

Data Framing Attack on State Estimation

Jinsub Kim, Lang Tong, *Fellow, IEEE*, and Robert J. Thomas, *Life Fellow, IEEE*

Abstract—A new mechanism aimed at misleading a power system control center about the source of a data attack is proposed. As a man-in-the-middle state attack, a data framing attack is proposed to exploit the bad data detection and identification mechanisms currently in use at most control centers. In particular, the proposed attack frames meters that are providing correct data as sources of bad data such that the control center will remove useful measurements that would otherwise be used by the state estimator. The optimal design of a data framing attack is formulated as a quadratically constrained quadratic program. It is shown that the proposed attack is capable of perturbing the power system state estimate by an arbitrary degree controlling only half of a critical set of measurements that are needed to make a system unobservable. Implications of this attack on power system operations are discussed, and the attack performance is evaluated using benchmark systems.

Index Terms—Power system state estimation, bad data test, data framing attack, cyber security, smart grid.

I. INTRODUCTION

A FEATURE of any future smart grid is the promise of a data-driven approach to automated monitoring, control, and decision as opposed to the current simulation-driven methods. The paradigm shift to a data-driven framework enables a deeper integration of data collection and sophisticated data processing into the monitoring and control process. While extracting actionable information from real-time sensor data can make a grid more efficient and adaptive to real-time operating conditions, it exposes the grid to possible cyber data attacks aimed at disrupting grid operations and potentially causing blackouts.

In [1], Liu, Ning, and Reiter presented perhaps the first framework for a man-in-the-middle (MiM) attack on the power system state estimation where an adversary would replace “normal” sensor data with “malicious data.” It was shown that, if the adversary could gain control of a sufficient number of meters, it could perturb the state estimate by an arbitrary amount without being detected by the bad data detector employed at the control center. Such undetectable attacks are referred to as *covert data attacks*.

The condition under which a covert data attack is possible was subsequently found in [2] to be equivalent to that of system

unobservability. In particular, a covert attack is possible if and only if the system becomes unobservable when the meters under attack are removed. (or equivalently, the adversary is able to control a critical set of meters.) The minimum number of meters that an adversary has to control in order to launch a covert data attack, referred to as a *security index*, is an important measure of security against a data attack. It represents a fundamental limit on the capability of an adversary to covertly disrupt the operation of a grid [2], [3].

In this paper, we show that a significant barrier on the capability of an adversary to mount an attack of the type described above can be circumvented by using a different form of attacks, that is, one that exploits the vulnerabilities of the existing bad data detection and removal mechanisms. In particular, we show that the adversary only needs to gain control of about half of the meters required by the security index while achieving the same objective of perturbing the state estimate by an arbitrary amount without being detected.

The attacks considered in this paper are referred to as *data framing attacks*, borrowing the notion of framing as that of providing false evidence to make someone innocent appear to be guilty of misconduct. In the context of state estimation, a data framing attack means that an adversary launches a data attack in such a way that the control center identifies properly functioning meters as sources of bad data. To this end, the attacker does not try to cause malicious data to pass the bad data detection without detection (as a covert attack tries to do). Instead, it purposely triggers the bad data detection mechanism and causes erroneous removal of good data. Unknown to the control center, the remaining data still contain adversary-injected malicious data, causing errors in the state estimate.

A. Related Work

There is an extensive literature on *covert data attacks*, following the work of Liu, Ning, and Reiter [1]. While the data framing attack mechanism proposed here is fundamentally different, insights gained from existing work are particularly relevant. Here, we highlight some of these ideas from the literature.

An explicit link between a covert attack on state estimation and system observability was made in [4], [5]. Consequently, classical observability conditions [6]–[8] can be modified for a covert attack and used to develop meter protection strategies [2], [5], [9]–[12]. A particularly important concept is the notion of a critical set of meters [6], [13], [14]. In assessing the vulnerability of the grid, the minimum number of meters necessary for a covert attack was suggested as a security index for the grid in [2], [3]. Subsequently, meter protection strategies were proposed in [15], [16] to optimize this security index.

Manuscript received March 27, 2014; accepted May 10, 2014. Date of publication June 19, 2014; date of current version August 13, 2014. This work was supported in part by the National Science Foundation under Grant CNS-1135844 and in part by the Army Research Office under Grant W911NF1010419. Part of this work was presented at the Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, November 2013.

The authors are with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853 USA (e-mail: jk752@cornell.edu; ltong@cornell.edu; rjt1@cornell.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSAC.2014.2332032

The framing attack strategy considered here relies on bad data identification and removal techniques that have long been the subject of study [13], [17]–[20]. See for example [21], [22] and the references therein. Typically, the residue vectors in normalized forms are widely used as statistics for the bad data test [17]. In particular, Mili *et al.* [20] proposed a hypothesis testing method, in which the set of suspect measurements are determined by the residue analysis in [17]. The use of non-quadratic cost functions in state estimation was also studied to enhance bad data identification performance. Especially, the weighted least absolute value estimation [23]–[26] and the least median of squares regression [27], [28] were considered as alternatives with comparably good performance. In this paper, we take the residue analysis in [17] as a representative bad data test and analyze the effect of a framing attack. However, the same analysis is applicable to other bad data tests.

Detection of data attacks on state estimation, referred to as *state attacks*, has been also studied in various frameworks. Kosut *et al.* [2] presented a generalized likelihood ratio test for detection. Morrow *et al.* [29] proposed the detection mechanism based on network parameter perturbation which deliberately modifies the line parameters and probes whether the measurements respond accordingly to the modification. Distributed detection and estimation of adversarial perturbation was also studied in [30]. In an effort to minimize the detection delay, the attack detection was also formulated as a quickest detection problem, and modified CUSUM algorithms were proposed in [31]–[33].

B. Summary of Results and Organization

We propose a data framing attack on power system state estimation. Specifically, we formulate the design of an optimal data framing attack as a quadratically constrained quadratic program (QCQP). Unlike general QCQPs, which are NP-hard, the proposed QCQP can be solved by finding a maximum eigenvalue of a matrix. To analyze the efficacy of our data framing attack, we present a sufficient condition under which the attack could achieve an arbitrary perturbation of the state estimate by controlling only half of the critical set of meters. We demonstrate the concept using both the IEEE 14-bus network and the IEEE 118-bus network and show that the sufficient condition holds for the critical sets associated with cuts.

The optimal design of our framing attack is based on a linearized system. In practice, a nonlinear state estimator is often used. We demonstrate that, under the usual nonlinear measurement model, a framing attack designed based on a linearized system model successfully perturbs the state estimate, and the attacker is able to control the degree of perturbation as desired.

The remainder of the paper is organized as follows. Section II introduces the measurement and adversary models including preliminaries related to state attacks. Section III presents the state estimation and bad data processing methodology. In Section IV, we present the main idea of the data framing attack and the QCQP framework for the attack design. Section V provides a theoretical justification of the efficacy of the data framing attack. In Section VI, we test the data framing attack

with the IEEE 14-bus network and the IEEE 118-bus network. Finally, Section VII provides concluding remarks.

II. MATHEMATICAL MODELS

This section introduces the topology and system state of a power network, the meter measurement model, and the adversary model. In addition, the covert state attack and its connection with network observability are explained. Throughout the paper, boldface lower case letters (e.g., \mathbf{x}) denote vectors, x_i denotes the i th entry of the vector \mathbf{x} , boldface upper case letters (e.g., \mathbf{H}) denote matrices, \mathbf{H}_{ij} denotes the (i, j) entry of \mathbf{H} , $\mathcal{R}(\mathbf{H})$ denotes the column space of \mathbf{H} , $\mathcal{N}(\mathbf{H})$ denotes the null space of \mathbf{H} , and script letters (e.g., \mathcal{S} , \mathcal{A}) denote sets. The multivariate normal distribution with the mean $\boldsymbol{\mu}$ and the covariance matrix $\boldsymbol{\Sigma}$ is denoted by $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

A. Network and Measurement Models

A power network is a network of buses connected by transmission lines, and thus the *topology* of the grid can be naturally defined as an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} is the set of buses, and \mathcal{E} is the set of lines connecting buses ($\{i, j\} \in \mathcal{E}$ if and only if bus i and bus j are connected.) The *system state* of a power network is defined as the vector of bus voltage magnitudes and phase angles, from which all the other quantities (e.g., power line flows, power injections, line currents) can be calculated.

In order to compute a real-time estimate of the system state, a control center collects measurements from line flow and bus injection meters¹ deployed throughout the grid. The meter measurements are related to the system state \mathbf{x} in a nonlinear fashion, and the relation is described by the AC power flow model [21]:

$$\mathbf{z} = h(\mathbf{x}) + \mathbf{e}, \quad (1)$$

where $h(\cdot)$ is the nonlinear measurement function, and \mathbf{e} is the Gaussian measurement noise with a diagonal covariance matrix.

If some of the meters malfunction or an adversary injects malicious data, the control center observes biased measurements,

$$\bar{\mathbf{z}} = h(\mathbf{x}) + \mathbf{e} + \mathbf{a}, \quad (2)$$

where \mathbf{a} represents a deterministic bias. In such a case, the data are said to be *bad*, and the biased meter entries are referred to as *bad data entries*. Note that even when a meter is protected from adversarial modification, it may still have a bias due to a physical malfunction or an improper parameter setting; filtering out the measurements from such malfunctioning meters was the original objective of the legacy bad data processing and is still in practice today [17].

