CPS: Market Analysis of Attacks Against Demand Response in the Smart Grid

Carlos Barreto
University of Texas at Dallas
Richardson, TX, USA

Alvaro A. Cárdenas
University of Texas at Dallas
Richardson, TX, USA

Nicanor Quijano
Universidad de Los Andes
Bogotá, Colombia

Eduardo Mojica-Nava
National University of Colombia
Bogotá, Colombia

ABSTRACT

Demand response systems assume an electricity retail-market with strategic electricity consuming agents. The goal in these systems is to design load shaping mechanisms to achieve efficiency of resources and customer satisfaction. Recent research efforts have studied the impact of integrity attacks in simplified versions of the demand response problem, where neither the load consuming agents nor the adversary are strategic.

In this paper, we study the impact of integrity attacks considering strategic players (a social planner or a consumer) and a strategic attacker. We identify two types of attackers: (1) a malicious attacker who wants to damage the equipment in the power grid by producing sudden overloads, and (2) a selfish attacker that wants to defraud the system by compromising and then manipulating control (load shaping) signals. We then explore the resiliency of two different demand response systems to these fraudsters and malicious attackers. Our results provide guidelines for system operators deciding which type of demand-response system they want to implement, how to secure them, and directions for detecting these attacks.

1. INTRODUCTION

The smart grid refers to the modernization of current electric power networks to achieve better reliability, efficiency of resources, and to provide consumers more information and choices in the way they use electricity.

Research efforts have been mainly focused on the technological side of the smart grid; however, particular attention should be placed on individual consumer incentives, since individual agents (firms or people) within the smart grid are one of the enabling factors that will make the grid smart [13]. One of the particular smart grid programs that will rely on individual interactions between consumers and producers of electricity is *Demand Response* (DR) [11], a program that tries to address the retail electricity market inefficiencies.

Currently, the electricity price in the wholesale market (the bulk power grid) is updated periodically to match generation with dynamic demand between bulk power generators and bulk power consumers. When the transmission system is congested (which is the default state), Locational Marginal Prices (LMPs) are computed at each load and at each generation point to determine how much distribution utilities will pay the system operator (per Megawatt), and how much will the system operator pay the generation points. LMPs are traditionally computed every 5 to 10 minutes, but there is recent work (e.g., New York power system) for computing LMPs in real-time. In contrast, retail markets (which consist of an electric utility interacting with factories, buildings, homes, etc.) adopt static pricing schemes such as fixed and time-of-use tariffs. Under these contracts, consumers have limited incentives to adapt their electricity consumption to market conditions.

The goal of a DR program is to control consumer loads that are responsive to conditions in the electric power system, in order to achieve better efficiency of the *retail* market. Currently, the majority of DR programs are used by large commercial consumers, and companies such as EnerNOC manage DR services for large corporations and several government agencies in the U.S. Most active DR programs are designed for grid stability; however the focus of DR programs in the future is expected to be on energy efficiency. One of the primary goals of the smart grid is to make DR programs available to a much broader range of consumers [10].

1.1 Previous Work on Integrity Attacks in the Power Grid Markets

Integrity attacks (or false-data injection attacks) have been recently proposed as a way to analyze the vulnerability of cyber-physical systems in general, and electric networks in particular [18].

The area that we are interested in, is how false data injection attacks can affect the markets used in the smart grid and how false data injection can use the markets to drive the power grid to unsafe states (e.g., malicious attackers). Negrete-Pincetic et al. [21] were one of the first to study how false control signals can affect the social welfare of the electricity market. Related work by Xie et al. [30] studied how false data injection attacks can be used to defraud bulk electricity markets by modifying LMPs, and work by Liyan et al.[19] studied how false meter data in the bulk of the power grid can be used to cause the largest errors in LMP estimation.

All this work of false data analysis focuses on the bulk electricity market; however, the retail DR market has different models. In addition, it is more likely that attacks will happen in the retail market as there are many more

ACSAC '14, December 08 - 12 2014, New Orleans, LA, USA

http://dx.doi.org/10.1145/2664243.2664284

participants in retail with highly varying levels of trustworthiness. Finally, attributing attacks will be more difficult given a large number of participants (and thus attackers will have higher incentives for attacking retail markets than bulk-electricity markets).

Work on the impact of integrity attacks on the retail DR market were recently analyzed by Tan et al. [28], where they showed that an attacker who can modify the pricing signal sent to electricity consumers will affect the system and could cause severe oscillations of electricity load. They presented a new DR model and then experimented with two different attack models, scaling attacks and delay attacks.

While their model is an important step towards understanding the resiliency of DR programs against attacks, it has two limitations. First, Tan et al. [28] introduce a new model which has not been validated by the smart grid community. The current consensus for modeling DR problems is to incorporate market interactions of a multi-agent system where each agent has a nonlinear valuation of electricity [25, 24, 14, 26, 7, 15, 12, 9, 17]. In this paper we use more representative DR models and study their security. Furthermore, in addition to dynamic pricing DR which is the focus of study by Tan et al., we study direct-load control DR models as well.

The second limitation is that the attacks considered by Tan et al. are limited to be parametric models of the pricing signal $u(t)$; these are delay attacks $u(t-\tau)$ and scaling attacks $\alpha u(t)$. One contribution in this paper is to model a more powerful attacker that is not constrained to only two possible attack strategies, but that can select an arbitrary attack signal $\hat{u}(t)$. In addition, the previous model of the attacker is not strategic; in this paper we model a *strategic* attacker that will select an attack strategy in a principled way, and in order to achieve a specific attack goal. For example, 1) a malicious attacker will have a goal of damaging the power grid by generating sudden overload spikes, whereas 2) a selfish attacker will try to defraud DR programs.

1.2 Contributions

In this paper we address the limitations of previous work and propose new contributions.

- We introduce two attacker models against DR programs: (1) A fraudster who tries to steal electricity without trying to damage the power grid, and (2) a malicious attacker that tries to damage the power grid. Our models assume non-parametric adversary models and are therefore more powerful that adversaries considered in previous work.
- We provide a formal security analysis for the two types of adversary models for two DR programs: dynamic pricing and direct-load control, using models previously proposed by DR communities.
- We show that dynamic pricing is more resilient to fraud and malicious attacks than direct load control mechanisms.
- Previous work analyzing DR [7, 15, 12, 9, 17] consider only equilibrium points (i.e., optimal steady states); they do not consider the transient dynamics of the agents adjusting to different market conditions and learning optimal outcomes. In this paper we design an evolutionary game-theory implementation of the transient dynamics of the DR problem. This is necessary for studying malicious attacks that need to create *sudden* electricity peaks, and therefore require the study of transient dynamics to their attacks (instead of finding the steady-state equilibrium). Using this transient analysis we show how a sophisticated attacker can manipulate the market to achieve better results than naive attacks. In an effort to facilitate future research, we are making the dynamic implementation for these simulations available online as an open source BSD project [6].

2. DEMAND RESPONSE MODELS

There are two main forms of DR programs: direct load control and dynamic pricing.

2.1 Direct Load Control

Direct Load Control (DLC) [11] is a demand control system in which the utility (or a DR broker) negotiates with consumers the ability to directly control flexible loads in their homes, buildings or industries. The utility company or companies like Trilliant [29] can use remote appliance controllers to turn specific appliances on and off during peak demand periods and critical events. These remote controllers can manage water heaters, pool pumps, and air conditioners (among others) and can be programmed to respond to time-of-use tiers, critical peak pricing events, and direct load control events.

