

Incentive Design for Direct Load Control Programs

Mahnoosh Alizadeh[†], Yuanzhang Xiao^{*}, Anna Scaglione[†], and Mihaela van der Schaar^{*}

[†]University of California Davis ^{*}University of California Los Angeles

Abstract—We study the problem of optimal incentive design for voluntary participation of electricity customers in a Direct Load Scheduling (DLS) program, a new form of Direct Load Control (DLC) based on a three way communication protocol between customers, embedded controls in flexible appliances, and the central entity in charge of the program. Participation decisions are made in real-time on an event-based basis, with every customer that needs to use a flexible appliance considering whether to join the program given current incentives. Customers have different interpretations of the level of risk associated with committing to pass over the control over the consumption schedule of their devices to an operator, and these risk levels are only privately known. The operator maximizes his expected profit of operating the DLS program by posting the right participation incentives for different appliance types, in a publicly available and dynamically updated table. Customers are then faced with the dynamic decision making problem of whether to take the incentives and participate or not. We define an optimization framework to determine the profit-maximizing incentives for the operator. In doing so, we also investigate the utility that the operator expects to gain from recruiting different types of devices. These utilities also provide an upper-bound on the benefits that can be attained from any type of demand response program.

I. INTRODUCTION

With the lack of utility-scale storage options in the power grid, the need to make electricity demand active is becoming more pressing each day. In the research community, the most favored option to make this vision happen is real-time pricing (RTP). Optimal real-time prices, if calculated correctly, would maximize the social surplus. However, there are several barriers that currently hinder the realization of this vision: 1) End-use customers need to have certainty in prices for a certain look-ahead horizon to plan consumption. With a wide-spread integration of renewables, calculating reliable clearing prices hours ahead of operation is challenging; 2) the lack of concrete models for the consumption flexibility of electricity consumers in today's market, specifically with the granularity needed to allow high penetration of renewables, and the time-inhomogeneity of these models due to variable appliance arrival and flexibility patterns; 3) the strict reliability requirements of power grid operations, allowing small error margins in price design. A line of iterative methods that actively ask for the customers' collective response to price signals are being proposed to address these problems [1], [2].

One the opposite side of the spectrum, Direct Load Control (DLC) has proven to be a popular type of demand control for grid operators, mostly due to the reliable and predictable nature of the demand's response to control signals. Currently, customers providing DLC services to the power grid sign long-term contracts that allow the provider to cut off the

supply of electricity to some of their appliances (e.g., air conditioners and pumps) during occasional contingencies. Commonly, DLC strategies take customer participation as a given, assuming small fixed ex-ante monetary payments for all participants, regardless of the level of service they provide or the discomfort they experience. However, participation in a DLC program presents risks for the customers and a rational intelligent customer would not provide this service extensively without appropriate financial incentives. In this work, we take a first step in addressing these inherent economic problems for designing day-to-day DLC incentives.

To avoid confusion with currently employed DLC programs, in which the individual consumption of appliances are not observable, and feedback control strategies are used [3], we will refer to our proposed program as Direct Load Scheduling (DLS). DLS does not merely cut off the electricity supply of appliances. Rather, a control center can optimally *plan* the consumption within consumer-specified laxity limits, e.g., scheduling the charge of an Electric Vehicle (EV) by a deadline. Contrary to common DLC practice, customers provide the DLS authority with an explicit expression of the service they need in an online fashion, leaving no uncertainty in how the demand responds to control signals. Appliances are recruited only on an event-based basis, i.e., every time they have to perform a task. Previous works have proposed various cost-minimizing scheduling algorithms for DLS, specifically for EVs, e.g. [4]–[9]. The possibility of using Vehicle-to-Grid (V2G) services for regulation and grid support has also been studied, e.g. [10], [11]. However, charge interruption and V2G could decrease EV battery life, delay full charge and present risks to the customer, and violate customer privacy. Thus, proper economic incentives for participation should be studied.

We assume that the DLS program is run by the same entity that provides electricity to customers, hereafter referred to as the *aggregator*. The aggregator is essentially an energy *trader*. Following current practice, we assume that the aggregator buys electricity at time-varying wholesale prices from the energy market and on sells this energy to end-use customers at flat rates, without being able to deny any electricity service requests. Thus, it essentially acts as intermediary node that shields the end-use customers from wholesale price fluctuations. With no demand flexibility, the aggregator has no control over the profitability of his venture in the short run. Rather, his profit is determined by 1) wholesale prices; 2) the predetermined billing tariffs, which are flat, regulated, and change very slowly; 3) the consumption behavior of customers, which is out of control of the aggregator. To

overcome this issue, the aggregator runs a DLS program and pays customers in return for directly scheduling their appliances, i.e., it effectively buys flexibility from customers while keeping them on flat rates. This paper focuses on the design of a market for trading flexibility between a single aggregator and a population of randomly arriving appliances. We model this market as a monopoly, in the sense that each customer has access to only one aggregator and is a price-taker. Competition between aggregators is left to future work.

II. MODEL

We consider a community with a large population of customers owning controllable appliances, and one aggregator. The basic observation of this work is that directly controlling various types of appliances has different utilities for the aggregator. This utility is a function of how flexible and how high the electricity consumption of the device is, and of the dynamic state of the grid at the time the appliance is used. Thus, it is best if the incentive that the aggregator pays to customers to recruit their appliances can vary dynamically with *time and appliance type*¹. To allow for this, we assume that these incentives are posted by the aggregator in dynamically updated and publicly available tables for all customers, which can be thought of as *menus* listing the different incentives customers can receive. Customers planning to use controllable appliances would then respond to these posted incentives by deciding whether they want to participate or not, and how much laxity they wish to offer the aggregator.

