A Game Theory Model for Electricity Theft Detection and Privacy-Aware Control in AMI Systems

Alvaro A. Cárdenas
University of Texas, Dallas

Saurabh Amin
MIT

Galina Schwartz, Roy Dong, Shankar Sastry
University of California, Berkeley

I. Abstract

We introduce a model for the operational costs of an electric distribution utility. The model focuses on two of the new services that are enabled by the Advanced Metering Infrastructure (AMI): (1) the fine-grained anomaly detection that is possible thanks to the frequent smart meter sampling rates (e.g., 15 minute sampling intervals of some smart meter deployments versus monthly-readings from old meters), and (2) the ability to shape the load thanks to advanced demand-response mechanisms that leverage AMI networks, such as direct-load control.

We then study two security problems in this context. (1) In the first part of the paper we formulate the problem of electricity theft detection (one of the use-cases of anomaly detection) as a game between the electric utility and the electricity thief. The goal of the electricity thief is to steal a predefined amount of electricity while minimizing the likelihood of being detected, while the electric utility wants to maximize the probability of detection and the degree of operational cost it will incur for managing this anomaly detection mechanism. (2) In the second part of the paper we formulate the problem of privacy-preserving demand response as a control theory problem, and show how to select the maximum sampling interval for smart meters in order to protect the privacy of consumers while maintaining the desired load shaping properties of demand-response programs.

II. Introduction

For most electric distribution utilities, creating a business case for improving computer security and supporting long-term security research is a difficult task because of the lack of risk models that capture the effects of security and privacy in their revenue and profit margins.

We consider the point of view of an electric distribution utility that needs to create a business case for improving their security posture by introducing an electricity-theft anomaly detection mechanism and a privacy-preserving demand response program.

We model the electricity-theft anomaly detection case as a game played between the utility and fraudulent consumers, and characterize the Nash equilibrium of the game.

In the second part of the paper we consider the privacy-preserving demand-response problem and using realistic values of a direct-load control example, we show how the peak shaving goal of the demand-response program depends on the privacy (sampling interval) of the Advanced Metering Infrastructure (AMI).

III. Background and Motivation

A. Electricity Theft

Energy theft in emerging economies has been a widespread practice. A World Bank report [1] found that up to 50% of electricity in developing countries is acquired via theft. Electricity theft can be caused by physical and cyber attacks. Physical security considerations range from defaulting on payments to directly connecting loads to the electricity distribution lines. A cyber attack against smart meters is also possible (and the focus of this paper). While some basic protective measures have been developed (tamper-evident seals, secure link communications), they are not enough to prevent successful attacks during the meter lifespan. In addition to vulnerabilities identified by security researchers [2], [3]—some of them allowing rogue remote firmware updates [4]—hacked smart meters have been used to steal electricity, costing a single U.S. electric utility hundreds of millions of dollars annually, as reported by a cyber-intelligence bulletin issued by the FBI [5]. The FBI report warns that insiders and individuals with only a moderate level of computer knowledge are likely able to compromise and reprogram meters with low-cost tools and software readily available on the Internet. The FBI report also assesses with medium confidence that as smart grid use continues to spread throughout the country, this type of fraud will also spread because of the ease of intrusion and the economic benefit to both the hacker and the electric customer.

Detecting electricity theft has traditionally been addressed by physical checks of tamper-evident seals by field personnel and by using balance meters [6]. While valuable, these techniques alone are not enough. Tamper evident seals can be easily defeated [7] and balance

meters can detect that some of the customers connected to it are misbehaving, but cannot identify exactly who they are. Despite the vulnerabilities of smart meters, the high-resolution data they collect is seen as a promising technology to improve electricity-theft detection. In general, utilities are gathering more data from many devices and they are leveraging *big data analytics* [8] to obtain better situational awareness of the health of their system. One of the key services offered by Meter Data Management (MDM) vendors for turning big data into actionable information is called *revenue assurance*, where data analytics software is used by the utility on the collected meter data to identify possible electricity theft situations and abnormal consumption trends [9]. Big data analytics is thus a new cost-effective way to complement the use of balance meters (which are still necessary to detect when electricity thieves connect directly to the power distribution lines instead of tampering with the meter) and physical personnel checking for tamper-evident seals.

In this paper we model of a utility using statistical anomaly detection in smart meter readings to identify potential electricity theft. Our work creates a game-theoretic formalism of recent research efforts in electricity theft detection [10], [11].

B. Privacy

Smart meters allow large-scale data collection, making individual household data available at unprecedented levels of granularity. Monitoring energy consumption at high granularity can allow the inference of detailed information about consumers' lives. Such behavioral data is highly valuable to advertising companies, law enforcement, and criminals. Hence, there is potential for erosion of individual privacy in the development of the smart grid, and we must ensure proper controls are in place.

Previous research has tried to mitigate these privacy concerns by power-mixing [12], [13], data aggregation [14], [15], and cryptographic techniques [16], [17]. While these approaches are promising, they do not address the privacy-by-design principle of data minimization; i.e., what is the minimum data collection frequency that still allows the utilities to efficiently perform advanced smart grid operations, including load management and demand-response?

In this paper we concentrate on the best principles for data collection of energy-use data. In particular, we formulate the problem as a discrete-time control sampling problem, and show what properties we need to study from this sampled system in order to maintain a satisfactory level of demand-response functionalities.

