Profiting from Attacks on Real-Time Price Communications in Smart Grids

Paul Wood and Saurabh Bagchi
School of Electrical and Computer Engineering
Purdue University
West Lafayette, Indiana, USA
Email: [pwood,sbagchi]@purdue.edu

Alefiya Hussain
ISI/USC
Marina del Rey, California, USA
Email: hussain@isi.edu

***Abstract*—The smart grid (SG) promises to revolutionize power grid efficiency and reliability by bringing wide-area control and coordination between both power producers and widely distributed consumers. Such improvements, however, depend on reliable communication infrastructures for cooperation, thus creating an interdependence between wide area networks and the power grid. Real-time pricing (RTP) systems coordinate producers and consumers via price signals, and recent research has shown that network disruptions in RTPs can significantly harm or disrupt power grid operation. In this paper, we theorize and demonstrate how strategic network disruptions can further disrupt grid operations in ways that are profitable to a strategic adversary.**

We quantify the economic impacts of a strategic adversary that utilizes denial of service (DoS) attacks to gain a financial advantage in the power market, without compromising the integrity of the RTP signals. The adversary develops a strategy of when and where to launch DoS attacks by utilizing our algorithm that optimizes prices in her favor. A defender minimizes these financial gains by obfuscating the network targets, reducing the effectiveness of attacks. Our results provide insights to the dependability of RTP when deployed across disruptable wide-area best-effort communication networks.

I. Introduction

Power system operators must integrate distributed and renewable energy resources (RER) into traditional power grids to meet global demand for clean, reliable energy. Variability in RER supply [4] and the resulting congestion in power transmission networks complicates this integration, however, as the uncertainty in power availability can destabilize energy markets and create wide price swings in energy markets [7]. For example, New York's Independent System Operator (NY-ISO) routinely experiences 10x price increases over a span of 5 minutes (from \$30 to \$300 per MWh [11]). Consumers have typically been shielded from these swings, but as RER penetration has increased, utilities have begun passing costs on to consumers [6] in the form of rate increases and rooftop solar penalties. In the emerging smart power grid, new control methods such as demand response (DR) and real time pricing (RTP) [15] are under development to minimize power market volatility through increased coordination [3]. These new systems, however, rely on wide area networks whose reliability may not be guaranteed. Consequently, RTP systems may fail during communication network outages [1], [9], [10], [16], [17], potentially resulting in a sub-optimal economic equilibrium. This drives up the cost of energy to legitimate users or can cause imbalance in the energy supply and demand resulting in wastage or instability to the power grid. In this paper, we analyze the interdependence between these communication failures and the power market through the lens of a strategic adversary who attempts to manipulate the RTP system for profit.

RTP systems facilitate producer-consumer coordination via networked incentive signals. A centrally coordinated market mechanism adjusts the price of power to balance supply and demand—if there is too much supply, the price is lowered and consumers will utilize the surplus power. RER fluctuations in supply are thus compensated for by adjustments in market price. Much like in other networked control systems, an RTP controller samples the flows of electric power and modulates the market price over a communication network. Unlike networked control systems, however, RTP systems must coordinate among devices that are owned by different economic entities, requiring pricing or incentive based control instead of direct modulation of loads and supplies. Since the price and power signals traverse geographically diverse end points, it becomes feasible for a strategic adversary to disrupt the power grid by disrupting the communication channels on which the RTP system relies [17].

Prior work has focused on two main classes of attacks: maximum disruption in real-time [16] and profit-driven attacks in day-ahead pricing [1]. In this paper we bridge these two concepts, creating profit-driven attacks in real-time pricing systems. Work in [16] measured the impact of delay and integrity attacks on RTP signals. The authors showed that specific attacks create large instabilities in market price, potentially crashing the market. In this paper, we show that profit-driven attacks can be successful without destabilizing the grid control loops or damaging grid equipment. The attacks that we study here are more feasible to be launched and are more likely to stay under the radar, thus having the potential for greater impact. Work in [1] analyzed the impact of arbitrary delay and modification attacks on the incentive signals in the RTP system. They modeled a strategic adversary that could manipulate an RTP system if all global communications were compromised. They did not, however, consider attacks that evolve in real time–instead they focused on day-ahead market planning techniques in a game theoretic framework. In this work, we focus on real-time attacks where the adversary is aware only of local market price–a much simpler attack since the integrity and confidentiality of 3rd party communications is still maintained. We also introduce a defense mechanism that distinguishes us from the two previous closely related approaches which focus on the attack modeling.

Our solution combines profit-driven attacks with real-time pricing systems and strategies. In this paper we develop a strategic adversary who utilizes a battery to capitalize on time-varying power prices that are caused by controlled disruptions in the RTP communication network. Power is purchased when it is cheap and resold later when the price rises—except the adversary influences price via attacks. Each disruption removes controllable loads and supplies from the system such that the market price trends higher or lower than in the attack-free case—supply and demand becomes less elastic. The attacker then selects network targets and disruption times with a goal of maximizing profit opportunity. For example, if an attack is predicted to raise the market price by an additional \$20, an adversary with a large-sized battery can buy power to charge the battery, launch the attack, and then sell it back for an additional profit.

In our simple attacks, the adversary launches denial of service (DoS) attacks against other end users by using a parametric algorithm trained and triggered on the real-time price signal. In our complex attacks, direct RTP price signal manipulation (e.g. man-in-the-middle attacks) is allowed for comparison with the simpler DoS strategies. Using these strategies, we demonstrate how an adversary can financially benefit from network attacks in the smart grid.