We would like to point out that having dynamically changing menus does not necessarily mean that the incentives have to be re-designed every hour of every day. In fact, incentives will exhibit similar daily or weekly cycles as market prices.

We define this market under the following assumptions:

Assumption 1: The aggregator's revenue from recruiting each appliance is additive and independent of other appliances' participation;

Assumption 2: A customer's initial choice to use an appliance is perfectly inelastic and not affected by incentives;

Assumption 3: The aggregator's load does not affect the wholesale market clearing prices;

Assumption 4: The aggregator has access to ex-ante forecasts of expected wholesale prices.

Assumptions 3 and 4 simplify the expression of the aggregator's utility when recruiting appliances in the DLS program. Next, we look at how customers respond to incentives.

A. Individual Appliance DLS Commitment Problem

We model the electricity consumption of customers using tasks that dynamically arrive, receive service, and depart at discrete time epochs indexed by $t = \{1, 2, 3, \dots\}$. An arrival event corresponds to the earliest time at which it is possible for an appliance to start its job. Each task, indexed

¹While, in theory, the incentive could vary across different customers offering the same service, here we assume that we want the incentives to be perceived as fair and not violate consumer privacy limits.

by i , has a so-called characteristic vector $\mathbf{v}_i \in \mathcal{C}$. The elements of \mathbf{v}_i fully describe the nature of the task, which could simply be the charge duration and rate for an EV, the desired temperature for a thermostatically controlled load (TCL). We further assume that \mathbf{v}_i can only be chosen from a finite codebook $\mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_Q\}$, designed to achieve a bounded and controllable load modeling error. This allows us to cluster similar energy requests in a finite number of classes, indexed by

$$q_i \in \mathcal{Q} = \{1, 2, 3, \dots, Q\},$$

which will help to highly reduce the computational effort of our algorithm. Thus, the consumption characteristics of task i are uniquely defined by its cluster index q_i .

The strategy set is defined by the laxity limits the customer commits to provide to the DLS program when recruited, described by the index

$$m_i \in \mathcal{M}_{q_i} = \{0, 1, 2, \dots, M_{q_i}\},$$

and referred to as the appliance's *mode*. Mode $m = 0$ corresponds to *no laxity*, i.e., the customer will not participate in the program. These laxity limits could, for example, correspond to the slack for an EV charge, or the width of the comfort band for a TCL. The appliance embedded controller is queried by the customer application, which maps the physical state of the appliance onto the set of modes that are available for the customer to choose from and, possibly, an indication of what is the flat rate cost with no laxity.

The customer then chooses the mode after observing the incentive menu. We denote by $I_q^t(m)$ the incentive the customer could receive for releasing the control of an appliance in cluster q in mode m at time t , with $I_q^t(0) = 0$. We denote the vector containing all (non-trivial) incentives available to appliances in cluster q under modes $m \geq 1$ as

$$\mathbf{I}_q^t = [I_q^t(1), I_q^t(2), \dots, I_q^t(M_q)]^T.$$

We expect the customers in charge of making commitment decisions to be only *boundedly rational*, i.e., they are likely to only spend a limited amount of effort in considering the economic utility of participation in the program or updating their home energy management system's parameters. However, in the rest of this section, we introduce an analytical model which is valid for rational customers. We do so with the disclaimer that this model is not a necessary element of our design and we present it solely to provide some intuition into the nature of this decision making problem.

Design constraint 1 (Diminishing Payoffs): Note that if an appliance does not join the DLS program right after its arrival time, it will *lose* some laxity. However, if consumers have reasons to believe that taking some risk and waiting to participate in the DLS program later could increase their expected payoff, even though the amount of laxity they can offer decreases, they will do so. To avoid this situation, we assume that the incentives will be designed such that the customer's payoff for the same level of risk (e.g., same

deadline) will be monotonically non-increasing in time. For example, if we assign a separate mode m every time the laxity offered by a deferrable load is increased by one unit, we require that

$$I_q^t(m) \leq I_q^{t-1}(m+1).$$

A Rational Customer Model: The utility gained by the customer from operating task i at time t in each mode is defined as a function

$$V_i^t(\cdot) : \mathcal{M}_{q_i} \rightarrow \mathbb{R},$$

which includes three terms: 1) The incentive $I_{q_i}^t(m_i)$ available for mode m_i ; 2) The commitment risk (disutility) $R_i^t(m_i)$ associated with agreeing to receive service under mode m_i , modeled through a privately known function

$$R_i^t(\cdot) : \mathcal{M}_{q_i} \rightarrow \mathbb{R}^+,$$

with $R_i^t(0) = 0$, i.e., there is no risk when the customer decides not to participate in the DLS program. Without loss of generality, we order the modes from low to high risk. This would result in monotonically non-decreasing individual risk functions $R_i^t(m_i)$; 3) the utility of receiving the standard service of using electricity and finishing a job. This term is a constant, since we assume it will eventually happen for every request, either through the DLS program under some mode $m_i \geq 1$, or through the standard service model of the power grid ($m_i = 0$). The disutility of not receiving this standard service in case of an emergency departure is captured through the risk term. Thus, we eliminate this term from the customer's decision making model.

