2019 57th Annual Allerton Conference on Communication, Control, and Computing
(Allerton)
Allerton Park and Retreat Center
Monticello, IL, USA, September 24-27, 2019

Attacks on Electricity Markets

Carlos Barreto and Xenofon Koutsoukos

Abstract— **In this work we analyze how an adversary (who participates in the electricity market) can manipulate the bids of other agents to change the market's equilibrium. Here the adversary attempts to profit without damaging the system. We formulate the adversary's goal as the solution of a biased efficiency metric and identify the precise attack that maximize the adversary's objective function. We propose a defense scheme that modifies the bids to mitigate the impact of the attack. We validate the results simulating a detailed electric distribution system equipped with a transactive energy market using GridLAB-D.**

Electricity market, power system, security.

I. Introduction

In recent years, cyber attacks targeting *critical infrastructures* became a real threat, rather than a remote possibility. Although other cyber attacks targeting information assets had a large history, e.g., theft of information, such actions didn't transcend the cyber-space. However, cyber attacks can affect the physical space by exploiting the vulnerabilities of automation systems. For example, Stuxnet, the first known computer malware designed to harm physical processes, targeted PLC's to sabotage the uranium enrichment process in Iran [1]. The efficacy of Stuxnet demonstrated the vulnerabilities of critical infrastructures against cyber attacks, which started a race to both develop cyber weapons and improve the security of critical infrastructures [2].

Cyber attacks on the power grid of Ukraine in 2015 and 2016 offered a glimpse of the devastating consequences cyber attacks on critical infrastructures [3]. This work is part of the efforts to improve the protection of critical infrastructures, in particular, the power grid. Here we analyze how an adversary that participates in the electricity market can exploit the vulnerabilities of markets to profit without damaging to the system.

In particular, we consider a scenario in which an adverse generator manipulates the bids of customers to profit without damaging the system. This can occur if the adversary compromises appliances that participate directly in electricity markets [4], [5], [6], [7]. Here we show the precise attacks that balances both the profit of the adversary and the damage to the system.

The system administrator can filter the bids to mitigate the impact of attacks. Concretely, we leverage knowledge of the attack strategy to estimate it's impact and modify the bids to compensate the attacker's actions. In this case, the defender manages to reduce the attack losses modifying some of the bids.

We validate the attack and defense strategies through a detailed simulation of an electric distribution system using GridLAB-D and the prototypical distribution feeder models provided by the Pacific Northwest National Laboratory (PNNL) [8]. GridLAB-D excels in modeling the distribution side of the power system, because it includes detailed end-use load models that incorporate weather and market behaviors. Thus, we validate our findings with a realistic electricity markets.

The paper is structured as follows. Section II makes an introduction of electricity markets. We presents some formal markets models in Section III. Section IV shows both the adversary's goal and the attack strategy. In Section V we propose a defense scheme to mitigate the impact of attacks and validate the proposed defense through simulations in Section VI. We conclude the paper in Section VII.

II. Background on Power Systems

The power system has three main components, namely generation, transmission, and distribution. The generation includes sources of electricity, such as hydro or thermal generators, often scattered in large geographical areas. The transmission and distribution infrastructures connect generators with customers, but differ in few aspects: the transmission system carries energy across large distances using high voltage transmission lines, while the distribution system reduces the voltage and delivers energy directly to customers.

The operation of the power system has both physical and economic constraints. On one hand, the system's components have limits in the power that they can generate or carry. Moreover, generators and transformers are designed to operate at specific frequencies (e.g., 50 Hz or 60 Hz). Hence, the power grid works at a fixed frequency and needs to maintaining a balance between generation and demand. However, this is a non-trivial task due to the uncertainties in the demand[1] and the restrictions of generators.[2]

Power systems try to allocate resources in an *efficient* way, that is, creating the highest social satisfaction. In this case, an allocation refers to the transaction of energy and capital among the generators and customers. The power system uses market mechanisms to guarantee an efficient operation. In particular, electricity markets use auctions to elicit private

C. Barreto and X. Koutsoukos are with the Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN 37235 USA. email: {carlos.a.barreto, xenofon.koutsoukos}@vanderbilt.edu.

[1]The demand changes in time and depends on many factors, such as the temperature, the time of the day, the day of the week, and the season, among others.

[2]In general, the main generators cannot make instantaneous changes in their state, except some fast, but expensive, generators.

978-1-7281-3151-1/19/$31.00 ©2019 IEEE

information from agents (e.g., cost or utility functions) and determine the most efficient allocations [9], [10].

Unlike other markets, electricity markets need a central authority that monitors the system and enforces *reliable allocations*, i.e., allocations that comply with the physical constraints of the system. In general, an *independent system operator* (ISO) plans the operation to the system in advance guaranteeing its efficiency and reliability. The process of finding the most economical way to supply the demand, avoiding violations of physical constraints, is called *economic dispatch*.