DLC has been a promising future direction for the smart grid for a variety of reasons. By controlling loads which can be modified without much impact on consumer satisfaction, we can allay many costs by shifting loads from peak demand and compensating for real-time load imbalances. For example, Pacific Gas & Electric deployed the SmartAC program in Spring 2007 [1]. Another provider of DR services has recruited over 1.25 million residential customers in DLC programs, and has deployed over 5 million DLC devices in the United States. In California, they have successfully curtailed over 25 MW of power consumption since 2007 [5]. In Hungary, for example, DLC accounts for 1600MW (25% of peak consumption).

2.2 Dynamic Pricing

Dynamic pricing programs [11] use incentives (e.g., via real-time pricing, rebates, etc.) to motivate consumers to reduce electricity consumption during peak hours. In contrast to direct-load control, consumers will be responsible for taking actions based on the incentive (control) signals; giving consumers a choice between cost and convenience. Therefore, while direct-load control is a *centralized* system, in dynamic pricing, agents make decisions in a *decentralized* way. By using dynamic pricing, utilities can create incentives for consumers to distribute their load more evenly—e.g., consume more energy when there is high wind or solar energy in the grid, and reduce consumption during peak demand times. This price-sensitive peak shaving will defer the need for grid expansion and will reduce the investments on generators that are only used for short peak demands.

2.3 Notation

We consider a population of N consumers of electricity. We divide a period of 24 hours in a set of T time intervals denoted $\tau = \{\tau_1, \ldots, \tau_T\}$; formally, we define the set τ as a partition of $[0, 24)$, where $\cup_{t\in\{1,\ldots,T\}}\tau_t = \tau$ and $\cap_{t\in\{1,\ldots,T\}}\tau_t = \varnothing$. Furthermore, we denote with q_i^t the electricity consumption of the i^{th} user in the t^{th} time interval. The daily electricity consumption of the i^{th} user is represented by the vector $\boldsymbol{q}_i = [q_i^1, \ldots, q_i^T]^\top \in \mathbb{R}^T_{\geq 0}$. The population consumption at a given time t is defined by the vector $\boldsymbol{q}^t = [q_1^t, , q_2^t \ldots, q_N^t]^\top \in \mathbb{R}^N_{\geq 0}$, and the joint electricity consumption of the whole population is denoted by $\boldsymbol{q} = [\boldsymbol{q}_1^\top, \ldots, \boldsymbol{q}_N^\top]^\top$. The aggregated consumption at a given time t is defined as $||\boldsymbol{q}^t||_1 = \sum_{j=1}^N q_j^t$, where $||\cdot||_1$ is the

1-norm.

A *valuation function* $v_i^t(q_i^t)$ models the *valuation* that the i^{th} user gives to an electricity consumption of q_i^t units in the t^{th} time interval. Finally, let $p(\cdot) : \mathbb{R} \to \mathbb{R}$ be the price of electricity charged to consumers.

2.4 Direct Load Control Model

DLC assumes a perfect competition market where a social planner wants to maximize the social welfare of a population. This problem can be represented by the following optimization problem [16, 8, 17]:

$$\begin{aligned} \underset{\boldsymbol{q}}{\text{maximize}} \quad & \sum_{i=1}^{N} U_i(\boldsymbol{q}) = \sum_{i=1}^{N}\sum_{t=1}^{T}\left(v_i^t(q_i^t) - q_i^t p(||\boldsymbol{q}^t||_1)\right) \\ \text{subject to} \quad & q_i^t \geq 0, i = \{1,\dots,N\}, t = \{1,\dots,T\}, \end{aligned} \tag{1}$$

where $U_i(\boldsymbol{q})$ represents the profit (valuation of electricity consumption minus the electricity bill) of the i^{th} customer in function of the population demand profile $\boldsymbol{q}$. Note that in this model users send their valuation for electricity to a central planner, and the central planner then decides the amount of electricity and price to charge to each agent. Here we make some assumptions on the problem characteristics in order to guarantee that the problem has a unique solution.

Assumption 1.

1. *The valuation function $v_i^t(\cdot)$ is differentiable, concave, and non-decreasing.*
2. *The price $p(\cdot)$ is differentiable, convex, and non-decreasing.*

Assumption 2. *The solution of the optimization problem in Eq. (1) is inside the feasible region, that is*

$$\frac{\partial}{\partial q_i^t} U_i(\mathbf{0}) > 0$$

Therefore, the First-Order Conditions (FOC) of this problem at the maximum, denoted by $\boldsymbol{\mu}$, are:

$$\begin{aligned} \frac{\partial}{\partial q_i^t}\sum_{i=1}^{N} U_i(\boldsymbol{q})\Big|_{\mathbf{q}=\boldsymbol{\mu}} &= \frac{\partial}{\partial q_i^t} v_i^t(q_i^t) \cdots \\ &\cdots - p(||\boldsymbol{q}^t||_1) - ||\boldsymbol{q}^t||_1 \frac{\partial}{\partial q_i^t} p(||\boldsymbol{q}^t||_1)\Big|_{\mathbf{q}=\boldsymbol{\mu}} = 0. \end{aligned} \tag{2}$$

The consumption profile $\boldsymbol{\mu}$ is efficient in the sense of Pareto.

2.5 Dynamic Pricing Model

In a decentralized version of the model proposed in the previous section, each agent would need to maximize individually their utility [16]:

$$\begin{aligned} \underset{\boldsymbol{q}_i}{\text{maximize}} \quad & U_i(\boldsymbol{q}_i, \boldsymbol{q}_{-i}) = \sum_{t=1}^{T}\left(v_i^t(q_i^t) - q_i^t p(||\boldsymbol{q}^t||_1)\right) \\ \text{subject to} \quad & q_i^t \geq 0, i = \{1,\dots,N\}, t = \{1,\dots,T\}. \end{aligned} \tag{3}$$

The selfish actions of each individual might lead to the outcome $\boldsymbol{\xi}$, the Nash equilibrium of the game. Under Assumptions 1 and 2, we can show that the first order conditions for this selfish optimization problem guarantee a unique maximum, i.e., a unique Nash equilibrium. In addition, it is easy to show that these first order conditions (Nash equilibrium) do not match those from Eq. (2) (Pareto equilibrium). Hence, the Nash equilibrium $\boldsymbol{\xi}$ of the game is not efficient in the sense of Pareto. In general, in a strategic environment

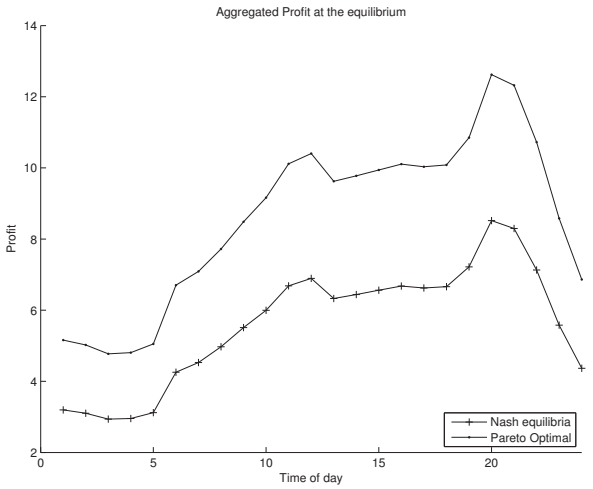

(a) Aggregated utility.

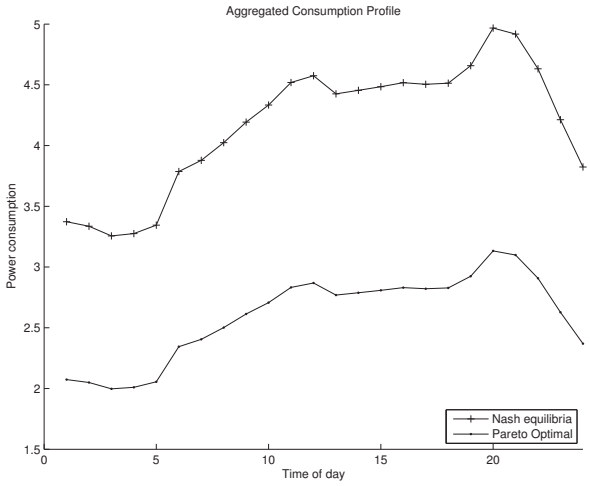

(b) Aggregated consumption.