Consequently, assuming that the risk function are chosen such that they have the same unit as the monetary incentive, the customer's utility is quasi-linear and is given by

$$V_i^t(m_i) = I_{q_i}^t(m_i) - R_i^t(m_i).$$

Upon receiving the incentive information $I_{q_i}^t(m_i)$ from the aggregator, the customer would solve the following optimization to determine the best operating mode of appliance i :

$$\max_{m_i \in \mathcal{M}_{q_i}} V_i^t(m_i). \quad (1)$$

The customer will not participate in the DLS program ($m_i = 0$) if none of the $I_{q_i}^t(m_i)$'s for $m_i \geq 1$ are at least marginally higher than the risk $R_i^t(m_i)$. Otherwise, the customer would pick the mode such that the margin between $I_{q_i}^t(m_i)$ and $R_i^t(m_i)$ is highest (highest residual worth). Ties are broken uniformly at random.

Even if we assume that this model perfectly describes the decision making procedure of all customers, since $R_i^t(\cdot)$ is only privately known, the aggregator cannot predict the outcome of (1) deterministically. Thus, to determine the incentives $I_q^t(m_i)$, the aggregator is faced with an optimization problem with incomplete information. The goal of the aggregator would be to maximize its expected profit, given aggregate statistics about the population's response to incentives.

B. The Aggregator Problem

In order to recruit directly controllable appliances, the aggregator needs to design appropriate incentives for every possible mode in all possible clusters, i.e., $I_q^t(m), \forall q \in \mathcal{Q}, \forall m \in \mathcal{M}_q$. The LSE participates in the electricity wholesale market run on an hourly basis, where it buys energy to serve its load, directly controllable or not. We will elaborate more on the nature of this market interaction later. For now, we take the the expected profit that the aggregator can make, through wholesale market transactions, by directly controlling an individual appliance from cluster q in mode m as known. For a recruitment of cluster q in mode m at time t , we denote this utility by $U_q^t(m)$. The utility of not recruiting an appliance is zero, i.e., $U_q^t(0) = 0$. As per Assumption 1, we define this recruitment utility as additive and independent, rendering the incentive design problem separable for individual appliances.

Since recruiting an appliance from cluster q in mode m at time t comes at a cost equal to $I_q^t(m)$, the net revenue of the aggregator from this recruitment, denoted by $N_q^t(m)$, is

$$N_q^t(m) = U_q^t(m) - I_q^t(m). \quad (2)$$

However, note that the mode m is chosen by the customer after seeing the incentives $I_q^t(m)$, through (1). Since only statistical information on the customers response strategy is available to the aggregator, only the *expected net revenue* can be maximized. Denote the event that any customer in cluster q picks mode m as $E_{q,m}(\mathbf{I}_q^t)$. This event happens if:

- Individual rationality constraint (**IR**):

$$I_q^t(m) - R_i^t(m) \geq 0,$$

- Incentive compatibility constraints (**IC**):

$$I_q^t(m) - R_i^t(m) \geq I_q^t(m') - R_i^t(m'), \quad \forall m' \in \mathcal{M}_q.$$

Consequently, the expected net revenue of recruiting an appliance of cluster q , simply denoted by N_q^t , is given by

$$N_q^t = \sum_{m \in \mathcal{M}_q} P(E_{q,m}(\mathbf{I}_q^t))(U_q^t(m) - I_q^t(m)), \quad (3)$$

which the aggregator would like to maximize; i.e. the aggregator would want to design the incentives as follows:

$$\max_{\mathbf{I}_q^t} \sum_t N_q^t. \quad (4)$$

The summation over time is required to find the optimal incentives in the presence of the diminishing payoff design constraint. The reader could envision that for a real-time implementation, the optimization (4) could be solved over a receding horizon.

In order to solve (4), it is essential to have an understanding of how $P(E_{q,m}(\mathbf{I}_q^t))$ changes with the incentives \mathbf{I}_q^t . Next, we propose two different views for approaching this problem.

1) *Bayesian Approach*: Here we assume that the aggregator has access to statistically learned prior knowledge on the risk levels of customers in different clusters. With this view, the incentive design problem would be similar in nature to optimal Bayesian unit-demand pricing, given that the customers' risk levels (valuations) for different modes (items) cannot be assumed to be drawn from independent distributions. The risk levels that a customer perceives for committing to the program with a single appliance under different modes are correlated. For example, in the case of EVs or any other deferrable loads, the mode corresponds to the amount of laxity that accompanies the request. Clearly, offering a higher consumption laxity entails an additional risk over that of offering a lower laxity, and these variables cannot be considered independent.

We assume that the aggregator parameterizes the risk function of each appliance i in cluster q according to:

$$R_i^t(m) = \gamma_i r_q^t(m), \quad (5)$$

where γ_i is a task-specific non-negative continuous random variable and represents the *type* of an individual task i , and the variables $r_q^t(m)$ are deterministic and shared by all appliances in the same cluster q , with $r_q^t(0) = 0$, and can model the average attitude of population towards risk. To ensure that this is a realistic parameterization of the risk function, the aggregator could suggest this specific structure as default to consumers when they pick their risk functions.

