Integrity Attacks on Real-Time Pricing in Smart Grids: Impact and Countermeasures

Jairo Giraldo, *Member, IEEE*, Alvaro Cárdenas, *Member, IEEE*, and Nicanor Quijano, *Senior Member, IEEE*

***Abstract*—Recent work has studied the impact caused by attackers that compromise pricing signals used in the emerging retail electricity market and send false prices to a subset of consumers. In this paper, we extend previous work by considering a more realistic adversary model that is not arbitrarily tied to scaling and delay attacks, but that can generate any arbitrary pricing signal and show how to keep the problem tractable with a new analysis based on sensitivity functions. In addition, we extend previous work by proposing countermeasures to mitigate the negative impact of these attacks. Countermeasures include selecting parameters of the controller, designing robust control algorithms, and by detecting anomalies in the behavior of the system.**

***Index Terms*—Real-time pricing, power systems security, control systems, attack detection, state estimation.**

I. Introduction

TO MAINTAIN a balance between optimizing the use of resources and the real-time control requirements for keeping the frequency and voltage of the power grid at their design levels, the power grid uses a daily and hourly scheduling of generation units to match the forecast electricity load via wholesale market transactions. A scheduling coordinator solicits generation through some form of auction where lowest bidders generate electricity and this in turn creates an economically optimal schedule of generators. In contrast to these traditional wholesale markets (e.g., between generation utilities and distribution utilities), many retail markets (e.g., between a distribution utility and an industry consumer of electricity) have traditionally adopted *static pricing schemes such as fixed and time-of-use tariffs, under which consumers have limited incentives to adapt their electricity consumption to market conditions. This lack of incentives results in high peak demands that strain infrastructure capacities and unnecessarily increase operational costs* [1]. This approach is inefficient, since the system infrastructure used to guarantee supply under peak hours is not completely used most of the time. According to the U.S. Department of Energy, 10% of the whole generating capacity and 25% of distribution capacity is used less than the 5% of the time.

In an effort to increase the efficiency of the power grid, many retail-markets are expanding the use of demand-response programs. In their basic form, demand-response programs are a control problem where the control signal are the incentives (e.g., real-time pricing), or direct-load control (e.g., the utility directly controlling the set-points of air conditioning systems in specific cases) for consumers to reduce electricity consumption during peak hours and to shift this load to off-peak hours. Currently most of the electricity consumers leveraging demand-response programs are large commercial consumers, but the market is expanding more and more to smaller industries and even residential consumers. As the number of smart devices necessary to manage this market expands, the potential attack surface of the market also increases, and therefore we need to begin considering the potential impact of attackers that compromise devices and communication channels used in this market. Clearly, any ISO understands that the pricing information is a core asset on its system and multiple layers of defense may be applied (e.g., cryptography methods); however, it is important to assume the worst case scenario where an attacker possesses the resources to go through the information technology (I.T) defense mechanisms and modify some of the transmitted information.

The security of demand response algorithms with real-time electricity pricing was recently explored by Tan *et al.* [1]. In their work, they consider an attacker that has compromised a portion of the communication channels used to send price information to consumers, and then study the effects to the power system from *scaling* and *delay* attacks, where the prices advertised to smart meters are compromised by a scaling factor (so consumers use the wrong prices) and by corrupted timing information (so consumers use old prices). While this previous work is an important step for initiating the discussion on how to analyze the impact of attacks on real-time pricing, this research has limitations on the way it modeled the adversary by limiting attacks to scaling and delays. In addition this previous work did not discuss any security countermeasures against attacks.

In this paper we extend the work of Tan *et al.* [1] in several directions:

- We model a more realistic attacker that can inject an arbitrary modification to the price received by the consumer,

Manuscript received July 9, 2015; revised October 29, 2015; accepted January 12, 2016. This work was supported in part by the National Institute of Standards and Technology (NIST) through the U.S. Department of Commerce under award 70NANB14H236. The work of J. Giraldo was supported by Colciencias under Grant 567. Paper no. TSG-00797-2015.

J. Giraldo was with the Universidad de los Andes, Bogotá 111711, Colombia. He is now with the Computer Science Department, University of Texas in Dallas, Dallas, TX 75080-3021 USA (e-mail: jairo.giraldo@utdallas.edu).

A. Cárdenas is with the Computer Science Department, University of Texas in Dallas, Dallas, TX 75080-3021 USA (e-mail: alvaro.cardenas@utdallas.edu).

N. Quijano is with the Departamento de Ingeniería Eléctrica y Electrónica, Universidad de los Andes, Bogotá 111711, Colombia (e-mail: nquijano@uniandes.edu.co).

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/TSG.2016.2521339

and is not constrained to scaling or delay attacks. The attacker aims to increase the difference between the generated and the consumed power by injecting small changes in the price signal.
- We use sensitivity analysis to quantify the impact of an attack. By using sensitivity analysis we can identify the attack signals that will be amplified and the ones that will be attenuated by the control loop, enabling us to determine how successful a given attack signal is.
- We propose countermeasures based on changing the parameter of the original controller by Tan *et al.* [1]. In addition, we propose an estimator and a new robust-control design that estimates the perturbation and computes a new price to attenuate the error between supply and demand caused by the attacker.
- We propose an attack-detection algorithm based on the CUSUM technique and evaluate its effectiveness to identify attacks for different controller parameters, and different attack frequencies. We are able to identify the trade-off between time of detection, frequency of the attack, and number of false alarms. Moreover, it is possible to define an attack that cannot be detected, but whose effect in the network is low due to the proposed detection method.
- All our results can be extended to linear (or linearized) feedback control systems.

II. Related Work: Impact of Integrity Attacks in the Power Grid

Our work falls within the scope of integrity attacks (or false-data injection attacks) to the sensor or control signals of a cyber-physical system. Integrity attacks have been proposed as a way to analyze the vulnerability of cyber-physical systems in general and the power grid in particular. Injecting false data to state estimation algorithms used in bulk of the power grid was first proposed by Liu *et al.* [2], and similar integrity attacks were proposed for compromised smart meters trying to defraud the electric utility [3].

The work on integrity attacks against bad data detectors for state estimation in the power grid has generated a significant body of follow up work; for example Dán and Sandberg [4], consider a defender that can secure individual sensor measurements by, for example, replacing an existing meter with another meter with better security mechanisms such as tamper resistance or hardware security support. Mo and Sinopoli [5] also extend the basic false data injection attack to consider attackers trying to maximize the error introduced in the estimate, and defenders with a new detection algorithm that attempts to detect false data injection attacks. Similar false-data injection attacks have been considered for specific devices in the power grid, such as integrity attacks against the Flexible Alternate Current Transmission System (FACTS) [6], [7], and Automatic Generator Control (AGC) [8]–[10]. All this related work has targeted operational data of the power grid, and is not related to electricity markets.