IV. A Model of Operational Costs and Profits for Distribution Utilities

We consider a cost model of a regional distribution utility that considers two important factors related to their profits.

First, we model the non-technical losses due to theft by a subset θ (from the total set of customers Θ) of consumers stealing electricity. To reduce the losses due to electricity theft the utility can invest in anti-fraud technologies and recover a part of the the electricity stolen by imposing fines on the consumers it has identified as committing fraud.

Second, we consider the cost the regional utility needs to pay to their provider of electricity in order to satisfy the demands of their consumers. To manage this cost the utility can deploy demand-response mechanisms.

In this paper we study how security and privacy affect the electricity distribution costs by using the following models:

1) To deal with electricity theft, we consider a game played between set θ of independent consumers stealing electricity and the electric utility. We find a Nash equilibrium of the game.
2) To manage the costs necessary to supply the demand of their consumers, we consider a direct load control demand response deployment. We formulate the *privacy-preserving demand-response problem* as the task of finding the maximum allowable sampling rate that keeps the demand lower than a predefined maximum value.

The privacy-preserving DR is a design consideration which can be imposed by the Government or a regulator on the electric distribution utility. Once the AMI sampling scheme is in place, our game-theoretic model permits us to consider strategic consumers who are interested in stealing electricity.

We assume an electric distribution utility who has an AMI deployment collecting a time-series of electric power consumption y_k^i for every time step k and every customer $i \in \Theta$.

Let q^i denote the expected total consumption of user i, and $q_U{}^i$ denote the expected unbilled part of the consumption of user i. Note that $q_U{}^i = 0$ for honest users $i \in \Theta - \theta$. Thus an electricity thief sends a signal y_k back to the utility that does not represent their true consumption.

We assume the distribution utility has three design variables: (1) the effort e invested in anti-fraud technologies, (2) the anomaly detection test $\mathcal{D}$ used to identify electricity theft, and (3) the sampling interval N (the time interval between measurement y_k and y_{k+N} taken by their smart meter deployments).

For a fixed N, the revenue of a distribution utility is the sum of the tariffs T from all customers plus the recovered fines F^r from the detected electricity theft:

$$R(e, \mathcal{D}) = \sum_{i \in \Theta} T(q^i - q_U{}^i) + \sum_{i \in \theta} \rho(e, q_U{}^i, \mathcal{D}) F^r(q_U{}^i), \quad (1)$$

where ρ represents the probability of detecting an electricity thief.

There are two main costs to the electric utility. The first is the investment in protecting their infrastructure against electricity theft $\psi(\mathrm{e})$, and the second one relates to the costs associated with meeting demand of all the consumers $\sum_k Y_k$, where $Y_k = \sum_{i\in\Theta} y_k^i + \mathrm{q_U}^i$ (i.e., Y_k is the total demand at time k including the unbilled demand $\mathrm{q_U}^i$).

The profit of the utility is thus:

$$\mathrm{R}(\mathrm{e}, \mathcal{D}) - \mathrm{C}(\{Y_k\}) - \psi(\mathrm{e}). \tag{2}$$

In the next section we study the terms $\mathrm{R}(\mathrm{e}, \mathcal{D})$ and $\psi(\mathrm{e})$ with a game theoretic model of electricity theft detection, and in Section VI we focus on the middle term of the equation $\mathrm{C}(\{Y_k\})$, formulating the problem of maximizing the privacy the utility provides to its consumers while keeping the same cost $\mathrm{C}(\{Y_k\})$ as other more privacy invasive AMI sampling rates.

V. A Game Theoretic Model of Electricity Theft Detection

There are many ways the utility can invest in protecting their infrastructure. They can:

1) invest in a centralized meter data management (MDM) solution that performs analysis of the time series received by consumers and comparing them to historical trends and correlate them to other customers in similar residences or businesses.
2) invest in increasing redundancy (balance meters) by adding redundant meters at different parts of their infrastructure.
3) invest in hardening the smart meters, by adding better tamper-resistant solutions, and embedded sensors in the meter that report reprogramming or tampering attempts.

In this paper we focus on the centralized MDM solution for many reasons: (1) it is the main focus of many AMI deployments, (2) it does not require the capital investments of the other technologies, (3) the operational cost of managing a meter data management solution fits better with the model we are going to introduce for the distributor in the next section, and (4) it is the only solution that can be retrofitted to an existing AMI deployment.

In the MDM security model we assume the distribution utility has an anomaly detection mechanism D that tries to identify if the received electricity consumption signal y is fraudulent or not.

We assume the period of study is from $k = 1$ to $k = n$, thus with the advanced metering infrastructure, the electric utility is able to collect for each user a vector $y^i \in \mathbb{R}^n$ (note that y_k can be negative, thus modeling consumers who can give electricity back to the grid by e.g., installing solar panels).

