Defending Against Strategic Adversaries in Dynamic Pricing Markets for 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—**Dynamic pricing markets in smart grids contain competitive environments in which strategic adversaries can launch cyber-attacks to extract profits from the system. When multiple actors are competing, the disruption of key assets can create large swings in the profitability of each actor by changing the supply and demand dynamics of the underlying system. These swings in profits can be leveraged by an attacker to extract profits from the system. In this paper, we explore the implications of attacks and defenses on market resilience when faced with a strategic profit-seeking adversary. Changes in the market may mitigate or exacerbate the likelihood of cyber-attacks on system assets, and we present a strategy for optimizing defenses to minimize attacks. We also explore the impact of information sharing on system behavior and the potential for collaboration by the system owners to improve resilience. These design principles are exercised on a model of an interconnected electric power grid to show the potential security improvements provided by architectural changes.**

I. Introduction

Dynamic pricing markets in the smart grid (SG) [1], [2] enable the optimization of physical resource allocation. NYISO, the power system operator for the state of New York, could realize as much as $400 million [3] in annual efficiency gains by leveraging wide-area real-time dynamic pricing systems. These gains, however, depend on consistent, reliable communication networks to facilitate control signal and measurement exchanges. The control signals comprise the price signals that the utility sends to the consumers and the measurement signals comprise the readings of the electricity usage at the consumers. Consumer-grade networks are often unreliable or congested at times, and they are highly susceptible to denial-of-service (DoS) attacks that disrupt communications entirely. As researchers pursue SG and other incentive-driven network control systems, they need tools to first understand what will be impact of outages of the network infrastructure on the demand-driven pricing mechanism on the SG and then how to mitigate the impact of this increasing attack surface and improve system resilience.

Electric power markets suffer from volatility because electricity is not easily stored. This volatility is expected to become much more acute with the increasing use of renewable energy sources, such as, solar and wind, that depend on the weather patterns. This creates a constant need to match supply with demand, but presently demand is inelastic and unaware of real-time market conditions. Supply and demand has historically been predictable which has limited this mismatch of supply and demand to a tolerable level. New renewable energy resources (RER) such as solar and wind driven supplies, however, have reduced this predictability to the point that technologies like roof-top residential solar are becoming cost-prohibitive [4] to integrate into the grid. The future smart grid is designed to bring elasticity to demand via techniques such as demand response (DR)[2] and transactive control (TC) [5] so that RER's can be better integrated and system efficiency improved. These techniques, however, rely on extensive communication infrastructures to coordinate wide-area energy consumption.

Transactive control enables distributed, independent control systems, operated by independent *actors* or market players, to coordinate via incentive-driven signals (prices). For example, the set point on an air conditioner may be sensitive to the cost of electricity in an automated way. This incentive can be set a priori via time-of-use pricing, but it is not sensitive to unpredictable changes in market conditions. Alternatively, a central market coordinator can negotiate with these automated loads by exchanging price and load information in real time with all of the actors. This negotiation process enables actors to rapidly respond to fluctuations in grid supply and by adjusting their energy usage based on price signals and their current exogenous needs. It is possible to disrupt the price signal negotiation, however, via attacks on the wide-area communication network. Since network attacks can influence the market price of energy directly, via control signal disruption, a *strategic adversary* (SA) can potentially launch attacks to manipulate prices in her favor.

When communications between market players are disrupted, the transactive control system becomes unable to influence consumption or production at those market players. For example if there is a spike in demand, the price signal should rise to curtail consumption and promote production. Producers who are aware of this signal increase their output and collect additional profits. An attack could disrupt the market signal at the producer, however, and as a result, power output would remain stagnant. Consequently, the market price may rise higher than it otherwise would have to sustain equivalent demand curtailment, and this may benefit the other producers in the market. It could also lead to blackouts if the disruption is severe enough. If the SA is a producer, then direct financial benefit in the power market can be gained from such an attack. The SA can also benefit from the profits of multiple actors through various means such as investing in these actors. The ability for a network attack to benefit the SA is called the *attacker's incentive*, and this paper focus on measuring and reducing that incentive via defensive maneuvers.

Prior work in [6], [7] has shown that network attacks in smart grid control systems can disrupt price signals and provide benefits to subsets of consumers. These techniques, however, do not consider defensive maneuvers that the market players can use to protect themselves. Additionally, they rely on a strong adversary that compromises the entire market communication infrastructure. In this paper, we only assume the attacker can disrupt network links. Additional work in [8], [9], [10], [11], [12] has created a game-theoretic structure around attack and defense in control systems. These works do not consider the financial incentives of the attacker, however. Instead they focus on overall system performance or lower level dynamics and model the attacker as benefiting from system disruption rather than profiteering. In our work, we combine game theoretic strategies for smart grids into an attacker/defender game, with multiple defenders, that relies on financial incentives to motivate attack and defense.