Negrete-Pincetic *et al.* [11] were one of the first to study how false control signals can affect the social welfare of the electricity market. Related work by Xie *et al.* [12] studied how false data injection attacks can be used to defraud bulk electricity markets by modifying Locational Marginal Prices (LMPs), and work by Jia *et al.* [13] studied how false meter data in the bulk of the power grid can be used to cause the largest errors in LMP estimation.

These integrity attacks have been studied in the bulk electricity market and specifically, the estimation problem alone; most previous work does not consider how the control algorithm can be designed to minimize the impact of integrity attacks, or studied the feedback control loop behavior of the system under attack [14].

III. Preliminaries

A. Demand Response Model

We follow the real-time pricing model from Tan *et al.* [1]. This model considers a market with consumers of electricity, suppliers of electricity, and a third party entity—an Independent System Operator (ISO)—with the goal of matching supply and demand by setting the market price for electricity. The general assumption is that the ISO determines, at each time instant $k \in \mathbb{N}_+$, a clearing price λ_k valid for the period of time $[k \cdot T, (k+1) \cdot T]$ (this is called an *ex-ante* market) every T hours (e.g., T=0.5h) and announces it to the suppliers and consumers.

The electricity demand is characterized by two components: a baseline electricity consumption b_k that captures the electricity consumption that is independent of the pricing mechanism, and a price-responsive demand $w(\lambda_k)$, which captures the amount of electricity consumption that can be controlled by the pricing signal λ_k.

The aggregated demand of all consumers is $d_k(\lambda_k) = b_k + w(\lambda_k)$. b_k can be considered as the necessary power to satisfy the main consumer needs at each instant k (e.g., refrigerator, cooking devices, light bulbs). $w_k \geq 0$ is then the amount of power that can be consumed depending on the price. For instance, doing laundry when the price is low, or turning off the lights of rooms that are not being used. If there is no real-time pricing, $d_k = b_k$.

As b_k is unknown, for simulation purposes it can be obtained from historical demand curves such as those from the New York ISO [15]. These demand curves have historical consumption traces from a specific population taken at 5 minute intervals and they can be used to predict future demand profiles. The Constant Elasticity of Own-price (CEO) has been commonly adopted to characterize the total price-responsive demand [16], [17]. The CEO model is defined by

$$w(\lambda_k) = D\lambda_k^{\epsilon} \tag{1}$$

where $D > 0$ is a constant that properly scales w_k and $\epsilon \in (-1, 0)$ is the *price elasticity demand* that captures how the demand is affected by a specific price λ_k [17].

Similarly, for the supply of electricity, Tan *et al.* [1], propose a linear regression between supply and cost, a model they validated from the Australian Energy Market Operator and the electricity market in California. Under these assumptions the

supply of electricity can be modeled by the following equation:

$$s(\lambda_k) = p\lambda_k + q, \tag{2}$$

where p and q are parameters estimated by the historical market data from the area of study.

B. Control Objective

The control objective of the ISO is to send price signals λ_k to keep the error between supply and demand of electric power $\mathcal{E}_k = s(\lambda_k) - d(k, \lambda_k)$ close to zero for every time instant k. This can be seen as a control problem in which the system to be controlled is the outcome of a market, the control variable is the price signal λ_k and the variable that can be measured is the error $\mathcal{E}_k$.

The price signal λ_k must be carefully designed because a direct feedback of the wholesale prices to the users might cause oscillations or even instability [1], [18].

C. Transfer Function Representation

Transfer functions are a mathematical representation of linear difference (or differential) equations that allow us to represent the system in a compact way and to evaluate the performance of the system in terms of the frequency components of the control signals—recall that every time series have an equivalent representation (a one to one mapping) to a function in the frequency domain given by the Fourier transform.

For our discrete-time system (where sensor and control actions are taken at given time steps k separated by the sampling period T (e.g., 30 minutes), the transfer function for the equations modeling the dynamics of the system can be obtained by using the z-transform (a transform similar to the Fourier transform).

In particular, we can define the transfer function of the price stabilization algorithm, the system, and the observation mechanism as $G_c(z)$, $G_p(z)$, and $H(z)$, respectively.

To express these transfer functions it is necessary to approximate the dynamics system at the operation point λ_0 to a linear system. Hence, following Tan *et al.* [1] we make the following approximations with the Taylor polynomials of the supply $s()$ and demand $w()$:

$$s(\lambda) \simeq \dot{s}(\lambda_0)(\lambda - \lambda_0) + s(\lambda_0) = \dot{s}(\lambda_0)\lambda + s_0$$
$$w(\lambda) \simeq \dot{w}(\lambda_0)(\lambda - \lambda_0) + w(\lambda_0) = \dot{w}(\lambda_0)\lambda + w_0$$

where $\dot{f} = \frac{df}{d\lambda}$, and where we define the constant (or endogenous) terms with $s_0 = s(\lambda_0) - \lambda_0\dot{s}(\lambda_0)$ and $w_0 = w(\lambda_0) - \lambda_0\dot{w}(\lambda_0)$. Therefore, the transfer functions can be defined as $G_s(z) = \dot{s}(\lambda_0) = p$, with initial condition s_0 and $G_w(z) = \dot{w}(\lambda_0) = D\epsilon(\lambda_0)^{\epsilon-1}$, with initial condition w_0. The outcome of the market can be expressed as $G_p(z) = G_s(z) - G_w(z) = \dot{s}(\lambda_0) - \dot{w}(\lambda_0)$ and the z-transform of the price signal is $\Lambda(z)$. Clearly, $s(z) = G_s(z)\Lambda(z) + s_0$ and $w(z) = G_w(z)\Lambda(z) + w_0$.

D. Control Algorithm for Setting Prices

The price setting control algorithm depends on the previous price λ_{k-1} and the previous observed error at $\mathcal{E}_{k-1}$, using a one step delay transfer $H(z) = z^{-1}$. If $\mathcal{E}_k$ is negative, it means that there was more power demanded than supplied, and thus the price will increase (to motivate consumers to decrease consumption), while if $\mathcal{E}_k$ is positive, then the price will decrease. The precise amount of increase and decrease of the prices at each time step should be selected carefully as inadequate price updates can make the system unstable. Tan *et al.* [1] found that when we design a proportional gain $\eta \in (0, 1)$ in the following price-setting algorithm:

$$\lambda_k = \lambda_{k-1} - \frac{2\eta}{\dot{s}(\lambda_0) - \dot{w}(\lambda_0)}\mathcal{E}_{k-1},$$

the system will remain stable. The transfer function representation of the controller is represented as

$$G_c(z) = \frac{2\eta}{\dot{s}(\lambda_0) - \dot{w}(\lambda_0)}\frac{1}{1 - z^{-1}}.$$

Note that the z-transform representation of the price is then $\Lambda(z) = -G_c(z)\mathcal{E}(z)$, and the supply-demand mismatch is $\mathcal{E}(z) = G_p\Lambda(z)$. Combining both expressions yields to the characteristic transfer function when there is no attack, which is given by $T_c(z) = \frac{2\eta}{z-1+2\eta}$ with a pole at $z = 1 - 2\eta$ (Fig. 1 depicts the block diagram of the suppliers, consumers $G_w(z) = G_w^1(z) + G_w^2$, and the price control strategy). Note that η is in fact an important design parameter for the control algorithm that affects the convergence rate of the price. For instance, for $\eta = 0.5$, the pole is at the origin of the z-plane and the system converges faster. As we will show, η can also determine the resiliency of the system under attacks. When properly selected, it can also attenuate the impact of attacks.