Electricity systems often use two markets, namely the *day ahead market* (DAM) and the *real time market* (RTM). The DAM plays a crucial role planning the future operation of the power system. In particular, the DAM accept bids of supply or demand for a future period (e.g., the following day) and produces commitments that Hence, buyers and sellers must fulfill. In this way, the system reduces future uncertainties and allows the generators to prepare in advance for their operation.[3]

The RTM complements the DAM correcting imbalances between demand and generation during the actual operation. For example, if a seller cannot provide the contracted energy, the system operator must purchase energy from other sellers that participate in the RTM. Likewise, if a buyer uses more (or less) energy, then the system operator buys (or sells) energy in the RTM. In general, the RTM accepts bids of supply or demand for the next hour and finds an efficient way to correct deviations in the commitments from the DAM.

Fig. 1 illustrates the operation of the RTM. First, the system operator collects the bids from the agents (e.g., demand and cost functions). Then, every few minutes the ISO measures the system's state and computes the market equilibrium (the best allocation of resources and the price) based on the bids. Once the one hour period finishes, the ISO calculates the payments for each agent that participated in the market.

A. Demand Management Systems

The power grid is going through a modernization process to improve its efficiency, resiliency, and reliability. In particular, some innovations focus on enhancing the participation of users. In general, the mechanisms to coordinate users, also called *demand management systems*, use economic incentives to shape the demand of users. For example, *direct load control* programs compensate users who turn off their loads when the system is under stress [11]. Other schemes, such as *real time pricing*, *time of use*, and *critical peak pricing*, design prices to induce the desired response of the users, e.g., reduction of demand peaks [11].

Transactive energy (TE) is a distributed management approach that implements a two way communication between suppliers and customers. Thus, users can participate in the market trading energy and other ancillary services [6], [7]. TE relies on *transactive controls*, which control appliances and make bids considering both the market's price and the owner's preferences. In particular, loads that convert electricity into thermal energy, such as water heaters and air conditioning systems, may have some flexibility to participate in a market. For example, transactive controllers can anticipate future market prices and decide whether to store energy in thermal form. In particular, the transactive controllers can prepare for periods with high demand (and high prices) turning on the air conditioning systems when the prices are low. In this way, users can reduce costs at the expenses of some degree of discomfort e.g., some deviation from the desired temperature.

[3]The ISO decides in the DAM the schedule of each generator, considering their physical constraints, e.g., the time that they need to turn on/off or how fast they raise or reduce their supply.

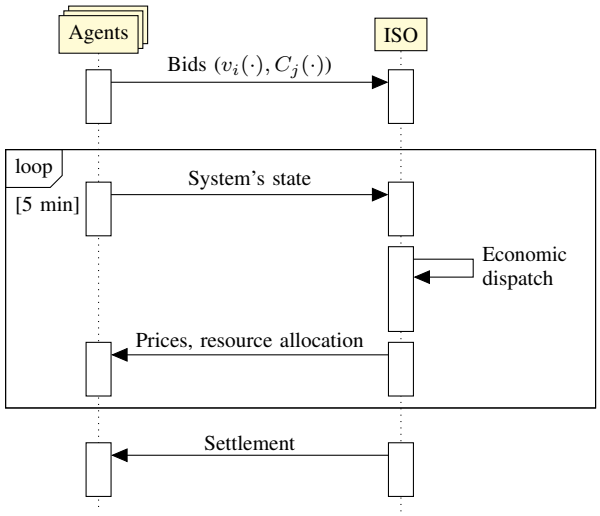

Fig. 1: Operation of a real time market.

III. System Model

In this section we introduce the objective of the market's participants, the system's efficiency metric (the social welfare), and the market's equilibrium. We use this market model to analyze how an adversary can attack the system to profit.

A. Ideal Market Model

Let us consider a market with n consumers and m generators, which conform the sets $\mathcal{C}$ and $\mathcal{G}$, respectively. Thus, the set of agents that participate in the market is $\mathcal{P} = \mathcal{C} \cup \mathcal{G}$. We denote with $q_i \geq 0$ the demand of the i^{th} customer and with $q_j \geq 0$ the production of the j^{th} generator. Furthermore, we represent the demand of both costumers and generators with the vector $\boldsymbol{q}$. For simplicity we ignore both the transmission and congestion losses; hence, the system's balance condition is

$$\sum\nolimits_{i\in\mathcal{C}} q_i = \sum\nolimits_{j\in\mathcal{G}} q_j. \tag{1}$$

Let us denote the surplus of the i^{th} customer as

$$u_i(q_i, p) = v_i(q_i) - q_i p,$$

where $v_i(q_i)$ represent the comfort of the i^{th} customer when it consumes q_i units of energy and $p \in \mathbb{R}$ represents the unitary price of energy. Roughly speaking, u_i quantifies the net benefit of consuming q_i units of energy. Moreover, we

define the surplus of the j^{th} generator (income minus the production costs) as

$$u_j(q_j, p) = q_j p - C_j(q_j), \tag{2}$$

where $C_j(q_j)$ represents the cost of producing q_j units of energy.