We then put forward a novel defense mechanism that randomizes the adversary's view of the network targets for the load end points. Consequently, the adversary's strategy can no longer be faithfully executed. The difference between the adversary's view and the actual system can be controlled by a defense parameter, namely, how many end point network addresses to randomize. This defense mechanism can be deployed through moving target defense mechanisms implemented by the RTP system operator. The network address assigned to the end load points is permuted either periodically, based on market fluctuations, or based on indications of attacks to minimize the impact of network disruptions on the RTP signal.

In our experiments, we show that an adversary could potentially profit from an RTP system with a simple rechargeable battery and low-level denial-of-service attacks. The adversary is able to extract up to \$119 per day from the RTP market using DoS methods, 69% higher than without attacks. If the adversary is able to compromise 20% of the devices in the complex integrity-based attack, revenue could be increased by 98% or more. In the future, if fluctuations increase and battery prices decline, the profit amount is expected to increase further. We then show that a defender, utilizing our shuffling and deception strategies against these attacks, can reduce the adversary's profitability by 29% with only 3 moving targets.

In this paper, we make the following novel contributions:

1) We present a strategy for a strategic adversary to illicitly profit from a real-time pricing mechanism in the smart grid. The attack relies on delaying communication on a subset of the network links and for a subset of time, given by the adversary's algorithm.
2) We present a defense strategy customized for protecting against such market manipulation attacks. The defense strategy can be customized to fit within a certain defense budget and the benefits scale proportionally to the defense cost incurred.
3) We quantify the cost to launch an attack of the type presented here, the economic advantage that can accrue to the attacker, and the cost of defense, all based on real-world scenarios and data.

The rest of the paper is organized as follows. Section II covers the market mechanism background for real-time pricing. Section III details the attacker's strategy in the RTP system. Section IV details some defensive techniques to stop attacks presented in this paper. Section V details the experimental setup, including the particular RTP system in use. Section VI has the experimental results, and Section VII concludes the paper.

II. Background: Real-Time Pricing Systems

The goal of RTP systems is to constantly match supply with demand, and any mismatch is known as *residual power (RP)*. Without widespread energy storage devices, it must be immediately corrected since positive RP (power surplus) is shunted and wasted while negative RP (power shortage) causes frequency droop, brownouts, and possible equipment damage.

A general dynamic pricing objective function [5], used by RTP, is shown in Equation 1. The objective is to minimize RP by controlling the market price, λ:

$$\arg\min_{\lambda} \left| \left(\sum_{i \in N} P_i(\lambda, t) \right) \right| \tag{1}$$

where N covers all the consumers/generators in the power grid, and $P_i(\lambda, t)$ is the power used or produced (negative) by each client at price λ for time t. Large systems may have multiple λ's for different locations in the power grid, but this paper is focused on a smaller market region with a single price signal.

A. Real-Time Communication

Real time communication provides the ability to incorporate dynamic pricing information at the consumer. Distributed consumers have access to changing information that is beneficially incorporated into the pricing optimization problem, *i.e.*

$$P(\lambda, t) = P_{\text{forecasted}}(t) + P_{\text{flex}}(\lambda, t) + P_{\text{unpredictable}}(t) \tag{2}$$

where $P_{\text{forecasted}}(t)$ are scheduled, predictable loads and supplies, $P_{\text{flex}}(\lambda, t)$ are price-sensitive and adjustable (RTP participants), and $P_{\text{unpredictable}}(t)$ are unexpected or unpredictable. Balancing Eq. (1) via (2) inherently requires constant communication between the RTP market players (via λ) to adapt to

TABLE I
LIST OF SYMBOLS

J	Set of targets
λ	Market clearing price ($)
T_w	Attack decision window (s)
$S(t)$	Energy strategy $\in [P_{\text{min}}, P_{\text{max}}]$
$P_i(\lambda)$	Power removed from grid (kW)
$A_j(t)$	Attack target j at time t, $\in 0, 1$
P_{atk}	Estimated price impact of attack ($)
D_j	Flexible load coefficient or gain for target j (kW/$)
$C_j(t)$	Cost to attack target j at time t ($)

changes in $P_{\text{unpredictable}}(t)$ as the system evolves in time. As RTP methods solve for new prices, those prices are broadcast system-wide to the market players. The market players respond by adjusting consumption and production values (i.e. Equation (2)) that the RTP algorithm samples, in the physical domain, for its next price calculation. Additional details can be found in [17].

B. RTP Controllers

The core control function samples Equation (2) and then solves Equation (1). Any number of solutions, from feedback controllers to gradient descent methods, can be used to solve for the price signal λ. The information flow in such algorithms may become irregular with imperfect networks, however. In such cases, the market may not perform as expected [17]. We utilize a technique described in Section V-A to facilitate delay-tolerant solutions to Equation (1).

C. Impact of Network Outages

Network outages disrupt the communication of λ from the RTP controller to the end users. The $P_i(\lambda, t)$ for each disrupted i becomes fixed with $\lambda =$ constant, frozen in a zero-order hold. This causes the impact of future price changes, *i.e.* $\frac{\delta P}{\delta \lambda}$, to decrease since less devices are aware of price changes and thus cannot respond with power adjustments. As a result, λ must go higher or lower to correct for the same amount of RP than in the perfect communication case. To maintain connection incentives, market players are charged retroactively for their power consumption based on the actual market price. This ensures that consumers do not avoid high prices by claiming that they were unaware of λ (*e.g.* via self-disconnection).

