Dynamic Incentive Design for Participation in Direct Load Scheduling Programs

Mahnoosh Alizadeh, *Student Member, IEEE*, Yuanzhang Xiao, *Student Member, IEEE*, Anna Scaglione, *Fellow, IEEE*, and Mihaela van der Schaar, *Fellow, IEEE*

***Abstract*—Interruptible Load (IL) programs have long been an accepted measure to intelligently and reliably shed demand in case of contingencies in the power grid. However, the emerging market for Electric Vehicles (EV) and the notion of providing non-emergency ancillary services through the demand side have sparked new interest in designing direct load scheduling programs that manage the consumption of appliances on a day-to-day basis. In this paper, we define a mechanism for a Load Serving Entity (LSE) to strategically compensate customers that allow the LSE to directly schedule their consumption, every time they want to use an eligible appliance. We study how the LSE can compute such incentives by forecasting its profits from shifting the load of recruited appliances to hours when electricity is cheap, or by providing ancillary services, such as regulation and load following. To make the problem scalable and tractable we use a novel clustering approach to describe appliance load and laxity. In our model, customers choose to participate in this program strategically, in response to incentives posted by the LSE in publicly available menus. Since 1) appliances have different levels of demand flexibility; and 2) demand flexibility has a time-varying value to the LSE due to changing wholesale prices, we allow the incentives to vary dynamically with time and appliance cluster. We study the economic effects of the implementation of such program on a population of EVs, using real-world data for vehicle arrival and charge patterns.**

***Index Terms*—Smart grids, energy, load management, retail market, economics.**

I. Introduction

UNDER pre-determined flat pricing tariffs, end-use customers are sheltered from the fluctuations of the wholesale price of electricity through a Load Serving Entity (LSE), leading to market inefficiencies [1]. The research on how to make electricity demand sensitive to the wholesale price of its supply can be divided into four major trends: 1) one option explored is to derive smart locational marginal pricing (LMP) strategies (see e.g., [2]–[4]) that capture the price-elasticity of a population of customers in the social welfare maximization solved by the system operator; 2) another line of research considers cooperating customers in a neighborhood that minimize their overall costs through decentralized iterative methods aimed at finding socially-optimal schedules, e.g., [5]–[7]; 3) a different direction is based on game-theory where the authors design billing mechanisms that allow for an incentive-compatible scheduling of the collective demand of a population of selfish customers, e.g., [8]–[12]; 4) last but not least, there is renewed interest in direct scheduling of the load of certain appliances by the LSE, to decrease market costs or to provide ancillary services. Interruptible Load (IL) programs started as a last-resort measure that allow for selective load shedding during contingencies, but recent papers have suggested more sophisticated schemes where the energy consumption of appliances is scheduled on a day-to-day basis to follow the variations of supply, e.g., [13]–[19]. These techniques could ease the transition towards day-to-day use of high levels of renewables in the grid. We refer to these schemes as Direct Load Scheduling (DLS).

Dynamic retail pricing methods are considered to be the most efficient and the least communication intensive option. However, there are certain factors that hinder the implementation of these programs: 1) the lack of reliable information on the price sensitivity of customers, especially since this sensitivity has a time and space in-homogeneous nature; 2) the lack of understanding of how, given the lack of perfect retail competition, a for-profit (and possibly regulated) LSE can be trusted to set fair dynamic prices, without exploiting the customers' lack of flexibility for certain tasks; 3) the strict reliability requirements of the grid, allowing only for small margins of operational error; 4) the public opinion considering electricity a basic need, requiring a price that is seemly and predictable well ahead of time. This expectation could prove to be problematic as renewable outputs are highly unpredictable.

DLS programs obtain a tight and reliable control the demand. This is considered a valuable trait in grid operations, especially when demand is used for non-emergency ancillary service provision, which requires a high-precision response with certain performances guarantees. DLS services can be coupled with any other Demand Response (DR) scheme to provide a smooth and reliable transition to active demand management in the grid. One important and well-studied aspect of this problem is *how* a certain authority can schedule the energy consumption of a large population of devices with a certain objective (module highlighted in grey in Fig. 1). While this part of the problem is of great importance, what is often neglected is *why* should a rational end-use customer voluntarily participate in such programs on a daily basis, as it could harm his/her privacy and

Manuscript received October 01, 2013; revised March 10, 2014; accepted July 28, 2014. Date of publication August 12, 2014; date of current version November 18, 2014. This work was supported by the Office of Electricity Delivery and Energy Reliability, Transmission Reliability Program of the U.S. Department of Energy, administered by the Lawrence Berkeley National Laboratory, and coordinated by the Consortium for Electric Reliability Technology Solutions. The guest editor coordinating the review of this manuscript and approving it for publication was Prof. Visa Koivunen.

M. Alizadeh and A. Scaglione are with the Department of Electrical Engineering, University of California Davis, Davis, CA 95616 USA (e-mail: malizadeh@ucdavis.edu).

Y. Xiao and M. van der Schaar are with the Department of Electrical Engineering, University of California Los Angeles, Los Angeles, CA 90095 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSTSP.2014.2347003

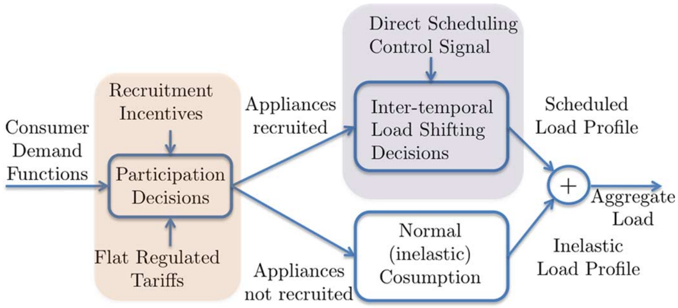

Fig. 1. Modules involved in Direct Load Scheduling programs.

comfort levels. Currently, IL programs provide small fixed one-time payments to customers, with the guarantee that load shedding will only happen infrequently and only during emergencies, and the notion that customers are working towards social good. However, for day-to-day direct scheduling, the mere promise of a "green future" may prove to be inadequate to incentivize high levels of customer participation.

A. Contribution

With this motivation, here we define a market that allows the LSE to strategically provide incentive payments to appliance owners that provide DLS services (module highlighted in pink in Fig. 1). Incentives change dynamically with grid conditions, and are provided every time an eligible appliance is plugged-in and needs to consume energy. To ensure fairness, we assume that incentives are posted in dynamically-updated publicly-available tables, and customers providing the same service receive the same payment. Recruited appliances release the control over their demand to the LSE, with the promise that they will receive the service they expect within their proposed laxity limits, but with the exact timing of energy delivery at the discretion of the LSE. After observing available incentives, owners of potential participating appliances strategically make contribution choices, i.e., they decide whether to participate in the DLS program and how much laxity they wish to offer to the LSE. Individual appliances have different flexibility, and customers can make different choices in how much laxity to offer. Thus, incentive payments vary accordingly to encourage higher contributions and maximize the profits of the LSE, soliciting the optimal levels of flexibility from appliances.

B. Related Works

Most works on DLS takes customers participation as a given. In [13], the customer response to price is modeled as a sigmoid function while the paper focuses on competitive pricing by a set of LSEs. A somewhat more related study is in [20], which proposed an online mechanism to incentivize Electric Vehicle (EV) owners to reveal their true flexibility to a scheduler by punishing misrepresentation under a soft Quality of Service (QoS) model and burning the extra units of power if necessary. Our model is different since it captures various types of appliances, provides hard QoS, and uses detailed costs based on electricity market structure. More recently, the authors in [21], [22] designed mechanisms to incentivize end-users to reveal their true laxities to an LSE. The load model used in both works assumes that each customer provides the LSE with the amount of energy that he/she needs by different deadlines, while abstracting out specific consumption profiles and constraints of different appliances inside each residence. This is different from our approach, which captures the specific characteristics of different appliances and their intrinsic flexibilities through a reduced-order load model defined in Section II. The authors in [23] provide different prices to each appliance that arrives based on the current system state and assuming that a welfare minimizing agent is in charge of load scheduling (payment of each user is minimized given the current load). This is different from our model, since the LSE is a profit maximizing entity here.

Building on the basic concepts proposed in our previous conference paper [24], here we expand the framework to study how an LSE can use recruited appliances to offer ancillary services on top of minimizing its energy market costs.

C. Synopsis

In Section II-A, we present our load clustering approach for recruiting heterogeneous flexible appliance in the DLS program. Section II-B discusses our assumptions. Section II-C describes the consumer choice model, with the LSE profit maximization problem to determine the optimal incentives following in Section II-D. We propose two approaches to solve the incentive design problem in Section II-D1 (Bayesian approach), and Section II-D2 (frequentist approach). Solving for the optimal incentives requires the LSE to know the utility of recruiting appliances in the DLS program, which we calculate for the case of deferrable loads in Section III. Finally, our numerical test cases are described in Section V.

II. Model

To make participation incentives commensurate with the service that is provided by the recruited appliances to the direct load scheduler, the first question that needs to be answered is how can the incentive designer sift through the value offered by large heterogeneous populations of appliances with variable load flexibilities in a scalable fashion.

A. Discretizing Flexibility

To put a value on the heterogeneous load flexibility offered by recruited appliances, the first step is to capture this flexibility mathematically. To give a simple example of how we approach this problem, let us look at an ideal battery load that is recruited at time t and is characterized by 4 parameters: a charge capacity, E, a rate constraint $g_{\max}$, a slack time to fully charge, μ, and the initial state of charge, S_a. We define slack time as the amount of time left between the end of an electric job and its deadline if the request is served immediately after arrival at the highest rate. Denote the state of charge of the battery at time ℓ by $s(\ell)$. Then, if the battery is recruited at time t, the possible load shapes $L(\ell)$ of this appliance in time is given by the set:

$$\mathcal{L}_{\boldsymbol{\theta}}(t) = \Bigg\{ L(\ell) | L(\ell) = \dot{\mathrm{s}}(\ell), \mathrm{s}(t) = \mathrm{S}_a, \dot{\mathrm{s}}(\ell) \in [0, g_{\max}], \\ \mathrm{s}\left(t + \frac{E - S_a}{g_{\max}} + \mu\right) = E, 0 \leq \mathrm{s}(\ell) \leq \mathrm{E} \Bigg\}. \quad (1)$$

Thus, the battery load's consumption characteristics are fully described by the parameters $\boldsymbol{\theta} = (S_a, E, g_{\max}, \mu)$. We divide these parameters into two vectors, $(\mathbf{v}, \mathbf{f})$ such that:

$$\mathbf{v} = (S_a, g_{\max}, E), \quad \mathbf{f} = \mu. \quad (2)$$

The first vector captures the inherent and unmodifiable properties of the request, and the second vector describes the level of flexibility that the customer *voluntarily* commits to provide along with this request. We assume that this division is possible for any type of flexible energy request. The vector $\mathbf{v}$ could describe the load shape of a dishwasher once on, or the charge duration and rate constraints for an Electric Vehicle (EV). On the other hand, $\mathbf{f}$ describes how *lax* the energy request is. For example, $\mathbf{f}$ can describe the slack time that accompanies a deferrable load such as dishwasher cycle, or the comfort band that dictates how far the temperature controlled by a Thermostatically Controlled Load (TCL) can swing around the desired value. It is useful to highlight the difference between these dependencies $\mathcal{L}_{\boldsymbol{\theta}}(t) \equiv \mathcal{L}^{\mathbf{v}}_{\mathbf{f}}(t)$ as changing $\mathbf{v}$ produces an entirely different set, while changing $\mathbf{f}$ has the effect of enlarging the set. In particular, for the battery load in (1), if $\mu' \geq \mu$, then $\mathcal{L}^{\mathbf{v}}_{\mu}(t) \subseteq \mathcal{L}^{\mathbf{v}}_{\mu'}(t)$.

Note that the parameters $(\mathbf{v}, \mathbf{f})$, e.g., $(S_a, E, g_{\max}, \mu)$ for the battery load, can take infinitely many values when dealing with a population of heterogeneous appliances. Thus, evaluating this set and the utility associated with having such flexibility at hand for every single appliance is impossible for the LSE.

To lower computational and communication requirements of the DLS program, we choose to quantize the vectors $\mathbf{v}$ and $\mathbf{f}$, a strategy most common in real-time and multimedia applications. This idea was first introduced in [25] to control large populations of appliances, which generalized the idea of quantizing the temperature range in bins proposed by Chong in [26] to analyze populations of interruptible TCLs. Consequently, we assume that eligible appliances can only choose the characteristic vector $\mathbf{v}$ from a finite codebook $\mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_Q\}$. Depending on the choice of $\mathbf{c}_q$, we assign appliances to *clusters*, indexed by

$$q \in \mathcal{Q} = \{1, 2, 3, \ldots, Q\}. \tag{3}$$

As an example, we can quantize the initial state of charge as well as the amount of charge required by EVs that receive home charging (1.1 kW rate) to 0.55 kWh intervals. Then, separate clusters are assigned to EVs that arrive with an empty battery and require 0.55 kWh, 1.1 kWh, 1.65 kWh, and so on.

