}$ ;
13:  Let  $\mathcal{A}$  be the set of lines joining  $\mathcal{V}_i$  to  $\mathcal{V}_j$  for any  $i, j$ .
14: end while

```

Fig. 2. Algorithm to find a set \mathcal{A} satisfying the conditions of Lemma 2.

bound as well. If it were not, there would exist a set of meters of size M that if removed, make the system unobservable, where $L < |g(\mathcal{A})| - h(\mathcal{A}) + 2$ for each \mathcal{A} . Consider the system after these M meters have been removed. Let $g'(\mathcal{A})$ be the value of g after removal of these meters. Note that h does not depend on the meters in the network so it does not change. Since L meters have been removed in total, $|g'(\mathcal{A})| \geq |g(\mathcal{A})| - M$. Hence

$$|g'(\mathcal{A})| - h(\mathcal{A}) + 2 \geq |g(\mathcal{A})| - h(\mathcal{A}) + 2 - M > 0. \quad (10)$$

This means that for the system after the meters are removed, the quantity in (8) is strictly positive, but the system is unobservable. The following lemma states that this is impossible, allowing us to conclude that there is a contradiction, and so (8) holds with equality.

Lemma 2: If the network is unobservable, then there exists a set of meters \mathcal{A} for which $|g(\mathcal{A})| - h(\mathcal{A}) + 2 \leq 0$.

Proof: Assume without loss of generality that adding a meter to a single line can recover observability. Let l^* be such a line. Again using the equivalent condition to observability proved in [15], there must exist a spanning tree \mathcal{F} and a function ϕ satisfying properties (1)–(3) listed above, when a meter is included on line l^* . It must be that $l^* \in \mathcal{F}$, because if not the network would be observable to begin with. Since \mathcal{F} is a spanning tree, removing any line from it splits the network into two connected components. In particular, let \mathcal{V}_1 and \mathcal{V}_2 be the sets of buses in the two connected components when l^* is removed from \mathcal{F} . Let \mathcal{A}_1 be the set of lines between nodes in \mathcal{V}_1 and nodes in \mathcal{V}_2 . Certainly then $h(\mathcal{A}_1) = 2$. Thus if $g(\mathcal{A}_1) = \emptyset$, then $|g(\mathcal{A}_1)| - h(\mathcal{A}_1) + 2 = 0$, so we are done. Otherwise, we claim that the algorithm given in Fig. 2 outputs a set \mathcal{A} for which $|g(\mathcal{A})| - h(\mathcal{A}) + 2 \leq 0$.

To prove correctness of this algorithm, we show the following loop invariants:

- 1) Every bus $v \in \mathcal{W}$ is incident to two lines each connected to a different \mathcal{V}_i set.
- 2) For every line $l \in \mathcal{A}$, if a meter were added to l , the system would become observable.
- 3) There are no line meters in the original network (without a meter on l^*) on \mathcal{A} .
- 4) For each $v \in g(\mathcal{A}) \setminus \mathcal{W}$, there exists a line l' in \mathcal{F} for which $\phi(l') = v$.

Proof at entrance to the loop: Statement (1) holds trivially, as \mathcal{W} is empty. At entrance to the loop, the only lines in \mathcal{A} are those connecting \mathcal{V}_1 to \mathcal{V}_2 . If there were a meter on a line connecting \mathcal{V}_1 to \mathcal{V}_2 , the network would be observable. This proves statement (2). Similarly, the network would be observable if there were a meter on any line connecting \mathcal{V}_1 to \mathcal{V}_2 , which proves statement (3). Since $\mathcal{W} = \emptyset$, $g(\mathcal{A}) \setminus \mathcal{W} = g(\mathcal{A})$. If the meter at some $v \in g(\mathcal{A})$ were assigned to no line in \mathcal{F} , then it could be assigned to its incident line in \mathcal{A} , in which case the network would be observable. This proves statement (4).

Induction step: We assume that statements (1)–(4) hold at the start of an iteration of the loop, and show that they hold again at the end.

When v is chosen in line 7, it is an element of $g(\mathcal{A})$. Therefore v must be incident to some line $l \in \mathcal{A}$. Moreover, since statement (4) holds at the end of the previous iteration, there exists an $l' \in \mathcal{F}$ with $\phi(l') = v$, which must also be incident to v . This l' is added to \mathcal{B} in line 10. Note that because only a single element is added to \mathcal{B} , when the \mathcal{V}_i are recalculated in line 12, one set splits into two; the rest stay as they were previously. Therefore the \mathcal{V}_i set to which l connects remains the same, where l' connects to a newly created \mathcal{V}_i set. Therefore statement (1) holds at the end of the loop iteration.

By statement (3), there are no line meters on \mathcal{A} , so in principle we may add a meter to any line in \mathcal{A} . Take any $l \in \mathcal{A}$ and add a meter to it. Line l connects two of the \mathcal{V}_i sets, which were previously connected by some line in \mathcal{F} . We alter \mathcal{F} and ϕ in the following way, to include l and restore their necessary properties. The bus meter that had previously been assigned to a line bridging the gulf between the same two \mathcal{V}_i sets bridged by l can instead be assigned to the other incident line to a different \mathcal{V}_i whose existence was proved in statement (1). This frees up a different bus meter, which then may be assigned to a line bridging another gulf, and so on until reaching a bus meter bridging the gulf originally bridged by l^* . This connects all the \mathcal{V}_i sets, and therefore the network would be observable. This proves statement (2).

Consider the lines newly added to \mathcal{A} in line 13. These lines bridge the same gulf as l' . If there were a meter on any of these lines, then the bus on meter v need not be assigned to l' . By a similar argument as above, it could be reassigned all the way back to l^* , so it may be removed and the network would still be observable. This proves statement (3).

Suppose there were a bus $v \in g(\mathcal{A}) \setminus \mathcal{W}$ for which there were no line l' satisfying $\phi(l') = v$. Bus v must be incident to a line $l \in \mathcal{A}$. By statement (3), l does not have a meter, so the meter on bus v could be used to simulate a meter on l . This effectively adds a meter to l ; but by statement (2), adding a meter to a line in \mathcal{A} brings observability. Since the

network is unobservable, there can be no such bus v . This proves statement (4).