IV. Attacker Model

In contrast to one-shot attacks, where the attacker provides false information only once [2], [19], in this work we consider that an attacker compromises a device or a communication channel, and has the capability to add false information at any moment and—more importantly—repeatedly over a long period of time.

For example, most of the work on false data injection in state estimation finds a value d^a to insert at an arbitrary point in time [2], however, this previous work does not consider the evolution of the system dynamics over time. In this context, the question we would like to pose from an adversarial point of view is the following:

- What is the worst attack time series d_k^a that can affect the system while keeping some bounds (prices will be bound by some maximum and minimum values: $\forall k\ d_k^a \in [d_{min}^a, d_{max}^a]$.

Tan *et al.* [1] proposed an adversary model where one attacker compromised the pricing communication channel between the ISO and a percentage ρ of consumers. They considered delay attacks and scaling attacks.

In a delay attack, the compromised price is an old version of the price, i.e., $\hat{\lambda}_k = \lambda_{k-\tau}$, and in a scaling attack, the compromised price is a scaled version of the true price, i.e., $\hat{\lambda}_k = \gamma\lambda_k$.

While the attacks defined above can be easily analyzed from a theoretical point of view, it is not clear why an attacker who has compromised a communications channel will select

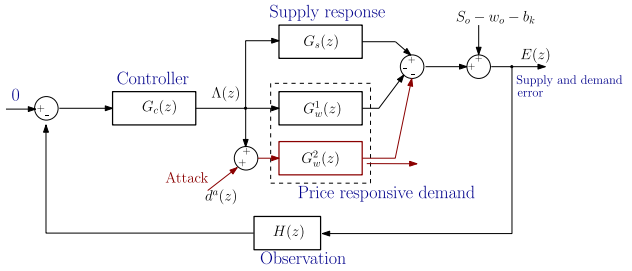

Fig. 1. Block diagram of the real-time pricing model under attack.

to launch these attacks when she has the flexibility of sending any arbitrary time series $\hat{\lambda}_k$ she wants, even one that bears no resemblance to the original time series λ_k.

Furthermore, scaling attacks and delay attacks are not strategic, and do not seek to maximize any objective function from the adversary. In this work we follow the generic and more powerful adversary model introduced by the false data injection paper [2], and we expand it to consider a time series. In particular, we model a compromised communication channel as $\hat{\lambda}_k = \lambda_k + d_k^a$, where $\hat{\lambda}_k$ is the price information received by the victim, and d_k^a is an arbitrary time signal that can take any value. It is clear now that scaling attacks and delay attacks are simple subsets of this new attack because for every scaling or delay attack possible producing a false price information $\hat{\lambda}_k$, there exists an arbitrary time signal d_k^a that will produce the same price $\hat{\lambda}_k$ received by the victim.

The question we now face is how to determine a strategic attack time series d_k^a that will try to cause as much damage as possible. In our case, the main objective of the attacker is to try to maximize the mismatch between power generated and power consumed. This mismatch can lead to over-generation, economical losses, and unstable behavior of the system, by intelligently inducing small changes in the price signal. One of our key insights into tackling this problem is the fact that for every time series, there is a one-to-one correspondence of the time series and its frequency (Fourier transform) representation. Therefore, instead of attempting to analyze the worst time series d_k^a in time, *we identify the worst-possible attacks in frequency space*.

In order to provide a mathematical tool that enables us to quantify the impact of the attack, we use sensitivity analysis. Sensitivity functions have been widely used to analyze the impact of external disturbances or parameter changes on the output of a feedback system. In systems and control theory, it is well known that feedback can attenuate or amplify disturbances; therefore, using the frequency representation of the system (the transfer function), it is possible to obtain the sensitivity function and observe the response of the system to a perturbation of a specific frequency ω [20].

In this work we focus our attention on additive attacks in the price information; however, our approach can be easily extended to analyze attacks over sensors.

In the next section we give the formal incorporation of the attacks against pricing signal, and in the section after that we use sensitivity analysis to identify the impact of the attacks.

A. Incorporating the Attack Into the Real-Time Pricing Model

One of the main assumptions in this work is that the attacker does not need to have complete knowledge of the system; however, if she does know how consumers and suppliers react with price changes, she can try to apply the worst attack. Our goal is to identify how modifying the prices with different frequencies, causes different effects in the system performance, and it does not require that the adversary possesses very specific knowledge of the system.

We assume that an amount ρ of communication channels are compromised, and each of these consumers receives the price value $\hat{\lambda}_k = \lambda_k + d_k^a$, where $d^a \in \mathbb{R}$ corresponds to the additional false information.

It is necessary to identify how the inclusion of this attack affects the system representation of the real-time problem. In particular, we need to identify how the attack changes the transfer functions of the model (i.e., we need to characterize the new transfer functions $G_w^1(z)$ for the consumers who are unaffected, and $G_w^2(z)$ for the consumers who receive false information, as shown in Fig. 1.).

Let us consider the price response demand based on the CEO model for the set of compromised nodes $\rho w_k(\lambda_k, d_k^a) = \rho D(\lambda_k + d_k^a)^\epsilon$. In order to linearize this model it is necessary to assume that $|d_k| << \lambda_k$ and $\lambda_k > 0$. As we will discuss towards the end of the paper (the attack-detection formulation), this is a perfect assumption for an attacker that wants to minimize its chances of being detected (by causing small changes to the price $|d_k| << \lambda_k$) but at the same time wants to find the best way to find a small signal deviation that will maximize the potential damage to the system.

The linearized model is described by:

$$w(\hat{\lambda}_k) = \rho w(\lambda_o + d_o^a) + \rho \dot{w}(\lambda_o + d_o^a)\big(\lambda_k + d_k^a - \lambda_o - d_o^a\big) + (1-\rho)(w(\lambda_o) + \dot{w}(\lambda_o)(\lambda_k - \lambda_o))$$

We can group the price-independent terms with b_k (the baseline consumption of electricity that is independent of the price), and then also group the price-dependent components for the transfer functions.

