

Provision of Regulation Service Reserves by Flexible Distributed Loads

Michael Caramanis, Ioannis Ch. Paschalidis, Christos Cassandra, Enes Bilgin and Elli Ntakou

Abstract—Following our previous work on the control of multiple appliances in response to Independent System Operator (ISO) Regulation Service Signals (RSS), we model the ISO’s RSS dynamics - evolving in a time scale of seconds - as a two level Markov process whose transition probabilities are calibrated on actual data. Appliance response is modeled as a Markov modulated process consistent with an exponentially distributed time to switch off and an expected utility that is concave in the price charged when an appliance turns on. Prices are broadcasted dynamically by a Smart Building Operator (SBO) with the objective of maximizing the time average of utility gained when appliances turn on minus the cost of imperfect RSS tracking. We prove certain properties of the stochastic Dynamic Programming (DP) policies that allow us to formulate the problem as an approximate Discrete State and Control Space DP and propose a reasonable approximation that renders the problem scalable to multiple appliance categories. The discretized state DP solution can be obtained as a solution to a Linear Program (LP). The LP provides the optimal dynamic price control policies and in addition yields the requisite information needed by the SBO to bid optimally for energy and Regulation Service Reserve (RSR) Capacity in the hour ahead balancing market. Solving for the discretized real time market optimal price policies, using them to calibrate a continuous analytic policy function, and extracting from the real time optimal policies the optimal bid to the hour ahead forward balancing market is the main contribution of this paper.

Index Terms—Distributed Demand Management, Regulation Service Provision, Real Time Stochastic Control, Hour-Ahead and Real Time Market Response

I. INTRODUCTION

The advent of the smart grid will undoubtedly enable broader participation in Electric Power Markets including Smart Building Operator (SBO) managed loads in commercial and residential buildings [3, 4, 5]. Power Market transactions cover a basket of related products that include stand-by reserve capacity in addition to energy [4, 5, 6, 7, 8, 9]. In fact, reserve requirements are likely to increase with the adoption of renewable generation which is clean but intermittent and volatile [2]. The cost of reserves, particularly fast up-and-down secondary reserves, is already substantial, in fact it is comparable to that of energy [4, 5]. Today,

secondary reserve requirements are approximately 1% of peak load, are procured in the day ahead or hour ahead forward markets, and utilized as needed by the Independent System Operators (ISOs) in real time to deal with stochastic demand and volatile uncontrollable generation fluctuations. Whereas today, secondary reserves commonly known as Regulation Service Reserves (RSRs) are provided for by flexible generators who are capable of modulating their output in positive as well as negative increments - hence the up-and-down nature of RSRs - expectations for increased RSR requirements as we go forward may pose unsustainable costs to substantive integration of renewable generation. Fortunately, the simultaneous increase in the adoption of flexible or deferrable loads with storage-like capabilities such as Electric Vehicle (EV) battery charging and duty cycle appliances, may provide synergistic reprieve [16,17] by augmenting the supply of RSRs and thereby controlling their cost. In this paper, we provide a decision support framework for flexible loads represented by SBOs to (i) participate in forward power markets through the flexible loads with RSR capabilities that the SBO represents, and, more importantly (ii) to track optimally the dynamic regulation signal that the ISO transmits in real-time. In fact, optimal bidding to the forward market is dependent on the expected costs associated with the optimal SBO dynamic policy used to respond to the ISO regulation signal. The coupling of cascaded markets such as the forward hour ahead market and the real time management of regulation signal tracking, is examined in a general manner in [18].

We start by describing the forward markets where RSR, Q^R , and Energy, Q^E , bids are made by the SBO, and then proceed to describe the task of the SBO in tracking the ISO dynamic regulation signal.

In the forward power market, the SBO secures a base Energy load Q^B , and offers a flexible load that can consume at a controllable rate ranging between 0 and $2R$. We model here the market rules adopted by Eastern US Interconnection ISOs such as PJM and NYISO, where the provider of RSR promises to respond to up or down real time dynamic signals broadcasted by the ISO in the aforementioned range, while the ISO commits to broadcasting dynamic up and down signals whose time average is 0. Because of the bidirectional nature of RSRs, the energy, Q^E , and bidirectional capacity, Q^R , that the ISO may schedule must conform to the following constraints: $0 \leq Q^R + Q^E \leq 2R$ and $Q^R \leq Q^E$. As described in more detail in Section IV, if the SBO