We also quantize the level of flexibility that accompanies a request, i.e., the vector $\mathbf{f}$. We assume that flexibility offered by appliances can be characterized by a discrete set of options, e.g., deadlines rounded to the hour. For appliances in cluster q, this choice is fully described by an index

$$m \in \mathcal{M}^q = \{0, 1, 2, \ldots, M^q\}, \tag{4}$$

and is referred to as the appliance's *mode*. Mode $m = 0$ corresponds to the default selection for an appliance that does not participate in the DLS program.

We can now write the load flexibility of *any* ideal battery in cluster q operated under mode m at time t as,

$$\mathcal{L}^q_m(t) \approx \Bigg\{ L(\ell) | L(\ell) = \dot{\mathrm{s}}(\ell), \mathrm{s}(t) = \mathrm{S}^q_a, \dot{\mathrm{s}}(\ell) \in \{0, \ldots, g^q_{\max}\}, \mathrm{s}\left(t + \frac{E^q - S^q_a}{g^q_{\max}} + m\right) = E^q, 0 \leq \mathrm{s}(\ell) \leq \mathrm{E}^q \Bigg\}. \tag{5}$$

Note that $\mathcal{L}^q_0(t) \subseteq \mathcal{L}^q_1(t) \cdots \subseteq \mathcal{L}^q_{M^q}(t)$, and that the set $\mathcal{L}^q_0(t)$ is a singleton, $\mathcal{L}^q_0(t) = \{L^q_0(t)\}$.

Given that $\mathcal{Q}$ and $\mathcal{M}^q$ are finite sets, the incentive design problem will now simply depend on the size of these quantization sets and not on the population size that the LSE is dealing with, allowing our DLS incentive design problem to be scalable in terms of communication and computation.

B. Basic Assumptions

The LSE wants to maximize its profit and the customer wants the service mode with the economic value. The functions that reflect their trade offs will be introduced in detail later in (9) and (11) respectively. In general, we assume that, given no incentives, the customers prefer $\mathcal{L}^q_m(t)$ to $\mathcal{L}^q_{m'}(t)$ iff $m < m'$, i.e., the customer chooses mode $m = 0$ if no incentives are available. Posting the incentives modifies the economic value of each mode, and introduces a new partial ordering across the sets $\mathcal{L}^q_m(t)$ for the customer. On the other hand, the LSE chooses the incentives for each m based on the total order introduced by market prices on the members of the set $\mathcal{L}^q_m(t)$, leading to a unique load profile $L^q_m(t)$ for each (q, m) that is most advantageous to the LSE. Here we assume a tariff structure imposed by the LSE that is backwards compatible with current practices. Specifically, customers are by default on a flat tariff. Thus, without DLS, the mode customers pick by default is mode $m = 0$, which is a singleton $L^q_0(t)$. For instance, for the ideal battery, the customer will choose $L^q_0(t) = g^q_{\max}$ for a time duration $T^q = (E^q - S^q_a)/g^q_{\max}$. We call the *recruitment utility* under mode m for the LSE (cf. Section II-D) the wholesale market savings or payments realized through serving a request through any of the profiles in $\mathcal{L}^q(t) \triangleq \bigcup_{m=0}^{M^q} \mathcal{L}^q_m(t)$ other than $L^q_0(t)$. Specifically, $U^q_m(t)$, the mth entry of the vector:

$$\mathbf{u}^q(t) = [U^q_0(t), U^q_1(t), \ldots, U^q_{M^q}(t)]^T \tag{6}$$

denotes the maximum savings/payment that the LSE can accrue in the wholesale market from choosing a load profile in $\mathcal{L}^q_m(t)$. Clearly, $U^q_0(t) = 0$ for all q. If $U^q_m(t) > 0$ for $m > 0$, then the LSE can offer discounts to the customers picking mode $m > 0$, as long as the revenue from $U^q_m(t)$ more than compensates for the paid incentives[1].

We denote the incentives available for cluster q at time t as

$$\mathbf{x}^q(t) = [x^q_1(t), x^q_2(t), \ldots, x^q_{M^q}(t)]^T, \tag{7}$$

where $x^q_m(t)$ denotes the incentive available to an appliance in cluster q that picks mode m and is recruited at time t.

To solve for $\mathbf{x}^q(t)$, there are two fundamental aspects that need to modeled. First, how reluctant customers are in committing to provide DLS services, i.e., what is the minimum payment they need to receive to commit. Second, what is $\mathbf{u}^q(t)$, which in turn determines how much is the LSE willing to discount the customer for offering the flexibility of appliances in cluster q to be operated under mode m, i.e., $x^q_m(t)$. We start by answering the first question in Section II-C and follow with addressing the second in Section II-D.

We now state the assumptions under which our proposed paradigm is defined (1, 3, 5, 6) and is analyzed (2, 4, 7):

Assumption 1: Appliances that decide not to participate in the DLS program are provided with the standard service of electricity and are billed on a flat rate. These 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;

Assumption 2: Each customer can only interact with one LSE, i.e., the market for obtaining DLS services is a monopoly

[1]Note that Time of Use (TOU) tariffs can be interpreted as a regularly occurring discounts that appears at off peak hours for all appliances in all clusters. Hence, the retail pricing model is sufficiently general to capture current practices that go beyond flat tariffs.

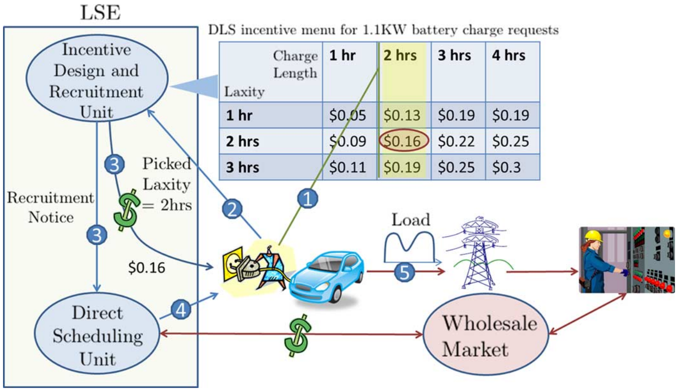

Fig. 2. Visual representation of the interactions between an EV that needs 2 hours of charge at 1.1 kWs and the DLS program. The order in which interactions happen are displayed on the arrows. Displayed incentive values are not optimized.

(extension to cover competition between LSEs to recruit appliances left to future work);

Assumption 3: The LSE's revenue from recruiting each appliance is additive and independent of other appliances;

Assumption 4: A customer's initial choice to use an appliance is perfectly inelastic and not affected by incentives;

Assumption 5: The LSE's load or ancillary service capacity offers do not affect the wholesale market clearing prices;

Assumption 6: The LSE has access to ex-ante forecasts of expected wholesale prices;

Assumption 7: Customers do not have any private information regarding wholesale market prices and are consequently myopic in deciding whether to participate in the DLS program[2].

We use the first four assumptions in the next Section, while Assumptions 5 and 6 simplify derivations of the expression of the LSE's utility when recruiting appliances in the DLS program in Section III.

C. Individual Appliance DLS Commitment Problem

A customer that needs to use an appliance belonging to cluster q would have two choices: 1) run the appliance under mode $m = 0$, and pay for the consumption at flat prices; 2) participate in the DLS program and make a commitment that the request will provide a certain level of slack time described by the mode $m > 0$. To make this decision, the customer looks at the current incentives posted by the LSE (see Fig. 2). If the customer chooses to take any incentive, he/she is bound to execute the consumption decisions dictated by the LSE, with deviations subject to penalties[3].

As previously mentioned, the first important effect that needs to be quantified by our model is the amount of incentive required for a customer to join the DLS program, and further, commit to a certain mode. For deferrable loads, which are the main focus of Section III, this commitment means that the appliance consumption request has to remain *active* until a certain deadline (the appliance should remain plugged in and available to be controlled). To capture this, one could think that committing to stay in the system presents a certain amount of *risk* for the customer, mainly due to an emergency/unplanned event that may require an earlier departure, and the penalties charged for not sticking to his/her commitment.

As commonly observed in many real world scenarios that involve human factors, the customer's participation decision, and consequently his/her choice of mode m, is most likely mediated by bounded rationality, and may not be correctly captured by a rational agent model. However, we provide such a model next, with the caveat that, even if the customer is not fully rational, our framework for incentive design can rely on model-free learning approaches (discussed in Section II-D2).

1) Rational Agent Decision Making Model: We model the customer as a utility-maximizing rational agent, and assume that the utility gained by the customer from operating appliance i in cluster q at time t under mode m is denoted by $V_m^{i,q}(t)$, which includes three terms:

1) the incentive $x_m^q(t)$ available for mode m;
2) the commitment risk (disutility) $r_m^{i,q}(t)$ associated with committing to receive service under mode m at time t, and relinquishing control over the consumption of the appliance to the LSE. This is modeled through a privately known non-negative vector
$$\mathbf{r}^{i,q}(t) = [r_0^{i,q}(t), r_1^{i,q}(t), \ldots, r_{M^q}^{i,q}(t)], \tag{8}$$
with $r_0^{i,q}(t) = 0$, i.e., there is no risk when the customer decides not to participate and picks mode $m = 0$. Without loss of generality, we order the modes from low to high risk, having monotonically non-decreasing risks with modes. The risk functions are also scaled to have the same unit as the monetary incentives $x_m^q(t)$;
3) the utility of receiving the standard service of using electricity and finishing a job.

The last term is a constant, since we assume consumption will eventually happen for every request, either through the DLS program under some mode $m \geq 1$, or through the standard service model of the power grid ($m = 0$). 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 and consider the following additive utility:
$$V_m^{i,q}(t) = x_m^q(t) - r_m^{i,q}(t). \tag{9}$$

Upon observing the incentives $x_m^q(t)$ posted by the LSE, the customer solves the following utility maximization problem to determine the best mode m for operating appliance i:
$$\max_{m \in \mathcal{M}^q} \quad V_m^{i,q}(t). \tag{10}$$

If no mode $m \geq 1$ offers an incentive that is marginally higher than the commitment risk $r_m^{i,q}(t)$ for that mode, the customer will not participate in the program (mode $m = 0$). In case of a tie, the customer picks the lowest risk mode.

The commitment risk captures the smallest price at which the customer is willing to sell DLS service under mode m and it is similar to *reservation prices* commonly used in microeconomics theory. Reservation prices vary between individuals according to their disposable income and how much they value the good they are offering. Here, this value could be a function of the task deadlines and of the penalties expected if they accept discounts for more slack time than they can provide and cannot honor their commitment.

[2]Assumption 7 can be replaced with a guarantee that payoffs will diminish as the customer holds back service that he/she could have provided. For deferrable loads with m corresponding to one time unit of slack time, this means $x_q^t(m) \leq x_q^{t-1}(m+1)$, i.e., highest incentive is offered when the deadline is farthest away from the commitment time. Consequently, the customer will either participate in the DLS program at the time of request arrival, or not at all.

[3]In this paper, we ignore the problem of designing the penalties and assume they are high enough that violations are negligible.

Note that the values of $r_m^{i,q}(t)$ are only privately-known to the customer. Thus, the LSE cannot deterministically model the outcome of (10) in response to different incentive vectors, and needs to design the incentives and maximize its profit with incomplete information. We address this problem next.

D. The LSE Problem

The LSE updates the incentives $\mathbf{x}^q(t)$ dynamically to maximize the expected profits of running the DLS program. We consider only profits obtained by using the demand flexibility of recruited appliances. They can be accrued in two possible ways. One is to minimize the costs of interacting with the wholesale electricity market, from which the LSE buys bulk energy to serve the load. The second one is by selling ancillary service capacity to the wholesale reserve market. We will elaborate more on the nature of these markets later in Section III. For now, we take the *DLS market utility* as the economic benefit that the LSE *expects* to gain in the wholesale market from recruiting an appliance in cluster q under mode m at time t as known, and denote it by $U_m^q(t)$, with $U_0^q(t) = 0$. Due to the seasonal nature of wholesale market prices, these expected utilities, as well as the incentives designed based on them, could exhibit a cyclic behavior and may only be periodically updated. As mentioned in Assumption 3, we define the appliance recruitment utility as additive and independent. Thus, we can focus on the profits gained from interacting with each individual appliance when designing the incentives. Since recruiting an appliance from cluster q in mode m at time t presents a cost equal to $x_m^q(t)$, the profit of the LSE from this recruitment is $U_m^q(t) - x_m^q(t)$.