$G_w^1(z) = (1-\rho)\dot{w}(\lambda_o)$ corresponds to the transfer function of consumers who receive unmodified price information, and $G_w^2(z) = \rho\dot{w}(\lambda_o + d_o)$, corresponds to the transfer function of the victims. Under the assumption that $|d_k| << \lambda_k$, we can neglect the term d_o in the linearization, such that $G_w^2(z) = \rho\dot{w}(\lambda_o)$.

V. Sensitivity Analysis

The sensitivity function models how one input to the system (in our case the attack) affects another signal in the system (we are mostly interested to see how the attack affects the error in power generated minus the demand, and to also see the impact on the prices).

We start by looking at the impact that a disturbance $d^a(z)$ (in the frequency space) can have on the error $\mathcal{E}(z)$. In particular, the sensitivity function for these two time series

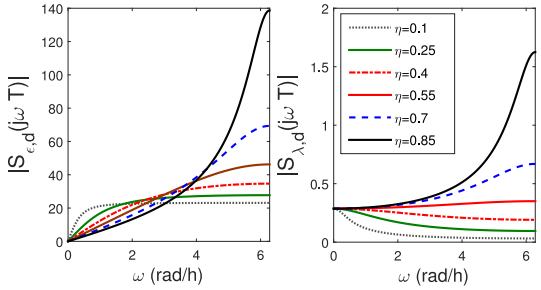

Fig. 2. Left: Sensitivity of the error $\mathcal{E}(z)$. Right: Sensitivity of the Price $\lambda(z)$. This sensitivity analysis uses parameters: $\rho = 0.5, p = 31, q = 917$, and $T = 0.5$h. The baseline consumption is $b = 400$ MW, which is proportional to 1 million households, and the base demand of each consumer is $b_i \in [2.8, 4.6]KW$.

(denoted as $S_{\mathcal{E},d}$) is the ratio $\mathcal{E}(z)/d^a(z)$:

$$S_{\mathcal{E},d} = -\frac{G_w^2(z)}{1 + G_c(z)H(z)G_p(z)} = -\frac{\rho\dot{w}(\lambda_0)(z-1)}{(z-1+2\eta)}. \quad (3)$$

As stated before, our interest is to analyze the effects of an additive attack in the frequency domain. We denote the angular frequency as ω. We then replace $z = e^{j\omega T}$ for T being the sampling period (the time interval between updating the sensor measurements and the prices). It is important to notice that the maximum frequency that an attacker can generate is limited by the sampling period, such that $\omega_{max} = \pi/T$. For instance, if the sampling period is $T = 0.5$ hours, then $\omega_{max} = 2\pi$.

From this equation we can see that the percentage of compromised channels ρ has a scaling effect on the sensitivity of the system. Moreover, the selection of the control parameter η proposed by Tan *et al.* [1] is fundamental for attenuating the effects of the attack. The left side of Fig. 2 shows how the attack can be amplified (or attenuated) as a function of the frequency of the attack signal. Clearly, the impact the supply-demand mismatch $\mathcal{E}$ is severe for most frequencies; however, we can also see how the control parameter η can be selected to attenuate the impact of high-frequency signals: smaller values of η will minimize the impact of high-frequency components of the attack time-series—this comes at the cost of a slower control action which might not be a bad idea, as changes in prices will remain small, giving consumers more predictability in their electricity consumption habits.

Recall that if the output $\mathcal{E}$ is different from zero, then there is over demand or over production of electricity, which can affect considerably the system (resulting in large frequency changes). Even if the price variations are small, the output amplifies the disturbance. There is a trade off between the η, ρ, and the frequency of the disturbance. An attacker can easily take advantage of this fact, and introduce intelligently false data to a portion of the users. This information can be of small amplitude, and hardly detected; however, the effects on the output can be catastrophic.

We can also obtain the sensitivity function with respect to the price. This function reveals how the attack modifies the real price calculated by the ISO. The function is described by

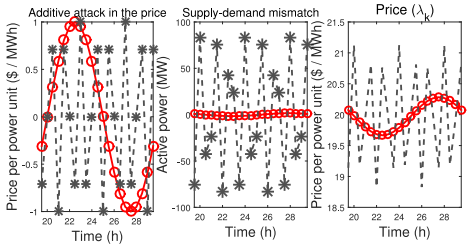

Fig. 3. Effects in the supply-demand mismatch (middle) and the price (right) for two attacks i) $d_k^a = \sin((3/2)\pi Tk)$ and ii)$d_k^a = \sin(\pi Tk/5)$ for $\eta = 0.8$.

$$S_{\lambda,d}(z) = -\frac{G_c(z)G_w^2(z)H(z)}{1 + G_c(z)H(z)G_p(z)} = -\frac{2\eta\rho\dot{w}(\lambda_0)}{(\dot{s}(\lambda_0) - \dot{w}(\lambda_0))(z-1+2\eta)}, \quad (4)$$

The left side of Fig. 2 shows the sensitivity function with respect to the price for different values of η, and $\rho = 0.5$. With this selection of ρ, the real price changes produced by the attack are attenuated for mostly of all η.

Now that we have gained some insight into how the "frequency components" of a time series can affect the system, we look at the "time domain" to apply these lessons in the analysis of attacks. In the foregoing of this work, we will consider sinusoidal attacks in order to illustrate the effects of attacks with an specific frequency. However, the analysis applies for any signal, from constant (i.e., $\omega = 0$) to random attacks.

As an example, Fig. 3 shows a high-frequency attack (black) and a low-frequency attack (red) on the left. The control algorithm is using $\eta = 0.8$. We can see a large error magnification caused by this control parameter (as predicted by Fig. 2). Similarly, the price signal is also amplified for the high frequency attack (as can be seen by the figure on the right).

VI. Designing an Attack-Resilient Controller

Previous work only studied the effects of the attack, but did not propose new control mechanisms to mitigate possible attacks. We know discuss how we can start designing attack-resilient controllers.

In order to design an attack-resilient controller, we can leverage the fact that the ISO has historical data showing the behavior of the system which can be used for learning the dynamics (parameters) of the system. Whenever the controller commands do not have the expected effect, or when the sensor signals do not reflect the normal evolution of the system we can try to identify these problems and design a controller that minimizes the impact of price or sensor attacks.

As the attack are unknown inputs into the system, we can use a type of disturbance estimators. Disturbance observers have been studied in literature but we focus our attention in the one introduced by Kim and Rew [21] for discrete-time systems.

We assume that the ISO possesses the information about the supply-demand error $\mathcal{E}_{k-1}$ and we try to detect an attack using the observer (an observer is another name for a "state estimator").

We first present the attack-resilient controller for a general discrete-time system, and in the next section we show how to apply it to our real-time pricing model.