The authors are with the Center for Information and Systems Engineering, Boston University, College of Engineering. Research partially supported by the NSF under grants EFRI-0735974 and EFRI-1038230, by the DOE under grant DE- FG52-06NA27490, by the ARO under grand W911NF-11-1-0227, AFOSR grant FA9550-12-1-0113, ONR grant N00014-09-1-1051 and by the ODDR& E MURI10 program under grant N00014-10-1-0952. Corresponding Author email mcaraman@bu.edu

reservation prices and the market clearing prices for energy and RSRs satisfy a market event that schedules positive Q^R , then, the maximum RSR is scheduled and the market clearing satisfies $Q^R = Q^E = R$. Denoting the market clearing prices for energy and RSRs respectively by Π^E and Π^R , the SBO (i) is charged by $(Q^B + R)\Pi^E - R\Pi^R$, but the cost reduction $R\Pi^R$ does not come for free, since (ii) the SBO will have to take control actions in real-time to respond to a regulation signal, $y(t)$, that the ISO updates every few seconds during the hour that follows the clearing of the hour ahead forward market. More precisely, the SBO must regulate the flexible load Q^E in real time, so that $Q^E(t) \approx R + Ry(t)$, where $y(t)$ is a regulation signal broadcasted dynamically by the ISO taking values in the interval $[-1, 1]$ with a time average of 0. This results in a time average of $Q^E(t)$ equal to R . While, there are no additional payments during the real time tracking period, there may be costs associated with $|Q^E(t) - R - Ry(t)| > 0$ i.e., there are performance-based-penalties¹ for the SBO. In addition, the SBO may inflict utility loss to the users of the appliances constituting the flexible load $Q^E(t)$.

This paper is a follow up on related past work [14, 15] focusing on the problem where ISO real time regulation signal is independent of SBO feedbacks, an assumption used in that earlier work, and addressing the dynamic control policies as opposed to static asymptotic control considered in that earlier work. In addition, in this paper we consider explicitly the interaction of optimal real time control and forward market bidding strategy. The rest of the paper is organized as follows. In Section II we develop the models used to represent flexible load response to a dynamic price signal and the stochastic dynamics of the ISO regulation signal. In Section III we present the optimal SBO tracking policy problem and propose several approximate SDP solution approaches, including one that is scalable to multiple flexible appliance classes since it decomposes the large multiple class problem to many single class problems. In Section IV we present the relationship between the hour ahead forward market optimal bid strategy and the real time tracking control policy, and sketch an algorithm for optimizing the hour ahead bid strategy. In Section V we present numerical results and conclude in Section VI.

II. CONSUMPTION AND REGULATION REQUEST SIGNAL MODELS

A. Duty Cycle Appliances (HVAC, Refrigerators, Hot Water Heaters)

Consider a building or a neighborhood with N appliances for example heating zones - of which, at time t , $n(t)$ are connected and consume actively at a rate of r kW each, while the rest, $N - n(t)$, are idle or disconnected and not consuming. We assume that when an appliance connects, i.e.

¹For example, the PJM ISO RSR market rules provide RSR providers with good performance based incentives or negative penalties that are conceptually equivalent.

starts to consume, it joins the infinite server queue of active appliances and continues to consume for a time period τ with $\mathbb{E}[\tau] = 1/\mu$. When it stops consuming it starts idling. The SBO acts as the appliance coordinator by broadcasting dynamically a price $u(t)$. Idle devices detect $u(t)$ with a given probability. An idle device that detects the SBO price $u(t)$ compares it to its utility and decides to either remain idle or connect, pay the price, and for that price consume at the rate r for a period τ . An idle device reconnects if its utility exceeds $u(t)$. The utility of idle appliances, $U(\theta(t), \tilde{v})$, is defined as a function of $\theta(t)$, the heating zone temperature at time t , and \tilde{v} a random variable² whose distribution depends on $\theta(t)$. That utility is associated with an average energy consumption of $r\mathbb{E}[\tau] = r/\mu$. which is for a bounded range space of \tilde{v} , it follows that $\theta(t)$ is also bounded. In general, for a given history of broadcasted prices $I_t = \{u(\xi), \forall \xi \leq t\}$, the conditional density function of the utility of an idle appliance $f(U(\theta(t), \tilde{v})|I_t)$ is obtainable in terms of the joint density function of $\theta(t)$ and \tilde{v} given I_t . [10, 11, 12, 13].

Assumption 1. The distribution, $f(U(\theta(t), \tilde{v})|I_t)$, depends only on a single sufficient statistic, the time average of $u(t)$, is uniformly distributed conditional upon the sufficient statistic, and is i.i.d. across appliances. Specifically, $f(U(\theta(t), \tilde{v})|I_t) = 1/U^M$, namely, the utility of idle appliances is uniformly distributed over $[0, U^M]$ where U^M is the maximal price broadcasted by the SBO which is consistent with the sufficient statistic and the quantity of regulation service that the SBO has sold to the ISO in the hour ahead market. This is a reasonable approximation for small variations in $u(t)$. We will later show that for large t and optimal selection of $u(t)$, the sufficient statistic is known, and equals Π^E , the clearing price of the hour ahead power market.

Assumption 2. Idle appliances detect the price broadcast by the SBO with a rate λ and the time elapsing between observation instances is exponential, or equivalently the number of observations over a period t_a is Poisson with parameter λt_a .

