

Optimal control of end-user energy storage

Peter M. van de Ven, Nidhi Hegde, Laurent Massoulié and Theodoros Salonidis, *Member, IEEE*

Abstract—An increasing number of retail energy markets show price fluctuations, providing users with the opportunity to buy energy at lower than average prices. We propose to temporarily store this inexpensive energy in a battery, and use it to satisfy demand when energy prices are high, thus allowing users to exploit the price variations without having to shift their demand to the low-price periods. We study the battery control policy that yields the best performance, i.e., minimizes the total discounted costs. The optimal policy is shown to have a threshold structure, and we derive these thresholds in a few special cases. The cost savings obtained from energy storage are demonstrated through extensive numerical experiments, and we offer various directions for future research.

Index Terms—Battery storage, dynamic pricing, dynamic programming, energy storage, Markov decision processes, threshold policy.

I. INTRODUCTION

Wholesale energy prices exhibit significant fluctuations during each day due to variations in demand and generator capacity. End users are traditionally not exposed to these fluctuations but pay a fixed retail energy price. Economists have long argued to remove the fixed retail prices in favor of prices that change during the day. Such *dynamic pricing* would better reflect the prices on the wholesale market and has been predicted to lead to lower demand peaks and to a lower and less volatile wholesale price [2]. Implementations of dynamic pricing have been enabled by recent developments in smart-grid technology such as smart meters.

An example of an increasingly popular dynamic pricing scheme is time-of-use pricing. Such pricing typically provides two or three price levels (e.g., ‘off-peak’, ‘mid-peak’ and ‘on-peak’) where the relevant level depends on the time of day. The price levels are determined well in advance and are only changed once or twice per year. A second example of dynamic pricing is real-time pricing where the retail energy price changes hourly or half-hourly and are based on the price on the wholesale energy market.

Dynamic pricing creates an opportunity for users such as households and data centers to exploit the price fluctuations and reduce energy costs. However, doing so would require users to shift their demands to low-price periods, and in practice users only show a minor shift in their demand in response to changes in the energy prices [3]–[6]. A possible solution is to equip users with a battery that can be used for energy storage; the battery can be charged when the energy price is low and the stored energy can then be used when the price is high. This allows users to benefit from the energy price

variations without having to adjust their consumption. Energy can be stored both by a dedicated battery, or by existing storage such as the battery pack of an electric car [7] (residential users) or a backup power supply [8] (data centers).

In this paper, we address the problem of organizing energy storage purchases to minimize long-term energy costs under variable demands and prices. This problem involves deciding whether to satisfy demand directly from the grid or from the battery, as well as up to what level to charge or discharge the battery. The resulting optimization problem is complicated by the stochastic nature of price and demand and due to the fact that we aim to minimize the long-term costs. We model the problem as a Markov Decision Process and show that there exists a two-threshold stationary cost-minimizing policy. When the battery level is below the lower threshold, the battery is charged up to it, and the battery is discharged when above the upper threshold. By comparing the costs incurred under this policy with the cost of satisfying all demand directly from the grid, we can show that energy storage may lead to significant cost savings.

Residential energy storage has been studied for the case of *arbitrage*, i.e., buying energy when it is inexpensive, and selling it later to the grid for a higher price [9]. This problem has been studied assuming that prices are known in advance in a finite horizon setting. These assumptions allow deterministic optimization problem formulations which can be solved using linear programming techniques [10], [11]. However, such approach does not take into account the stochasticity in prices and demands and does not allow for long-term cost optimization. In [12] the authors consider the problem of energy storage from the point of view of the grid operator, and propose a threshold policy that is shown to be asymptotically optimal as the size of the storage unit increases.

A model similar to ours was used to investigate control of energy storage in the context of *data centers* [8]. The model in [8] assumes that the battery is fully efficient and the proposed scheduling algorithm is a sub-optimal heuristic, whose gap from optimality increases as storage size decreases. In [13] this approach is extended to multiple data centers, each with different time-varying prices. In contrast to [8] and [13], our model incorporates battery inefficiencies and we investigate the optimal scheduling policy.

A related problem is optimal control of energy storage for renewable energy. The two-price case is considered in [14], while the more general case is discussed in [15] and [16]. The case without transmission costs in [15] is closely related to our setting, and a similar two-threshold policy is shown to be optimal. In recent work [16], the authors consider a similar model to ours to address storage control for renewable energy generation for general prices, and show that the average cost optimal policy has a similar threshold structure. The authors

Peter M. van de Ven and Theodoros Salonidis are with the IBM Thomas J. Watson Research Center, Nidhi Hegde is with Technicolor Paris Research Lab and Laurent Massoulié is with INRIA, MSR-Inria Joint Centre.

A preliminary version of this paper has appeared as [1].

of [15] and [16] use a finite-horizon and infinite horizon average criterion, respectively. The battery and price models differ slightly, and the work in [16] accounts for dissipation losses, but does not allow state-dependent charging constraints. In contrast to [16], our approach incorporates periodicity, which allows us to model daily fluctuations in price and demand.