III. STRATEGIC ADVERSARY

A. Strategy Summary

The adversary owns a energy storage device (*e.g.* a rechargeable battery) and profits by purchasing power at a low price and selling it back into the market at a higher price. To maximize profit, first the adversary monitors the real-time pricing signal and establishes charge and discharge price thresholds. Then she attempts to increase profit by estimating the price impact of a DoS attack by monitoring the gradient of the market price history. Peaks in attack impact are identified as they evolve in real-time, and attacks are launched whenever the peak estimated market price exceeds the charge or discharge thresholds. In this way, the opportunity for arbitrage is maximized, and the adversary increases profitability.

B. Capabilities and Resources

In this paper, we model the strategic adversary as a single end user in an RTP system. The adversary can view the price signal and launch attacks against other users in the system.

1) User Discovery: The attacker is assumed to know the IP addresses of clients participating in RTP market in one of three ways. First, since the market operates on a local level and the potential addresses of Internet-facing devices are geographically correlated for ease of routing, reviewing public IP address registries could reveal targets, especially if RTP participation is widespread. Second, many microgrid applications could utilize peer-to-peer services, especially for islanded operation. These applications could require peer advertisements or open ports that would reveal addresses and service locations. Third, many last-mile network connections utilize shared infrastructure such as cable modem services or passive optical networks. Promiscuous modems could reveal periodic access patterns that are unique to RTP devices, for example. Other alternatives include hacking the RTP controller or other man-in-the-middle security breaches.

2) DDoS Attack Capability: The adversary has access to a DDoS-as-a-Service providers or ”booter/stressers”. Such services offer chunks of attack time for a nominal monthly fee [13]. Armed with a target IP address, the adversary can simply pass it via a web interface and start an attack. It is assumed that consumer grade connections are of sufficiently low capacity such that multiple users can easily be taken offline simultaneously. Additionally, since the attacks are deep in distributed last-mile networks, filtering costs may be prohibitively high. Other smart grid vulnerabilities listed by NESCOR [10] include easy to jam wireless communication channels and physical or logical access to communication channels for entities that do not require it.

3) Energy Storage: The adversary is armed with a rechargeable battery that can charge and discharge at a particular rate, has a limited useful lifetime, and a maximum capacity. The battery is assumed to be 100% efficient such that no energy is lost in the charge and discharge process.

C. Strategy Definition

The adversary’s arbitrage strategy is to charge the battery when energy is inexpensive and discharge when the price becomes higher:

$$\text{Maximum Revenue} = \arg\max_{S(t)} \sum_{t \in T} T_w \lambda(t) S(t) \tag{3}$$

where T is the set of market clearing windows of negotiation, $\lambda(t)$ is the market price at time t, and T_w is the attack strategy window width that breaks the strategy space into chunks to aid in solution calculation. $S(t)$ is the adversary’s energy strategy–charge or discharge at time t. The adversary’s goal is to maximize profit by manipulating $\lambda(t)$ via DoS attacks.

D. Price Manipulation

Section II-C described that whenever clients are disconnected from the marketplace, *e.g.* via DoS, the effective gain of the price signal decreases. For example, if $P_i(\lambda) = C\lambda$ for some constant C, and 10 clients are connected, $P(\lambda) = 10C\lambda$. If a DoS attack removes 5 clients, then $P(\lambda) = 5C\lambda$, for an attack impact gain of $5C$. To achieve the same ΔP, λ would need to change twice as much during the attack. This means that a high RP coupled with client disconnections leads to more dramatic price swings, and the adversary can leverage these swings to increase revenue. The price manipulation strategy can be broken into estimating RP and estimating the change in the gain due to a DoS attack to calculate the manipulative power of an attack.

1) Target Gain Estimation: In order to influence λ, the adversary needs to know the price gradients of each target's $P_i(\lambda)$ function, in other words, how elastic is the user's load with respect to the price she pays for energy. This function is private for each user and unknown even by the RTP service, so the adversary must estimate the gain $\frac{\delta P_i}{\delta \lambda}$ for each user. One method is to compromise the Internet-connected devices in users home, such as by default passwords or weak encryption, and directly reveal the functions to the adversary. Alternatively the adversary can estimate the gain as D_j for target j in the following way. First the adversary measures the gradient of price as a moving average over k timesteps, *e.g.* as $\frac{1}{10C}$. Then the adversary attacks target j and measures the new gradient over an additional k timesteps, *e.g.* as $\frac{1}{9C}$. The individual client gain is then calculated as $D_j = 10C - 9C = 1C$. The goal is to observe increases in market prices (or similarly, decreases) and if the attack causes a market participant to go offline, then the rate that the price changes will increase *i.e.* become convex temporarily $(|\frac{\delta\lambda}{\delta t}_{\text{pre-atk}}| < |\frac{\delta\lambda}{\delta t}_{\text{post-atk}}|)$. This type of approach is not perfect–it is quite possible that targets will be miss-classified due to external market conditions such that the price may be concave even without an attack. This classification error simply erodes the adversary's ability to accurately decide which end points to target (a parameter in experimentation).