Using Assumptions 1 and 2, when the SBO broadcasts price $u(t)$ during $[t, t + \Delta t]$, the number of any of the $[N - n(t)]$ appliances that connect during $[t, t + \Delta t]$ is Poisson distributed with parameter $\lambda \Delta t (N - n(t))(U^M - u(t))/U^M$.³

Assumption 3. We assume that the controlled system will exhibit $n(t)$ such that $n_1 \leq n(t) \leq n_2$ with $n_2 - n_1$ is small.

²The random variable represents different preferences amongst appliance users, different local conditions e.g., windows are open, a heating zone is or is not occupied -, or different future plans, for example an EV expects to depart sooner, needs a high State of Charge etc.

³For $\Delta t \rightarrow 0$, the same expression becomes the probability that a single appliance will connect while the probability of additional connections is negligible.

Using the above assumptions, letting $\Delta t = 1$ by appropriate selection of time unit and defining $\lambda^M = \lambda(N - (n_1 + n_2)/2)$ expressed in units of $\Delta t = 1$, we can reasonably approximate the Poisson arrival rate over the period $[t, t + 1]$ as $\lambda^M (U_i^M - u_i(t))/U_i^M$.

We have chosen to burden the reader with the details of three Assumptions underlying the concluding statement above in order to make transparent the shortcomings of Assumption 1 which has been commonly adopted in dynamic pricing models applied by Kelly and others to telecommunications bandwidth allocation and more recently to electricity demand management [14, 20]. The shortcoming of this Assumption, particularly when a non-trivial portion of market participants (whether mobile telephone users or electricity consuming appliances) is active, is that it introduces the notion that the utility of connecting is state independent. In fact, that utility depends on the past history of price control. For example, receiving repeatedly a busy signal may result in a higher utility associated with placing a phone call, or observing a persistently high consumption price that has repeatedly discouraged a heating zone from turning the heater on will almost certainly result in a lower ambient temperature and increase the utility to consume. Although we do not attempt to correct the shortcomings of assumption 1 in this paper, we wish to identify the issue and encourage future work that relaxes this restrictive Assumption. We are in fact engaged in such future work which we are eager to report on in the near future.

B. Electric Vehicles (EVs)

Defining (i) by λ_{ev}^M the rate at which an EV arrives in the Smart Buildings charging station, or a plugged in but non-charging EV monitors the price broadcasted by the SBO, and (ii) by $U(\text{SoC}(t), T)$ the utility of charging r_{ev}/μ_{ev} kW Δt when the State of Charge of their battery is $\text{SoC}(t)$, $0 \leq \text{SoC}(t) \leq 100\%$, and the desired departure time is T , (see [16]) we can similarly write the Poisson arrival rate as $\lambda_{ev}^M (U_{ev}^M - e_{ev}(t))/U_{ev}^M$.

C. Generalization of Appliance Classes

Denoting different classes of appliances by the subscript i , the above can be generalized as Poisson arrival rate of connections by appliances of class i when the SBO broadcasts to that appliance class a price $u_i(t)$ over the time period $[t, t + 1]$ is $\lambda_i^M (U_i^M - u_i(t))/U_i^M$. The expected utility rate that appliances of class i experience at time t when the SBO broadcasts price $u_i(t)$ is

$$\begin{aligned} & \lambda_i^M (U_i^M - u_i(t))/U_i^M [(U_i^M + u_i(t))/2] \\ & = \left(\lambda_i^M / (2U_i^M) \right) \left((U_i^M)^2 - (u_i(t))^2 \right) \end{aligned}$$

Noting that when an appliance of class i connects it pays $u_i(t)$ to consume on average r_i/μ_i kW Δt , it is instructive to think that when a price of $u_i(t)\mu_i/r_i$ per kW Δt is broadcasted, for a period Δt , an additional amount of energy equal to r_i/μ_i kW Δt is purchased over the period $[t, t + \Delta t]$

regardless of the fact that it will be actually consumed over the period $[t, t + \tau]$ where τ is an exponentially distributed random variable with mean $1/\mu$. Therefore, the expected purchase rate of energy at time t when the broadcasted price is $u_i(t)$ equals the arrival rate times times the expected purchase by a single arrival, namely

$$\lambda_i^M [(U_i^M - u_i(t))/U_i^M] r_i/\mu_i \quad (1)$$

When a constant price \bar{u}_i is broadcasted, then noting that $n(t)$ is an $M/M/\infty$ queue whose length is a Poisson distribution with rate

$$\lambda_i^M [(U_i^M - \bar{u}_i)/U_i^M] / \mu_i$$

and hence

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}[n_i(t)] &= \bar{n}_i(\bar{u}_i) \\ &= \lambda_i^M [(U_i^M - \bar{u}_i)/U_i^M] / \mu_i \end{aligned}$$

Hence, the average consumption rate when the constant price \bar{u}_i is broadcasted is

$$\bar{n}_i(\bar{u}_i) r_i = \lambda_i^M [(U_i^M - \bar{u}_i)/U_i^M] r_i / \mu_i \quad (2)$$

Comparing equations 1 and 2 it may be seen that the expected rate at which energy is *purchased* under dynamic pricing is equal to the average rate at which energy will be *consumed* when the same price is broadcasted for a long period of time.