Our model is closely related to periodic-review, single-item inventory models, and the optimal policy mirrors the optimality of the base-stock policy for the backlog case. However, in our case the demand is known before the purchasing decision, and we require that all demand is met in the time slot that it arises. In addition, the state description is continuous, and the battery inefficiency fundamentally changes the structure and analysis of the optimal policy.

The remainder of the paper is structured as follows. In Section II we introduce the model and describe the decision variables. In Section III we demonstrate the optimality of a threshold policy, and in Section IV we derive some properties of the thresholds. Section V discusses various numerical examples and Sections VI and VII describe some directions for future research and provides concluding remarks, respectively.

II. MODEL

Consider a user with certain energy requirements and a battery that can be used for energy storage. Time is slotted, and we denote by $B(t)$ the battery level (state of charge) in kWh at time t , $t = 0, 1, \dots$. Let \bar{B} represent the maximum battery level, and $\mathcal{B} = [0, \bar{B}]$ the range of all possible battery levels, so

$$B(t) \in \mathcal{B}. \quad (1)$$

In each time slot t some demand $D(t)$ arises, and we may purchase energy at a price of $P(t)$ per unit. The demand has some compact support $D(t) \in \mathcal{D}$, as does the price $P(t) \in \mathcal{P}$. Both are bounded, and we denote by \bar{D} and \bar{P} the maximum demand and price, respectively, so $\mathcal{D} \subseteq [0, \bar{D}]$ and $\mathcal{P} \subseteq [p_{\min}, \bar{P}]$, with p_{\min} the minimum price. Finally, we denote by \mathcal{M} the set of modulating states, that influence the price and demand transitions. For example the time of day or the season.

Denote by $\Omega = \mathcal{D} \times \mathcal{P} \times \mathcal{M}$ the set of possible realizations of demand, price and modulating state, and for any $\mathbf{x} \in \Omega$, denote by $d(\mathbf{x})$ and $p(\mathbf{x})$ the corresponding price and demand, respectively. Demand and price may be correlated, and we denote by $f_{\mathbf{x}}(\mathbf{y})$ the probability density function of moving from state \mathbf{x} to state \mathbf{y} in the next slot, for any $\mathbf{x}, \mathbf{y} \in \Omega$.

The battery may not be completely efficient, and its performance is affected by the charging efficiency $\eta_c \in (0, 1]$ and discharging efficiency $\eta_d \in (0, 1]$. Energy purchased to charge the battery is reduced by a factor η_c , and only a fraction η_d of the discharged energy is converted into electricity. This model is general and encompasses for example batteries of electric vehicles, uninterrupted power supplies of data centers and batteries dedicated to end user storage. The model is not intended to capture all the subtleties of battery behavior in each of these applications but rather the essential tradeoffs and phenomena that arise in practice. In Section VI we

discuss various model refinements that can be made without fundamentally altering the results or the derivations.

In addition to satisfying the demand from the battery, we also allow demand to be met directly from the grid, bypassing the battery. Let $A_1(t)$ denote the amount of energy purchased directly from the grid in slot t , $A_2(t)$ the amount of energy bought to charge the battery, and $A_3(t)$ the energy discharged from the battery used towards satisfying demand, see Figure 1.




Figure 1. A graphical representation of the model.

We assume

$$A_1(t), A_2(t), A_3(t) \geq 0, \quad (2)$$

and require that all demand must be met, i.e.,

$$D(t) = A_1(t) + \eta_d A_3(t). \quad (3)$$

The battery has state-dependent charging constraints $\bar{A}_c(b)$ and $\bar{A}_d(b)$, so the amount of energy that can be charged to and discharged from the battery is bounded as

$$A_2(t) \leq \bar{A}_c(b), \quad A_3(t) \leq \bar{A}_d(b). \quad (4)$$

The battery level of the battery evolves according to

$$B(t+1) = B(t) + \eta_c A_2(t) - A_3(t), \quad (5)$$

and the energy costs in slot t are given by

$$g(t) = (A_1(t) + A_2(t))P(t). \quad (6)$$

We are interested in the total discounted costs $\sum_{t=0}^{\infty} g(t)\alpha^t$, with $0 < \alpha < 1$ the discount factor that represents value reduction over time. The reason for considering discounted rather than average costs is to emphasize early decisions and costs, in order to emulate the effect of reduced battery efficiency over time. Note that the total discounted costs are finite, since the per-slot costs are bounded. Our goal is now to choose in each slot the $A_1(t)$, $A_2(t)$ and $A_3(t)$ that solve the following optimization problem:

$$\begin{aligned} \min \mathbb{E} \left\{ \sum_{t=0}^{\infty} g(t)\alpha^t \right\} \\ \text{subject to : (1), (2), (3), (4), (5)} \end{aligned} \quad (7)$$

The infinite horizon problem belongs to a class of stochastic optimization problems, which in general are difficult to solve. A first step towards dealing with (7) is to see that it is never optimal to charge and discharge the battery in the same slot. This is intuitively clear, because charging and discharging the battery simultaneously corresponds to routing $\min\{A_2(t), \eta_d A_3(t)\}$ energy from the grid to the user, through the battery. Because of the battery inefficiency it is beneficial to instead circumvent the battery, and satisfy the demand directly from the grid.