We have established all four loop invariants. In particular, statement (4) implies that the loop never terminates prematurely at line 9, so when the loop concludes $g(\mathcal{A}) = \mathcal{W}$. Note that after every iteration of the loop, $|\mathcal{W}| = m - 1$ and $h(\mathcal{A}) = m + 1$. Therefore

$$|g(\mathcal{A})| - h(\mathcal{A}) + 2 = |\mathcal{W}| - h(\mathcal{A}) + 2 \quad (11)$$

$$= (m - 1) - (m + 1) + 2 = 0. \quad (12)$$

□

All that remains to prove Theorem 2 is to show that the quantity on the right hand side of (8) can be calculated in polynomial time. We make use of the theory of submodular functions. A submodular function is a real-valued function f defined on the collection of subsets of a set W such that for any $A, B \subseteq W$,

$$f(A \cup B) + f(A \cap B) \leq f(A) + f(B). \quad (13)$$

Moreover, a function f is supermodular if $-f$ is submodular. There are several known techniques to find the set $A \subseteq W$ minimizing $f(A)$ in time polynomial in the size of W [26], [27], [28]. It is not hard to see that $|g(\mathcal{A})|$ is submodular in \mathcal{A} , and $h(\mathcal{A})$ is supermodular. Therefore, their difference is submodular, so it can be efficiently minimized. This concludes the proof of Theorem 2.

IV. THE WEAK ATTACK REGIME

In this section, we study the problem in the weak attack regime, where the adversary cannot or does not perform an unobservable attack as described in Sec. III. In this regime, it is sometimes possible to detect the adversary's presence, and so we study detectors to do so.

In our analysis of the weak attack regime, we use mean square error as our metric for the the amount of damage done by an attack. In particular, we will seek strategies by the control center that minimize the mean square error for arbitrary attacks by the adversary. Mean square error is a very generic measure that gives a sense of how far the control center's estimate is from the truth. However, for certain specific uses of the state estimate, it may be not give a precise picture of the impact of the adversary. For example, in Sec. VI, we study adversarial attacks on electricity markets, and we use revenue change as our measure of adversary impact, which is more relevant to that application of the state estimate. Here, we want a measure that is more broadly applicable, if possibly less precise, so we adopt mean square error.

A. A Bayesian Framework and MMSE Estimation

We consider in this paper a Bayesian framework where the state variables are random vectors with Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_x, \boldsymbol{\Sigma}_x)$. We assume that, in practice, the mean $\boldsymbol{\mu}_x$ and covariance $\boldsymbol{\Sigma}_x$ can be estimated from historical data. By subtracting the mean from the data, we can assume without loss of generality that $\boldsymbol{\mu}_x = \mathbf{0}$.

In the absence of an attack, i.e. $\mathbf{a} = \mathbf{0}$ in (4), (\mathbf{z}, \mathbf{x}) are jointly Gaussian. The minimum mean square error (MMSE) estimator of the state vector \mathbf{x} is a linear estimator given by

$$\hat{\mathbf{x}}(\mathbf{z}) = \underset{\hat{\mathbf{x}}}{\operatorname{argmin}} \mathbb{E}(\|\mathbf{x} - \hat{\mathbf{x}}(\mathbf{z})\|^2) = \mathbf{K}\mathbf{z} \quad (14)$$

where

$$\mathbf{K} = \boldsymbol{\Sigma}_x \mathbf{H}^T (\mathbf{H} \boldsymbol{\Sigma}_x \mathbf{H}^T + \boldsymbol{\Sigma}_e)^{-1}. \quad (15)$$

The minimum mean square error, in the absence of attack, is given by

$$\mathcal{E}_0 = \min_{\hat{\mathbf{x}}} \mathbb{E}(\|\mathbf{x} - \hat{\mathbf{x}}(\mathbf{z})\|^2) = \operatorname{Tr}(\boldsymbol{\Sigma}_x - \mathbf{K}\mathbf{H}\boldsymbol{\Sigma}_x).$$

If an adversary injects malicious data $\mathbf{a} \in \mathcal{A}_k$ but the control center is unaware of it, then the state estimator defined in (14) is no longer the true MMSE estimator (in the presence of attack); the estimator $\hat{\mathbf{x}} = \mathbf{K}\mathbf{z}$ is a ‘‘naive’’ MMSE estimator that ignores the possibility of attack, and it will incur a higher mean square error (MSE). In particular, it is not hard to see that the MSE in the presence of \mathbf{a} is given by

$$\mathcal{E}_0 + \|\mathbf{K}\mathbf{a}\|_2^2. \quad (16)$$

The impact on the estimator from a particular attack \mathbf{a} is given by the second term in (16). To increase the MSE at the state estimator, the adversary necessarily has to increase the ‘‘energy’’ of attack, which increases the probability of being detected at the control center.

B. Statistical Model and Attack Hypotheses

We now present a formulation of the detection problem at the control center. We assume a Bayesian model where the state variables are random with a multivariate Gaussian distribution $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_x)$. Our detection model, on the other hand, is not Bayesian in the sense that we do not assume any prior probability of the attack nor do we assume any statistical model for the attack vector \mathbf{a} .

Under the observation model (4), we consider the following composite binary hypothesis:

$$\mathcal{H}_0 : \mathbf{a} = \mathbf{0} \quad \text{versus} \quad \mathcal{H}_1 : \mathbf{a} \in \mathcal{A}_k \setminus \{\mathbf{0}\}. \quad (17)$$

Given observation $\mathbf{z} \in \mathbb{R}^m$, we wish to design a detector $\delta : \mathbb{R}^m \rightarrow \{0, 1\}$ with $\delta(\mathbf{z}) = 1$ indicating a detection of attack (\mathcal{H}_1) and $\delta(\mathbf{z}) = 0$ the null hypothesis.

An alternative formulation, one we will not pursue here, is based on the extra MSE $\|\mathbf{K}\mathbf{a}\|_2^2$ at the state estimator. See (16). In particular, we may want to distinguish, for $\|\mathbf{a}\|_0 \leq k$,

$$\mathcal{H}'_0 : \|\mathbf{K}\mathbf{a}\|_2^2 \leq C, \quad \text{versus} \quad \mathcal{H}'_1 : \|\mathbf{K}\mathbf{a}\|_2^2 > C. \quad (18)$$

Here both null and alternative hypotheses are composite and the problem is more complicated. The operational interpretation, however, is significant because one may not care in practice about small attacks that only marginally increase the MSE of the state estimator.

C. Generalized Likelihood Ratio Detector with L_1 Norm Regularization

For the hypotheses test given in (17), the uniformly most powerful test does not exist. We propose a detector based on the generalized likelihood ratio test (GLRT). We note in particular that, if we have multiple measurements under the same \mathbf{a} , the GLRT proposed here is asymptotically optimal in the sense that it offers the fastest decay rate of miss detection probability [29].

