MaMIoT: Manipulation of Energy Market Leveraging High Wattage IoT Botnets

Tohid Shekari*, Celine Irvene*, Alvaro A. Cardenas†, and Raheem Beyah*
t.shekari@gatech.edu,cirvene3@gatech.edu,alacarde@ucsc.edu,rbeyah@ece.gatech.edu
*Georgia Institute of Technology
†University of California, Santa Cruz

ABSTRACT

If a trader could predict price changes in the stock market better than other traders, she would make a fortune. Similarly in the electricity market, a trader that could predict changes in the electricity load, and thus electricity prices, would be able to make large profits. Predicting price changes in the electricity market better than other market participants is hard, but in this paper, we show that attackers can manipulate the electricity prices in small but predictable ways, giving them a competitive advantage in the market.

Our attack is possible when the adversary controls a botnet of high wattage devices such as air conditioning units, which are able to abruptly change the total demand of the power grid. Such attacks are called Manipulation of Demand via IoT (MaDIoT) attacks. In this paper, we present a new variant of MaDIoT and name it Manipulation of Market via IoT (MaMIoT). MaMIoT is the first energy market manipulation cyberattack that leverages high wattage IoT botnets to slightly change the total demand of the power grid with the aim of affecting the electricity prices in the favor of specific market players. Using real-world data obtained from two major energy markets, we show that MaMIoT can significantly increase the profit of particular market players or financially damage a group of players depending on the motivation of the attacker.

CCS CONCEPTS

• **Security and privacy → Economics of security and privacy**; *Distributed systems security*.

KEYWORDS

Electricity market manipulation, financial profit/damage, high wattage IoT botnet

ACM Reference Format:
Tohid Shekari*, Celine Irvene*, Alvaro A. Cardenas†, and Raheem Beyah*. 2021. MaMIoT: Manipulation of Energy Market Leveraging High Wattage IoT Botnets. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21), November 15–19, 2021, Virtual Event, Republic of Korea.* ACM, New York, NY, USA, 19 pages. https://doi.org/10.1145/3460120.3484581

1 INTRODUCTION

Real-world attacks, as well as demonstration projects, have shown the effectiveness of cyberattacks against the power grid [31, 35, 63, 64]. These are *direct* attacks, meaning that they directly target the critical components (e.g., generators) or the supervision and control system of the power grids. Recent works, however, have shown how to attack the power grid *indirectly*, by compromising consumer devices (and not devices in the grid) [28, 51]. In particular, the adversary creates or rents a botnet of high wattage IoT devices (e.g., an Internet-connected EV charger or water heater), and then, collectively and abruptly changes the electricity demand of thousands of these devices (via simultaneously turning them on/off), creating an unanticipated sudden power surge which can potentially result in local or regional blackouts [28, 51].

In this paper, we analyze a new unexplored threat from high wattage IoT botnets: attacks to the deregulated wholesale electricity market [50]. According to the U.S. Energy Information Administration (EIA), the average price of electricity and total energy consumption in the U.S. was 75 USD/MWh and 2.935×10^9 MWh, respectively [55, 56], with approximately 220 billion USD transactions. Such markets can be attractive targets for cybercriminals around the world and selfish traders who are willing to manipulate the market.

Market manipulation (creating artificial prices) is not a new problem. In the U.S., the primary purpose of the Securities and Exchange Commission (SEC) is to enforce the law against stock market manipulation. Recently, security researchers started to study how botnets can facilitate stock market manipulation [62]. In this paper, we perform a similar study but in the electricity market.

In a role similar to that of the SEC for the stock market, the Federal Energy Regulatory Commission (FERC) has oversight on electricity markets in the U.S. and can impose penalties on entities that manipulate the prices. While there have been multiple electricity market manipulation cases over the years, none of the discovered cases so far has been enabled by cyberattacks [45].

The most visible example of electricity market manipulation is the case of Enron [38]: Enron traders had names for strategies that they used to manipulate the market. Some of these include "Death Star," where traders filed nonexistent transmission schedules in order to get paid to alleviate congestion that did not exist; "Fat boy", where traders overscheduled power transmission reflecting nonexistent demand; "Get Shorty" focused on selling power and services it did not have with the expectation that they would not be asked to fulfill the contract, and "Ricochet" focused on exporting electricity outside of California, to be latter bought at higher prices

CCS '21, November 15–19, 2021, Virtual Event, Republic of Korea

ACM ISBN 978-1-4503-8454-4/21/11…$15.00
https://doi.org/10.1145/3460120.3484581

by circumventing the local price caps [38]. While Enron was a high-profile case, there are several other traders that have been fined for manipulating the market over the years, including JPMorgan [16, 21], Louis Dreyfus Energy Services [14], and Barclays [47]. Still, there may be several other cases of market manipulation that are not detected [14]. The difficulty in proving market manipulation cases in the power grid might motivate attackers to use cyberattacks in their efforts to profit from (or shock) the electricity market.

Our proposed attack, which we call Manipulation of Market via IoT (MaMIoT), exploits the relationship between demand and price fluctuations [8] and manipulates the market prices by slightly altering the total power consumption of the grid through a high wattage IoT botnet. This botnet can give a huge advantage to the malicious participants in the market, as they can predict sudden (but small) changes in the demand for electricity (changes created by the botnet). Being able to predict electricity demand changes is akin to a stockbroker who could predict small fluctuations of the stock prices in advance.

The market manipulation through MaMIoT can be implemented in two general ways based on the ultimate goal and motivation of the attacker: i) to provide additional financial profits for one of the market players (i.e., the attacker is one of the market players such as the previously discovered market manipulation cases by FERC); ii) to cause economic damage to the entire market (i.e., attacker is a nation-state actor who is doing this as a part of a trade/cold war). For each of the cases, we develop an optimization model to maximize the profit (or damage) of a specific market player (or to the entire market) while keeping the attack as stealthy as possible. The input data for the optimization models are obtained by crawling and processing publicly available datasets from official electricity market websites (they can be similarly obtained through a trading tool called Bloomberg terminal or similar trading software). The outputs of the optimization models are the timeline of the botnet activation/deactivation (to realize the manipulated prices) along with the malicious bids/offers in the electricity market (to realize the additional attack gain).

The main contributions of this paper are summarized as follows:

- This is the first paper in the literature that identifies and analyzes the emerging threat from the high wattage IoT botnets to the wholesale electricity markets.
- In order to develop successful attacks, we develop optimization algorithms to decide when and how to attack, subject to the constraints of the market, and the power constraints of the system. Using the optimization models helps us maximize the attacks' gains.
- We evaluate and test the effectiveness of the attacks with real-world traces.
- We propose a set of practical countermeasures to considerably limit the damaging consequences and severity of the studied attacks.

The rest of this paper is organized as follows. In Section 2, we explain the basic structure of electricity markets and their various players. We then present the threat model and attack feasibility in Section 3. We develop a formulation of the attack model for different attackers in Section 4. In Section 5, we evaluate the performance of the proposed approach with real-world case studies. We then propose a set of practical countermeasures in Section 6. Finally, we conclude and discuss open research questions in Section 7.

2 BACKGROUND

2.1 Structure of the Electricity Market

There are two main markets for electricity. The wholesale market focuses on the bulk power grid, while the retail market is where individual consumers (e.g., homeowners) interact with electric utilities. In this paper, we focus on the wholesale market.

Before deregulation of the wholesale market in the 1980s and 90s, the electricity industry operated as a monopoly, which meant that generators, transmission lines, substations, and distribution lines were owned and operated by monopolistic (sometimes government-owned) utilities. Proposers of deregulation argued that rising electricity costs were due to the lack of an efficient market.

With deregulation, electric utilities were forced to sell their generation plants and became wholesale consumers, having to purchase electricity on the spot market everyday. Deregulated markets also allowed new participants (outside of electric utilities) to join the wholesale markets such as banks, financial firms, and smaller traders; in fact, regulators of the electricity market encourage traders to join these markets in the hopes of making them more efficient. Deregulated electricity markets allow the participation and competition of multiple energy producers and utilities in the market providing customers with efficient, cheap, and more reliable energy [6]. There are in general four major players in the market: producers (generators), consumers (retailers), a market operator, and a regulator.

2.1.1 Producers. Generation companies such as nuclear or coal power plants, hydropower plants, and wind farms mainly fall into this category where their basic goal is to produce and sell electric energy. They may also sell services such as frequency regulation, voltage control, and reserves to help the system operators maintain the reliability of the power grid. A generation company can own a single generator or a portfolio of generators with different technologies [32, 50]. In some cases, financial companies such as JPMorgan rent a power plant with multiple generators to participate in the market and make profits from their trading strategies [16, 47]. Other traders can also buy electricity from producers and then sell them in the wholesale market [38]. Electricity prices on the supply side are highly affected by fuel prices.

2.1.2 Retailers. Retailers buy electrical energy from the wholesale energy markets and resell it to consumers (e.g., homeowners). Electric utilities and electric vehicle (EV) aggregators[1] are two examples of such retailers [32, 58], but again, other traders can join the market and purchase electricity [38]. Consumer prices are highly affected by weather and economic activity.

2.1.3 Market Operator (MO) or Independent System Operator (ISO). Market operators (MO), and independent system operators (ISO), run a computer program to match the bids and offers submitted by producers and retailers [32]. A responsibility of the ISO is to clear the market in such a way that it preserves the reliability of

[1] An EV aggregator is a market player who participates in the wholesale market on behalf of a certain number of EVs and charges the batteries of these EVs based on a signed contract.

the power grid. For example, if all producers of electricity are in one geographical area and all consumers in another, the ISO has to make sure that the power transmission lines have the capacity to transfer the amount of energy. Therefore, specific bids and offers that violate the limitations of the power grid, will be removed from the market to maintain the grid's stability [32].

2.1.4 Regulator. A regulator is a government organization responsible for ensuring the fair and efficient operation of market players. This organization monitors the market, studies its environment, and determines a set of rules to prevent abuse, manipulation, and fraud by the market players. The regulator also sets the prices for the products and services that are provided by monopolies or single parties to preserve fairness in the market [32].

2.2 Day-Ahead and Real-Time Markets

The wholesale market is different than various other markets in that the products cannot be stored, so the production of electricity has to match the demand for electricity at every point in time, which in turn can lead to high volatility of electricity prices. To hedge this price volatility, the market is divided into two parts: the day-ahead market (which helps stabilize the prices of electricity) and the real-time market [50].

In the day-ahead market, all players in the market make forecasts of how much electricity will be needed for the next day, and then at 12pm, they make offers for the amount of electricity they will produce (or buy) for every hour of the 24 hours of the next day. About four hours later the market is cleared by the ISO, and it releases the specific commitments for each player. For example, if player 1 submitted a bid for consuming 2MWh for a price of $15 from 3 pm to 4 pm, player 1 has to do that, otherwise, she will be penalized financially.

Since predicting the exact energy demand a day in advance is impossible, the market needs to have a real-time component to correct prediction errors from the day-ahead market. If the day-ahead market committed to less generation than what is currently in demand, players make new bids and offers for electricity. If the day-ahead market is committed to more generation than what is currently in demand, the prices of electricity in the real-time market can plummet and in some cases can become negative (asking industries to consume electricity and being rewarded for that).

Both markets work the same way. Bids/offers submitted to the ISO (for the day-ahead or real-time market) at a specific time slot are shown in Figure 1. As illustrated in the figure, each player of the market submits a quantity-price pair to the ISO for each time interval. The ISO sorts the bids/offers based on the suggested prices and solves the optimization problem expressed in equations (1)–(4) to maximize the social welfare of the market players and determine the optimal price of the market at each time slot while satisfying the power system physical constraints.

$$\text{maximize} \quad welfare = \sum_{d\in\Omega_D} P_d^D \lambda_d^D - \sum_{s\in\Omega_S} P_s^S \lambda_s^S \qquad (1)$$

subject to

$$0 \le P_d^D \le P_d^{D,\max}, \forall d \in \Omega_D \qquad (2)$$

$$0 \le P_s^S \le P_s^{S,\max}, \forall s \in \Omega_S \qquad (3)$$

$$\textit{system reliability constraints} \qquad (4)$$

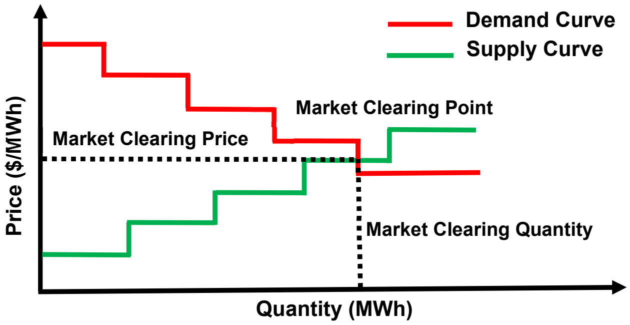

Figure 1: Illustration of a typical bid/offer in the market and its settlement mechanism.