Let $E_m^{i,q}(\mathbf{x}^q(t);t) = 1$, if customer i in cluster q picks mode m at time t and $E_m^{i,q}(\mathbf{x}^q(t);t) = 0$ else. Also, let $\mathcal{P}^q(t)$ be the population of DLS subscribers at time t in cluster q, with a size $|\mathcal{P}^q(t)|$. We assume $|\mathcal{P}^q(t)|$ is known to the LSE as the customers that subscribe have to refresh the incentive table and can, therefore, be counted. The LSE ex-post profit is:

$$Y^q(\mathbf{x}^q(t);t) = \sum_{m\in\mathcal{M}^q} (U_m^q(t) - x_m^q(t)) \sum_{i\in\mathcal{P}^q(t)} E_m^{i,q}(\mathbf{x}^q(t);t). \tag{11}$$

Remember that after observing the available incentives $x_m^q(t)$ posted by the LSE, each participating appliance endogenously chooses their mode m, through (10). Specifically, $E_m^{i,q}(\mathbf{x}^q(t);t) = 1$ happens if the following conditions are met:

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

$$x_m^q(t) - r_m^{i,q}(t) \geq 0, \tag{12}$$

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

$$x_m^q(t) - r_m^{i,q}(t) \geq x_{m'}^q(t) - r_{m'}^{i,q}(t), \ \ \forall m' \in \mathcal{M}^q. \tag{13}$$

Thus, ideally, the LSE would like to solve the following optimization problem to maximize his profits from interacting with all cluster q appliances:

$$\begin{aligned} \max_{\mathbf{x}^q(t)} \quad & \sum_{m\in\mathcal{M}^q} (U_m^q(t) - x_m^q(t)) \sum_{i\in\mathcal{P}^q(t)} E_{m_i}^{i,q}(\mathbf{x}^q(t);t) \\ \text{s.t.} \quad & x_{m_i}^q(t) - r_{m_i}^{i,q}(t) \geq 0, \\ & x_{m_i}^q(t) - r_{m_i}^{i,q}(t) \geq x_{m'}^q(t) - r_{m'}^{i,q}(t), \ \ \forall m' \in \mathcal{M}^q. \end{aligned} \tag{14}$$

However, note that the LSE can at best have access to statistical information on the customers' preferences, i.e., $r_{m'}^{i,q}(t)$, and can only maximize the *expected profit*. Let the $P_m^{i,q}(\mathbf{x}^q(t);t)$ denote the probability of this $E_m^{i,q}(\mathbf{x}^q(t);t) = 1$, i.e.,:

$$P_m^{i,q}(\mathbf{x}^q(t);t) = \mathbb{E}\{E_m^{i,q}(\mathbf{x}^q(t);t)\}. \tag{15}$$

The expected profit per appliance is:

$$\rho^q(\mathbf{x}^q(t);t) = \frac{\mathbb{E}\{Y^q(\mathbf{x}^q(t);t)\}}{|\mathcal{P}^q(t)|}. \tag{16}$$

We define the *mode selection average probability* and *mode selection average probability vector* as:

$$P_m^q(\mathbf{x}^q(t);t) = \frac{\sum_{i\in\mathcal{P}^q(t)} P_m^{i,q}(\mathbf{x}^q(t);t)}{|\mathcal{P}^q(t)|} \tag{17}$$

$$\mathbf{p}^q(\mathbf{x}^q(t);t) = [P_0^q(\mathbf{x}^q(t);t), \ldots, P_{M^q}^q(\mathbf{x}^q(t);t)]^T. \tag{18}$$

The LSE only needs information about these average quantities to maximize the expected profit. In fact:

$$\begin{aligned} \rho^q(\mathbf{x}^q(t);t) &= \sum_{m\in\mathcal{M}^q} (U_m^q(t) - x_m^q(t)) P_m^q(\mathbf{x}^q(t);t), && (19) \\ &= (\mathbf{u}^q(t) - \mathbf{x}^q(t))^T \mathbf{p}^q(\mathbf{x}^q(t);t), && (20) \end{aligned}$$

where he definition of $\mathbf{u}^q(t)$ and $\mathbf{x}^q(t)$ are in (6) and (7). Thus, the LSE designs the incentives at time t by solving

$$\max_{\mathbf{x}^q(t)\succeq \mathbf{0}} (\mathbf{u}^q(t) - \mathbf{x}^q(t))^T \mathbf{p}^q(\mathbf{x}^q(t);t) \tag{21}$$

Note that solving (21) requires the LSE to model how $P_m^q(\mathbf{x}^q(t);t)$ changes with the incentive profile $\mathbf{x}^q(t)$. It is also a non-linear problem, given that $\mathbf{p}^q(\mathbf{x}^q(t);t)$ is a non-linear vector function of the incentive vector $\mathbf{x}^q(t)$. We propose two different approaches next.

Remark II.1: Clearly, the profit maximizing incentives designed through solving (21) *are only optimal is the LSE is operating in a monopolistic retail market, where the only options available to the customer are to either accept the flat regulated tariffs or opt into the direct load scheduling program for a time-varying discount on the regulated tariffs. Naturally, as we will see in the numerical experiment, the retail market will not operate at its most efficient point with the lack of competition. However, an extension of this problem to capture retail competition is not trivial and it is one that we plan to address in future work.*

1) Bayesian Approach: Given statistically learned prior information on the risk levels, the LSE can view the incentive design problem as being close to that of optimal Bayesian unit-demand pricing. However, what makes the problem difficult is that the customers' risk levels for different modes cannot be assumed to be independent values, given that increasing flexibility amounts to expanding the set $\mathcal{L}_m^q(t)$. That means that risk levels that a customer perceives for committing to the program with a single appliance under different modes are correlated. Clearly, offering a higher slack time entails an *additional risk* over that of offering a lower slack time, and these variables cannot be considered independent.

The first step to model the correlation between the risk values for different modes is to define how the LSE parametrizes the risk functions $r_m^{i,q}(t)$. For appliance cluster q, we choose

$$r_m^{i,q}(t) = \gamma_i j_m^q(t), \tag{22}$$

where γ_i is a task-specific, non-negative continuous random variable representing the so-called *type* of an individual task i, and the deterministic variables $j_m^q(t)$ quantify the general attitude of customers using appliances in cluster q towards committing to mode m at time t, with $j_0^q(t) = 0$. The interpretation could be that the probability of an emergency event happening is homogeneous across the population. However, how customers *feel* about this emergency event (value of lost load) varies and is mirrored in their type. To ensure that this parameterization of the risk function is not unrealistic, the LSE could suggest this specific structure as default to consumers when they pick their risk functions. We assume that the LSE has access to statistical priors of the γ_i's, with $F_\gamma^q(g)$ denoting the cumulative distribution function of γ for cluster q.

Remark II.2: The reader can note the incentive design problem (21) *can be solved separately for different values of t and q. Consequently, for brevity of notation in this Section, hereafter we drop the indices q and t and write* (21) *as:*

$$\max_{\mathbf{x}\succeq\mathbf{0}} (\mathbf{u}-\mathbf{x})^T \mathbf{p}(\mathbf{x}) \tag{23}$$

From (12) and (13), the customer picks mode $m \geq 1$ iff:

$$\mathbf{IR}: \gamma_i \leq \frac{x_m}{j_m} = \frac{x_m - x_0}{j_m - j_0}, \tag{24}$$

$$\mathbf{IC}1: \gamma_i \leq \frac{x_m - x_{m'}}{j_m - j_{m'}}, \ \ 1 \leq m' < m, \tag{25}$$

$$\mathbf{IC}2: \gamma_i \geq \frac{x_{m'} - x_m}{j_{m'} - j_m}, \ \ m' > m, \tag{26}$$

so we should have $\gamma_i \in [l_m(\mathbf{x}), h_m(\mathbf{x})]$, with

$$h_m(\mathbf{x}) = \min \left\{ \left. \frac{x_m - x_{m'}}{j_m - j_{m'}} \right|_{0 \leq m' < m} \right\}, \tag{27}$$

$$l_m(\mathbf{x}) = \max \left\{ \left. \frac{x_{m'} - x_m}{j_{m'} - j_m} \right|_{m' > m} \right\}. \tag{28}$$

which gives,

$$P_m(\mathbf{x}) = F_\gamma(h_m(\mathbf{x})) - F_\gamma(l_m(\mathbf{x})). \tag{29}$$

However, due to the absence of any natural ordering, these constraints will render the optimization problem (21) rather complex. Thus, next, we will impose a design constraint that ensures that *local incentive compatibility* 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 $m' < m$.

2) Design Constraint 1: (Single-Crossing Incentive Profile): we will design the incentive profile such that $\forall m \in \mathcal{M}$, the ratio $(x_{m+1} - x_m)/(j_{m+1} - j_m)$ is non-increasing, i.e., incentives grow slower than risks for higher modes m.

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 LSE. However, it considerably lowers the numerical effort required to solve (21). To see a numerical study of the level of sub-optimality introduced by this constraint, see [24]. The main advantage of this design constraint is that it leads to a significant simplification for the case of uniform distribution for the γ_i, discussed in the following lemma, which we use later for a more general numerical approach to solve (23).

Let ∂ denote the first difference operator, i.e.,

$$\partial x_m = x_m - x_{m-1}, \tag{30}$$

with $x_0 \equiv 0$.

Lemma II.3: Under design constraint 1, if $F_\gamma(g)$ is a uniform distribution over $[0, \gamma_{\max}]$, with $x_1/j_1 \leq \gamma_{\max}$,

$$\mathbf{p}(\mathbf{x}) = -\frac{1}{\gamma_{\max}} \mathbf{Q}\mathbf{x}, \tag{31}$$

where

$$\mathbf{Q} = \begin{bmatrix} \phi_1 & \frac{1}{\partial j_2} & & & & \\ \frac{1}{\partial j_2} & \phi_2 & \frac{1}{\partial j_3} & & & \\ & \frac{1}{\partial j_3} & \phi_3 & \frac{1}{\partial j_4} & & \\ & & & & \ddots & \frac{1}{\partial j_M} \\ & & & & \frac{1}{\partial j_M} & \phi_M \end{bmatrix}, \tag{32}$$

with $\phi_m \triangleq -1/\partial j_m - 1/\partial j_{m+1}$.

Proof: If the incentive profile is single-crossing, it is easy to verify that customer i will pick mode $m \geq 1$ simply iff

$$\frac{x_{m+1} - x_m}{j_{m+1} - j_m} \leq \gamma_i \leq \frac{x_m - x_{m-1}}{j_m - j_{m-1}}, \tag{33}$$

where the right hand inequality ensures (24) and (25), and the left hand inequality ensures (26). To keep expression (33) compact, we use a dummy mode $m = M+1$, with $x_{m+1} = x_m$, and $j_{m+1} = j_m + 1$.

Consequently, the probability that a customer picks mode m is response to the incentives is,

$$P_m(\mathbf{x}) = \frac{1}{\gamma_{\max}} \left(\frac{\partial x_m}{\partial j_m} - \frac{\partial x_{m+1}}{\partial j_{m+1}} \right), \tag{34}$$

as long as all of the ratios $\partial x_m / \partial j_m$ fall in the support of γ_i. Writing this equation in matrix form gives (31). ∎

Given this, the expected profit we want to maximize in (23) can be written as:

$$\rho(\mathbf{x}) = (\mathbf{u}-\mathbf{x})^T \mathbf{p}(\mathbf{x}) = -\frac{1}{\gamma_{\max}} (\mathbf{u}-\mathbf{x})^T \mathbf{Q}\mathbf{x}, \tag{35}$$

This allows us to state the following result:

Proposition II.4: Under Design constraint 1, if γ_i introduced to define the random customer risk r^i in (22) *is a Uniform distribution, then* (23) *has the following standard quadratic programming (QP) form:*

$$\begin{aligned} \max_{\mathbf{x}\succeq\mathbf{0}} \quad & \mathbf{x}^t \mathbf{Q}\mathbf{x} + \mathbf{q}^T \mathbf{x} \\ \text{s.t.} \quad & \mathbf{Q}\mathbf{x} \preceq 0, \\ & \partial \mathbf{x} \succeq 0, \\ & \mathbf{b}(r)\mathbf{x} \leq \gamma_{\max}, \end{aligned} \tag{36}$$

where we have where

$$\mathbf{b}(r) = \left[\frac{1}{\partial j_1}, \overbrace{0, 0, \ldots, 0}^{M-1} \right] \tag{37}$$

$$\mathbf{q} = -\mathbf{u}^T\mathbf{Q} \tag{38}$$

The first constraint ensures that the optimized incentives are single-crossing and the second ensures that the incentives are non-increasing with modes.

Corollary II.5: the optimization problem in (36) *convex.*

Proof: The constraints are all linear. We only need to check whether the objective function is concave, or equivalently, whether $\mathbf{Q}$ is a negative semidefinite matrix.