Finally, it is clear that the following propositions hold:

- Proposition 1.** 1.1 Both the average consumption rate and the expected purchase rates are linear functions of the price charged for a packet of energy” of size r_i/μ_i , and
- 1.2 The expected utility rate realized when a price per packet of energy u_i is broadcasted during a period $[t, t + dt]$, is a concave function of u_i .
- 1.3 When dynamic packet prices are charged whose time average is $\bar{u}_i = \Pi^E r_i/\mu_i$ where Π^E is the hour ahead clearing price, the linearity stated above will guarantee that the average purchase rate will be

$$\lambda_i^M [(U_i^M - u_i(t))/U_i^M] r_i/\mu_i$$

- 1.4 The time average of the expected utility rate under the dynamic prices described above, i.e. prices with time average $\lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T u_i(t) dt = \bar{u}_i = \Pi^E r_i/\mu_i$ is smaller than the average utility under a static constant price, namely:

$$\begin{aligned} & \lambda_i^M [(U_i^M)^2 - (\bar{u}_i)^2] / (2U_i^M) > \\ \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T & \left\{ \lambda_i^M [(U_i^M)^2 - (u_i(t))^2] / (2U_i^M) \right\} dt \\ & = \mathbb{E}_{u_i} [\lambda_i^M [(U_i^M)^2 - (u_i)^2] / (2U_i^M)] \end{aligned}$$

and in fact their difference is proportional to the variance of $u_i(t)$.

Proof: Suppose that a dynamic pricing policy is applied over the period $[0, T]$ in such a way that the constant price $u_i(k)$ is broadcasted over the interval

$[(k-1)\Delta t, k\Delta t]$, such that, $k = 1, \dots, K$ and $K\Delta t = T$. The average utility achieved over $[0, T]$ is given by

$$\frac{1}{K\Delta t} \sum_{k=1}^K \left[\frac{\lambda_i^M}{2U_i^M} ((U_i^M)^2 - (u_i(k))^2) \Delta t \right] \quad (3)$$

Represent the price using the definition $u_i(k) = \bar{u}_i + \epsilon_i(k)$, where \bar{u}_i is the average price over the period, i.e., $\frac{1}{K} \sum_{k=1}^K u_i(k) = \bar{u}_i$, and $\epsilon_i(k)$ is the deviation from this average at time k . Recalling that the statistical behavior of $y(t)$ guarantees energy neutrality in the request of RSRs and the linearity of energy consumption w.r.t. price, note that the time average of $\epsilon_i(k)$ equals zero. For $\Delta t \rightarrow 0$, we can conclude that the average utility over the period is represented by the expression

$$\frac{\lambda_i^M}{2U_i^M} \left[(U_i^M)^2 - (\bar{u}_i)^2 - \frac{1}{T} \int_0^T (\epsilon_i(t))^2 dt \right] \quad (4)$$

This clearly shows that price deviations from the average price at any time t , i.e., all instances of $\epsilon_i(t) \neq 0$, for any $t \in [0, T]$ will subtract from the constant price utility. An important implication of this result is that the term $(\lambda_i^M)/(2U_i^M T) \int_0^T (\epsilon_i(t))^2 dt$ in fact quantifies the utility loss due to the fluctuation of $u_i(t)$ about the static average price that would have been broadcasted in the absence of RSR provisioning. In conclusion, as $T \rightarrow \infty$ the time average of ϵ_i goes to zero and the average utility loss is directly proportional to the variance of the dynamic price broadcasted in order to track the regulation signal $y(t)$. ■

D. The Regulation Signal and Tracking costs

The ISO manages the regulation service reserves procured in the hour ahead market by broadcasting a regulation signal $y(t)$ at time intervals of 4 seconds [4]. The regulation signal $y(t)$ indicates the portion of R that should be offered at time $(t + 4sec)$. $y(t)$ takes values in $[-1, 1]$ and is allowed to change by approximately 0, 0.033, or -0.033^4 . Using the real historical data, we estimated transition probabilities of a two layer ($D = up, down$) Markov chain. For $D = up$ there is a higher probability that $y(t+4sec) = y(t) + 0.033$ and a lower probability that $y(t+4sec) = y(t) - 0.033$. The opposite trend holds when $D = down$. The state of the regulation signal is denoted by the pair $D(t), y(t)$. When $y(t + 4sec) > y(t)$ we set $D(t + 4sec) = up$. When $y(t + 4sec) < y(t)$ we set $D(t + 4sec) = down$. When $y(t + 4sec) = y(t)$ we set $D(t + 4sec) = D(t)$.