The intuition behind equation (1) is to maximize the area between the red and the green curve in Figure 1. P_d^D is the power demand (in MWh) by player d and λ_d^D is the price player d is willing to pay to buy that amount of power. In the figure, P_d^D is one of the steps in the x-axis of the red curve and λ_d^D is one of the steps in the y-axis of the red curve. Similarly, P_s^S is the amount of power supplier s is willing to provide at price λ_s^S. At the market-clearing price, all players are happy because consumers are buying for less than (or equal) to their bid, and suppliers are receiving more (or equal) for the generation they promised. Equations (2) and (3) denote that one of the bids or offers is not going to be accepted in its totality (e.g., that is why in Figure 1 the red line intersects the green line, meaning that one of the supply offers is cut shorter than what the supplier was offering). Finally, Equation (4) is beyond the scope of this paper, but it basically deals with the physical topology of the grid and makes sure that the scheduled supply and demand do not violate any capacity constraints of the transmission lines in the power grid.

3 THREAT MODEL

We assume our attacker has a high wattage botnet, as proposed in recent work [15, 28, 51]. The difference with previous work on high wattage botnets is that we are not using the botnet in an attempt to cause electricity blackouts, instead, we study how an attacker can profit from the botnet by manipulating the electricity market.

For example, one of the possible ways to profit from the electricity markets is by creating congestion. Power companies that buy or sell in the wholesale market can get hurt by sudden price spikes, but they can buy a financial instrument known as a congestion contract, which acts as a hedge against losses. Financial firms such as *DC Energy* or *Saracen Energy* can also buy these contracts, and then profit when the grid becomes overburdened [14]. An attacker with a high wattage botnet can attempt to create congestion in specific areas of the grid, or specific times.

Market manipulation in the wholesale electricity market is not new. Perhaps the most popular case of wholesale electricity market manipulation is the case of Enron, a company that claimed revenues of over 100 billion during 2000 according to Fortune magazine, and who Fortune magazine named *America's Most Innovative Company* for six consecutive years. In the deregulated wholesale electricity market, traders–often pure middlemen who do not own power plants–began to apply their experience in market trading. Enron used a variety of strategies to manipulate the electricity market in California. These strategies included offering to sell electricity but schedule it in a way that cannot be delivered (e.g., through a low capacity line), scheduling too much electricity to flow on some lines so the ISO would pay to relieve that congestion, and urging operators to remove power generation plants to perform unnecessary maintenance, in order to cut the supply and share the profits of higher prices for generators [19]. Enron traders even labeled these strategies, with names such as "Death Star," "Fat Boy," "Get Shorty," "Ricochet," and "Mega Watt Laundering". Enron later claimed that all competitors were employing similar strategies but avoided these names [38].

There are dozens of investigations for market manipulation every year. One high-profile case happened when FERC found evidence of manipulative bidding by JPMorgan in the California electricity market back in 2013 [21]. After a long fight in court, JPMorgan agreed to pay $410 million USD to settle allegations [16]. The company had rented two power plants and used manipulative bidding strategies in the market by creating artificial conditions (e.g., temporary power shortage in the grid) to sell the generated power at expensive premium rates [16]. Another market manipulation case occurred when Louis Dreyfus Energy Services began buying cheap congestion contracts in an area with a lot of generation (wind turbines) and then a second trader created the impression that congestion was hitting the desired area, thus profiting by nonexistent congestion [14]. More recently in 2017, FERC approved a $105 million settlement with the British bank Barclays for market manipulation [47].

In this paper, we focus on attackers with access to a high wattage botnet that can manipulate the market. We consider two different types of attackers:

Attacker Type I: The first attacker is a fraudulent trader, similar to one of the cases identified in the last two paragraphs. The goal of this trader is to use the high wattage botnet to her advantage, manipulating the electricity market and profiting financially from the attack.

Attacker Type II: The second attacker does not participate in the market, but instead uses the high wattage botnet to make the market as inefficient as possible, and thus cause widespread economic damage to operators of the power grid.

The overall structure of the threat model for these attackers is shown in Figure 2. Attackers first crawl the historical and real-time market data from available online sources to obtain the optimization parameters that are necessary for designing the attack scenarios (⓪). An Attacker Type I (fraudulent insider) then submits bids or purchase orders, and then also submits commands to the botnet (①). An Attacker Type II does not participate in the market, and simply sends commands to the botnet to cause market inefficiencies (②).

3.1 Basics of MaMIoT

The intuition behind the MaMIoT attack is the following: with a high wattage IoT botnet, the attacker can predict better the real-time demand than other peers in the market, because the high

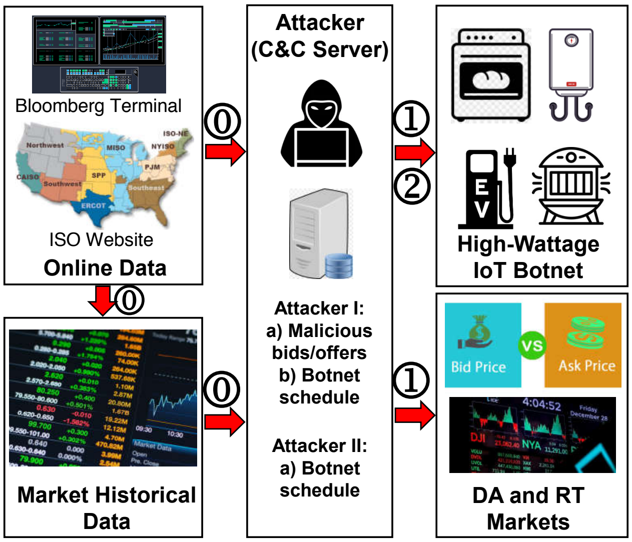

Figure 2: The overall view of the threat model and attack scenarios. ⓪ Crawler: Crawling the historical and real-time market data to be used for designing the attack scenario, ① Attacker Type I: Submitting the malicious bids/offers to the day-ahead and real-time markets and modifying the grid demand with the available botnet, ② Attacker Type II: Modifying the grid demand with the available botnet.

wattage IoT botnet can allow the attacker to increase or decrease the electricity load slightly at will.

While not entirely an accurate analogy, using an example from the airplane industry can provide insights into how the electricity market can be manipulated: suppose you book an airline ticket for a flight you do not intend to board: it is a waste of time and money unless you are sure the flight will be overbooked and the airline will have to dish out rewards to passengers who agree to stay home [40]. Similarly, if you commit to producing electric power in the day-ahead market but the load does not materialize in the real-time market (e.g., by turning off several high wattage IoT bots), you will get rewarded for not producing the power you did not have in the first place. On the other hand, an attacker can increase the load on a given day by turning on several high wattage IoT bots. If the attacker is prepared (e.g., putting two generators in service for the day, instead of only one), it can deliver electricity in the real-time market at lower prices than other generators who did not anticipate this extra demand (and who did not turn on reserves).

More concretely, an adversary can manipulate the real-time market prices by slightly changing the total demand of the power grid through a high wattage IoT botnet. This observation can be mathematically represented as:

$$\lambda_k^{RT} = \lambda_k^{RT0} + \alpha_k \Delta D_k^{System}, \forall k \in \Omega_K \tag{5}$$

where λ_k^{RT} is the manipulated real-time market price, λ_k^{RT0} is the original market price, ΔD_k^{System} is the power grid demand manipulation, and α_k is a constant number that can be obtained by analyzing market historical data. Additionally, k and Ω_K are the indexes and set of time intervals (e.g., 15 min.) in the market. According to this equation, an attacker can manipulate the real-time

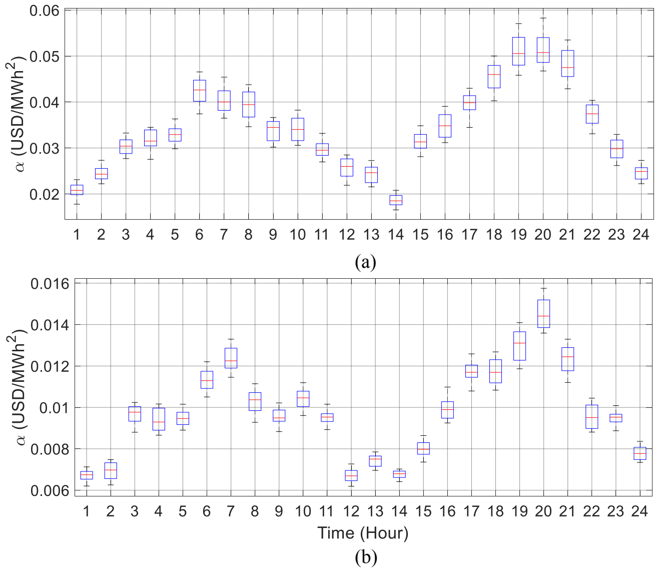

(a)

(b)

Figure 3: The coefficient representing the price-load sensitivity in the real-time market obtained from analyzing the market historical data for one month. a) New York ISO; b) California ISO.

market price in his own favor by slightly changing the total demand of the power grid through high wattage IoT botnets (ΔD_k^{System}). Based on our analysis, α_k changes considerably at every hour in a given market. Therefore the attacker needs to be strategic and find the *optimal* time to attack, as changing the load at different times will give different benefits.

By analyzing the historical data of two large electricity markets (New York and California) [11, 12, 42, 43] during one month period, we can estimate the value of α_k at each time interval; this is illustrated in Figure 3. According to this figure, the real-time market price in the New York market is more sensitive to demand manipulation compared to the California market. As we can see, price manipulation at certain hours (19-21) can be done with a fewer number of high wattage IoT bots because of the higher price-load sensitivity factor (α_k). For example, a high wattage IoT botnet with 100,000 bots can change the system demand by 1% and this could result in +15 USD (~30% increase) in New York and +5 USD (~20% increase) in California.

Launching successful market manipulation attacks requires sophisticated strategies for maximizing the objective function while maintaining the committed resources cleared in the market, and a low profile to avoid being detected by the market regulator. Before we discuss sophisticated optimization strategies, we first describe a naive baseline attack.

3.2 Baseline Attack

A naive attack strategy for a consumer to get lower electricity prices would be to turn off all high wattage devices in the botnet. With lower demand, the price of electricity will fall and the consumer will pay less for electricity. The equivalent naive attack strategy for a generator is to turn on all high wattage devices in the botnet, increasing the demand, and thus increasing electricity prices. The algorithm for the baseline attack is outlined in Algorithm 1.

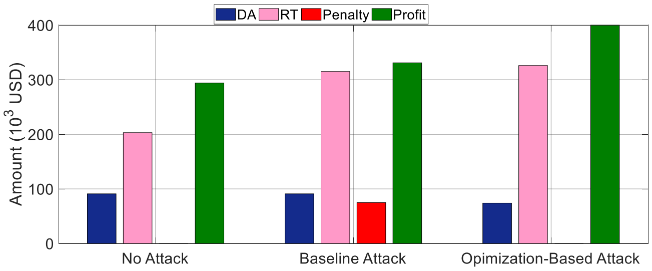

Figure 4: The profit breakdown of the simulated market player in a single day with different bidding strategies in the New York market. DA: day-ahead profit, RT: real-time profit, Penalty: market penalties, Profit: overall profit.

Algorithm 1 Baseline Attacker

```
function BASELINE($URL_{ISO}$)
    $History = Crawl(URL_{ISO})$   ▷ Read market historical data
    for $k = 1$ to $K$ do
        $\alpha_k = Statistics(History)$   ▷ Estimate price-load sensitivity at each time slot
        $botnet_k = Maximize(\alpha_k)$   ▷ Maximize the price at each time slot and find the relevant botnet attack
return $botnet_k$
```

Although the baseline attack may seem reasonable and effective at first glance, our analysis shows that it has two major weaknesses. First, if the adversary tries to benefit a single market player, this price manipulation must be accompanied by the consideration of the player's physical constraints; otherwise, this strategy will lead to lower attack gains because of the inevitable market penalties (making promises to produce or consume electricity, and then not being able to fulfill these promises). Figure 4 illustrates the profit breakdown of a typical market player in a single day with different bidding strategies. As we can see, the overall profit of the player increases in the baseline attack scenario compared to when there is no attacks. However, there are some penalties in the baseline attack scenario because of the violation of the market limitations and the exclusion of the player's physical constraints (breaking the promises, as explained above). To prevent these penalties, we need a more sophisticated attacker, which we introduce in the next section. In the more sophisticated attack, the adversary gains less profit in the day-ahead market; however, he obtains a large profit in the real-time market with no market penalty. The small day-ahead profit reduction can be regarded as the preparation cost for gaining the maximum profit in the real-time market with no penalties.