Our solution encompasses a method for estimating the attacker's incentive through attack strategies, mapping them to a game, and playing the game from a defender's perspective to minimize the attacker's incentive. First, we create a model for translating attacks and impacts on a smart grid to a strategy space for the attacker. A dynamic market is implemented with communication links that can be disrupted via denial of service attacks to capture the attacker/defender strategies. From this space, we optimize the attacks to maximize the attacker's incentive (profits) by attacking communication links that distort the market to benefit the adversary. This is done via mixed integer linear programming (MILP). We then model a defender that attempts to minimize the attacker's incentive by blocking certain attack strategies via defensive investments (i.e. DDoS protection). Since we model multiple defenders (actors), we also explore the impact of information sharing among the defenders on the reduction in the attacker's incentive.

We test our solution with a smart-grid based transactive control system. The baseline system optimizes power consumption by controlling the market price signal. A simulated communication network facilitates the exchange of price and load information. The attacker can choose which communication links to disrupt with a DoS attack, and the defender can choose some links to protect. We show that the baseline attacker incentives can be as high as 51% of overall operating profits. When the defender and adversary's budget are equal, the attacker's incentive is reduced by up to 70%. These results validate the utility of this paper's technique in optimizing defensive investments. It points the way forward for practitioners (such as, utilities) looking to deploy demand-driven pricing for electricity by showing how much resilience in the networking infrastructure is needed to assure a certain level of economic profit from the system.

The rest of the paper is organized as follows. Section II covers the background in dynamic pricing markets and how they can be manipulated. Section III outlines the basic attack/defense strategy, and Section IV expands the strategy to include information sharing among market players. Section V evaluates the strategies against an example dynamic pricing market, and the related work is discussed in Section VI. The paper is concluded in Section VII.

II. Preliminaries

A. Electric Power Grids

Power grids are complex, interconnected systems composed of generators (sources) and loads (sinks). Each generator and load is connected to a series of transmission links (edges). A simple approximation of the energy system is a DC-load flow model which can be represented as a flow graph [13] where each asset (load, generator, edge) in the physical system is an edge or node in the graph. Profit-seeking actors sell energy above cost (generators) or transform that energy into something more useful (loads). The system is most efficient when supply and demand are equalized since a surplus of power is dissipated as waste heat and a shortage of power causes brownouts, blackouts, and other grid stability issues. The imbalance of supply and demand is known as residual power (RP), and power grid operators strive to minimize this value. Dynamic market mechanisms [14] and demand response (DR) [2] minimize RP by either direct load control (DLC) or dynamic real-time markets. The work in this paper focuses on power markets rather than DLC since the markets have a direct impact to profitability and thus attacker's incentive.

B. Power Markets

Power markets utilize a variety of economic strategies to minimize RP (1) and maximize the system's social welfare (SW). The SW defines the global system benefit from energy transactions as shown in (2), where ω_i is the value (consumers) or cost (producer) of power at each actor or market player i, and P_i is the amount of power consumed or produced by that market player. The parameter C penalizes the system for residual power with $C \gg \overline{\omega_i}$. While somewhat simple on the surface, the problem of maximizing SW is complicated by time-varying changes in ω and the constraints on P that arise from power grid topologies and physical power constraints. To address these challenges, new smart grid models [14], [15] allow real-time power markets to evolve with changing system conditions such as outages or unpredictable RERs.

The power market utilizes (3) to minimize RP, effectively maximizing SW. Each actor is exposed to the price λ, and they adjust their power output/input to optimize their individual economic situations. For consumers, their individual profit is

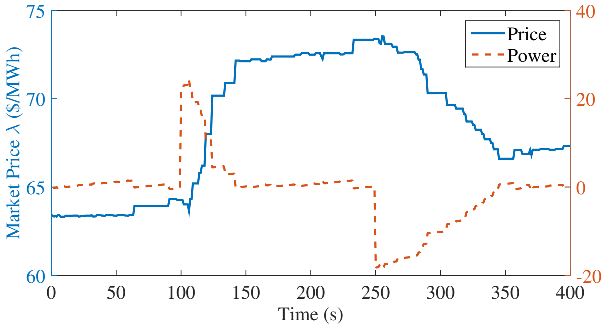

Fig. 1. The residual power (RP) and market price (λ) are shown for an example grid scenario based on prior market solution work, later described in Section V-A. The market experiences a demand surge in $\overline{P_i}$ for consumers at t=100 s followed by a reduction at t=250 s. During the surge, residual power spikes until the market price is corrected.

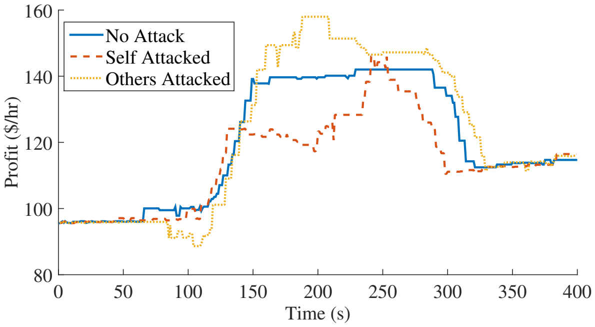

Fig. 2. The profitability for a generator is shown for two attack scenarios. At t=50 s a DoS attack is launched on the communication link connecting the generator to the market (self-attacked) or another market player (others attacked), and it lasts until t=200 s.