Let us consider a generic linear discrete-time system for a sampling period $T > 0$ of the form

$$x_{k+1} = Ax_k + Bu_k + \Gamma d_k, y_k = Cx_k \tag{5}$$

where $x_k \in \mathbb{R}^n$, $u_k \in \mathbb{R}^m$, $d_k \in \mathbb{R}^q$, and $y_k \in \mathbb{R}^l$ are the state variable, the control input, the disturbance, and the measurement output, respectively. The matrices A, B, Γ, C are of adequate dimension.

For $d_k = (d_k^1, \dots d_k^q)$, the disturbance is slowly time-varying, such that $d_{k+1}^i - d_k^i < T\mu_i$, $\forall i = 1, \dots, q$. Given a $K \in \mathbb{R}^{q\times n}$ and $C = I_n$, the observer is described as follows

$$\begin{aligned} z_{k+1} &= z_k + K\Big((A - I_n)x_k + Bu_k + \Gamma\hat{d}_k\Big) \\ \hat{d}_k &= Kx_k - z_k \end{aligned} \tag{6}$$

Under the assumption that Γ is invertible, we can choose $K = (I_q - \Phi)\Gamma^{-1}$ for $\Phi = [\phi_1, \dots, \phi_q]^\top$, and $\phi_i \in (-1, 1)$. The estimation error $e_k = d_k - \hat{d}_k$ is then bounded by $e_\infty = \frac{T\mu_i}{1-|\phi_i|}$ for $\phi_i \in (0, 1)$, and $\mu_i > 0$.

A. Estimation of Price Attacks

Let $G_p = \dot{s}(\lambda_0) - \dot{w}(\lambda_0)$ and $d_k = d_k^a$ to simplify notation. We can write the feedback real-time pricing problem using a discrete-time state space representation as follows

$$\mathcal{E}_k = G_p u_k - \rho\dot{w}(\lambda_0)d_k \tag{7}$$

Note that comparing (7) with (5), we have $A = 0$, $B = G_p$, $\Gamma = -\rho\dot{w}(\lambda_0)$, $x_{k+1} = \mathcal{E}_k$ and $u_k = \lambda_k$.

Note that to compute the state estimation, it is necessary to know Γ, which means that we would need prior knowledge about the amount of compromised nodes. Obviously, this requirement seems unrealistic as ρ will remain unknown to the defender; however, we can exploit a very interesting property of the estimator we found to perform state estimation without knowing ρ, as stated in the following proposition.

Proposition 1: Let us consider the disturbance estimator described in (6) for the real-time pricing model in (7). The rate of change of the disturbance $\Delta d_k = d_k - d_{k-1}$ is bounded such that $|\Delta d_k| \le T\mu$ for some constant μ and T the sampling period. We define $\hat{\Gamma}$ as an approximate value of Γ and $\hat{e}_k = \Gamma d_k - \hat{\Gamma}\hat{d}_k$ as an error between the real effect of the disturbance and its estimate. If $K = \hat{\Gamma}^{-1}(1-\phi)$ for $\phi \in (-1, 1)$, the error converges and is bounded by

$$|\hat{e}_\infty| \le \frac{|\Gamma|T\mu}{1 - |\phi|}. \tag{8}$$

Proof: The error evolution is

$$\begin{aligned} \hat{e}_{k+1} &= \Gamma d_{k+1} - \hat{\Gamma}\hat{d}_{k+1} \\ &= \Gamma d_{k+1} - \hat{\Gamma}K(G_p u_k + \Gamma d_k) \\ &\quad + \hat{\Gamma}(z_k - Kx_k + KG_p u_k + K\hat{\Gamma}\hat{d}_k) \\ &= \Gamma d_{k+1} - \hat{\Gamma}K\Gamma d_k - \hat{\Gamma}d_k + \hat{\Gamma}K\hat{\Gamma}\hat{d}_k \\ &= \Gamma\Delta d_{k+1} + (1 - \hat{\Gamma}K)\hat{e}_k \end{aligned}$$

As $K = (1-\phi)/\hat{\Gamma}$, in the equilibrium when $\hat{e}_{k+1} = \hat{e}_k$, $\hat{e}_\infty$ is bounded by

$$|\hat{e}_\infty| = \frac{|\Gamma|\Delta d_{k+1}}{1 - |\phi|} \le \frac{|\Gamma|T\mu}{1 - |\phi|}. \qquad \blacksquare$$

Remark 2: μ is directly related to the maximum change of the attack signal, in other words, its frequency. Clearly, large μ implies large estimation errors. We will explore the effect of high frequency attacks on estimation and on a robust control strategy proposed in the next section.

Remark 3: If the portion of compromised nodes is identified, then the estimation error $e_k = d_k - \hat{d}_k$ converges and is bounded by $|e_\infty| \le \frac{T\mu}{1-|\phi|}$.

B. Robust Control Algorithm

It is possible to modify the disturbance rejection using an *add-on compensator* in the controller of the form

$$u_k = u_{nom} - B^{-1}\hat{\Gamma}\hat{d}_k = \lambda_k - Gp^{-1}\hat{\Gamma}\hat{d}_k$$

where u_{nom} is the controller under normal conditions.

The mismatch between the supply and the demand is then described by $\mathcal{E}_k = G_p\lambda_k + \Gamma d_k - \hat{\Gamma}\hat{d}_k$. Clearly, if $\hat{e}_k = \Gamma d_k - \hat{\Gamma}\hat{d}_k$ is small, disturbances are attenuated.

Including the robust controller in the system produces an improvement in the estimation, leading to the following result.

Proposition 4: For the RTP system under additive attack, and the proposed robust controller $\hat{\lambda}_k = \lambda_k - G_p^{-1}\hat{\Gamma}\hat{d}_k$, where $\hat{d}_k$ is estimated according to (6), the estimation error is bounded by (8) and the new sensitivity function $\hat{S}_{\mathcal{E},d}$ is described by

$$\hat{S}_{\mathcal{E},d} = \frac{\Gamma(z-1)^2}{(z-\phi)(z-1+2\eta)}. \tag{9}$$

Proof: The first part of the proof is similar to Proposition 1, but because $\hat{\lambda}_k = \lambda_k - G_p^{-1}\hat{\Gamma}\hat{d}_k$ is the input, it leads to

$$\begin{aligned} \hat{e}_{k+1} &= \Gamma d_{k+1} - \hat{\Gamma}K\hat{e}_k - \hat{\Gamma}\hat{d}_k \\ &= \Gamma\Delta d_{k+1} + \phi\hat{e}_k \end{aligned}$$

As $K = (1-\phi)/\hat{\Gamma}$, in the equilibrium when $\hat{e}_{k+1} = \hat{e}_{k-1}$, $\hat{e}_\infty$ is bounded by

$$|\hat{e}_\infty| = \frac{|\Gamma|\Delta d_{k+1}}{1 - |\phi|} \le \frac{|\Gamma|T\mu}{1 - |\phi|}$$