It turns out that this observation significantly simplifies the optimization problem, by reducing the number of decision variables. Specifically, in view of the restriction on simultaneous charging and discharging:

$$\begin{aligned} A_1(t) &= D(t) + (B(t+1) - B(t))\eta_d \mathbf{1}_{\{B(t+1) < B(t)\}}, \\ A_2(t) &= (B(t+1) - B(t))\eta_c^{-1} \mathbf{1}_{\{B(t+1) > B(t)\}}, \\ A_3(t) &= -(B(t+1) - B(t))\eta_d \mathbf{1}_{\{B(t+1) < B(t)\}}. \end{aligned}$$

Thus, the choice for $B(t+1)$ fixes $A_1(t)$, $A_2(t)$ and $A_3(t)$, and (7) reduces to a single-variable decision problem. The per-slot costs may be rewritten in terms of $B(t+1)$ as

$$\begin{aligned} g(t) &= (D(t) + (B(t+1) - B(t))^+ \eta_c^{-1} \\ &\quad + (B(t+1) - B(t))^- \eta_d) P(t), \end{aligned}$$

with $(x)^+ = \max\{x, 0\}$ and $(x)^- = -\max\{-x, 0\}$. Note that (4) can be expressed in $B(t+1)$ as

$$-\bar{A}_d(B(t)) \leq B(t+1) - B(t) \leq \bar{A}_c(B(t))\eta_c.$$

III. THE STRUCTURE OF THE OPTIMAL POLICY

In this section we discuss how to choose in each slot the value for $\beta = B(t+1)$ that minimizes the total discounted costs. To this end, we rewrite our model as a Markov decision process. We denote by $J_{\mathbf{x}}(b)$ the minimal total discounted costs for a battery differential δ , starting from state $\mathbf{x} \in \Omega$, and battery level $b \in \mathcal{B}$. Let $\gamma_{\mathbf{x}}(\delta) = (d(\mathbf{x}) + (\delta)^+ \eta_c^{-1} + (\delta)^- \eta_d) p(\mathbf{x})$ denote the immediate costs, $G_{\mathbf{x}}(\beta) = \int_{\mathbf{y} \in \Omega} f_{\mathbf{x}}(\mathbf{y}) J_{\mathbf{y}}(\beta) d\mathbf{y}$ and define

$$H_{\mathbf{x}}(\beta, b) = \gamma_{\mathbf{x}}(\beta - b) + \alpha G_{\mathbf{x}}(\beta),$$

the total discounted costs given battery level b , state \mathbf{x} and action β . Then the cost function satisfies the Bellman equation

$$J_{\mathbf{x}}(b) = \inf_{\beta \in U_{\mathbf{x}}(b)} H_{\mathbf{x}}(\beta, b), \quad (8)$$

with $U_{\mathbf{x}}(b)$ the control set that contains all feasible decisions for $B(t+1)$. This set may be written as

$$U_{\mathbf{x}}(b) = [U_{\mathbf{x}}^-(b), U_{\mathbf{x}}^+(b)],$$

where $U_{\mathbf{x}}^-(b) = \max\{0, b - d(\mathbf{x}), b - \bar{A}_d(b)\}$ and $U_{\mathbf{x}}^+(b) = \min\{\bar{B}, b + \bar{A}_c(b)\}$.

It is readily verified that the infimum in (8) can be attained by a stationary optimal policy, see, e.g., [17, Proposition 4.4]. We shall demonstrate that this optimal policy specifies for each state $\mathbf{x} \in \Omega$ two battery thresholds $\beta_{\mathbf{x}}^-, \beta_{\mathbf{x}}^+ \in \mathcal{B}$, $\beta_{\mathbf{x}}^- \leq \beta_{\mathbf{x}}^+$ such that the cost-minimizing choice for the battery level $\beta_{\mathbf{x}}^*(b)$ is given by

$$\beta_{\mathbf{x}}^*(b) = \begin{cases} \min\{\beta_{\mathbf{x}}^-, U_{\mathbf{x}}^+(b)\}, & b \leq \beta_{\mathbf{x}}^-, \\ \max\{\beta_{\mathbf{x}}^+, U_{\mathbf{x}}^-(b)\}, & b \geq \beta_{\mathbf{x}}^+, \\ b, & \text{otherwise.} \end{cases} \quad (9)$$

So if $b \leq \beta_{\mathbf{x}}^-$, then the optimal policy is to charge the battery as close to $\beta_{\mathbf{x}}^-$ as the control set allows. If $b \geq \beta_{\mathbf{x}}^+$, then the battery should be discharged down to $\beta_{\mathbf{x}}^+$ within the boundaries of the control set. In case $\beta_{\mathbf{x}}^- < b < \beta_{\mathbf{x}}^+$, it is optimal to neither charge nor discharge the battery. So if the battery level is sufficiently low, the battery is charged, and all demand in

satisfied directly from the grid, while for high battery level, demand is (partially) satisfied from the battery. When the battery level is between both thresholds, the battery is neither charged nor discharged, and all demand is met from the grid.