Figure 1: Aggregated utility and consumption of the population at the inefficient outcome (Nash equilibrium) and the optimal outcome (Pareto equilibrium).

that satisfies Assumptions 1 and 2, the outcome is inefficient and requires an external intervention, such as economic incentives [3, 4].

We now show a numerical example of how the Nash equilibrium is different than the Pareto equilibrium. We select some typical numerical values and functions previously used in the literature [23, 20]:

$$v_i^t(q_i^t) = \alpha_i^t \log(1 + q_i^t),\ \alpha_i^t > 0, \tag{4}$$

$$p(\|\boldsymbol{q}\|_1) = \beta\|\boldsymbol{q}\|_1,\ \beta > 0. \tag{5}$$

These functions satisfy Assumptions 1 and 2. In this case, we consider $T = 24$ time periods and define the valuations of each individual using consumption measurements provided by the Colombian electricity system administrator [**?**] (a detailed implementation of the simulations can be found in [4]). We can see how the Nash equilibrium produces less total utility for all parties (Fig. 1a) and produces more power consumption (Fig. 1b) than the Pareto equilibrium.

While the Nash equilibrium $\boldsymbol{\xi}$ is suboptimal, we can show,

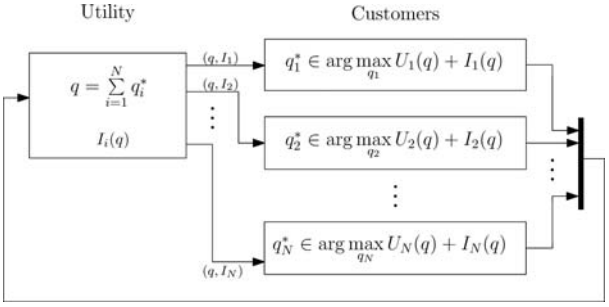

Figure 2: Dynamic pricing DR model.

however, that if we consider an added incentive to the individual cost function of each player, the Nash equilibrium of the game with incentives can be made efficient in the sense of Pareto. In order to incentivize agents to modify their behavior for the good of the population, the social planner sends them an incentive (e.g., a dynamic price signal or reward) to indirectly control their load; therefore the new cost function for the i^{th} agent is

$$W_i(\boldsymbol{q}_i^t, \boldsymbol{q}_{-i}^t) = \sum_{t=1}^{T} \left(v_i^t(q_i^t) - q_i^t p\left(\left\| \boldsymbol{q}^t \right\|_1 \right) + I_i(\boldsymbol{q}^t) \right). \quad (6)$$

where incentives are of the form:

$$I_i(\boldsymbol{q}^t) = \left\| \boldsymbol{q}_{-i}^t \right\|_1 \left(p\left(\frac{N}{N-1} \left\| \boldsymbol{q}_{-i}^t \right\|_1 \right) - p\left(\left\| \boldsymbol{q}^t \right\|_1 \right) \right) \quad (7)$$

The form of this incentive is related to the price used in the Vickrey-Clarke-Groves mechanism [22] and some utility functions used in potential games [2]. Note that with these incentives it can be shown that the first-order conditions of Eq. (6) are the same to the first-order conditions of Eq. (2); and therefore, the Nash equilibrium of the system with incentives is equal to the optimal outcome of the DLC model.

This dynamic pricing model is depicted in Fig. 2, where the incentive I can represent the dynamic pricing signal. In this DR approach we consider that the utility sends a two dimensional signal to each customer, namely the total consumption and the incentive $(\boldsymbol{q}, I_i)$—in a practical implementation the utility would send the consumer the price of electricity at the current time interval: $p(\boldsymbol{q})$ instead of the consumption $\boldsymbol{q}$ but this does not affect our analysis—and each customer responds with some consumption $\boldsymbol{q}_i$. Note that the incentives modify the price paid by each user according to their relative consumption. Hence, two different users receive different incentives as long as their consumption is different. Specifically, users who introduce less externalities in the system receive larger incentives or rewards.

2.6 Transient (Evolutionary) Analysis

In the previous two sections we introduced solutions used in game theory to find equilibrium points (Pareto, Nash). These solutions provide information about the system in steady state, but overlook the trajectory followed to reach such solutions. Furthermore, while finding the steady state equilibrium is good for modeling the final outcome of an attack, we also need to consider the transient dynamics that show how consumers will behave to changes in the market and to maliciously injected signals. Thus, we can model malicious attackers that will try to create sudden electricity overloads in the system and damage power distribution equipment or produce cascading failures.

Population dynamics [4] can be used to model negotiation approaches between players of a game and are used to model the transient dynamics of the system evolution before it reaches a steady state. From the perspective of population games [27], we have a multi-population game, with the following characteristics:

- There are N populations. Each population is associated with an agent.
- The resource to be allocated in each population is the daily power consumption $||\boldsymbol{q}_i||_1$.
- The strategy of each user is the consumption at T time intervals plus the consumption q_i^{T+1} that represents a slack variable, i.e., the power not consumed in any time interval is represented as a consumption in the $(T+1)^{th}$ time interval. Hence, in each population there are $T+1$ possible strategies.
- The fitness functions are defined as the derivative of the utility function $U_i(\boldsymbol{q})$, e.g., the fitness (under normal operation) is the marginal utility, that is defined as:

$$f_i^t(\boldsymbol{q}^t) = \frac{\partial U_i^t(\boldsymbol{q}^t)}{\partial q_i^t} \quad (8)$$

  and $f_i^{T+1} = 0$.

We use replicator dynamics to solve the resource allocation problem in Eq. (3). These dynamics might be seen as a set of deterministic rules that guide the resource allocation process to find an outcome that maximizes the utility of each agent. In particular, the allocation is carried out by evaluating the convenience of consuming some resources at a given hour. An equilibrium is achieved when an agent cannot increase its profit by redistributing its resources (i.e., a Nash equilibrium).

Replicator dynamics are described by the differential equation

$$\dot{q}_i^t = q_i^t \left(f_i^t(\boldsymbol{q}^t) - \bar{f}_i^t(\boldsymbol{q}^t) \right), \quad (9)$$

where $\bar{f}_i^t(\boldsymbol{q}^t) = \sum_{t=1}^{T} q_i^t f_i^t(\boldsymbol{q}^t)$ is the average payoff the population i.

While any actor (including attackers) can change their actions arbitrarily, we consider that changes in power consumption are bounded; hence, we model smooth demand changes. A continuity notion that satisfies this requirement is Lipschitz continuity, which is also a requirement for the existence of a solution to a differential equation.

3. ADVERSARY MODEL

We assume an adversary model that compromises the central system where the control signals are computed. For 1) DLC, this means the attacker can arbitrarily send commands to curtail loads to consumers, and for 2) dynamic pricing, this means that the attacker can send arbitrary incentive price signals to consumers.

In addition, we assume two types of attackers: 1) a fraudster, whose objective is to defraud the system and pay less for electricity (a version of electricity theft, only that instead of falsifying their electricity consumption, it falsifies information sent to consumers), and 2) a malicious attacker, whose objective is to damage the electricity distribution system.