An SBO who has sold R kW of regulation service reserves, is obliged to regulate the number of its connected appliances to track $\bar{R} + y(t)R$, where \bar{R} is the average consumption rate in the absence of regulation reserve offerings⁵. The SBO is

⁴The update time period in different ISO balancing areas ranges in the US from 4 to 8 seconds depending on the type of regulation service. Similarly, the maximal increment in $y(t)$ between consecutive broadcasts conforms to a maximal ramp of $1/150$ to $1/300$.

⁵more explicitly, $\bar{R} = Q^B + R = \sum_i \bar{n}_i r_i$ where over bar denotes time average and Q^B the base load secured by the SBO

assessed a tracking cost over the period $\Delta t = 4sec$ of

$$K \left[\sum_i n_i(t + 4sec) r_i - \bar{R} - y(t)R \right]^2 \Delta t$$

Proposition 2. The SBO broadcasts prices so as to maximize the utility enjoyed by all classes minus the tracking costs. The linearity of Proposition 1.1 suggests that marginal utility of consumption is constant. In addition, the SBO must be broadcasting the same price per kW Δt to all classes. Thus for a price of $u(t)$ \$/kW Δt , the price $u_i(t)$ broadcasted to class i for consuming $r_i \Delta t / \mu_i$ kW Δt satisfies $u_i(t) = u(t) r_i \Delta t / \mu_i$, which for $\Delta t = 1$ implies $u_i(t) = u(t) r_i / \mu_i$.

III. OPTIMAL REAL TIME TRACKING POLICY

A. Formulation of Optimal Stochastic Control Problem

The SBOs task is to broadcast dynamically class specific prices that maximize expected total class utility minus tracking costs. For a two class case, given the class specific utility and arrival (connection) and departure (disconnection) rates, and the stochastic dynamics of regulation signal, we have a uniformized, infinite horizon average cost problem formulation with all rates expressed in units that correspond to $\Delta t = 1$, the following Bellman Equation

$$\begin{aligned} & h(n_1, n_2, y, D_{old}) + \bar{J} = \\ \min_u & \mathbb{E}_{\Delta n_1, \Delta n_2, \Delta y, D_{new} | u} \left[K \left(\sum_i (n_i + \Delta n_i) r_i - \bar{R} - y(t)R \right)^2 \right. \\ & \left. - \sum_i \left(\frac{\lambda_i^M}{2U_i^M} ((U_i^M)^2 - (\frac{ur_i}{\mu_i})^2) \right) \right. \\ & \left. + h(n_1 + \Delta n_1, n_2 + \Delta n_2, y + \Delta y, D_{new}) \right] \end{aligned}$$

Using Proposition 2, the probability distribution of $\Delta n_1, \Delta n_2, \Delta y, D_{new}$ is given in terms of the arrival rates $(\lambda_i^M / U_i^M) [U_i^M - u(t) r_i / \mu_i]$, the departure rates, $1 / \mu_i$, and the regulation signal stochastic dynamics described in Section II.

B. Solution Approach

We first solved the optimal stochastic problem described above by discretizing the set of allowable prices $0 \leq u \leq U^M$, assuming that the positive probability states in the optimally controlled system limits n_i in the vicinity of $(\lambda_i^M / U_i^M) (U_i^M - \Pi^E r_i / \mu_i) / \mu_i$ as is suggested by Proposition 1.3 and equation 2. The maximum ramp rate allows us also to discretize ISO regulation signals $y(t)$ rendering the resulting state space discrete. Given the discrete state and control spaces, the optimal differential cost satisfying the infinite horizon average cost DP Bellman equation can be obtained by a an LP problem with constraints equal to the product of state and control space magnitudes [1]. The optimal solution for the discretized control approximate DP provides the optimal price for each state (n_1, n_2, y, D) as well as the associated probability that the state will be visited and the associated optimal price broadcasted [1]. Numerical solution (see Section V) revealed the following properties:

- 1) The time averages of prices, and states are: $\mathbb{E}[u_i] = \Pi^E r_i / \mu_i$, $\mathbb{E}[n_i] = (\lambda_i^M / U_i^M)(U_i^M - \Pi^E r_i / \mu_i)$, $\mathbb{E}[y] = 0$ and $Pr(D = up) = Pr(D = down) = 0.5$
- 2) The optimal prices are a function of three features of the state, (i) $x = \sum_i (n_i) r_i - y(t)R$, (ii) $y(t)$ and (iii) D
- 3) The optimal price u and hence $u_i = wr_i / \mu_i$ appear to be described by a sigmoid function $u = \pi(n_1, n_2, y, D) = [U^M / (1 + e^{ax - b(y, D)})]$. Controlling for the value of D and for appropriate calibration of parameter a and the linear function b , the sigmoid approximation is observed to perform well. Further discussion is provided in the numerical results section below.