In order to show that (9) indeed describes the structure of the optimal policy, we require the following lemma.

Lemma 1: The cost function $J_{\mathbf{x}}(b)$ is convex and non-increasing in b .

The proof of Lemma 1 relies on the fact that the total discounted costs can be viewed as the limit of finite-horizon discounted costs, which may be shown to possess these properties by induction on the horizon. The proof can be found in the appendix, along with the other proofs.

We are now in position to state and prove our main result.

Theorem 1: The policy that solves the minimization problem (7) is of the form $\beta = \beta_{\mathbf{x}}^*$ as in (9).

By Lemma 1 we know that the right-hand side of the Bellman equation (8) in fact defines a convex optimization problem, the solution of which can be found by finding the $\beta \in U_{\mathbf{x}}(b)$ for the which the subdifferential $\partial H_{\mathbf{x}}(\beta, b)$ has the proper shape. The proof of Theorem 1 then relies on close inspection of this subdifferential.

In case the battery is fully efficient, the charging threshold and discharging threshold are identical, as is shown in the next corollary.

Corollary 1: Let $\eta_c = \eta_d = 1$, then the optimal policy is as in (9), with $\beta_{\mathbf{x}}^- = \beta_{\mathbf{x}}^+$.

In Section VI we present several model extensions under which Theorem 1 remains valid.

The optimal policy (9) is illustrated in Figure 2, which shows $\beta_{\mathbf{x}}^*(b)$ plotted against b . The diagonal segment in the middle corresponds to $\beta_{\mathbf{x}}^*(b) = b$, while the two horizontal lines represents the thresholds $\beta_{\mathbf{x}}^-$ and $\beta_{\mathbf{x}}^+$. The outer diagonal segments represent the boundaries of the control space $U_{\mathbf{x}}(b)$. Both horizontal lines coincide in the scenario with a completely efficient battery ($\eta_d \eta_c = 1$).




Figure 2. The structure of the optimal policy $\beta_{\mathbf{x}}^*$ as a function of b .

IV. BATTERY LEVEL THRESHOLDS

In the previous section we have established the threshold-based structure of the optimal policy (9). Analytically deter-

mining the thresholds β_x^- and β_x^+ is a difficult problem in general, and in this section we present some structural results and results for special cases. For simplicity, we limit ourselves to the case $\eta_d = \eta_c = 1$, and denote $\beta_x = \beta_x^- = \beta_x^+$. Similar properties can be shown to hold for the inefficient case.

We first present sufficient conditions for the thresholds to be equal to either 0 or \bar{B} .

Proposition 1: Let $\mathbf{x} \in \Omega$. If for all $b_1, b_2 \in \mathcal{B}$, $b_1 < b_2$, $\mathbf{y} \in \Omega$

$$J_{\mathbf{y}}(b_1) - J_{\mathbf{y}}(b_2) \leq (b_2 - b_1) \frac{p(\mathbf{x})}{\alpha},$$

then $\beta_x = 0$. If for all $b_1, b_2 \in \mathcal{B}$, $b_1 < b_2$, $\mathbf{y} \in \Omega$

$$J_{\mathbf{y}}(b_1) - J_{\mathbf{y}}(b_2) \geq (b_2 - b_1) \frac{p(\mathbf{x})}{\alpha},$$

then $\beta_x = \bar{B}$.

This result can for example be used to show that $\beta_x = 0$ in certain special cases, as is described in the following proposition.

Proposition 2: Let $\mathbf{x} \in \Omega$ and denote $p_{\min} = \min_{p \in \mathcal{P}} p$, then $\beta_x = 0$ if either:

- (i) $p(\mathbf{x}) = \bar{P}$;
- (ii) $p_{\min} > 0$ and $\alpha < p_{\min}/\bar{P}$.

Proposition 2 states that if the price is very high, or the discount factor is sufficiently low, it is optimal not to charge the battery at all, and to try to satisfy the demand from the battery as much as possible.

In case that the state transitions are i.i.d. ($f_x \equiv f$) or if the transition probabilities are determined by the price level ($f_x \equiv f_{p(x)}$), the thresholds β_x depend on the price only and are independent from the demand and modulating state. In this case, write $\beta_x = \beta_{p(x)}$. We can show that for i.i.d. prices, the thresholds are decreasing in the price level.

Proposition 3: Assume that the prices and demands are i.i.d. across time, then $\beta_{p(x)}$ is decreasing in p .

The monotonicity observed in Proposition 3 is very intuitive, but does not extend to Markovian prices. The reason is that threshold values are partially determined by the evolution of the price: even for a low price it might be beneficial to set a low threshold, if the price in the next slot is even lower. Such an example is presented below. For ease of presentation we use a discrete probability distribution f in this example, noting that similar examples can be constructed using continuous transition probabilities.