The z transform of the error $\hat{e}_{k+1}$ is

$$\hat{e}(z) = \frac{\Gamma(z-1)}{(z-\phi)}d(z)$$

and of the power mismatch is $\mathcal{E}(z) = G_p\Lambda(z) + \hat{e}(z)$. Replacing $\Lambda(z) = -G_c(z)H(z)\mathcal{E}(z)$ yields to

$$\mathcal{E}(z) = -\frac{2\eta\mathcal{E}(z)}{(1-z^{-1})z} + \hat{e}(z)$$

Dividing by $d(z)$ and factorizing we obtain (9). $\blacksquare$

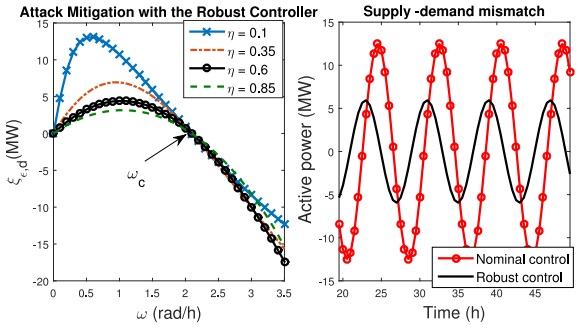

Fig. 4. Left: Performance metric of the robust controller for different η; right: example of the supply-demand mismatch for an attack $d_k = \sin(\pi/4kT)$. Clearly the robust control mitigates the effects of the attack better than the nominal control.

C. Robust Control Performance

Equation (9) describes the new sensitivity function including the robust controller. In order to quantify the performance of the proposed control strategy we define $\xi_{\mathcal{E},d}(\omega T, \phi) = |S_{\mathcal{E},d}(j\omega T)| - |\hat{S}_{\mathcal{E},d}(j\omega T, \phi)|$, which is the attack mitigation metric that compares the nominal (the case without the add-on compensation) and the robust control strategies. $\xi_{\mathcal{E},d}(\omega T, \phi) > 0$ for the frequencies at which the attack is mitigated with the robust control, i.e., when the attack amplification is smaller than for the case with no compensation. When $\xi_{\mathcal{E},d}(\omega T, \phi) < 0$ the system performance without compensation is better than the performance of system with compensation. We can obtain the frequencies at which the robust controller stops improving the system response under attacks. To do this, we need to find $\Omega_c = \{\omega : \xi_{\mathcal{E},d}(\omega T, \phi) < 0\}$, which corresponds to the case when $|S_{\epsilon,d}(j\omega T)| < |\hat{S}_{\epsilon,d}(j\omega T, \phi)|$. Taking (3) and (9), the condition reduces to the frequencies at which $|z - \phi| < |z - 1|$. Replacing $z = e^{j\omega T}$ and solving for ω, we obtain

$$\Omega_c = \left\{\omega : \omega > \frac{1}{T}\arccos\left(\frac{\phi + 1}{2}\right)\right\}.$$

Let us define ω_c as the lower bound of Ω_c, which corresponds to the frequency that makes $|z - \phi| = |z - 1|$. ω_c depends on ϕ and is larger when ϕ approaches -1. However, the pole corresponding to $z - \phi$ would approach to the unit circle, compromising the exponential stability of the system.

Fig. 4 (left) illustrates attack mitigation using the robust control when $\phi = 0$. Note that $\omega_c = 2.1$ *rad*, such that for all $\omega < \omega_c$ the attack attenuation is better than without the add-on compensator; however, for high frequencies, the inclusion of the compensator increases the impact of the attack. As an example, Fig. 4 (right) depicts the supply-demand mismatch for an attack $d_k = \sin(\pi/4kT)$ and $\eta = 0.5$ using both controllers. Clearly the robust controller attenuates the impact of the attack better than the nominal control.

The impossibility of attenuating high frequency attacks is a disadvantage of our proposed control algorithm. Therefore, in order to deal with this issue we can use the estimated disturbance $\hat{d}_k$ to calculate the frequency of the attack signal and: i) choose when to connect or disconnect the add-on compensator, and/or ii) change the parameter η in order to decrease the impact of the attack. Clearly, estimating the frequency of the attack and using the sensitivity analysis give us a powerful tool to take decisions that minimize the impact of the attack.

VII. Detection Mechanism

We have designed a new real time pricing algorithm that not only assures stability, but also minimizes the impact of attacks. However, in practice, while we have attenuated the attack, it would still be desirable to know if we are under attack or not, so we can remove compromised devices from our system.

The ISO calculates a clearing price each time period, but even in the presence of an attack, the price changes are small (see Fig. 3). However, the state estimator used in our robust controller can give information about the presence of an attacker, by analyzing the statistical behavior of the state estimator over long periods of time.

The detection mechanism that we propose is based on the accumulation of the rate of change of the estimated signal $\hat{d}_k$, $r_k = |\hat{\Gamma}\hat{d}_k - \hat{\Gamma}\hat{d}_{k-1}|$. This is known as the non-parametric CUSUM detection statistic [22]. The CUSUM statistic is a representative test belonging to the field of "Change detection theory", and it allows to detect persistent attacks even when they are of small amplitude. Besides, it requires low computational capacities (e.g., it does not need to store all previous values). The CUSUM statistic is defined as:

$$\begin{aligned} S_0 &= 0 \\ S_{k+1} &= (S_k + r_k - \alpha_k)^+ \end{aligned} \tag{10}$$

where S_k is the accumulated impact of the disturbance, and α_k is selected in such a way that $E[rk - \alpha_k] < 0$ when there is no attack, i.e., the rate of change of S_k under normal conditions (without attacks) remains close to zero or increases slowly. In our proposed strategy, the use of the error $\hat{\Gamma}\hat{d}_k$ is due to the fact that the ISO does not have knowledge about Γ. An attack is detected when $S_k > \delta$, and S_k is reset (set to 0). δ has to be selected such that the number of false alarms is low and it will be explained in the following section. It is important to know how long will take to the detection algorithm to detect and attack with certain frequency. It is possible to obtain some bounds for the time of detection as follows.

We can define $\Delta\hat{d}_k = \hat{d}_k - \hat{d}_{k-1}$. Replacing $\hat{d}_k$ from (5) and (6) we obtain $\Delta\hat{d}_k = K\hat{e}_{k-1}$. Using (8) and due to the fact that $K = (1 - \phi)/\hat{\Gamma}$ we find that

$$|\Delta\hat{d}_k| \leq \frac{|1 - \phi||\Gamma|T\mu}{(1 - |\phi|)|\hat{\Gamma}|}.$$

Replacing in our detection algorithm, we obtain

$$S_{k+1} \leq \left(S_k + \frac{|1 - \phi||\Gamma|T\mu}{1 - |\phi|} - \alpha_k\right)^+.$$

Let us define k^* as the detection time, i.e., the time at which $S_{k^*} = \delta$, and we assume that $\alpha_k = \alpha$ is a constant value properly selected such that the CUSUM tends to zero when there is no attack. Thus, the time of detection is no worst than

$$k^* \geq \frac{\delta}{\left(\frac{|1-\phi||\Gamma|\mu T}{1-|\phi|} - \alpha\right)^+} \tag{11}$$

Clearly, high frequency disturbances (i.e., large μ) or a large amount of compromised nodes (i.e., large $\Gamma = -\rho\dot{w}(\lambda_0)$) lead to fast detection. *Note that the detection is independent on* η.