The distribution of the measurement \mathbf{z} under the two hypotheses differ only in their means:

$$\begin{aligned} \mathcal{H}_0 &: \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \Sigma_z) \\ \mathcal{H}_1 &: \mathbf{z} \sim \mathcal{N}(\mathbf{a}, \Sigma_z), \mathbf{a} \in \mathcal{A}_k \setminus \{\mathbf{0}\} \end{aligned}$$

where $\Sigma_z \triangleq \mathbf{H}\Sigma_x\mathbf{H}^T + \Sigma_e$. The GLRT is given by

$$L(\mathbf{z}) \triangleq \frac{\max_{\mathbf{a} \in \mathcal{A}_k} f(\mathbf{z}|\mathbf{a})}{f(\mathbf{z}|\mathbf{a}=\mathbf{0})} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \tau, \quad (19)$$

where $f(\mathbf{z}|\mathbf{a})$ is the Gaussian density function with mean \mathbf{a} and covariance Σ_z , and the threshold τ is chosen from considering the null hypothesis at a certain false alarm rate. This is equivalent to

$$\min_{\mathbf{a} \in \mathcal{A}_k} \mathbf{a}^T \Sigma_z^{-1} \mathbf{a} - 2\mathbf{z}^T \Sigma_z^{-1} \mathbf{a} \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\geq}} \tau. \quad (20)$$

Thus the GLRT reduces to solving

$$\begin{aligned} &\text{minimize} && \mathbf{a}^T \Sigma_z^{-1} \mathbf{a} - 2\mathbf{z}^T \Sigma_z^{-1} \mathbf{a} \\ &\text{subject to} && \|\mathbf{a}\|_0 \leq k. \end{aligned} \quad (21)$$

For a fixed sparsity pattern, i.e. if we know the support but not necessarily the actual values of \mathbf{a} , the above optimization is easy to solve. In other words, if we know a small set of suspect meters from which malicious may be injected, the above test is easily computable. The sparsity condition on \mathbf{a} makes the above optimization problem non-convex, but for small k it can be solved exactly simply by exhaustively searching through all sparsity patterns. For larger k , this is not feasible. It is a well known technique that (21) can be approximated by a convex optimization:

$$\begin{aligned} &\text{minimize} && \mathbf{a}^T \Sigma_z^{-1} \mathbf{a} - 2\mathbf{z}^T \Sigma_z^{-1} \mathbf{a} \\ &\text{subject to} && \|\mathbf{a}\|_1 \leq \nu \end{aligned} \quad (22)$$

where the L_1 norm constraint is a heuristic for the sparsity of \mathbf{a} . The constant ν needs to be adjusted until the solution involves an \mathbf{a} with sparsity k . This requires solving (22) several times. A similar approach was taken in [30].

D. Classical Detectors with MMSE State Estimation

We will compare the performance of the GLRT detector with two classical bad data detectors [8], [9], both based on the residual error $\mathbf{r} = \mathbf{z} - \mathbf{H}\hat{\mathbf{x}}$ resulted from the MMSE state estimator.

The first is the $J(\hat{\mathbf{x}})$ detector, given by

$$\mathbf{r}^T \Sigma_e^{-1} \mathbf{r} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \tau. \quad (23)$$

The second is the largest normalized residue (LNR) test given by

$$\max_i \frac{|r_i|}{\sigma_{r_i}} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \tau, \quad (24)$$

where σ_{r_i} is the standard deviation of the i th residual error r_i . We may regard this is a test on the l_∞ -norm of the measurement residual, normalized so that each element has unit variance.

The asymptotic optimality of the GLRT detector implies a better performance of GLRT over the above two detectors when the sample size is large. For the finite sample case, numerical simulations shown in Sec VII confirm that the GLRT detector improves the performance of the $J(\hat{\mathbf{x}})$ and LNR detectors. The interesting exception is the case when only one meter is under attack, i.e. $\|\mathbf{a}\|_0 = 1$ and $\Sigma_e = \sigma_e^2 \mathbf{I}$. In this case, the GLRT turns out to be identical to the LNR detector. Therefore, the GLRT can be viewed as a generalization of the LNR detector, in that it can be tuned to any sparsity level. Moreover, this provides some theoretical justification for the LNR detector. The equivalence of the two detectors is stated and proved in the following Proposition.

Proposition 1: When $k = 1$, the GLRT detector given in (20) is equivalent to the LNR detector given in (24).

Proof: If $k = 1$, the left hand side of (20) becomes

$$\min_i \min_{a_i} (\Sigma_z^{-1})_{ii} a_i^2 - 2\mathbf{z}^T (\Sigma_z^{-1})_i a_i \quad (25)$$

where $(\Sigma_z^{-1})_{ii}$ is the i th diagonal element of Σ_z^{-1} , and $(\Sigma_z^{-1})_i$ is the i th row of Σ_z^{-1} . The second minimization can be solved in closed form, so (25) becomes

$$-\max_i \frac{[\mathbf{z}^T (\Sigma_z^{-1})_i]^2}{(\Sigma_z^{-1})_{ii}}. \quad (26)$$

We may therefore write the GLRT as

$$\max_i \frac{|\mathbf{z}^T (\Sigma_z^{-1})_i|}{\sqrt{(\Sigma_z^{-1})_{ii}}} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \tau. \quad (27)$$

The vector of numerators in (27) is given by $\mathbf{r}' = \Sigma_z^{-1} \mathbf{z}$. Note that the covariance matrix of \mathbf{r}' is simply Σ_z^{-1} . Therefore we may regard (27) as a test on the maximum element of the \mathbf{r}' after each element is normalized to unit variance.

We now show that \mathbf{r}' is just a constant multiple of \mathbf{r} , meaning that (27) is identical to (24), saving a constant factor. Recall that $\mathbf{r} = (\mathbf{I} - \mathbf{H}\mathbf{K})\mathbf{z}$, where

$$\begin{aligned} \mathbf{I} - \mathbf{H}\mathbf{K} &= \mathbf{I} - \mathbf{H}\Sigma_x\mathbf{H}^T (\mathbf{H}\Sigma_x\mathbf{H}^T + \Sigma_e)^{-1} \\ &= (\mathbf{H}\Sigma_x\mathbf{H}^T + \Sigma_e - \mathbf{H}\Sigma_x\mathbf{H}^T) (\mathbf{H}\Sigma_x\mathbf{H}^T + \Sigma_e)^{-1} \\ &= \Sigma_e \Sigma_z^{-1} = \sigma_e^2 \Sigma_z^{-1}. \end{aligned}$$

Thus $\mathbf{r} = \sigma_e^2 \mathbf{r}'$; the two detectors are identical. \blacksquare

V. ATTACK OPERATING CHARACTERISTICS AND OPTIMAL ATTACKS

We now study the impact of malicious data attack from the perspective of an attacker. We assume that the attacker knows the (MMSE) state estimator and the (GLRT) detector used by the control center. We also assume that the attacker can

choose k meters arbitrarily in which to inject malicious data. In practice, however, the attacker may be much more limited. Thus our results here are perhaps more pessimistic than in reality.