We assume that electricity thieves have compromised smart meters and can thus send falsified meter measurements y_k. (Attackers that steal electricity by connecting directly to the distribution lines are outside the scope of this paper.) We assume each measurement y_k^i is the result of a random process driven by a probability density function $f_1^i(y^i)$. An honest user will have a probability density $f_0^i(y^i)$ different from the density of an attacker $f_1^i(y)$ and satisfying the following constraints:

$$\mathbf{1}^T \mathbb{E}_0^i[Y] = \mathrm{q}^i \quad \text{and} \quad \mathbf{1}^T \mathbb{E}_1^i[Y] \leq \mathrm{q}^i - \mathrm{q_U}^i \tag{3}$$

We assume that the utility knows f_0^i for each user (e.g., by historical profiles the utility can estimate the normal electricity consumption distribution f_0^i before the user compromises and reprograms the smart meter to start giving fake signals).

We assume consumers have already established an average consumption pattern q^i (a fixed value) and that attackers have a minimum amount of electricity it wants to steal $\mathrm{q_U}^i$.

For notational simplicity we drop in the following analysis the superscript i denoting the individual consumer from $f_0, f_1, y, y_k, \mathrm{q_U}$, and q. However, we note that all the results are valid for the general case when each consumer i is different from the other consumers as we will show at the end of this section.

From Eq. (2) we observe that **the goal of the utility is given by**:

$$\max_{\mathrm{e}\geq 0, \mathcal{D}} \sum_{\Theta} \mathrm{T}(\mathrm{q} - \mathrm{q_U}) + \sum_{\theta} \rho(\mathrm{e}, \mathrm{q_U}, \mathcal{D}) \mathrm{F^r}(\mathrm{q_U}) - \psi(\mathrm{e}). \tag{4}$$

In this paper we assume the tariff T is given by a regulator and is not controlled by the distribution utility, therefore the parameters that the utility can control affect only the following terms:

$$\max_{\mathrm{e}\geq 0, \mathcal{D}} \sum_{\theta} \rho(\mathrm{e}, \mathrm{q_U}, \mathcal{D}) \mathrm{F^r}(\mathrm{q_U}) - \psi(\mathrm{e}). \tag{5}$$

The operational cost ψ of managing the anomaly detector $\mathcal{D}$ is quantified by the resources (effort e) the distribution utility assigns for dealing with false alarms (e.g., the number of analysts and field engineers responding to false events). The probability of detecting fraud increases with the effort e dedicated by the utility in anti-fraud mechanisms, and with the amount of electricity stolen. As we will show later in this section, it also depends on the density function f_1 the attacker uses. $\mathrm{q_U}$ and can be modeled as

$$\rho : \mathbb{R}_+ \times \mathbb{R} \times \mathcal{A}_{\mathrm{q_U}} \to [0, 1] \tag{6}$$

where

$$\mathcal{A}_{\mathrm{q_U}} = \left\{ f_1 : \mathbf{1}^T \mathbb{E}_1[Y] \leq \mathrm{q} - \mathrm{q_U} \right\}. \tag{7}$$

Thus ρ assigns for to each investment level e, stolen electricity $\mathrm{q_U}$, and pdf f_1, a probability of detection.

Given this problem definition, the optimal anomaly detection test $\mathcal{D}$ is the one that maximizes the probability of detecting a fraudster ρ subject to an upper bound on the false alarm rate e (the investment).

From Neyman-Pearson theory we know that the optimal detection test $\mathcal{D}$ (the test that maximizes the probability of detection given a constraint in the number of false alarms) is the likelihood-ratio test:

$$\mathcal{D}(y) = \ln \frac{f_1(y)}{f_0(y)} \underset{H_0}{\overset{H_1}{\gtrless}} \tau \tag{8}$$

Thus

$$\rho(\mathrm{e}, \mathrm{q_U}, \mathcal{D}) = P_1\left[\ln \frac{f_1(y)}{f_0(y)} > \tau\right] + \gamma P_1\left[\ln \frac{f_1(y)}{f_0(y)} = \tau\right] \tag{9}$$

where τ and γ are selected such that

$$P_0\left[\ln \frac{f_1(y)}{f_0(y)} > \tau\right] + \gamma P_0\left[\ln \frac{f_1(y)}{f_0(y)} = \tau\right] = \mathrm{e} \tag{10}$$

We note however that we do not know f_1, as it is selected by an attacker.

We assume the attacker knows the anomaly detection test used by the utility company is a likelihood-ratio test, but we assume the attacker does not know e (i.e. the attacker does not know during operation the threshold τ and randomization γ used by the utility, which can be selected and changed online depending on the availability of analysts investigating alarm reports). Therefore we assume the attacker wants to minimize the likelihood ratio value.

In other words, **The goal of the attacker** is to find $f_2 \in \mathcal{A}_{\mathrm{q_U}}$ (i.e., an f_2 that satisfies the constraint on the amount of electricity stolen $\mathrm{q_U}$) while minimizing the expected likelihood ratio function:

$$\min_{f_2 \in \mathcal{A}_{\mathrm{q_U}}} \mathbb{E}_2\left[\ln \frac{f_1(y)}{f_0(y)}\right] \tag{11}$$

$$= \min_{f_2 \in \mathcal{A}_{\mathrm{q_U}}} \int f_2(y) \ln \frac{f_1(y)}{f_0(y)} dy \tag{12}$$

Note that $f_1(y)$ is chosen by the defender as part of its likelihood ratio test $\mathcal{D}(y)$, while $f_2(y)$ is chosen by the attacker.