The second weakness of the baseline attack model is that the adversary might be detected by FERC fairly easily. Stealth is a key point for the success of MaMIoT attacks as this will allow the attacks to be repeatable (otherwise short-term gains will be small). Figure 5 shows the system load profile associated with different bidding strategies mentioned in Figure 4. As can be seen, the load profile of the system during the baseline attack exceeds the upper limit of a typical load forecasting error. Therefore the system operator

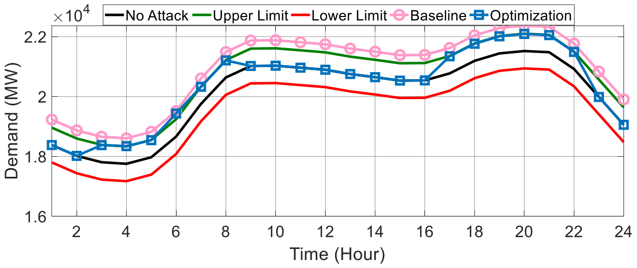

Figure 5: This figure shows that the optimized attack is less disruptive to the grid than the baseline attack. The optimized attack only activates the botnet at certain times, and with fewer active bots.

can easily differentiate and detect this as an anomaly. Conversely, the load profile of the optimization-based attack remains within the lower and upper error limit band, and hence, it will be hard to differentiate these small electricity changes from the general daily errors in load forecasting.

In short, while the naive attack may be better for the adversary than not launching attacks, the gains will be short-lived. There are too many variables and constraints (physical constraints of the player, market constraints, and stealth constraints) that the baseline attack does not consider. Therefore, we introduce more sophisticated adversaries who leverage mathematical optimization frameworks to maximize the attack gains.

3.3 Attacker Type I

There are two main decision variables for this type of attacker: i) malicious bids/offers made by the market player (attacker), and ii) system demand alteration at each time interval through the high wattage IoT botnet (see ① in Figure 2).

To determine the key parameters (e.g., price-load sensitivity (α) as shown in Figure 3) for the optimization model, the adversary first analyzes publicly available historical market data from the market's website [12, 43, 46] or on a Bloomberg terminal [1]. Next, the attacker runs an optimization problem to determine the malicious day-ahead and real-time bids/offers in the electricity market and the required system demand change of each time slot (this will be realized through the high wattage IoT botnet). In addition, we constrain attacks to be stealthy so that it is hard to accuse a specific market player of abuse. The algorithm for the first attacker type is outlined as follows (we will give details in Section 4.1):

Algorithm 2 Attacker Type I

```
function ATTACKI($URL_{ISO}$)
    $History = Crawl(URL_{ISO})$  ▷ Read market historical data
    for $k = 1$ to $K$ do
        $\alpha_k = Statistics(History)$  ▷ Estimate price-load sensitivity at each time slot
        $D_k^{stealthy,max} = Statistics(History)$  ▷ Estimate stealth parameter at each time slot
    $botnet_k, Bid_t^{DA}, Bid_k^{RT} = Optimization(\alpha_k, D_k^{stealthy,max}, physics)$  ▷ Maximize the player's gain subject to player's physical constraints, stealth constraints, and market constraints
return $botnet_k, Bid_t^{DA}, Bid_k^{RT}$
```

3.4 Attacker Type II

In this case, the attacker is a nation-state actor whose goal is to maximize the economic damage to a group of market players by manipulating real-time market prices through high wattage IoT botnets. Because this attacker is external to the system, the only decision variable that is needed to be implemented in the market is the power demand changes at each time interval through the available high wattage IoT botnet (see ② in Figure 2).

Financial markets have already seen nation-state attacks as part of cold/trade wars and MaMIoT is the first cyber-based energy market manipulation that could damage the electric industry generation/demand of a targeted country [2]. A nation-state attacker could even be a foreign investor in generation/demand companies who wants to alter the total revenue of the electricity generation/consumption corporations to affect their stock shares in his favor.

Similar to the previous attacker, the nation-state actor analyzes the historical market data to obtain price-load sensitivity at each time slot (price-load sensitivity (α) as shown in Figure 3). Then, the attacker solves an optimization problem to determine the optimal attack vector to be implemented with the botnet of high wattage IoT devices at each time interval. As mentioned earlier, we design the attack mechanism to be stealthy. The algorithm for the second attacker type is outlined as follows (we will provide details in Section 4.2):

Algorithm 3 Attacker type II

```
function ATTACKII($URL_{ISO}$)
    $History = Crawl(URL_{ISO})$  ▷ Read market historical data
    for $k = 1$ to $K$ do
        $\alpha_k = Statistics(History)$  ▷ Estimate price-load sensitivity at each time slot
        $D_k^{stealthy,max} = Statistics(History)$  ▷ Estimate stealth parameter at each time slot
    $botnet_k = Optimization(\alpha_k, D_k^{stealthy,max})$  ▷ Maximize the attack's gain subject to stealth constraints and market constraints
return $botnet_k$
```

3.5 Attack Feasibility

When we consider the feasibility of the MaMIoT attack there are two questions that come up, i) Will this attack work in practice? and ii) Can one acquire, compromise, and control a large botnet of high wattage IoT devices located within certain geographic boundaries (e.g., within the state of California)?

We argue the answer to both of these questions is yes. To start, the command and control of IoT botnets is not new [49]. As IoT devices have grown in complexity and become more widely deployed, their power consumption has increased accordingly. This is emphasized in [23] where we see the average power consumption of an air purifier is 200W, making the premise of a high wattage IoT botnet fairly reasonable.

3.5.1 Number of Available High Wattage IoT Bots. The number of high wattage IoT devices that an attacker can use in a MaMIoT attack is growing. The number of houses with smart thermostats in North America alone has increased at an unprecedented scale, representing a small fraction of the total high wattage IoT devices in the automation field (see Figure A1 in Appendix) [53]. EV chargers are another big source of high wattage devices. Concerning the matter of location, attackers can trivially determine whether a compromised device is within a certain geographical area through the device's IP address.

A MaMIoT attack does not need a significant number of compromised high wattage IoT devices to be effective, but as the size of the botnet increases so does the economic impact of the attack (discussed at length in Section 5). Even with a small botnet of high wattage devices, the attack can be extremely devastating as illustrated in Sections 5.3 and 5.4. All things considered if we take into account that IoT botnets, such as Mirai, are capable of containing over six hundred thousand compromised devices [5], a future implementation of MaMIoT with a high wattage botnet of 100,000 bots would be a common scenario. Now, we discuss how this botnet could be obtained.

3.5.2 IoT Botnet Acquisition. Since the release of its source code in 2016, variants of Mirai have run rampant [34] and have been credited with several attacks including assaults against OVH (French cloud computing company), Dyn (DNS service provider), and the Liberian Internet infrastructure. These and other IoT malware such as, Bashlite, Reaper, Satori, and Linux.Aidra, have been able to infect IoT devices through primarily known and patchable vulnerabilities [49] resulting in a low barrier to entry for the supply and demand of botnet for hire services.

Botnet rental services level the playing field for entities that are unable to create/deploy malware for building their own army of bots. On the dark web, buyers can obtain access to DDoS services for periods ranging from days to several months [13]. Within their service period, clients can launch a limited or unlimited (for a premium) number of attacks per day with a guaranteed minimum duration ranging from minutes to hours. Some of the available botnet rental services can be found in the Appendix. Based on the presented results in Section 5, even if the cost of building/renting a high wattage IoT botnet is ten times bigger than what is mentioned in a realistic botnet rental service, this cost is still negligible compared with the attack gain.

3.5.3 Effect of the Attack on the End User's Billing Statement. The financial effect of the proposed attacks on each end-user depends on their monthly total power consumption as well as the duration of the attack. According to the EIA, the average electricity consumption of Americans is 914 kWh per month. Tennessee has the highest electricity consumption at 1,282 kWh per residential customer, and Hawaii has the lowest at 517.75 kWh per residential customer [57]. Assuming that each of the high wattage IoT bots consumes 3 kW electricity and considering the stealth strategies explained in Sections Appendix II and 5 (the attack is carried out 100 days per year (8 days/month) and each bot is turned on for 3 hours on average during the daily attack), each compromised home would consume 72 kWh more electricity in each month. This means a 7.8% increase in the billing statement in the attacked residents, which will likely be unnoticeable. For example, a typical customer who pays $120 monthly for his electricity bill in the US, will pay $129 if he is attacked. Note that the considered numbers are associated with the most severe IoT botnet attack on the electricity market (see NY3 and CA3 in Figures 7, 13, and 14). For example, replacing 3 kW with 1 kW will lead to a trivial 2.6% increase in the monthly electricity bill.

4 FORMULATION OF THE ATTACK MODEL

In this section, we explain the optimization models that adversaries can employ to determine the attack scenarios as explained in Section 3.

4.1 Attacker Type I

As mentioned in Section 3, this type of attacker is one of the market players whose goal is to maximize his own profit by manipulating the real-time system demand through the strategic use of high wattage IoT botnets. To show the effectiveness of the MaMIoT attack, we present one sample optimization model for a common market player: a generation company. Note that without loss of generality, the proposed optimization framework with slight changes can be leveraged to model the other types of market players. It is worth mentioning that the detailed explanations of the notation used in the following equations are given in Appendix I for quick referencing.

We assume that a conventional power plant, which includes multiple steam turbines and generators, can control a botnet of high wattage IoT devices to make profit from the energy market. The following optimization problem is developed to determine the optimal offers in the day-ahead and real-time markets along with the attack vector to be sent to the bots in the botnet. The objective function of the model is defined as:

$$\begin{aligned}\text{maximize}\quad profit^{G} = \sum_{g\in\Omega_G}\sum_{t\in\Omega_T} profit_{gt}^{DA,G}\\ + \sum_{g\in\Omega_G}\sum_{k\in\Omega_K} profit_{gk}^{RT,G} - \sum_{k\in\Omega_K} Cost_{k}^{Botnet}\end{aligned} \tag{6}$$

where $profit^{G}$ is the total profit of the generation company. Similarly, $profit_{gt}^{DA,G}$ and $profit_{gk}^{RT,G}$ denote the profit of unit g at hourly (sub-hourly) time interval t (k) in the day-ahead and real-time markets, respectively. Also, $Cost_{k}^{Botnet}$ represent the cost of building/renting the required botnet for the desired attack. These variables can be calculated as follows:

$$\begin{aligned}profit_{gt}^{DA,G} = \lambda_{t}^{DA} P_{gt}^{DA,G} - \left(\lambda_{g}^{SU} u_{gt} + \lambda_{g}^{SD} v_{gt}\right)\\ -\lambda_{g}^{G,Constant} x_{gt}^{G}, \forall g \in \Omega_G, t \in \Omega_T\end{aligned} \tag{7}$$

$$\begin{aligned}profit_{gk}^{RT} = \lambda_{k}^{RT} P_{gk}^{RT,G} - \left(\lambda_{g}^{G,Fuel} P_{gk}^{RT}\right)\\ -\frac{\lambda_{t}^{DA,Dev}}{\mathcal{K}}\left(P_{gt}^{DA,Dev+,G} + P_{gt}^{DA,Dev-,G}\right),\\ \forall g \in \Omega_G, k \in \Omega_K, t \in \Omega_k,\end{aligned} \tag{8}$$

$$Cost_k^{Botnet} = \lambda_k^{Botnet} D_k^{attack}, \forall k \in \Omega_K. \quad (9)$$