Denoting the elements of the matrix $\mathbf{Q}$ as Q_{ij}, the Gershgorin circle theorem tells us that every eigenvalue of $\mathbf{Q}$ lies within at least one of the Gershgorin discs $D(Q_{ii}, \sum_{j\neq i}|Q_{ij}|)$, where $D(Q_{ii}, \sum_{j\neq i}|Q_{ij}|)$ is a closed disc on the real axis that centers at q_{ii} and has a radius $\sum_{j\neq i}|Q_{ij}|$. Hence, since all the Gershgorin discs of $\mathbf{Q}$ lie on the non-positive half of the real axis, we have that all the eigenvalues are non-positive, and $\mathbf{Q}$ is negative semidefinite. ∎

For a general choice of $F_\gamma(g)$, the Bayesian incentive design optimization (21) may not necessarily be convex. Next, we provide an approximate solution for sufficiently smooth choices of $F_\gamma(g)$, without any convergence guarantees. We choose to use a first-order Taylor approximation of $F_\gamma(g)$, which is inexpensive to evaluate. We can use such approximation to sequentially solve for $\mathbf{x}$ through a sequence of approximate subproblems. Assume that we have an estimate of $\mathbf{x}$ at iteration k, denoted $\mathbf{x}^{(k)}$. Then, at iteration $k+1$, we can approximately write the probability of a customer picking mode m, i.e., the probability that γ_i satisfies (33), as

$$\begin{aligned} P_m(\mathbf{x}) &= F_\gamma\left(\frac{\partial x_m}{\partial j_m}\right) - F_\gamma\left(\frac{\partial x_{m+1}}{\partial j_{m+1}}\right) \\ &\approx \left(\frac{\partial x_m}{\partial j_m} - \frac{\partial x_{m+1}}{\partial j_{m+1}}\right) f_\gamma(y_m^{(k)}) \end{aligned} \tag{39}$$

where $f_\gamma(g)$ denotes the probability density function (PDF) corresponding to $F_\gamma(g)$, and

$$y_m^{(k)} = \frac{1}{2}\left(\frac{\partial x_m^{(k)}}{\partial j_m} + \frac{\partial x_{m+1}^{(k)}}{\partial j_{m+1}}\right) \tag{40}$$

which can be calculated after iteration k and be taken as a constant in iteration $k+1$. Introducing (cf. (32)):

$$\begin{aligned} \mathbf{Q}'_k &= \mathbf{Y}_k\mathbf{Q}, \quad \mathbf{q}'_k = -\mathbf{u}^T\mathbf{Q}'_k \\ \mathbf{Y}_k &= \mathrm{Diag}\left[f_\gamma(y_1^{(k)}), f_\gamma(y_2^{(k)}), \ldots, f_\gamma(y_M^{(k)})\right] \end{aligned} \tag{41}$$

Equation (35) becomes:

$$\rho_k(\mathbf{x}) = -(\mathbf{u}-\mathbf{x})^T\mathbf{Q}'_k(r)\mathbf{x}, \tag{42}$$

and the quadratic subproblem at iteration $k+1$ is:

$$\begin{aligned} \mathbf{x}^{(k+1)} &= \mathrm{argmax}_{\mathbf{x}\succeq\mathbf{0}} \quad \mathbf{x}^T\mathbf{Q}'_k\mathbf{x} + {\mathbf{q}'_k}^T\mathbf{x} \\ \text{s.t.} \quad & \mathbf{Q}\mathbf{x} \preceq 0, \\ & \partial\mathbf{x} \succeq 0,. \end{aligned} \tag{43}$$

In our numerical experiments, we look at the Gaussian distribution as an example of general risk type PDFs.

A question that naturally comes to mind at this point is how good these Bayesian solutions are, given the assumptions involved about the customers and the limitations imposed by the design constraints. Given the lack of access to the optimal solution, to answer this question, we look at the performance of a model-free learning approach that can get arbitrarily close to the optimal solution given *a long enough* learning phase.

3) The Black-Box Approach: Here we take a frequentist view, taking the mode selection probability vector $\mathbf{p}(\mathbf{x}(t);t)$ to be a *deterministic* but unknown function of $\mathbf{x}$. Given no single underlying model that characterizes $\mathbf{p}(\mathbf{x}(t);t)$, the LSE sees this function as a *black-box* and can use learning techniques to directly learn the profit-maximizing incentive profile. Since learning happens over time, we bring back the superscript t to denote variables associated with time index t. We still omit q. The LSE's expected total profit at time t is given by[4]:

$$\begin{aligned} \rho_{\text{total}}(\mathbf{x}(t);t) &= |\mathcal{P}(t)|\rho(\mathbf{x}(t);t), \\ &= |\mathcal{P}(t)|(\mathbf{u}(t)-\mathbf{x}(t))^T\mathbf{p}(\mathbf{x}(t);t) \end{aligned} \tag{44}$$

where $|\mathcal{P}(t)|$ is the total appliance population interacting with the LSE at time t, and $\rho(\mathbf{x}(t);t)$ is the expected profit per appliance as given by (21).

At every t, the LSE knows $|\mathcal{P}(t)|$[5] and $\mathbf{u}(t)$. However, given the unknown nature of $\mathbf{p}(\mathbf{x}(t);t)$, the LSE cannot find the global maximum of (44) directly. Rather, he will have to try posting different incentives $\mathbf{x}(t)$ in order to *learn* $\mathbf{p}(\mathbf{x}(t);t)$ and minimize the *regret* with respect to the optimal but unknown incentive $\mathbf{x}^*(t)$:

$$\min_{\mathbf{x}(t)} \ \rho_{\text{total}}(\mathbf{x}^*(t);t) - \rho_{\text{total}}(\mathbf{x}(t);t) \tag{45}$$

In doing so, the LSE is faced with a fundamental trade-off: he will have to pick whether to *exploit* his current knowledge of $\mathbf{p}(\mathbf{x}(t);t)$ and post the incentive that maximizes his profit from the set of previously tested incentives; or to *explore* by posting a new $\mathbf{x}(t)$ and learning more about $\mathbf{p}(\mathbf{x}(t);t)$ through sampling the customer response. Next, we address how we approach this learning problem.

Here we assume that due to large aggregation levels, an accurate estimate of $\mathbf{p}(\mathbf{x}(t);t)$ of any posted $\mathbf{x}(t)$ is available to the LSE, i.e., if $n_m(t)$ customers selected mode m in response to incentive $\mathbf{x}(t)$,

$$\mathbf{n}(t) = [n_0(t), n_1(t), \ldots, n_M(t)]^T \tag{46}$$

$$\mathbf{p}(\mathbf{x}(t);t) \approx \frac{\mathbf{n}(t)}{|\mathcal{P}(t)|} \tag{47}$$

where clearly, $\|\mathbf{n}(t)\|_1 = |\mathcal{P}(t)|$.

Consequently, the LSE is faced with a sequential maximization problem of an expensive time-varying black-box function $\mathbf{p}(\mathbf{x}(t);t)$ multiplied by a time-varying *context*[6], i.e., $|\mathcal{P}(t)|(\mathbf{u}(t)-\mathbf{x}(t))$. Generally, the problem of how to *optimally* maximize the profit when learning a time-varying yet structured reward function remains open [27]. However, here we continue by imposing a very reasonable constraint on (45): we assume

[4]Note that this profit is associated with a specific cluster q. However, we continue to refrain from burdening the notation with the cluster index, since the learning problem is independent across different clusters.

[5]As mentioned before, we take $|\mathcal{P}(t)|$ to be known before posting $\mathbf{x}(t)$, which can be realized if customers make an *anonymous query* every time they need to see the incentives.

[6]The term context is commonly used in the bandit literature to refer to the time-varying exogenous effect of the *environment* on the reward associated with each arm's selection at time t. The environment here is the energy market and the customer population size.

that the mode selection probability $\mathbf{p}(\mathbf{x}(t);t)$, the expected recruitment utility $\mathbf{u}(t)$, and the population size $|\mathcal{P}(t)|$ are daily periodic functions, i.e., $\forall t = iH+h, h = 0,\ldots,H-1, i \in \mathbb{Z}$:

$$\rho_{\rm total}(\mathbf{x}(t);t)|_{t=iH+h} = |\mathcal{P}(h)|(\mathbf{u}(h)-\mathbf{x}(t))^T\mathbf{p}(\mathbf{x}(t);h). \tag{48}$$

If we treat $\{\mathbf{p}(\mathbf{x}(t);h), h = 0,\ldots,H-1\}$ as H unrelated black box functions, this allows us to divide the hard sequential learning problem (45) into H parallel sequential learning subproblems that are more benign since the reward function does not change dynamically. The objective of the $h^{\rm th}$ sub-problem on day i is:

$$\max_{\mathbf{x}(iH+h)\succeq\mathbf{0}} \rho_{\rm total}(\mathbf{x}(iH+h);h). \tag{49}$$

There is a large body of literature on global optimization with expensive black-box functions. Most algorithms for black-box optimization are derivative-free and fall into the category of randomized search methods. The brute-force way to approach this random search is to have a pure exploration phase with randomly drawn $\mathbf{x}(t)$'s, following with a pure exploitation phase. However, one could clearly use the (mostly) continuous structure of the reward to sample the function more smartly, and have a shorter exploration phase. The important observation is that gathering overall knowledge about the black-box is irrelevant, especially in the suboptimal areas of the search space. Works following this observation are referred to as *response surface methods* and contain 2 steps:

1) Based on the data collected thus far, a surrogate model of the reward at previously untried points is built. This model is used to provide an estimate of the reward, possibly with the error level associated with the estimate;
2) A design method chooses the best next search point based on the constructed model. This method is commonly heuristic and belongs to one of two categories: either the point where the surrogate function is maximized is picked, or with one more layer of sophistication, the point where it is most *likely* that a high reward can be observed, see [28]–[30].

Here we use the popular kriging interpolator, which can be thought of as a way of modeling a deterministic black-box function as a realization of a stochastic process [31]. Assume that we have observed the value of $\mathbf{p}(.;h)$ at w incentive values $[\mathbf{x}_1,\ldots,\mathbf{x}_w]$. Given these observations, we model the value of $\mathbf{p}(\mathbf{x};h)$ for a new incentive $\mathbf{x}$ as

$$\hat{\mathbf{p}}(\mathbf{x};h) = \mathbf{B}^h\mathbf{f}(\mathbf{x}) + \mathbf{z}(\mathbf{x}), \tag{50}$$

where $\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}),\ldots,f_P(\mathbf{x})]$ is a vector of P chosen regression functions, and the coefficient in the $M \times P$ matrix $\mathbf{B}^h$ are regression parameters trained based on the w observed points. The vector $\mathbf{z}(\mathbf{x})$ contains M zero mean independent random processes with a covariance $C_m^h(\mathbf{x},\mathbf{y})$ between $z_m(\mathbf{x})$ and $z_m(\mathbf{y})$. This random process models the prediction error of the regression model $\mathbf{B}^h\mathbf{f}(\mathbf{x})$ given the errors at previously observed sites. Consequently, if we assume a Gaussian noise model, we can write

$$\hat{\mathbf{p}}(\mathbf{x};h) \sim \mathcal{N}(\mathbf{B}^h\mathbf{f}(\mathbf{x}), \Sigma(\mathbf{x})) \tag{51}$$

where $\Sigma_{m,m'}(\mathbf{x}) = 0$ for $m \neq m'$. See [32] for details on the kriging interpolator.

The availability of a standard error in kriging allows to calculate a 'statistical' measure for selecting the best new sampling point $\mathbf{x}$. Here, we use one of the most popular approaches for selecting an iterate, which is to find the point where the probability of improving the function beyond its current best value T is highest, i.e.,

$$\max_{\mathbf{x}} Q\left(\frac{T-|\mathcal{P}(h)|(\mathbf{u}(h)-\mathbf{x})^T\mathbf{B}^h\mathbf{f}(\mathbf{x})}{\sigma}\right), \tag{52}$$

where $\sigma^2 = |\mathcal{P}(h)|^2(\mathbf{u}(h)-\mathbf{x})^T\Sigma(\mathbf{x})(\mathbf{u}(h)-\mathbf{x})$ and $Q(.)$ denotes the Gaussian Q function. We refer the reader to [33] for a rigorous analysis of the performance of this method.

Remark II.6: Note that in reality the expected utility at hour h, $\mathbf{u}(h)$, and the population size, $|\mathcal{P}(h)|$, are not a deterministic functions and are random. Although the general global optimization of stochastic functions is fairly well studied, there is far less literature on the problem when function evaluations are expensive, and there are no rigorous extensions of sample selection techniques to account for stochasticity [34]. In practice most algorithms ignore this stochasticity when designing the sampling method, which is what we do here. However, observing that these variables commonly vary within a small set, we simulated this stochasticity in the numerical experiment and observed encouraging empirical performance, clearly without discussing any guarantees.

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

III. The Recruitment Utilities for the LSE

In this section we derive different possible expressions for the *DLS market utility* function $U_m^q(t)$ that quantifies the profit the LSE can accrue in its wholesale market interactions, from recruiting an appliance in a certain cluster q and mode m at a certain time t. We consider two ways in which the LSE can benefit from a recruitment transaction, and provide examples of flexible loads that would allow to realize the corresponding DLS market utility in Section III-B and in Section III-C respectively.

Here we only look at one energy market settlement plus the reserve market. Under multiple economic settlements, the presence of a forward purchase would make the recruitment utilities on different appliances to be not additive, and could stimulate competition between customer. These are complications in the problem that will be considered in future work.

A. Multi-Stage Decisions of the Electricity Market

To understand how the LSE can benefit from having flexible loads in the wholesale energy market, it is useful to capture how generation dispatch decisions are typically handled, and mirror these control nobs in the decision model available to the LSE. Let t be the continuous time index. Wholesale market dispatch decisions prescribe the power output of participating generators in multiple settlements. However, these prescriptions cannot be chosen arbitrarily. Denoting by T the duration of an hour, dispatch decisions can be interpreted as assigning the coefficients of a basis expansion of the generation profile as follows:

$$\begin{aligned}\mathrm{G}(\mathrm{t}) = &\sum_{\ell=0}^{H-1} G(\ell)\phi_0(\mathrm{t}-\ell T) + \sum_{\ell=0}^{H-1} y(\ell)\phi_1(\mathrm{t}-\ell T)\\ &+ \sum_{\ell=0}^{H-1}\sum_{s=0}^{S-1} d(\ell S+s)\phi_2\left(\mathrm{t}-\ell T-\frac{sT}{S}\right),\end{aligned} \tag{53}$$

where the dispatched shape for the generation profile $\mathrm{G}(\mathrm{t})$ is decomposed in three components: the first term whose coefficients $G(\ell)$ are settled in the *day-ahead market* for hour ℓ, the second term whose coefficients $y(\ell)$ are hourly adjustments to the day-ahead settlements, and the third representing sub-hourly adjustments controlled by the coefficient $d(\ell S+s)$. This last term fine-tunes the balance of power in the grid.