A. AOC and Optimal Attack Formulations

The attacker faces two conflicting objectives: maximizing the MSE by choosing the best data injection \mathbf{a} vs. avoiding being detected by the control center. The tradeoff between increasing MSE of the state estimator and lower the probability of detection is characterized by *attacker operating characteristics* (AOC), analogous to the receiver operating characteristics (ROC) at the control center. Specifically, AOC is the probability of detection of the detector $\Pr(\delta(\mathbf{z}) = 1 | \mathbf{a})$ as a function of the extra MSE $\mathcal{E}(\mathbf{a}) = \mathcal{E}_0 + \|\mathbf{K}\mathbf{a}\|_2^2$ (16) at the state estimator, where \mathcal{E}_0 is the MMSE in the absence of attack.

The optimal attack in the sense of maximizing the MSE while limiting the probability of detection can be formulated as the following constrained optimization

$$\max_{\mathbf{a} \in \mathcal{A}_k} \|\mathbf{K}\mathbf{a}\|_2^2 \quad \text{subject to} \quad \Pr(\delta(\mathbf{z}) = 1 | \mathbf{a}) \leq \beta, \quad (28)$$

or equivalently,

$$\min_{\mathbf{a} \in \mathcal{A}_k} \Pr(\delta(\mathbf{z}) = 1 | \mathbf{a}) \quad \text{subject to} \quad \|\mathbf{K}\mathbf{a}\|_2^2 \geq C. \quad (29)$$

In order to evaluate the true worst-case performance for any detector, (28) or (29) would need to be solved. This is very difficult, due to the lack of analytical expressions for the detection error probability $\Pr(\delta(\mathbf{z}) = 1 | \mathbf{a})$. We propose a heuristic for $\Pr(\delta(\mathbf{z}) = 1 | \mathbf{a})$, which will allow us to approximate the above optimization with one that is easier to solve.

B. Minimum Residue Energy Attack

Given the naive MMSE state estimator $\hat{\mathbf{x}} = \mathbf{K}\mathbf{z}$ (14-15), the estimation residue error is given by

$$\mathbf{r} = \mathbf{G}\mathbf{z}, \quad \mathbf{G} \triangleq \mathbf{I} - \mathbf{H}\mathbf{K} \quad (30)$$

Substituting the measurement model, we have

$$\mathbf{r} = \mathbf{G}\mathbf{H}\mathbf{x} + \mathbf{G}\mathbf{a} + \mathbf{G}\mathbf{e}.$$

where $\mathbf{G}\mathbf{a}$ is the only term from the attack. Therefore, an attack vector \mathbf{a} will be more difficult to detect at the control center if $\mathbf{G}\mathbf{a}$ is small. Recall from (16) that the damage in MSE done by injecting \mathbf{a} is $\|\mathbf{K}\mathbf{a}\|_2^2$. We therefore consider the following equivalent problems:

$$\max_{\mathbf{a} \in \mathcal{A}_k} \|\mathbf{K}\mathbf{a}\|_2^2 \quad \text{subject to} \quad \|\mathbf{G}\mathbf{a}\|_2^2 \leq \eta, \quad (31)$$

or equivalently,

$$\min_{\mathbf{a} \in \mathcal{A}_k} \|\mathbf{G}\mathbf{a}\|_2^2 \quad \text{subject to} \quad \|\mathbf{K}\mathbf{a}\|_2^2 \geq C. \quad (32)$$

The above optimizations remain difficult due to the constraint $\mathbf{a} \in \mathcal{A}_k$. However, given a specific sparsity pattern $\mathcal{S} \subset \{1, \dots, n\}$ for which $a_i = 0$ for all $i \notin \mathcal{S}$, solving the optimal attack vector \mathbf{a} for the above two formulations amounts to a standard generalized eigenvalue problem.

In particular, for fixed sparsity pattern \mathcal{S} , let $\mathbf{a}_{\mathcal{S}}$ be the nonzero subvector of \mathbf{a} , $\mathbf{K}_{\mathcal{S}}$ the corresponding submatrix of \mathbf{K} , and $\mathbf{G}_{\mathcal{S}}$ similarly defined. The problem (32) becomes

$$\min_{\mathbf{u} \in \mathbb{R}^k} \|\mathbf{G}_{\mathcal{S}}\mathbf{u}\|_2^2 \quad \text{subject to} \quad \|\mathbf{K}_{\mathcal{S}}\mathbf{u}\|_2^2 \geq C. \quad (33)$$

Let $\mathbf{Q}_G \triangleq \mathbf{G}_{\mathcal{S}}^T \mathbf{G}_{\mathcal{S}}$, $\mathbf{Q}_K \triangleq \mathbf{K}_{\mathcal{S}}^T \mathbf{K}_{\mathcal{S}}$. It can be shown that the optimal attack pattern has the form

$$\mathbf{a}_{\mathcal{S}}^* = \sqrt{\frac{C}{\|\mathbf{K}_{\mathcal{S}}\mathbf{v}\|_2^2}} \mathbf{v} \quad (34)$$

where \mathbf{v} is the generalized eigenvector corresponding to the smallest generalized eigenvalue λ_{\min} of the following matrix pencil

$$\mathbf{Q}_G \mathbf{v} - \lambda_{\min} \mathbf{Q}_K \mathbf{v} = \mathbf{0}.$$

This k dimensional symmetrical generalized eigenvalue problem can be solved using the QZ algorithm [31].

VI. ATTACKS ON ELECTRICITY MARKETS

Since malicious attacks can change the state estimation significantly even in the weak regime, it is natural to consider the impact of an attack on the electricity market, since it makes use of state estimation to set prices and calculate payment.

Most deregulated electricity markets in the United States consist of two components: a day-ahead market and a real-time market. In the the day-ahead market, given the load forecast \mathbf{L} , based on the DC lossless model, the following problem is solved to find \mathbf{P}^* , the vector of predicted power generated at each bus:

$$\begin{aligned} & \text{minimize} && \sum_i C_i P_i \\ & \text{subject to} && \sum_i P_i - \sum_j L_j = 0 \\ & && P_i^{\min} \leq P_i \leq P_i^{\max} \\ & && \sum_i S_{ki} P_i - \sum_j S_{kj} L_j \leq T_k^{\max} \end{aligned} \quad (35)$$

where L_j is the forecast load at bus j , P_i^{\min} and P_i^{\max} are the lower and upper capacity bound for the generator at bus i , T_k^{\max} is the line flow limit for branch k , and S_{ki} is the shift factor of branch k to bus i with respect to the reference bus. In particular, S_{ki} is the amount the flow in branch k would change if one additional MW were transferred from bus i to the reference bus. The solution P^* to (35) is the economic dispatch, and the locational marginal price (LMP) at bus i is given by

$$\lambda_i^* = \lambda - \sum_k S_{ki} \mu_k \quad (36)$$

where λ , μ are the dual variables corresponding to the equation and line flow constraints respectively.