The day-ahead profit for each unit, $profit_{gt}^{DA,G}$, includes the revenue from the day-ahead market participation ($\lambda_t^{DA} P_{gt}^{DA,G}$) minus the costs associated with the unit start-up, shut-down ($\lambda_g^{SU} u_{gt} + \lambda_g^{SD} v_{gt}$), and its constant operation ($\lambda_g^{G,Constant} x_{gt}^G$). The real-time profit for each unit, $profit_{gk}^{RT,G}$, includes the revenue from the real-time market participation ($\lambda_k^{RT} P_{gk}^{RT,G}$) minus the fuel cost of the unit ($\lambda_g^{G,Fuel} P_{gk}^{RT}$) along with the cost associated with the penalty for deviating from the day-ahead bid in real-time operation. According to our analysis, the real-time market price (i.e., λ_k^{RT}) in (8) can be notably affected by the real-time power mismatch between the system generation and demand. This property can be effectively used by the adversary to change the profit which can be obtained from the real-time market. The attacker can change the real-time system demand through the high wattage IoT botnets and affect the real-time market price in his favor. By analyzing the historical data of the market (which is publicly available on the official websites of ISOs and Bloomberg terminal [1, 11, 12, 42, 43, 46]), we can extract the relationship between the system real-time power mismatch and the real-time market price. In this paper, we assumed a linear model for this change as follows:

$$\lambda_k^{RT} = \lambda_k^{RT0} + \alpha_k \Delta D_k^{System}, \forall k \in \Omega_K \quad (10)$$

where λ_k^{RT} is the real-time market price after the attack, λ_k^{RT0} is the expected real-time market price before the attack, ΔD_k^{System} is the total change in the system demand which can be done through a high wattage IoT botnet, and α_k is a constant number which can be obtained by analyzing market historical data. According to (10), the attacker can significantly alter the real-time market price in his favor if he has access to a large number of compromised IoT devices. However, if the attacker changes the system demand significantly, it can be easily detected by the ISO in the market as an anomaly. Therefore, in order to keep the attack stealthy and undetectable, we need to limit the system demand change to stay within the normal load forecasting error (as determined from historical market data). The mathematical representation of this limitation can be defined as:

$$\begin{aligned} -\Delta D_k^{stealthy,\max} \leq \Delta D_k^{System} = D_k^{attack} - D_k^{actual} \\ \leq \Delta D_k^{stealthy,\max}, \forall k \in \Omega_K \end{aligned} \quad (11)$$

in which $\Delta D_k^{stealthy,\max}$ is the average of the load forecasting error at time slot k which is determined by analyzing the market historical data from the ISO's public website. Another point that should be considered here is that the system demand alteration should be capped with the maximum capability of the high wattage IoT botnet, that is,

$$\begin{aligned} -\Delta D_k^{botnet,\max} \leq \Delta D_k^{System} = D_k^{attack} - D_k^{actual} \\ \leq \Delta D_k^{botnet,\max}, \forall k \in \Omega_K \end{aligned} \quad (12)$$

where $\Delta D_k^{botnet,\max}$ is the maximum capability of the IoT botnet at time slot k. This parameter represents the maximum capability of the attacker in changing the total demand of the power grid. It should be noted that additional strategies, such as limiting the number of hours for the demand alteration, can be embedded in (11) to maintain attack stealth. The physical constraints associated with the power plant are listed as follows:

$$\begin{aligned} P_{gk}^{Act,G} = P_{gt}^{DA,G} + \left(P_{gt}^{DA,Dev+,G} - P_{gt}^{DA,Dev-,G}\right) \\ + P_{gk}^{RT,G}, \forall g \in \Omega_G, k \in \Omega_K, t \in \Omega_k \end{aligned} \quad (13)$$

$$\begin{aligned} x_{gt}^G P_g^{\min} \leq P_{gk}^{Act,G} \leq x_{gt}^G P_g^{\max}, \\ \forall g \in \Omega_G, k \in \Omega_K, t \in \Omega_k \end{aligned} \quad (14)$$

$$\begin{aligned} -R_g^D \leq P_{gk}^{Act,G} - P_{g(k-1)}^{Act,G} \leq R_g^U, \\ \forall g \in \Omega_G, k \in \Omega_K \end{aligned} \quad (15)$$

$$x_{g(t-1)}^G - x_{gt}^G + u_{gt}^G \geq 0, \forall g \in \Omega_G, t \in \Omega_T \quad (16)$$

$$x_{gt}^G - x_{g(t-1)}^G + v_{gt}^G \geq 0, \forall g \in \Omega_G, t \in \Omega_T \quad (17)$$

$$\begin{aligned} x_{gt}^G - x_{g(t-1)}^G \geq x_{g\tau}^G, \forall g \in \Omega_G, t \in \Omega_T, t \neq t_1, \\ \tau \in \left[t+1, \min(t + T_g^{U,G} - 1, T)\right], \end{aligned} \quad (18)$$

$$\begin{aligned} x_{g(t-1)}^G - x_{gt}^G \geq 1 - x_{g\tau}^G, \forall g \in \Omega_G, t \in \Omega_T, t \neq t_1, \\ \tau \in \left[t+1, \min(t + T_g^{D,G} - 1, T)\right], \end{aligned} \quad (19)$$

The group of (13)–(19) is related to the physical constraints of every power plant including various types of units. More specifically, the real-time output power of each generating unit at each time slot can be calculated through (13). Equation (14) describes the constraint in which the output power of a generator must be between its minimum and maximum amount when it is running (i.e., $x_{gt} = 1$). Also, equation (15) defines the ramp limit on the increase/decrease of the output power of each generator. Generator start-up and shut-down constraints are modeled through (16)–(17). Finally, depending on the type of the unit, it has minimum up and down time limitations which are mathematically represented via (18)–(19).

Ultimately, most electricity markets do not allow the players to deviate too much from their submitted bids in the day-ahead market [12, 43, 46]. The mathematical model of this constraint is given as:

$$0 \leq P_{gt}^{DA,Dev+,G} \leq \kappa P_{gt}^{DA,G}, \forall g \in \Omega_G, t \in \Omega_T \quad (20)$$

$$0 \leq P_{gt}^{DA,Dev-,G} \leq \kappa P_{gt}^{DA,G}, \forall g \in \Omega_G, t \in \Omega_T \quad (21)$$

where κ (e.g., %20) is the percentage that allows the players to deviate from their day-ahead bids subject to a specified penalty. Different markets may have various regulations which can be mathematically incorporated in the optimization model without the loss of generality. It should be noted that the proposed optimization formulation considers the integrated behavior of all market players including the malicious one. The effect of the attack on the other market players is discussed in Section 5.

4.2 Attacker Type II

As pointed out in Section 3, this type of attacker is a nation state actor whose goal is to maximize the economic damage to the market players by manipulating the system real-time demand through high wattage IoT devices. This attack can target either the generation side or the demand side depending on the ultimate goal of the attacker. The optimization model for attacking the demand side companies (i.e., retailers) is as:

$$\text{maximize} \quad \textit{economic damage} = \sum_{k\in\Omega_K}\left(D_k^{attack}\lambda_k^{RT} - D_k^{actual}\lambda_k^{RT0}\right) - \sum_{k\in\Omega_K}\lambda_k^{Botnet}D_k^{attack} \tag{22}$$

subject to
(10)–(12).

According to this model, the attacker seeks to maximize the economic damage to the retailers through affecting the real-time market prices while keeping his attack stealthy. Similar to this case, the model for attacking the generation side can be defined as:

$$\text{maximize} \quad \textit{economic damage} = \sum_{k\in\Omega_K}\left(G_k^{actual}\lambda_k^{RT0} - G_k^{attack}\lambda_k^{RT}\right) - \sum_{k\in\Omega_K}\lambda_k^{Botnet}D_k^{attack} \tag{23}$$

subject to
(10)–(12).

Note that in both of the aforementioned models we assumed that the attacker can attack either the generation side or the demand side in one day. We can easily modify this assumption by changing the limits of the sums in the objective functions.

5 NUMERICAL ANALYSIS AND DISCUSSION

5.1 Description of the Studied Test Cases

To evaluate the attack scenarios with real-world datasets, we collected market data associated with New York and California ISOs during a one-year (May 2018 – May 2019) period [11, 12, 42, 43]. The historical data is presently available on the ISOs websites and on Bloomberg terminal, and are updated every 5 minutes. Publicly available historical datasets are also typically available in the other electricity markets around the world which makes these markets vulnerable to attacks such as MaMIoT. The California ISO is one of the largest ISOs in the world, which is responsible for delivering roughly 0.300×10^9 MWh of electricity each year to its customers [10]. Similarly, the New York ISO is another large electricity market in the US with 0.156×10^9 MWh of total annual energy consumption [41]. In the following subsections, we will present our analysis of the aforementioned markets. Since the direct implementation of this attack in electricity markets can have huge financial consequences (e.g., 2 million USD per day with a relatively small botnet), we have used the real-world market data to simulate the attack with reasonable and detailed models. This helped us avoid any law-related repercussions while investigating the attack consequences with real-world data.

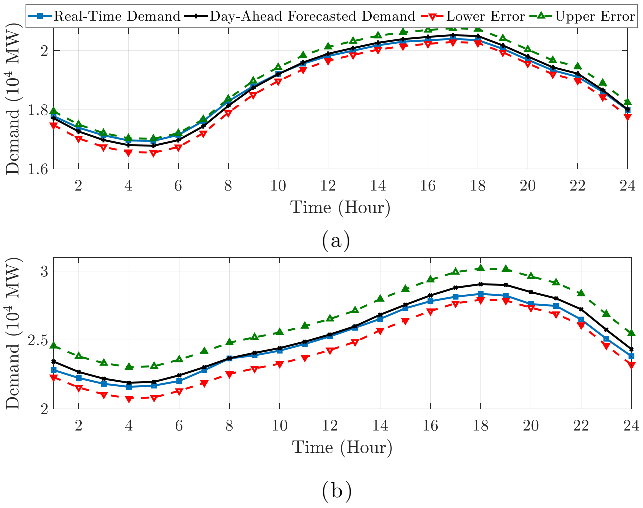

(a)

(b)

Figure 6: Typical load forecasting error band: a) New York ISO (580 MW); b) California ISO (2265 MW).

5.2 Determining the Input Parameters of the Optimization Models

As explained in Section 4.1, a slight deviation of the system's real-time loads from their forecasted value has a linear effect on the real-time market price (see (10)). In order to launch a successful MaMIoT attack, the adversary must first obtain this relationship from the market historical data. In fact, the goal is to determine α_k for the market under investigation. Since the trends in load profiles and market prices change every month, the α_k parameter must also be updated every month. Figure 3 shows the value of this parameter for the California and New York markets for each time interval from the market data on June 2019. This figure was acquired through analyzing the historical data of these markets.

Another important parameter that plays a key role in keeping the attack stealthy is $\Delta D_k^{stealthy,\max}$. According to our analysis, the average prediction error of the system real-time demand at different time slots is 580 MW and 2265 MW in the New York and California ISOs. Figure 6 shows a typical day-ahead forecast and the real-time demand associated with each of the analyzed markets. The dashed lines in the figure indicate the upper and lower prediction errors for each market. The figure illustrates that the load forecasting error band for the California market is higher than that of the New York market. Some reasons for this are: i) the California market is a bigger market and has more maximum power capacity, and ii) the share of flexible loads in the California market is larger than that of the New York market. By limiting the system demand change to the specified error range (typical prediction error), the attacker can make the attack look similar to normal real-time system demand, thereby keeping the attack stealthy and repeatable. Note that in the simulated cases, we considered three different average power consumption for each bot within the botnet (see Figures 7, 13, and 14). The subscripts 1, 2, and 3 of each bar plot in the figures represent 1kW, 2kW, and 3kW for the two markets (NY and CA), respectively.

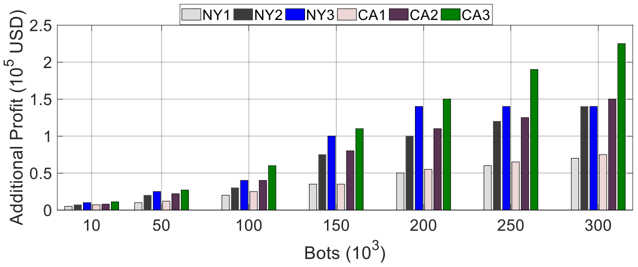

Figure 7: Total additional daily profit of the power plant owner versus the number of compromised high wattage IoT devices.

5.3 Market Player Attacker Results

In this section, we assume that the attacker owns a power plant and can participate in the day-ahead and real-time electricity markets. This power plant includes ten different units whose technical characteristics are given in Appendix IV. The maximum power generation of the power plant is 1990 MW. We simulated the participation of this power plant in the New York and California markets and assumed that the player had control over a high wattage IoT botnet. Figure 7 illustrates the total additional daily profit the power plant owner stands to gain versus the varying numbers of compromised high wattage IoT devices in the botnet. According to this figure, as the size of the botnet increases, the total additional profit increases in both of the studied markets. Our analysis revealed that the power plant owner can gain up to 326,000 USD daily profit (in NYISO) without the implementation of MaMIoT attack ($\Delta D_k^{botnet,\max} = 0$), but with only 200,000 compromised IoT devices (with an average power consumption of 3 kW per device), they could gain an additional 150,000 USD in profit. This is 30% more than the base case without any attacks. By implementing the MaMIoT attack for only 100 days in a year, the studied market player would be able to gain an additional 15 million USD in profit from the electricity market. Interestingly, MaMIoT does not require any specific number of compromised IoT devices to launch a successful attack. This means that the success rate for the attack is 100% with any given botnet size. However, working with a smaller-sized botnet simply results in less additional profit.

Figure 7 shows that with a larger number of compromised IoT devices, the attacker can gain more economic profit from the bigger electricity markets. Another interesting observation from Figure 7 is that the daily additional profit of the power plant owner in the New York market saturates once the botnet size exceeds 200,000 bots. The reason being the attacker can only control 600 MW of system demand with 200,000 bots (with an average power consumption of 3 kW per device). However, according to Figure 6, a stealth attack on the New York ISO can alter a maximum 580 MW of the system's total demand in real-time. So, although the attacker controls over 580 MW with more than 200,000 bots, he is limited to the allowable range (below 580 MW) to keep the attack stealthy. With the maximum demand alteration for botnets greater than 200,000 bots capped at 580 MW, the attack's effect will be the same in all the cases where the botnet size is greater than 200,000 bots.