Therefore, market operator dispatch decisions on $\mathrm{G}(\mathrm{t})$ are limited to a certain signal space (53), defined by the independent waveforms $\{\phi_0(\mathrm{t}-\ell T), \phi_1(\mathrm{t}-\ell T)\}_{\ell=0}^{H-1}$ and $\{\phi_2(\mathrm{t}-jT/S)\}_{j=0}^{SH-1}$ to determine $\mathrm{G}(\mathrm{t})$[7]. The decision is to determine the most economic generation units to operate and how much should each one generate depending on the further physical limitations of each unit. In fact, capacity and ramping constraints bound the set $\mathcal{G}(t)$ of available shapes $\mathrm{G}(\mathrm{t})$ and cost curves introduce a partial ordering of $\mathrm{G}(\mathrm{t})$ that ultimately leads to the choice of certain vectors of coefficients $\boldsymbol{G}=[G(0),\ldots,G(H-1)]^T$, $\boldsymbol{y}=[y(0),\ldots,y(H-1)]^T$ and $\boldsymbol{d}=[d(0),\ldots,d(SH-1)]^T$ for each generator (possibly 0).

The situation of a retailer that buys power from the market can mirror that of a generator, as far as the structure of the market purchase is concerned, i.e.,

$$\mathrm{L}(\mathrm{t}) = \sum_{\ell=0}^{H-1} L(\ell)\phi_0(\mathrm{t}-\ell T) + \sum_{\ell=0}^{H-1} y(\ell)\phi_1(\mathrm{t}-\ell T) + \sum_{\ell=0}^{H-1}\sum_{s=0}^{S-1} d(\ell S+s)\phi_2\left(\mathrm{t}-\ell T-\frac{sT}{S}\right). \quad (54)$$

The basic difference from the generation side is that demand retailers with inflexible demand approach the early market settlements passively with a forecasted load profile they have to procure, i.e., they request to the market to meet a *base load* equal to $\boldsymbol{L}=[L(0),\ldots,L(H-1)]^T$. During the day, the other two terms $y(\ell)$ and $d(\ell S+s)$ go to random variations around this base load due to forecast error and the retailer is bound to pay for them. However, with flexible demand, the retailer can start to actively manage its own $L(\ell), y(\ell), d(\ell S+s)$. Following DR initiatives, in today's markets, industrial loads with mostly constant base loads $\boldsymbol{L}$ can offer to alter their real-time load around the base load $\boldsymbol{L}$ and act like a generator in the hour-ahead and subhourly dispatch decisions[8]. This allows them to accrue benefits from having dispatched values for $\boldsymbol{y}=[y(0),\ldots,y(H-1)]^T$ and $\boldsymbol{d}=[d(0),\ldots,d(SH-1)]^T$. This is also an option available to the LSE (see Section III-C). In addition, an LSE with flexible load can actively decide the vector $\boldsymbol{L}$ instead of providing passive load forecasts. Thus, it can procure an $\boldsymbol{L}$ which would increase its profit margins (see Section III-B). Clearly, if the LSE were to control a sizable amount of demand, it could potentially choose $\boldsymbol{L}$ to affect market prices in its favor, as discussed in [35]. However, in the following, we will implicitly assume that the LSE is sufficiently small so that it can be treated as a price-taker and cannot affect market prices.

[7]The waveforms shapes are pretty elementary: $\phi_0(\mathrm{t})$ is rectangular, $\phi_1(\mathrm{t})$ is trapezoidal and $\phi_2(\mathrm{t})$ is a straight line.

[8]The current practice is that the base load is actually an average of $\boldsymbol{L}$ over many days where the flexibility is not exercised.

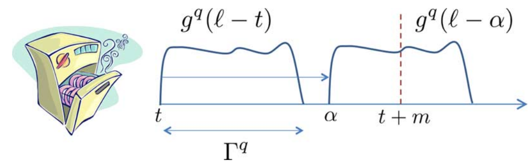

Fig. 3. Deferrable Non-interruptible Appliances.

B. Utility Gained From Inter-Hourly Load Shifting

Typically electricity markets in the US trade energy on an hourly basis. Denote by $\pi^e(\ell)$ the *expected* market clearing price of the energy market at a future hour ℓ.

Remember that we denote the set of all possible shapes for the electricity consumption profile of an appliance in cluster q that is recruited in mode m at time t as the set $\mathcal{L}_m^q(t)$. On the other hand, if the appliance is not recruited, its mode is $m=0$. Thus, the load of an unrecruited appliance starting at time t is $\mathcal{L}_0^q(t)=\{L_0^q(\ell)\}$, which we refer to as the *intrinsic base load*, since it is the ideal base-load current market standards for DR try to capture. By recruiting an appliance in cluster q under mode m at time t, and shifting its load to the best (cost minimizing) time period, the LSE can save:

$$U_m^q(t) = \sum_\ell \pi^e(\ell)L_0^q(\ell) - \min_{L(\ell)\in\mathcal{L}_m^q(t)} \sum_\ell \pi^e(\ell)L(\ell). \quad (55)$$

To provide the reader with practical examples of recruitment utilities, here we look at deferrable loads (long-duration loads with flexible consumption schedules) such as EVs, washer/dryers, etc. Once these types of appliance are turned on, their load profile can be easily predicted by the customer, either from factory provided documentation, or by metering previous consumption profiles. We divide these appliances in 3 major categories, which we look at separately.

First, we discuss the category that is well-suited to accrue utility from inter-hourly load shifting.

1) Design Constraint 2: Here we assume that for all deferrable loads, the mode index $m\in\mathcal{M}^q$ is directly equivalent to the amount of slack time that accompanies the request.

2) Deferrable Non-Interruptible Appliances: This category best models appliances such as washer/dryers, and noninterruptible EV battery chargers. For these types of loads, the only degree of freedom to manage demand is through scheduling the time at which the appliance starts its job. Once scheduled, the appliance cannot be interrupted. These loads are characterized by a certain load phasor $g(\ell-t;\mathbf{v})$, where t is the arrival time of the request and $\mathbf{v}$ are the phasor parameters (rate, duration, samples, etc). Denoting the time at which the appliance is scheduled to start its job as α, the phasor is shifted to be $g(\ell-\alpha;\mathbf{v})$ where $t<\alpha<t+\mu$ and μ is the slack time. Following the discussion in Section II-A, we quantize $(\mathbf{v},\mu)$ so that: 1) to meet the deadline constraint imposed by the customer, the set of acceptable values for α is $t\le\alpha\le t+m$ where m is the quantized value of μ; 2) once the appliance starts its job, the energy consumption is characterized by a pulse $g^q(\ell-\alpha)$, which is the quantized version of $g(\ell-\alpha;\mathbf{v})$. The total length of the task is denoted by Γ^q (see Fig. 3). To relate this with the variables required to solve (55), note that in case the appliance is not recruited ($m=0$), the load is:

$$L_0^q(\ell) = g^q(\ell-t), \quad (56)$$

and the set of possible load shapes under mode m is:

$$\mathcal{L}_m^q(t) = \{L(\ell) | L(\ell) = g^q(\ell - \alpha),\ t \leq \alpha \leq t + m\}. \quad (57)$$

Thus, the DLS market utility for the LSE of non-interruptible appliances in cluster q under mode m for load shifting at time t is:

$$U_m^q(t) = \sum_{\ell=t}^{t+\Gamma^q} \pi^e(\ell) g^q(\ell - t) - \min_{\alpha} \left[\sum_{\ell=\alpha}^{\alpha+\Gamma^q} \pi^e(\ell) g^q(\ell - \alpha) \right]$$
$$\text{s.t. } t \leq \alpha \leq t + m. \quad (58)$$

Denoting the optimum shift by α_m, i.e.,

$$\alpha_m = \underset{\alpha}{\operatorname{argmin}} \sum_{\ell=\alpha}^{\alpha+\Gamma^q} \pi^e(\ell) g^q(\ell - \alpha) \ \text{ s.t. } t \leq \alpha \leq t + m,$$

the load profiled to be scheduled to accrue this utility is

$$L_m^q(\ell) = g^q(\ell - \alpha_m). \quad (59)$$

C. Utility Gained From Selling Ancillary Services

Since electricity demand and supply vary within the hour, a combination of *ancillary services* are used to compensate for any intra-hour deviation between supply and demand, to maintain the stability of the grid. Ancillary services are procured in a *reserve market*, which is run in parallel with the energy market everyday. Traditionally, fast-ramping generators like gas and hydro provide these services, due to the reliable, clean, and highly controllable nature of their outputs. However, with high levels of wind and solar integration being considered for the grid, the possibility of using the demand side for ancillary services provision is being considered [36]–[38]. Here we study the possibility of using direct load scheduling for ancillary services that maintain the daily balance of supply and demand in the grid, i.e., regulation and load following.

When using flexible loads to provide ancillary services, we assume that the LSE is able to manage its load on an intra-hourly bases, and expects to be paid at an hourly-variable rate of $\pi^{a,r}(\ell)$ and $\pi^{a,l}(\ell)$ per unit capacity of regulation and load following services respectively. In expressing the utility, the LSE needs to map the structure of the flexibility specified in the sets $\mathcal{L}_m^q(t)$ onto an expression for the capacity offered for that particular service. These services are defined next[9].

Definition III.1. Load Following Services on an Hourly Basis: These services are used to balance the grid supply and demand on an hourly basis. The LSE will need to adjust its hourly demand to deviate from the expected base load by a value $y(\ell)$ dispatched by the system operator. If a capacity $C^{LF}(\ell)$ of load following services were offered for hour ℓ, we have

$$|y(\ell)| \leq C^{LF}(\ell). \quad (60)$$

Definition III.2. Regulation Services at Sub-Hourly Epochs: In order to have dispatched variations on the sub-hourly load, we assume that each hour ℓ consist of S sub-hourly epochs at which regulation signals are sent out by the grid operator. We index these epochs by $j = \ell S + s, s = 0, \ldots, S-1, \ell = 1, 2, 3, \ldots$ and denote the regulation capacity offered for hour ℓ as $C^R(\ell)$. The LSE is required to deviate from its hourly base load $L(\ell)$, at the grid operator's request, by a value $d(j)$ at each epoch j, with

$$|d(j)| \leq C^R(\ell) \ \text{ for } \ \ell S \leq j < (\ell + 1)S. \quad (61)$$

[9]For brevity of notation, some aspects of the service are not completely realistic. Load following services are commonly offered on a 20–30 minute basis rather than hourly. However, more realistic definitions are compatible with our design framework.

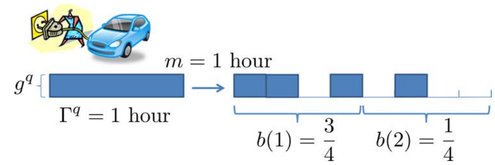

Fig. 4. Interruptible EV Charging with a Non-controllable Rate.

Since the ancillary services capacity $C(\ell)$ that can be offered at hour ℓ is tied with the optimized base load $L(\ell)$ purchased in the energy market, the utility of load shifting in the hourly market is not independent of the utility of selling ancillary services, and the two should be considered concurrently. This point will be further elaborated in the following examples on specific types of appliances.

1) Interruptible EV Charging With a Non-Controllable Rate: Here we look at the specific case of electric vehicles that charge at a constant and non-controllable rate, specified in $\mathbf{v}$, but have the capability for multiple automated interruptions in power consumption during the charging process. Instead of considering all possible arbitrary profile of interruption, given by $\mathbf{f}$, once again we quantize the possible ways in which the charge can be stopped and resumed. More specifically, for an EV in cluster q, with battery capacity E^q, non-controllable rate g^q kWs and a slack time (mode) of m time units, the degrees of freedom for the LSE would be to pick disjoint time intervals with a total length of $\Gamma^q = E^q/g^q$ to charge the vehicle in the interval $[t, t + m + \Gamma^q]$. They can be thought of a box with a length of Γ^q and height g^q, that need to be disassembled in hour long portions, which are re-distributed between slots, leaving some slots *full* and some slots *empty* (see Fig. 4). The choice of the LSE is mirrored by a decision variable $b(\ell)$, describing the fraction of hour ℓ for which the EV should be charged (at the constant rate g^q). Clearly,

$$0 \leq b(\ell) \leq 1. \quad (62)$$

Given this, the set of possible hourly load shapes for an interruptible EV recruited at time t is:

$$\mathcal{L}_m^q(t) = \left\{ L(\ell) \,\middle|\, L(\ell) = g^q b(\ell), \sum_{\ell=t}^{t+m+\Gamma^q} g^q b(\ell) = E^q \right\}. \quad (63)$$

On the other hand, if $m = 0$, the charge is never interrupted (intrinsic base load):

$$L_0^q(\ell) = g^q \operatorname{rect}\left(\frac{\ell - t}{\Gamma^q} \right), \quad (64)$$

where $\operatorname{rect}(.)$ denotes the unit pulse between $[0, 1]$.