$\text{SW}_a = P_i \cdot (\omega_i - \lambda)$. If $\lambda > \omega$ for a consumer, then $f(\lambda) = 0$ since the consumer would experience a net loss by consuming energy. Changes in energy needs or production are captured by (4). For example, a wind power producer has a very low ω since wind is free and are thus driven by constraints in P.

$$\text{RP} = \sum_{\forall i} P_i(t) \tag{1}$$

$$\text{SW(t)} = \sum_{\forall i} \omega_i(t) P_i(t) - C \cdot |\text{RP(t)}| \tag{2}$$

$$P_i(t) = f_i(\lambda(t)) \tag{3}$$

$$\underline{P_i(t)} < P_i(t) < \overline{P_i(t)} \tag{4}$$

Online power markets optimize (2) by repetitively sampling P_i and updating λ. Each consumer receives a message containing λ and replies with $P_i = f(\lambda)$. Fig. 1 shows how a market evolves during a step transient in (4). The implementation of these systems, however, exposes security vulnerabilities that can be utilized by strategic adversaries to extract profit from the system. For example, the value of $f(\lambda)$ may be based on an outdated λ during a communication outage thus reducing the SW. The defensive strategies presented in this work curtail impacts to SW in a cost-effective manner.

C. Profit Manipulation

Network disruptions have a direct impact on market price (λ). Whenever a disruption occurs, the market player enters a zero-order hold mode. For market players, the price λ is fixed while P may change based on time-varying constraints. Consequently, the market loses its influence on power usage for a subset of market players whenever the network is disrupted. Fig. 2 demonstrates how the profits of a generator can be influenced by attacks on its communication link during the scenario shown in Fig. 1 and detailed in Section V-A. The generator loses money if its own link is disrupted and can gain additional profits when some competitors' links are disrupted. Market players are retroactively charged the actual λ market price to promote market participation—otherwise self-disconnection would be a valid strategy.

III. Attack and Defense Strategy

Definitions:

A set of market players
I set of target network links
SW social welfare or profitability of the system
SW_a profitability of market player a
$\text{IM}[a, i]$ impact or change in profit realized by market player a when network link i is attacked
P_i^{atk} probability of network link i being attacked
$P[a, i]$ probability of network link i being attacked, as estimated by market player a
D_i boolean indicating if network link i is defended
C_i cost to defend asset i
A_i boolean indicating if network link i is attacked by the SA
$A[i, n]$ two dimensional A_i for n iterations in a game with imperfect information at the various actors
O_a, O_i set of network links owned by market player a, owner of asset i
σ_a knowledge level of market player a

A. Players and Definitions

The set of market players A in the power market wish to maximize their profit SW_a, as defined in Section II-B. The power system is comprised of a set of assets and their communication links in I. The term asset refers to both the physical system consuming or producing energy and it's associated target, the network link. There is a one-to-one mapping of assets (and thus targets) to market players defined as *ownership* such that one market player may own multiple assets. Each actor has a defensive decision to make for each asset that it owns– whether or not to invest in its defense $D_i \in 0, 1$. This decision has a cost of defense C_i. If the asset is attacked, $A_i \in 0, 1$, then the system experiences the impact IM from the attack, unless $D_i = 1$ in which case the attack is assumed to fail via perfect defense.

B. Attacker's Incentive

The strategic adversary (SA) attempts to profit from the manipulations described in Section II-C by launching network attacks on the links that interconnect the market players with the market mechanism and the price signal λ. Each attack results in a change in the profitability of each market player, and this is captured in the impact matrix IM$[a, i]$ [13]. In this paper, IM$[a, i]$ is estimated via dynamic market simulations (Section V-A) by approximating the market conditions for each player and evaluating the resulting changes in profit in the market. For example, the impact of attacking each network link on the generator in Fig. 2 is summarized by IM. The SA wishes to maximize the gain in profit for some market players with whom she has a financial interest, shown in (5). A, I is the set of actors A and network links I to attack and profit from as the attacker's strategy.

$$\text{argmax}_{A,I} \sum_{a\in A, i\in I} \text{IM}[a, i] \tag{5}$$

The probability of an attack on target i is proportional to the attacker's incentive gained from that attack. Abstractly, the target i could be any perturbation in the system–network outages, power plant disruptions, transmission line faults, etc. In this paper, however, we are focusing on a dynamic pricing system for the smart grid, and the targets are limited to network link disruptions. Similarly, the actors that benefit from the attack A are collections of consumers and/or generators participating in the dynamic market.

C. Defensive Maneuvers

The market players in the system can estimate the IM via their own impact analysis. Using their individual IM, they can also estimate the attacker's strategy and use it to construct a corresponding defensive strategy. Without any budgetary constraints, the defenders will protect all the targets in I by investing in high capacity, secured network links. Budgets are limited, however, so defenders must optimally select targets to defend. Section IV describes how the defenders can have different views on the system parameters and still coordinate a defense.

1) Underlying Game: The impact matrix IM is computed by assessing the *underlying game*, i.e. the power market, with successful attacks, as described in Section II-C. Two versions of the system are compared–in one version, the attack was successful and in the other no attack is present. The resulting change in profitability for each actor is summarised by IM as the difference between the profits for each market player in the two scenarios.