A much smaller LP can now be formulated and solved since the number of constraints is now proportional to the number of sigmoid functional policy approximations that are of the order of 2 to 3 as opposed to the 11 different prices in the discretized control set. More importantly, the average cost obtained for the best sigmoid analytic approximation selected by the LP solution, compared favorably to that obtained by solving the discretized policy set LP. Additional work that we do not report here due to space limitations, indicates that the sigmoid function representation of optimal pricing policies enables the use of promising stochastic gradient estimation approaches [19] to optimize directly the parameters of the function without discretization restrictions.

C. A Scalable Approximate SDP Formulation

Although the sigmoid function approximation of the control policy enhances computational efficiency, the state space increases so fast that the problem becomes intractable when the number of states increases due to existence of multiple classes. A reasonable decomposition is to split the regulation service, R , that the SBO sells in the hour ahead market, to class specific R_i such that $\sum_i R_i = R$, and select R_i in proportion to $\bar{n}_i r_i = (\lambda_i^M / U_i^M)(U_i^M - \Pi^E r_i / \mu_i) r_i / \mu_i$. Solutions of the decomposed one class sub-problems were obtained and simulated revealing comparable average cost to the multiclass solution.

IV. HOUR AHEAD AND REAL TIME DECISION RELATIONS

The SBOs participation in the hour-ahead Power market consists of submitting its demand function to purchase energy that corresponds to all demand with utility larger than the hour ahead clearing market Π^E . At the same time it bids a flexible load that does not exceed $2R$, and whose utility, U^E , is in the vicinity of the SBOs estimate of the power markets clearing price⁶, Π^E . The flexible load bid is accompanied by a reservation price for regulation service,

⁶See [16] for an example of how EVs with a given future departure time can determine their utility of charging energy at a time preceding their departure time as the incremental change in the expected cost of charging fully by the departure time with respect to an incremental change in the SoC at an earlier time.




Fig. 1: Market Event Outcomes and Associated Conditional Costs

U^R , and the following constraints are placed on the energy Q^E and regulation service Q^R that the market can schedule: $0 \leq Q^R \leq Q^E \leq 2R$. The ISO clears market participant bids by solving the following optimization problem:

$$\max U^E Q^E - U^R Q^R + \text{other participant terms.}$$

Subject to

- Energy Balance equality constraint
- Regulation Service Reserve inequality constraint
- Capacity constraints taking into consideration the up and down nature of regulation service reserves: $0 \leq Q^R + Q^E \leq 2R$, $Q^R \leq Q^E$

Denoting by Π^E the shadow price of the energy balance constraint by Π^R the shadow price of the R reserve requirements constraint, we have the following market events: If $|U^E - \Pi^E| + U^R \leq \Pi^R$ then the market clears at $Q^R = Q^E = R$. Otherwise, if $U^E \geq \Pi^E$ then $Q^R = 0$, $Q^E = 2R$, while if $U^E \leq \Pi^E$ then $Q^R = 0$, $Q^E = 0$. These events are shown in the Figure 1.

In order to bid appropriately, the SBO must know its unit costs U^R , and it must also select R so as to maximize its hour ahead income $U^R R$ minus its losses during the real time tracking of the regulation signal. More specifically it wants to maximize: $\Pi^R R \Delta t - \{ [(\text{Expected Utility when } Q^R = R) - (\text{Utility rate when } Q^R = 0)] + (\text{Expected Tracking costs when } Q^R = R) \}$ or in terms of Section III results and notation, we want to Maximize $\Pi^R R \Delta t - Z(R)$ over R , where

$$Z(R) = \sum_i \left[\lambda_i^M [(U_i^M)^2 - (\Pi^E r_i / \mu_i)^2] / (2U_i^M) \right. \\ \left. - \mathbb{E}_{u_i} [\lambda_i^M [(U_i^M)^2 - (u_i)^2] / (2U_i^M)] \right] \\ + \mathbb{E} \left[K \left(\sum_i (n_i + \Delta n_i) r_i - \bar{R} - y(t)R \right)^2 \right]$$

The optimization described above can be done systematically for concave $Z(R)$ by finding the value of R at which $\partial Z(R) / \partial R = \Pi^R$. The derivative above cannot be estimated explicitly but can be estimated numerically. In Section V we show numerical evidence that supports the theoretical expectation that $Z(R)$ is indeed convex. The theoretical expectation

of convexity derives from the dependence of $Z(R)$ on the variance of $u_i(t)$ see Proposition 1.4 - and the variance of the tracking error, both increasing with R .