We assume that the aggregator has access to statistical priors for the type γ_i , and we denote by $F_\gamma^q(g)$ the cumulative distribution of γ_i for appliances in cluster q . With this new notation, $P(E_{q,m}(\mathbf{I}_q^t))$ for all $m \geq 1$ is the probability of the following event:

$$\mathbf{IR} - \gamma_i \leq \frac{I_q^t(m)}{r_q^t(m)} = \frac{I_q^t(m) - I_q(0)}{r_q^t(m) - r_q^t(0)}, \quad (6)$$

$$\mathbf{IC1} - \gamma_i \leq \frac{I_q^t(m) - I_q^t(m')}{r_q^t(m) - r_q^t(m')}, \quad \forall 1 \leq m' < m \quad (7)$$

$$\mathbf{IC2} - \gamma_i \geq \frac{I_q^t(m') - I_q^t(m)}{r_q^t(m') - r_q^t(m)}, \quad \forall m' > m, \quad (8)$$

so we should have $\gamma_i \in [l_{q,m}^\gamma(\mathbf{I}_q^t), h_{q,m}^\gamma(\mathbf{I}_q^t)]$, with

$$h_{q,m}^\gamma(\mathbf{I}_q^t) = \min \left\{ \frac{I_q^t(m) - I_q^t(m')}{r_q^t(m) - r_q^t(m')} \Big|_{0 \leq m' < m} \right\}, \quad (9)$$

$$l_{q,m}^\gamma(\mathbf{I}_q^t) = \max \left\{ \frac{I_q^t(m') - I_q^t(m)}{r_q^t(m') - r_q^t(m)} \Big|_{m' > m} \right\}. \quad (10)$$

which gives,

$$P(E_{q,m}(\mathbf{I}_q^t)) = F_\gamma^q(h_{q,m}^\gamma(\mathbf{I}_q^t)) - F_\gamma^q(l_{q,m}^\gamma(\mathbf{I}_q^t)) \quad (11)$$

However, due to the absence of any natural ordering, these constraints will render the optimization problem (4) rather complex. Thus, next, we will impose a design constraint that ensures that *local incentive compaibility* is sufficient for decision making, i.e., if the customer prefers mode m over

adjacent modes $m + 1$ and $m - 1$, he/she will also prefer mode m over all modes $m' > m$ and $m' < m$.

Design constraint 2 (Single-Crossing Incentive Profile): we will design the incentive profile such that $\forall m \in \mathcal{M}_q$, the ratio $\frac{I_q^t(m+1) - I_q^t(m)}{r_q^t(m+1) - r_q^t(m)}$ is non-increasing, i.e., incentives grow slower than risks for higher modes m .

Proposition 2.1: If the incentive profile is single-crossing, customer i will pick mode $m \geq 1$ simply iff

$$\frac{I_q^t(m+1) - I_q^t(m)}{r_q^t(m+1) - r_q^t(m)} \leq \gamma_i \leq \frac{I_q^t(m) - I_q^t(m-1)}{r_q^t(m) - r_q^t(m-1)}, \quad (12)$$

where the right hand inequality ensures (6) and (7), and the left hand inequality ensures (8). To keep expression (12) compact, we use a dummy mode $m = M_q + 1$, with $I_q(M_q + 1) = I_q(M_q)$, and $r_q^t(M_q + 1) = r_q^t(M_q) + 1$.

We acknowledge that the single-crossing condition restricts the values that the incentive profile can take into a region that may be suboptimal for the aggregator. However, it considerably lowers the numerical effort to solve (4).

To illustrate the next steps required to solve the problem, specifically required for our numerical experiments, we assume that the types γ_i for cluster q are drawn from a uniform distribution over $[0, \gamma_{\max}^q]$. We design the incentives such that all of the ratios $\frac{I_q^t(m+1) - I_q^t(m)}{r_q^t(m+1) - r_q^t(m)}$ fall in the probability space of γ_i . This happens if the incentives are positive and non-decreasing with m , and that the ratio $I_q^t(1)/r_q^t(1)$ is not above γ_{\max}^q , i.e., someone might pick mode $m = 1$. Consequently, imposing the diminishing pay-off constraints for deferrable loads, (4) is written as,

$$\begin{aligned} \max_{\mathbf{I}_q^t} \sum_t N_q^t &= \frac{1}{\gamma_{\max}^q} \sum_t \sum_{m=1}^{M_q} \left(U_q^t(m) - I_q^t(m) \right) \times \\ &\quad \left(\frac{I_q^t(m) - I_q^t(m-1)}{r_q^t(m) - r_q^t(m-1)} - \frac{I_q^t(m+1) - I_q^t(m)}{r_q^t(m+1) - r_q^t(m)} \right), \\ \text{s.t. } 0 &\leq I_q^t(m) \leq I_q^{t-1}(m+1), \quad 1 \leq m \leq M_q - 1, \\ \frac{I_q^t(m+1) - I_q^t(m)}{r_q^t(m+1) - r_q^t(m)} &\leq \frac{I_q^t(m) - I_q^t(m-1)}{r_q^t(m) - r_q^t(m-1)}, \quad m \geq 1, \\ I_q^t(m) &\geq I_q^t(m-1), \quad m \geq 1, \\ \frac{I_q^t(1)}{r_q^t(1)} &\leq \gamma_{\max}^q, \end{aligned} \quad (13)$$

for which the objective function is quadratic in \mathbf{I}_q^t , and the constraints are all affine.

The statistics on the risk functions can be obtained by conducting market surveys or using learning techniques. The details are out of the scope of this work.

2) *Model-free Learning Approach*: If there is no single underlying model that characterizes how costumers respond to participation incentives, or no information is available on the private risk functions of customers, model-free online learning techniques can be used to directly learn the probabilities $P(E_{q,m}(\mathbf{I}_q^t))$. These approaches explore different alternatives for the incentive signals \mathbf{I}_q^t , observe the response

of the population, an update their estimate of $P(E_{q,m}(\mathbf{I}_q^t))$ accordingly. This requires that the event of a customer *considering* to join the program be observable to the aggregator, which is not the case in our current design. This issue can be addressed by asking the customers to make an *anonymous query* every time they need to look at the incentive menus for a specific cluster.