We will prove later in the paper that the solution f_1^* to the following equation is such that the optimal move for the defender is to choose $f_1 = f_1^*$ and the optimal move for the attacker is to choose $f_2 = f_1^*$.

$$\min_{f_1 \in \mathcal{A}_{\mathrm{q_U}}} \int f_1(y) \ln \frac{f_1(y)}{f_0(y)} dy \tag{13}$$

Notice first that the objective function is convex in f_1. We let $q^\epsilon(y) = f_1^*(y) + \epsilon h(y)$ and construct the Lagrangian of the objective function and the constraints

$$\int q^\epsilon(y) \ln \frac{q^\epsilon(y)}{f_0(y)} dy + \mu_1 \left(\int q^\epsilon(y) dy - 1\right) + \mu_2 \left(\mathbf{1}^T \int y q^\epsilon(y) dy - (\mathrm{q} - \mathrm{q_U})\right) \tag{14}$$

By taking the derivative with respect to ϵ and equating this quantity to zero for all possible $h(y)$, we find that the optimal f_1^* has to be of the form:

$$f_1^*(y) = f_0(y) e^{-\mu_2 y - \mu_0} \tag{15}$$

where $\mu_0 = \mu_1 + 1$. In order to obtain the values of the Lagrange multipliers μ_0 and μ_2 we use the constraints of f_1. The first constraint states that f_1^* must be a pdf and therefore

$$\int f_0(y) e^{-\mu_2 y - \mu_0} dy = 1 \tag{16}$$

solving for μ_0 we have

$$\mu_0 = \ln \int f_0(y) e^{-\mu_2 y} dy \tag{17}$$

Replacing this solution in Eq. (15) we get

$$f_1^*(y) = \frac{f_0(y) e^{-\mu_2 y}}{\int f_0(y) e^{-\mu_2 y} dy} \tag{18}$$

The second constraint in $\mathcal{A}_{\mathrm{q_U}}$ is rewritten in terms of Eq.(18) as

$$\mathbf{1}^T \int y \frac{f_0(y) e^{-\mu_2 y}}{\int f_0(y) e^{-\mu_2 y} dy} dy = \mathrm{q} - \mathrm{q_U} \tag{19}$$

from where we can obtain μ_2 (once we know f_0).

We now show that f_1^* is a *Nash equilibrium* between the attacker and the defender. First, we show that an attacker has no incentive to deviate from f_1^* if the defender selects f_1^* for its likelihood ratio test:

Assume an attacker selects a pdf

$$f_2 \in \mathcal{A}_{\mathrm{q_U}} \tag{20}$$

then, the expected value of the likelihood ratio test under f_2 is:

$$\begin{aligned}
\mathbb{E}_2\left[\ln \frac{f_1^*(y)}{f_0(y)}\right] &= \int \ln \frac{f_1^*(y)}{f_0(y)} f_2(y) dy \\
&= \int f_2(y) \ln \frac{e^{-\mu_2 y}}{\int f_0(y) e^{-\mu_2 y} dy} dy \\
&= \int f_2(y) \ln e^{-\mu_2 y} - \ln \int f_0(y) e^{-\mu_2 y} dy \int f_2(y) dy \\
&= \int (-\mu_2 y) f_2(y) dy - \ln \int f_0(y) e^{-\mu_2 y} dy \\
&\geq -\mu_2(\mathrm{q} - \mathrm{q_U}) - \ln \int f_0(y) e^{-\mu_2 y} dy \\
&= \int -\mu_2 y f_1^*(y) dy - \int f_1^*(y) dy \ln \int f_0(y) e^{-\mu_2 y} dy \\
&= \int f_1^*(y) \ln \frac{e^{-\mu_2 y}}{\int f_0(y) e^{-\mu_2 y} dy} dy \\
&= \int f_1^*(y) \ln \frac{f_1^*(y)}{f_0(y)} dy
\end{aligned} \tag{21}$$

Similarly, the defender has no incentive on selecting a pdf f_3 different from f_1^* if the attacker selects f_1^*. The

proof of this statement is a direct result of the Neyman-Pearson lemma, which in our case implies that the test that maximizes the probability of detection ρ for any fixed false alarm rate is the likelihood ratio test with with f_1^* as the alternate hypothesis and f_0 as the null hypothesis.

Theorem 1. *Let*

$$\mathcal{D}^{i,*}(y) = \ln \frac{f_1^{i,*}(y)}{f_0^i(y)} \underset{H_0}{\overset{H_1}{\gtrless}} \tau \tag{22}$$

then $f_1^{i,*}$ is a Nash equilibrium between a distribution utility with the following objective:

$$\max_{\mathcal{D}^i} \sum_{i\in\theta} \rho(\mathrm{e}, \mathrm{q_U}^i, \mathcal{D}^i)\mathrm{F^r}(\mathrm{q_U}^i) - \psi(\mathrm{e}). \tag{23}$$

and players $i \in \theta$ with the following objective:

$$\min_{f_2^i \in \mathcal{A}^i_{\mathrm{q_U}}} \int f_2^i(y) \ln \frac{f_1^{*,i}(y)}{f_0^i(y)} dy \tag{24}$$

The final part of the optimization problem for the utility company is the selection of e.