2) Residual Power Estimation: In the RTP system, the residual power is only known by the RTP controller, as the output of Equation (2) and the input of Equation (1). The RTP controller decreases and increases λ as a function of RP. Therefore the gradient of λ is loosely proportional to the amount of RP in the system. If RP is negative (shortage), then $\frac{d\lambda}{dt}$ will be positive, and vice versa for a positive RP. If the gains for every client (D_j) are known, then the cumulative system gain ($D = \sum D_j$) can be used to estimate RP as $D\frac{\delta\lambda}{\delta t}$. The attack power is then estimated as

$$P_{\text{atk}}(t) = \frac{\delta\lambda(t)}{\delta t} \sum_{j \in J} D_j \tag{4}$$

where J is the set of valid targets, $\frac{\delta\lambda(t)}{\delta t}$ is the current price gradient smoothed over k timesteps, and D_j is the estimate for $\frac{\delta P_j(\lambda)}{\delta\lambda}$. It is assumed $P_{\text{atk}}(t) = 0 \; \forall \; S(t) \neq 0$ since the attack is already ongoing. If a net imbalance of power exists, and the RTP signal is actively correcting this by increasing prices, for example, then the attack will cause the price to overshoot by approximately P_{atk}.

Algorithm 1: Basic DoS Strategy

```
1 if $\lambda(t) + P_{atk} < \bar{\lambda}_{buy}$ then
2     $A_{j \in J}(t) \leftarrow 1$
3     $S(t) \leftarrow P_{\text{max}}$
4 else if $\lambda(t) + P_{atk} > \bar{\lambda}_{sell}$ then
5     $A_{j \in J}(t) \leftarrow 1$
6     $S(t) \leftarrow -P_{\text{max}}$
7 else
8     $A_{j \in J}(t) \leftarrow 0$
9     $S(t) \leftarrow 0$
```

E. Attack Strategy

Once a set of viable targets and their gains are known, then the adversary may use them to influence $\lambda(t)$ in an attempt to improve Equation (3). First, the baseline strategy is developed with $\bar{\lambda}_{\text{buy}}$ as the target price for charging periods and $\bar{\lambda}_{\text{sell}}$ as the target price for discharge. This price is established by a-priori observations of RTP price trends.

Algorithm 1 contains the strategy for attacking targets. Lines 1-3 represent the buying strategy and lines 4-6 the selling one. Line 1 states that whenever the current market price plus the power of attack (which can be negative) is less than the buying price threshold, then the attack should be launched and the battery should charge. Similarly, line 4 sells when the estimated price after attack is higher than the selling threshold. Lines 7-9 stop the attack if the price is not within the buy or sell thresholds. Additional constraints (not shown) keep the battery's energy within capacity. The net benefit from the attack is measured by comparing (3) with and without attack for the same scenario.

Algorithm 1 can be further enhanced by peak detection on Lines 1 and 4 rather than operating on the first point that meets the attack standards. The authors of [8] map real-time peak detection to best choice and optimal stopping problem, and we further enhance the algorithm by selecting $\bar{\lambda}$ based upon outlier detection on the $\lambda(t) + P_{\text{atk}}$ signal, as described in [8] Section 3.3.

F. Integrity Attacks

If λ values sent to some subset of clients can be manipulated, then a new attack strategy can be implemented to further defraud the market, where $P_{\text{atk+}}$ is for selling and $P_{\text{atk-}}$ for buying:

$$P_{\text{atk+}}(t) = \sum_{j \in J} \frac{\delta\lambda(t)}{\delta t}(P_j(\lambda) - \overline{P_j}) \tag{5}$$

$$P_{\text{atk-}}(t) = \sum_{j \in J} \frac{\delta\lambda(t)}{\delta t}(P_j(\lambda) - \underline{P_j}) \tag{6}$$

where $\overline{P_j}(\lambda)$,$\underline{P_j}(\lambda)$ represent the maximum and minimum power output for each target j, and $\frac{\delta\lambda(t)}{\delta t}$ is the current price gradient. Algorithm 1 is supplemented by these strategies where Line 1 gets $P_{\text{atk-}}(t)$ and Line 4 gets $P_{\text{atk+}}(t)$.

Using this strategy, whenever the market price is decreasing due to positive residual power, even more positive residual power is added by further reducing the load (atk-), causing the market price to plummet further. Similarly, whenever the market price is increasing due to negative residual power, even more load is placed on the system thus increasing price. The net effect is that compromised devices make poor market decisions that benefit the adversary.

1) Cost of Attack: The cost of attack can be incorporated into a modified version of (3):

$$\arg\max_{S(t)} \sum_{t \in T} \left(T_w \lambda(t) S(t) - \sum_{jinJ} T_w A_j(t) C_j(t) \right) \tag{7}$$

where $A_j(t)$ is the binary attack indicator and $C_j(t)$ is the cost per second of attacking target j. The adversary is still attempting to maximize profits in the left term, but each attack that influences $\lambda(t)$ also has a cost $= A_j(t)C_j(t)$ in the right term, which can be constrained by a budget (cost $\leq$ budget).

The costs can be optimized by sorting targets by their cost-impact factors $C_j D_j$ and prioritizing target above a threshold D_{thresh}:

$$A_j = 0 \; \forall \; C_j D_j < D_{\text{thresh}}, j \in J \tag{8}$$

where D_{thresh} eliminates cost-ineffective targets. This restricts the attack strategy by reducing P_{atk} to constrain costs.

IV. Defender Strategies

This section covers defensive strategies that can minimize the attacker's profit.

A. Moving Target Defense

The adversary can be countered in two ways. First, the ability to manipulate RP directly can be removed through device security and is left to other research. Second, the RTP service could remove the ability to disrupt communication links, however the RTP operator would need to harden hundreds or thousands of links to distributed end homes instead of just the network's edges.