A fraudster does not necessarily want to attack the electric grid and the consumers of electricity (if the grid is down the fraudster would get no utility), but wants to exploit the system into behaving in unanticipated ways for personal gain, such as paying less electricity than others. We assume the attacker will still be charged at the correct price

for the electricity she consumes (that is the attacker has not compromised the metering system). Defrauding the utility company by compromising the control signals from DR algorithms might be even more beneficial for the attacker than compromising the meter readings, because if an attacker compromises the meter installed in its neighborhood and is detected, then the utility has evidence to attribute the attack; however, if the attack to the control signals is detected, the fraudster can still claim deniability of this attack as it is not immediately obvious who the culprit of the attack is (in particular, the attacker can mask itself in a large group of beneficiaries with a parameter γ that we will introduce later).

We define a malicious attacker as an adversary whose goal is to cause damage to the system and all their players. One practical way to achieve this is to cause the maximum sudden overload in the power distribution network, which can potentially cause blackouts because of equipment failures (e.g. burnt transformers) and circuit breakers opening.

We note that in contrast to previous work, our attack can be arbitrary, and thus can model realistic attackers not tied to pre-specified attacks (such as delays). In addition, we model two different strategic attackers that will try to achieve an objective, and will select the attack signal to achieve this objective. For example, in the dynamic pricing model the attacker will use mechanism design to generate the incentives that will achieve its own goals.

4. FRAUDSTER (SELFISH) ATTACKER

4.1 Direct Load Control

The goal of the fraudster is to achieve the best possible selfish utility by manipulating all direct-load control signals $\mathbf{q}$. We consider a scenario with full information. The best possible outcome for a fraudster with control of all control signals can be represented by the following optimization problem:

$$\begin{aligned} &\underset{\boldsymbol{q}_i, \boldsymbol{q}_{-i}}{\text{maximize}} \quad U_i(\boldsymbol{q}_i, \boldsymbol{q}_{-i}) \\ &\text{subject to} \quad q_i^t \geq 0, i = \{1, \ldots, N\}, t = \{1, \ldots, T\}. \end{aligned} \tag{10}$$

If Assumptions 1 and 2 are satisfied, then there is a unique solution to Eq. (10), denoted by q^*, that satisfies the following first order conditions:

$$\nabla U_i(\boldsymbol{q})\Big|_{\boldsymbol{q}=q^*} = 0.$$

Specifically, the gradient of the utility for an attacker (the i^{th} agent) is

$$\frac{\partial}{\partial q_i^t} U_i(\boldsymbol{q}) = \frac{\partial}{\partial q_i^t} v_i^t(q_i^t) - p(||\boldsymbol{q}^t||_1) - q_i^t \frac{\partial}{\partial q_i^t} p(||\boldsymbol{q}^t||_1), \tag{11}$$

and the gradient for a victim j is

$$\frac{\partial}{\partial q_j^t} U_i(\boldsymbol{q}) = -q_i^t \frac{\partial}{\partial q_j^t} p(||\boldsymbol{q}^t||_1), \tag{12}$$

for all $i, j \in \mathcal{P}$, such that $j \neq i$ and $t \in \{1, \ldots, T\}$.

To implement these results, we consider a random state $\boldsymbol{q}$, from which the population evolves with the dynamics of a uncompromised system (see Eq. (2)). At simulation time 2, an attacker compromises the system, causing a switch in the system dynamics. The fitness functions implemented in the replicator dynamics are then changed from Eq. (2) to those defined by Eq. (11) and (12). Fig. 3a shows the evolution of both utility and consumption of the population. Fig. 3b show the final state of the attacker in a compromised and uncompromised system. From these figures we can see that the attack is successful in obtaining more electricity consumption while maximizing its utility; however, we can see that the average utility function of the population decreases while the aggregated consumption is decreased and the population is no longer in a social optimum.

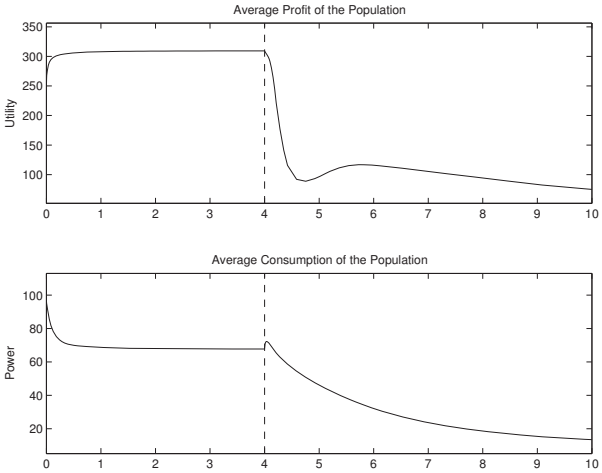

(a) The utility function of the population decreases when the Fraudster launches its attack.

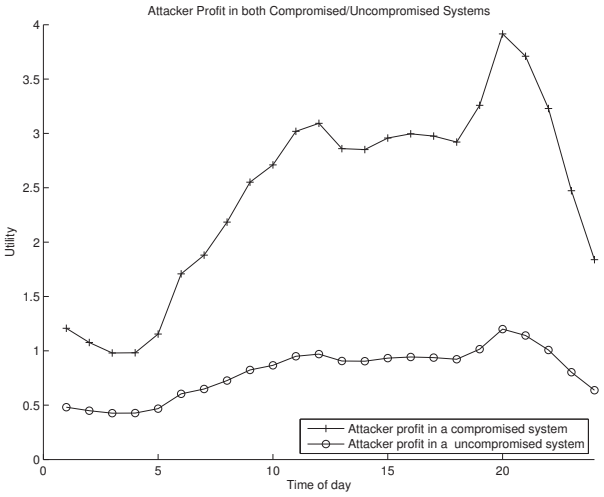

(b) The Fraudster can create an attack that maximizes its utility and allows it to consume more power than with no attack.

Figure 3: Fraudster attack for DLC model.

By maximizing Eq. (10), the fraudster would be the only one benefiting from the system, and the identity of the attacker might be uncovered by the central authority if it monitors the term q_i^t in the FOC of some victims–Eq. (12). Therefore, the attacker might want to mask its actions by increasing the utility of a subset $\mathcal{S}$ of the population (including the attacker); in this way, the attacker can gain plausible deniability (i.e., lack of evidence proving an allegation) as it is only one among a set of beneficiaries. In addition to make the attack more subtle (and thus harder to detect), the attacker can use a parameter λ to quantify how much of the utility function it wants to maximize compared to the utility that others will receive:

$$\begin{aligned} \underset{\boldsymbol{q}}{\text{maximize}} \quad & \lambda \sum_{h\in\mathcal{S}} U_h(\boldsymbol{q}) + \sum_{h\in\mathcal{V}} U_h(\boldsymbol{q}) \\ \text{subject to} \quad & q_i^t \geq 0, i = \{1,\dots,N\}, t = \{1,\dots,T\}, \end{aligned} \tag{13}$$

where $\lambda \geq 1$ represents the severity of the attack and $\mathcal{V}$ and $\mathcal{S}$ are two disjoint nonempty sets of consumers. The cardinality of each set is denoted as $N_v = |\mathcal{V}|$ and $N_S = |\mathcal{S}|$. Therefore, the FOC of the problem in Eq. (13) are:

$$\lambda \left(\frac{\partial}{\partial q_i^t} v_i^t(q_i^t) - p(||\boldsymbol{q}^t||_1) - \sum_{h\in\mathcal{S}} q_h^t \frac{\partial}{\partial q_i^t} p(||\boldsymbol{q}^t||_1) \right) \cdots$$
$$\cdots - \sum_{h\in\mathcal{V}} q_h^t \frac{\partial}{\partial q_i^t} p(||\boldsymbol{q}^t||_1) = 0. \tag{14}$$

for an agent $i \in \mathcal{S}$, and

$$\frac{\partial}{\partial q_j^t} v_j^t(q_j^t) - p(||\boldsymbol{q}^t||_1) - \sum_{h\in\mathcal{V}} q_h^t \frac{\partial}{\partial q_j^t} p(||\boldsymbol{q}^t||_1) \cdots$$
$$\cdots - \lambda \sum_{h\in\mathcal{S}} q_h^t \frac{\partial}{\partial q_j^t} p(||\boldsymbol{q}^t||_1) = 0. \tag{15}$$

for an agent (victim) $j \in \mathcal{V}$.