Even though the model in (1) is nonlinear, the state estimate is generally obtained by iterations of weighted linear least squares estimation with the locally linearized model [21].

¹Other types of meters can also be considered, but we restrict our attention to line flow and bus injection meters for simplicity.




Fig. 1. Adversary model with state estimation and bad data test.

Therefore, it is reasonable to analyze the performance of state estimation using the locally linearized model around the system operating point. To this end, in analyzing the impact of an attack on state estimation, we adopt the so-called DC model [21]. In the DC model, for ease of analysis, the AC model (1) is linearized around the system state where all voltage phasors are equal to $1\angle 0$, and only the real part of the measurements are retained:

$$\mathbf{z} = \mathbf{H}\mathbf{x} + \mathbf{e}, \quad (3)$$

where $\mathbf{z} \in \mathbb{R}^m$ is the measurement vector consisting of real part of line flow and bus injection measurements, the system state $\mathbf{x} \in \mathbb{R}^n$ is the vector of voltage phase angles at all buses except the reference bus (\mathbf{x} is unknown, but deterministic), $\mathbf{H} \in \mathbb{R}^{m \times n}$ is the DC measurement matrix that relates the system state to bus injection and line flow amounts, and \mathbf{e} is the Gaussian measurement noise with a diagonal covariance matrix Σ . We represent the noise covariance matrix Σ as $\Sigma = \sigma^2 \bar{\Sigma}$, where $\bar{\Sigma}$ is a diagonal matrix representing the variation of noise variances across different meters ($\sum_{i=1}^m \bar{\Sigma}_{ii} = 1$), and σ^2 is a scaling factor.

Each row of \mathbf{H} has a special structure depending on the type of the meter [21]. For ease of presentation, consider the noiseless measurement $\mathbf{z} = \mathbf{H}\mathbf{x}$. If an entry z_k of \mathbf{z} is the measurement of the line flow from bus i to bus j , z_k is $B_{ij}(x_i - x_j)$ where B_{ij} is the line susceptance and x_i is the voltage phase angle at bus i [21]. If z_k is the measurement of bus injection at i , it is the sum of all the outgoing line flows from i , and the corresponding row of \mathbf{H} is the sum of the row vectors corresponding to all the outgoing line flows.

Any analysis based on the DC model needs to be verified using realistic AC model simulations; we demonstrate in Section VI that the proposed attack strategy is effective using AC model simulations.

B. An Adversary Model

We consider a man-in-the-middle attack on power system state estimation, where, as described in Fig. 1, an adversary is assumed to be capable of modifying the data from a subset of analog meters \mathcal{S}_A . We refer to the meters in \mathcal{S}_A as *adversary meters*.

The control center observes corrupted measurements $\bar{\mathbf{z}}$ instead of the actual measurements \mathbf{z} in (1). We assume that the adversary knows the line parameters (i.e., the measurement function h and the measurement matrix \mathbf{H}).

The adversarial modification is mathematically modeled as follows:

$$\bar{\mathbf{z}} = \mathbf{z} + \mathbf{a}, \quad \mathbf{a} \in \mathcal{A}, \quad (4)$$

where \mathbf{a} is an attack vector, and \mathcal{A} is the set of feasible attack vectors defined as

$$\mathcal{A} \triangleq \{\mathbf{a} \in \mathbb{R}^m : a_i = 0, \forall i \notin \mathcal{S}_A\}. \quad (5)$$

Note that \mathcal{A} fully characterizes the ability of the adversary. In addition, the adversary is assumed to design a vector \mathbf{a} without observing any entry of \mathbf{z} , i.e., the attack does not require any real-time observation.

C. Network Observability and Covert State Attacks

For state estimation to be feasible, the control center needs to have enough meter measurements so that the system state can be uniquely determined. Formally, a power network is said to be *locally observable* at a state \mathbf{x}_0 if the system state can be uniquely determined from the noiseless meter measurements $h(\mathbf{x})$ in a neighborhood of \mathbf{x}_0 . This implies that the Jacobian of h at \mathbf{x}_0 has full rank. However, due to the intractability of checking local observability for all feasible operating points, the DC model (3) is generally adopted for observability analysis [6]: the network is said to be *observable* if the DC measurement matrix \mathbf{H} has full rank. In practice, power networks should be designed to satisfy the observability requirement. Hence, we assume that the network of interest is observable (i.e., \mathbf{H} has full rank.)

The concept of network observability is closely related to the feasibility of a covert state attack. In particular, we need to introduce the concept of a critical set of meters, formally defined as follows.

Definition 2.1: A subset \mathcal{S}_0 of meters is said to be a *critical set* if removing all meters in \mathcal{S}_0 from the network makes the network unobservable whereas removing any strict subset of \mathcal{S}_0 does not.

The covert state attack was first proposed in [1] for the DC model, and it is formally defined as follows.

Definition 2.2: Given a measurement vector $\mathbf{z} = \mathbf{H}\mathbf{x} + \mathbf{e}$, an attack \mathbf{a} is said to be *covert* if $\bar{\mathbf{z}}$ is equal to $\mathbf{H}\bar{\mathbf{x}} + \mathbf{e}$ for some $\bar{\mathbf{x}} \neq \mathbf{x}$.

Note that if measurements are perturbed by a covert attack, the corrupt measurements appear to be normal measurements from the state $\bar{\mathbf{x}}$. From the above definition, an attack \mathbf{a} is covert if and only if \mathbf{a} is equal to $\mathbf{H}\mathbf{y}$ for some nonzero $\mathbf{y} \in \mathbb{R}^n$. In [2], the condition for existence of a covert attack was characterized as a network unobservability condition, as stated in the following theorem.

Theorem 2.1 ([2]): A covert attack exists if and only if removing the adversary meters renders the network unobservable (or equivalently, the attacker can control at least a critical set of meters.) In addition, if \mathbf{a} is covert, then so is $\gamma \cdot \mathbf{a}$, and $\|\bar{\mathbf{x}} - \mathbf{x}\|_2$ increases to infinity as γ grows.

III. STATE ESTIMATION AND BAD DATA PROCESSING

This section introduces a popular approach to state estimation and bad data processing, which we assume to be employed by the control center. Once the control center receives measurements \mathbf{z} , it aims to obtain an estimate $\hat{\mathbf{x}}$ of the system state \mathbf{x} .

Because bad data entries in \mathbf{z} may result in a bias in the state estimate, the control center employs a mechanism to filter out possible bad data entries in \mathbf{z} .

Fig. 1 illustrates an iterative scheme for obtaining $\hat{\mathbf{x}}$, which consists of three functional blocks: state estimation, bad data detection, and bad data identification [17], [21]. The iteration begins with the initial measurement vector $\mathbf{z}^{(1)} \triangleq \mathbf{z}$ and the initial measurement function $h^{(1)} \triangleq h$ where the superscript denotes the index for the current iteration.

A. State Estimation and Bad Data Detection

At the k th iteration, the state estimator uses $(\mathbf{z}^{(k)}, h^{(k)})$ as an input, and obtains the weighted least squares (WLS) estimate of the system state:

$$\hat{\mathbf{x}}^{(k)} \triangleq \arg \min_{\mathbf{x}} \left(\mathbf{z}^{(k)} - h^{(k)}(\mathbf{x}) \right)^T \left(\Sigma^{(k)} \right)^{-1} \left(\mathbf{z}^{(k)} - h^{(k)}(\mathbf{x}) \right), \quad (6)$$

where $\Sigma^{(k)}$ is the covariance matrix of the corresponding noise vector. Based on the state estimate, the residue vector is also evaluated:

$$\mathbf{r}^{(k)} \triangleq \mathbf{z}^{(k)} - h^{(k)}(\hat{\mathbf{x}}^{(k)}). \quad (7)$$

We assume that the $J(\hat{\mathbf{x}})$ -test [17], [21] is employed for bad data detection: the bad data detector makes a decision based on the sum of weighted squared residues:

$$\begin{cases} \text{bad data} & \text{if } (\mathbf{r}^{(k)})^T \left(\Sigma^{(k)} \right)^{-1} \mathbf{r}^{(k)} > \tau^{(k)}; \\ \text{good data} & \text{if } (\mathbf{r}^{(k)})^T \left(\Sigma^{(k)} \right)^{-1} \mathbf{r}^{(k)} \leq \tau^{(k)}. \end{cases} \quad (8)$$

The $J(\hat{\mathbf{x}})$ -test is widely used due to its low complexity and the fact that the test statistic has a χ^2 distribution if the data are good [17]. The latter fact is used to set the threshold $\tau^{(k)}$ for a given false alarm constraint.

B. Iterative Bad Data Identification and Removal

If the bad data detector (8) declares that data are good, the algorithm returns the state estimate $\hat{\mathbf{x}}^{(k)}$ and terminates. However, if the bad data detector declares that the data are bad, bad data identification is invoked to identify and remove one bad data entry from the measurement vector.