Thus, if no ancillary services are provided, the utility gained by recruiting an appliance of cluster q in mode m at time t is

$$U^q_{m,s}(t) = \sum_{\ell=t}^{t+\Gamma^q} \pi^e(\ell)g^q - \min_{b(\ell)} \left[\sum_{\ell=t}^{t+m+\Gamma^q} \pi^e(\ell)g^q b(\ell)\right]$$
$$\text{s.t.} \quad 0 \le b(\ell) \le 1, \quad \sum_{\ell=t}^{t+m+\Gamma^q} b(\ell) = \Gamma^q,$$

Denoting the optimal $b(\ell)$ in (65) as $b_m(\ell)$, the optimal load profile the LSE will pick for recruits under mode m is:

$$L^q_m(\ell) = g^q b_m(\ell), \quad t \le \ell \le t+m+\Gamma^q. \tag{65}$$

Notice that, unlike the previous category, here the LSE has the flexibility to pause and restart the charge which, in large scales, could allow the LSE to offer ancillary services. To avoid an unreasonable number of interruptions, we refrain from offering regulation services. Instead, we decide to offer load following services, described in Definition III.1.

The questions is: how much load following capacity can the LSE offer with a single interruptible EV? An interruptible EV with a charge time of $b(\ell)$ assigned to hour ℓ can at most decrease its demand by $g^q b(\ell)$, or increase it by $g^q(1-b(\ell))$. Thus, any ancillary service dispatch signal within hour ℓ that is outside this range cannot be honored by this EV. Consequently, the capacity $C^{LF}(\ell)$ for load following offered by the specific EV in cluster q is smaller than both $g^q(1-b(\ell))$ and $g^q b(\ell)$.

Since the total charging load of the EV is constant, any increase in the load at a certain hour has to happen concurrently with a decrease in the load at another hour. Thus, following a dispatch signal $y(\ell)$ is plausible *if and only if*,

$$\sum_{\ell=t}^{t+m+\Gamma^q} y(\ell) = 0, \quad |y(\ell)| \le C^{LF}(\ell).$$

However, since the load following signals $y(\ell)$ are exogenous to the operation of the LSE, this constraint might not always hold. However, given a certain constant capacity C^{LF}, it is common to assume that load following signals are energy neutral. So we refrain from providing a time-varying load following capacity $C^{LF}(\ell)$ using a single appliance.

Assumption 7: The energy neutrality requirement holds for load following dispatch signals[10], i.e.,

$$\sum_{\ell=t}^{t+m+\Gamma^q} y(\ell) = 0, \quad |y(\ell)| \le C^{LF}, \tag{66}$$

Consequently, we have

$$U^q_m(t) = \max_{b(\ell),C^{LF}} \left\{ \sum_{\ell=t}^{t+\Gamma^q} \pi^e(\ell)g^q - \sum_{\ell=t}^{t+m+\Gamma^q} \pi^e(\ell)b(\ell)g^q + \sum_{\ell=t}^{t+m+\Gamma^q} \pi^{a,l}(\ell)C^{LF} \right\}$$
$$\text{s.t.} \quad 0 \le b(\ell) \le 1, \quad \sum_{\ell=t}^{t+m+\Gamma^q} g^q b(\ell) = E^q,$$
$$C^{LF} \le g^q b(\ell), \qquad t \le \ell \le t+m+\Gamma^q,$$
$$C^{LF} \le g^q(1-b(\ell)), \quad t \le \ell \le t+m+\Gamma^q. \tag{67}$$

[10]This is usually not the case for load following services on shorter time scales. However, a large-scale aggregation of DR resources could allow for extensive load defferal to resolve this issue.

Using the optimal value of $b(\ell)$ in (67), we can again write the load that the LSE will pick for mode m recruits, this time using them to provide ancillary services, as (65).

2) EVs With a Controllable Rate of Charge: Finally, with a fully controllable charging rate $g(\ell)$, a vehicle of cluster q in mode m requires to receive a total charge of E_q energy units anytime between $[t, t+m+\Gamma^q_{\min}]$ at the discretion of the LSE, with $\Gamma^q_{\min}$ denoting the minimum charge length for appliances in cluster q. If charging can happen at a maximum rate of $g^q_{\max}$ per hour, $\Gamma^q_{\min} = E_q S/g^q_{\max}$. Thus, we can write set of possible load shapes as,

$$\mathcal{L}^q_m(t) = \left\{ L(\ell) \middle| L(\ell) = g(\ell),\ g(\ell) \in [0, g^q_{\max}],\ \sum_{\ell=t}^{t+m+\Gamma^{\min}_q} g(\ell) = E_q \right\} \tag{68}$$

and the intrinsic base load as,

$$L^q_0(\ell) = g^q_{\max}\mathrm{rect}\left(\frac{\ell-t}{\Gamma^{\min}_q}\right), \tag{69}$$

where we have assumed that if the EV is not recruited by the DLS program, it will charge at the highest possible rate.

Given this, the recruitment utility for load shifting is:

$$U^q_{m,s}(t) = \sum_{\ell=t}^{t+\Gamma^q_{\min}} \pi^e(\ell)g^q_{\max} - \min_{g(\ell)} \left[\sum_{\ell=t}^{t+m+\Gamma^q_{\min}} \pi^e(\ell)g(\ell)\right]$$
$$\text{s.t.} \quad 0 \le g(\ell) \le g^q_{\max}, \quad \sum_{\ell=t}^{t+m+\Gamma^q_{\min}} g(\ell) = E_q.$$

Same as the previous category, the next step here is to study the potential for these EVs to provide ancillary services. Due to increased degrees of demand flexibility available through a controllable rate of charge, we choose to provide regulation services using these appliances, a scenario considered before in the literature [39]–[42].

At an epoch within hour ℓ, these EVs have the potential to increase their load from $g(\ell)$ to any amount up to $g^q_{\max}$, or decrease it down to 0. However, like the previous case, the resources to perform this load modification are limited, and to ensure that the load modification is feasible, we assume again that we offer a homogeneous regulation capacity C^R based on one EV, and that regulation signals are independent random variables and energy neutral in time.

Assumption 8: The energy neutrality requirement holds for regulation dispatch signals, i.e.,

$$\sum_{j=(\ell^a+1)S}^{(\ell^d-1)S-1} d(j) = 0, \quad |d(j)| \le C^R. \tag{70}$$

Then, we can calculate the recruitment utility of EVs with a controllable rate of charge in cluster q arriving at time t as:

$$\begin{aligned} U_m^q(t) = \max_{g(\ell), C^R} & \left\{ \sum_{\ell=t}^{t+\Gamma_{\min}^q} \pi^e(\ell) g_{\max}^q - \sum_{\ell=t}^{t+m+\Gamma_{\min}^q} \pi^e(\ell) g(\ell) \right. \\ & \left. + \sum_{\ell=t}^{t+m+\Gamma_{\min}^q} \pi^{a,r}(\ell) C^R \right\} \\ \text{s.t.} \quad & 0 \le g(\ell) \le g_{\max}^q, \quad \sum_{\ell=t}^{t+m+\Gamma_{\min}^q} g(\ell) = E_q, \\ & C^R \le g_{\max}^q - g(\ell), \quad \forall t \le \ell \le t+m+\Gamma^q, \\ & C^R \le g(\ell), \qquad \forall t \le \ell \le t+m+\Gamma_{\min}^q. \end{aligned} \tag{71}$$

Denoting the optimal value of $g(\ell)$ maximizes (71) as $g_m(\ell)$, the load profile that the LSE will pick for mode m recruits in this case is given by:

$$L_m^q(\ell) = g_m(\ell), \quad t \le \ell \le t+m+\Gamma_{\min}^q. \tag{72}$$

IV. Addressing Privacy Concerns Regarding Direct Load Scheduling

A major concern commonly associated with direct load scheduling programs is that customers do not like to reveal their detailed consumption characteristics to a for-profit load serving entity. Even though retail dynamic pricing is also not immune from this problem, because smart meters convey a lot of information about the electric appliance in use, the information for direct load scheduling is disaggregated at appliance level already. However, as shown in [25], [43], the clustering approach we take in this paper can help alleviate this issue by allowing the LSE to run the direct load scheduling program without knowing the identity of the owners of recruited appliances. Here we briefly explain how this is possible by presenting the gist of our results in [43]. We focus on the case of non-interruptible deferrable loads, and we refer the reader to our previous work for other load types. A reader not interested in privacy issues can simply skip this Section.

A. Cluster-Based Load Management

Since all appliances that share the same cluster q and mode (laxity) m have the same quantized load characteristics and flexibility, the LSE can bundle all of them in a single service queue, indexed by (q, m). Appliances that are recruited under queue (q, m) can be dispatched following a First-In-First-Out (FIFO) discipline and their ownership is irrelevant for the sake of the control. Also, our incentive design in Section II-D is based on the aggregate behavior (cf. (17)) and, thus, individual preferences are not exploited or accounted for in the price design. This means that in principle control and pricing design could be anonymous. The question is how to convey the control decision to the appliances when you do not know who specifically is available. We explain next how to do that.

Mathematically, we define:

- the arrival process $a_m^q(t)$, denoting the aggregate number of appliances (irrespective of who owns them) that are recruited in cluster q under mode m *at or before* t;
- the activation process $d_m^q(t)$, denoting the aggregate number of appliances in cluster q under mode m that are scheduled to start their job *at or before* t.

Hence, following the load model described in Section III-B1 for deferrable non-interruptible appliances, the aggregate system load due to this type of appliances at time ℓ is given in terms of $d_m^q(t), t \le \ell$, as:

$$L(\ell) = \sum_{t=0}^{\ell} \sum_{q=1}^{Q} \sum_{m=1}^{M^q} (d_m^q(t) - d_m^q(t-1)) g^q(\ell - t) \tag{73}$$

The choice of $d_m^q(t)$ to manage load is clearly restricted by the appliance arrival process $a_m^q(t)$:

1) *Causality*: only appliances that have already arrived in the system can be scheduled, i.e., $\forall q, m$:

$$d_m^q(t) \le a_m^q(t). \tag{74}$$

2) *Deadline constraint*: appliances that arrive with delay laxity m need to be scheduled within m epochs of their arrival. Under a FIFO discipline, this means that:

$$d_m^q(t) \ge a_m^q(t-m). \tag{75}$$

Since past activations are irreversible, $d_m^q(t) \ge d_m^q(t-1)$.

B. Uplink

As we saw, for scheduling appliances, the LSE will only need the values of $a_m^q(t)$'s. So the aggregate arrival numbers in different queues need to be gathered from the appliance population, and no individual identifiers are required for load management. Each appliance (or a respective Home Energy Management System) possesses a gateway assumed to be part of a larger wireless AMI. At the neighborhood level, the appliances communicate their request identifiers (q, m) to *collector nodes* through a network operated by a third non-interested party, and get recruited. The values of $a_m^q(t)$'s for each neighborhood is then calculated by the respective collector node and communicated to the LSE at pre-determined intervals, through a secure out-of-band channel (e.g., over the Internet). The third-party entity that owns the collector nodes is also in charge of billing and verification that recruited appliances follow the LSE dispatch signals.

Note that in the uplink it is not possible to preserve anonymity if the AMI is owned and operated by the LSE itself. Also, it is clear that anonymity is not an exact synonym of privacy and that information leakage can lead the LSE to infer private information about customers from the aggregate queries. Studying this aspect is beyond the scope of our work, and we invite the reader to see [44], [45].

The remaining question is how to design a functional downlink channel so that the LSE can notify the specific appliances it needs to activate. We explain this next.

C. Downlink

When *real-time* decisions $d_m^q(t)$ are made, the FIFO rule has to be implemented to directly address the $d_m^q(t)$ appliances that arrived earlier than the appliances waiting in cluster q under mode m, and activate them. But how can the LSE address in the downlink these specific appliances when their identity is unknown? The idea is simple: at time t, the LSE needs to find the time epoch $T_m^q \le t$ at which the number of appliances that arrived in queue (q, m) since the origin of time, i.e., $a_m^q(T_m^q)$, was marginally less than the number of appliances the LSE plans to have activated since time epoch 0 until time t, i.e., $d_m^q(t)$:

$$T_m^q(t) = \max\{\ell \le t : a_m^q(\ell) \le d_m^q(t)\}. \tag{76}$$

If this time index $T_m^q(t)$ is broadcast to *all appliances*, then the appliances belonging to queue (q, m) that submitted requests before time $T_m^q(t)$ can start their jobs. The rest wait until they are authorized later.

V. Numerical Experiment

A. Setting and Data

The numerical experiment we conduct is based on real Plug-in Hybrid EV (PHEV) charge data we studied in [46]. The database includes 620 charge events spread across a few months. However, for the purposes of this experiment, we aggregate all these events in one single day, taken to be September 1st 2013. Each entry of this database includes the time of plug-in (arrival), charge amount between 0 to 5 kWhs, and departure time (taken as an approximate deadline). We divide these requests in 10 clusters, each representing different charge amounts with 0.5 kWh quantization steps. To demonstrate the possibility of providing ancillary services through DLS, we take the PHEVs to be rate-flexible with a maximum rate of $g_{\max}^q = 3$ kWs, even though in reality all charge events used a constant rate of 1.1 kWs (home charging). Wholesale energy and regulation capacity prices are taken to be equal to the ISO New England's day-ahead market clearing prices in the Maine load zone on September 1st.