For illustration purposes, let us assume an homogeneous population in which agents have the same consumption preferences. Since the population is homogeneous, we know that the consumption of all the members of a set (either $\mathcal{S}$ or $\mathcal{V}$) is the same. That is, users that are in the same conditions must have the same consumption at the solution of Eq. (13), denoted by $\boldsymbol{x}$. In this case, we denote by x_s and x_v the consumption of users from either $\mathcal{S}$ or $\mathcal{V}$, respectively. We can take into account this property, as well as the form of the price function $p(z) = \beta z + b$, to obtain the following expressions from Eq (14) and (15):

$$\frac{x_s}{f_v(x_v)} = \frac{1}{\beta N_s(1+\lambda)}, \tag{16}$$

$$\frac{x_v}{f_s(x_s)} = \frac{\lambda}{\beta N_v(1+\lambda)} \tag{17}$$

where $f_s(x_s) = \dot{v}_i(x_s) - 2\beta N_s x_s - b$ and $f_v(x_v) = \dot{v}_j(x_v) - 2\beta N_v x_v - b$. Now, if we divide Eq. (16) and (17) we obtain

$$\frac{x_s f_s(x_s)}{x_v f_v(x_v)} = \frac{N_v}{\lambda N_s}. \tag{18}$$

Note that $f_\omega(x_\omega)$ is equivalent to the derivative of

$$\hat{U}_\omega(\boldsymbol{q}) = \sum_{h\in\omega} \left(v_h(q_h) - q_h p \left(\sum_{h\in\omega} q_h \right) \right) \tag{19}$$

with respect to some q_h, and evaluated at the equilibrium $\boldsymbol{x}$, for $\omega = \{\mathcal{S}, \mathcal{V}\}$. We can interpret $\hat{U}_\omega$ as the welfare of agents belonging to the subset ω. Hence, the term $x_\omega f_\omega(x_\omega)/2$ can be seen as an approximation of the utility of an isolated population in Eq. (19).

Summarizing, Eq. (18) gives information about the utility ratio between an attacker and a victim in function of λ and the number of agents in each subpopulaiton. If we consider $N_s = \gamma N$ and $N_v = N - N_s$, then Eq. (18) can be rewritten as

$$\frac{x_s f_s(x_s)}{x_v f_v(x_v)} = \frac{1}{\lambda} \frac{1-\gamma}{\gamma}. \tag{20}$$

Eq. (20) shows the attacker's utility as a function of the proportion of agents that benefit from the attack (γ). This relation holds as long as the consumption of the victims is different from zero.

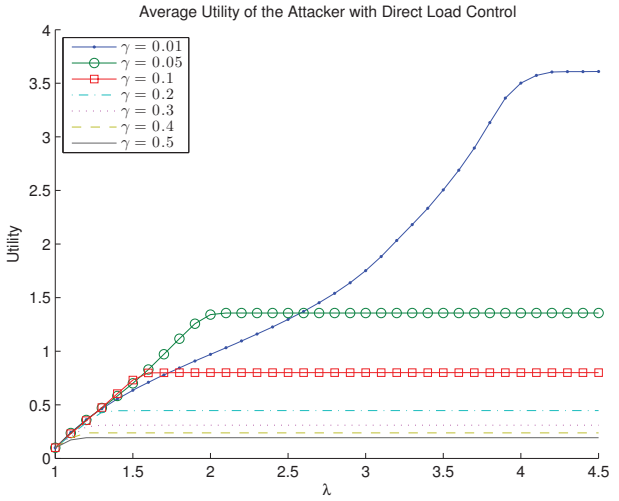

Figure 4: Utility of fraudster as a function of the parameter λ for different values of γ in the DLC model.

Fig. 4 shows the utility of the attacker for different values of λ (severity of attack) and γ (percentage of benefited consumers). Now, observe how the maximum benefit of the attacker decreases significantly with the proportion of agents benefiting from the attack (γ): in other words, the price that the attacker pays in order to maintain plausible deniabiliy if the attack is detected increases significantly. Fig. 4 also shows that the utility of the attacker does not increase indefinitely with λ but reaches a saturation point where the electricity consumption by the affected population is so low, that the attacker cannot gain any more by sending them control signals requesting lower consumption. It is interesting to note also that for small values of λ, the attacker does not gain much by being the single user benefiting from the attacks; in fact, for twice the importance of the selfish utility ($\lambda = 2$) and a population of 5% consumers benefiting form the attack ($\gamma = 0.05$), the attacker is better off than when it is the only recipient of the benefits ($\gamma = 0.01$ in a population of 100).

4.2 Dynamic Pricing

Recall that in the dynamic pricing model the central authority sends incentives I to drive the agents towards a Pareto optimal point.

$$\max_{\boldsymbol{q}_i} W_i(\boldsymbol{q}_i, \boldsymbol{q}_{-i}) = U_i(\boldsymbol{q}) + I_i(\boldsymbol{q}).$$

with respect to $\boldsymbol{q}_i$ (where I_i is the incentive signal sent by the utility to drive the system to a Pareto equilibrium). We formulate the goal of the fraudster in the DR with incentives case as the following optimization problem:

$$\underset{I, \boldsymbol{q}_i}{\text{maximize}} \quad U_i(\boldsymbol{q}) + I_i(\boldsymbol{q}), \tag{21}$$

Because in contrast to the DLC model, in the dynamic pricing model the central authority does not know the valuation functions v_i, the objective function in Eq. (21) cannot be optimized by an attacker, even if it compromises the central system. Therefore, we first use an approximation to this objective function that can be solved and then we compare how close it achieves the real objective in Eq. (21).

We now define the goal of the attacker in the dynamic pricing model to find an incentive signal I to drive the system towards the solution of Eq. (13). λ and γ are again parameters that the attacker can select.

Let $\boldsymbol{q}^s$ be the solution to Eq. (13). Recall that in this case the attacker cannot control $\boldsymbol{q}$ but instead controls I. Leveraging the theory of mechanism design we can show that an attacker can incentivize all agents to adopt $\boldsymbol{q}^s$ by sending the following false incentives:

$$I_j(\boldsymbol{q}) = \left(\sum_{h \in \mathcal{V}-j} q_h + \lambda \sum_{h \in \mathcal{S}} q_h \right) \left(\frac{N}{N-1} p(||\boldsymbol{q}_{-j}||_1) - p(||\boldsymbol{q}||_1) \right),$$

for all $j \in \mathcal{V}$ and

$$I_i(\boldsymbol{q}) = \left(\frac{1}{\lambda} \sum_{h \in \mathcal{V}} q_h + \sum_{h \in \mathcal{S}-i} q_h \right) \left(\frac{N}{N-1} p(||\boldsymbol{q}_{-i}||_1) - p(||\boldsymbol{q}||_1) \right),$$

for all $i \in \mathcal{S}$.

For λ large enough, this attack will (for practical purposes) maximize $U_i(\boldsymbol{q})$. Note that this attack does not require the valuation function of each user, but it needs to know the total consumption of either the users that benefit the attack or those who do not.