A widely used criterion for identifying a bad data entry is the normalized residue [17], [21], which is considered one of the most reliable criteria [19]. In the normalized residue analysis, each $r_i^{(k)}$ is divided by its standard deviation under the good data hypothesis (i.e., the standard deviation of $r_i^{(k)}$ when there exists no bad data entry in $\mathbf{z}^{(k)}$.) If there exists no bad data entry in $\mathbf{z}^{(k)}$, and the state estimate $\hat{\mathbf{x}}^{(k)}$ is close to the actual state \mathbf{x} , the distribution of $\mathbf{r}^{(k)}$ can be approximated by $\mathcal{N}(\mathbf{0}, \mathbf{W}^{(k)} \Sigma^{(k)})$ where

$$\mathbf{W}^{(k)} \triangleq \mathbf{I} - \mathbf{H}^{(k)} \left(\left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \left(\mathbf{H}^{(k)} \right) \right)^{-1} \left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \quad (9)$$

TABLE I
ITERATIVE STATE ESTIMATION

Iterative-State-Estimation(\mathbf{z}, h, Σ)	
1:	$\mathbf{z}^{(1)} \leftarrow \mathbf{z}; \quad h^{(1)} \leftarrow h; \quad k \leftarrow 1;$
2:	while (<i>true</i>)
3:	$(\hat{\mathbf{x}}^{(k)}, \mathbf{r}^{(k)}) \leftarrow \text{State-Estimation}(\mathbf{z}^{(k)}, h^{(k)});$
4:	$result \leftarrow \text{Bad-Data-Detection}(\mathbf{r}^{(k)});$
5:	if $result == \text{good}$
6:	break;
7:	else
8:	$(\mathbf{z}^{(k+1)}, h^{(k+1)}) \leftarrow \text{Bad-Data-ID}(\mathbf{r}^{(k)}, \mathbf{z}^{(k)}, h^{(k)});$
9:	end
10:	$k \leftarrow k + 1;$
11:	end
12:	return $\hat{\mathbf{x}}^{(k)}$;

with $\mathbf{H}^{(k)}$ denoting the Jacobian of $h^{(k)}$ at $\hat{\mathbf{x}}^{(k)}$ and \mathbf{I} denoting the identity matrix with the appropriate size (see Appendix of [17] for the detail.) Hence, the normalized residue is calculated as

$$\tilde{\mathbf{r}}^{(k)} = \Omega^{(k)} \mathbf{r}^{(k)}, \quad (10)$$

where $\Omega^{(k)}$ is a diagonal matrix with

$$\Omega_{ii}^{(k)} = \begin{cases} 0 & \text{if } \{i\} \text{ is a critical set,}^2 \\ \frac{1}{\sqrt{(\mathbf{W}^{(k)} \Sigma^{(k)})_{ii}}} & \text{otherwise.} \end{cases} \quad (11)$$

Once the normalized residue $\tilde{\mathbf{r}}^{(k)}$ is calculated, the meter with the largest $|\tilde{r}_i^{(k)}|$ is identified as a bad meter. The bad data identification unit removes the row of $\mathbf{z}^{(k)}$ and the row of $h^{(k)}$ that correspond to the bad meter and returns the updated measurement vector and measurement function for the next iteration, denoted by $\mathbf{z}^{(k+1)}$ and $h^{(k+1)}$.

Table I provides the pseudocode for the overall procedure of the iterative state estimation and bad data test.

Using the DC model (3), state estimation, bad data detection, and bad data identification are the same with that in the AC model, except that the nonlinear measurement function $h^{(k)}(\mathbf{x})$ is replaced with the linear function $\mathbf{H}^{(k)} \mathbf{x}$ (so, the Jacobian is the same everywhere.) Note that the WLS state estimate (6) is replaced with a simple linear WLS solution:

$$\hat{\mathbf{x}}^{(k)} = \left(\left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \left(\mathbf{H}^{(k)} \right) \right)^{-1} \left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \mathbf{z}^{(k)}, \quad (12)$$

and thus

$$\mathbf{r}^{(k)} = \mathbf{z}^{(k)} - \mathbf{H}^{(k)} \hat{\mathbf{x}}^{(k)} = \mathbf{W}^{(k)} \mathbf{z}^{(k)}. \quad (13)$$

²If $\{i\}$ is a critical set (i.e., removing the meter i makes the grid unobservable), its residue is always equal to zero [21], and the corresponding diagonal entry of $\mathbf{W}^{(k)} \Sigma^{(k)}$ is zero. For such a meter, the normalizing factor is 0 such that its normalized residue is equal to 0.

IV. DATA FRAMING ATTACK

In this section, we present a new attack strategy on state estimation, referred to as a *data framing attack*, which exploits the bad data processing to remove data from some normally operating meters and make the adversary meters appear to be trustworthy. We present the main idea and the QCQP framework for an optimal attack strategy.

We focus our attention on the case where the adversary cannot control enough meters to launch a covert attack. A framing attack starts with setting the set of *framed meters*, denoted by \mathcal{S}_F , which it aims to frame as sources of bad data. The framed meter set \mathcal{S}_F is chosen such that after the meters in \mathcal{S}_F are removed from the network, a covert attack with the adversary meters in \mathcal{S}_A becomes feasible. For instance, suppose that \mathcal{S} is a critical set. The feasibility condition of a covert attack in Theorem 2.1 implies that if $\mathcal{S} \setminus \mathcal{S}_A$ is removed from the network, then the adversary with \mathcal{S}_A can launch a covert attack; because, further removing all meters in \mathcal{S}_A makes the network unobservable. Therefore, a framing attack can set \mathcal{S}_F to be $\mathcal{S} \setminus \mathcal{S}_A$.

The resulting perturbation of the state estimate by a framing attack does depend on the choice of \mathcal{S}_F . Finding an optimal \mathcal{S}_F for a given attack objective is certainly an important problem. However, it is out of scope of this paper. We focus on the design of the attack vector for a fixed \mathcal{S}_F .

A. Effect of Attack on Normalized Residues

To analyze how an attack affects the bad data processing, we analyze, using the DC model (3), the adversarial effect on the normalized residue vector at the first iteration. In this subsection, we omit the superscript to simplify notation: all the quantities we consider are associated with the first iteration unless otherwise specified.

Suppose that \mathbf{z} is a measurement vector *without* bad data. The normalized residue in the first iteration is obtained as

$$\tilde{\mathbf{r}} = \mathbf{\Omega}\mathbf{r} = \mathbf{\Omega}\mathbf{W}\mathbf{z}, \quad (14)$$

where $\mathbf{\Omega} = \mathbf{\Omega}^{(1)}$ is defined as in (11).

Due to the normalization, each entry \tilde{r}_i is distributed as $\mathcal{N}(0, 1)$ unless $\{i\}$ is a critical set [21]; if $\{i\}$ is a critical set, \tilde{r}_i is equal to zero for any \mathbf{z} .

If an attack vector \mathbf{a} is added, the resulting normalized residue is

$$\tilde{\mathbf{r}} = \mathbf{\Omega}\mathbf{W}(\mathbf{z} + \mathbf{a}) = \mathbf{\Omega}\mathbf{W}\mathbf{z} + \mathbf{\Omega}\mathbf{W}\mathbf{a}. \quad (15)$$

Thus, if $\{i\}$ is not a critical set, \tilde{r}_i is distributed as $\mathcal{N}((\mathbf{\Omega}\mathbf{W}\mathbf{a})_i, 1)$; if $\{i\}$ is critical, $\tilde{r}_i = (\mathbf{\Omega}\mathbf{W}\mathbf{a})_i$ surely.

Recalling that the absolute normalized residues (i.e., $|\tilde{r}_i|$) are the statistics used for identifying the bad data entries, one intuitive heuristic to frame the meters in \mathcal{S}_F as bad is to make the mean energy of the normalized residues at the framed meters as large as possible. Making the framed meters have large normalized residues at the first iteration is of course not a guarantee that their data will be identified as bad and removed in the subsequent iterations. Nevertheless, this is a

reasonable heuristic to avoid the difficult task of analyzing the dynamic adversarial effect in subsequent iterations. Note that the expected energy of the normalized residues at the framed meters is

$$\mathbb{E} \left[\sum_{i \in \mathcal{S}_F} (\tilde{r}_i)^2 \right] = \sum_{i \in \mathcal{S}_F} \mathbb{E} [(\tilde{r}_i)^2] = \sum_{i \in \mathcal{S}_F} (\mathbf{\Omega}\mathbf{W}\mathbf{a})_i^2 + C, \quad (16)$$

where C is the number of the meters in \mathcal{S}_F that do not form a single-element critical set. Therefore, maximizing the mean energy of the normalized residues at \mathcal{S}_F is equivalent to maximizing $\sum_{i \in \mathcal{S}_F} (\mathbf{\Omega}\mathbf{W}\mathbf{a})_i^2 = \|\mathbf{R}_F \mathbf{\Omega}\mathbf{W}\mathbf{a}\|_2^2$ where $\mathbf{R}_F \in \mathbb{R}^{|\mathcal{S}_F| \times m}$ is the row-selection matrix that retains only the rows corresponding to the framed meters.

B. Optimal Framing Attack via QCQP

The ultimate objective of a framing attack is to gain an ability to perturb the state estimate by an arbitrary degree. To this end, a framing attack aims to accomplish two tasks.

The first is to make the bad data processing remove the framed meters such that a covert attack exists in the network with the remaining meters. As discussed in Section IV-A, we attempt to achieve this goal by maximizing the mean energy of the normalized residues at \mathcal{S}_F , which is equivalent to maximizing $\|\mathbf{R}_F \mathbf{\Omega}\mathbf{W}\mathbf{a}\|_2^2$.