A parameter needed to conduct this experiment is the user risk function (i.e., type), which cannot be directly extracted from the data. So we try to generate estimates of the customers types based on the data, with the steps taken described next. For simplicity, and due to lack of any meaningful alternatives, we take the commitment risk function in (22) to be linearly increasing with the type (i.e., $j_m^q(t) = m$). We then assume that $\gamma_i \sim U(0, 0.1)$, and solve for a noisy estimate of the incentives, denoted by $\hat{x}_m^q(t)$, through (36). From these sample incentives, we reiterate back and find the γ_i for each customer such that for the mode corresponding to the deadline in the database, denoted by m_d, we have

$$(m_d + 1)\gamma_i \geq \hat{x}_{m_d+1}(t).$$

which provides us with a preliminary estimate of the γ_i's. The uniform distribution that best fits these type values is then determined to be $U(0, 0.08)$.

B. Experiments With Bayesian Incentive Design Under Uniform Risk Parameters

Here we determine the incentives using (36) and simulate the customer response. Fig. 5 shows the number of PHEVs that joined the program at different hours of the day, and which mode they picked. The reader can observe that the number of arrivals, along with the number of recruitments with high slack times, is considerably higher during evening and early night hours.

Next, we study the economic effects of the implementation of the DLS program. To do so, we compare the customer savings from upgrading to DLS from the normal flat pricing service paradigm of the grid. This saving is captured by the customer maximized utility in (10). We sum the value of this utility for all the plug-in events that happened at a certain hour, and show this aggregate savings as a function of time in Fig. 6. Note that higher savings at evening hours are due to a combination of three effects, that are intuitive by looking at Fig. 5: higher arrival numbers, higher incentives, and customers having more laxity to offer (lower risk types). To provide insight into the magnitude

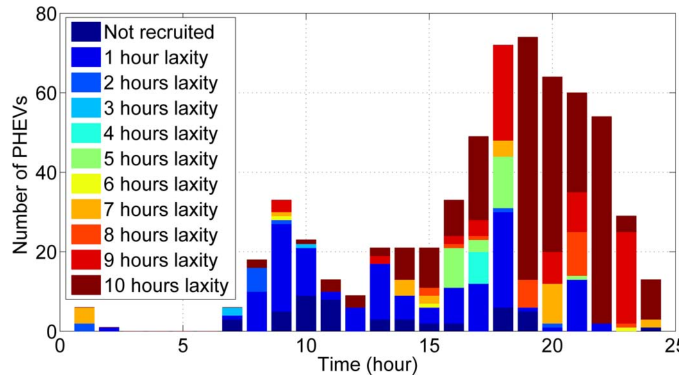

Fig. 5. The arrival time of the 620 PHEVs, along with the mode under which they were recruited.

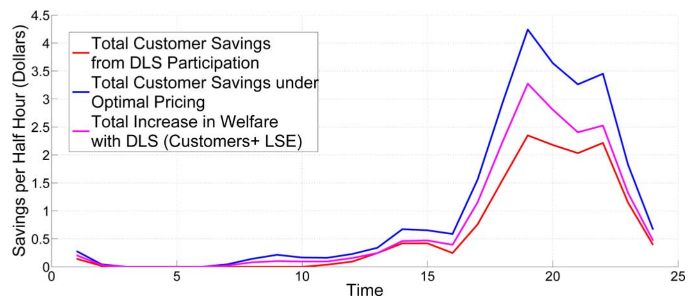

Fig. 6. Welfare effects of DLS compared with a fictitious optimal pricing scheme.

of these savings, we compare this scenario to one that provides the highest possible savings that could be achieved by the customer through any demand management scheme. The savings could be realized if customers were engaged in a ideal demand management scheme where they can directly talk to the main operator that runs the energy and ancillary services markets, communicate their quantized consumption flexibility to the operator, find the least cost solution to serve their demand, and get dispatched to provide ancillary services. To provide an absolute highest bound on the savings, we assume that, since there are no commitments to a DLS unit in this ideal scheme, the customer feels no *commitment risk*, and the market profit gained is equivalent to that gained by the LSE through recruiting the appliances with their maximum laxity, i.e., their individual $U_q^t(m_d)$ given the m_d. The value of this utility is summed up across the population and is shown as a function of time in Fig. 6. Note that part of this difference in customer surplus between the hypothetical optimal demand management scheme (which is equivalent to optimal pricing with perfect retail competition) and the DLS scheme is due to the presence of a monopolist that facilitates the interactions between the customers and the wholesale market, i.e., the LSE. To showcase this effect, we show the total increase in the community's welfare (customers+LSE) with the implementation of DLS in Fig. 6. The difference between this value and that of optimal pricing is due to the presence of the commitment risk in our model.

Remark V.1: We would like to point out that the welfare effects studied here are purely on the retail side, assuming that our LSE will not affect energy market prices. Clearly, with a large integration of active demand resources, this will not be the case and energy market prices will start to be affected, most likely decreasing the generators' welfare.

Next, we demonstrate the ability of the LSE to follow the regulation dispatch signals sent by the system operator in real-time.

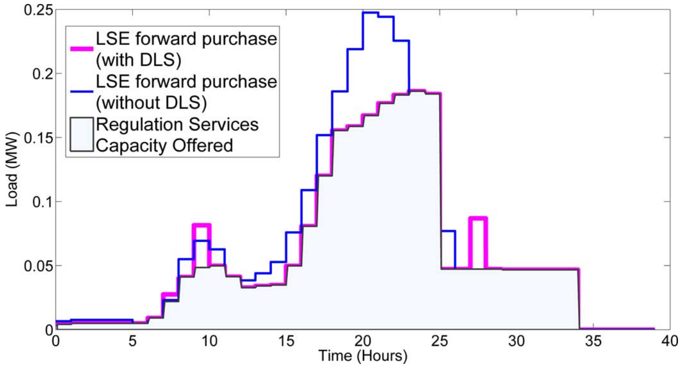

Fig. 7. The energy market purchase of the LSE, with and without DLS.

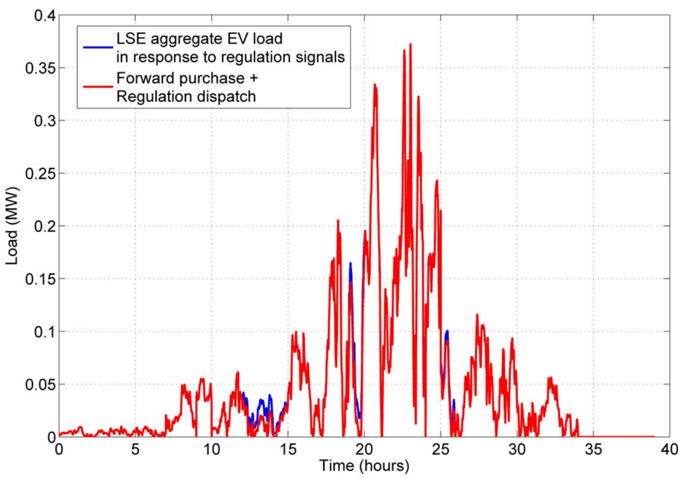

Fig. 8. The performance of 620 PHEVs in following regulation dispatch signals sent to the LSE.

To do so, we fist show the hourly bulk purchase that the LSE makes from the wholesale market, with and without the DLS program, in Fig. 7, along with the regulation capacity offered by the LSE to the ancillary services market by recruiting the PHEVs. Real regulation signals are scaled down in magnitude to be below the hourly capacity offered. The scaled down signal is then added on top of the LSE's optimal hourly bulk purchase with DLS to provide the profile that the LSE should follow in real-time, shown in the red line in Fig. 8. We assume that the LSE uses a Least Laxity First (LLF) policy (similar to [19], [47]) to follow this profile in real-time by charging the PHEVs with the least amount of slack time first. The reader can observe occasional deviations from the dispatch, which is mainly due to the violations of the energy neutrality assumption for real regulation signals.

In our last figure, we numerically study the welfare effects of the introduction of the following concepts when designing DLS program incentives: 1) classification of requests of different capability in different clusters; 2) allowing the incentives to vary with the laxity offered, i.e., modes; and 3) allowing the incentives to dynamically change with time (See Fig. 9). This study is done by respectively constraining the incentive profiles to be equal for all EVs offering the same laxity, be equal for all EVs recruited at the same time, and be a constant that does not change with cluster, mode, or time. The reader can observe that clustering, even though effective in increasing the community's welfare, does not play as important of a role as the two other concepts, i.e., allowing the incentives to change dynamically with time and laxity offered. We think this is mainly because of the

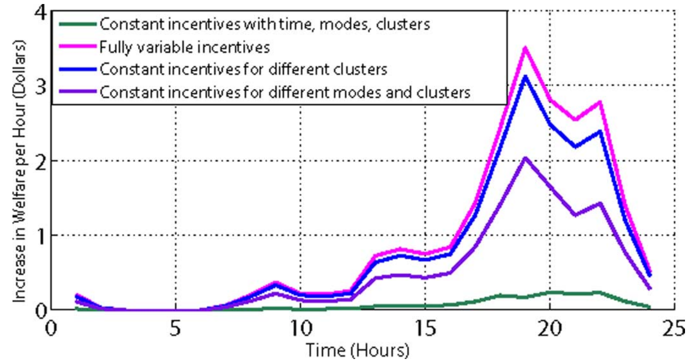

Fig. 9. Comparison of the increase in community (LSE+customers) welfare with DLS, when several parameters of the model are gradually eliminated.

uneven distribution of charge lengths for events studied in [46], with a large proportion of the PHEV requiring full charge when plugged in.

C. Performance Comparison

While the previous experiments were all carried out by posting incentives from (36), the reader should remember that we have not claimed that these incentives are optimal, due to two effects: 1) the Bayesian modeling of customer using the uniform distribution could be erroneous; 2) the single-crossing design constraint limits the search space. Here, we attempt to assess the performance of multiple proposed solutions for the incentive design problem. Since the incentives are cluster-specific, to have a meaningful comparison, we focus on a specific cluster (Rate-flexible EVs with 3 hours of charge), and a specific hour (arrivals between 6–7 pm). Due to the small size of the EV data set, we increase the population of arriving customers in this cluster at this time following the statistical model in [46]. Thus, we generate requests from a Poisson distribution with an average arrival rate of 1000 EVs, and laxity values from the following mixed log-normal and exponential distribution:

$$f(m_d) = 0.7 \ln \mathcal{N}(2.25, 0.4^2) + 0.3(1.089 e^{-1.089 m_d}).$$

Individual risk parameters are then found using the procedure described in Section V-A.

Fig. 10 and Table I compare the profit of the LSE in 74 consecutive days (Sep 1st to Nov 13th 2013) when the incentives are derived from the following methods: 1) A Bayesian setting with a uniform prior, i.e., (36); 2) A Bayesian setting with a Gaussian prior, i.e., (43); 3) A black-box setting with a learning phase of 100 days and the sample selection rule in (52), and a first order polynomial regressor; and 4) A black-box setting with a learning phase of 100 days and purely randomized sample selection. The results show that almost consistently, method 3 has the best performance, followed by method 4, method 1, and method 2. What we would like to point out is that even though none of these techniques are optimal, we claim that the margin of error of method 4 from the optimal (but unknown) incentive is rather small. To illustrate this point, we derive an upper-bound for the LSE profit given any incentive design scheme: we calculate the highest possible profit the LSE can extract if it knew the risk parameter of every participating customer, and was allowed to discriminate between customers, i.e., pay each customer at their reservation price (risk value) for the all recruitment modes,

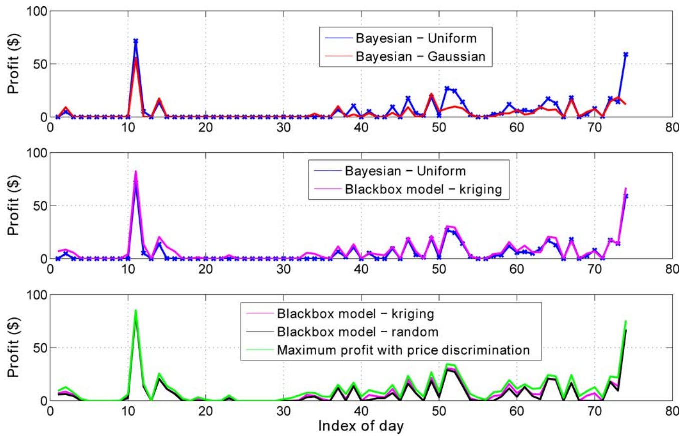

Fig. 10. Comparison of the LSE's profit under various incentive design schemes.

TABLE I
COMPARISON OF THE 4 STUDIED INCENTIVE DESIGN SCHEMES.

Method	LSE profit	Avg daily recruitment # and payment
Bayesian - Uniform	$493	707 EVs - 3.2c per EV
Bayesian - Gaussian	$281	555 EVs - 1c per EV
Black box - kriging	$653	560 EVs - 2c per EV
Black box- random	$501	754 EVs - 3.8c per EV
Upper bound	$774	708 EVs - 2.1c per EV

and in an arbitrarily small (ε) higher amount for the most profitable mode:

$$\rho_{\text{total}}(\mathbf{x}(t);t) \leq \sum_{i=1}^{|\mathcal{P}(t)|} \max_{m}(U_m(t) - \gamma_i j_m(t) - \varepsilon). \quad (77)$$

This value of this upper bound is shown in Fig. 10. The reader should note that this upper-bound cannot be achieved under any publicly posted incentive design scheme like ours.