$$\max_{1\geq \mathrm{e} \geq 0} \sum_{i\in\theta} \rho(\mathrm{e}, \mathrm{q_U}^i, \mathcal{D}^{i,*})\mathrm{F^r}(\mathrm{q_U}^i) - \psi(\mathrm{e}). \tag{25}$$

Assuming ψ is a linear function and ρ is differentiable with e we obtain the following first order condition:

$$\sum_{i\in\theta} \partial_{\mathrm{e}}\rho(\mathrm{e}, \mathrm{q_U}^i, \mathcal{D}^{i,*})\mathrm{F^r}(\mathrm{q_U}^i) = \psi. \tag{26}$$

To understand the interpretation of this result, assume all customers steal the same amount of electricity $\mathrm{q_U}$. Then Eq. (26) simplifies to

$$\sum_{i\in\theta} \partial_{\mathrm{e}}\rho(\mathrm{e}, \mathrm{q_U}, \mathcal{D}^{i,*}) = \frac{\psi}{\mathrm{F^r}(\mathrm{q_U})} \tag{27}$$

Now notice how $\rho(\mathrm{e}, \mathrm{q_U}, \mathcal{D}^i)$ is the Receiver Operating Characteristic (ROC) curve of $\mathcal{D}^i$, therefore the above equation simply means that the optimal false alarm rate e can be identified as the place in the ROC curve where the sum of the slopes is equal to $\frac{\psi}{\mathrm{F^r}(\mathrm{q_U})}$. Since the ROC is continuous and the slope starts at ∞ and goes to 0 as e grows, there exists such a point satisfying the first order condition.

VI. Privacy-Preserving Demand-Response

In the last section we focused on the first and last terms of Eq. (2). In this section we investigate the middle term $\mathrm{C}(\{Y_k\})$ by formulating the problem of maximizing the privacy the utility provides to its consumers while keeping the same cost $\mathrm{C}(\{Y_k\})$ as other more privacy invasive AMI sampling rates.

To understand how the cost $\mathrm{C}(\{Y_k\})$ is related to demand-response, we summarize a cost function we have encountered with some of the distribution utilities we have talked with. A large portion of distributors in the U.S. buy the energy from larger regional transmission utilities. Their cost function C depends not only on the amount of power required to meet the demand of all their consumers $\mathrm{Q} = \sum Y_k$, but also on the time-properties of the demand $Y_k, Y_{k+1}, \ldots$. A common cost function (e.g., for a monthly period) is the following:

$$\mathrm{C}(Y) = r\mathrm{Q} + p \max_k Y_k \tag{28}$$

where usually $p \gg r$. The first part of the equation represents the total amount of energy bought during the period at rate r (this price might change over the month, but for simplicity we assume it is fixed), while the latter part represents the maximum amount of power that was required by the utility in the one-month period. This latter part takes into consideration the costs for capacity planning, to make sure providers have enough resources to supply the maximum demand of the distribution utility.

This cost structure is also a major incentive to implement demand-response programs targeting the reduction of $\max_k Y_k$. In this section we study how privacy-preserving sampling of smart meter users y_k impacts demand-response programs, and provide a set of metrics to study in order to find the optimal privacy-preserving sampling that keeps the maximum demand $\max_k Y_k$ low with a high probability.

For simplicity we assume that the set of attackers is empty ($\theta = \emptyset$), leaving the joint problem of electricity theft and its impact in demand response for future work.

A. Individual user models

There are many (envisioned or deployed) demand response programs aimed to shaping the load Y_k and lowering the peaks at the request of the utility company. Most of the current deployments are based on messages sent by the electric utility either by phone or email to the energy administrator of a facility or a house owner, informing of an incoming event, and asking them to lower their electricity consumption during a certain period of time in return or reduced electricity bills.

This method is inefficient and therefore there are many ongoing programs trying to create novel demand-response programs that can reach a larger set of customers and achieve better control of the load. One popular case study is the use of real-time price incentives delivered by the smart meter to automatic appliances in the user home that respond to these price signals according to some preference of the user. Another common demand-response program is direct-load control [18], where the utility or demand-response provider controls a load in a consumer premise (typically a thermostat within some predefined bounds).

In this section we consider a simple direct-load control example. For an individual user i, we model her power

consumption dynamics with the following model:

$$(y^i_{k+1} - \overline{y}^i_{k+1}) = \alpha^i(y^i_k - \overline{y}^i_k) + \beta^i u^i_k + w^i_k \tag{29}$$

where $y^i_k \in \mathbb{R}$ represents the power consumption of user i at time k, and $u^i_k \in \mathbb{R}$ represents the control signal given to user i at time k. Here, we assume that α^i, β^i are given and $\overline{y}^i_k$ is given for all k. Furthermore, we assume that w^i_k is normally distributed with zero mean and variance $\sigma^2_{d,i}$. The w^i_k are mutually independent across time k. Also, we assume we can measure x^i_k directly.

We provide some justification for such a model, as well as some physical intuition for the significance of the parameters α^i and β^i. First, assume that w^i_k and u^i_k are simply zero for all k. Then, we can see that the trajectory will simply be $y^i_k = \overline{y}^i_k$; thus, $\overline{y}^i_k$ represents the uncontrolled trajectory of power consumption. Of course, people's behavior is not deterministic, so the w^i_k term models some of the uncertainty in human behavior. Now, suppose some external disturbance perturbs the user's power consumption away from this default preference; if $|\alpha^i| < 1$, then this perturbation will eventually die out and the user will resume her previous power consumption patterns. The α^i parameter models the sensitivity of the user's preferences to perturbations.