The second task is to ensure that the attack becomes covert once the framed meters are removed, thereby making the attack as effective as a covert attack (i.e., enable perturbation of the state estimate by an arbitrary degree.) Let \mathbf{H}_0 denote the $m \times n$ measurement matrix obtained from \mathbf{H} by replacing the rows corresponding to the framed meters with zero row vectors. Then, the attack becomes covert after the framed meters are removed (i.e., the attack vector lies in the column space of the remaining measurement matrix,) if and only if \mathbf{a} is in $\mathcal{R}(\mathbf{H}_0)$. Therefore, we restrict the attack vector \mathbf{a} to be not only in the feasible set \mathcal{A} but also in $\mathcal{R}(\mathbf{H}_0)$.

Based on the aforementioned intuition, we solve the following optimization to find the optimal *direction* to align the attack vector:

$$\begin{aligned} \max_{\mathbf{a}} \quad & \|\mathbf{R}_F \mathbf{\Omega}\mathbf{W}\mathbf{a}\|_2^2 \\ \text{subj.} \quad & \|\mathbf{a}\|_2^2 = 1, \quad \mathbf{a} \in \mathcal{R}(\mathbf{H}_0) \cap \mathcal{A}. \end{aligned} \quad (17)$$

The optimization (17) gives the optimal direction \mathbf{a}^* of an attack vector that maximizes the mean energy of the normalized residues at \mathcal{S}_F , among the feasible directions that render the attack covert after the framed meters are removed.

To provide a more intuitive description of the feasible set of (17), we introduce the $(m - |\mathcal{S}_A| - |\mathcal{S}_F|) \times n$ matrix $\tilde{\mathbf{H}}$ obtained from \mathbf{H} by removing the rows corresponding to the adversary meters and the framed meters. It can be easily seen that $\mathbf{a} \in \mathcal{R}(\mathbf{H}_0) \cap \mathcal{A}$ if and only if $\mathbf{a} = \tilde{\mathbf{H}}\mathbf{x}_0$ for some $\mathbf{x}_0 \in \mathcal{N}(\tilde{\mathbf{H}})$. Therefore, the dimension of $\mathcal{R}(\mathbf{H}_0) \cap \mathcal{A}$ is equal to the dimension of $\mathcal{N}(\tilde{\mathbf{H}})$. For instance, if $\mathcal{S}_A \cup \mathcal{S}_F$ is a critical set, $\tilde{\mathbf{H}}$ has rank $n - 1$, and its null space has dimension one. Therefore, in this case, $\mathcal{R}(\mathbf{H}_0) \cap \mathcal{A}$ is a one-dimensional space, and there is no need to search for the optimal direction. On the other hand,

if $\mathcal{S}_A \cup \mathcal{S}_F$ contains more than one critical sets, the dimension of $\mathcal{N}(\bar{\mathbf{H}})$ is greater than one, and the optimization (17) searches for the optimal direction in the infinite set of feasible directions.

Finally, we set an attack vector \mathbf{a} as $\eta \cdot \mathbf{a}^*$ where $\eta \in \mathbb{R}$ is a parameter that adjusts the direction (i.e., positive or negative depending on the sign of η) and the magnitude of the resulting perturbation of the state estimate. The injected attack vector will cause the bad data detector to invoke the bad data identification mechanism for several iterations.

It is important to point out that removal of all the framed meters is *not* necessary although it does guarantee the successful perturbation of the state estimate. Let N denote the total number of iterations in state estimation. Then, the state estimate is obtained by finding the linear WLS estimate (12) to fit $\mathbf{z}^{(N)}$ using the columns of $\mathbf{H}^{(N)}$. And, as long as $\mathbf{z}^{(N)}$ contains a meter data modified by the attacker, the state estimate will be distorted to fit the adversarially modified data. In some circumstances, an attack vector designed by (17) might cause removal of only a part of the framed meters and some of the adversary meters. Nevertheless, such removal may still be able to make some adversary meter data remain in $\mathbf{z}^{(N)}$ thereby leading to a successful attack.

The optimization (17) can be written as a QCQP:

$$\begin{aligned} \min_{\mathbf{q}} \quad & \mathbf{q}^T \mathbf{P} \mathbf{q} \\ \text{subj.} \quad & \mathbf{q}^T \mathbf{Q} \mathbf{q} - 1 = 0, \quad \mathbf{q} \in \mathbb{R}^p, \end{aligned} \quad (18)$$

where

$$\mathbf{P} \triangleq -(\mathbf{R}_F \Omega \mathbf{W} \mathbf{B})^T (\mathbf{R}_F \Omega \mathbf{W} \mathbf{B}), \quad \mathbf{Q} \triangleq \mathbf{B}^T \mathbf{B}, \quad (19)$$

and $\mathbf{B} \in \mathbb{R}^{m \times p}$ is the basis matrix of the p -dimensional vector space $\mathcal{R}(\mathbf{H}_0) \cap \mathcal{A}$. Note that the dimension p is nonzero because \mathcal{S}_F is set such that the feasible set of (17) is nonempty. In addition, \mathbf{P} is negative semidefinite, and \mathbf{Q} is positive definite since \mathbf{B} has full column rank. The positive definiteness of \mathbf{Q} implies that a solution exists (i.e., the objective function is bounded below.)

The KKT conditions for (18) are as follows:

$$\mathbf{P} \mathbf{q} + \lambda (\mathbf{Q} \mathbf{q}) = 0, \quad \mathbf{q}^T \mathbf{Q} \mathbf{q} - 1 = 0, \quad (20)$$

where λ is the Lagrange multiplier for the equality constraint. The optimal solution \mathbf{q}^* of (18) is the one that results in the minimum objective function value among all (λ, \mathbf{q}) pairs satisfying the KKT conditions (20).

The KKT conditions (20) imply that

$$\begin{aligned} \mathbf{Q}^{-1} \mathbf{P} \mathbf{q} &= \lambda \mathbf{q}; \\ \mathbf{q}^T \mathbf{P} \mathbf{q} &= \mathbf{q}^T (-\lambda \mathbf{Q} \mathbf{q}) = -\lambda \mathbf{q}^T \mathbf{Q} \mathbf{q} = -\lambda. \end{aligned} \quad (21)$$

For any solution (λ, \mathbf{q}) of (20), the first equation means that λ should be an eigenvalue of $\mathbf{Q}^{-1} \mathbf{P}$, and \mathbf{q} should be in the corresponding eigenspace. The second equation means that the objective function value at \mathbf{q} is equal to $-\lambda$. Therefore, we can find an optimal solution \mathbf{q}^* of (18) as follows: (i) find the maximum eigenvalue of $\mathbf{Q}^{-1} \mathbf{P}$, and (ii) find an eigenvector \mathbf{q}^* in the corresponding eigenspace that satisfies $(\mathbf{q}^*)^T \mathbf{Q} \mathbf{q}^* - 1 = 0$.

Once \mathbf{q}^* is found, an optimal solution \mathbf{a}^* of the original problem (17) is obtained as $\mathbf{a}^* = \mathbf{B} \mathbf{q}^*$.

V. FACTOR-OF-TWO RESULT

In this section, we demonstrate that a framing attack enables the attacker controlling only a half of a critical set of meters to perturb the state estimate by an arbitrary degree. Specifically, given a partition $\{\mathcal{S}_1, \mathcal{S}_2\}$ of a critical set of meters, we present a sufficient condition under which the attacker can control one of \mathcal{S}_1 or \mathcal{S}_2 to perturb the state estimate by an arbitrary degree. We provide the numerical evidence based on the IEEE benchmark networks that for a critical set associated with a cut, we can find a partition with $|\mathcal{S}_1| \simeq |\mathcal{S}_2|$ satisfying the sufficient condition.

A. Estimation of Adversarial State Estimate Perturbation

The exact analysis of how a framing attack would perturb the state estimate is a difficult task due to the iterative nature of state estimation and bad data processing. However, assuming that meter SNRs are high, we can estimate the effect of a framing attack as follows. Since SNRs of most practical meters tend to be higher than 46 dB [34], the high meter SNR assumption is reasonable.