The real-time market prices are based on the state estimation, which yields an estimate \hat{P}_i for the power generated at bus i and \hat{L}_i for the power load at bus i . Based on these, we can calculate the power flow through each line; a line is said to be *congested* if the estimated power flow through it exceeds the line flow limit T_k^{\max} . Let \hat{C} be the set of congested

lines. The real-time market uses the following incremental OPF formulation around the operating point found in (35):

$$\begin{aligned} & \text{minimize} && \sum C_i \Delta P_i - \sum C_j \Delta L_j \\ & \text{subject to} && \sum \Delta P_i = \sum \Delta L_j \\ & && \Delta P_i^{\min} \leq \Delta P_i \leq \Delta P_i^{\max} \\ & && \sum_i S_{ki} \Delta P_i + \sum_j S_{kj} \Delta L_j \leq 0, \text{ for all } k \in \hat{C}. \end{aligned} \quad (37)$$

In our simulations, the upper and lower bound of ΔP_i are chosen as 0.1MW and -2MW. The real-time LMP is calculated as

$$\hat{\lambda}_i := \hat{\lambda} - \sum_{j \in \hat{C}} A_{ji} \hat{\mu}_j \quad (38)$$

where $\hat{\lambda}$ and $\hat{\mu}_j$ are the dual variables corresponding to the linear constraint and line flow constraints respectively.

In the day-ahead market, the operator calculates the economic dispatch, yielding \mathbf{P}^* and λ^* ; the generator at bus i receives $P_i^* \lambda_i^*$, and the customer at bus j pays $L_j \lambda_j^*$. In the real time market, the operator uses the state estimate to calculate the real-time LMP $\hat{\lambda}$, then the generator at bus i receives $(\hat{P}_i - P_i^*) \hat{\lambda}_i$ and the customer at bus j pays $(\hat{L}_j - L_j) \hat{\lambda}_j$. Note that \hat{P}_i and \hat{L}_j are calculated from the state estimate, and so may be influenced by the adversary.

Hence, in the real-time market, we see that state estimation is involved in two parts: the calculation of the congestion pattern, and the estimation of generations and loads. The real-time LMP is determined entirely by the congestion pattern. Under the simple scenario that every participant follows the day-ahead optimal dispatch, the differences between the day-ahead dispatch and the estimation of generations and loads are,

$$\hat{\mathbf{P}} - \mathbf{P}^* = \mathbf{H}_P \mathbf{K} \mathbf{z} - \mathbf{P}^* = (\mathbf{H}_P \mathbf{K})(\mathbf{e} + \mathbf{a}) \quad (39)$$

$$\hat{\mathbf{L}} - \mathbf{L} = \mathbf{H}_L \mathbf{K} \mathbf{z} - \mathbf{L}^* = (\mathbf{H}_L \mathbf{K})(\mathbf{e} + \mathbf{a}) \quad (40)$$

where \mathbf{H} , \mathbf{K} , \mathbf{z} , \mathbf{e} , and \mathbf{a} are the same as in previous sections, and \mathbf{H}_P and \mathbf{H}_L are the corresponding part in \mathbf{H} to generation and load respectively.

Finally, the real-time revenue for a generator at bus i is given by

$$(\mathbf{H}_P \mathbf{K})_i (\mathbf{e} + \mathbf{a}) \hat{\lambda}_i \quad (41)$$

where $(\mathbf{H}_P \mathbf{K})_i$ is the i th row of $\mathbf{H}_P \mathbf{K}$.

In Sec. VII, we present numerical simulations of malicious data attacks in the weak regime on the IEEE 14-bus test system. We illustrate that with some probability of detection, the adversary can inject an attack vector to alter real-time LMP and potentially make a profit.

VII. NUMERICAL SIMULATIONS

A. GLRT Performance

We present some simulation results on the IEEE 14 bus system shown in Fig. 1 to compare the performance of the GLRT with the $J(\hat{x})$ test and the LNR test [8], [9]. For various sparsity levels, we find the minimum energy residue attack as discussed in Sec. V-B. The adversary may then scale this attack vector depending on how much it wishes to influence the mean square error. We plot both the ROC and AOC curves

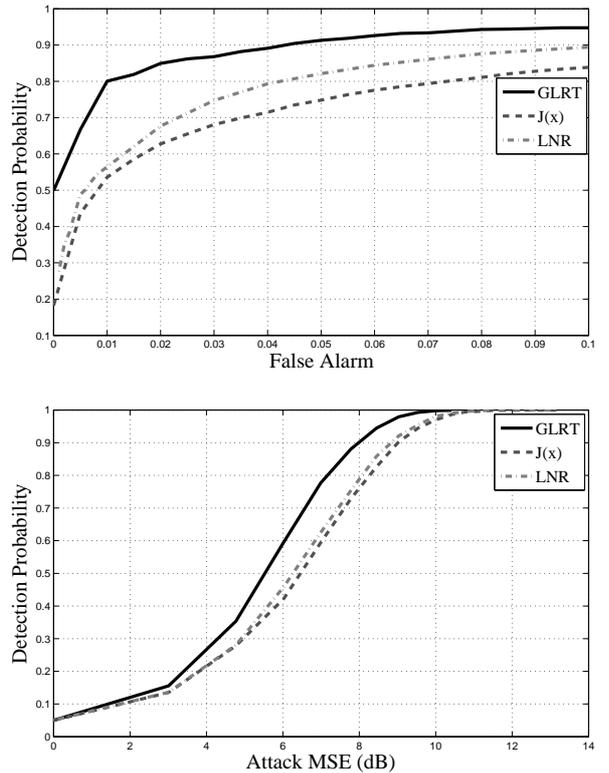


Fig. 3. Above: ROC Performance of GLRT for the 2 sparsity case. MSE with attack is 8db. SNR=10db. Below: AOC Performance of GLRT for the 2 sparsity case. False alarm rate is 0.05. SNR=10db.

for various sparsity levels and all three detectors. For the AOC curve, we fix a probability of false alarm and vary the length of the attack vector along the direction minimizing the energy residue, plotting the MSE vs. the probability of detection. For the ROC curve, we fix the length of the attack vector, but vary the detector's threshold and plot the probability of false alarm vs. probability of detector. In our simulations, we characterize the mean square error increase at the control center using the ratio between the resulting MSE from the attack and the MSE under no attack (i.e. $\mathbf{a} = 0$) in dB. We assume that the prior distribution on the state is given by $\Sigma_x = I\sigma_x^2$, and measurement noise distribution is given by $\Sigma_e = I\sigma_e^2$. We characterize the noise level in terms of SNR, defined as $10 \log \frac{\sigma_x^2}{\sigma_e^2}$.

Fig. 3 shows the ROC and AOC curves for the worst-case 2-sparse attack. We implement the GLRT using exhaustive search over all possible sparsity patterns. This is feasible because of the low sparsity level, so we need not resort to the L_1 minimization as in (22). Observe that the GLRT performs consistently better than the other two conventional detectors.