Each attack on the system causes an overall net-negative impact on profitability. The system operates at a global-optimal whenever communications are uninterrupted. Any perturbations that disrupt communications result in decreased efficiency because of suboptimal responses to market prices (λ). Therefore, the sum across all actors for any given target is always zero or negative. Some actors, however, may benefit from competitor elimination, which is the basis for the strategic adversary's profit model. Table I shows an example impact matrix for three market players and three targets. Ai owns target Ti.

TABLE I
EXAMPLE IMPACT MATRIX (IM)

	T1	T2	T3
A1	-2	-2	3
A2	4	-4	-2
A3	-4	2	-4

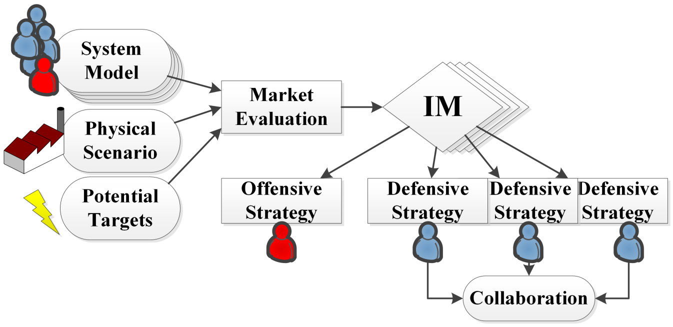

Fig. 3. In the overall system flow, set of defenders and a strategic adversary each have a view of the system and its market. The system is exposed to a physical scenario and analyzed for a set of potential targets via a market mechanism simulation. From this simulation, the profitability of each market player is captured in a set of impact matrices (IM). Each market player has an independently calculated IM and thus a different defense strategy that can be rectified via collaboration.

D. Defensive System Overview

The defensive investment optimization problem is designed to minimize the attacker's incentive thus reducing the probability of attack and denying profits to the adversary (resource exhaustion). An impact model is analyzed for each target and assessed as an impact to the profitability of each market player (IM). Once the matrix is calculated, it can be analyzed strategically to determine the best defensive action for each market player as in [13]. Fig. 3 shows the system layout.

1) Defensive Investments: Each market player in the system has a choice to defend self-owned targets from attacks at a cost C_i. If this cost is less than the expected reduction in profits, then it is in the actor's best interest to invest in defensive measures. The expected impact is IM$[a, i]P[a, i]$ where $P[a, i]$ is actor a's expectation that asset i will be attacked, based on the SA's optimal strategy. In game theory terms, we follow the Stackelberg model where the different parties move one after the other. In our case, the attacker's move is estimated and the defenders decide to defend the assets appropriately. The attacker does not have the ability to come up with a repeat attack after observing the defensive actions, so the Nash equilibrium point is not analyzed in this model. Due to the slow-moving nature of defensive investments, the defender's strategy is not immediately observable by the adversary.

IV. Multiple Knowledge Levels

A. Multiple Underlying Games

Each actor a and the strategic adversary have their own underlying game G_a, G_{SA}, respectively, from which IM is calculated. This arises because the parameters $\omega_i, f_i, \underline{P_i}, \overline{P_i}$ from Section II-B must be estimated by actors who do not own those assets, including the strategic adversary. Each actor assesses the impacts from their independent viewpoints of the system. The game G_a is derived from the ground truth game GA by adding noise to the above parameters. In the perfect knowledge model, all of the games are identical, $G_{a_1} = G_{a_2} \forall a_1, a_2$. Imperfect information is modeled by allowing the underlying games to diverge by sampling i.e. fictitious play.

The game G itself contains a set of dynamic parameters x ($\omega_i, f_i, \underline{P_i}, \overline{P_i}$) that are used to determine optimal market price. Fixed components of the game are the ownership and the network structure of the system. Each market player wants to keep its parameters secret to maintain a competitive edge in the marketplace. The dynamic parameters, however, can be estimated by observing market conditions and surveying physical equipment infrastructures. Each market player therefore can establish a "noisy" view of the underlying game, as defined by (6). The dynamic parameters are sampled from a normal distribution of the ground truth game. Sign changes are not allowed because it is assumed that each market player knows if an asset is a producer or consumer. The parameter σ_a defines the knowledge level of the actor a, and it is applied to all parameters except parameters for assets that the actor itself owns O_a. Intuitively, this parameter models the amount of information shared among each other. Greater is σ_a, less is the information that actor a has.

$$\begin{gathered} x' = \mathcal{N}(x, \sigma_a^2) \; \forall \, x \in GA, \, x \notin O_a, \\ x' = x \;\; \forall \, x \in GA, x \in O_a \\ x' \in (-\infty, 0] \text{ if } x' < 0, \text{ else } x' \in [0, \infty) \end{gathered} \tag{6}$$