Suppose that the attacker adds an attack vector \mathbf{a} to \mathbf{z} , and the iterative state estimation is executed on $\bar{\mathbf{z}}$. The measurement vector at the k th iteration is

$$\bar{\mathbf{z}}^{(k)} = \mathbf{H}^{(k)} \mathbf{x} + \mathbf{a}^{(k)} + \mathbf{e}^{(k)}, \quad (22)$$

where $\mathbf{H}^{(k)}$, $\mathbf{a}^{(k)}$, and $\mathbf{e}^{(k)}$ are obtained from \mathbf{H} , \mathbf{a} , and \mathbf{e} by removing the $(k-1)$ rows corresponding to the meters identified as bad until the $(k-1)$ st iteration. The state estimate $\hat{\mathbf{x}}^{(k)}$ at the k th iteration is

$$\begin{aligned} & \left[\left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \mathbf{H}^{(k)} \right]^{-1} \left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \bar{\mathbf{z}}^{(k)} \\ &= \mathbf{x} + \left[\left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \mathbf{H}^{(k)} \right]^{-1} \left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \\ & \quad \cdot \left(\mathbf{a}^{(k)} + \mathbf{e}^{(k)} \right). \end{aligned} \quad (23)$$

Hence, the state estimate error at the k th iteration is

$$\begin{aligned} \hat{\mathbf{x}}^{(k)} - \mathbf{x} &= \left[\left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \mathbf{H}^{(k)} \right]^{-1} \\ & \quad \times \left(\mathbf{H}^{(k)} \right)^T \left(\Sigma^{(k)} \right)^{-1} \cdot \left(\mathbf{a}^{(k)} + \mathbf{e}^{(k)} \right). \end{aligned} \quad (24)$$

In addition, the residue vector is

$$\begin{aligned} \mathbf{r}^{(k)} &= \mathbf{W}^{(k)} \bar{\mathbf{z}}^{(k)} = \mathbf{W}^{(k)} \left(\mathbf{H}^{(k)} \mathbf{x} + \mathbf{a}^{(k)} + \mathbf{e}^{(k)} \right) \\ &= \mathbf{W}^{(k)} \left(\mathbf{a}^{(k)} + \mathbf{e}^{(k)} \right). \end{aligned} \quad (25)$$

From (24) and (25), we can see that both the state estimate error and the residue vector do not depend on the actual state \mathbf{x} . Considering that bad data detection and identification at each

iteration exclusively rely on the residue vector, the observation from (24) and (25) implies that if we aim to analyze how much the attack would perturb the final state estimate, i.e., $\hat{\mathbf{x}}^{(N)} - \mathbf{x}$, where N denotes the total number of iterations, we can simply work with $\mathbf{a} + \mathbf{e}$ by assuming that \mathbf{x} is equal to $\mathbf{0}$.

Furthermore, if meter SNRs are significantly large (i.e., $\sigma^2 \ll 1$), we can estimate the resulting perturbation of the state estimate by running the *noiseless* version of the iterative state estimation on the attack vector \mathbf{a} and checking the resulting $\hat{\mathbf{x}}^{(N)}$. The noiseless version means the algorithm which the iterative state estimation converges to as σ^2 decays to 0. Specifically, Σ is replaced³ by $\bar{\Sigma}$, and at the k th iteration, the bad data detector declares presence of bad data if and only if $(\mathbf{r}^{(k)})^T (\bar{\Sigma}^{(k)})^{-1} \mathbf{r}^{(k)} > 0$ (i.e., the data are declared to be good if and only if state estimation results in a zero residue vector.)

B. Factor-of-Two Theorem for Critical Sets

Suppose that $\{\mathcal{S}_1, \mathcal{S}_2\}$ is a partition of a critical set, and let $\bar{\mathbf{H}}$ denote the measurement matrix after removing the meters in $\mathcal{S}_1 \cup \mathcal{S}_2$ from the network. Since $\mathcal{S}_1 \cup \mathcal{S}_2$ is a critical set, $\bar{\mathbf{H}}$ has rank $n - 1$, and the dimension of its null space is one. Let $\Delta \mathbf{x}$ denote a unit basis vector of the null space of $\bar{\mathbf{H}}$. Recalling the discussion in Section IV-B, if \mathcal{S}_1 is the set of adversary meters, and \mathcal{S}_2 is the set of framed meters, then the framing attack aligns the attack vector along $\mathbf{H}_1 \Delta \mathbf{x}$, where \mathbf{H}_1 is the $m \times n$ matrix obtained from \mathbf{H} by replacing the rows corresponding to the meters in \mathcal{S}_2 with zero row vectors (\mathbf{H}_2 is defined in the same way by replacing the rows of \mathbf{H} corresponding to \mathcal{S}_1 with zero row vectors.)

The following theorem provides a sufficient condition that guarantees that a framing attack can use one of \mathcal{S}_1 and \mathcal{S}_2 to perturb the state estimate by an arbitrary degree under the high SNR setting. The condition is based on the result of executing the analysis described in Section V-A.

Theorem 5.1: Suppose that if we run the noiseless version of the iterative state estimation on $\mathbf{H}_1 \Delta \mathbf{x}$, then there exists a unique state $\mathbf{y} \in \mathbb{R}^n$ such that the final state estimate is always equal to \mathbf{y} (i.e., $\hat{\mathbf{x}}^{(N)} = \mathbf{y}$) regardless of whatever decisions are made under tie⁴ situations in bad data identification. Under this condition, the following statements hold for any true state $\mathbf{x} \in \mathbb{R}^n$:

- 1) If a framing attack using \mathcal{S}_1 and \mathcal{S}_2 as \mathcal{S}_A and \mathcal{S}_F respectively (i.e., $\mathbf{a} = \eta \cdot \mathbf{H}_1 \Delta \mathbf{x}$ where $\eta \in \mathbb{R}$ is a scaling factor) is launched, then

$$\lim_{\sigma^2 \rightarrow 0} \Pr \left(\bar{\mathbf{z}}^{(N)} = \mathbf{H}^{(N)} (\mathbf{x} + \eta \cdot \mathbf{y}) + \mathbf{e}^{(N)} \right) = 1, \quad (26)$$

where N is the random variable representing the total number of iterations in the iterative state estimation.

³Note that state estimation and bad data identification are not affected by the value of σ^2 . Because, σ^2 gets cancelled out in the state estimate expression (12), and bad data identification relies on the relative order of the normalized residue magnitudes, which are not affected by the value of σ^2 . Only bad data detection is affected by the decaying σ^2 .

⁴It is possible that a *tie* may occur in bad data identification at some iteration: i.e., the largest absolute normalized residue is assumed by more than one meters. In a tie situation, we assume that bad data identification chooses an arbitrary meter with the largest absolute normalized residue.

- 2) If a framing attack using \mathcal{S}_2 and \mathcal{S}_1 as \mathcal{S}_A and \mathcal{S}_F respectively (i.e., $\mathbf{a} = \eta \cdot \mathbf{H}_2 \Delta \mathbf{x}$) is launched,

$$\lim_{\sigma^2 \rightarrow 0} \Pr \left(\bar{\mathbf{z}}^{(N)} = \mathbf{H}^{(N)} (\mathbf{x} + \eta \cdot (\Delta \mathbf{x} - \mathbf{y})) + \mathbf{e}^{(N)} \right) = 1. \quad (27)$$

Proof: See Appendix. ■

Remark: If no tie occurs in bad data identification, the condition of Theorem 5.1 is naturally satisfied. Note that a tie in bad data identification is a rare event, so the condition is likely to hold for a general partition of a critical set.

The event $\{\bar{\mathbf{z}}^{(N)} = \mathbf{H}^{(N)} (\mathbf{x} + \eta \cdot \mathbf{y}) + \mathbf{e}^{(N)}\}$ in (26) means that the remaining measurements at the final iteration of state estimation appear to be normal measurements from the perturbed state $\mathbf{x} + \eta \cdot \mathbf{y}$. Theorem 5.1 implies that if the condition is met, then at least one of \mathcal{S}_1 and \mathcal{S}_2 can be used as the set of adversary meters by a framing attack to perturb the state estimate by an arbitrary degree, because \mathbf{y} and $\Delta \mathbf{x} - \mathbf{y}$ cannot be simultaneously $\mathbf{0}$. Especially, if the condition holds for a partition with $|\mathcal{S}_1| = |\mathcal{S}_2|$, then the adversary controlling only a half of the critical set can perturb the state estimate by an arbitrary degree.

One important question is whether a partition $\{\mathcal{S}_1, \mathcal{S}_2\}$ with $|\mathcal{S}_1| \simeq |\mathcal{S}_2|$ that satisfies the condition of Theorem 5.1 can be found in general. To answer this question, we investigated critical sets *associated with cuts*⁵ in the IEEE 14-bus network and the IEEE 118-bus network, where every bus has an injection meter and every line has line meters for both directions. The spanning tree observability criterion in [6] implies that the set \mathcal{S} of the meters associated with a cut (i.e., the set of all line meters on the cut-set lines and all injection meters at the endpoints of the cut-set lines) is a critical set if the cut decomposes the network topology into two connected graphs. For instance, the cut in Fig. 2 disconnects the bus 3 from the rest of the network, and $\{\{2,3\}, \{3,4\}\}$ is the associated cut-set. The set of circled red meters is the critical set associated with the cut.

We executed 20 000 runs of the random contraction algorithm by Karger and Stein [35]—a randomized algorithm for finding a cut—and found 118 cuts in the 14-bus network and 290 cuts in the 118-bus network. For each cut, we found a partition $\{\mathcal{S}_1, \mathcal{S}_2\}$ of the critical set \mathcal{S} associated with the cut such that $|\mathcal{S}_1| \simeq (|\mathcal{S}|/2)$: \mathcal{S}_1 consists of the line meters (both directions) on a subset of cut-set lines such that $||\mathcal{S}_1| - (|\mathcal{S}|/2)| \leq 1$, and \mathcal{S}_2 is set to be $\mathcal{S} \setminus \mathcal{S}_1$. In both networks, for every cut we considered, the partition constructed in the aforementioned manner satisfied the condition of Theorem 5.1; this suggests that controlling about a half of a critical set associated with a cut, a framing attack may perturb the state estimate by an arbitrary degree.⁶

VI. NUMERICAL RESULTS

We tested the performance of framing attacks with the IEEE 14-bus network and the IEEE 118-bus network using the AC

⁵A cut of an undirected graph $(\mathcal{V}, \mathcal{E})$ is defined as a partition $\{\mathcal{V}_1, \mathcal{V}_2\}$ of \mathcal{V} consisting of two nonempty subsets, and the associated cut-set is the subset of lines connecting two vertices in different partitions.