Defensive maneuvers can still be made, however. Intuitively, some targets are "safe" from attack because (8) marks them inefficient. The targets A_j represent IP addresses that will be DDoS'd by the adversary. The defender can mitigate attacks by shuffling the targets ($A_x \rightarrow A_y$) so that the attacker's efforts to attack the IP of target x actually disable target y. High value targets can then be swapped for low value targets so that the attacker's profits are minimized. This can be done by synchronizing dynamic IP assignment operations with regional Internet service providers (ISPs)–the RTP operator requests the ISP swap the addresses of x and y. The impact is not negligible, however. A forced IP reassignment will cause temporary client interruption, and the creation and support of infrastructure to perform such reassignments would require at least some engineering support. The attacked client would also effectively pay penalty rates for power, so the RTP operator would need to properly incentivize participation.

The defender can strategize about which targets' IP's to swap. Optimally, the highest-value targets would be swapped for the lowest-value targets, and this is what Algorithm 2 performs. The list of targets is sorted by their estimated impact D_j and the lowest value k targets are swapped with the highest value targets. This minimizes the change in λ due to A_j

Algorithm 2: Defender Moving Target Strategy

```
Sort $J$ by $D_j$ descending
$k \leftarrow 0$
while $k < |J|/2$ AND $D_{J(k)} >$ Threshold do
    Swap IP of $J(k)$ with $J(|J| - k)$
end
```

B. Detection via Deception

The strategies in Section III-C are all driven from the end-user's observable incentive signal. A false incentive signal could be sent to suspects in the system in order to trigger false attacks on the system. Correlation can be drawn between false signals and corresponding DoS attacks to identify the adversary. For example, a deceptive, high λ_x value could trigger Line 4 in Algorithm 1, and the RTP operator could send this false signal to a potential adversary and observe A_j via heartbeat signals. This strategy has a cost in that if the signal is sent to a non-adversary, the market efficiency decreases because load will be added or removed contradictory to the current market price λ:

$$\text{Cost of Deception} = |T_w \lambda(P_i(\lambda, t) - P_i(\lambda_x, t))| \tag{9}$$

where λ_x is the false price and T_w is the duration of the false price signal. If deception occurs for one T_w then the cost is effectively the change in revenue that the client i was providing to the market. For example, a client uses 1 kWh of energy during T_w. When the adversary check is run, λ_x is set to 2λ and the client consumes 1/4 kWh. The cost is then $3/4 \cdot \lambda$.

V. Experimental Setup

This section covers the real time pricing mechanism used for experimentation and the load/supply models for evaluation.

A. Overview of RTP Controller

The price setting algorithm inside the RTP controller minimizes an implicitly non-stationary objective function, the residual power (from Equation (1)). Traditional market solutions in [14], [18] utilize gradient descent and interior point methods to solve Equation (1), but these techniques require convex objective functions and gradients. A stationary λ with a varying t will change with the supply and demand of electric power, and during transient events such as faults or surge in demand, the change between $P(\lambda, t - \epsilon)$ and $P(\lambda, t + \epsilon)$ can be very large, thus violating the convex requirements. To overcome this limitation, we utilize a modified Nelder-Mead (NM) [2] algorithm that does not make assumptions about the of the objective function and can adapt to large $\Delta P()$.

Three modifications of NM are completed to enable the algorithm to perform online optimizations on non-stationary functions. First, the search space used by NM is modified to

TABLE II

PARAMETER DISTRIBUTIONS USED IN EXPERIMENTS, RESPECTIVELY FOR THE CONSUMER (C) AND THE GENERATOR (G)

	P_{min} (kW)	P_{max} (kW)	λ_{min} (\$)	λ_{max} (\$)
C	$\lvert\mathcal{N}(0, 0.5^2)\rvert$	$\lvert\mathcal{N}(3, 1^2)\rvert$	0	$\lvert\mathcal{N}(250, 75^2)\rvert$
G	-$\lvert\mathcal{N}(150, 50^2)\rvert$	0	$\lvert\mathcal{N}(30, 5^2)\rvert$	$\lvert\mathcal{N}(80, 5^2)\rvert$

prevent simplex collapse so that transients can be detected. This is done by adding noise to the points of the simplex so that it maintains a minimum size. Second, the cached function values $P(\lambda)$ are updated periodically to reflect the current value, and this is used for relative point ranking. This enables the algorithm to adapt to large $\Delta P()$. Finally, the algorithm is modified to perform re-evaluation of the simplex space during its shrink operation. Source code for this algorithm is available publicly[1], including all code used to generate data for this paper.

B. Load and Generation Model

We model supply and demand ($P_j(\lambda)$) as scaled sigmoid functions ($\frac{1}{1+e^{\lambda}}$) where maximum power (P_{max}) is achieved at maximum price (λ_{max}) and *vice versa*. For the experiments conducted in this paper, unless specified otherwise, the parameters used are listed in Table II. There are 2 generators G, $N_G = 2$, and 100 consumers C, $N_C = 100$. The adversary's battery has a capacity of 1200 kWh and a charge/discharge rate of 600 kW/h, and T_w is 5 minutes. The capacity is selected to supplement one generator in the system for 8 hours (*e.g.* a solar farm during the night). The residual power measurement (the objective function of Eq. (1)) is assumed reliable and sampled out of band, *e.g.* via a dedicated state estimation system.