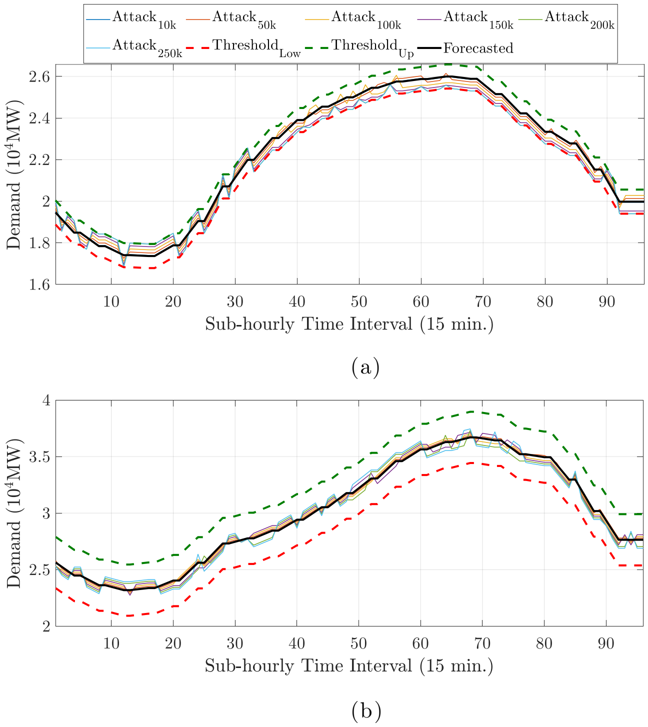

Figure 8: Load profile of the power grid at each time interval associated with attacks launched by the power plant owner with different botnet sizes: a) New York ISO; b) California ISO. Notice how the attack increases and decreases the consumption of energy.

Figure 8 shows the load profile of the system at each time interval associated with different botnet sizes. In this figure, attackers 10k, 50k, 100k, 150k, 200k, and 250k are associated with botnets with 10,000, 50,000, 100,000, 150,000, 200,000, and 250,000 compromised high wattage IoT devices. The figure shows the attacked load profiles are within the specified load forecasting error range and therefore maintain stealth in the proposed attack model. The manipulated system load profile is very similar to typical real-time system demand. This makes it very hard for the market regulator or ISO to detect one player is abusing the market mechanism in his own favor. Such stealth strategies enable the adversary to repeat his attack and multiple times per month and make significant additional profits from the electricity markets.

Figure 9 illustrates the profit breakdown of the adversary in the New York and California markets with different attack scenarios. As can be seen, the overall profit of the attacker in both New York and California markets is maximum when the adversary uses the optimization-based attack. The baseline attack excludes the key constraints in the attack scenario, and hence, causes monetary penalties from the market. The optimization-based attack on the other hand has zero penalties in both markets, which leads to the maximum profit for the malicious market player.

To illustrate the interaction between multiple market players in the New York market, we considered 21 generation players and 20 consumer players. For the first case, let's assume one of the generation players is malicious and can control a botnet of high

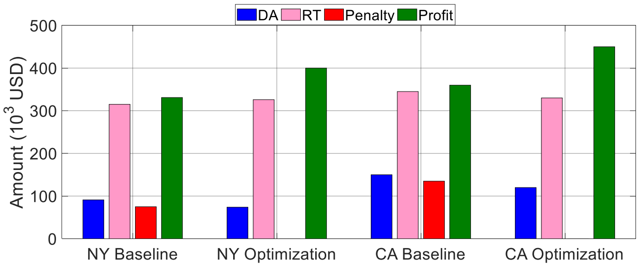

Figure 9: The profit breakdown of the simulated market player in a single day with different bidding strategies in the New York and California markets. DA: day-ahead profit, RT: real-time profit, Penalty: market penalties, Profit: overall profit.

wattage devices. Figure 10 shows the overall daily profit of the market players during the no attack and optimization-based attack scenarios. As it can be seen, the manipulations of the malicious market player increase the gain of the other generation players in the market as well. However, since the adversary knows about the manipulated real-time prices in advance, he prepares for the manipulated situation and obtains the maximum profit out of that. The consumer market players lose small profits because of this market manipulation.

In the other simulated case, we assumed that the adversary is one of the consumer players in the New York market. Accordingly, one of the 20 players on the consumer side is malicious and can control a high wattage IoT botnet. Figure 11 shows the overall daily profit of the market players during the no attack and optimization-based attack scenarios. As it can be seen, the malicious market player gains the maximum profit from the attack while the other consumer players gain marginal profit from the manipulations. Conversely, the benign generation players lose small profits because of the price manipulations.

Finally, as it was discussed in Section 4.1, the day-ahead price forecasts are used to determine the optimal attack scenario by the malicious market player. Here, we aim to analyze the effect of prediction error in this parameter on the attack's gain. Figure 12 shows the daily profit of the attacker in both markets versus the estimation error in the day-ahead market prices. according to this figure, the adversary's gain does not change significantly with the increase in the estimation error. This observation illustrates that we made a reasonable assumption in our formulation to consider this parameter in the optimization model.

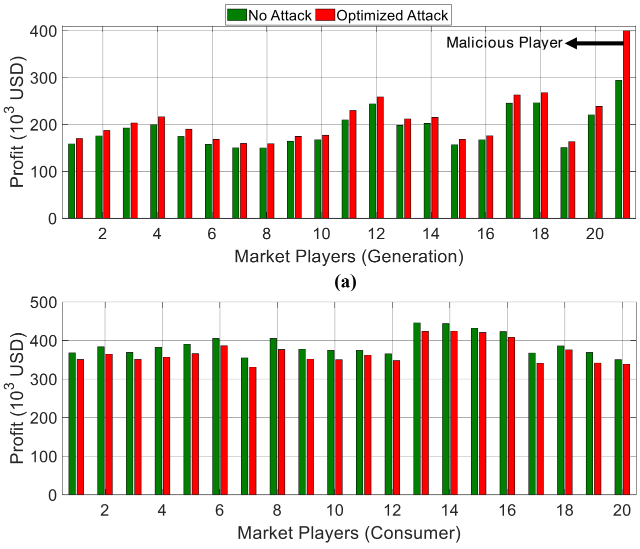

(a)

(b)

Figure 10: The daily profit of the market players in the New York ISO (only 21st generation player is malicious). a) Generation players, and b) Consumer players.

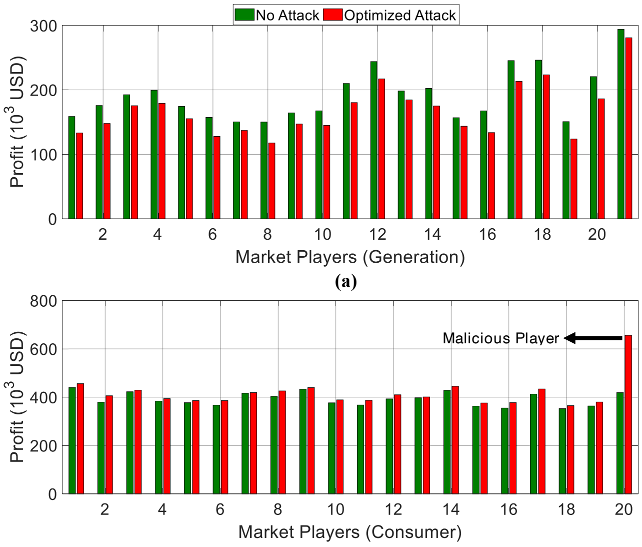

(a)

(b)

Figure 11: The daily profit of the market players in the New York ISO (only 20th consumer player is malicious). a) Generation players, and b) Consumer players.

5.4 Nation-State Attacker Results

As explained in Section 4.2, this type of attacker is a nation-state actor who can target the generation or demand-side players in a specific electricity market. To attack the demand side, we executed the first optimization model with the objective function given in (22). Figure 13 shows the total daily economic damage that the attacker can impose on the demand side players of the studied markets versus the number of compromised high wattage IoT devices. According to this figure, with only 200,000 compromised IoT devices, the attacker can impose 3.5 million USD and 5 million USD daily economic damage to the California and New York markets, respectively. If we simulate the attacker performing the attack 100 days per year, the annual economic damage would be 350 million USD and 500 million USD for the California and New York markets. From the figure, we see the economic damage to the California market is higher than that of the New York market when the size of the botnet is big enough (more than 270,000 compromised devices). Note that the attacker can impose this huge economic damage on

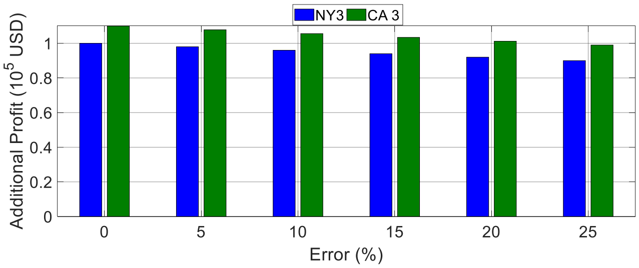

Figure 12: Total additional daily profit of the malicious market player versus the estimation error in the day-ahead market price. This plot shows that the effect of the prediction error in the attack is not significant.

the studied markets while his attack is still stealthy. Interested readers can refer to Appendix IV.1 to see the load profile of the system during the attack scenarios.

The nation-state attacker can also target the players in the generation side of the market. To evaluate this attack on the studied markets, we executed the proposed optimization model with the objective function given in (23). Figure 14 shows the total daily economic damage to the generation companies in each of the studied markets versus the number of compromised IoT devices that the attacker controls. According to this figure, with only 200,000 compromised IoT devices, the attacker can impose 2.8 million USD and 2.9 million USD economic damage to the generation companies in the California and New York ISOs, respectively. Similar to the demand side attack and with the assumption that the attacker will launch MaMIoT attack on the studied markets 100 days per year, the total annual economic damage will be 280 million USD and 290 million USD in the California and New York markets, respectively. The attacker can cause greater damage in the California market than the New York market once the botnet size exceeds 220,000 compromised devices. Even with a small number of compromised IoT bots, the attacker can still cause notable damage to the studied markets. For example, if the botnet includes 10,000 bots (with 3 kW average power consumption for each bot), the annual economic damage to the generation companies will be 1.75 million USD and 2.5 million USD in the California and New York markets, respectively. To achieve this, we assume that the attacker will launch MaMIoT attack 100 days per year. It is worth mentioning that the SO is not able to detect the attack in any of the simulated scenarios as the system load profile is very similar to typical real-time system demand. Interested readers can refer to Appendix IV.1 for further detailed analysis on the stealthiness of the MaMIoT attack in the generation side companies of the studied markets.

6 COUNTERMEASURES

While currently there is no single effective countermeasure to prevent the MaMIoT attack, a combination of the following strategies could be employed to reduce its damaging consequences.

Section 5 illustrates the economic consequence from attacker II is much more detrimental than attacker I. Attackers in class II are more likely to occur in real-world scenarios because of the reduced concern for negative legal repercussions, such as prosecution.

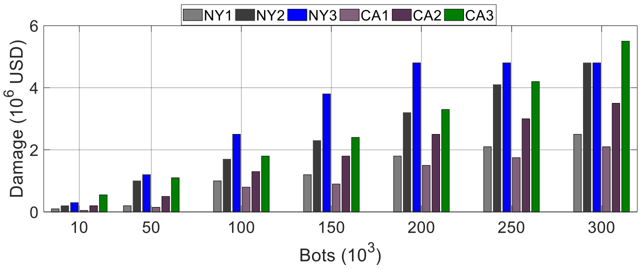

Figure 13: Total daily economic damage that the nation state attacker can impose on the demand side of the studied markets versus the number of compromised high wattage IoT devices.

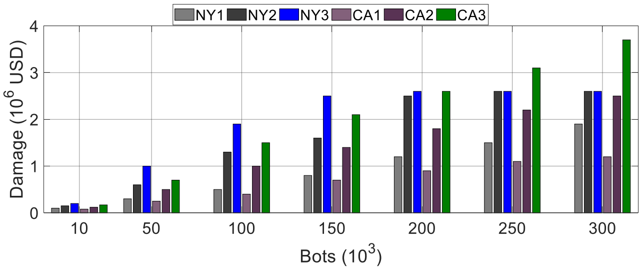

Figure 14: Total daily economic damage that the nation state attacker can impose on the generation side of the studied markets versus the number of compromised high wattage IoT devices.