Assumption 1.

- *The valuation function v_i is continuous, derivable, and concave increasing.*
- *The cost function C_j is continuous, derivable, and convex increasing.*
- *The aggregated generation cost function is quadratic.*

We define the social welfare as the sum of the customer surplus and the aggregate seller surplus, that is,

$$f(\boldsymbol{q}, p) = \sum\nolimits_{i\in\mathcal{C}} u_i(q_i, p) + \sum\nolimits_{j\in\mathcal{G}} u_j(q_j, p) = \sum\nolimits_{i\in\mathcal{C}} v_i(q_i) - \sum\nolimits_{j\in\mathcal{G}} C_j(q_j). \tag{3}$$

The pair $(\boldsymbol{q}, p)$ conforms a market equilibrium, which determines the efficiency of the system. The market operator ignores the costs of necessities of the participants; hence, it cannot find directly the tuple $(\boldsymbol{q}^{op}, p^{op})$ that maximizes the social welfare f in Eq. (3). Nonetheless, it can elicit the private information through an auction.

An auction has the following steps:

1) The market operator commits to an auction, which describes the type of bids allowed and the procedure to calculate the equilibrium $(\boldsymbol{q}, p)$.
2) The participants choose their bids according to their interests, e.g., their surplus function and potential competition with other agents.
3) The market operator computes the equilibrium, that is, the unitary energy price p and the total amount of energy bought or sold.

In this case, the market operator requests directly the private information of each agent (but they can report false information). Thus, customers report a function $\hat{v}_i(\cdot)$ and generators report a function $\hat{C}_j(\cdot)$. In turn, the market operator finds the allocation $\boldsymbol{q}^*$ that solves the following optimization problem

$$\begin{aligned} \underset{\boldsymbol{q}}{\text{maximize}} \quad & \hat{f}(\boldsymbol{q}, p) = \sum\nolimits_{i\in\mathcal{C}} \hat{v}_i(q_i) - \sum\nolimits_{j\in\mathcal{G}} \hat{C}_j(q_j) \\ \text{subject to} \quad & \sum\nolimits_{i\in\mathcal{C}} q_i = \sum\nolimits_{j\in\mathcal{G}} q_j. \end{aligned} \tag{4}$$

Moreover, the clearing price p^* is equal to the marginal cost, that is, $p^* = \frac{d}{dx}\hat{C}_j(x)\Big|_{x=q_j}$, which hold for every generator $j \in \mathcal{G}$.

Some celebrated results from *mechanism design* [9] show that the auction mechanisms creates incentives to report truthfully private information. Hence, $\hat{v}_i = v_i$ and $\hat{C}_j = C_j$ for all $i \in \mathcal{C}$ and $j \in \mathcal{G}$. Therefore, the solution to Eq. (4) maximizes Eq. (3), i.e., $(\boldsymbol{q}^*, p^*) = (\boldsymbol{q}^{op}, p^{op})$.

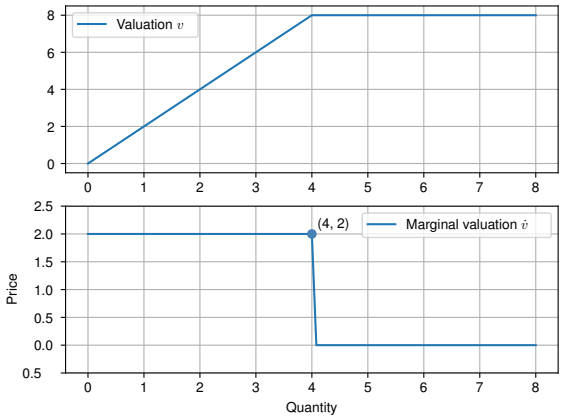

Fig. 2: Piecewise linear valuation associated to a bid with $Q_i = 4$ and $\sigma_i = 2$.

B. Simplified Market Model

In an ideal market the participants bid their supply or demand function, which describes the trades that they would accept. Although some electricity markets use these type of auctions, they often restrict the form of the functions. For example, some markets assume that the cost functions are quadratic and allow the bidders to choose some coefficients [12].

We use a simplified market model that accepts bids of the form (Q_k, σ_k), where Q_k represents the maximum capacity to either consume of produce energy and σ_k represents the price accepted (or charged) for Q_k or less units of energy. In this case, the bids represent piecewise linear functions that satisfy [13] (Fig. 2 shows the example of a valuation function)

$$\dot{v}_i(x) = \begin{cases} \sigma_i & \text{if } 0 \le x \le Q_i \\ 0 & \text{if } x > Q_i \end{cases} \tag{5}$$

for each customer $i \in \mathcal{C}$ and

$$\dot{C}_j(x) = \begin{cases} \sigma_j & \text{if } 0 \le x \le Q_j \\ 0 & \text{if } x > Q_j \end{cases}$$

for each generator $j \in \mathcal{G}$.