Fig. 4 shows the ROC and AOC curves for the worst-case 3-sparse attack, again using exhaustive search for the GLRT. Interestingly, the LNR test outperforms the GLRT at this sparsity level. We believe the reason for this is that the GLRT has little recourse when there is significant uncertainty in the sparsity pattern of the attack. In particular, the meters being controlled by the adversary here are the bus injection

meter at bus 1, and the two meters on the transmission line between bus 1 and 2. These constitute three of the seven meters that hold any information about the state at bus 1. Thus, it may be difficult for the detector to determine which of the several meters around bus 1 are the true adversarial meters. The GLRT does not react to this uncertainty: it can only choose the most likely sparsity pattern, which is often wrong. Indeed, in our simulations the GLRT identified the correct sparsity pattern only 4.2% of the time.

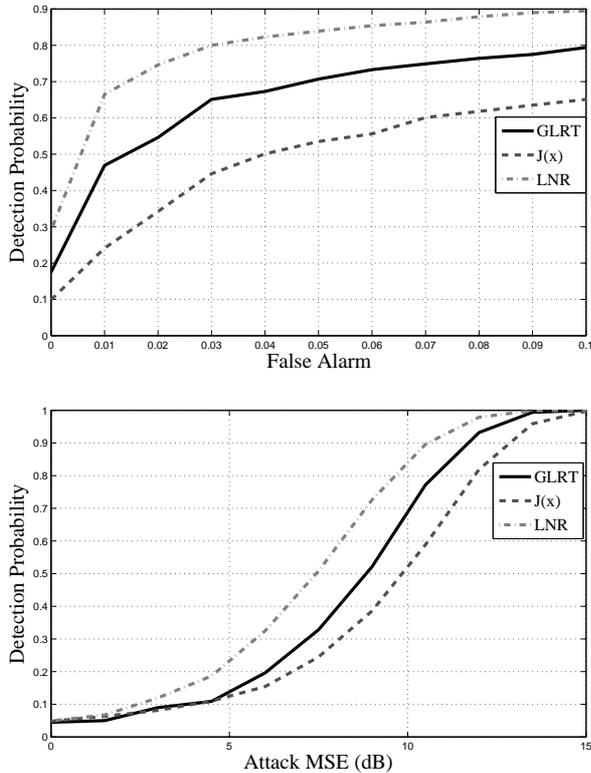


Fig. 4. Above: ROC Performance of GLRT for the 3 sparsity case. MSE with attack is 10db. SNR=10db. Below: AOC Performance of GLRT for 3 sparsity case. False alarm rate is 0.05. SNR=10db

Continuing our analysis of 3-sparsity attacks, we conduct simulations when the adversaries are placed randomly in the network, instead of at the worst-case meters. Once their random meters are chosen, we find the worst-case attack vector using the energy residual heuristic. This simulates the situation that the adversaries cannot choose their locations, but are intelligent and cooperative in their attack. The resulting performance of the three detectors is shown in Fig. 5. Observe that we have recovered the outperformance of the GLRT as compared to the conventional detectors, if only slightly. When the placement of the adversaries is random, they are not as capable of cooperating with one another, therefore their attack is easier to detect.

We increase the sparsity level to 6, at which it is impossible to perform exhaustive search for the GLRT. At this sparsity level, it becomes possible to perform an unobservable attack, so it is not as illuminating to choose the worst-case sparsity pattern, as that would be very difficult to detect. Instead, we again choose the sparsity pattern randomly but optimize

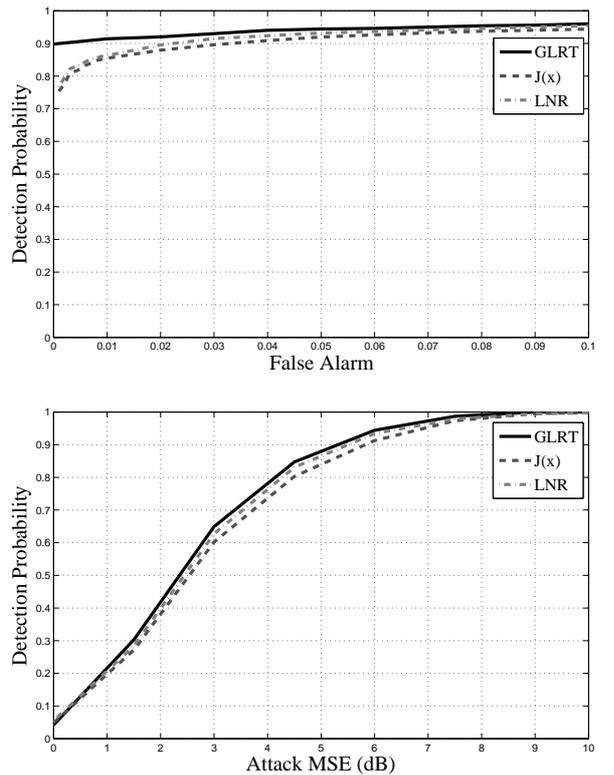


Fig. 5. Above: ROC Performance of GLRT under random attack for 3 sparsity case. MSE with attack is 6db. SNR=10db. Below: AOC Performance of GLRT under random attack for 3 sparsity case. False alarm rate is 0.05. SNR=10db

the attack within it. Fig. 6 compares the performance of the GLRT implemented via L_1 minimization as in (22) to the two conventional detectors. Note again that the GLRT outperforms the others.

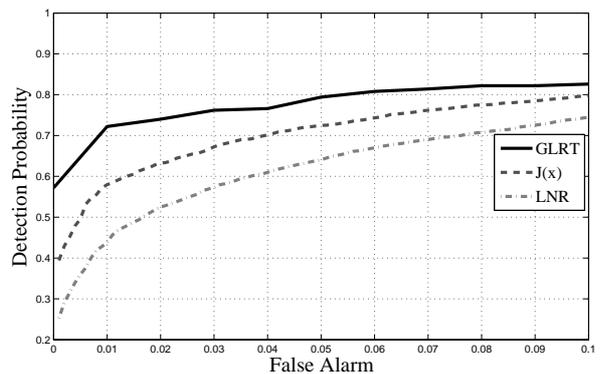


Fig. 6. ROC Performance of GLRT under random attack for 6 sparsity case. MSE with attack is 6db. SNR=10db.