Therefore reducing the effect and possibility of nation-state attackers is the first priority in determining countermeasures. Publicly available historical market data is one of the biggest contributing factors for making the MaMIoT attack possible. To eliminate the risk of nation state attackers, the ISOs should only release detailed market data to market players. This new data privacy plan would add the first barrier for nation state attackers to get access to recent historical market data for estimating price sensitivity and other crucial parameters required to launch a successful stealth attack. Without this information (α_k and $\Delta D_k^{stealthy,\max}$), the economic consequence of an undetectable attack is limited. An intelligent attacker would be forced to launch an overly conservative attack to maintain stealth, causing minimal demand changes.

Figures 15 and 16 show the daily economic damage of the attacker type II on both simulated markets versus the estimation error in the stealth and price-load sensitivity parameters, respectively. As can be seen, the influence of the attack severely declines following the increase in the estimation error of the key parameters (~50% influence decline when there is 25% estimation error). These results verify the partial effectiveness of the data privacy countermeasure discussed in the previous paragraph.

While tightening access to historical market data will thwart many attackers, it may also prevent researchers and market analysts from performing analyses on these markets. To avoid this, a more practical solution would be releasing redacted or altered versions of the market data or even delaying the release of the full datasets, such that it cannot be used in real-time. This would significantly reduce the effectiveness of the MaMIoT attack by a nation-state actor. This strategy would make it very hard for the attacker to estimate

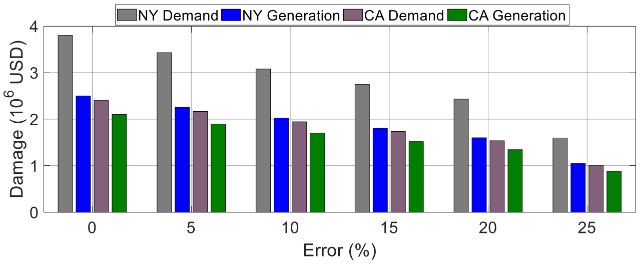

Figure 15: Total daily economic damage of the nation state adversary in the simulated markets versus the estimation error in the stealth parameter ($D^{stealthy,\mathbf{max}}$).

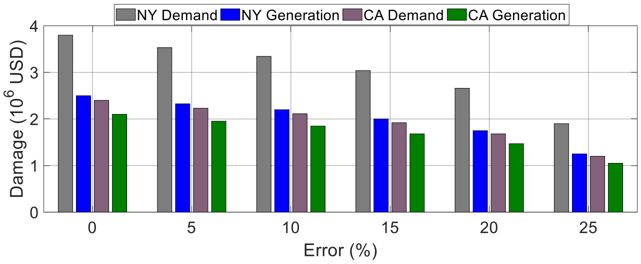

Figure 16: Total daily economic damage of the nation state adversary in the simulated markets versus the estimation error in the price-load sensitivity parameter (α).

the crucial parameters of the optimization models reliably. As an illustrative example, our analysis shows that releasing the down-sampled (i.e., every 2 hours instead of every 5 minutes) version of the market data with a month delay can decrease the attack economic damage up to 87%.

The most effective and practical countermeasure against MaMIoT attacks is to develop and install non-intrusive load monitoring (NILM) or non-intrusive appliance load monitoring (NIALM) algorithms on the electricity meters of homes in the power grid. NILM and NIALM can be defined as the process of analyzing voltage and current going into a house (through the electricity meters) and deducing what appliances are used at which times in the house as well as their individual energy consumption [3]. These algorithms have been traditionally developed to help the home owners and/or utility companies optimize the energy usage of the home and minimize their monthly electricity bill. It goes without saying that NILM is considered a low-cost alternative to attaching individual monitors on each appliance. With the recent advancements in the field of machine learning, especially with the introduction of deep learning, reliable NILM algorithms can be developed to quickly detect the MaMIoT attacks and inform the suspicious activities to the home owner and utility companies. For example, the NILM can easily reveal the suspicious use of electric oven in the morning when the home owner is at work and detect it as an anomaly in the meter's data. Of course, further detailed analysis is needed to design and tune reliable and state-of-the-art NILM algorithms to be used in practice. A sample data of a residential customer which can be used in the NILM attack detection is shown in Figure 17 [33]. To address the privacy concerns of the customers, the developed machine learning can learn about the energy usage pattern without a specific reference to the used devices in the house. In such cases, the issued alert by the trained model will let the home owner to know there is an authorized use of the devices in the house without point to a specific device.

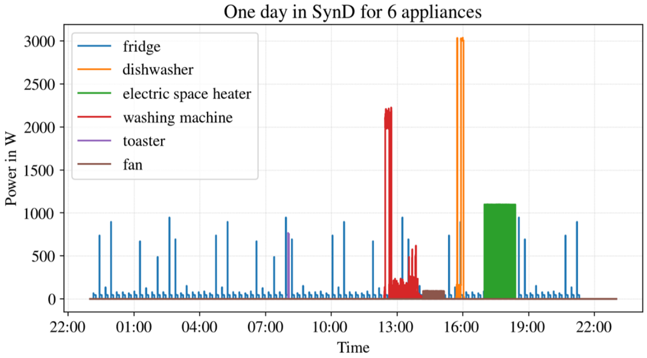

Figure 17: A sample data of a residential customer which can be used in the NILM attack detection [33].

7 CONCLUSIONS

In this paper, we introduced MaMIoT, the first energy market manipulation cyberattack in which an adversary can slightly alter the power system real-time demand through a botnet of high wattage IoT devices to help market players gain additional profit from the electricity market or cause major economic damage to a set of market players. We evaluated the attack models on real datasets from the two big electricity markets in the U.S., the California and New York markets. The simulation results revealed that with only 200,000 bots in a botnet, the attacker can cause 2.8 (2.1) million USD and 3.8 (2.2) million USD worth of economic damage to the demand (generation) side players of the California and New York markets, respectively. We also showed that the MaMIoT attack can help a typical power plant owner gain an additional 30% in profit from the energy market, all while maintaining attack stealth for increased repeatability.

We hope that this paper raises awareness of the significance of MaMIoT attacks to the market operators, ISOs, IoT manufacturers, and system security experts to make the electricity markets more secure against cyberattacks. In the near future, this problem will be even more critical as the number of smart appliances with Internet connectivity continues to grow.

ACKNOWLEDGEMENTS

The authors would like to sincerely thank FERC, NERC, Reliability First Corporation, and anonymous reviewers for their insightful comments and fruitful discussions which significantly improved the quality of the paper. This work was partially supported by the following grants: NSF CNS-1929580, NSF CNS-1929406, NSF CNS-1929410, and NSF CNS-1931573.

REFERENCES

[1] -. 2019. Bloomberg Terminal. https://en.wikipedia.org/wiki/Bloomberg_Terminal

[2] . 2019. DDOS Attacks against Global Markets. https://www.akamai.com/us/en/multimedia/documents/secure/ddos-attacks-against-global-markets-white-paper.pdf

[3] EJ Aladesanmi and KA Folly. 2015. Overview of non-intrusive load monitoring and identification techniques. *IFAC-PapersOnLine* 48, 30 (2015), 415–420.

[4] Sajjad Amini, Fabio Pasqualetti, and Hamed Mohsenian-Rad. 2016. Dynamic load altering attacks against power system stability: Attack models and protection schemes. *IEEE Trans. Smart Grid* 9, 4 (2016), 2862–2872.

[5] Manos Antonakakis et al. 2017. Understanding the Mirai botnet. In *26th USENIX Security Symp.* 1093–1110.

[6] Kankar Bhattacharya, Math HJ Bollen, and Jaap E Daalder. 2012. *Operation of restructured power systems.* Springer Science & Business Media.

[7] Security Boulevard. 2018. *Here's how anyone with $20 can hire an IoT botnet to blast out a week-long DDoS attack.* https://securityboulevard.com/2018/08/heres-how-anyone-with-20-can-hire-an-iot-botnet-to-blast-out-a-week-long-ddos-attack/

[8] Paul J Burke and Ashani Abayasekara. 2018. The price elasticity of electricity demand in the United States: A three-dimensional analysis. *The Energy Journal* 39, 2 (2018).

[9] Buyexerciser. 2020. *Treadmill workout tips: How long should I run on the treadmill?*

[10] California Independent System Operator. 2019. California Independent System Operator. https://en.wikipedia.org/wiki/California_Independent_System_Operator

[11] California Independent System Operator. 2019. Energy Market & Operation Data. http://oasis.caiso.com/mrioasis/logon.do

[12] California Independent System Operator. 2019. Reliability Requirements. http://www.caiso.com/planning/Pages/ReliabilityRequirements/Default.aspx#Historical

[13] Catalin Cimpanu. 2016. *You Can Now Rent a Mirai Botnet of 400,000 Bots.* https://www.bleepingcomputer.com/news/security/you-can-now-rent-a-mirai-botnet-of-400-000-bots/

[14] Julie Creswell and Robert Gebeloff. 2014. Traders profit as power grid is overworked. *The New York Times* (2014).

[15] Adrian Dabrowski, Johanna Ullrich, and Edgar R Weippl. 2017. Grid shock: Coordinated load-changing attacks on power grids: The non-smart power grid is vulnerable to cyber attacks as well. In *Proc. of the 33rd Ann. Computer Security Applications Conf. (ACSAC).* 303–314.

[16] Scott DiSavino. July 2013. *JPMorgan to pay $410 million to settle power market case.* https://www.reuters.com/article/us-jpmorgan-ferc/jpmorgan-to-pay-410-million-to-settle-power-market-case-idUSBRE96T0NA20130730

[17] Yury Dvorkin and Siddharth Garg. 2017. IoT-enabled distributed cyber-attacks on transmission and distribution grids. In *North American Power Symp. (NAPS).* 1–6.

[18] Energy Efficiency and Renewable Energy Clearinghouse. 2020. *Energy Use of Some Typical Home Appliances.* http://sites.science.oregonstate.edu/~hetheriw/energy/quick/eff/EREC_Brief_Energy_Use_of_Some_Typical_Home_Appliances.htm

[19] Timothy Egan. 2005. Tapes show Enron arranged plant shutdown. *New York Times* (2005).

[20] We Energies. 2020. *Appliance savings with Time-of-Use.* https://www.we-energies.com/residential/acctoptions/tou_wi_shiftappli.htm

[21] Maureen Farrell. July 2013. *JPMorgan settles electricity manipulation case for $410 million.* https://money.cnn.com/2013/07/30/investing/jp-morgan-electricity-fines/index.html

[22] Laundry Butler for You. [n.d.]. *How Much Laundry Does the Average Person Do?*

[23] GE. [n.d.]. GE Wi-Fi connect appliances. https://www.geappliances.com/ge/connected-appliances/

[24] Dan Goodin. 2017. *Assessing the threat the Reaper botnet poses to the Internet—what we know now.* https://arstechnica.com/information-technology/2017/10/assessing-the-threat-the-reaper-botnet-poses-to-the-internet-what-we-know-now/

[25] Dan Goodin. 2018. *New IoT botnet offers DDoSes of once-unimaginable sizes for $20.* https://arstechnica.com/information-technology/2018/02/for-sale-ddoses-guaranteed-to-take-down-gaming-servers-just-20/

[26] Dan Goodin. December 2017. *100,000-strong botnet built on router 0-day could strike at any time.* https://arstechnica.com/information-technology/2017/12/100000-strong-botnet-built-on-router-0-day-could-strike-at-any-time/

[27] Martin Holladay. 2013. *Garage Door Openers Are Always On.*

[28] Bing Huang, Alvaro A Cardenas, and Ross Baldick. 2019. Not everything is dark and gloomy: Power grid protections against IoT demand attacks. In *28th USENIX Security Symp.* 1115–1132.

[29] imperva. 2019. *Booters, Stressers and DDoSers.* https://www.imperva.com/learn/application-security/booters-stressers-ddosers/

[30] Rommel Joven and Evgeny Ananin. 2018. *DDoS-for-Hire Service Powered by Bushido Botnet.* https://www.fortinet.com/blog/threat-research/ddos-for-hire-service-powered-by-bushido-botnet-.html

[31] John Kennedy. [n.d.]. https://www.siliconrepublic.com/enterprise/dragonfly-us-russia-energy-grid-hackers. https://www.symantec.com/blogs/threat-intelligence/dragonfly-energy-sector-cyber-attacks

[32] Daniel Sadi Kirschen and Goran Strbac. 2004. *Fundamentals of power system economics.* Vol. 1. Wiley Online Library.

[33] Christoph Klemenjak, Christoph Kovatsch, Manuel Herold, and Wilfried Elmenreich. 2020. A synthetic energy dataset for non-intrusive load monitoring in households. *Scientific Data* 7, 1 (2020), 1–17.