⁶The average size of the critical sets we considered is 15.7 for the 14-bus network and 12.7 for the 118-bus network.




Fig. 2. IEEE 14-bus network: the rectangles on lines and buses represent line flow meters and bus injection meters respectively. The line meter on the line $\{i, j\}$, that is closer to i , measures the power flow from i to j . The red dashed line describes a cut, and the circled meters are the meters associated with the cut.

model and the nonlinear iterative state estimation described in Section III. The simulation results demonstrate the efficacy of framing attacks under the real-world power system setting. Because the ultimate goal of the attack is to perturb the state estimate, we measure the mean l_2 -norm of the resulting state estimate error:

$$\mathbb{E} [\|\hat{\mathbf{x}} - \mathbf{x}\|_2]$$

where $\hat{\mathbf{x}}$ is the state estimate, and \mathbf{x} is the true state.

A. Simulation Setting

For each test network, we chose representative attack scenarios (i.e., \mathcal{S}_A and \mathcal{S}_F) and tested the performance of framing attacks. For each case, we ran Monte Carlo simulations to evaluate the mean state estimate error. In each Monte Carlo run, the true state \mathbf{x} was generated by a multivariate Gaussian distribution with small variances.⁷ Its mean was set as the operating state given by the IEEE data [36], which is far from the nominal state used to obtain the DC model. Based on the state \mathbf{x} , noisy measurements were generated by the AC model (i.e., $h(\mathbf{x}) + \mathbf{e}$). The attack vector was constructed based on the DC measurement matrix \mathbf{H} as described in Section IV. Once constructed, the attack vector was added to the noisy measurements, and the iterative state estimation⁸ was executed on the corrupted measurements. After the iterative state estimation finished, we measured $\|\hat{\mathbf{x}}^{(N)} - \mathbf{x}\|_2$.

Note that the design of a framing attack was studied using the DC model which has only the real part of the measurements. For the simulations, we designed an attack vector based on the DC model, and the attack modified only the corresponding

⁷The standard deviation of each phase angle is set to be 1.15 degree. The standard deviation of each voltage magnitude is set to be 0.01 p.u.

⁸The false alarm rate of the bad data detector is set to be 0.04 throughout all the simulations.

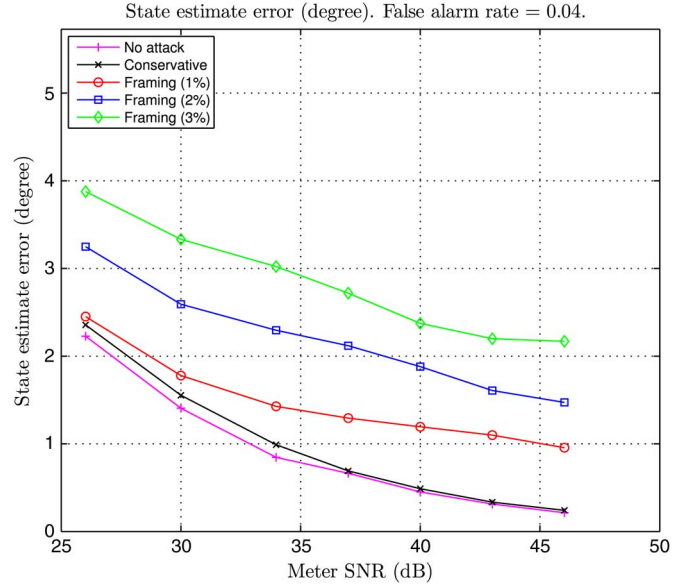


Fig. 3. The 14-bus network: 1,000 Monte Carlo runs. The adversary meters are (2, 3), (3, 4), and (4, 3), and the framed meters are (2), (3), (4), (3, 2).

real part of the measurements. Considering the linear decoupled model (see Chapter 2.7 in [21]), such an attack is expected to modify primarily the bus voltage phase angles and have little effect on the bus voltage magnitudes. Hence, in interpreting the results, we focus on the perturbation of the phase angle part of the state estimate.

For comparison, we also executed the conservative scheme in [2], which aims to perturb the state estimate by the maximum degree while avoiding detection by the bad data detector. In the conservative scheme, the attack vector was designed as a solution to

$$\begin{aligned} \max_{\mathbf{a} \in \mathcal{A}} & \left\| (\mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{a} \right\|_2^2 \\ \text{subj. } & \mathbf{r}^T \boldsymbol{\Sigma}^{-1} \mathbf{r} \leq \tau, \end{aligned} \quad (28)$$

where the constraint guarantees that the alarm is not raised at all, and the objective function is the resulting perturbation of the state estimate due to the attack vector.

B. Simulation Results With 14-Bus Network

We first tested the case where the adversary can control only a half of a critical set. Specifically, we considered the adversary who can control (2, 3), (3, 4), and (4, 3): (i, j) denotes the line meter for the power flow from i to j , and (i) denotes the injection meter at bus i . The framed meters were set to be (3, 2), (2), (3), and (4) such that the set of adversary meters and framed meters is the critical set associated with the cut in Fig. 2. We tested framing attacks with three different attack magnitudes: $\|\mathbf{a}\|_1$ is 1%, 2%, or 3% of $\|\mathbf{z}\|_1$.

Fig. 3 shows the resulting state estimate error versus the meter SNR. The meter SNR ranges from 26 dB to 46 dB (equivalently, the noise-to-signal amplitude ratio ranges from 5% to 0.5%.) The normal state estimate error and the state estimate error under the conservative attack are very close, and both decay to zero as the SNR increases. However, the state estimate




Fig. 4. The 14-bus network: 1,000 Monte Carlo runs. The adversary meters are (2, 3), (3, 4), (4, 3), (6, 12), (12, 6), and (12, 13), and the framed meters are (2), (3), (4), (3, 2), (6), (12), (13), and (13, 12).

TABLE II
THE THREE BUSES WHOSE PHASE ANGLES ARE MOST
SIGNIFICANTLY PERTURBED BY EACH ATTACK:
1,000 MONTE CARLO RUNS, SNR = 46 dB

(2), (3), (4), (3, 2), (6), (12), (13), (13, 12)	(2), (3) (4), (3, 2)	(6), (12), (13), (13, 12)
1) bus 12: 2.075° 2) bus 3: 0.272° 3) bus 14: -0.180°	1) bus 3: -2.183° 2) bus 14: 0.182° 3) bus 9: 0.168°	1) bus 12: 2.878° 2) bus 14: 0.005° 3) bus 9: 0.004°

errors in the presence of framing attacks converge to constants as the SNR increases, and the constants are proportional to the attack magnitudes especially in the high SNR region. The result implies that a framing attack can adjust the state estimate perturbation by choosing a proper attack magnitude.

Second, we demonstrate that a framing attack may perturb the state estimate in various directions depending on the choice of the set of framed meters. We considered the case that the adversary controls (2, 3), (3, 4), (4, 3), (6, 12), (12, 6), and (12, 13). Note that the adversary still cannot control any critical set, and thus a covert attack is infeasible. The framing attack with \mathcal{S}_F equal to any of the following three sets successfully perturbed the state estimate: (i) (2), (3), (4), (3, 2), (6), (12), (13), and (13, 12); (ii) (2), (3), (4), and (3, 2); (iii) (6), (12), (13), and (13, 12). or instance, Fig. 4 shows the state estimate error versus the meter SNR for the first set. While the three sets all resulted in successful state estimate perturbation, each resulted in a different direction of perturbation. For each \mathcal{S}_F , Table II shows the three buses, whose phase angle estimates were most significantly perturbed, and the mean perturbation of their phase angle estimates; positive perturbation means overestimation, and negative perturbation means underestimation. The table demonstrates that the adversary controlling a large number of meters may adjust the direction of perturbation by setting \mathcal{S}_F properly. Note that the impact of a framing attack with a specific \mathcal{S}_F can be estimated based on the analysis in Section V-A.

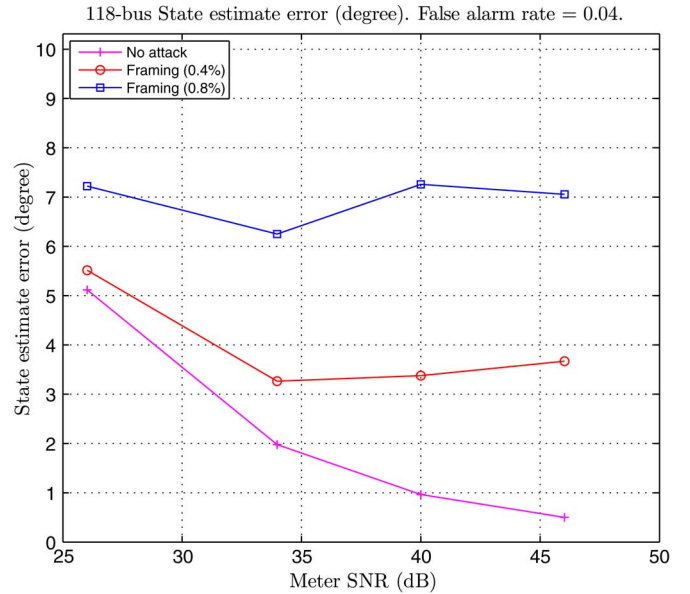


Fig. 5. The 118-bus network: 250 Monte Carlo runs. The adversary meters are (20,21), (21, 20), and (21, 22), and the framed meters are (20), (21), (22), and (22, 21).