The auctioneer accepts bids periodically (e.g, 5 min intervals) and creates demand and offer curves ordering the bids in descending and ascending price, respectively. Thus, we can approximate the demand and offer curves with a piecewise linear function. In particular, consider the ordered bids of customers $\sigma_{c_1} \ge \sigma_{c_2} \ge \ldots \ge \sigma_{c_n}$, with $\{c_1, \ldots, c_n\} \equiv \mathcal{C}$. Then, the demand curve is the piecewise linear function (Fig. 3 shows an example of demand curve formed with four bids)

$$\dot{v}(q) = \begin{cases} \sigma_{c_k} & \text{if } \sum_{i=1}^{k} Q_{c_i} \le q < \sum_{i=1}^{k+1} Q_{c_i} \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Likewise, consider the ordered bids of sellers $\sigma_{s_1} \le \sigma_{s_2} \le \ldots \le \sigma_{s_n}$, with $\{s_1, \ldots, s_n\} \equiv \mathcal{G}$, which leads to the

following offer curve

$$\dot{C}(q) = \begin{cases} \sigma_{s_k} & \text{if } \sum_{i=1}^{k} Q_{s_i} \leq q < \sum_{i=1}^{k+1} Q_{s_i} \\ 0 & \text{otherwise.} \end{cases}$$

The market equilibrium corresponds to the intersection between the demand and the offer curves, which guarantee that the total demand equals the total production.

IV. Attack model

In this section we formulate the goal of an adverse generator and show how it can modify the bids of consumers to achieve its objective. We formulate the attack using the model in Section III-A, assuming a single buyer and seller. Later we extend the attack to the simplified market model of Section III-B considering multiple bidders.

A. Adversary's goal

Here we assume that an adversary participates in the market and manipulates the bids of other agents to profit. In particular, we leverage the attack model introduced in [14] to describe the objective of an adversary that attempts to profit without damaging the system.

Let $\mathcal{A} \subseteq \mathcal{P}$ be the set of adversaries and $\mathcal{V} = \mathcal{P} - \mathcal{A}$ the set of victims (agents who suffer the consequences of the attack). In this case, the adversaries design the attack so that the market maximizes an alternative objective function f_A, rather than the social welfare f. We define the attacker's objective function as

$$f_A(\boldsymbol{q}, p) = \lambda \sum\nolimits_{a \in \mathcal{A}} u_a(q_a, p) + \sum\nolimits_{v \in \mathcal{V}} u_v(q_v, p) = f(\boldsymbol{q}, p) + (\lambda - 1) \sum\nolimits_{a \in \mathcal{A}} u_a(q_a, p), \quad (7)$$

where $\lambda \geq 1$ represents the intensity of the attack.

If $\lambda = 1$, then f_A is equal to the social welfare f; however, if $\lambda > 1$, then the allocation that maximizes f_A will benefit the agents of the set $\mathcal{A}$ at the expenses of these belonging to $\mathcal{V}$.

The next result shows that the adversary benefits with higher market prices.

Proposition 1. *Let us consider two market prices p_1, p_2 such that $p_1 \geq p_2$. Then, each generator has larger profits with larger market prices. that is,*

$$u_j(q_1, p_1) \geq u_j(q_2, p_2).$$

Proof: With each price the j^{th} generator will produce the amount that maximizes its profit (see Eq. (2)). Hence, the generator will produce q_1 and q_2 units of energy, such that $p_1 = \dot{C}_j(q_1)$ and $p_2 = \dot{C}_j(q_2)$. Since the generation cost is concave and increasing we have $q_1 \leq q_2$. Furthermore, the cost function satisfies

$$C_j(q_2) - C_j(q_1) \geq \dot{C}_j(q_2)(q_2 - q_1) \quad (8)$$

Now, let us consider the impact of the prices in the generator's profit

$$u_j(q_1, p_1) - u_j(q_2, p_2) = q_1 p_1 - q_2 p_2 + C_j(q_2) - C_j(q_1).$$

Since $p_1 > p_2$, we an rewrite the previous expression as

$$u_j(q_1, p_1) - u_j(q_2, p_2) \geq (q_1 - q_2) p_2 + C_j(q_2) - C_j(q_1). \quad (9)$$

Now, we can use Eq. (8) to rewrite Eq. (9) as

$$u_j(q_1, p_1) - u_j(q_2, p_2) \geq 0.$$

B. Attack strategy

Here we consider a market with only one seller and one buyer, which will be useful to design attack on the simplified market in Section III-B. Let us approximate the aggregate valuation function of customers as $v(d) = \delta \log(1+d)$ and the aggregate generation cost function as $C(g) = \sigma g^2/2$, where d and g represent the total demand and generation, respectively. From the balance condition in Eq. (1) the generation equals to the demand; here for simplicity we denote $q = d = g$.