On the other hand, β^i represents the efficacy of control. For example, the input signal, u^i_k, which modifies power consumption could be the price of electricity. In such a scenario, β^i would represent the user's price elasticity of demand near the operating point. Another possible signal could be direct-load control, either through some Advanced Metering Infrastructure (AMI) or even, as in current practice, phone calls to commercial plants. Yet another possibility is thermostatically controlled loads where the input is a thermostat setpoint.

We also note that, defining $\mu^i_k = \overline{y}^i_{k+1} - \alpha^i \overline{y}^i_k$, this model is equivalent to:

$$y^i_{k+1} = \alpha^i y^i_k + \beta^i u^i_k + d^i_k \tag{30}$$

where d^i_k is normally distributed with mean μ^i_k and variance $\sigma^2_{d,i}$. Additionally, by recursion, we can see that:

$$y^i_{k+N} = (\alpha^i)^N y^i_k + \sum_{j=0}^{N-1} (\alpha^i)^{(N-1)-j} \beta^i u^i_{k+j} + \sum_{j=0}^{N-1} (\alpha^i)^{(N-1)-j} d^i_{k+j}. \tag{31}$$

We note that, given x^i_k and u^i_j for $j \in \{k, k+1, \ldots, k+(N-1)\}$, that y^i_{k+N} is a normally distributed random variable with mean:

$$(\alpha^i)^N y^i_k + \sum_{j=0}^{N-1} (\alpha^i)^{(N-1)-j} \beta^i u^i_{k+j} + \sum_{j=0}^{N-1} (\alpha^i)^{(N-1)-j} \mu^i_{k+j} \tag{32}$$

and variance: $\sum_{j=0}^{N-1} ((\alpha^i)^{(N-1)-j})^2 \sigma^2_{d,i}$.

B. Aggregated model

Now, suppose we have n users whose dynamics are de-coupled and can be modeled by Eq. (30). As done in previous section we define the state to be the vector containing the power consumption of all n users: $y_k = (y^1_k, y^2_k, \ldots, y^n_k)$. This yields the following dynamics:

$$y_{k+1} = Ay_k + Bu_k + d_k \tag{33}$$

where $A = \text{diag}(\alpha^1, \alpha^2, \ldots, \alpha^n)$, $B = \text{diag}(\beta^1, \beta^2, \ldots, \beta^n)$ are known, and d_k follows a multivariate normal distribution with known mean $\mu_k = (\mu^1_k, \mu^2_k, \ldots, \mu^n_k)$ and known variance $\Sigma_d = \text{diag}(\sigma^2_{d,1}, \sigma^2_{d,2}, \ldots, \sigma^2_{d,n})$. The d_k are mutually independent across time k.

Once again using recursion, we can see:

$$y_{k+N} = A^N y_k + \sum_{j=0}^{N-1} A^{(N-1)-j} B u_{k+j} + \sum_{j=0}^{N-1} A^{(N-1)-j} d_{k+j}. \tag{34}$$

Say y_k and u_j for $j \in \{k, k+1, \ldots, k+(N-1)\}$ are given. Then y_{k+N} is a multivariate normal random variable with mean:

$$A^N y_k + \sum_{j=0}^{N-1} A^{(N-1)-j} B u_{k+j} + \sum_{j=0}^{N-1} A^{(N-1)-j} \mu_{k+j} \tag{35}$$

and variance: $\sum_{j=0}^{N-1} A^{(N-1)-j} \Sigma_d A^{(N-1)-j}$.

With this model, we can begin to perform some analysis. Specifically, we can formulate demand-response programs as control policies. Once this is done, we can quantify the effects of sub-sampling on the performance of such control policies. The end result is that we will be able to state what sampling rate is needed to achieve certain performance criteria; in a sense, this will allow us to quantify how much the utility company should be willing to 'pay' for the user's power consumption data.

C. Example control policy

Note that the total power consumption at time k is given by $Y_k = \mathbf{1}^T x_k$, where $\mathbf{1}$ is a vector of ones. Furthermore, note that the system is controllable if $\beta^i \neq 0$ for all users.

We provide an example of a control policy. Suppose the system is controllable. Now, let this be the control scheme: given a measurement of the complete state, x_k, it calculates the total demand: $Y_k = \mathbf{1}^T y_k$. It also has a given target value for Y_k, call this Y^*_k. Then, it tries to set each y^j_{k+1} to be $y^{j*}_k = \frac{y_k}{Y_k} Y^*_k$. Thus, it attempts to maintain proportions as well as set a target consumption. If the state is not measured at time k, it will just use the expected value $\mathbb{E}[y_k]$, given available measurements, as an estimator. In both cases, it is easy to see that:

$$u^i_k = \frac{1}{\beta^i}(y^{j*} - \alpha^i \mathbb{E}[y_k] - \mu^i_k) \tag{36}$$

for all users i. This control is deterministic given a sampling rate, target demand, and measurements.