C. Real-World Dataset

To analyze the effectiveness of the attack strategy at handling unexpected fluctuations in power, the forecasting error from several days of New York Independent System Operator (NYISO) data is used to generate an error function. The difference in the day-ahead forecast and actual load model for June 19-26, 2015 are used to create this signal [11]. The uncontrollable load signal is scaled so that the highest and lowest values are no more than 50% of the maximum and minimum amount that can be absorbed by the generators and consumers in the system. In training, where necessary, the first 7 days are used while the last day is used as the test in all of the experimental results.

VI. EXPERIMENTAL RESULTS

This section covers the results of our evaluation. We first identify the baseline arbitrage opportunities in the market. Then we consider the economic advantage that an adversary can achieve by blocking the communication to a subset of the consumers, first with a naïve attack and then with an attack that tracks the real time price fluctuations. We then evaluate to what degree the attack can be reduced through our defense mechanism. We conclude by considering the financial gains or investments in dollar terms for the attack.

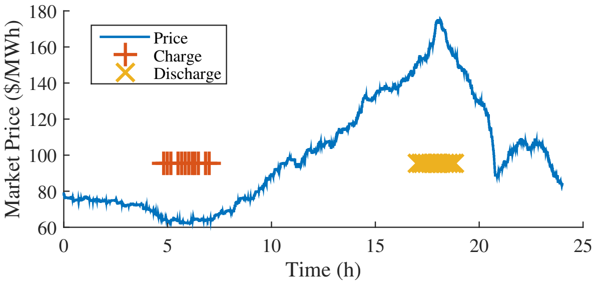

Fig. 1. The market price during optimal baseline operation is shown. The charge and discharge markers indicate the adversary's optimal charging strategy.

A. Experiment I: Baseline Profits

In this experiment, the baseline profits are established for any consumer in the attack-free case. As the experiment day progresses, unpredictable changes in consumption cause the price of power to increase and decrease, as shown by the price in Fig. 1. The adversary's rechargeable device participates in this market, attempting to minimize the average buy-price and maximize the sell-price to turn a profit.

The charge and discharge duration is limited to two hours by the capacity and charge rate of the battery described in Section V-B, and in this scenario, the adversary is able to profit \$116.48 for the 24 hour period from buying low and selling high. For this strategy, $\bar{\lambda}_{\text{buy}} = \$64.48, \bar{\lambda}_{\text{sell}} = \$146.62$ were chosen via repeated search optimization on the test day. This level of precision is not attainable in practice because the market is assumed unpredictable, so the revenue here is a maximum value using the attack-free strategy in Algorithm 1. This value is not the true maximum possible, since additional arbitrage opportunity exists between hours 20 and 24, but it shows the maximum effectiveness of the heuristic. A more realistic value of $\bar{\lambda}_{\text{buy}} = \$65.18, \bar{\lambda}_{\text{sell}} = \$111.20$, selected by the top 15% and bottom 15% quantile of prices over the training period, yields a reduced profit of \$70.88. Note that the attack-free strategy is not a harmful event for the power grid–this stabilizes market price and grid loading which is beneficial to consumers and grid operators.

B. Experiment II: Impact of DoS Attacks

In this experiment, the adversary is given the ability to disrupt communication with $|J| = 20$ users connected to the market. These disruptions increase market volatility by forcing the attacked users to enter a holding pattern in energy consumption. To maximize profits, the adversary implements Algorithm 1. First the adversary selects the parameters $\bar{\lambda}_{\text{buy}}, \bar{\lambda}_{\text{sell}}$, and she does this by observing $\lambda(t) + P_{\text{atk}}$ during the training phase. It is assumed that D_j is known by the adversary, and $\bar{\lambda}_{\text{buy}} = \$65.00, \bar{\lambda}_{\text{sell}} = \$111.35$ are selected by quantiles of 15% and 85% respectively.

The values for D_j, P_{atk} during the test phase are shown in Fig. 2. The value of D_j, shown as the sum over all J targets, peaks when the market price is relatively low and is suppressed during peak prices. This is because the two largest market players, the generators, are producing maximal output

[1] https://github.com/pcwood21/RTP_DoS_Simulation

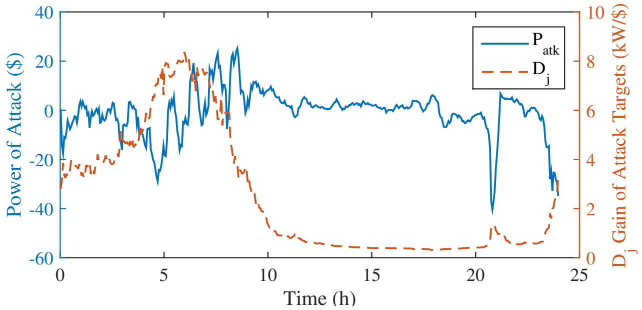

Fig. 2. The power of attack and the gain for 20 targets is shown for the day. The D_j term becomes saturated at high market prices due to output saturation at the largest market players.

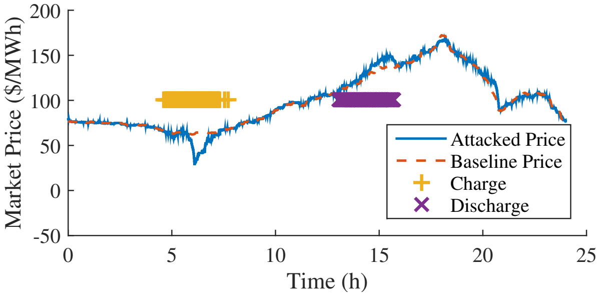

Fig. 3. The market price is shown when the attacker implements a DoS attack strategy on 20 targets.

after about $80. After the $100 price range, there is only a gradual change in consumption with price increases, mainly by small consumers. In this experiment, the cost of each target is assumed equal, and therefore the power of attack is dominated by the generators.