Next, we study the recruitment utility of different devices for the aggregator.

III. THE RECRUITMENT UTILITY FUNCTIONS

There are several options for an aggregator to profit from recruiting flexible appliances. Here we only study the possibility of load shifting based on the recruitment of long-duration non interruptible deferrable appliances (e.g., EVs, washer/dryers) or preheating and precooling of TCLs.

In order to solve (4), the aggregator needs to know beforehand how much utility it could expect from recruiting an appliance in cluster q in mode m at time t . We denoted this value by $U_q^t(m)$. In this Section, we specifically focus on calculating this payoff for the case of load shifting with non-interruptible deferrable loads and Thermostatically Controlled Loads (TCL), with the numerical results focusing specifically on Plug-in Hybrid Electric Vehicles (PHEV). Due to limited space, we leave the discussion of interruptible and sheddable loads to future work.

Energy is traded on an hourly basis, and the demand is modeled as constant in hourly intervals in the energy market. The task of ensuring the sub-hourly balance of demand and supply is left to ancillary service providers, which are procured and dispatched by the grid operator to respond quickly (in a matter of minutes or seconds) to variations of the demand. Traditionally, ancillary services are offered by fast ramping generators.

In order to serve its' load, the aggregator needs to purchase a certain amount of energy for every hour ℓ from the wholesale energy market, which we denote by $L(\ell)$, and refer to as the *base load*. We assume that the aggregator has access to ex-ante forecasts of the wholesale energy market clearing prices, and we denote the expected value of the price for hour ℓ as $\pi^e(\ell)$. We assume the aggregator's load is small enough to not affect the price. The aggregator can save money in the wholesale energy market if it recruits flexible appliances and shifts their load to hours at which energy is cheaper.

Denote the set of all feasible *hourly load traces* of an individual task in cluster q under mode m recruited at time t as $\mathcal{L}_{q,m}^t(\ell) \in \mathcal{L}_{q,m}^t$. To obtain an hourly load profile for an appliance with sub-hourly consumption variations, we simply average out the total consumption of the device within each hour. We denote as $L_{q,0}^t(\ell)$ the consumption profile of the appliance if it is not recruited by the DLS program and starts its consumption at time t . Serving this power to the consumer through buying energy from the wholesale market presents a cost for the aggregator. By having direct control over this appliance, the aggregator can expect to gain the following

payoff by shifting the load away from peak hours:

$$U_q^t(m) = \sum_{\ell} \pi^e(\ell) L_{q,0}^t(\ell) - \min_{L_{q,m}^t(\ell) \in \mathcal{L}_{q,m}^t} \sum_{\ell} \pi^e(\ell) L_{q,m}^t(\ell). \quad (14)$$

Remark 3.1: The consumption profile of flexible appliances that do not choose to join the DLS program, i.e., $L_{q,0}^t(\ell)$, is a function of the tariff that the customers are billed on. On a flat tariff, the customers do not need to spend any effort to find the best time at which they should consume electricity, and simply plug in an appliance at the request arrival time.

Design constraint 3: Here we assume that for all deferrable loads, the mode index $m \in \mathcal{M}_q$ is directly equivalent to the amount of laxity (slack time) that accompanies the request.

A. Deferrable Non-interruptible Appliances

Here we look at appliances whose consumption profile can merely be shifted in time, but cannot be modified in any way once started. A classic example is that of a washing machine cycle. In this case, the control variable is the *activation time* of the task, denoted by α (See Fig. 1). The set of possible values of α depends on the initial time at which the request arrives, and the laxity that accompanies the request, i.e., m .

Once activated, the appliance consumes a predetermined amount of power, denoted by the time-shifted pulse $g_q(j - \alpha)$ for appliances in cluster q , with a length of Γ_q epochs. This pulse varies on a sub-hourly basis as a function of the time epochs j . To relate this with the variables required to solve the optimization (15), we need to map this sub-hourly varying load to an hourly load profile, resulting in

$$\mathcal{L}_{q,m}^t = \left\{ L_{q,m}^t(\ell) = \sum_{j=\ell S}^{(\ell+1)S-1} g_q(j - \alpha) | t \leq \alpha \leq t + m \right\}.$$

Notice the summations required to calculate the average hourly consumption of the device from the sub-hourly load. To avoid the inconvenience of going from $g_q(j - \alpha)$ to hourly consumption values $L_{q,m}^t(\ell)$ when solving (15), we simply expand the expected wholesale energy price vector $\pi^e(\ell)$ to define virtual sub-hourly wholesale prices $\pi^p(j)$,

$$\pi^p(j) = \sum_{\ell} \frac{1}{S} \pi^e(\ell) \Pi\left(\frac{j - \ell S}{S}\right),$$

where $\Pi(\cdot)$ denotes the unit pulse function between $[0, 1)$. With this new definition, we have

$$\begin{aligned} \sum_{\ell=\lfloor \frac{\alpha}{S} \rfloor}^{\lfloor \frac{\alpha+\Gamma_q}{S} \rfloor} \pi^e(\ell) L_{q,m}^t(\ell) &= \sum_{\ell=\lfloor \frac{\alpha}{S} \rfloor}^{\lfloor \frac{\alpha+\Gamma_q}{S} \rfloor} \pi^e(\ell) \sum_{j=\ell S}^{(\ell+1)S-1} g_q(j - \alpha) \\ &= \sum_{j=\alpha}^{\alpha+\Gamma_q} \pi^p(j) g_q(j - \alpha). \end{aligned} \quad (15)$$

Consequently, the expected utility of recruiting an appliance in cluster q under mode m at time t is

$$U_q^t(m) = \sum_{j=t}^{t+\Gamma_q} \pi^p(j)g_q(j-t) - \min_{\alpha} \left[\sum_{j=\alpha}^{\alpha+\Gamma_q} \pi^p(j)g_q(j-\alpha) \right] \quad \text{s.t.} \quad t \leq \alpha \leq t+m, \quad (16)$$

where the first term refers to the cost incurred by the aggregator if no demand control is exercised, whereas the second term mirrors the lowest possible cost with which the appliance can be served, considering a demand laxity of m time units.