Example 1: Consider an example with four price levels $p_i = i$, $i = 1, \dots, 4$. We assume that $\alpha \geq 3/4$, $\bar{B} = 1$ and $D \equiv 1$, and we choose the following price transition probabilities:

$$f_1(1) = f_1(3) = 1/2, \quad f_2(1) = f_3(4) = f_4(2) = 1.$$

The corresponding thresholds are $\beta_1 = \beta_3 = 1$ and $\beta_2 = \beta_4 = 0$, which are clearly not monotone. The derivation of these thresholds is presented in Appendix G.

V. NUMERICAL EVALUATION

In this section we evaluate the operation of the optimal energy storage management policy (9) in residential environments and real-time pricing (RTP) scenarios. Our goal is

to demonstrate the practical feasibility of the optimal policy and the extent of its cost savings under RTP and demand fluctuations that might arise in real life. We first describe the price and demand datasets used for the evaluation and then outline our low complexity implementation of the optimal policy. We evaluate the cost savings of the optimal policy in scenarios of individual home storage units and shared energy storage.

A. Price and demand datasets

We emulate residential RTP data with historical hourly spot prices of the Ontario energy market [18]. Although Ontario currently does not use RTP, the spot prices provide a reasonable RTP estimate. The residential demand data is synthetically generated using the tool in [19]. This tool uses a high-resolution model of domestic electricity usage based on patterns of home occupancy and appliance usage, weather conditions and characteristics of all major appliances commonly found in the domestic environment.

In our approach, the optimal thresholds can be dynamically determined using empirical distributions of historical price and demand data for each time slot of the day. The smallest time slot is determined by the coarsest granularity between price and demand data. In our experiments, we have hourly price data and 1-minute demand data, therefore we use time slots of one hour. We also use a training window of one month. We found that this window size is long enough to provide adequate characterization of the distribution of each hour of the day, and short enough to use the optimal thresholds for the prices and demands that appear in the next window.

For concise presentation, we show results for a representative scenario where a home of four occupants equipped with a battery receives hourly prices from the Ontario energy market during January and February 2011.

Figure 3(a) plots the average, minimum and maximum hourly price in Ontario during January 2011. We observe one active phase with multiple price peaks between 9 a.m. and 10 p.m., and a low price at night. Note that the energy prices are lower at night and display multiple peaks per day. Figure 3(b) shows that the February prices follow a similar trend.




Figure 3. Hourly average, minimum and maximum Ontario spot market prices (Ontario, 2011 dataset).

We first determine the empirical distribution of the January 2011 Ontario hourly energy prices and demand data. The empirical price (demand) distribution for each hour of the day is computed by dividing the number of observations of a specific price (demand) level by the total number (31) of observations. These distributions are used to determine the

thresholds of the optimal policy. Then, we use the February 2011 prices and demands to emulate the operation of the optimal policy and compute the resulting electricity cost.

B. Implementation

Analytical determination of the optimal thresholds is possible for special cases but is difficult in general (see Section IV). Instead, we compute the thresholds numerically using policy iteration. Since policy iteration can be computationally intensive in practice, we have reduced the complexity of our implementation as follows. First, we discretize the state space by rounding the demand data and energy storage level to multiples of 0.5 kWh, and the price data to multiples of 5 ct. Second, we use the hour of the day as the modulating state and assume that prices and demand are independent from hour to hour (but not necessarily identically distributed). Thus, the optimal thresholds depend on price realization and the hour of the day, but are independent of the demand realization. The modulating variable \mathcal{M} could also be used to differentiate between different days and months and to take into account weather conditions. Such detailed state description would likely yield higher cost savings as it would allow more accurate price and demand predictions. On the other hand, it might generate too large a state space for policy iteration to be feasible. The following experiments demonstrate that using our simple choice of the modulating state yields very high cost savings.

We evaluate the optimal policy for a fully efficient battery ($\eta_d = \eta_c = 1$) and no charging constraints ($\bar{A}_c(b) = \bar{A}_d(b) = \infty$) thus obtaining an upper bound on the potential cost savings of energy storage. We implement the policy iteration algorithm in Matlab; each slot in the policy iteration algorithm corresponds to one hour. For the above parameters and a discount factor $\alpha = 0.99$, a laptop PC with a quad-core 2.2 GHz Intel processor and 8 GB RAM typically requires approximately 5 min to compute the optimal thresholds.

C. Energy savings

Figure 4 shows the relative energy savings of the optimal policy over the setting without storage, as a function of the battery size B_{\max} . Three important observations are in place. First, the savings increase with battery size to up to 38%, which is significant. Second, the savings reach their maximum at $B_{\max} = 16$ (kWh); increasing battery size beyond this point does not increase savings, as the optimal policy will not utilize any battery capacity beyond 16 kWh. This saturation point can be explained by the fact that the value of stored energy decreases over time due to the discounting of the costs and the cyclic price and demand levels. Third, the size of a typical hybrid vehicle battery pack is in the order of 16 kWh [7]. This suggests that car battery packs are well-suited for home energy storage since their size corresponds to the amount of storage required by the optimal energy storage policy.