B. Perfect Information Game

In the perfect information game, the strategic adversary and all actors have a perfect view of the system, $G_{\text{SA}} = GA$, $G_a = GA \; \forall \, a \in A$. In this form of the model, there is a single, optimal outcome for the defenders. Since $P[a_1, i] = P[a_2, i] \forall a_1, a_2 \in A$, the defensive decision is the same for each actor, and if the costs are correctly distributed among the defenders, then there is a single globally optimal defense strategy. The maximization problem (7) is solved by the defenders via mixed integer linear programming (MILP). The maximum value of this equation is zero because if no target is attacked, then no defense is necessary. Practically, protecting a network link (e.g. via DDoS protection) has some cost C_i for establishing a more reliable communication channel. The defender that owns each link, must decide to invest in its protection or not based on the likelihood of attack and the financial impact of the link outage.

$$\max \sum_{\forall a \in A} \sum_{i \in I} A_i \text{IM}[a, i](1 - D_i) - D_i C_i \tag{7}$$

The strategic adversary, the driving force behind A_i, is playing a similar game in (5). Since everyone shares the same knowledge, the perceived impact at each market player is the same, and all actors agree on which targets should be defended. Both the attacker and defender may have constraints on $\sum A_i$ and $\sum D_i C_i$ due to budget constraints on how many assets can be attacked and defended, respectively.

C. Imperfect Attack Strategies

The adversary is assumed to be perfectly rational (no anarchy) but may not have perfect knowledge of the system and subsequently makes suboptimal decisions. To capture this, A_i is evolved into a mixed probability-based strategy across several underlying games for the adversary. The SA has a single, optimal (pure) strategy per (5), and a mixed strategy is created by combining multiple pure strategies into a single mixed strategy. Multiple IM' are calculated for the strategic adversary's underlying game G_{SA} that are derived from GA as define in (6), with the caveat that the SA owns no assets. Equation (5) is optimized for each IM' across N fictitious games, each with a knowledge level σ, as a noise ratio. This results in N strategies for each asset i as $A[i, n]$. Equation (8) is the calculation for the probability of attack on target i given the N fictitious games for the adversary. The outcome P_i^a is an average of the boolean strategies for each of the SA's hypothetical games.

$$P_i^a = \frac{\sum_{n \, \in \, N} A[i, n]}{N} \tag{8}$$

Defense with Mixed Attack Strategies: The defenders strategy, as captured in (7), is modified below in (9) to account for the fact that the SA may have a non-boolean attack plan. Previously, A_i was binary and now P_i^a is a rational number so that the defender is operating on a mixed strategy.

$$\max \sum_{\forall a \in A} \sum_{i \in I} P_i^a \text{IM}[a, i](1 - D_i) - D_i C_i \tag{9}$$

D. Multiple Defender Optimization

The maximization problems presented earlier for optimizing defensive investments do not consider the scenario where multiple defenders do not have the same information level and are optimizing around different underlying games. Each defender's underlying game, G_a, is used in place of GA to calculate a mixed attacker strategy using (8). Each actor then has a different threat model $P[a, i]$ based on G_a instead of GA. (10) is performed by each actor to complete the optimization of D_i. Only the owner of asset i can determine the value of D_i. This approach enables no single actor to have a global view of the system which accurately models how a large interdependent system would operate.

$$\max \sum_{\forall a \in A} \sum_{\forall i \in O_a} P[a,i]\mathrm{IM}[a,i](1-D_i) - D_i C_i \qquad (10)$$

Cost Collaboration: This problem is supplemented with a collaboration method. The cost of defense of target i is proportionally shared among benefiting actors. Since defensive decisions are segmented by asset owner, and attacks against owned-assets are always damaging, there is no Price of Anarchy (PoA) in this defensive model.

V. Experimentation

A. Experimental Setup

The underlying game, as described in Section II-B, is solved via an online Nelder-Meade (NM) [16] optimization technique. Each iteration of NM is assumed to take one second and requires one round-trip communication of λ and P_i. The model for f_i is given in (12). In the case of a consumer, $\omega_i = \overline{P_i}$. For a producer, $\omega_i = \underline{P_i}$. The source code and corresponding market model details are available at [17].

$$\lambda_s = 6 * \frac{\lambda - \underline{\lambda_i}}{\overline{\lambda_i} - \underline{\lambda_i}} - 3 \qquad (11)$$

$$f_i(\lambda) = \frac{\overline{P_i} - \underline{P_i}}{1 + e^{\lambda_s}} + \underline{P_i} \qquad (12)$$

The model has 20 generators with an average $\overline{P_i} = 0, \underline{P_i} = 5, \underline{\lambda_i} = 30, \overline{\lambda_i} = 80$, and there are 100 consumers with $\overline{P_i} = 1.5, \underline{P_i} = 0.3, \underline{\lambda_i} = 0, \overline{\lambda_i} = 100$ for a total of 120 market players. The consumers $\underline{P_i}$ is modified as $\underline{P_i}'(t) = \underline{P_i} + 0.30 \cdot sin(\frac{\pi \cdot t}{3600})$, and all the other parameters are agnostic to time. At t=50s, the targeted network links are disrupted such that the λ term is fixed for those assets. At t=200s, the links are restored and communication is resumed. At t=100s a step load is introduced by setting $\overline{P_i} = 2, \underline{P_i} = 0.4$ for all consumers. The parameters are restored to the default values at t=250s. The simulation is executed for 400 seconds. This scenario can be seen in Figs. 1 and 2.