B. Residue Energy Heuristic

We present some numerical evidence that the residue energy described in Sec. V-B works well as a heuristic in that it is roughly increasing with the probability of detection $\Pr(\delta(\mathbf{z}) = 1 | \mathbf{a})$ no matter what detector is used. For the $J(\hat{\mathbf{x}})$

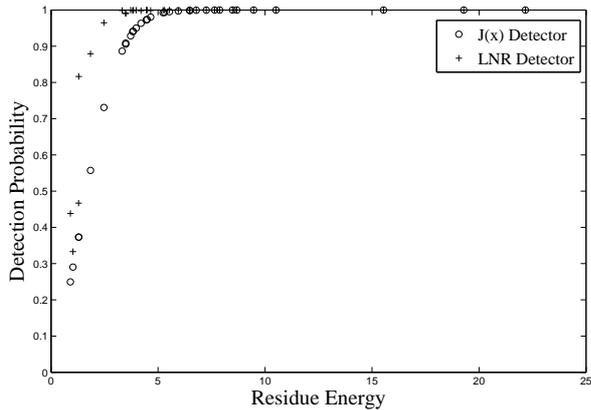


Fig. 7. Comparison of the residue energy heuristic with the true detection probability for 1-sparse attack vectors for both $J(\hat{\mathbf{x}})$ and LNR detectors. False alarm rate is 0.05. SNR is 10dB. Attack MSE is 3dB.

and LNR detectors, we consider the detection probability for all 1-sparse vectors \mathbf{a} satisfying $\|\mathbf{K}\mathbf{a}\|_2^2 = C$ on the 14-bus test system. We plot in Fig. 7 the value of the residue energy vs. the true probability of detector of \mathbf{a} for both detectors. Observe that the scatter plots are roughly increasing.

We evaluate the performance of the residue energy heuristic on 2-sparse vectors in the following way. For each pair of entries i, j of \mathbf{a} , we optimize (32) where \mathbf{a} is constrained to have sparsity pattern $\{i, j\}$. We then evaluate the true probability of detection for the two detectors, with the same parameter values as above. The results are shown in Fig. 8 for the $J(\hat{\mathbf{x}})$ and LNR detectors. Again, the heuristic appears to track the true probabilities reasonably well. This provides some justification for our use earlier in the ROC and AOC curves of approximating the worst-case performance of these detectors by assume the maximum residue energy attack.

C. Electricity Markets

In IEEE 14-bus system, we assume the generators at 1, 2, 3, 6 and 8 can generate real power, with the cost 15, 31, 30, 10, and 20 respectively. We assume the forecasted load at every buses is 100MW. In the day-ahead market, the congestion pattern is branch 1-2 and 8-7. The LMPs at bus 1, 2, 3, 6, 8 are 15, 31, 29.25, 27.02, and 20. Table. I shows that the adversary can disturb the real-time LMP. The leftmost column is the possible congestion pattern the adversary can induce by injecting an attack vector under detection probability 0.5. Each row shows the real-time LMPs with the corresponding congestion pattern. Note that if the congestion pattern is the same as the economic dispatch, the LMPs are the same as the day-ahead market.

Fig. 9 plots the real-time revenue of generation at bus 1 versus detection probability. We compare the effect of an attack on a single meter and an attack on two meters.

VIII. CONCLUSIONS

We studied both adversarial schemes and countermeasures for the control center for malicious data attacks. The problem

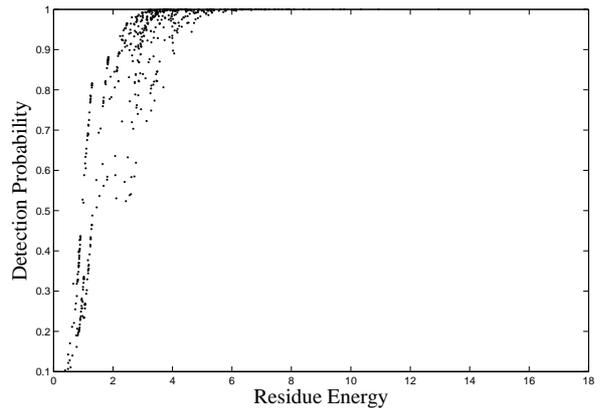
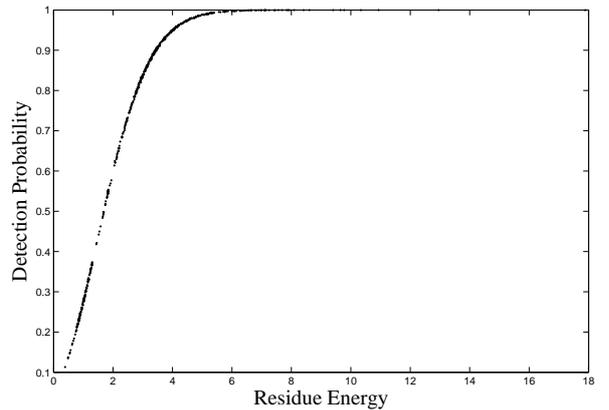


Fig. 8. Comparison of the residue energy heuristic with the true detection probability for 2-sparse attack vectors. False alarm rate is 0.05. SNR is 10dB. Attack MSE is 3dB. Above: Scatter plot for the $J(\hat{\mathbf{x}})$ detector. Below: Scatter plot for the LNR detector.

TABLE I
THE EFFECT OF MALICIOUS DATA ATTACK ON REAL-TIME MARKET PRICE

congestion pattern	λ_1	λ_2	λ_3	λ_6	λ_8
none	31	31	31	31	31
8-7	31	31	31	31	20
1-2	15	31	29.25	27.03	27.55
1-2 and 8-7	15	31	29.25	27.03	20

was divided into two regimes: the strong attack regime, in which so-call unobservable attacks exist, and the weak attack regime, in which they do not. The boundary between these regimes is χ^* , the size of the smallest unobservable malicious data attack, which can be considered a security index for a power system. We provided a characterization of χ^* , which allows for it to be calculated in a computationally efficient manner making use of submodular function minimization. In the weak attack regime, we studied the generalized likelihood ratio test as a detector for this problem; in particular, this detector was implemented using convex optimization via L_1 norm regularization. We also provided a residue energy heuristic to find particularly damaging attacks in this regime. We applied these weak regime techniques to determine the effect of malicious data attacks on prices in electricity markets.

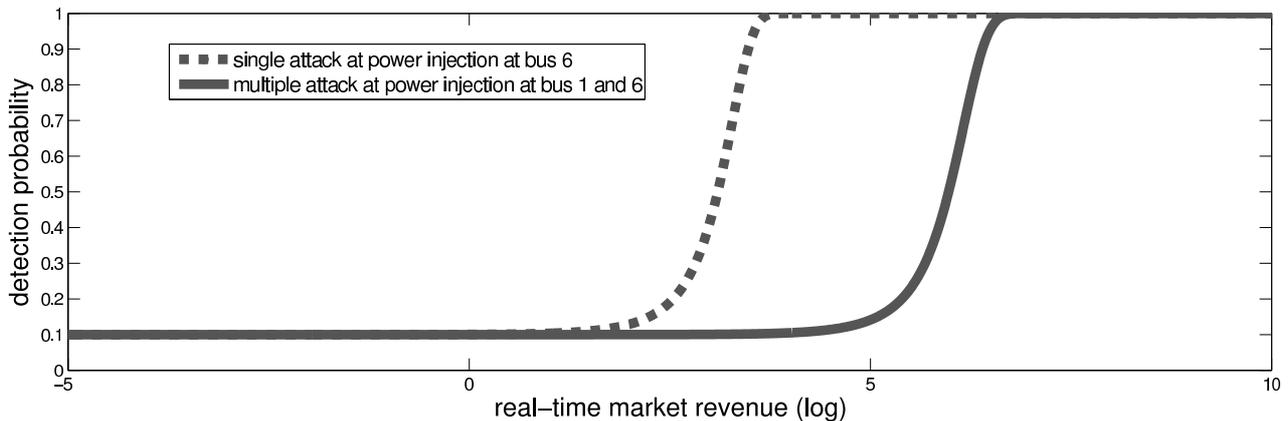


Fig. 9. Real-time revenue of generation at bus 1 versus the detection probability. We compare the effect of single attack at bus 6 and a dual attack on buses 1 and 6.