D. Structure of the optimal policy

In the remainder of this section we consider a battery with capacity of 16 kWh. To illustrate how the optimal storage




Figure 4. Energy cost savings of optimal storage (over no energy storage) vs. battery size (Ontario, February 2011 dataset).

policy works, we show in Figure 5 the thresholds as a function of time, for two price levels (15 ct and 25 ct). We observe that the thresholds peak early in the morning when the price is low. Moreover, the lower 15 ct price level yields lower thresholds at all hours, in accordance with Proposition 3. Due to low demand and several consecutive hours with low prices, the 15 ct threshold drops to zero at 11 p.m., and increases again the next hour.




Figure 5. Optimal thresholds vs hour of day, for 15 ct/kWh (black) and 25 ct/kWh (gray) price levels.

Figure 6 shows the average amount of energy bought, plotted against the hour of the day, with optimal energy storage (black line) and without storage (gray line). We observe that storage allows users to purchase energy early in the morning when the price is low, while users without storage are forced to buy energy during peak hours, at higher prices.

E. Resource pooling

An alternative to energy storage for individual households is to pool storage capacity: rather than individual users each having a small battery, further cost savings might be achieved by multiple users sharing a single large battery. Figure 7 compares the case of n users each with a 16 kWh battery to the case of a shared $16 \times n$ kWh battery. We use the tool from [19] to generate distinct demand data for each individual home user. In the scenario with storage but without pooling, each home has its own storage, and each set of demand data corresponds to one energy storage unit. In this case,




Figure 6. The average amount of energy bought, with (black) and without (gray) storage.

the thresholds are computed for each individual user. In the shared storage scenario, the aggregate demand data of all n homes is input to the shared storage unit. Since this scenario only has a single storage unit (of size $16n$), we only have to compute the optimal thresholds once. Figure 7 compares the aggregate monthly costs without storage (black), storage without pooling (gray) and storage with pooling (blue) plotted against the number of users.




Figure 7. The effect of resource pooling.

We observed similar performance by running the same experiment with artificially time-shifting the demands of these users to be negatively correlated (results not shown for brevity). This can be explained by the fact that the prices are the same for all users. Thus, irrespective of the size of the battery, the behavior of the optimal policy is primarily influenced by the common pricing signal, eliminating the potential benefits of pooling.

VI. EXTENSIONS & OUTLOOK

Our model is a first step on energy storage management and focuses on a single user that utilizes storage to minimize its own costs. The ideas and results can also be applied in multi-user settings where the objective is to minimize the grid peak load. For example, in a scenario with a large population,

a certain fraction of battery-equipped users could apply the optimal control policy (9) to shift their demand from peak to off-peak hours. Assessing the impact of such an approach is an interesting topic of future research.

The model can be extended in several directions to cover a broader variety of applications, and more realistic battery models. Below we briefly describe three such extensions.

A. Battery replacement costs

In order to assess the cost effectiveness of energy storage, one has to incorporate battery replacement costs. Determining the optimal policy in this case requires solving a joint problem of battery dimensioning and energy storage management for cost minimization or battery lifetime maximization. Smaller batteries are cheaper but, as shown in Figure 4, they may provide less opportunity to exploit price fluctuations. The optimal size of the battery will most likely depend on the spread and volatility of energy prices and may differ between energy markets [20].

As a first-order approximation of the battery replacement costs, we may assume that the battery breaks down after a geometric number of operations (charging/discharging), at which point the battery needs replacement at costs C . Under these assumptions, the analysis presented in this paper largely holds, with some minor modifications. Assuming that battery lifetime has mean $1/q$, the immediate costs can be rewritten as

$$\gamma_{\mathbf{x}}(\delta) = (d(\mathbf{x}) + (\delta)^+ \eta_c^{-1} + (\delta)^- \eta_d) p(\mathbf{x}) + qC \mathbf{1}_{\{\delta \neq 0\}}.$$

This new cost function is convex everywhere but in $\delta = 0$, and Lemma 1 and Theorem 1 must be modified accordingly.

B. Self-discharge and varying efficiency

We may extend the battery model by including the effects of self-discharge and time-varying efficiency. That is, every time slot we assume that some fraction ξ of the stored energy dissipates, and the storage evolution (5) is modified accordingly:

$$B(t+1) = \xi B(t) + \eta_c(t) A_2(t) - A_3(t).$$

The rest of the paper can be adjusted in a similar fashion, and all results and derivations continue to hold. The case $\xi = 1$ and $\eta_c(t) \equiv \eta_c$ corresponds to the present model.

C. Local energy generation

Another straightforward extension of our model is to assume that the end user itself generates some energy. This energy is either used to satisfy demand, is sold back to the grid or stored in the battery. This extension allows us to model for example wind energy farms, where energy storage can be used to successfully incorporate the increasing amount of renewable energy into the grid. Our results then extend to describe an optimal policy for the joint decision on buying, selling and storing (renewable) energy.