The social welfare based on the reported functions ($\hat{v}$ and $\hat{C}$) is (see Eq. (4))

$$\hat{f}(q, p) = \hat{v}(q) - \hat{C}(q).$$

Moreover, the goal of an adverse generator is (see Eq. (7))

$$f_A(q, p) = v(q) - C(q) + (\lambda - 1)(qp - C(q)). \quad (10)$$

Although Eq. (12) depends on the market's price p, which is unknown at the moment of the attack, it is possible to leverage the equilibrium conditions to express the fake bids as a function of known information. In particular, since $C(q)$ is quadratic and $p = \dot{C}(q)$ we have

$$qp - C(q) = \sigma q^2 - \frac{\sigma}{2} q^2 = \frac{\sigma}{2} q^2 = C(q).$$

Hence, we can rewrite Eq. (10) as

$$f_A(q, p) = v(q) - C(q) + (\lambda - 1) C(q).$$

The generator can transform $\hat{f}$ into f_A manipulating the bid of the customer as

$$\hat{v}^a(q) = v(q) + (\lambda - 1) C(q). \quad (11)$$

Alternatively, the if the market requests the customer's marginal valuation, the adversary can submit

$$\frac{\partial \hat{v}^a}{\partial q}(q) = \dot{v}(q) + (\lambda - 1)\dot{C}(q), \quad (12)$$

which follows from Eq. (11).

From Eq. (6) we can express the fake bid Eq. (12) on the individual bids as

$$\frac{\partial \hat{v}_i^a}{\partial q_i}(q_i) = \dot{v}_i(q_i) + (\lambda - 1)\dot{C}(q)$$

However, the marginal cost depends on the total demand q, but we define the valuation with respect to the individual demand q_i. For this reason, we approximate the marginal cost with an estimation of the future price $\tilde{p} \approx \dot{C}(q)$, resulting

$$\frac{\partial \hat{v}_i^a}{\partial q_i}(q_i) = \dot{v}_i(q_i) + (\lambda - 1)\tilde{p}.$$

If the market clears within short periods, then we can approximate the future price $\tilde{p}$ using the clearing price from previous periods.

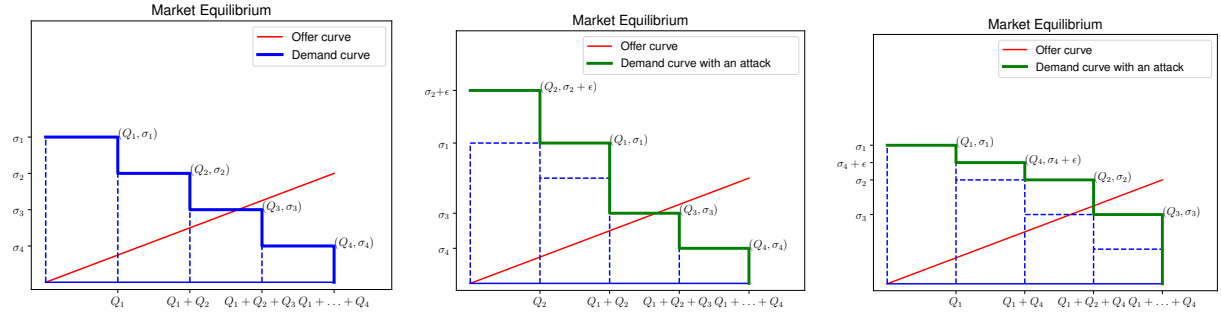

(a) Nominal case. (b) Attack that modifies the 2nd bid. (c) Attack that modifies the 4th bid.

Fig. 3: An attack that injects a bias ε in the bids may change the equilibrium (intersection of the offer and demand curves) if it affects bids with prices below the equilibrium price.

Unlike other attacks that target sensors [15], the adversary does not need detailed information about the physical structure of the power system (e.g., grid's topology), because the economic dispatch has that information into account to find the optimal operation.

V. Defense

Here we analyze how to modify the bids to mitigate the impact of attacks. In particular, we leverage knowledge of the attack strategy to estimate and mitigate it's impact.

We denote with (q^{op}, p^{op}) and (q^a, p^a) the optimal equilibrium and the equilibrium with an attack, respectively. Let us assume that the adversary compromises $n^a = n\gamma$ bids of buyers, where $\gamma \in [0, 1]$ represents the degree of the attack. Let us denote with $\mathcal{C}^a$ the set of meters compromised, where $n^a = |\mathcal{C}^a|$. Thus, we denote an attack with the tuple (λ, γ).

The adversary can change the equilibria if it manages to show that the buyers have a higher willingness to pay for energy. In particular, the adversary needs to change the bids that don't accept the equilibrium price. According to Eq. (5), the adversary must target bids belonging to the set $\mathcal{C}^L \equiv \{k|\sigma_k \leq p^{op}, k \in \mathcal{C}\}$, which has total demand $Q^L = \sum_{i\in\mathcal{C}^L} Q_i$. Conversely, modifying bids that accept the market price won't change the equilibrium. Fig. 3 shows an example of attacks on bids and their impact in the market's equilibria.