Remark 5: If the attack is of $\omega = 0$, i.e., a step input, the sudden change will cause a single increase of S_k. If the amplitude is smaller than δ, the attack cannot be detected.

Remark 6: If the attacker designs an attack such that $\frac{|1-\phi||\Gamma|\mu T}{1-|\phi|} \leq \alpha$ the attack could not be detected at any time. However, α is typically small, so undetected attacks do not have a strong effect in the system.

A. Performance of the Detection Mechanism

The performance of a detection mechanism can be evaluated by the time to detect an attack (i.e., the time at which $S_k = \delta$) and the number of false alarms; however, the number of false alarms requires the specification of a time-interval, which makes it difficult to obtain a general performance metric. As δ depends directly on the number of false alarms when the system is not under attack, large δ implies that it will require more time to detect an attack, but it will decrease the false alarms caused by normal changes in the system. On the other hand, small δ increases false alarms but it is possible to detect attacks faster. As a consequence, there exists a trade-off between the number of false alarms and δ. Using some ideas from *Sequential detection theory*, and due to the fact that we want to monitor real-time control systems where there is not a fixed amount of time to observe, we propose the average time between false alarms T_{FA}, or more precisely, the expected time between false alarms $\mathbb{E}[T_{FA}]$. Clearly, $\mathbb{E}[T_{FA}]$ depends on the selection of δ. Therefore, for a given δ, we can take a series of measures during $\mathfrak{T}$ hours in a secure environment with no attacks and find the number of false alarms *nFA* during that time $\mathfrak{T}$. One way to calculate the expected time between false alarms is $\mathbb{E}[T_{FA}] = \mathfrak{T}/nFA$.

For a given δ, we calculate $\mathbb{E}[T_{FA}]$ and the minimum time for detection obtained in (11), which allow us to evaluate how decreasing the false alarms (increasing $\mathbb{E}[T_{FA}]$) affects the speed to detect an attack. An example is provided in the next section.

B. Simulations: Detecting Attacks

We assume a populated area with 1 million households, each one receiving information about the price every 30 minutes. To improve the realism of the simulations, we assume that the parameters D and b_k change each time period according to a half-hourly baseline demand profile provided by NYISO from the New York city from June 15th to June 30th. The baseline load per house is a scaled version of the real one. The parameters of the linear CEO model are $p = 31$ and $q = 917$ during the simulation time. We estimate the expected time of a false alarm $E[T_{fa}]$ for a time interval of 7 days for $\delta \in [0.01, 10]$ when there are no attacks. A large δ means that it will take us longer to raise a false alarm. Using this metric we can visualize the trade-off between δ and the time of detection.

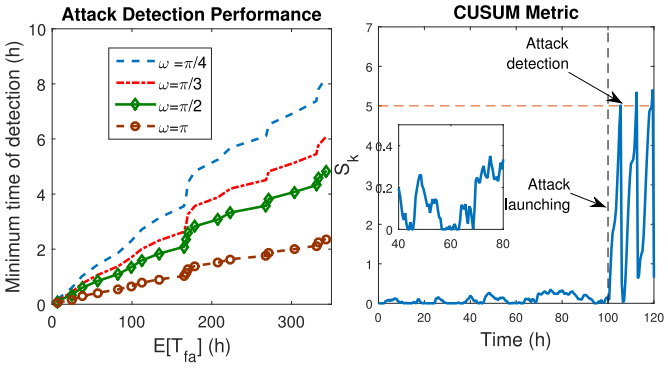

Fig. 5. Left: Minimum time of detection with respect to the estimated time of false alarms; right: example of the detection mechanism for $\delta = 5$, $\alpha = 0.06$, and an attack $d_k^a = 0.1\sin((\pi/4)kT)$ after 10 h. The minimum time for detection is 3.9 h and $\mathbb{E}[T_{fa}] = 150$ h.

We assume that an attack is launched and modifies the price information of 50% of the households. The attack is of amplitude 0.1 \$ /MWh, and frequency ω.

The estimation is based on prior information of the baseline load. However, we assume an error in the real-time baseline consumption, such that the ISO calculates the estimation and the robust control based on an approximate load profile, and not the real time consumption. Despite that limitation, the detection algorithm is able to detect an attack when a threshold is achieved.

Fig. 5 (left) illustrates the minimum time of detection with respect to $E[T_{fa}]$ for attacks of different frequencies. The ISO can choose a small δ to increase the detection speed but it would cause an increase in the number of false alarms. For high frequencies, the time of detection is low, which is an advantage in order to start a scan in the smart meters and find the victims of the attack. For instance, Fig. 5 (right) depicts the detection metric for $\delta = 5$, $\alpha = 0.06$, and an attack $d_k^a = 0.1\sin((\pi/4)kT)$ after 10 h. According to (11), the minimum time of detection is 3.9 h and $\mathbb{E}[T_{fa}] = 150$ h. After the attack is launched at $t = 100$ h, it takes approximately 4 h to detect it. If we select $\delta = 0.2$, the detection is almost instantaneous, but several false alarms will be generated.

Our work on detection is preliminary, and in future work we plan to identify the tradeoffs the attacker will face when deciding to launch attacks that maximize the error between power generated and consumed while also maintaining the attack undetected. Moreover, we plan to compare the CUSUM statistic with other detection mechanisms such as the Shiryaev-Roberts test and the bad-data detection.

VIII. Conclusion

In this work we used the theory from sensitivity analysis to understand how previously proposed attacks could be generalized and evaluated in a formal setting. In particular we showed how to find better attacks than previously proposed, and how to design robust control systems that can mitigate a large number of attacks.

We also found that the design of the price adjustment mechanism is fundamental in the resiliency of the system.

In particular, low values of η reduce the effect of the attacks on both the prices and sensors.

We also proposed an attack-resilient controller and an attack detection mechanisms. We believe we are one of the few research papers focusing on the important aspect of designing robust control algorithms against false data injection, as much of the previous work tends to focus on state estimation but does not consider the control actions of the system under attack, and how to design a controller that mitigates these attacks.

Our results show principled ways to use control theory in the design of attack-resilient cyber-physical systems. In general we believe that a well-designed defense-in-depth mechanism for cyber-physical systems will have to leverage not only information security expertise, but control theory to detect, respond, and reconfigure systems that can survive partial compromises.