1) Communication Topology: In the experimental model, the communication paths between the market organizer (NM algorithm) and the individual market players are independent. In practice [18], however, there will be interdependence between communication failures across the different market players as many of them will share common links at some point in the communication path. For this reason, we have distributed the 120 market players on a tree topology with 4 top tier network links, 12 mid-tier links (3 for each top tier link), and 120 leaf links to better capture the interdependent networking impacts on smart grid topologies as shown in Fig. 4. Since future communication topologies have not yet been determined, and because dynamic markets may not operate on the same infrastructure as existing smart metering technologies, the topology used is purely speculative. As concrete topologies evolve, they can be substituted into this experimentation framework to identify changes in strategy and crucial network links.

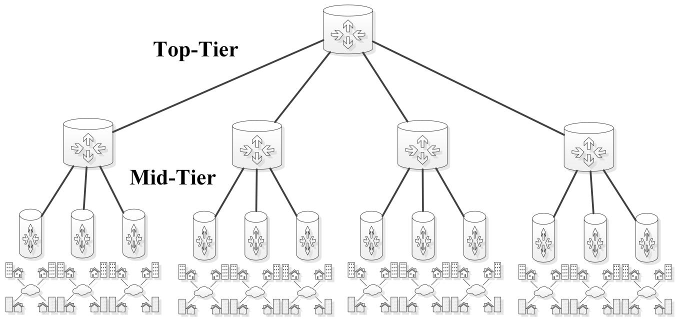

Fig. 4. The simulated network has four top-tier links, twelve mid-tier links, and one hundred twenty leaf links.

B. Experiment 1: The Attacker's Incentive

In this experiment, the strategic adversary attempts to maximize her incentive, as described in Section III-B. The strategy space I is the selection of links in the communication topology to disrupt. Each disruption results in a particular IM that is used to calculate the maximum attacker's incentive. Fig. 5 shows the attacker's incentive as a function of the number of links that she can attack simultaneously. As the number of simultaneously disrupt-able links increases, the attacker's incentive also increases. The attacker's incentive plateaus, however, when the overall system performance degradation becomes the dominating factor due to large numbers of link outages. The AI plateaus because the system as a whole becomes less profitable whenever most network links are disrupted (e.g. the top-level link attacks).

C. Experiment 2: Collaborating Defenders

In this experiment, the defense strategy presented in Section IV is evaluated. The defenders attempt to reduce the attacker's incentive shown in Fig. 5 by securing particular network links, thus eliminating them from the attacker's profit pool. The market players at each mid-tier communication hub are joined together so that there are 12 owners with 10 assets each. Collaboration is then possible on the mid-tier and top-tier network links, and they are the focus of this experiment.

Fig. 6 shows the reduction in attacker's incentive for different target budgets and a fixed $\sigma = 0.1$ for each owner. The defensive budget is progressively reduced relative to the number of attacked links. The defenders are able to significantly reduce the attacker's incentive in most large-attack cases at the leaf links.

Fig. 7 shows the reduction in AI for a range of knowledge levels across the defenders. In this case, 75 links are attacked and 75 links can be defended. The AI is maximally reduced when the owners knowledge levels are maximized ($\sigma \rightarrow 0$) indicating that collaboration can improve overall defense effectiveness.

VI. Related Work

Game theory applications for the smart grid [8] have become an increasingly important component of power system optimization. The core goals of these games, along with other

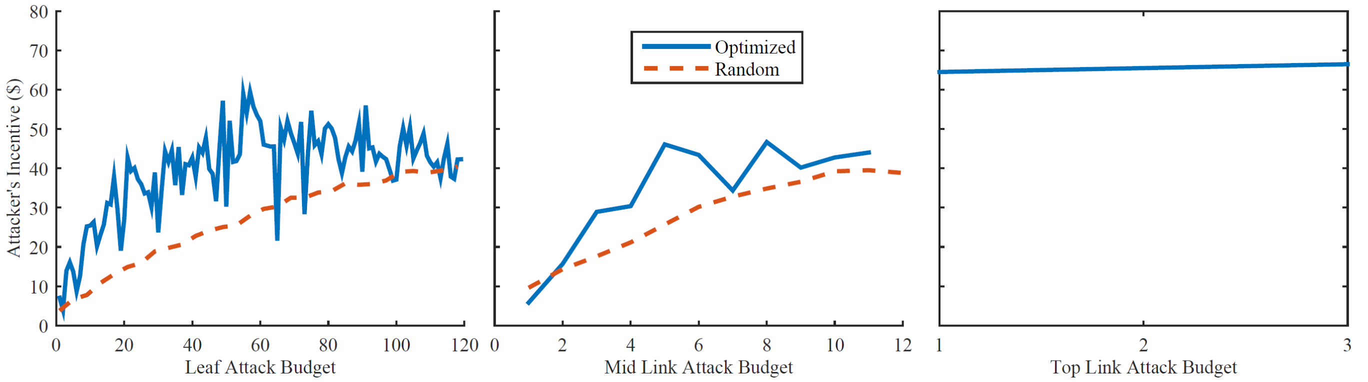

Fig. 5. The strategic adversary attempts to maximize her incentive by disrupting network links. In the graph on the left, the adversary is disrupting leaf-links in the communication topology. In the middle graph, the adversary is disrupting mid-tier links, and in the graph on the right, the top tier links are disrupted. In each case, the attacker's incentive is maximized for the given targets attacked. The strategy shown is maximized from (5) and compared to the mean of a random target selection.