In conclusion, being able to solve the real time control problem is useful in providing the means for solving the hour ahead optimal bid problem. For example, U^R must be selected so that it is equal to $Z(R)/R$

V. NUMERICAL RESULTS

As already mentioned, the Stochastic DP defined in Section III is equivalent after discretization of the control set to an LP [1]. We calibrated the transition rates of the Two-Layer Markov Chain that models ISO signals using actual PJM $y(t)$ time series data available in [4]. As this solution approach requires a discrete control set, the continuous control space, $u \in [0, U^M]$, was discretized into 11 price levels, i.e., $u \in \{0, 0.1U^M, \dots, U^M\}$. Various instances of the single class problem were then solved for. As mentioned in Section III-B, the optimal price policy appears to fit in all cases a sigmoid function representation. We thus proceeded to represent the control policies with a small set of parameters, each representing a continuous sigmoid policy function. The corresponding sigmoid policy function approximations were then allowed to compete in the LP solution. In fact, the LP was formulated with each constraint for each state repeated for each of the sigmoid policy function approximations and the best sigmoid approximation was selected by the optimal LP solution.

In all of the single class problem instantiations/cases reported on in this Section, we used the following parameters values: $K = 0.5$, $r = 1$, $\mu = 1$, $\lambda^M = 150$, $U^M = 50$, $\Pi^E = 1/3U^M$ hence $\bar{R} = \bar{n}r = \bar{n} = 100$.

A. Accuracy of the Two-Layer Markov Chain model

An important factor in our model that makes it more realistic is the accurate performance of the Two-Layer Markov Chain used to model the ISO RS signal. As can be seen in Table I, the averages of the RS signal is invariably very close to zero matching the behavior exhibited in the real data. The variance of the signals is observed to be in the range $[0.143, 0.161]$ in our model, while the variance is 0.154 in the PJM data. Moreover, the fraction of the time that the ISO signal was in the $D = up$ and $D = down$ states was almost equal in all cases (not shown in the table). This fidelity is also observed in various Monte Carlo simulations of the RS signal. In these simulations, the time statistics of the signal states almost exactly matched the real PJM data.

TABLE I: Results of the LP Approach with the Discretized Control Space

Case	R	$\mathbb{E}[u]$	σ_u^2	$\mathbb{E}[r]$	$\mathbb{E}[y]$	σ_y^2
1	5	0.33 U^M	0.104 $(U^M)^2$	100.07	0.014	0.161
2	10	0.34 U^M	0.114 $(U^M)^2$	99.95	-0.005	0.160
3	15	0.33 U^M	0.125 $(U^M)^2$	100.12	0.008	0.154
4	20	0.35 U^M	0.140 $(U^M)^2$	99.8	-0.01	0.155
5	25	0.33 U^M	0.150 $(U^M)^2$	99.8	0.01	0.143




Fig. 2: Optimal Price Structure: Price versus tracking error

B. Structure of the optimal policies

A very important result displayed in Table I is that the average price in the LP solution was almost equal to the energy clearing price $\Pi^E = 1/3U^M$. With this average price, the system would consume the same average energy as purchased in the hour ahead market. In addition, it was observed that the optimal policies exhibit a sigmoidal structure. In Figure 2, the Y axis represents the optimal price levels that correspond to the difference between the number of the customers in the system n and the obligation $\bar{n} + Ry$, represented on the X axis. Each line in the figure represents the optimal price policy for different obligation levels associated with the direction state $D = up$. When the system is in the direction state $D = down$, a similar but yet different graph is obtained. This structure of the optimal policies shows that the optimal control depends on all on n , y and D state components and helps us to determine how to represent the control policies via appropriately parameterized analytic sigmoid functional approximations.

C. Performance of the sigmoid functional representation of control policies

Table II compares LP solutions corresponding to discretized control spaces relative to sigmoid functional approximation of control policies. The columns of the table represents the number of constraints, the number of iterations/base changes to reach the optimal solution and the resulting objective values. It is noteworthy that sigmoid function control policy approximation reduces computational burden significantly through the employment of a smaller number of constraints. In addition the sigmoid function policy representation matches and in some cases better the optimal average cost of the policy space discretization solution. It is important to emphasize that the competing sigmoid function parameter sets were selected judgmentally. Later work has shown that a formal parameter selection algorithm can indeed increase the performance and desirability of the analytic functional policy representation. The superior performance of the continuous functional approximation of the control policy is not surprising and can be understood to result from the removal of the discretization error.