B. Thermostatically Controlled Loads

Another category of appliances that can help the aggregator save money in the energy market are TCLs, which can use the inherent energy storage property of building thermal mass to shift their load to cheaper hours. A substantial potential for intrahourly load shifting with heating and cooling devices is through preheating and precooling of the air-conditioned space. To preheat a space, the temperature is increased above the comfort band of the user at off-peak hours, before residents/employees arrive at home/work. The preheating should be scheduled such that by the time the residents arrive, the temperature is close to the highest acceptable temperature in the comfort band. While some commercial building utility managers currently exercise this option during night hours to save on energy, with the integration of renewable resources, the *best time and amount* of preheating would be variable on a daily basis and needs coordination with the aggregator.

We assume that each TCL is equipped with a thermostat that keeps the temperature in a certain comfort band $[x_q^{\min}, x_q^{\max}]$, with x_q^{\min} and x_q^{\max} representing the lower and upper thresholds chosen by the customers in cluster q . Here, we adopt a simple first-order model proposed in [12] to explain the state dynamics of temperature. Let $x_i(j)$ (a representative temperature) characterize the state of the i -th heating device; $x_i(j)$ evolves according to the linear stochastic difference equation

$$x_i(j) - x_i(j-1) = -k_q(x_i(j-1) - x_a(j-1)) + W_q b_i(j) + \nu_i(j), \quad (17)$$

where we denote by

- k_q the average loss rate for buildings in cluster q ;
- $x_a(j)$ the ambient temperature;
- W_q the average rate of heat gain supplied by devices in cluster q ;
- $b_i(j)$ the operating state of the device at time j (1 for "on" or 0 for "off");
- $\nu_i(j)$ a zero mean Wiener noise process.

We assume that, once on, the unit consumes a constant power approximately equal to P_q for cluster q . Thus, the only variable that needs to be chosen and can affect the

aggregator's cost is the times at which to run the unit for preheating. Due to the constant power of the device, the amount of preheating is a direct consequence of this choice. However, we would like to point out that more sophisticated methods for determining the trajectory of the temperature for precooling/preheating exists in the literature, both based on optimal control theory [13]–[16], or learning and simulation [17]. The calculations in the section could be updated to account for any of these techniques.

In order for the temperature to be around x_q^{\max} at the time that the building occupants arrive, we must have previously increased it enough in the preheating period. This preheating process can be done over multiple disjoint time intervals, whenever it is deemed as a cost-effective action.

For appliances recruited at time t , preheating can start right away and can continue to the time at which building occupants arrive, which is equal to $t+m$, with the mode index m directly mirroring the time laxity offered for preheating by the customer. Thus, the optimization that needs to be solved to minimize the expected cost for this is

$$C_q^{\text{preheat}}(m) = \min_{b_i(j)} \left[\sum_{j=t}^{t+m} \pi^p(j)b_i(j) \right] \quad \text{s.t.} \quad x_q^{\max} - \delta \leq x_q(t+m) \leq x_q^{\max} + \delta, \quad b_i(j) \in [0, 1] \quad (18)$$

where $\delta \geq 0$ should be picked such that at least one number in the set $[x_q^{\max} - \delta, x_q^{\max} + \delta]$ lies within the controllability subspace of the dynamics that govern $x_i(t+m)$. To write these dynamics in closed form, we assume that $x_i(t) = x_a(t)$, i.e., we assume that the unit stays off long enough when the building is vacant, and the temperature $x(t)$ has already reached the ambient temperature $x_a(t)$ when preheating starts. Then, we can write the temporal evolution of the expected value of $x_i(j)$ as

$$x_i(j) = c_i(j) + \sum_{w=t}^{j-1} (1 - k_q)^{j-w-1} W_q b_i(w) \quad (19)$$

with $c_i(j) = \sum_{w=t}^{j-1} (1 - k_q)^{j-w-1} (x_i(w)(1 - k_q) + k_q x_a(w))$. The $b_i(w)$'s are the decision variables.

In order to see how much benefit direct scheduling can provide to the aggregator, we need to compare the cost to the case when no preheating is exercised. In that case, heating starts at $t+m$, and it will take the TCL τ_q^{up} units of time to continuously work and get the temperature to x_q^{\max} , where

$$\tau_q^{up} \approx \frac{1}{k_q} \ln \left(\frac{\frac{W_q}{k_q}}{x_q^{\max} - \frac{W_q}{k_q} - x_a(t+m)} \right). \quad (20)$$

This assumes that the dynamics of the ambient temperature is much slower than that of the TCL. The first heating cycle right after turning on in the morning will have an expected

cost equal to

$$C_q^{\text{normal}} = \sum_{j=t+m}^{t+m+\tau_q^{up}-1} \pi^p(j)$$

for the aggregator. Thus, the utility of recruiting a TCL in cluster q in mode m at time t is given by:

$$U_q^t(m) = \max \{ C_q^{\text{normal}} - C_q^{\text{preheat}}(m), 0 \}$$

Note that providing any DLS service requires that the appliance is submetered. Here this is essential since the overall consumption of the device could be fairly increased and the consumer should not be billed accordingly.