In order to make these adjustments, we must reinterpret $D(t)$ as the demand minus the amount of energy generated in a slot, and remove the constraints (2) that state that $A_1(t)$, $A_2(t)$ and $A_3(t)$ are non-negative.

VII. CONCLUSIONS

In this paper we studied the control of end-user energy storage under price fluctuations. We derived the structure of the cost-minimizing storage policy, which turns out to be a simple threshold-based policy (9). We described the behavior of these thresholds for some special cases, and showed by means of a numerical study that energy storage can lead to significant cost savings. We discussed various model extensions and generalizations that would broaden the scope of this work, without fundamentally altering the results.

APPENDIX

A. Proof of Lemma 1

Proof: Let $J_{\mathbf{x},n}(b)$ denote the minimal n -step discounted costs, starting from state \mathbf{x} and battery level b . Let $H_{\mathbf{x},n}(\beta, b) = \gamma_{\mathbf{x}}(\beta - b) + \alpha G_{\mathbf{x},n}(\beta)$ and $G_{\mathbf{x},n}(\beta) = \int_{\mathbf{y} \in \Omega} f_{\mathbf{x}}(\mathbf{y}) J_{\mathbf{y},n}(\beta) d\mathbf{y}$, then $J_{\mathbf{x},n}(b)$ satisfies

$$J_{\mathbf{x},n}(b) = \min_{\beta \in U_{\mathbf{x}}(b)} H_{\mathbf{x},n-1}(\beta, b). \quad (10)$$

These finite-horizon costs converge as $\lim_{n \rightarrow \infty} J_{\mathbf{x},n} = J_{\mathbf{x}}$, cf. [17, Proposition 4.4]. Thus, in order to show that $J_{\mathbf{x}}$ is convex and non-increasing, it suffices to show by induction that this holds for all $J_{\mathbf{x},n}$.

This statement trivially holds for $n = 0$, since $J_{\mathbf{x},0} \equiv 0$. Now let $n \in \mathbb{N}$, and assume that $J_{\mathbf{x},n-1}$ is convex and non-increasing. The operator on the right-hand side of (10) can be identified as the infimal convolution operator, which preserves convexity [21, Theorem 5.4]. We have that $G_{\mathbf{x},n-1}$ is convex by the induction hypothesis, and since it can be readily verified that $\gamma_{\mathbf{x}}$ is convex as well, so is $J_{\mathbf{x},n}$.

Let $b_1, b_2 \in \mathcal{B}$, $b_1 \leq b_2$, then in order to establish that $J_{\mathbf{x},n}$ is non-increasing, we need to show that $J_{\mathbf{x},n}(b_1) \geq J_{\mathbf{x},n}(b_2)$. Denote by $\beta_{\mathbf{x}}^*$ the choice for β he achieves the minimum in (10). We distinguish two cases: (i) $\beta_{\mathbf{x}}^*(b_1) \in U_{\mathbf{x}}(b_2)$; and (ii) $\beta_{\mathbf{x}}^*(b_1) < U_{\mathbf{x}}(b_2)$. In case (i):

$$\begin{aligned} J_{\mathbf{x},n}(b_2) &\leq H_{\mathbf{x},n-1}(\beta_{\mathbf{x}}^*(b_1), b_2) \\ &= \gamma_{\mathbf{x}}(\beta_{\mathbf{x}}^*(b_1) - b_2) + \alpha G_{\mathbf{x},n-1}(\beta_{\mathbf{x}}^*(b_1)) \\ &\leq \gamma_{\mathbf{x}}(\beta_{\mathbf{x}}^*(b_1) - b_1) + \alpha G_{\mathbf{x},n-1}(\beta_{\mathbf{x}}^*(b_1)) = J_{\mathbf{x},n}(b_1), \end{aligned}$$

by the fact that $\gamma_{\mathbf{x}}$ is increasing.

In case (ii) we have by the induction hypothesis

$$\begin{aligned} J_{\mathbf{x},n}(b_2) &\leq H_{\mathbf{x},n-1}(U_{\mathbf{x}}(b_2), b_2) \\ &= \gamma_{\mathbf{x}}(-d(\mathbf{x})) + H_{\mathbf{x},n-1}(U_{\mathbf{x}}(b_2)) \\ &\leq \gamma_{\mathbf{x}}(\beta_{\mathbf{x}}^*(b_1) - b_1) + H_{\mathbf{x},n-1}(\beta_{\mathbf{x}}^*(b_1)) = J_{\mathbf{x},n}(b_1), \end{aligned}$$

completing the proof. \blacksquare

B. Proof of Theorem 1

Proof: Since $\gamma_{\mathbf{x}}$ and $J_{\mathbf{x}}$ are convex, so is $H_{\mathbf{x}}$, and (8) can be recognized as a convex optimization problem. Thus, $\beta^* \in U_{\mathbf{x}}(b)$ is a global minimizer if and only if there exists some subgradient $g \in \partial H_{\mathbf{x}}(\beta^*, b)$ such that for all $\beta \in U_{\mathbf{x}}(b)$,

$$g(\beta^* - \beta) \geq 0.$$

Thus, $\beta^* \in (U_{\mathbf{x}}(b), U_{\mathbf{x}}^+(b))$ is a minimizer iff $0 \in \partial H_{\mathbf{x}}(\beta^*, b)$, while $\beta^* = U_{\mathbf{x}}(b)$ and $\beta^* = U_{\mathbf{x}}^+(b)$ are minimizers iff there exists some $g \in \partial H_{\mathbf{x}}(\beta^*, b)$ such that $g \geq 0$ and $g \leq 0$, respectively.