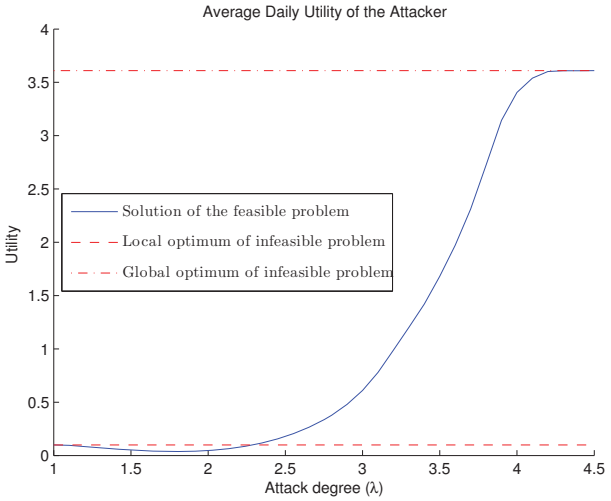

Figure 5: Impact of the attack in the fraudster's utility as a function of λ.

Now, let us study how well our approximate problem solves the ideal (but infeasible) objective function of the attacker. First, we note that Eq. (21) has multiple solutions, and only one local maximum corresponds to the Pareto optimal outcome. On the other hand, the optimization problem in Eq. (13) is feasible and has a unique solution, but might lead to a suboptimal attack.

Fig. 5 shows the utility of an attacker as a function of λ. With $\lambda = 1$ the attacker does not have any impact on the system, and the system is in the Pareto optimal outcome. As λ increases slightly, the utility of the attacker decreases as a consequence of incentives, that can be seen as taxes; however, the attack is profitable for larger values of λ, converging to the value of the infeasible problem in Eq. (10).

The attack has negative impact on the population, which is forced to reduce its consumption and consequently its utility (see Fig. 6).

A drawback of the dynamic pricing scheme is that it does not satisfy the budget balance property [22]. In other words,

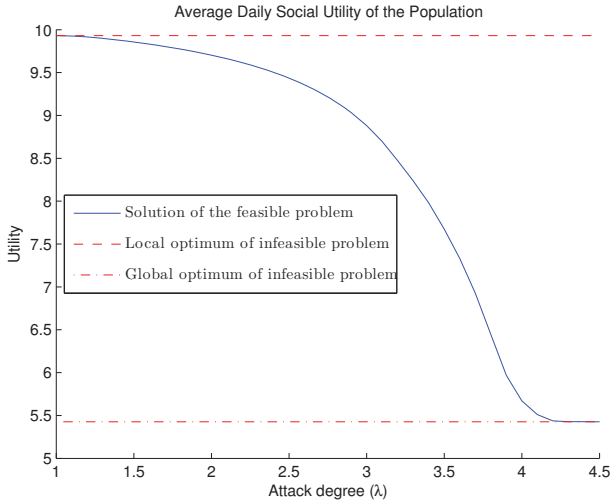

Figure 6: Impact of the attack in the social welfare utility in function of the parameter λ.

the scheme can be implemented if there is a source of external subsidies to fund part of the incentives. Note that an homogeneous population at the Pareto optimal outcome does not require external subsidies; however, if an attacker disrupts the equilibrium, the agents with low and large consumption receive positive and negative rewards, respectively. As the attack increases (λ), the demand imbalance is higher and consequently, the social planer has to increase the magnitude of the incentives. However, the taxes imposed to the attacker (due to its large consumption) are not enough to sustain the rewards to the victims and the amount of external subsidies tend to increase.

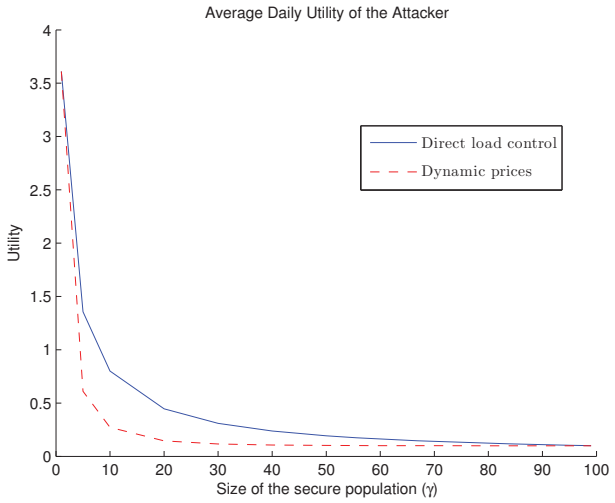

Figure 7: Fraudsters obtain more benefits from attacking DLC systems when compared to dynamic pricing.

We now compare in more detail the impact of fraudster attacks to DLC and dynamic pricing schemes. Fig. 7 and 8 show the impact of the proportion of agents that benefit from the attack γ (the ones that mask the identity of the attacker) in the utility of the attacker and the population, respectively. Note that the attacker earns less benefits as the size of the benefited population grows; however, it is

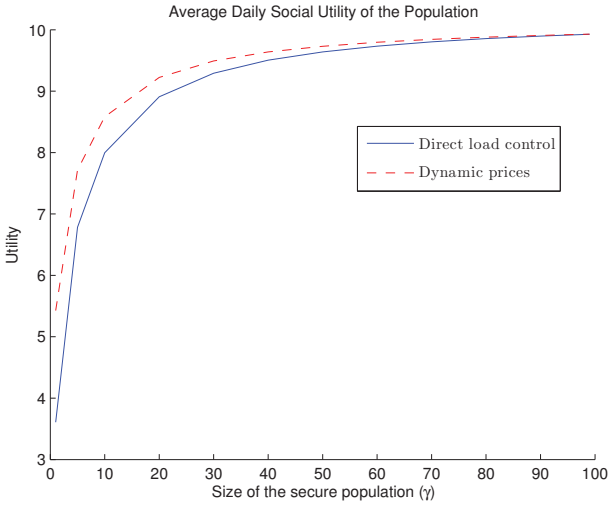

Figure 8: Consumers are better off by using dynamic pricing instead of DLC systems in the case of fraudster attacks.

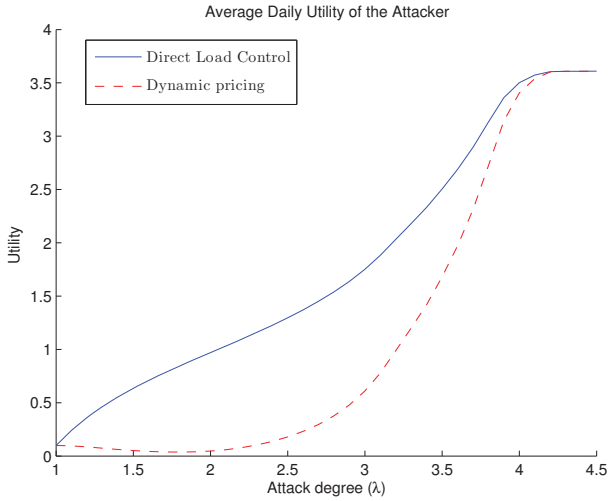

Figure 9: Fraudsters obtain more benefits from attacking DLC systems when compared to dynamic pricing.

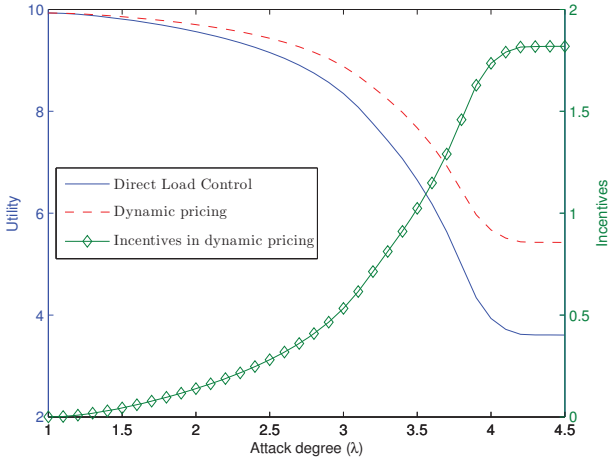

Figure 10: Impact of the attack in the social welfare utility and global incentives as a function of the attack severity λ for both the DLC and dynamic pricing schemes with $\gamma = 0.01$.

important to note that an attacker obtains less utility in a system with dynamic prices when compared to a DLC model.