REFERENCES

- [1] S. Gorman, "Electricity grid in U.S. penetrated by spies," *The Wall Street Journal*, April 2009.
- [2] A. L. Ott, "Experience with pjm market operation, system design, and implementation," *IEEE Trans. Power Systems*, vol. 18, no. 2, pp. 528–534, May 2003.
- [3] T. Zhang and E. Litvinov, "Ex-post pricing in the co-optimized energy and reserv markets," *IEEE Tran. on Power Systems*, vol. 21, no. 4, pp. 1528 – 1538, Nov. 2006.
- [4] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," in *ACM Conference on Computer and Communications Security*, 2009, pp. 21–32.
- [5] O. Zeitouni, J. Ziv, and N. Merhav, "When is the generalized likelihood ratio test optimal," *IEEE Tran. Information Theory*, vol. 38, no. 3, pp. 1597–1602, Mar. 1991.
- [6] B. C. Levy, *Principles of Signal Detection and Parameter Estimation*. New York, NY: Springer, 2008.
- [7] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications, 2nd ed.* NY: Springer, 1998.
- [8] F. C. Schweppe, J. Wildes, and D. P. Rom, "Power system static state estimation, Parts I, II, III," *IEEE Tran. on Power Appar. & Syst.*, vol. PAS-89, pp. 120–135, 1970.
- [9] E. Handschin, F. C. Schweppe, J. Kohlas, and A. Fiechter, "Bad data analysis for power system state estimation," *IEEE Trans. Power Apparatus and Systems*, vol. PAS-94, no. 2, pp. 329–337, Mar/Apr 1975.
- [10] A. Monticelli, "Electric power system state estimation," *Proceedings of the IEEE*, vol. 88, no. 2, pp. 262–282, 2000.
- [11] O. Kosut, L. Jia, R. J. Thomas, and L. Tong, "Malicious data attacks on smart grid state estimation: attack strategies and countermeasures," in *Proc. IEEE 2010 SmartGridComm*, Gaithersburg, MD, USA, Oct 2010.
- [12] H. Sandberg, A. Teixeira, and K. H. Johansson, "On security indices for state estimators in power networks," in *First Workshop on Secure Control Systems, CPSWEEK 2010*, Stockholm, Sweden, Apr 2010.
- [13] G. Dán and H. Sandberg, "Stealth attacks and protection schemes for state estimators in power systems," in *Proc. IEEE 2010 SmartGridComm*, Gaithersburg, MD, USA., Oct 2010.
- [14] H. Sandberg, A. Teixeira, and K. H. Johansson, "On security indices for state estimators in power networks," in *Proc. of the First Workshop on Secure Control Systems, CPSWEEK2010*, Stockholm, Switzerland, April 2010.
- [15] G. R. Krumpholz, K. A. Clements, and P. W. Davis, "Power system observability: a practical algorithm using network topology," *IEEE Trans. Power Apparatus and Systems*, vol. 99, no. 4, pp. 1534–1542, July 1980.
- [16] 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 *First Workshop on Secure Control Systems, CPSWEEK 2010*, Stockholm, Sweden, Apr 2010.
- [17] O. Kosut, L. Jia, R. J. Thomas, and L. Tong, "Limiting false data attacks on power system state estimation," in *Proc. 2010 Conference on Information Sciences and Systems*, Mar 2010.
- [18] —, "On malicious data attacks on power system state estimation," in *Proc. 45th Intl. Univ. Power Engineering Conf.*, Cardiff, Wales, UK, Aug 2010.
- [19] L. Mili, T. V. Cutsem, and M. Ribbens-Pavalla, "Bad data identification methods in power system state estimation—A comparative study," *IEEE J. Select. Areas Commun.*, vol. 16, no. 6, pp. 953–972, August 1998.
- [20] A. Monticelli, F. Wu, and M. Yen, "Multiple bad data identification for state estimation by combinatorial optimization," *IEEE Trans. Power Systems*, vol. PWRD-1, no. 3, pp. 361–369, July 1986.
- [21] M. G. Cheniae, L. Mili, and P. Rousseeuw, "Identification of multiple interacting bad data via power system decomposition," *IEEE Trans. Power Systems*, vol. 11, no. 3, pp. 1555–1563, August 1996.
- [22] A. Abur and A. G. Expósito, *Power System State Estimation: Theory and Implementation*. CRC, 2000.
- [23] R. J. Thomas, L. Tong, L. Jia, and O. E. Kosut, "Some economic impacts of bad and malicious data," in *PSerc 2010 Workshop*, vol. 1, Portland Maine, July 2010.
- [24] L. Xie, Y. Mo, and B. Sinopoli, "False data injection attacks in electricity markets," in *Proc. IEEE 2010 SmartGridComm*, Gaithersburg, MD, USA., Oct 2010.
- [25] A. Monticelli and F. Wu, "Network observability: Theory," *IEEE Trans. Power Apparatus and Systems*, vol. PAS-104, no. 5, pp. 1042–1048, May 1985.
- [26] A. S. M. Grötschel, L. Lovász, "The ellipsoid method and its consequences in combinatorial optimization," *Combinatorica*, vol. 1, no. 2, pp. 169–197, June 1981.
- [27] W. H. Cunningham, "On submodular function minimization," *Combinatorica*, vol. 5, no. 3, pp. 185–192, September 1985.
- [28] A. Schrijver, "A combinatorial algorithm minimizing submodular functions in strongly polynomial time," *Journal of Combinatorial Theory Series B*, vol. 80, no. 2, pp. 346–355, November 2000.
- [29] S. Kourouklis, "A large deviation result for the likelihood ratio statistic in exponential families," *The Annals of Statistics*, vol. 12, no. 4, pp. 1510–1521, 1984.
- [30] D. Gorinevsky, S. Boyd, and S. Poll, "Estimation of faults in DC electrical power systems," in *Proc. 2009 American Control Conf.*, St. Louis, MO., June 2009, pp. 4334–4339.
- [31] G. Golub and C. V. Loan, *Matrix Computations*. Baltimore, Maryland: The Johns Hopkins University Press, 1990.