TABLE II: Comparison of Discretized and Sigmoidal Control

Case	R	NoC_D	NoI_D	J_D^*	NoC_S	NoI_S	J_S^*
1	5	19624	6526	1802	7164	2927	1755
2	10	39244	16986	1484	14324	6585	1498
3	15	58864	30291	1120	21484	20492	1081
4	20	78484	55655	560	28644	24527	514
5	25	98104	+150000	-110	35804	39985	-147




Fig. 3: Utility, Price Variance and RS Provision Relationship

Y1 axis: Exp. Utility and Exp. Tracking Cost, Y2 axis: Price Variance, X axis: Reservation Service Provision R

D. Relationship among price variance, utility and RS provision

As shown in Proposition 1.3, an increase in the price variance decreases the utility gain and the loss is $\lambda^M \sigma_u^2 / (2U^M)$. Figure 3 shows that the variance of the price increases as the amount of the RS provision increases (also shown in Table I in exact numerical values). As our Proposition suggests, utility gain is less than what SBO would obtain if it broadcasted a constant price (diamond curve in Figure 3) and the loss exactly matches with the theoretical derivation noted above. Moreover, Figure 4 represents the convexity of $Z(R)$ as described in Section IV.




Fig. 4: $Z(R)$ vs. Reserve Service Provision R

VI. CONCLUSION

We have demonstrated that with a reasonable model of the statistical behavior of each one of multiple load classes, a distributed Load Aggregator can act as a SBO and effectively control the degrees of freedom that flexible loads avail themselves to. The benefits of such control are mutual to the SBO as well as the System Operator and hence to society. We showed that optimal control is computationally tractable and can provide support to the design of optimal bidding strategies to the forward markets where stand by regulation service reserve capacity is transacted.

Work presented here indicates that future work is reasonably expected to be able to address pooling and economies of scale issues, as well as more realistic appliance cluster behavior modeling.

REFERENCES

- [1] D.P. Bertsekas, "Dynamic Controlling and Optimal Control", *Athena Scientific*, 1995, vol. II
- [2] Y. V. Makarov, C. Loutan, J. Ma, and P. de Mello, "Operational impacts of wind generation on California power systems", *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 10391050, 2009.
- [3] *White Paper on Integrating Demand and Response into the PJM Ancillary Service Markets*, February 2005.
- [4] <http://pjm.com/markets-and-operations/ancillary-services/mkt-based-regulation.aspx>
- [5] NYISO Ancillary Services Manual, 2011, www.nyiso.com
- [6] M. Bryson, *PJM Manual 12: Balancing Operations, Revision 16*, November 2007.
- [7] B. Kranz, R. Pike, and E. Hirst, "Integrated electricity markets in New York", *The Electricity Journal*, vol. 16, no. 2, pp. 5465, 2003.
- [8] *NYISO Day-Ahead Scheduling Manual 11*, June 2001
- [9] A. L. Ott, "Experience with PJM market operator, system design, and implementation", *IEEE Transactions on Power Systems*, no. 2, pp. 528 534, 2003.
- [10] D.L. Hammerstrom et al. "Pacific Northwest GridWise Testbed Demonstration Projects Part I. Olympic Peninsula Project"
- [11] D.P. Chassin and J.C. Fuller, "On the Equilibrium Dynamics of Demand Response", *Procs. of HICSS 44*, January 2011.
- [12] R.G. Pratt and Z.T. Taylor, "Development and testing of equivalent thermal parameters calculated from hourly building performance data", *In procs. of 1988 ACEEE Summer Study on Energy Efficiency in Buildings*, 10, 268-285, 1988.
- [13] N. Lu; D.P. Chassin, "A state-queueing model of thermostatically controlled appliances", *IEEE Transactions on Power Systems*, 19:3, 1666-1673, Aug. 2004.
- [14] Y.C. Paschalidis, B. Li, M. C. Caramanis "Demand-Side Management for Regulation Service Provisioning through Internal Pricing", *IEEE Transactions on Power Systems*, In Press
- [15] Y.C. Paschalidis, B. Li, M. C. Caramanis "Market-Based Mechanism for Providing Demand-Side Regulation Service Reserves", *Proceedings of the 50th CDC*, pp. 21-26, Dec 2011.
- [16] J.M. Foster, M.C. Caramanis, "Energy Reserves and Clearing in Stochastic Power Markets: The Case of Plug-In-Hybrid Electric Vehicle Battery Charging, Proceedings of 49th IEEE Conference on Decision and Control, pp.1037-1044, Dec 2010
- [17] M. C. Caramanis, J. M. Foster, "Uniform and Complex Bids for Demand Response and Wind Generation Scheduling in Multi-Period Linked Transmission and Distribution Markets", *Proceedings of the 50th CDC*, pp. 4340-4347, Dec. 2011
- [18] M.C. Caramanis, J.M. Foster, "Coupling of Day Ahead and Real Time Power Markets for Energy and Reserves Incorporating Local Distribution Network Costs and Congestion", *Proceedings, 48th annual Allerton Conference on Communication, Control and Computing*, September 28-October 1, 2010, pp. 42-49.
- [19] C.G. Cassandras and S. Lafortune, "Introduction to Discrete Event Systems", *New York: Springer Science and Business Media*, 2010
- [20] I.C. Paschalidis and J.N. Tsitsiklis, "Congestiondependent pricing of network services", *IEEE ACM Trans. Networking*, vol. 8, no. 2, pp.171184, 2000