In contrast, the attack causes lower losses to all consumers with dynamic pricing because the victims receive rewards for their low consumption, as seen in Fig. 8. Furthermore, as γ tends to 1, the population welfare increases because more agents obtain benefits of the attack. Thus, with $\gamma = 1$ the population has no victims and the outcome is optimal in the sense of Pareto.

Similar behavior can be seen when we evaluate the severity of the attack λ in terms of the utility gained by the attacker in Fig. 9 and the average utility of the population in Fig. 10.

Fig. 9 shows that the attack is not optimal for a system with dynamic prices, and consequently, the attacker might do worse that the Pareto optimal outcome for $\lambda < 2.3$. This happens because the attacker is penalized due to its large consumption; however, this penalization is reduced as the other consumers are at their minimum consumption when λ is large enough.

Fig. 10 shows the social welfare under an attack in both DLC and dynamic pricing systems. The attack is less detrimental with a dynamic pricing mechanism because the social planner provides incentives to consumers that reduce their consumption. Fig. 10 also shows the total amount of incentives granted to the population; in particular, the rewards and taxes are not in balance, and consequently, the system requires external subsidies to sustain the DR program. Therefore, these subsidies can be seen as losses to the social planner caused by the attack.

We can conclude that dynamic pricing schemes are more resilient than DLC models because consumers do not have to bear all the looses; the majority of the losses are taken by the social planner (electric utility of DR company). Therefore, in terms of securing the systems, we can say that dynamic pricing schemes are incentive-compatible by placing the burden of the losses to the entity that knows better how to secure these systems, rather than placing the burden of the attack on consumers.

5. MALICIOUS ATTACKER

We now consider an attacker whose goal is to cause damage to the power system, and in order to do that it tries to create a sudden (unanticipated) spike in electricity consumption. Even if the distribution system has protection mechanisms, a sudden spike that has not been forecasted will be very difficult to protect against, and it might cause local blackouts (by tripping a distribution fuse or circuit breaker) or kickstarting blackstart generators.

Creating this consequence with a DLC model is straightforward, since in a DLC model the central entity controls electricity consumption signals, the attacker can directly send all electricity consumption signals to their maximum value at the same time instant. Unless consumers somehow override the control commands to increase electricity consumption, this will have a sudden and clear impact.

For the dynamic pricing model the attack strategy is not that easy, as agents make their optimization in an individual, distributed, and rational way.

Following the basic idea behind the attack to DLC models, the most intuitive (but **naive**) attack against dynamic pricing is to select a very low price for electricity (give rewards

via the incentive price signals) for electricity consumption at the time where the highest forecast of demand is set to happen (typically in the evening). Let us called the identified time for the sudden electricity consumption spike t_{attack}.

We assume that the attacker is able to compromise the incentive signal and send the following malicious incentives:

$$I_i^m(\boldsymbol{q}) = \begin{cases} I_i(\boldsymbol{q}^t) + \sigma_1 ||\boldsymbol{q}||_1 & \text{if } t = t_{attack}, \\ I_i(\boldsymbol{q}^t) & \text{otherwise,} \end{cases}$$

where $\sigma_1 > 0$ is a design parameter. This attack creates a sudden increase in the incentives. As a consequence, the consumption of the population will increase at t_{attack}. Again, because this attack takes place in the distributed optimization setting, note that the attack can be implemented without information of the valuation functions by each consumer.

We now show how the previous intuitive (but naive) attack, is not optimal for dynamic pricing schemes. In particular, dynamic pricing models consumers that can defer the consumption of electricity when market conditions are unfavorable but then return to consuming electricity en masse when market conditions change abruptly (in this case this sudden shift in market conditions is created by the malicious attacker).

A strategic attack can be implemented by carefully increasing the prices of electricity in hours previous to the attack (so consumers start deferring electricity consumption to a later time), and then, at the time where the attacker wants to create the spike of consumption, immediately lowering the prices of electricity to their lowest values. This attack causes loads that can be deferred to shift their consumption to a later time, accumulating the need to use electricity until the price is favorable.

The attack is implemented with the following incentives:

$$I_i^m(\boldsymbol{q}) = \begin{cases} I_i(\boldsymbol{q}^t) + \sigma_1 ||\boldsymbol{q}||_1 & \text{if } t = t_{attack}, \\ I_i(\boldsymbol{q}^t) - \sigma_2 ||\boldsymbol{q}||_1 & \text{if } t \in [t_a, t_b], \\ I_i(\boldsymbol{q}^t) & \text{otherwise,} \end{cases}$$

where σ_i are positive real numbers, $[t_a, t_b]$ is the time period in which the attack focuses on reducing the demand, and t_{attack} is the time at which the peak is caused, with $t_a < t_b < t_{attack}$.

In this case, the fitness function (marginal utility) used for population dynamics (to model the transient evolution of the behavior of market participants) of the i^{th} individual is:

$$f_i^t(\boldsymbol{q}) = \begin{cases} \frac{\partial}{\partial q_i^t}\left(\sum_{h=1}^N W_h(\boldsymbol{q})\right) + \sigma_1, & \text{if } t = t_{attack}, \\ \frac{\partial}{\partial q_i^t}\left(\sum_{h=1}^N W_h(\boldsymbol{q})\right) - \sigma_2, & \text{if } t \in [t_a, t_b], \\ \frac{\partial}{\partial q_i^t}\left(\sum_{h=1}^N W_h(\boldsymbol{q})\right), & \text{otherwise,} \end{cases}$$

Note that the parameters σ_1 and σ_2 either increase or decrease the marginal utility associated with a consumption at different times of the day. Thus, the attacker can deceive users by making them believe that it is not convenient to allocate resources (consume electricity) during $[t_a, t_b]$. Also, the attacker can increase the consumption at t_{attack} by spreading the belief that users can increase their utility by using more resources at that time.

Simulations are made with $\sigma_1 = 50$, $\sigma_2 = 100$, $t_a = 0$hrs, $t_b = 17$hrs, $t_{attack} = 20$hrs. In particular, the attack time coincides with the demand peak in the Pareto optimal outcome.

Fig. 11 shows the impact of both the naive and the strategic attack after the attack is launched, that is, during the initial transition period. In particular, the naive attack succeeds in causing an increase of the demand at t_{attack}. However, the main impact of the attack is produced only during

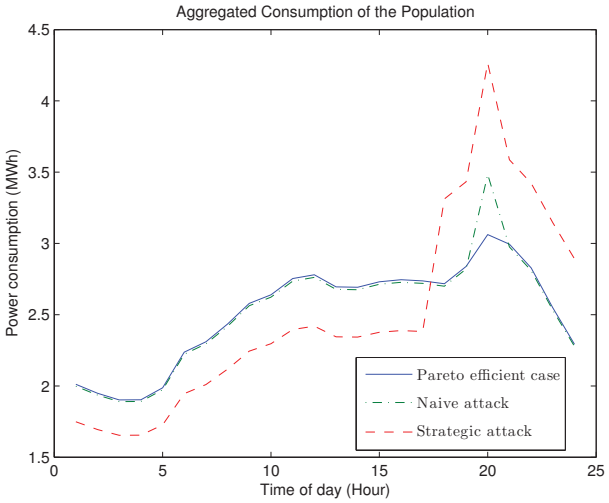

Figure 11: Impact of a malicious attack on the populaiton demand for two different attacks 1) attack on a single hour and 2) coordinated attack on various hours of the day.

t_{attack}. On the other hand, the strategic attack achieves a greater peak by causing demand reduction prior to the attack. Roughly speaking, the strategic attack sets conditions so that the population has more resources to consume when prices are low.