IV. NUMERICAL CASE STUDY USING EV DATA

We simulate the interactions of one aggregator with a population of Plug-in Hybrid EVs arriving at random to receive level-1 battery charge (1.1 kW instantaneous rate). The arrival time, charge duration, and laxity data are taken from real PHEV charge events recorded and studied in [18]. This database includes 620 charge events that happen over the length of several months. However, we ignore the dates and treat the plug-in events as if they happened on the same day. The aggregator interacts with the wholesale energy market. Hourly LMPs for the one day of operation simulated here are taken from ISO New England's Maine load zone on September 1st, 2013 to noon of September 2nd.

We translate the absolute charge laxity values in [18] into *risk functions* by tuning the type γ_i in (5) such that it would not be individually rational for the customer to offer laxities that exceed their real departure time, given one sample of the daily incentive profile. For lack of any meaningful alternative, we take $r_q^t(m) = m$. Denote the mode corresponding to the maximum laxity that could be offered while allowing a timely departure for the customer by m_d . Then we assume that for each charging event in cluster q at time t , given sample incentives $\hat{I}_q^t(m)$, the risk type γ_i of the user i is such that

$$(m_d + 1)\gamma_i \geq \hat{I}_q^t(m_d + 1).$$

Next, in order to solve the optimization (13), we fit a uniform distribution on the customer types, resulting in

$$\gamma_i \sim U(0, 0.0721).$$

One can use the Gershgorin circle theorem to show that the optimization problem (13) is concave (with a negative semidefinite quadratic term). Incentives are designed for half-hour intervals, i.e., $t = 1, \dots, 48$ for one full day. Remember that the solution to this optimization is sub-optimal due to the single-crossing condition and the uniform prior assumed for the customer risk types γ_i .

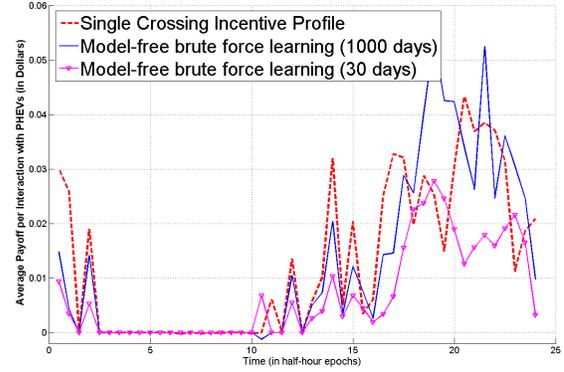


Fig. 1. Comparison of aggregator payoff when offering incentives designed through different schemes to a PHEV population

A. Comparison of Aggregator Payoff for Different Incentive Design Schemes

The average payoff (N_q^t) that the aggregator receives at different times of the day t when interacting with the above-described customers under the single-crossing incentives is shown in Fig. 1 (dashed red curve). To study the level of sub-optimality imposed by the single-crossing design constraint, we compare the outcome of the model-based incentive design technique to a brute-forced learning method, where several scenarios are generated for the daily incentive profile and tested on the population to observe the response. With more trials, the aggregator would find better incentive profiles. Fig. 1 compares the aggregator's profit with the single-crossing profile with that of the brute-force learning method after 30 and 1000 days. The reader can see that the learning method is performing better than the single-crossing profile after 1000 days, but not in 30 days.

B. Welfare Effects

We define the consumer savings through a DSM technique as the expected change in the monetary value of maximum utility across different alternatives for receiving electricity service. As explained before, since the basic service that is being received/provided through any demand management program is always the same as the normal operation mode of the grid today, i.e., there is no change in overall consumption and only the timing of consumption changes, we ignore the utility of receiving the standard service of electricity consumption in these calculations. The consumer savings by participating in the DLS program (under brute-force optimized incentives after 1000 days) is equal to $I_{q_i}^t(m_i) - R_i^t(m_i)$. Saving impacts are individually calculated for each charge request, and the sum of this value over the 620 charging events is shown in Fig. 2 as a function of time of plug-in.

Now assume that, instead of going through a intermediary node like the aggregator, the customer could have, through a hypothetical demand management scheme, directly interacted

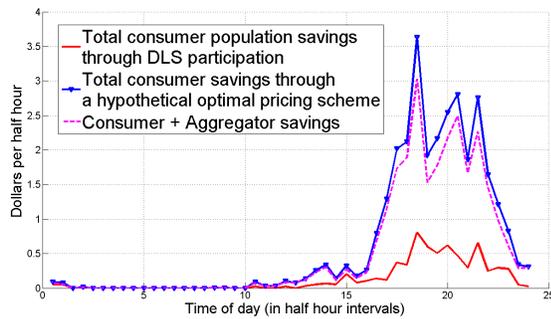


Fig. 2. Aggregate Consumer savings for the 620 Charging Events

with the electricity grid. Hence, customers could submit their consumption flexibility as bids to the Independent System Operator (ISO) that runs the market, and are billed based on optimal market clearing prices. This is the ultimate goal of dynamic pricing schemes and could realize the most efficient outcome of the market. Alternatively, one could imagine any scheme that uses the full potential of customers for load shifting, and relays all the wholesale market savings to them. For now, assume that today's regulated flat tariffs are designed such that customers are billed at the averaged wholesale cost of the electricity they consume (ignoring aggregator mark up). Then, the consumer savings through the implementation of such dynamic pricing program would have been equal to the recruitment utility, i.e., $U_q^t(m_d)$, assuming that customers would use their maximum possible laxity with no reservation when making individual scheduling decisions in response to a price (no commitment risk function). The value of this quantity is summed up across the population and is compared to customer savings gained from the DLS scheme in Fig. 2.