We assume that the adversary does not compromise the quantity requested in each bid Q_i, because the utility company measures the actual demand for billing purposes. Moreover, let us assume that the bids of buyers have roughly the same quantity, that is, $Q_i \approx \bar{Q}$, for $i \in \mathcal{C}$. Hence, the attack changes the equilibrium quantity in proportion to the compromised bids with price below p^{op}, that is, bids from the set $\mathcal{C}^L$. Thus, we can approximate the impact of an attack (change in the equilibrium quantity) as

$$q^a - q^{op} \approx \bar{Q}|\mathcal{C}^L \cap \mathcal{C}^a| \tag{13}$$

If the adversary selects the bids randomly, then the expected number of compromised bids belonging to $\mathcal{C}^L$ is

$$\mathrm{E}[|\mathcal{C}^L \cap \mathcal{C}^a|] = \eta = |\mathcal{C}^L|\ n^a/n\ .$$

We estimate the number of bids in $\mathcal{C}^L$ as

$$|\mathcal{C}^L| \approx \frac{\sum_{i=1}^{n} Q_i - q^{op}}{\bar{Q}}.$$

With the above expressions we rewrite Eq. (13) as

$$\mathrm{E}[q^a - q^{op}] \approx \bar{Q}\eta.$$

Thus, we can estimate the equilibria quantity without attacks q^{op} using the observed equilibria q^a and an estimation of n^a

$$q^{op} \approx \frac{nq^a - n^a \sum_{i=1}^{n} Q_i}{n - n^a} \tag{14}$$

Moreover, we can use Eq. (14) to estimate the number of bids compromised

$$n^c \approx (q^a - q^{op})/\bar{Q}$$

The defender can mitigate the attack's impact moving the equilibria quantity close to q^{op}. In particular, the defender can reduce the price of n^c of the bids with largest prices, which has the effect of moving the attacked demand curve to the left (closer to p^{op}).

VI. Experimental results

A. Power Grid Model

We make a detailed simulation of an electric distribution system using GridLAB-D and the prototypical distribution feeder models provided by the Pacific Northwest National Laboratory (PNNL) [8]. GridLAB-D excels the modeling the distribution side of the power system, because it includes detailed end-use load models that incorporate weather and market behaviors. In particular, GridLAB-D models retail markets through auctions, in which both sellers and appliances participate [4]. In this case, we use the prototypical feeder *R1-12.47-3*, which represents a moderately populated area. We added representative residential loads to the distribution model using the script in [16]. Thus, our distribution model has 118 commercial and residential loads, which in turn incorporate appliances such as heating, ventilation, and air conditioning (HVAC) systems, water heaters, pool pumps, among others.

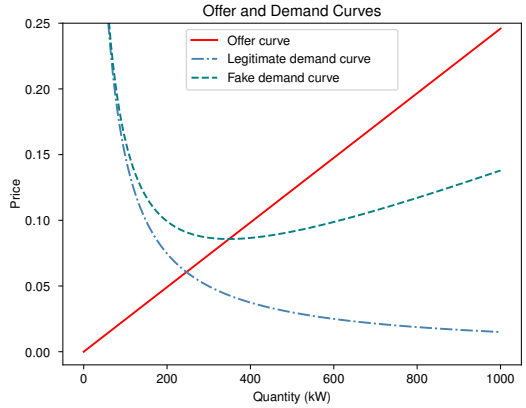

(a) Ideal model.

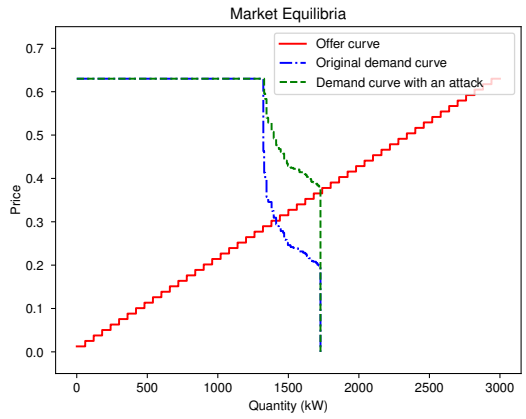

(b) GridLAB-D's model.

Fig. 4: Market equilibrium with an attack of intensity $\lambda = 1.5$. The attack increases both the price and the production at the equilibrium.

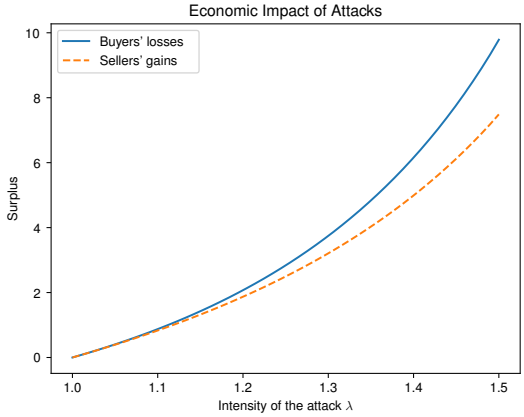

(a) Ideal model.