The adversary implements Algorithm 1 and launches attacks during market operations. Fig. 3 shows how the market responds during the attacks. At around 5 hours in, the adversary begins to launch her attacks, and the market price begins dropping in response to these attacks. Once the battery has charged, the attacks end and the market begins to behave normally. After the price rises, the adversary again attacks to increase the market price further. The increase at this price level is smaller due to the low D_j values in this price range. The attack yields $119.77 of profit for the day, an increase of 69% over the baseline charging profile.

C. Experiment III: Impact of Integrity Attacks

In this experiment, the adversary is given the capability to manipulate individual price signals. The underlying methodology behind the attack is identical to the denial of service case, but the loads are provided a manipulated price signal instead of a stale one, and the P_{atk} is calculated appropriately.

Fig. 4 shows impact of compromising devices in the RTP system. The adversary is able to impact price and extract additional profits totalling $140.36, 98% higher than the baseline. For this attack, the adversary compromises the λ signal as it is sent to the consumer device. The P_{atk}, $\overline{P_j}$ is achieved by sending $\lambda = 1000$, when the adversary wants to buy, and $\underline{P_j}$ with $\lambda = -1000$, when the adversary wants to sell

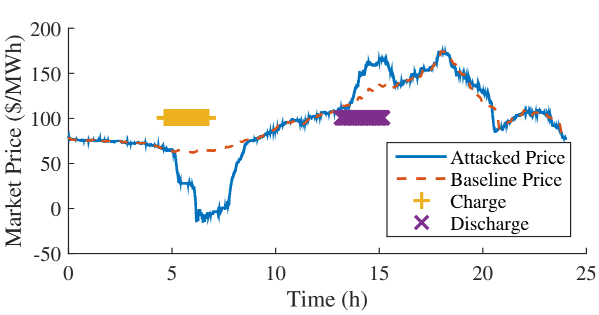

Fig. 4. The market price is shown when the attacker implements a integrity attack strategy on 20 targets. During the charging phase, consumers are mislead into conserving power, and during discharge, consumers are misled into over-purchasing power, and this results in price increases and decreases for the adversary to leverage.

to the grid. This causes the loads imbalance power directly, greatly increasing the effectiveness of the adversary. A cost comparison with DoS attacks is difficult to achieve, however, since the cost of compromising encryption or passwords on consumer devices is not easily quantified.

The effectiveness of these attacks suggests that compromised devices could significantly impede RTP system deployment. Effective defenses, however, are known and need to be deployed more widely, such as, the use of strong authentication scheme and enforcing non-default, and strong passwords. This type of attack can also be very damaging to grid equipment since coordinated loads can cause large transients in voltage and current to occur in the grid, along with instability of RTP systems as has been shown convincingly in [16]. However, that goal is not the focus of this paper.

D. Experiment IV: Defensive Strategies

Experiment IV analyzes the defensive strategies presented in Section IV. The defender's goal is to reduce or eliminate the adversary's profits. The first defense that the defender implements is protecting the information about the individual consumer loads, D_j. Investing in stronger end-device encryption and protections, for example, can protect this information. Fig. 5 shows how the effectiveness of the attack in Experiment II decreases as the accuracy of the adversary's D_j terms also decreases. Random noise is added to the D_j values used in the attacker's strategy to reflect inaccurate collection techniques (Section III-D1): $D_j^* = \mathcal{N}(D_j, \sigma^2)$. The lack of good target information significantly reduces the effectiveness of the adversary, making attacks less profitable. Initially, the adversary's profit drops sharply and then the law of diminishing returns kicks in and the curve flattens out. In this part of the curve, the adversary's estimates are already quite inaccurate and additional noise does not make a significant difference.

Another defensive technique is to swap targets D_x with D_y, as described in Section IV-A where the targets are rearranged using Algorithm 2. Fig. 6 shows the profit of the adversary versus the number of swaps that the defender is allowed. As the number of swaps increases, the profit decreases but at a lesser rate–the adversary routinely targets the most valuable assets, and since these are first swapped with the least valuable,

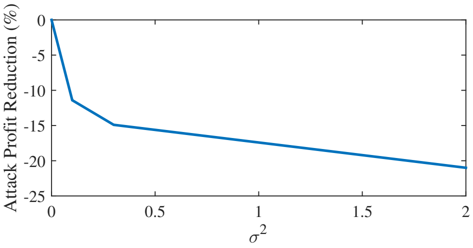

Fig. 5. The reduction of the adversary's attack-induced profit is shown as her information about the targets decreases. Errors in target value effectively reduce the profit of the adversary.

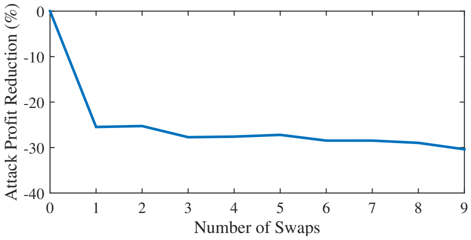

Fig. 6. The reduction of the adversary's attack-induced profit is shown as targets are swapped by the defender. Since the swaps are optimized on target impact, there is a diminishing return on investment for swapping all of the assets. The first swap protects a large generator with the most impact.

the impact of the swaps drops off rapidly. When 8 targets are swapped, the adversary's profit is reduced by 30%.