6. CONCLUSIONS AND FUTURE WORK

We have introduced two new attack models for DR programs. In contrast to previous work, our attackers are strategic adversaries with clearly defined goals (objective functions), furthermore we proved the optimality of the fraud strategies for attacks against DR with DLC, and proved an optimality property of the fraudster attack against DR with dynamic pricing. We also showed how dynamic pricing is more resilient to attacks than DLC mechanisms.

In addition we introduced population dynamics to model the transient behavior of DR systems and in particular, created a model with consumers that can defer the consumption of electricity based on incentives. We showed how this behavior can be exploited by malicious attackers to cause a sudden spike in electricity load. We introduced two threat scenarios and showed that a strategic attack can perform better than the initial intuitive attacks. In practice we believe that this is the worst possible attack to the power distribution system as utilities will not be able to forecast (and therefore plan contingencies to) these changes orchestrated by attackers.

The main goal of this work was to understand the vulnerability of DR systems to market manipulators with access to the control signals sent by the central authority to consumers. In future work we plan to study mechanisms to improve the security of these systems and to minimize the effects of attacks.

One interesting problem would be to design anomaly detection schemes to detect fraudsters (depending on the parameters γ and λ), and then evaluate them against fraudsters that will try to maximize their gains while avoiding detection.

While our DR models are an improvement over previous work studying the security of DR systems; the accuracy of our analysis for practical applications will still depend on how well our DR models match real deployed systems. As

we continue to deploy trial DR systems around the world we will obtain more data on several properties of the models such as the average elasticity of electricity consumption and how much can incentives control the overall electricity load.

Acknowledgments

This work was supported in part by NIST award 70NANB-14H236 from the U.S. Department of Commerce.

7. REFERENCES

[1] Alexander, M., Agnew, K., Goldberg, M.: New approaches to residential direct load control in california. In: 2008 ACEEE Summer Study on Energy Efficiency in Buildings (2008)

[2] Arslan, G., Marden, J.R., Shamma, J.S.: Autonomous vehicle-target assignment: A game-theoretical formulation. Journal of Dynamic Systems Measurement and Control 129(5), 584 (2007)

[3] Barreto, C., Mojica-Nava, E., Quijano, N.: Design of mechanisms for demand response programs. In: Proceedings of the 2013 IEEE 52nd Annual Conference on Decision and Control (CDC). pp. 1828–1833 (2013)

[4] Barreto, C., Mojica-Nava, E., Quijano, N.: Incentives-based mechanism for efficient demand response programs. arXiv preprint arXiv:1408.5366 (2014)

[5] California Energy Commission: Docket No. 13-IEP-1F: Increasing demand response capabilities in California (2013)

[6] Carlos Barreto: Population dynamics Toolbox (PDToolbox) (2014), `https://github.com/carlobar/PDToolbox_matlab`

[7] Chen, L., Li, N., Low, S.H., Doyle, J.C.: Two market models for demand response in power networks. In: Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on. pp. 397–402. IEEE (2010)

[8] Chen, L., Li, N., Low, S.H., Doyle, J.C.: Two market models for demand response in power networks. In: Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on. pp. 397–402. IEEE (2010)

[9] Fahrioglu, M., Alvarado, F.L.: Designing cost effective demand management contracts using game theory. In: Power Engineering Society 1999 Winter Meeting, IEEE. vol. 1, pp. 427–432. IEEE (1998)

[10] Federal Energy Regulatory Commission: 2011 assessment of demand response and advanced metering (November 2011)

[11] FERC: FERC Staff Issue Assessment of Demand Response and Advanced Metering (October 2013)

[12] Gellings, C.W.: The smart grid: enabling energy efficiency and demand response. The Fairmont Press, Inc. (2009)

[13] Honebein, P.C., Cammarano, R.F., Boice, C.: Building a social roadmap for the smart grid. The Electricity Journal 24(4), 78–85 (2011)

[14] Huang, L., Walrand, J., Ramchandran, K.: Optimal smart grid tariffs. In: Information Theory and Applications Workshop (ITA), 2012. pp. 212–220. IEEE (2012)

[15] Ibars, C., Navarro, M., Giupponi, L.: Distributed demand management in smart grid with a congestion game. In: Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on. pp. 495–500. IEEE (2010)

[16] Johari, R., Tsitsiklis, J.N.: A scalable network resource allocation mechanism with bounded efficiency loss. Selected Areas in Communications, IEEE Journal on 24(5), 992–999 (2006)

[17] Li, N., Chen, L., Low, S.H.: Optimal demand response based on utility maximization in power networks. In: Power and Energy Society General Meeting, 2011 IEEE. pp. 1–8. IEEE (2011)

[18] Liu, Y., Reiter, M.K., Ning, P.: False data injection attacks against state estimation in electric power grids. In: CCS '09: Proceedings of the 16th ACM conference on Computer and communications security. pp. 21–32. ACM, New York, NY, USA (2009)

[19] Liyan, J., Thomas, R.J., Tong, L.: Impacts of malicious data on real-time price of electricity market operations. In: 45th Hawaii International Conference on System Sciences. pp. pp.1907–1914 (January 2012)

[20] Mohsenian-Rad, A., Wong, V., Jatskevich, J., Schober, R., Leon-Garcia, A.: Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. IEEE Transactions on Smart Grid 1(3), 320–331 (dec 2010)

[21] Negrete-Pincetic, M., Yoshida, F., Gross, G.: Towards quantifying the impacts of cyber attacks in the competitive electricity market environment. In: 2009 IEEE PowerTech (June 2009)

[22] Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press, 32 Avenue of the Americas, New York, NY 10013-2473, USA (2007)

[23] Roozbehani, M., Rinehart, M., Dahleh, M., Mitter, S., Obradovic, D., Mangesius, H.: Analysis of competitive electricity markets under a new model of real-time retail pricing. In: Energy Market (EEM), 2011 8th International Conference on the European. pp. 250–255 (may 2011)

[24] Roozbehani, M., Dahleh, M., Mitter, S.: Dynamic Pricing and Stabilization of Supply and Demand in Modern Electric Power Grids. In: First IEEE Smart Grid Commnunications Conference (SmartGridComm) (October 2010)

[25] Roozbehani, M., Dahleh, M.A., Mitter, S.K.: Volatility of power grids under real-time pricing. Power Systems, IEEE Transactions on 27(4), 1926–1940 (2012)

[26] Samadi, P., Schober, R., Wong, V.W.: Optimal energy consumption scheduling using mechanism design for the future smart grid. In: Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on. pp. 369–374. IEEE (2011)

[27] Sandholm, W.H.: Population Games and Evolutionary Dynamics (Economic Learning and Social Evolution). The MIT Press (2011)

[28] Tan, R., Krishna, V.B., Yau, D.K., Kalbarczyk, Z.: Impact of integrity attacks on real-time pricing in smart grids. In: ACM Conference on Computer and Communications Security (CCS 2013) (2013)

[29] Trilliant: Direct load control (2014), `http://trilliantinc.com/solutions/consumer/direct-load-control`

[30] Xie, L., Mo, Y., Sinopoli, B.: False Data Injection Attacks in Electricity Markets. In: First IEEE Smart Grid Commnunications Conference (SmartGridComm) (October 2010)