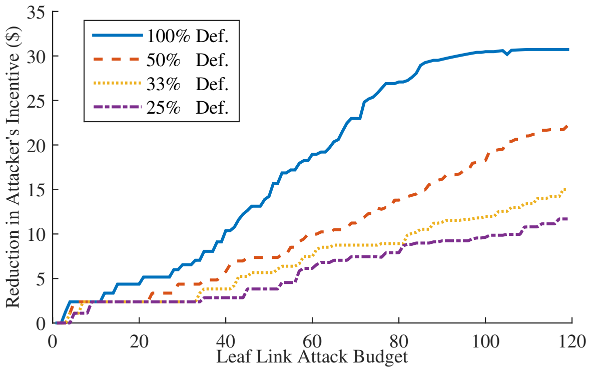

Fig. 6. The attacker's incentive is reduced by defensive investments in communication links. As the number of links attacked increases, the number of links defended also increases. The effectiveness of the defense, however, is reduced by imperfect knowledge levels among the defenders ($\sigma = 0.1$). Each line represents a different amount of aggregate defense budget, relative to the number of links attacked.

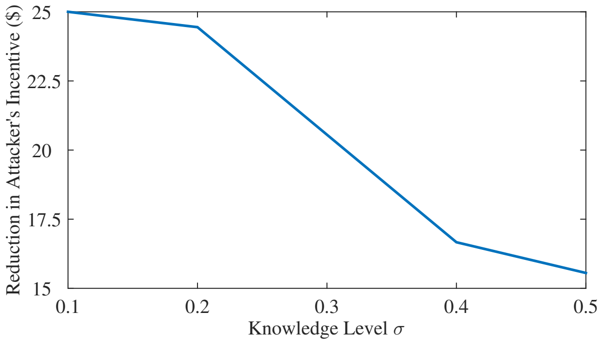

Fig. 7. This figure shows the cumulative reduction in attacker's incentive for different defender knowledge levels when 75 targets are attacked and defended. Decreased knowledge levels (high σ values) results in ineffective defense. As defenders are unwilling and unable to collaborate on defensive investments, the system suffers overall from poor defensive strategies.

approaches such as dynamic market mechanisms [14], are to improve the social welfare of the smart grid by utilizing market forces to balance demand with supply. The usefulness of these games, however, has not been well studied in the context of a strategic adversary that seeks to maliciously profit from the system by launching attacks.

A separate but related set of game theories optimize the security of information systems [19], [20] by playing attacker/defender games designed to create a defensive strategy that is optimal for a given adversarial model. Most of these approaches, however, utilize qualitative metrics for target valuation, costs, and attack success since it is difficult to value computer system breaches. The work presented in this paper applies attacker/defender concepts to a concrete smart grid cyber-physical system in which the utilities of attack and defense are derived from their actual operational influences. This also allows information exchanges to have quantitative impacts on success metrics.

Several attacker/defender games or security games have been constructed around Stackelberg games [21] for solving defensive investment optimization problems and scheduling patrols [22]. These games solve an attacker/defender model where the defender moves first in response to a perceived adversary and have been extended to support multiple human-modeled adversaries [23] in a computationally efficient way. These models, however, do not address defenders who exist in a competitive environment. The work presented in this paper analyzes attacker/defender games in a competitive environment.

The long-term financial impacts of attacks have been studied in [24]. Adversaries are modeled in [24] as has having budgets that deteriorate with unsuccessful attacks, resulting in reduced attack viability. The model captures some of the economic factors we have in this paper, but it does not make a connection between the physical system's behavior and the resulting financial outcome of attacks. Similarly, game theoretic techniques in [8], [9], [10], [11], [12] have proposed methods for determine how adversaries might manipulate the physical control systems via attacks, but they do not draw the financial connection between physical perturbations and adversarial profits. The work presented in this paper focuses on the financial motivations of attackers and defenders, resultant from system perturbations, as an attack and defense planning tool.

VII. Conclusion

In this paper, we presented a modeling technique to connect the networking components of a dynamic pricing market with a security strategy to defend against a profit-motivated adversary. This model allows the economics of cyberattacks on power markets to be captured and used to assess the risk to assets in the system. The amount of information that competing market players share about assets in the system is also modeled and used to analyze the benefits of collaboration in a defensive context. We then apply techniques to mitigate the attacker's incentive thus improving system resilience. We show that the baseline attacker incentives can be as high as 51% of baseline operating profits. When the defender and adversary's budget are equal, the attacker's incentive is reduced by up to 70%. These results validate the utility of this paper's technique in optimizing defensive investments. The model presented in the paper and the simulation results show promising approaches to countering the growing threat of cyberattacks in smart grids.