C. Simulation Results With 118-Bus Network

Through the simulations with the 118-bus network, we aim to demonstrate the effect of framing attacks on a larger network. We considered the scenario where the adversary controls (20, 21), (21, 20), and (21, 22), and the framed meters are (20), (21), (22), and (22, 21); i.e., the set of the adversary meters and the framed meters is the critical set associated with the cut isolating the bus 21 from the rest of the network. Fig. 5 shows the state estimate errors under the non-attack scenario and the framing attacks with different attack magnitudes. The results imply that the framing attacks successfully perturbed the state estimate in a large network.

VII. CONCLUSION

We have presented a data framing attack on power system state estimation. Controlling only a half of a critical set, the data framing attack can perturb the state estimate by an arbitrary degree. A theoretical justification was provided, and numerical experiments demonstrated the efficacy of the framing attack.

Our results indicate that most known countermeasures, that are aimed at merely preventing covert state attacks, are not sufficient for protection against attacks on state estimation. The proposed framing attacks can successfully perturb the state estimate even when those countermeasures are employed. In designing a countermeasure, the possibility of a framing attack needs to be taken into account.

An important direction for future work is to design a mechanism that can nullify the attack impact once the presence of an attack is detected. In particular, we need a robust state estimation mechanism that can produce an unbiased state estimate with reasonable accuracy even when some data entries are untrustworthy due to the attacker's modification.

APPENDIX
PROOF OF THEOREM 5.1

Let \mathcal{S} denote the set of sequences of meter removals that can possibly happen when the noiseless version of the iterative state estimation is executed on $\mathbf{H}_1\Delta\mathbf{x}$. In other words, $(a_1, \dots, a_M) \in \mathcal{S}$ if and only if some decisions under tie situations may result in the removal of the meters $\{a_1, \dots, a_M\}$ in the order of a_1, \dots, a_M . The cardinality of \mathcal{S} can be greater than 1 since different decisions under tie situations may result in different sequences of meter removals.

For any sequence $(a_1, \dots, a_M) \in \mathcal{S}$, the existence of such \mathbf{y} —as described in the condition—implies that if all the meters in the sequence are removed, the remaining part of $\mathbf{H}_1\Delta\mathbf{x}$, denoted by $\mathbf{H}_1^{(M)}\Delta\mathbf{x}$, is equal to $\mathbf{H}^{(M)}\mathbf{y}$, where $\mathbf{H}_1^{(M)}$ and $\mathbf{H}^{(M)}$ are obtained from \mathbf{H}_1 and \mathbf{H} respectively, by removing the rows corresponding to all meters in the sequence.

Now, consider running the iterative state estimation on $\mathbf{H}\mathbf{x} + \mathbf{H}_1\Delta\mathbf{x} + \mathbf{e}$, which is the resulting measurement vector when the framing attack is launched with \mathcal{S}_A and \mathcal{S}_F equal to \mathcal{S}_1 and \mathcal{S}_2 respectively. The (25) implies that the residue vector in each iteration only depends on $\mathbf{H}_1\Delta\mathbf{x} + \mathbf{e}$. In addition, as σ^2 decreases to zero, the results of bad data detection and identification heavily depend on $\mathbf{H}_1\Delta\mathbf{x}$, and thus the sequence of meter removals becomes highly likely to be in \mathcal{S} . Formally,

$$\lim_{\sigma^2 \rightarrow 0} \Pr((a_1, \dots, a_N) \in \mathcal{S}) = 1, \quad (29)$$

where (a_1, \dots, a_N) is a sequence of meters removed by the bad data identification rule. Let $\mathbf{H}^{(N)}$ and $\mathbf{e}^{(N)}$ denote the matrix and vector obtained from \mathbf{H} and \mathbf{e} respectively by removing the rows corresponding to $\{a_1, \dots, a_N\}$. Note that N , (a_1, \dots, a_N) , $\mathbf{H}^{(N)}$, and $\mathbf{e}^{(N)}$ are all random.

The event $\{(a_1, \dots, a_N) \in \mathcal{S}\}$ implies that $\mathbf{H}_1^{(N)}\Delta\mathbf{x} = \mathbf{H}^{(N)}\mathbf{y}$, and thus

$$\bar{\mathbf{z}}^{(N)} = (\mathbf{H}^{(N)}\mathbf{x} + \mathbf{e}^{(N)}) + \mathbf{H}_1^{(N)}\Delta\mathbf{x} = (\mathbf{H}^{(N)}\mathbf{x} + \mathbf{e}^{(N)}) + \mathbf{H}^{(N)}\mathbf{y}. \quad (30)$$

Therefore,

$$\lim_{\sigma^2 \rightarrow 0} \Pr(\bar{\mathbf{z}}^{(N)} = (\mathbf{H}^{(N)}\mathbf{x} + \mathbf{e}^{(N)}) + \mathbf{H}^{(N)}\mathbf{y}) = 1. \quad (31)$$

When the attack vector $\mathbf{H}_1\Delta\mathbf{x}$ is replaced with $\eta \cdot \mathbf{H}_1\Delta\mathbf{x}$, we can use the exactly same logic to derive the following.

$$\lim_{\sigma^2 \rightarrow 0} \Pr(\bar{\mathbf{z}}^{(N)} = (\mathbf{H}^{(N)}\mathbf{x} + \mathbf{e}^{(N)}) + \eta \cdot \mathbf{H}^{(N)}\mathbf{y}) = 1. \quad (32)$$

Now, consider running the iterative state estimation over $\mathbf{H}\mathbf{x} + \mathbf{H}_2\Delta\mathbf{x} + \mathbf{e}$; this is the case when the framing attack is launched with \mathcal{S}_A and \mathcal{S}_F equal to \mathcal{S}_2 and \mathcal{S}_1 respectively. First, note that

$$\mathbf{H}\Delta\mathbf{x} = \mathbf{H}_1\Delta\mathbf{x} + \mathbf{H}_2\Delta\mathbf{x}. \quad (33)$$

Therefore, $\mathbf{H}\mathbf{x} + \mathbf{H}_2\Delta\mathbf{x} + \mathbf{e}$ is equivalent to $\mathbf{H}(\mathbf{x} + \Delta\mathbf{x}) - \mathbf{H}_1\Delta\mathbf{x} + \mathbf{e}$.

Suppose we run the noiseless version of the iterative state estimation on $-\mathbf{H}_1\Delta\mathbf{x}$. The set of sequences of meter removals that can possibly happen is equivalent to \mathcal{S} , because the sign change only flips the signs of residue entries; it does not affect

their absolute values, which are the statistics used for detection and identification of bad data entries. Furthermore, it can be easily seen that the final state estimate is always equal to $-\mathbf{y}$ regardless of whatever decisions are made under the tie situations.

The above paragraph implies that we can analyze the result of running the iterative state estimation on $\mathbf{H}(\mathbf{x} + \Delta\mathbf{x}) - \mathbf{H}_1\Delta\mathbf{x} + \mathbf{e}$ by following the same procedure of deriving (31). In particular, one can easily derive the following:

$$\lim_{\sigma^2 \rightarrow 0} \Pr(\bar{\mathbf{z}}^{(N)} = \mathbf{H}^{(N)}(\mathbf{x} + \Delta\mathbf{x}) + \mathbf{e}^{(N)} + \mathbf{H}^{(N)}(-\mathbf{y})) = 1, \quad (34)$$

or equivalently,

$$\lim_{\sigma^2 \rightarrow 0} \Pr(\bar{\mathbf{z}}^{(N)} = \mathbf{H}^{(N)}\mathbf{x} + \mathbf{e}^{(N)} + \mathbf{H}^{(N)}(\Delta\mathbf{x} - \mathbf{y})) = 1. \quad (35)$$

When the attack vector $\mathbf{H}_2\Delta\mathbf{x}$ is scaled by η (i.e., $\mathbf{a} = \eta \cdot \mathbf{H}_2\Delta\mathbf{x}$), we can derive the following in a similar manner:

$$\lim_{\sigma^2 \rightarrow 0} \Pr(\bar{\mathbf{z}}^{(N)} = \mathbf{H}^{(N)}\mathbf{x} + \mathbf{e}^{(N)} + \eta \cdot \mathbf{H}^{(N)}(\Delta\mathbf{x} - \mathbf{y})) = 1. \quad (36)$$

Therefore, the proof is complete. \blacksquare


ACKNOWLEDGMENT

We thank anonymous reviewers for considerate and helpful comments. Their comments helped us improve the manuscript.