Notice that the problem of privacy-preserving control is different from other problems in networked control systems because while we might not be able to sample the electricity consumption y_k as frequently as we would like, our control signal does not need to be sampled, as more frequent controls u_k do not compromise privacy. In future work we plan to place additional realistic constraints to the demand-response control signal such as minimizing the number of times it is used (i.e., maximize the amount of time $u_k = 0$), or placing bounds u_{min}

Now, say we have the following criteria: the system should be within an interval $[Y_k^* - Y_b, Y_k^* + Y_b]$ for all k with probability $1-\epsilon$, where y_b is some pre-specified constant. We can simulate the efficacy of this control while varying sampling rates.

D. Simulation results

We simulate the model and control policy for a two-user system. Our goal is to show an example of the type of studies an AMI deployment might consider in order to select the sampling rate of their smart meters. In particular, our goal is to define an upper bound $Y_k^* + Y_b$ that will keep the second term of our desired target cost $C(\{Y_k\})$ bounded.

We then show two metrics of interest as functions of the sampling interval N:

1) $\Pr\left[Y_{k+N} \in [Y_k^* - Y_b, Y_k^* + Y_b\right]$ in Figure 2,
2) The interval $S = [Y_{min}, Y_{max}]$ such that $\Pr[Y_{k+N} \in S] \geq 0.99$ in Figure 3.

While demand-response programs only care about the highest interval bound Y_{max} (not the lower one), if we don't place the lower bound in our problem formulation, the optimal control signal would attempt to drive the demand to be as low as possible. In future work we plan to address this drawback by adding a control cost function that includes a penalty cost each time we use the control signal.

Our simulation uses the following parameters:

$$\begin{aligned} A &= \text{diag}(0.64, -0.32) \\ B &= \text{diag}(2.54, 1.27) \\ \Sigma_d &= \text{diag}(10, 20) \\ Y^* &= 94.9482 \\ Y_b &= 25 \end{aligned} \tag{37}$$

where $Y_k^* = Y^*$ for all k. These parameters are influenced by results in statistical estimation of thermostatically controlled loads [19]. We use the time window $K = 128$. For μ_k, we plot the nominal total demand and desired control interval in Figure 1.

To see the efficacy of our control policy and how it varies with different sampling rates, we used Monte Carlo simulation methods. We have the following results. First, we plot the probability our control will keep the total demand in the desired interval, as a function of sampling intervals, in Figure 2. Next, we show the demand intervals we can maintain with 99% probability in Figure 3.

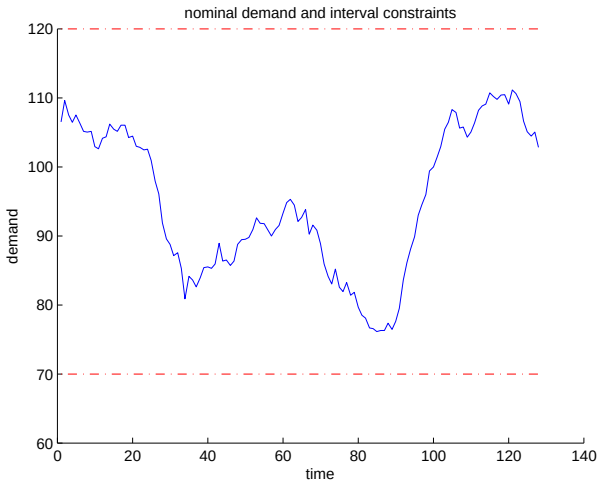

Fig. 1. The nominal total demand and the desired demand bounds.

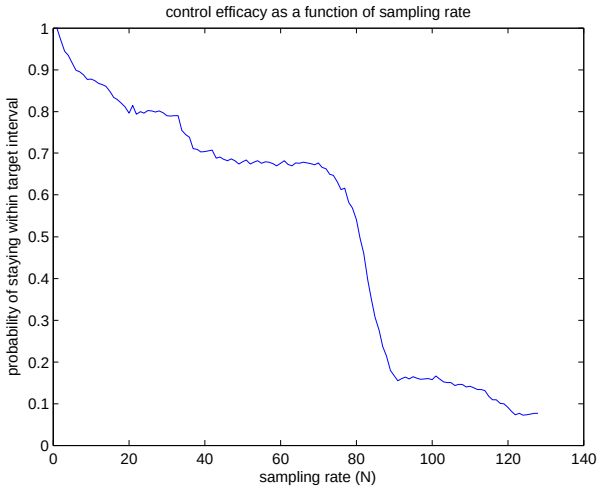

Fig. 2. The probability our control will keep us in the desired interval, as a function of sampling rate.

With these simulations our goal was to perform a preliminary exploratory study of the type of problem formulation and the type of properties we would like to maintain in a demand-response system. In future work we plan to explore the analytical properties of privacy-preserving demand-response in AMI networks.

VII. Conclusions and Future Work

In this paper we presented a unified cost model that allowed us to study the impacts of electricity theft detection and privacy for the bottom line of a distribution utility (their profits).