[34] KrebsonSecurity. [n.d.]. *Did the Mirai Botnet Really Take Liberia Offline?* https://krebsonsecurity.com/2016/11/did-the-mirai-botnet-really-take-liberia-offline/

[35] Robert M Lee, Michael J Assante, and Tim Conway. 2016. ICS Defense Use Case: Analysis of the cyber attack on the Ukrainian power grid. *Electricity Information Sharing and Analysis Center, SANS ICS* (2016).

[36] Jessica Lietz. 2018. *How Much Does the Hot Water Heater Affect an Electric Bill?* https://homeguides.sfgate.com/much-hot-water-heater-affect-electric-bill-88704.html

[37] Craig Lloyds. 2018. *How Much Electricity Do All Your Appliances Use?*

[38] Bethany McLean and Peter Elkind. 2013. *The smartest guys in the room: The amazing rise and scandalous fall of Enron.* Penguin.

[39] Amir-Hamed Mohsenian-Rad and Alberto Leon-Garcia. 2011. Distributed internet-based load altering attacks against smart power grids. *IEEE Trans. Smart Grid* 2, 4 (2011), 667–674.

[40] T Mulligan. 2002. How Enron Manipulated State's Power Market. *Los Angeles Times* (2002).

[41] New York Independent System Operator. [n.d.]. Annual Report. https://www.nyiso.com/documents/20142/2223020/2018-Power-Trends.pdf/4cd3a2a6-838a-bb54-f631-8982a7bdfa7a

[42] New York Independent System Operator. 2019. Energy Market & Operation Data. https://www.nyiso.com/energy-market-operational-data

[43] New York Independent System Operator. 2019. Load Data. https://www.nyiso.com/load-data

[44] Union of Concerned Scientists. [n.d.]. *Electric Vehicle Charging Types, Time, Cost and Savings.*

[45] Office of Enforcement Federal Energy Regulatory Commission Washington, D.C. 2019. 2018 Report on Enforcement. https://www.ferc.gov/legal/staff-reports/2018/11-15-18-enforcement.pdf?csrt=4611620575164854265

[46] Pennsylvania and New Jersey Independent System Operator. 2019. Energy Market. https://www.pjm.com/markets-and-operations/energy.aspx

[47] Troutman Pepper. November 2017. *FERC Approves $105 Million Settlement with Barclays for Market Manipulation.* https://www.lexology.com/library/detail.aspx?g=79b6712f-2db8-415e-9a93-6307c086d5a6

[48] Payless Power. 2019. *HOW MANY WATTS DOES A REFRIGERATOR USE.*

[49] Radware. 2018. *A Quick History of IoT Botnets.* https://blog.radware.com/uncategorized/2018/03/history-of-iot-botnets/

[50] Mohammad Shahidehpour, Hatim Yamin, and Zuyi Li. 2003. *Market operations in electric power systems: forecasting, scheduling, and risk management.* John Wiley & Sons.

[51] Saleh Soltan, Prateek Mittal, and H Vincent Poor. 2018. BlackIoT: IoT botnet of high wattage devices can disrupt the power grid. In *27th USENIX Security Symp.* 15–32.

[52] Alireza Soroudi. [n.d.]. *Power system optimization modeling in GAMS.* Springer.

[53] Statistica. 2019. Number of homes with smart thermostats in North America from 2014 to 2020 (in millions). https://www.statista.com/statistics/625868/homes-with-smart-thermostats-in-north-america/

[54] HVAC Talk. 2019. *How many hours should the AC run during the hottest days of the year?*

[55] US Energy Information Administration. 2019. U.S. energy facts explained . https://www.eia.gov/energyexplained/us-energy-facts/

[56] US Energy Information Administration. 2019. Wholesale electricity prices were generally lower in 2019, except in Texas. https://www.eia.gov/todayinenergy/detail.php?id=42456#

[57] US Energy Information Administration. 2020. 2018 Average Monthly Bill- Residential. https://www.eia.gov/electricity/sales_revenue_price/pdf/table5_a.pdf

[58] Stylianos I Vagropoulos and Anastasios G Bakirtzis. 2013. Optimal bidding strategy for electric vehicle aggregators in electricity markets. *IEEE Trans. Power Syst.* 28, 4 (2013), 4031–4041.

[59] Christian Vasquez. June 2020. *'Major vulnerability': EV hacks could threaten power grid.* https://www.eenews.net/stories/1063401375

[60] Whirlpool. 2020. *How long do dishwashers run?* https://www.whirlpool.com/blog/kitchen/how-long-do-dishwashers-run.html

[61] Zhang Xu, Haining Wang, Zichen Xu, and Xiaorui Wang. 2014. Power Attack: An Increasing Threat to Data Centers. In *Network and Distributed System Security (NDSS) Symp.* 1–15.

[62] Carter Yagemann, Simon P Chung, Erkam Uzun, Sai Ragam, Brendan Saltaformaggio, and Wenke Lee. 2020. On the Feasibility of Automating Stock Market Manipulation. In *Annual Computer Security Applications Conference.* 277–290.

[63] Mark Zeller. 2011. Myth or reality – Does the Aurora vulnerability pose a risk to my generator?. In *64th Ann. Conf. for Protective Relay Engineers.* 130–136.

[64] Kim Zetter. July 2018. *Inside the Cunning, Unprecedented Hack of Ukraine's Power Grid.* https://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/

Appendix I NOMENCLATURE

The notation used throughout the paper is reproduced below for quick reference.

Indices and Sets

d, Ω_D	Index and set of submitted demand bids to the electricity market.
s, Ω_S	Index and set of submitted generation offers to the electricity market.
g, Ω_G	Index and set of generating units within the power plant.
t, Ω_T	Index and set of hourly time intervals.
k, Ω_K	Index and set of sub-hourly time intervals (e.g., 15 min.).
Ω_k	Set of hourly time intervals associated with k^{th} sub-hourly time interval. For example, $\Omega_3 = 1$ or $\Omega_5 = 2$ for 15 min. sub-hourly time intervals.
$\mathcal{K}$	Number of sub-hourly time slots within one hour, $\mathcal{K} = \frac{card(\Omega_K)}{card(\Omega_T)}$ (e.g., $\mathcal{K} = 4$ for 15 min. sub-hourly time intervals).

Parameters

D^{actual}	Actual demand of the system before attack.
$P^{D,\mathbf{max}}/P^{S,\mathbf{max}}$	Submitted energy quantity bids/offers to the market by its players.
$P^{\mathbf{min}}/P^{\mathbf{max}}$	Minimum/maximum output power of generating units.
R^U/R^D	Ramp-up/down limit of generating units.
$T^{U,G}/T^{D,G}$	Minimum up/down time of generating units within the power plant.
α	Constant coefficient defining the dependency of real-time market price on the system demand change at each time slot.
κ	The percentage of the day-ahead bids which are allowed to be deviated in real-time operation with specific penalties.
λ^D/λ^S	Submitted demand/supply price to the market.
λ^{DA}	Hourly day-ahead market price.
$\lambda^{RT,0}$	Sub-hourly expected real-time market price before the load alteration attack.
$\lambda^{DA,Dev}$	Penalty price for the day-ahead bid/offer deviation.
$\lambda^{SU}/\lambda^{SD}$	Start-up/shut-down cost of generating units.
$\lambda^{G,Constant}$	Constant running cost of generating units.
$\lambda^{G,Fuel}$	Fuel price of generating units.
Δt	Duration of each sub-hourly time interval within one hour (i.e., 0.25 for 15 min. sub-hourly intervals).
$\Delta D^{stealthy,\mathbf{max}}$	Average of the load forecasting error at each time slot.
$\Delta D^{botnet,\mathbf{max}}$	Maximum capability of the attacker in changing system demand at each time slot.

Variables

P^D/P^S	Accepted energy bids/offers in the market.
$profit^G$	Total profit of the power plant participating in the day-ahead and real-time markets.
$profit^{DA,G}$	Hourly profit of the generating units obtained from the day-ahead market.
$profit^{RT,G}$	Sub-hourly profit of the generating units obtained from the real-time market.
$P^{DA,G}$	Hourly energy quantity offers of generating units in the day-ahead market.
$P^{RT,G}$	Sub-hourly energy quantity offers of generating units in the real-time market.
$P^{Dev+/-,G}$	Hourly positive/negative deviation of generating unit output power from the accepted day-ahead offers.
$P^{Act,G}$	Sub-hourly real-time output power of generating units.
λ^{RT}	Sub-hourly real-time market price affected by MaMIoT.
ΔD^{System}	Sub-hourly system demand alteration through MaMIoT.
D^{attack}	Sub-hourly system demand following MaMIoT attack.
u^G/v^G	Binary indicator for start-up/shut-down of generating units (e.g., $u^G_{gt} = 1$ means generator g starts up at hour t and $v^G_{gt} = 1$ means generator g shuts down at hour t).
x^G	Scheduled status of generating units (e.g, $x^G_{gt} = 1$ denotes unit g is running at hour t).

Appendix II STEALTH STRATEGIES

In order to make the MaMIoT attack repeatable and add to the motivation of the attackers, the adversary can employ several strategies, alone or in combination. Some of the practical strategies are outlined as follows:

Appendix II.0.1 From the End-User's Perspective. It goes without saying that the attacker should try to hide his activity from the compromised homes. To achieve this goal, one effective strategy would be the use of compromised high wattage IoT devices when the awareness of the home owner is very low. According to the typical time of use for some popular categories of high wattage home IoT devices summerized in the Appendix, it can be surmised that there are many opportunities for botnet attacks outside of the normal time of use which would be undetected by an end user. While some HVAC devices such as AC and heaters tend to run on/off all day, others such as an EV charger may only consume power during "after work" hours when end users are home.

In order to conceal additional device usage for limited period of time (i.e., 1-3 hours on average), the attacker can classify the compromised IoT devices and leverage their potential based on their availability time. For example devices such as ovens are used during hours when presumably no one is in the kitchen (1-4AM) while devices such as EV chargers can be used during the night when the EV is connected to the grid. Some of these devices such the EV charger have been proven to have a great potential in these attacks [59].

Appendix II.0.2 From the Market Operator's Perspective. Additionally, the attacker needs to hide his activity from the market operator. The following items list some of the practical strategies in this category.

I) Smooth Load Profile Changes: The main way the system operator (SO) can detect the MaMIoT attack, is to analyze the daily load profile of the system. A naive attacker changes the system demand without considering any limitations, which might lead to a noticeable difference between the attacked load profile and a typical benign one. In this paper, we formulate the model such that the attacked load profile of the system becomes very similar to a typical daily load profile, making it very challenging for the SO to detect any abnormalities in the system (see Section 5 for numerical results).

II) The Frequency of Attack: As the frequency of the attack increases, the possibility of it being caught by SO increases as well. A smart attacker will launch the MaMIoT attack only for a certain number of days (e.g., 100) in each year. By doing this the attack days can be determined randomly, making it hard for the SO to determine which days are normal and which days the market is attacked.

III) Choosing a Suboptimal Attack Scenario: In this strategy, the attacker does not implement the optimal attack scenario on the market. Instead, he sacrifices a portion of his profit to make his attack stealthier. To achieve this, the attacker runs the proposed optimization model and chooses a suboptimal point (e.g., 80% of the optimal point).

IV) Targeting Other Players: In this strategy, the attacker occasionally maximizes the profit of the other players in the market to defer the suspicion of the SO onto them. These players can be the competitors of the attacker or the entities whose loss result in economic benefit for the attacker.

Appendix III ADDITIONAL EVIDENCE ON HIGH WATTAGE IOT BOTNET

The domain space of high wattage IoT devices for use in the IoT Skimmer attack is very large. Fig. A1 shows the trend of the growing number of houses with smart thermostats in North America alone, representing a fraction of the total high wattage IoT devices in the automation field [53]. Concerning the matter of location, attackers can trivially determine whether a compromised device is within a certain geographical area through the device's IP address.

Table AI gives a breakdown of some advertised and estimated costs for utilizing DDoS for hire and IoT botnet rental services. This table shows how the commoditization of cybercrime has made it feasible to launch attacks for less than the cost of most cyber certifications. It is worth mentioning that although the botnets presented in Table AI are not necessarily built from high wattage IoT devices, the given numbers in the table can still be used for estimating the cost of building/renting a typical high wattage IoT botnet.

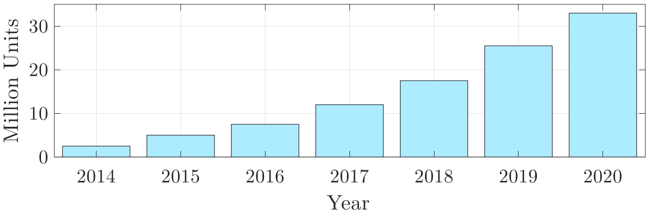

Figure A1: The growing trend of homes with smart thermostats in the North America region [53].

Table AII shows the typical time of use for some popular categories of high wattage home IoT devices. From the table it can be surmised that there are many opportunities for botnet attacks outside of the normal time of use which would be undetected by an end user. While some HVAC devices such as AC and heaters tend to run on/off all day, others such as an EV charger may only consume power during "after work" hours when end users are home.