VI. CONCLUSIONS

Direct load scheduling, mainly due to its reliable nature, can pave the way towards more active management of demand in the power grid, particularly with the granularity required to offer ancillary services through the demand side. Here we provided a framework to quantify the costs and benefits of implementing such a program by a for-profit entity in the short run. However, the viability of such ventures in the long run will depend on the level of retail market competition, the initial infrastructure investment to build communication networks (or upgrade existing networks to comply with the requirements) and to retrofit eligible appliances with sub-metering equipments, and the operational and maintenance costs associated with the direct scheduling scheme.

REFERENCES

[1] S. Borenstein and S. P. Holland, "On the efficiency of competitive electricity markets with time-invariant retail prices," Tech. Rep. National Bureau of Economic Research, 2003.

[2] R. Li, Q. Wu, and S. S. Oren, "Distribution locational marginal pricing for optimal electric vehicle charging management," *IEEE Trans. Power Syst.*, vol. 29, no. 1, pp. 203–211, Jan. 2013.

[3] M. G. Vayá and G. Andersson, "Locational marginal pricing based smart charging of plug-in hybrid vehicle fleets," in *Proc. Smart Energy Strategies Conf.*, 2011, pp. 21–23.

[4] N. O'Connell, Q. Wu, J. Ostergaard, A. H. Nielsen, S. T. Cha, and Y. Ding, "Electric vehicle (ev) charging management with dynamic distribution system tariff," in *Proc. 2nd IEEE PES Int. Conf. and Exhib. Innovative Smart Grid Technol. (ISGT Europe),* , 2011, pp. 1–7.

[5] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *Proc. IEEE Power Energy Soc. General Meeting*, 2011, pp. 1–8.

[6] T.-H. Chang, M. Alizadeh, and A. Scaglione, "Real-time power balancing via decentralized coordinated home energy scheduling," *IEEE Trans. Smart Grid*, vol. 4, no. 3, pp. 1490–1504, Sep. 2013.

[7] S.-J. Kim and G. Giannakis, "Scalable and robust demand response with mixed-integer constraints," *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 2089–2099, Dec. 2013.

[8] A.-H. Mohsenian-Rad, V. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid," *IEEE Trans. Smart Grid*, vol. 1, no. 3, pp. 320–331, Dec. 2010.

[9] Z. Ma, D. Callaway, and I. Hiskens, "Decentralized charging control for large populations of plug-in electric vehicles," in *Proc. 49th IEEE Conf. Decision Control (CDC)*, 2010, pp. 206–212.

[10] W. Tushar, W. Saad, H. Poor, and D. Smith, "Economics of electric vehicle charging: A game theoretic approach," *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 1767–1778, Dec. 2012.

[11] C. Chen, S. Kishore, and L. V. Snyder, "An innovative rtp-based residential power scheduling scheme for smart grids," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, 2011, pp. 5956–5959.

[12] P. Samadi, H. Mohsenian-Rad, R. Schober, and V. Wong, "Advanced demand side management for the future smart grid using mechanism design," *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1170–1180, Sep. 2012.

[13] S. Chen, T. Mount, and L. Tong, "Optimizing operations for large scale charging of electric vehicles," in *Proc. 46th Hawaii Int. Conf. Syst. Sci.*, 2013, vol. 0, pp. 2319–2326.

[14] O. Sundstrom and C. Binding, "Planning electric-drive vehicle charging under constrained grid conditions," in *Proc. Int. Conf. Power Sys. Technol. (POWERCON)*, 2010, pp. 1–6.

[15] M. Caramanis and J. M. Foster, "Management of electric vehicle charging to mitigate renewable generation intermittency and distribution network congestion," in *Proc. 48th IEEE Conf. Decision Control, held Jointly with the 28th Chinese Control Conf. CDC/CCC 2009*, 2009, pp. 4717–4722.

[16] K. Mets, T. Verschueren, W. Haerick, C. Develder, and F. De Turck, "Optimizing smart energy control strategies for plug-in hybrid electric vehicle charging," in *Proc. IEEE/IFIP Network Operat. Management Symp. Workshops (NOMS Wksps),* , 2010, pp. 293–299.

[17] S. Deilami, A. Masoum, P. Moses, and M. A. S. Masoum, "Real-time coordination of plug-in electric vehicle charging in smart grids to minimize power losses and improve voltage profile," *IEEE Trans. Smart Grid*, vol. 2, no. 3, pp. 456–467, Sep. 2011.

[18] D. Wu, D. Aliprantis, and L. Ying, "Load scheduling and dispatch for aggregators of plug-in electric vehicles," *IEEE Trans. Smart Grid*, vol. 3, no. 1, pp. 368–376, Mar. 2012.

[19] A. Subramanian, M. Garcia, A. Dominguez-Garcia, D. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling of deferrable electric loads," in *Proc. Amer. Control Conf. (ACC)*, 2012, pp. 3643–3650.

[20] E. H. Gerding, V. Robu, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings, "Online mechanism design for electric vehicle charging," in *Proc. 10th Int. Conf. Auton. Agents Multiagent Syst.*, Richland, SC, USA, 2011, vol. 2, AAMAS '11, pp. 811–818.

[21] E. Bitar and S. Low, "Deadline differentiated pricing of deferrable electric power service," in *Proc. IEEE 51st Annu. Conf. Decision Control (CDC)*, 2012, pp. 4991–4997.

[22] M. Kefayati and R. Baldick, "Energy delivery transaction pricing for flexible electrical loads," in *Proc. IEEE Int. Conf. Smart Grid Commun. (SmartGridComm)*, 2011, pp. 363–368.

[23] P. Samadi, H. Mohsenian-Rad, V. Wong, and R. Schober, "Tackling the load uncertainty challenges for energy consumption scheduling in smart grid," *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 1007–1016, Jun. 2013.

[24] M. Alizadeh, Y. Xiao, A. Scaglione, and M. van der Schaar, "Incentive design for direct load control programs," in *Proc. 51th Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, pp. 1029–1036.

[25] M. Alizadeh, A. Scaglione, and R. J. Thomas, "From packet to power switching: Digital direct load scheduling," *IEEE J. Select. Areas Commun.*, vol. 30, no. 6, pp. 1027–1036, Jul. 2012.

[26] C.-Y. Chong and R. P. Malhami, “Statistical synthesis of physically based load models with applications to cold load pickup,” *IEEE Trans. Power Apparatus Syst.*, vol. PAS-103, no. 7, pp. 1621–1628, Jul. 1984.
[27] T. A. Desautels, “Spinal cord injury therapy through active learning,” Ph.D. dissertation, California Inst. of Technol., Pasadena, CA, USA, 2014.
[28] D. R. Jones, M. Schonlau, and W. J. Welch, “Efficient global optimization of expensive black-box functions,” *J. Global Optimiz.*, vol. 13, no. 4, pp. 455–492, 1998.
[29] A. Cassioli and F. Schoen, “Global optimization of expensive black box problems with a known lower bound,” *J. Global Optimiz.*, vol. 57, no. 1, pp. 177–190, 2013.
[30] D. D. Cox and S. John, “Sdo: A statistical method for global optimization,” in *Multidisciplinary Design Optimization: State-of-the-Art*. Philadelphia, PA, USA: SIAM, 1997, pp. 315–329.
[31] J. Sacks *et al.*, “Design and analysis of computer experiments,” *Statist. Sci.*, vol. 4, no. 4, pp. 409–423, 1989.
[32] D. R. Jones, “A taxonomy of global optimization methods based on response surfaces,” *J. Global Optimiz.*, vol. 21, no. 4, pp. 345–383, 2001.
[33] H.-M. Gutmann, “A radial basis function method for global optimization,” *J. Global Optimiz.*, vol. 19, no. 3, pp. 201–227, 2001.
[34] M. Tesch, “Improving robot locomotion through learning methods for expensive black-box systems,” Ph.D. dissertation, Robotics Inst., Carnegie Mellon Univ., Pittsburgh, PA, USA, 2014.
[35] M. G. Vaya and G. Andersson, “Optimal bidding strategy of a plug-in electric vehicle aggregator in day-ahead electricity markets,” in *Proc. 10th Int. Conf. Eur. Energy Market (EEM)*, 2013, pp. 1–6.
[36] P. Cappers, J. MacDonald, C. Goldman, and O. Ma, “An assessment of market and policy barriers for demand response providing ancillary services in U.S. electricity markets,” *Energy Policy*, vol. 62, pp. 1031–1039, 2013.
[37] M. Albadi and E. El-Saadany, “A summary of demand response in electricity markets,” *Elect. Power Syst. Res.*, vol. 78, no. 11, pp. 1989–1996, 2008.
[38] M. Parvania and M. Fotuhi-Firuzabad, “Demand response scheduling by stochastic SCUC,” *IEEE Trans. Smart Grid*, vol. 1, no. 1, pp. 89–98, Jun. 2010.
[39] J. Kang, S. J. Duncan, and D. N. Mavris, “Real-time scheduling techniques for electric vehicle charging in support of frequency regulation,” *Proc. Comput. Sci.*, vol. 16, no. 0, pp. 767–775, 2013.
[40] S. Han, S. Han, and K. Sezaki, “Development of an optimal vehicle-to-grid aggregator for frequency regulation,” *IEEE Trans. Smart Grid*, vol. 1, no. 1, pp. 65–72, Jun. 2010.
[41] E. Sortomme and K. Cheung, “Intelligent dispatch of electric vehicles performing vehicle-to-grid regulation,” in *Proc. IEEE Int. Elect. Veh. Conf. (IEVC)*, 2012, pp. 1–6.
[42] J. Donadee and M. Ilic, “Stochastic co-optimization of charging and frequency regulation by electric vehicles,” in *Proc. North Amer. Power Symp. (NAPS)*, 2012, pp. 1–6.
[43] M. Alizadeh, A. Scaglione, A. Applebaum, G. Kesidis, and K. Levitt, “Reduced-order load models for large populations of flexible appliances,” *IEEE Trans. Power Syst.*, 2014, to be published.
[44] S. R. Rajagopalan, L. Sankar, S. Mohajer, and H. V. Poor, “Smart meter privacy: A utility-privacy framework,” in *Proc.. IEEE Int. Conf. Smart Grid Commun. (SmartGridComm)*, 2011, pp. 190–195.
[45] C. Dwork, “Differential privacy: A survey of results,” in *Theory and Applications of Models of Computation*. New York, NY, USA: Springer, 2008, pp. 1–19.
[46] M. Alizadeh, A. Scaglione, J. Davies, and K. Kurani, “A scalable stochastic model for the electricity demand of electric and plug-in hybrid vehicles,” *IEEE Trans. Smart Grid*, vol. 5, no. 2, pp. 848–860, Mar. 2014.
[47] S. Chen, L. Tong, and T. He, “Optimal deadline scheduling with commitment,” in *Proc. 49th Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, 2011, pp. 111–118.

Mahnoosh Alizadeh is a Ph.D. candidate at the University of California Davis. She received the B.Sc. degree in electrical engineering from Sharif University of Technology in 2009. Her research interests mainly lie in the area of real-time control and optimization for cyber-physical systems, with a special focus on designing scalable and economic demand side management frameworks for the Smart Grid.

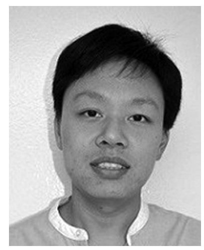

Yuanzhang Xiao received the B.E. and M.E. degree in electrical engineering from Tsinghua University, Beijing, China, in 2006 and 2009, respectively. He is currently pursuing the Ph.D. degree in the Electrical Engineering Department at University of California, Los Angeles. His research interests include multi-agent sequential decision making, game theory, and optimization.

Anna Scaglione is currently Professor in Electrical and Computer Engineering at the University of California Davis. She served as the Editor in Chief of the IEEE SIGNAL PROCESSING LETTERS in 2012–2013, the Associate Editor for the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS from 2002 to 2005, and from 2008 to 2011 in the Editorial Board of the IEEE TRANSACTIONS ON SIGNAL PROCESSING, where she was Area Editor in 2010–11. She has been in the Signal Processing for Communication Committee from 2004 to 2009 and is in the steering committee for the conference Smartgridcomm since 2010. She was general chair of the workshop SPAWC 2005. She is the first author of the paper that received the 2000 IEEE SIGNAL PROCESSING TRANSACTIONS Best Paper Award; she has also received the NSF Career Award in 2002, and is a co-recipient of the Ellersick Best Paper Award (MILCOM 2005), and the 2013 IEEE Donald G. Fink Award. Her expertise is in the broad area of signal processing for communication systems and networks. Her current research focuses on signal processing algorithms for networks and for sensors systems, with specific focus on Smart Grid, demand side management and reliable energy delivery. She has been a Fellow of the IEEE since 2011.

Mihaela van der Schaar is Chancellor’s Professor of Electrical Engineering at University of California, Los Angeles. She is a Fellow of IEEE. She received an NSF CAREER Award (2004), the Best Paper Award from IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY (2005), the Okawa Foundation Award (2006), the IBM Faculty Award (2005, 2007, 2008), the Most Cited Paper Award from EURASIP: Image Communications Journal (2006), the Gamenets Conference Best Paper Award (2011) and the 2011 IEEE Circuits and Systems Society Darlington Award Best Paper Award. For more information about her research visit: http://medianetlab.ee.ucla.edu/