REFERENCES

- [1] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," in *Proc. 16th ACM Conf. Comput. Commun. Security*, 2009, pp. 21–32.
- [2] O. Kosut, L. Jia, R. J. Thomas, and L. Tong, "Malicious data attacks on the smart grid," *IEEE Trans. Smart Grid*, vol. 2, no. 4, pp. 645–658, Dec. 2011.
- [3] H. Sandberg, A. Teixeira, and K. H. Johansson, "On security indices for state estimators in power networks," in *Proc. 1st Workshop Secure Control Syst. CPSWEEK*, Stockholm, Sweden, Apr. 2010, pp. 1–6.
- [4] 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 Smart Grid Commun.*, Gaithersburg, MD, USA, Oct. 2010, pp. 220–225.
- [5] R. B. Bobba *et al.*, "Detecting false data injection attacks on dc state estimation," in *Proc. 1st Workshop Secure Control Syst. CPSWEEK*, Stockholm, Sweden, Apr. 2010, pp. 1–9.
- [6] G. R. Krumpholz, K. A. Clements, and P. W. Davis, "Power system observability: A practical algorithm using network topology," *IEEE Trans. Power App. Syst.*, vol. PAS-99, no. 4, pp. 1534–1542, Jul. 1980.
- [7] A. Monticelli and F. F. Wu, "Network observability: Theory," *IEEE Trans. Power App. Syst.*, vol. PAS-104, no. 5, pp. 1042–1048, May 1985.
- [8] A. Monticelli and F. F. Wu, "Network observability: Identification of observable islands and measurement placement," *IEEE Trans. Power App. Syst.*, vol. PAS-104, no. 5, pp. 1035–1041, May 1985.
- [9] T. Kim and H. Poor, "Strategic protection against data injection attacks on power grids," *IEEE Trans. Smart Grid*, vol. 2, no. 2, pp. 326–333, Jun. 2011.
- [10] S. Bi and Y. Zhang, "Defending mechanisms against false-data injection attacks in the power system state estimation," in *Proc. IEEE GLOBECOM Workshops*, Houston, TX, USA, Dec. 2011, pp. 1162–1167.
- [11] A. Giani *et al.*, "Smart grid data integrity attacks: Characterizations and countermeasures," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Oct. 2011, pp. 232–237.
- [12] J. Kim and L. Tong, "On topology attack of a smart grid: Undetectable attacks and countermeasures," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 7, pp. 1294–1305, Jul. 2013.
- [13] K. Clements and P. Davis, "Multiple bad data detectability and identifiability: A geometric approach," *IEEE Trans. Power Del.*, vol. PWRD-1, no. 3, pp. 355–360, Jul. 1986.

- [14] G. Korres and G. Contaxis, "Identification and updating of minimally dependent sets of measurements in state estimation," *IEEE Trans. Power Syst.*, vol. 6, no. 3, pp. 999–1005, Aug. 1991.
- [15] G. Dán and H. Sandberg, "Stealth attacks and protection schemes for state estimators in power systems," in *Proc. IEEE Smart Grid Commun.*, Gaithersburg, MD, USA, Oct. 2010, pp. 214–219.
- [16] O. Vukovic, K. C. Sou, G. Dan, and H. Sandberg, "Network-layer protection schemes against stealth attacks on state estimators in power systems," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Oct. 2011, pp. 184–189.
- [17] E. Handschin, F. C. Schweppe, J. Kohlas, and A. Fiechter, "Bad data analysis for power system state estimation," *IEEE Trans. Power App. Syst.*, vol. PAS-94, no. 2, pp. 329–337, Mar. 1975.
- [18] A. Monticelli and A. Garcia, "Reliable bad data processing for real-time state estimation," *IEEE Trans. Power App. Syst.*, vol. PAS-102, no. 5, pp. 1126–1139, May 1983.
- [19] T. Van Cutsem, M. Ribbens-Pavella, and L. Mili, "Bad data identification methods in power system state estimation—A comparative study," *IEEE Trans. Power App. Syst.*, vol. PAS-104, no. 11, pp. 3037–3049, Nov. 1985.
- [20] L. Mili, T. Van Cutsem, and M. Ribbens-Pavella, "Hypothesis testing identification: A new method for bad data analysis in power system state estimation," *IEEE Trans. Power App. Syst.*, vol. PAS-103, no. 11, pp. 3239–3252, Nov. 1984.
- [21] A. Abur and A. G. Expósito, *Power System State Estimation: Theory and Implementation*. Boca Raton, FL, USA: CRC Press, 2000.
- [22] A. Monticelli, *State Estimation in Electric Power Systems: A Generalized Approach (Power Electronics and Power Systems)*. New York, NY, USA: Springer-Verlag, 1999.
- [23] W. W. Kotiuga and M. Vidyasagar, "Bad data rejection properties of weighted least absolute value techniques applied to static state estimation," *IEEE Trans. Power App. Syst.*, vol. PAS-101, no. 4, pp. 844–853, Apr. 1982.
- [24] A. Abur and M. Celik, "A fast algorithm for the weighted least absolute value state estimation [for power systems]," *IEEE Trans. Power Syst.*, vol. 6, no. 1, pp. 1–8, Feb. 1991.
- [25] M. Celik and A. Abur, "A robust WLAV state estimator using transformations," *IEEE Trans. Power Syst.*, vol. 7, no. 1, pp. 106–113, Feb. 1992.
- [26] H. Singh and F. Alvarado, "Weighted least absolute value state estimation using interior point methods," *IEEE Trans. Power Syst.*, vol. 9, no. 3, pp. 1478–1484, Aug. 1994.
- [27] L. Mili, M. Cheniae, and P. Rousseeuw, "Robust state estimation of electric power systems," *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.*, vol. 41, no. 5, pp. 349–358, May 1994.
- [28] M. Cheniae, L. Mili, and P. Rousseeuw, "Identification of multiple interacting bad data via power system decomposition," *IEEE Trans. Power Syst.*, vol. 11, no. 3, pp. 1555–1563, Aug. 1996.
- [29] K. Morrow, E. Heine, K. Rogers, R. Bobba, and T. Overbye, "Topology perturbation for detecting malicious data injection," in *Proc. 45th HICSS*, Jan. 2012, pp. 2104–2113.
- [30] A. Tajer, S. Kar, H. Poor, and S. Cui, "Distributed joint cyber attack detection and state recovery in smart grids," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Oct. 2011, pp. 202–207.
- [31] Y. Huang, H. Li, K. Campbell, and Z. Han, "Defending false data injection attack on smart grid network using adaptive CUSUM test," in *Proc. 45th Annu. CISS*, Mar. 2011, pp. 1–6.
- [32] S. Cui *et al.*, "Coordinated data-injection attack and detection in the smart grid: A detailed look at enriching detection solutions," *IEEE Signal Process. Mag.*, vol. 29, no. 5, pp. 106–115, Sep. 2012.
- [33] Y. Huang *et al.*, "Bad data injection in smart grid: Attack and defense mechanisms," *IEEE Commun. Mag.*, vol. 51, no. 1, pp. 27–33, Jan. 2013.
- [34] *Accuracy of Digital Electricity Meters*, Electric Power Research Institute, Palo Alto, CA, USA, May 2010, white paper.
- [35] D. R. Karger and C. Stein, "A new approach to the minimum cut problem," *J. ACM*, vol. 43, no. 4, pp. 601–640, Jul. 1996.
- [36] Power Systems Test Case Archive. [Online]. Available: <http://www.ee.washington.edu/research/pstca/>



Jinsub Kim received the Ph.D. degree in electrical and computer engineering with minors in applied mathematics and statistics from Cornell University, Ithaca, NY, USA. He is currently a Postdoctoral Associate with Cornell University. His research interest spans statistical inference and optimization for a smart grid and network security. He has studied data falsification attacks on a power network that may cause an incorrect topology or state estimate. His graduate study was partially supported through Samsung Scholarship.



Lang Tong (S'87–M'91–SM'01–F'05) received the B.E. degree in automation from Tsinghua University, Beijing, China, in 1985 and the M.S. and Ph.D. degrees in electrical engineering from the University of Notre Dame, Notre Dame, IN, USA, in 1987 and 1991, respectively.

He is currently the Irwin and Joan Jacobs Professor in Engineering at Cornell University, Ithaca, NY, USA, where he is also the Site Director of the Power Systems Engineering Research Center. He was a Postdoctoral Research Affiliate at the Information Systems Laboratory, Stanford University, in 1991. He was the 2001 Cor Wit Visiting Professor at Delft University of Technology and had held visiting positions at Stanford University and the University of California at Berkeley. His current research focuses on statistical inference, optimization, and economic problems in energy systems.

Prof. Tong was a Distinguished Lecturer of the IEEE Signal Processing Society. He received the 1993 Outstanding Young Author Award from the IEEE Circuits and Systems Society, the 2004 Best Paper Award from the IEEE Signal Processing Society, and the 2004 Leonard G. Abraham Prize Paper Award from the IEEE Communications Society. He is also a coauthor of seven student paper awards. He received the Young Investigator Award from the Office of Naval Research.

He is also a coauthor of seven student paper awards. He received the Young Investigator Award from the Office of Naval Research.



Robert J. Thomas (M'73–SM'83–F'93–LF'08) currently holds the position of Professor Emeritus of Electrical and Computer Engineering at Cornell University, Ithaca, NY, USA. He has had assignments with the Office of Electric Energy Systems, U.S. Department of Energy (USDOE), Washington, DC, USA, and with the National Science Foundation as the first Program Director for the Power Systems Program at the Engineering Directorate's Division of Electrical Systems Engineering. He is the author of more than 100 technical papers and two book chapters.

He has published in the areas of transient control and voltage collapse problems, as well as technical, economic, and institutional impacts of restructuring. Prof. Thomas is a member of Tau Beta Pi, Eta Kappa Nu, Sigma Xi, and ASEE; an ACSF Fellow of Cornell University. He is the Founding Director of the 13-university-member National Science Foundation Center, Power Systems Engineering Research Center. He was a member of the USDOE Secretary's Power Outage Study Team and is a Founding Member of the Coalition for Electric Reliability Solutions. He has received five teaching awards and the IEEE Centennial and Millennium medals.