In order to conceal additional device usage for limited period of time (i.e., 1-3 hours on average), the attacker can classify the compromised IoT devices and leverage their potential based on their availability time. For example devices such as ovens are used during hours when presumably no one is in the kitchen (1-4AM) while devices such as EV chargers can be used during the night when the EV is connected to the grid (see Table AII). Some of these devices such the EV charger have been proven to have a great potential in these attacks [59].

Appendix IV POWER PLANT SIMULATION DATA

The simulated power plant consists of ten different units (generators) with the technical characteristics given in Table AIII [52]. In this table, the units of the given parameters in the first row from left to right are USD/MWh, USD, USD, USD, MW/hr, MW/hr, hr, hr, MW, MW, hr, nothing, and hr. Also, U_g^0 denotes time periods unit g has been on at the beginning of the planning horizon (end of hour 0). Similarly, S_g^0 represents the time periods that unit g has been shut-down at the beginning of the planning horizon.

Appendix IV.1 Stealthiness of the Attack

In this section, we added the additional simulation results illustrating the stealth of the MaMIoT attack for attacker II (nation state attacker). Figure A2 depicts the load profile of the studied electricity markets under different levels of MaMIoT attacks on the demand side companies and further illustrates how all of the attack scenarios stay within a normal load forecasting error range. As can be seen in the figure, since the system demand change in the California ISO is much less sensible than the New York ISO, the attack detection in the California market will be a harder process.

The load profile of the California and New York ISOs under different levels of MaMIoT attacks on the generation side companies is represented in Figure A3. Similar to the demand side attack, the load profile of different attacks are within the normal load forecasting error range. As a general rule, which is true in most of the time intervals, the nation state attacker can harm the demand side companies by increasing the real-time market system demand. On the other hand, decreasing the system real-time demand will

Table AI: IoT Botnet Rental and DDoS for Hire Cost Breakdown

Name	Botnet Size	Rental Cost	Duration	Bandwidth	Type of Bots
JenX [25]	-	\$20/target	-	295Gbps	small/office routers
Mirai variant [13]	50k	\$3-4000/2 weeks	1 hour	-	cameras, routers, DVRs, etc.
Bushido [30]	20k	\$20-150/month	-	500Gbps	cameras, routers, DVRs, etc.
Reaper [24]	30k	-	-	-	cameras, routers, DVRs, etc.
Satori [26]	100k	-	-	-	small/office routers
Estimate for IoT Botnet Services [7]	-	~\$15/week	-	300Gbps	-
Estimate for DDoS Services [29]	-	\$20-45/month	1 hour	220Gbps	-

Table AII: High Wattage consumer IoT Device Availability [18]. Wattage represents maximum per device.

Smart IoT Device	Energy Consumption (W)	Peak Use Time	Avg Use Length	Time to Attack
Water Heater [36]	5000	Morning	3h/day	Early Morning
AC [54]	1000	All-day	9h/day	Anytime
Garage Opener [27]	1100	All-day	3min/day	Midday
Fridge [48]	900	All-day	24h/day	Midday
Heater [20]	1500	Evening	3h/day	Anytime
EV charger [44]	12000	Evening	8h/day	Early Morning
Oven and Stove [37]	4000	Evening	1h/day	Early Morning
Washer [22]	1200	Sporadic	2h/wk	Early Morning
Dryer [22]	1800	Sporadic	2h/wk	Early Morning
Dishwasher [60]	852	Sporadic	120min/day	Early Morning
Treadmill [9]	735	Sporadic	90min/wk	Early Morning

Table AIII: Technical Data of the Simulated Power Plant Units [52]

Unit	$\lambda_g^{G,Fuel}$	$\lambda_g^{G,Constant}$	λ_g^{SU}	λ_g^{SD}	R_g^U	R_g^D	$T_g^{U,G}$	$T_g^{D,G}$	$P_g^{\min}$	$P_g^{\max}$	U_g^0	$x_{g(t=0)}^G$	S_g^0
g_1	12.1	82	42.6	42.6	80	80	3	2	80	200	1	0	1
g_2	12.6	49	50.6	50.6	120	120	4	2	120	320	2	0	0
g_3	13.2	100	57.1	57.1	50	50	3	2	50	150	3	0	3
g_4	13.9	105	47.1	47.9	250	250	5	3	250	520	1	1	0
g_5	13.5	72	56.6	56.9	80	80	4	2	80	280	1	1	0
g_6	15.4	29	141.5	141.5	50	50	3	2	50	150	0	0	0
g_7	14	32	113.5	113.5	30	30	3	2	30	120	0	1	0
g_8	13.5	40	42.6	42.6	30	30	3	2	30	110	0	0	0
g_9	15	25	50.6	50.6	20	20	0	0	20	80	0	0	0
g_{10}	14.3	15	57.1	57.1	20	20	0	0	20	60	0	0	0

lead to economic damage to the generation side companies in the electricity markets.

Appendix V RELATED WORK

Appendix V.1 Attacks on Financial Markets and Historical Electricity Market Manipulation Cases

Financial markets have been recently a popular target for cybercriminals around the world. In this line, hackers have leveraged the concept of market manipulation to affect the specific market players or the entire market with the aim of gaining monetary profits or causing financial damage to the market players. Market manipulation can be defined as the deliberate and malicious interference with the market values to create an artificial price for a tradable entity [2]. One of the main ways employed by cybercriminals to implement the market manipulation attack in financial markets is the DDoS attack. In this attack the adversary deliberately reduces the availability of products and/or services from a targeted company or even an entire financial exchange platform, to affect the associated stock prices. Many companies which deliver services to their clients via online or web applications could fall victim. In this type of attack, while the victim does not experience physical loss,

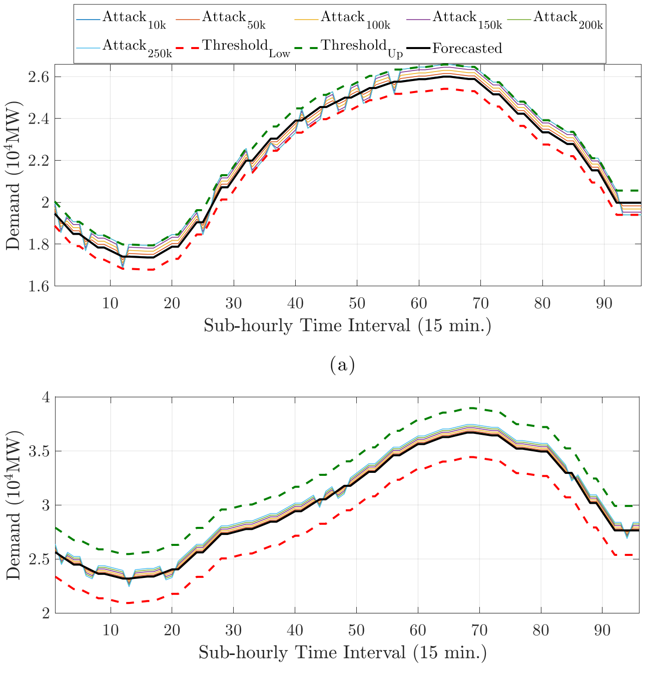

(a)

(b)

Figure A2: Load profile of the power grid at each time interval associated with attacks on the demand side companies with different botnet sizes: a) New York ISO; b) California ISO.

they could be severely affected by the negative consequences of service unavailability and reduced investor confidence.

The biggest market manipulation attack campaign which leveraged the DDoS attack against US financial markets to date was the Operation Digital Tornado campaign organized by a group called L0ngWave99. Between February and April 2012, this campaign launched over six DDoS attacks against US securities and commodities exchange [2]. The Al-Qassam Cyber Fighters, known as QCF, was an attack campaign supported by anti-Western rhetoric group Hamas that claimed responsibility for Operation Ababil, a series of DDoS attacks against US financial institutions between 2012 and 2013 [2]. The full list and detailed explanation of attacks in this category can be found in [2].

In the electricity market domain, since the passage of the Energy Policy Act of 2005, fraud and market manipulation have been the top enforcement priority of the Federal Energy Regulatory Commission (FERC). For fiscal year 2018, FERC reported 16 potential market manipulation cases, 14 of which were closed with no action [45]. The reason for most of these no action closures was that no evidence was discovered on the detail and mechanism of the attacks which greatly undermined the credibility of allegations. From this we see that market manipulation attack in electricity markets is an emerging field which needs significant research and investigation.

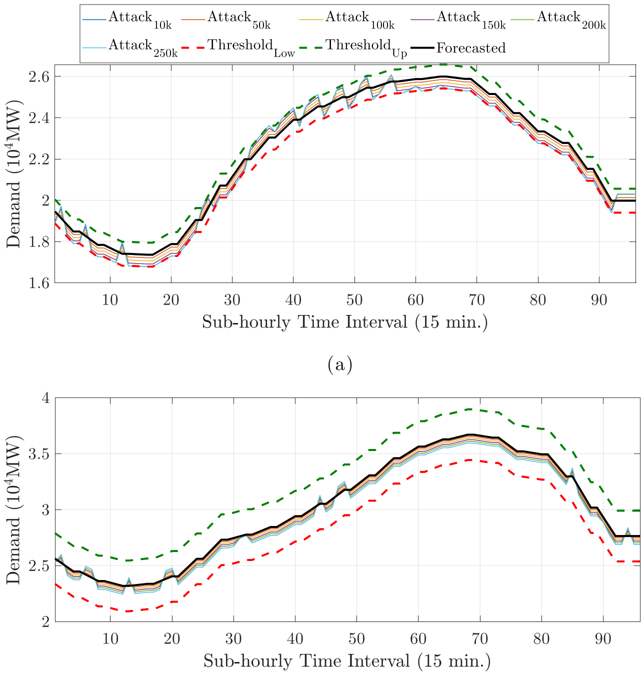

(a)

(b)

Figure A3: Load profile of the power grid at each time interval associated with attacks on the generation side companies with different botnet sizes: a) New York ISO; b) California ISO.

Appendix V.2 Indirect Attacks on Power Systems

In recent years, many researchers have studied the effect of indirect cyberattacks on different sectors of the power grid. In these attacks, the adversaries try to indirectly affect the normal operation of the system to sabotage stand-alone components or cause blackout in the entire grid [4, 15, 17, 39, 51, 61]. This class of attacks was first introduced in [39] where the system total demand was altered by the intruders to cause overflow in the power transmission lines and other system components, pushing the grid towards instability. The attack stems from compromising the load control signals associated with big industrial loads and data centers. By securing the communication channels between the control center and controllable loads, the risk of this attack is greatly reduced. The possibility of load altering to attack big data centers with the aim of causing power outages was studied in [61]. The authors showed that exploiting the attack vectors in cloud environments (platform as a service (PaaS), infrastructure as a service (IaaS), and software as a service (SaaS)) can be effectively used for taking down big data centers. According to this paper, defense and prevention mechanisms for such attacks are either impractical or extremely expensive.

The authors in [4] developed a software-based protection scheme to detect and protect against the load altering attacks introduced in [39]. This protection system is purely software and does not require any changes in the traditional communication channels/protocols.

In [15, 17, 51], the authors studied the possibility of exploiting compromised IoT devices to alter the total demand of the power grid and cause instability in the system. More specifically, the method developed in [17] is an optimization-based approach which requires a complete knowledge about the power grid (topology of the grid, detailed parameters of the transmission lines/generators, and real-time regional generation/demand). However, implementation of this attack is very challenging in practice as the required information may not be readily available to attackers. To overcome this challenge, Dabrowski et al. proposed a new method to increase the total system demand through remotely activating CPUs, GPUs, hard disks, screen brightness, and printers to cause frequency instability in the European power grid [15]. Although the new approach did not require as much detailed information about the system components, the number of compromised IoT devices needed for a successful attack is quite high because the devices do not consume a lot of power. Soltan et al. proposed the use of high wattage IoT devices to launch various types of attacks (frequency instability, power line cascade tripping, and black start restoration interruption) on a power grid to cause blackouts in the entire system [51]. More recently, Haung et al. conducted in-depth analysis on the impact of high wattage IoT attacks on the power grid and illustrated that launching random attacks on the grid may not lead to large scale blackouts if the embedded protection schemes in the system work properly [28].

Despite the numerous improvements of these works [28, 51] over the previous ones, they all still suffer from the following weaknesses: i) they require a large number of compromised IoT devices to launch a successful attack, ii) the proposed attacks are not stealthy, and iii) there is no direct economic profit for the attacker to launch these attacks. Motivated by these points, this paper presents a new attack mechanism based on a botnet of high wattage IoT devices to attack deregulated electricity markets while requiring only a minimal number of bots, maintaining stealth, and offering financial gain.