In future work, we are looking to model online learning aspects of dynamic pricing markets. In such a model, the attacker seeks to learn through iterative attacks, which also reveal more information about the system and the defensive strategies. Conversely, the defender also seeks to learn of the attack strategy through a multi-round strategy.

Acknowledgment

The authors would like to thank Christopher N. Gutierrez (gutier20@purdue.edu) for his assistance in reviewing the paper and for his feedback and discussion.

References

[1] "Ieee vision for smart grid communications: 2030 and beyond," *IEEE Vision for Smart Grid Communications: 2030 and Beyond*, pp. 1–390, May 2013.

[2] P. Siano, "Demand response and smart gridsa survey," *Renewable and Sustainable Energy Reviews*, vol. 30, pp. 461–478, 2014.

[3] S. Newell and A. Faruqui, "Dynamic pricing: Potential wholesale market benefits in new york state," *The Brattle Group*, 2009.

[4] P. Fairley, "Innovation amid a raucous rooftop solar squabble," *Spectrum, IEEE*, vol. 52, no. 7, pp. 14–15, July 2015.

[5] S. Katipamula, D. P. Chassin, D. D. Hatley, R. G. Pratt, and D. J. Hammerstrom, "Transactive controls: Market-based gridwisetm controls for building systems," *Pacific Northwest National Laboratory, Richland, WA.[Online]. Available: http://www. pnl. gov/main/publications/external/technical_reports/PNNL-15921. pdf*, 2006.

[6] R. Tan, V. Badrinath Krishna, D. K. Yau, and Z. Kalbarczyk, "Impact of integrity attacks on real-time pricing in smart grids," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 439–450.

[7] C. Barreto, A. A. Cárdenas, N. Quijano, and E. Mojica-Nava, "Cps: market analysis of attacks against demand response in the smart grid," in *Proceedings of the 30th Annual Computer Security Applications Conference*. ACM, 2014, pp. 136–145.

[8] W. Saad, Z. Han, H. V. Poor, and T. Basar, "Game-theoretic methods for the smart grid: an overview of microgrid systems, demand-side management, and smart grid communications," *Signal Processing Magazine, IEEE*, vol. 29, no. 5, pp. 86–105, 2012.

[9] K. J. Ross, "Application of game theory to improve the defense of the smart grid," DTIC Document, Tech. Rep., 2012.

[10] S. Backhaus, R. Bent, J. Bono, R. Lee, B. Tracey, D. Wolpert, D. Xie, and Y. Yildiz, "Cyber-physical security: A game theory model of humans interacting over control systems," *Smart Grid, IEEE Transactions on*, vol. 4, no. 4, pp. 2320–2327, 2013.

[11] S. Amin, G. A. Schwartz, and A. Hussain, "In quest of benchmarking security risks to cyber-physical systems," *Network, IEEE*, vol. 27, no. 1, pp. 19–24, 2013.

[12] Y. Mo, T.-H. Kim, K. Brancik, D. Dickinson, H. Lee, A. Perrig, and B. Sinopoli, "Cyber–physical security of a smart grid infrastructure," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 195–209, 2012.

[13] P. Wood, S. Bagchi, and A. Hussain, "Optimizing defensive investments in energy-based cyber-physical systems," in *Parallel & Distributed Processing Symposium Workshops (IPDPSW), 2015 IEEE International*. IEEE, 2015.

[14] D. Shiltz, M. Cvetkovic, and A. M. Annaswamy, "An integrated dynamic market mechanism for real-time markets and frequency regulation," 2015.

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

[16] R. R. Barton and J. S. Ivey Jr, "Nelder-mead simplex modifications for simulation optimization," *Management Science*, vol. 42, no. 7, pp. 954–973, 1996.

[17] P. Wood, "Source code for model and simulation," 2015. [Online]. Available: https://github.com/pcwood21/CPS_Model_Comsnet16

[18] Z. Wang, A. Scaglione, and R. J. Thomas, "Generating statistically correct random topologies for testing smart grid communication and control networks," *Smart Grid, IEEE Transactions on*, vol. 1, no. 1, pp. 28–39, 2010.

[19] A. Laszka, M. Felegyhazi, and L. Buttyan, "A survey of interdependent information security games," *ACM Computing Surveys (CSUR)*, vol. 47, no. 2, p. 23, 2014.

[20] M. H. Manshaei, Q. Zhu, T. Alpcan, T. Bacşar, and J.-P. Hubaux, "Game theory meets network security and privacy," *ACM Computing Surveys (CSUR)*, vol. 45, no. 3, p. 25, 2013.

[21] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus, "Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 895–902.

[22] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, M. Tambe, and F. Ordóñez, "Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service," *Interfaces*, vol. 40, no. 4, pp. 267–290, 2010.

[23] M. Brown, W. B. Haskell, and M. Tambe, "Addressing scalability and robustness in security games with multiple boundedly rational adversaries," in *Decision and Game Theory for Security*. Springer, 2014, pp. 23–42.

[24] K. Hausken, "Income, interdependence, and substitution effects affecting incentives for security investment," *Journal of Accounting and Public Policy*, vol. 25, no. 6, pp. 629–665, 2006.