We first formulated a game between the distribution utility and electricity thieves, and found the Nash equilibrium of the game as a probability density function that

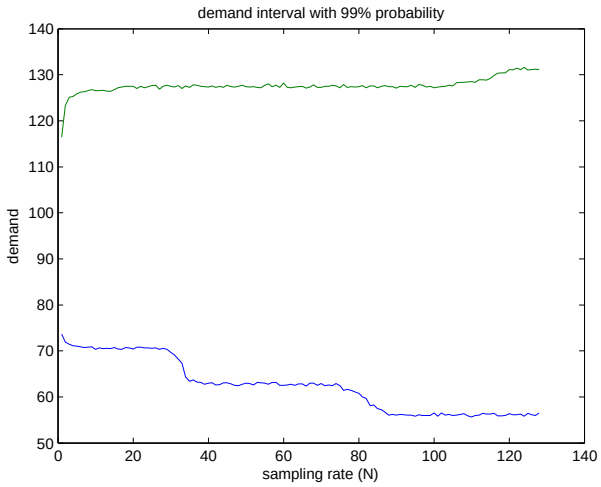

Fig. 3. The demand interval we can maintain with probability 99%, given a sampling rate.

attackers and defenders must choose in order to send AMI measurements y_k.

We then performed a preliminary analysis of how to achieve the maximum level of privacy possible subject to a bound on the maximum load.

In future work we plan to explore more directions. We are particularly interested in a study on *privacy-preserving electricity theft detection*. In other words, we would like to explore how the sampling interval of smart meters affects the ability of the distribution utility to identify anomalies and electricity theft. According to the central-limit theorem, for large sampling intervals and with an i.i.d assumption, the distribution of attackers and defenders will follow a Gaussian distribution. The mean and the variance of an honest user will be known in advance, and the mean of the attacker will also be known thanks to the constraint imposed by q_U; therefore the goal of the attacker will be to find the variance that minimizes their probability of detection:

Another research direction is to place more realistic demand-response constraints in the problem formulation of the control signal as well as to the objective in the controller. Such constraints include a cost for using the control signal, a control signal that is bounded, users that are not elastic to the control, etc.

Finally, we would also like to jointly study the problem of privacy preserving demand response under electricity theft.

References

[1] P. Antmann, "Reducing technical and non-technical losses in the power sector," World Bank, Tech. Rep., July 2009.

[2] S. McLaughlin, D. Podkuiko, S. Miadzvezhanka, A. Delozier, and P. McDaniel, "Multi-vendor penetration testing in the advanced metering infrastructure," in *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*, Dec. 2010.

[3] M. Davis, "Smartgrid device security. adventures in a new medium," http://www.blackhat.com/presentations/bh-usa-09/MDAVIS/BHUSA09-Davis-AMI-SLIDES.pdf, July 2009.

[4] D. Peterson, "AppSecDC in review: Real-world backdoors on industrial devices," http://www.digitalbond.com/2012/04/11/appsecdc-in-review/, April 2012.

[5] B. Krebs, "FBI: smart meter hacks likely to spread," http://krebsonsecurity.com/2012/04/fbi-smart-meter-hacks-likely-to-spread/, April 2012.

[6] E. De Buda, "System for accurately detecting electricity theft," US Patent Application 12/351978, Jan. 2010.

[7] A. Appel, "Security seals on voting machines: A case study," *ACM Transactions on Information and Systems Security*, vol. 14, pp. 1–29, 2011.

[8] A. Lesser, "When big IT goes after big data on the smart grid," http://gigaom.com/cleantech/when-big-it-goes-after-big-data-on-the-smart-grid-2/, March 2012.

[9] C. Geschickter, *The Emergence of Meter Data Management (MDM): A Smart Grid Information Strategy Report*. GTM Research, 2010.

[10] D. Mashima and A. A. Cardenas, "Evaluating electricity theft detectors in smart grid networks," in *Proceedings of Research in Attacks, Intrusions and Defenses (RAID) Symposium.*, September 2012.

[11] S. McLaughlin, B. Holbert, S. Zonouz, and R. Berthier, "AMIDS: A multi-sensor energy theft detection framework for advanced metering infrastructures," in *Proceedings of the third IEEE International Conference on Smart Grid Communications (SmartGridComm)*, November 2012.

[12] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *First IEEE Smart Grid Commnunications Conference (SmartGridComm)*, October 2010.

[13] S. E. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *ACM Conference on Computer and Communications Security*, 2011, pp. 87–98.

[14] F. Li, B. Luo, and P. Liu, "Secure Information Aggregation for Smart Grids Using Homomorphic Encryption," in *First IEEE Smart Grid Commnunications Conference (SmartGridComm)*, October 2010.

[15] G. Taban and V. D. Gligor, "Privacy-preserving integrity-assured data aggregation in sensor networks," in *The 2009 International Symposium on Secure Computing (SecureCom) 2009*, 2009, pp. 168–175.

[16] K. Kursawe, G. Danezis, and M. Kohlweiss, "Privacy-friendly aggregation for the smart-grid," in *Privacy Enhancing Technologies - 11th International Symposium, PETS*, July 2011.

[17] A. Rial and G. Danezis, "Privacy-preserving smart metering," in *Proceedings of the 2011 ACM Workshop on Privacy in the Electronic Society, WPES*, October 2011.

[18] D. Callaway and I. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, January 2011.

[19] A. Aswani, N. Master, J. Taneja, V. Smith, A. Krioukov, D. Culler, and C. Tomlin, "Identifying models of hvac systems using semiparametric regression," in *Proceedings of the American Control Conference*. IEEE, 2012.