One interesting area of future research that we did not address in this paper are the possible attack strategies that can be achieved by combining attacks to both: sensors and control signals. All our models assumed the attacker compromised the price signals, but not both. It is clear that if the attacker controls all control signals and all sensor signals then there is nothing we can do, but if the attacker has partial compromise of controllers and sensors, then the defender might still be able to design a robust algorithm that attenuates the attacks. We plan to look into this area in future work.

REFERENCES

[1] R. Tan, V. B. Krishna, D. K. Y. Yau, and Z. Kalbarczyk, "Impact of integrity attacks on real-time pricing in smart grids," in *Proc. ACM SIGSAC Conf. Comput. Commun. Security*, Berlin, Germany, 2013, pp. 439–450.

[2] Y. Liu, M. K. Reiter, and P. Ning, "False data injection attacks against state estimation in electric power grids," in *Proc. 16th ACM Conf. Comput. Commun. Security*, Chicago, IL, USA, 2009, pp. 21–32.

[3] D. Mashima and A. A. Cárdenas, "Evaluating electricity theft detectors in smart grid networks," in *Research in Attacks, Intrusions, and Defenses (RAID)*. Berlin, Germany: Springer-Verlag, 2012, pp. 210–229.

[4] G. Dán and H. Sandberg, "Stealth attacks and protection schemes for state estimators in power systems," in *Proc. IEEE 1st Smart Grid Commun. Conf. (SmartGridComm)*, Gaithersburg, MD, USA, 2010, pp. 214–219.

[5] Y. Mo and B. Sinopoli, "Secure estimation in the presence of integrity attacks," *IEEE Trans. Autom. Control*, vol. 60, no. 4, pp. 1145–1151, Apr. 2015.

[6] L. R. Phillips *et al.*, "Analysis of operations and cyber security policies for a system of cooperating flexible alternating current transmission system," Sandia Nat. Lab., Albuquerque, NM, USA, Tech. Rep. SAND2005-7301, Dec. 2005.

[7] S. Sridhar and G. Manimaran, "Data integrity attack and its impacts on voltage control loop in power grid," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, San Diego, CA, USA, 2011, pp. 1–6.

[8] M. Vrakopoulou, P. M. Esfahani, K. Margellos, J. Lygeros, and G. Andersson, "Cyber-attacks in the automatic generation control," in *Cyber Physical Systems Approach to Smart Electric Power Grid*. Berlin, Germany: Springer-Verlag, 2015, pp. 303–328.

[9] S. Sridhar and M. Govindarasu, "Model-based attack detection and mitigation for automatic generation control," *IEEE Trans. Smart Grid*, vol. 5, no. 2, pp. 580–591, Mar. 2014.

[10] S. Sridhar and G. Manimaran, "Data integrity attacks and their impacts on SCADA control system," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, Minneapolis, MN, USA, Jul. 2010, pp. 1–6.

[11] M. Negrete-Pincetic, F. Yoshida, and G. Gross, "Towards quantifying the impacts of cyber attacks in the competitive electricity market environment," in *Proc. IEEE PowerTech*, Bucharest, Romania, Jun. 2009, pp. 1–8.

[12] L. Xie, Y. Mo, and B. Sinopoli, "Integrity data attacks in power market operations," *IEEE Trans. Smart Grid*, vol. 2, no. 4, pp. 659–666, Dec. 2011.

[13] L. Jia, R. J. Thomas, and L. Tong, "Impacts of malicious data on real-time price of electricity market operations," in *Proc. 45th Hawaii Int. Conf. Syst. Sci.*, Maui, HI, USA, Jan. 2012, pp. 1907–1914.

[14] M. Esmalifalak, G. Shi, Z. Han, and L. Song, "Bad data injection attack and defense in electricity market using game theory study," *IEEE Trans. Smart Grid*, vol. 4, no. 1, pp. 160–169, Mar. 2013.

[15] NYISO. (2014). *NYISO Markets and Operations—Market Data—Load Data*. [Online]. Available: http://www.nyiso.com/public/markets_operations/market_data/load_data/index.jsp

[16] S.-E. Fleten and E. Pettersen, "Constructing bidding curves for a price-taking retailer in the Norwegian electricity market," *IEEE Trans. Power Syst.*, vol. 20, no. 2, pp. 701–708, May 2005.

[17] M. G. Lijesen, "The real-time price elasticity of electricity," *Energy Econ.*, vol. 29, no. 2, pp. 249–258, 2007.

[18] M. Roozbehani *et al.*, "Analysis of competitive electricity markets under a new model of real-time retail pricing," in *Proc. Int. Conf. Energy Market Eur. (EEM)*, Zagreb, Croatia, May 2011, pp. 250–255.

[19] A. Teixeira *et al.*, "Security of smart distribution grids: Data integrity attacks on integrated volt/VAR control and countermeasures," in *Proc. Amer. Control Conf. (ACC)*, Portland, OR, USA, 2014, pp. 4372–4378.

[20] J. C. Doyle, B. A. Francis, and A. R. Tannenbaum, *Feedback Control Theory*. Mineola, NY, USA: Courier Dover, 2013.

[21] K.-S. Kim and K.-H. Rew, "Reduced order disturbance observer for discrete-time linear systems," *Automatica*, vol. 49, no. 4, pp. 968–975, 2013.

[22] A. A. Cárdenas *et al.*, "Attacks against process control systems: Risk assessment, detection, and response," in *Proc. 6th ACM Symp. Inf. Comput. Commun. Security*, Hong Kong, 2011, pp. 355–366.

Jairo Giraldo (GSM'12–M'12) received the B.S. degree in electronic engineering from the National University of Colombia, Manizales, in 2010, and the M.S. and Ph.D. degrees from the Universidad de los Andes, Bogotá, in 2012 and 2015, respectively. He is currently a Research Associate with the Computer Science Department, University of Texas at Dallas. His research interests include distributed control algorithms for the power grid and their security and privacy.

Alvaro Cárdenas (M'06) received the B.S. degree in electrical engineering with a minor in mathematics from the Universidad de los Andes, Bogotá, Colombia, in 2000, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Maryland, College Park, MD, in 2002 and 2006, respectively. He is currently an Assistant Professor of Computer Science with the University of Texas at Dallas. His research interests include cyber-physical systems security and network security.

Nicanor Quijano (S'03–M'06–SM'13) received the B.S. degree in electronics engineering from Pontificia Universidad Javeriana, Bogotá, Colombia, in 1999, and the M.S. and Ph.D. degrees in electrical and computer engineering from Ohio State University, Columbus, OH, USA, in 2002 and 2006, respectively. He is currently a Full Professor with the Universidad de los Andes, where he is also the Director of the Research Group in Control and Automation Systems. His current research interests include hierarchical and distributed optimization methods, and using bio-inspired and game-theoretical techniques for dynamic resource allocation applied to problems in energy, water, and transportation.