The considerable difference in consumer savings between the DLS program and that of optimal pricing is due to two factors: 1) the presence of an intermediary for-profit node, i.e., the aggregator; 2) the commitment risk function $R_i^t(m_i)$, preventing the customers from realizing all their load shifting potentials. To observe the welfare effects of the second phenomena, we show the sum of the aggregator and consumer savings under DLS in Fig. 2. The reader can observe that the overall decrease in the community's (aggregator + consumers) welfare is minimal. Remember that this decrease in welfare comes with the promise of reliability and controllability, which are the essential attributes that drive electricity market design problems and are required for safe grid operation.

V. CONCLUSIONS AND FUTURE WORK

Due to the popularity of direct control mechanisms in the power grid, designing appropriate incentives for customer participation in these programs seems like an inevitable issue. Here we approached this issue as a market design problem to trade flexibility with an aggregator of electricity services. Extension of this work to look at the possibility

of providing ancillary services through collective effort of several appliances, and to study competition between several aggregators that can serve the same population, is left to future work.

REFERENCES

- [1] C. Chen, S. Kishore, and L. V. Snyder, "An innovative rtp-based residential power scheduling scheme for smart grids," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*. IEEE, 2011, pp. 5956–5959.
- [2] K. Shaloudegi, N. Madinehi, S. Hosseini, and H. Abyaneh, "A novel policy for locational marginal price calculation in distribution systems based on loss reduction allocation using game theory," *Power Systems, IEEE Transactions on*, vol. 27, no. 2, pp. 811–820, 2012.
- [3] D.-C. Wei and N. Chen, "Air conditioner direct load control by multi-pass dynamic programming," *Power Systems, IEEE Transactions on*, vol. 10, no. 1, pp. 307–313, 1995.
- [4] P. Sanchez-Martin, G. Sanchez, and G. Morales-Espana, "Direct load control decision model for aggregated ev charging points," *Power Systems, IEEE Transactions on*, vol. 27, no. 3, pp. 1577–1584, 2012.
- [5] W. Su and M.-Y. Chow, "Performance evaluation of an eda-based large-scale plug-in hybrid electric vehicle charging algorithm," *Smart Grid, IEEE Transactions on*, vol. 3, no. 1, pp. 308–315, 2012.
- [6] M. Caramanis and J. Foster, "Management of electric vehicle charging to mitigate renewable generation intermittency and distribution network congestion," in *Decision and Control, Proceedings of the 48th IEEE Conference on*, 2009, pp. 4717–4722.
- [7] Y. He, B. Venkatesh, and L. Guan, "Optimal scheduling for charging and discharging of electric vehicles," *Smart Grid, IEEE Transactions on*, vol. 3, no. 3, pp. 1095–1105, 2012.
- [8] D. Wu, D. Aliprantis, and L. Ying, "Load scheduling and dispatch for aggregators of plug-in electric vehicles," *Smart Grid, IEEE Transactions on*, vol. 3, no. 1, pp. 368–376, 2012.
- [9] S. Chen, Y. Ji, and L. Tong, "Large scale charging of electric vehicles," in *Power and Energy Society General Meeting, 2012 IEEE*, 2012, pp. 1–9.
- [10] J. Tomi and W. Kempton, "Using fleets of electric-drive vehicles for grid support," *Journal of Power Sources*, vol. 168, no. 2, pp. 459 – 468, 2007. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378775307005575>
- [11] M. G. Vaya and G. Andersson, "Combined smart-charging and frequency regulation for fleets of plug-in electric vehicles," in *Power and Energy Society General Meeting, 2013 IEEE*, 2013.
- [12] C. Alvarez, R. Malhame, and A. Gabaldon, "A class of models for load management application and evaluation revisited," *Power Systems, IEEE Transactions on*, vol. 7, no. 4, pp. 1435–1443, 1992.
- [13] J. Ma, J. Qin, T. Salsbury, and P. Xu, "Demand reduction in building energy systems based on economic model predictive control," *Chemical Engineering Science*, vol. 67, no. 1, pp. 92 – 100, 2012, dynamics, Control and Optimization of Energy Systems. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0009250911005240>
- [14] K. ho Lee and J. E. Braun, "Development of methods for determining demand-limiting setpoint trajectories in buildings using short-term measurements," *Building and Environment*, vol. 43, no. 10, pp. 1755 – 1768, 2008. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0360132307002168>
- [15] Y. Ma, F. Borrelli, B. Hencsey, A. Packard, and S. Bortoff, "Model predictive control of thermal energy storage in building cooling systems," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, 2009, pp. 392–397.
- [16] M. Hasan, "Aggregator-assisted residential participation in demand response program," *Master's thesis, Virginia Polytechnic Institute and State University*, 2012.
- [17] R. Yin, P. Xu, M. A. Piette, and S. Kiliccote, "Study on auto-dr and pre-cooling of commercial buildings with thermal mass in california," *Energy and Buildings*, vol. 42, no. 7, pp. 967–975, 2010.
- [18] K. Kurani et al., "Learning from consumers: Plug-in hybrid electric vehicle (phev) demonstration and consumer education, outreach, market research program, volumes i and ii," *Institute of Transportation Studies, UC Davis, Research Report UCD-ITS-RR-10-21*, 2010.