E. Return on Investment

The energy storage device used by the adversary, as described in Section III-B3 has a particular cost to install and maintain, and the attacks have a particular cost, based on ”DDoS/Booter” service pricing [13]. These absolute values can factor in to Equation 7 to determine the economic viability of the attack strategies. For revenue, if the adversary repeatedly executes the strategy in Experiment II, amounts could be as much as $43,000 per year. For cost, storage device prices are expected to fall to $200 per kWh by 2020 [12] and continue to fall with increases in production, so we estimate a yearly battery cost of $24,000 amortized over a 10 year lifetime. The net profit is then $\$43,000 - \$24,000 = \$19,000$ per year. This cost analysis does not include residual value or added benefits of a distributed battery system such as improved reliability during grid failures. An important note is that if the strategy is profitable, then more devices can yield more profit, or groups of attackers could form battery-consortiums for example. Booter service costs can vary, but average residential Internet connections are low-bandwidth and easily disabled with attacks relative to high-visibility targets like news websites, which keeps these costs low. Based upon the results in [13], a few hundred dollars can maintain a botnet for launching these attacks. For example, 212-booter launched 1993 attacks over 57 days with profits of $509, suggesting 1,000 attacks could be purchased for about $250 per month. The end result is that the adversary in this experiment could come out $15,000 ahead each year with a 1.2 MWh battery.

VII. Conclusion

In this paper, we presented how a strategic adversary could profit from a real-time pricing system in the smart grid by launching denial of service attacks on consumers connected to the pricing system, *i.e.* , delaying the price signal being sent by the system operator to the consumers. We showed that an adversary could increase revenues by 69% by disrupting up to 20 clients or by as much as 98% if the integrity of the pricing signal is compromised. We then showed how a RTP system operator could mitigate network attacks by strategically reconfiguring the device network. In this way, a defender is able to reduce the adversary's profits from DoS attacks by 30% with 8 IP address swaps. This work exposes some risks to real-time pricing systems in smart grids and provides a novel technique for defending against these attacks.

In future work, we will assess more stateful attack models where the adversary plans attacks over longer periods of time. When combined with the real-time pricing market, these attacks should further increase the adversary's profitability at the expense of other RTP users.

References

[1] C. Barreto, A. A. Cárdenas, N. Quijano, and E. Mojica-Nava. CPS: Market analysis of attacks against demand response in the smart grid. In *ACSAC '14*, 2014.

[2] R. R. Barton and J. S. Ivey Jr. Nelder-Mead simplex modifications for simulation optimization. *Management Science*, 1996.

[3] L. Bird, M. Milligan, and D. Lew. Integrating variable renewable energy: Challenges and solutions. *National Renewable Energy Laboratory*, 2013.

[4] G. Boyle. *Renewable electricity and the grid: the challenge of variability*. Earthscan, 2012.

[5] L. Chen, N. Li, S. Low, and J. Doyle. Two market models for demand response in power networks. In *SmartGridComm*, Oct 2010.

[6] P. Fairley. Innovation amid a raucous rooftop solar squabble. *Spectrum, IEEE*, 52(7):14–15, July 2015.

[7] J. C. Ketterer. The impact of wind power generation on the electricity price in germany. *Energy Economics*, 44:270–280, 2014.

[8] M. Krotofil, A. A. Cárdenas, B. Manning, and J. Larsen. CPS: driving cyber-physical systems to unsafe operating conditions by timing DoS attacks on sensor signals. In *Proceedings of the 30th Annual Computer Security Applications Conference*, 2014.

[9] Y. Liu, P. Ning, and M. K. Reiter. False data injection attacks against state estimation in electric power grids. *ACM Transactions on Information and System Security (TISSEC)*, 14(1):13, 2011.

[10] NESCOR. Electric sector failure scenarios and impact analyses. Technical report, Electric Power Research Institute, Incorporated, 2013.

[11] NYISO. NYC LBMP http://www.nyiso.com/public/markets_operations/market_data/custom_report/index.jsp?report=rt_lbmp_zonal, 2015.

[12] J. N. Russell Hensley and M. Rogers. Battery technology charges ahead http://www.mckinsey.com/insights/energy_resources_materials/battery_technology_charges_ahead, 2012.

[13] J. J. Santanna, R. Durban, A. Sperotto, and A. Pras. Inside booters: An analysis on operational databases. In *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, pages 432–440, May 2015.

[14] D. Shiltz, M. Cvetkovic, and A. M. Annaswamy. An integrated dynamic market mechanism for real-time markets and frequency regulation http://hdl.handle.net/1721.1/96683, 2015.

[15] P. Siano. Demand response and smart grids—a survey. *Renewable and Sustainable Energy Reviews*, 2014.

[16] R. Tan, V. Badrinath Krishna, D. K. Yau, and Z. Kalbarczyk. Impact of integrity attacks on real-time pricing in smart grids. In *CCS*, 2013.

[17] P. Wood, D. Shiltz, T. R. Nudell, A. Hussain, and A. M. Annaswamy. A framework for evaluating the resilience of dynamic real-time market mechanisms. *IEEE Transactions on Smart Grid*, PP(99):1–1, 2016.

[18] R. Zhou, Z. Li, and C. Wu. An online procurement auction for power demand response in storage-assisted smart grids. *INFOCOM*, 2015.