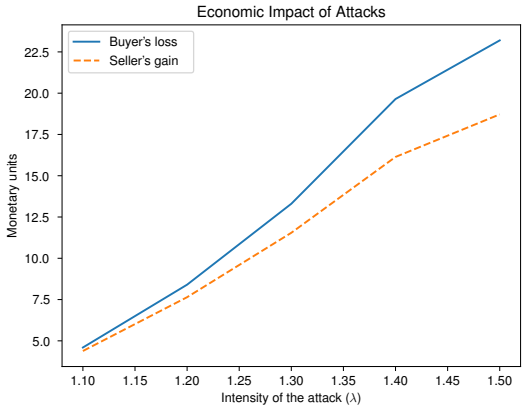

(b) GridLAB-D' model.

Fig. 5: Economic impact of the attack for both customers and sellers as a function of the attack intensity λ. The losses of the customer exceed the benefit of the seller.

GridLAB-D models the response of the loads to weather and market's prices, giving realism to the simulations. In particular, we simulate weather from summer time in Nashville, TN, and use auctions that collect bids and decide the equilibrium each 5 minutes. The details of the market's structure are available in [17].

The transactive controllers report in their bids their current state. In particular, the controllers choose as the bid's quantity the current demand of their appliances, which approximates the future demand. Furthermore, the bid's price is an estimation of the price necessary to maintain the current demand. GridLAB-D assumes that the non-responsive loads bids the maximum price allowed in the market (set as 0.63 in the simulations). In this way, each bid will determine a segment of the aggregate curve and the intersection of the curves determines the equilibrium, that is, the total demand $\boldsymbol{q}$ and price p.

B. Impact of Attacks

Fig. 4 shows an example of the market's curves with an attack with $\lambda = 1.5$. The offer curve corresponds to the marginal cost, while the demand curve corresponds to the marginal valuation. The curves of the GridLAB-D's model are truncated because some loads are inelastic (the flat segment with maximum price represent inelastic loads). In this case the attack changes all the bids submitted by controllers and manages to rise the prices; however, the equilibrium cannot surpass the total demand requested by the customers.

Fig. 5 shows the economic impact of the attack as a function of the attack's intensity λ. Concretely, we define the benefit of the adversary as

$$\sum\nolimits_{i \in \mathcal{G}} u_i(q_i^a, p^a) - u_i(q_i^{op}, p^{op})$$

and the losses of the customers as

$$\sum\nolimits_{i \in \mathcal{C}} u_i(q_i^{op}, p^{op}) - u_i(q_i^a, p^a).$$

Observe that the seller has positive gains with $\lambda > 1$; however, the damage on the customers exceeds the profit of the adversary. As expected, the attack harms the social welfare, because the adversary can benefit only by causing losses to other agents.

C. Efficacy of the Defense

Fig. 6 shows an example of the market equilibrium with the proposed defense scheme. Here we use $\lambda = 1.5$ and

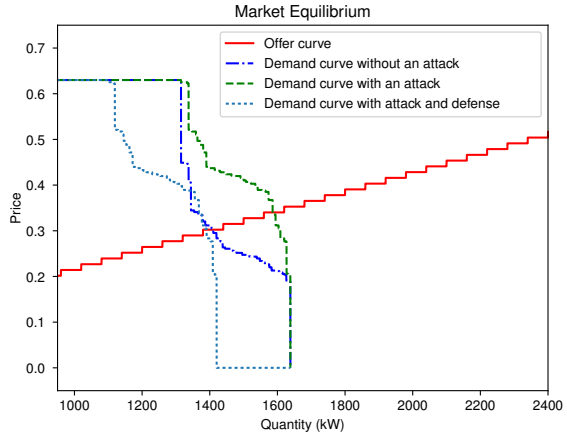

Fig. 6: Example of the market equilibrium with and without an attack. The attack has intensity $\lambda = 1.5$ and increases both the price and the generation at the equilibrium.

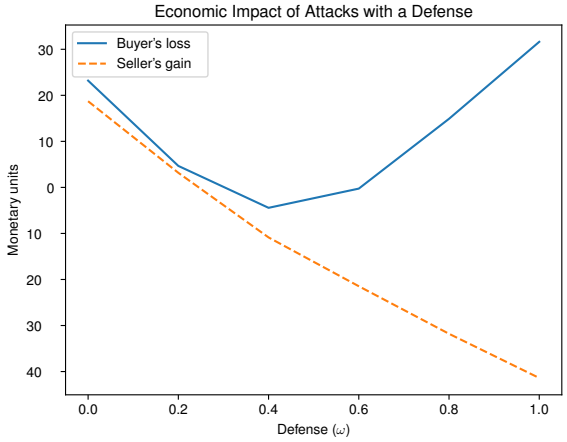

Fig. 7: Efficacy of the defense scheme as a function of the estimated proportion of compromised meters ω. The efficacy improves with the accuracy of the estimations, i.e., as ω approaches $\gamma = 0.4$.

$\omega = \gamma = 0.8$. This example illustrates that the defense moves the demand curve to the left to compensate the actions of the adversary.

The proposed defense scheme requires an estimation of $\gamma = n^a/n$, the proportion of meters compromised. Fig. 7 shows the efficacy of the proposed defense when the defender uses $\omega = n^c/n$ as an estimation of γ (the proportion of meters compromised). Here we measure the efficacy of the defense as

$$f(q^{op}, p^{op}) - f(q^d, p^d),$$

where (q^d, p^d) represents the equilibrium with the proposed defense. The efficacy improves as ω approaches γ; however, the defense cannot prevent completely the damage of the attack. In this case, the defense is more efficient when ω falls below γ, hence, errors estimating the attack parameters may harm the system's efficiency.

VII. Conclusions and Future Work

This work shows how an adversary can manipulate bids to profit regulating the damage on the system. The attacker leverages the market infrastructure; hence, the attack's design does not need information about the system's state or its topology. Moreover, the adversary can succeed despite of the restrictions imposed by the auctioneer regarding the type of bids accepted. We propose a defense strategy that compensates the impact of the attack by modifying some bids; however, inaccurate estimations of the attack's parameters can harm both buyers and sellers.

This work analyzes only one type of adversary (a generator or seller); however, we plan to analyze how other agents (or group of agents) may attack the system to profit. We also plan to design defense mechanisms anticipating possible responses of the adversary.

References

[1] K. Zetter, "An unprecedented look at stuxnet, the world's first digital weapon," WIRED magazine, 2014. [Online]. Available: http://www.wired.com/2014/11/countdown-to-zero-day-stuxnet/

[2] D. E. Sanger, *The Perfect Weapon: War, Sabotage, and Fear in the Cyber Age*. Crown, 2018.

[3] K. Zetter, "Inside the cunning, unprecedented hack of ukraine's power grid," http://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/, WIRED magazine, mar 2016, accessed October 16, 2017.

[4] J. C. Fuller, K. P. Schneider, and D. Chassin, "Analysis of residential demand response and double-auction markets," in *2011 IEEE Power and Energy Society General Meeting*, July 2011, pp. 1–7.

[5] S. Soltan, P. Mittal, and H. V. Poor, "Blackiot: Iot botnet of high wattage devices can disrupt the power grid," in *27th USENIX Security Symposium (USENIX Security 18)*. Baltimore, MD: USENIX Association, 2018, pp. 15–32.

[6] K. Kok and S. Widergren, "A society of devices: Integrating intelligent distributed resources with transactive energy," *IEEE Power and Energy Magazine*, vol. 14, no. 3, pp. 34–45, May 2016.

[7] R. Masiello and J. R. Aguero, "Sharing the ride of power: Understanding transactive energy in the ecosystem of energy economics," *IEEE Power and Energy Magazine*, vol. 14, no. 3, pp. 70–78, 2016.

[8] K. P. Schneider, Y. Chen, D. P. Chassin, R. G. Pratt, D. W. Engel, and S. E. Thompson, "Modern grid initiative distribution taxonomy final report," Pacific Northwest National Laboratory, Tech. Rep., 2008.

[9] R. B. Myerson, "Optimal auction design," *Mathematics of operations research*, vol. 6, no. 1, pp. 58–73, 1981.

[10] P. Klemperer, "Why every economist should learn some auction theory," in *Advances in Economics and Econometrics: Invited Lectures to 8th World Congress of the Econometric Society*, M. Dewatripont, L. Hansen, and S. Turnovsky, Eds. Cambridge University Press, 2003.

[11] P. Siano, "Demand response and smart grids—a survey," *Renewable and sustainable energy reviews*, vol. 30, pp. 461–478, 2014.

[12] R. Baldick, "Electricity market equilibrium models: The effect of parametrization," *IEEE Transactions on Power Systems*, vol. 17, no. 4, pp. 1170–1176, 2002.

[13] S. Li, W. Zhang, J. Lian, and K. Kalsi, "Market-based coordination of thermostatically controlled loads—part i: A mechanism design formulation," *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 1170–1178, 2016.

[14] C. Barreto and A. Cardenas, "Impact of the market infrastructure on the security of smart grids," *IEEE Transactions on Industrial Informatics*, pp. 1–1, 2018.

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

[16] https://github.com/gridlab-d/Taxonomy_Feeders/tree/master/PopulationScript, 2015, accessed: January 12, 2019.

[17] H. Neema, H. Vardhan, C. Barreto, and X. Koutsoukos, "Web-based platform for evaluation of resilient and transactive smart-grids," 03 2019.