The subdifferentials of $\gamma_{\mathbf{x}}$ and $G_{\mathbf{x}}$ are given by

$$\partial \gamma_{\mathbf{x}}(\delta) = \begin{cases} \eta_d p(\mathbf{x}), & \delta < 0, \\ [\eta_d p(\mathbf{x}), \eta_c^{-1} p(\mathbf{x})], & \delta = 0, \\ \eta_c^{-1} p(\mathbf{x}), & \delta > 0, \end{cases}$$

and

$$\partial G_{\mathbf{x}}(\beta) = [\sigma_{\mathbf{x}}^-(\beta), \sigma_{\mathbf{x}}^+(\beta)],$$

for some $-\infty < \sigma_{\mathbf{x}}^-(\beta) \leq \sigma_{\mathbf{x}}^+(\beta) \leq 0$. The subdifferential of $H_{\mathbf{x}}$ can then be written as (cf. [21, Theorem 23.8]):

$$\begin{aligned} \partial H_{\mathbf{x}}(\beta, b) &= \begin{cases} [\eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^-(\beta), \eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^+(\beta)], & \beta < b, \\ [\eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^-(\beta), \eta_c^{-1} p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^+(\beta)], & \beta = b, \\ [\eta_c^{-1} p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^-(\beta), \eta_c^{-1} p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^+(\beta)], & \beta > b. \end{cases} \end{aligned}$$

Consequently, $\beta^* = U_{\mathbf{x}}(b)$ is optimal if

$$\eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^+(\beta^*) \geq 0.$$

Let $\beta^* \in (U_{\mathbf{x}}(b), b)$, then β^* is optimal if

$$\eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^-(\beta^*) \leq 0 \leq \eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^+(\beta^*).$$

Let $\beta^* = b$, then β^* is optimal if

$$\eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^-(\beta^*) \leq 0 \leq \eta_c^{-1} p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^+(\beta^*).$$

Let $\beta^* \in (b, U_{\mathbf{x}}^+(b))$, then β^* is optimal if

$$\eta_c^{-1} p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^-(\beta^*) \leq 0 \leq \eta_c^{-1} p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^+(\beta^*).$$

Finally, let $\beta^* = U_{\mathbf{x}}^+(b)$, then β^* is optimal if

$$\eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^-(\beta^*) \leq 0.$$

Denote

$$\begin{aligned} B_{1,\mathbf{x}}^- &= \{\beta \in [0, \bar{B}] : \eta_c^{-1} p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^+(\beta) \geq 0\}, \\ B_{2,\mathbf{x}}^- &= \{\beta \in [0, \bar{B}] : \eta_c^{-1} p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^-(\beta) \leq 0\}, \\ B_{1,\mathbf{x}}^+ &= \{\beta \in [0, \bar{B}] : \eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^+(\beta) \geq 0\}, \\ B_{2,\mathbf{x}}^+ &= \{\beta \in [0, \bar{B}] : \eta_d p(\mathbf{x}) + \alpha \sigma_{\mathbf{x}}^-(\beta) \leq 0\}, \end{aligned} \quad (11)$$

and define

$$\begin{aligned} \beta_{1,\mathbf{x}}^- &= \begin{cases} \min B_{1,\mathbf{x}}^-, & \text{if } B_{1,\mathbf{x}}^- \neq \emptyset, \\ \bar{B}, & \text{otherwise,} \end{cases} \\ \beta_{2,\mathbf{x}}^- &= \begin{cases} \max B_{2,\mathbf{x}}^-, & \text{if } B_{2,\mathbf{x}}^- \neq \emptyset, \\ 0, & \text{otherwise,} \end{cases} \\ \beta_{1,\mathbf{x}}^+ &= \begin{cases} \min B_{1,\mathbf{x}}^+, & \text{if } B_{1,\mathbf{x}}^+ \neq \emptyset, \\ \bar{B}, & \text{otherwise,} \end{cases} \\ \beta_{2,\mathbf{x}}^+ &= \begin{cases} \max B_{2,\mathbf{x}}^+, & \text{if } B_{2,\mathbf{x}}^+ \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Since $\sigma_{\mathbf{x}}^-$ and $\sigma_{\mathbf{x}}^+$ are non-decreasing (due to the fact that $G_{\mathbf{x}}$ is convex and non-increasing), we have that $\beta_{1,\mathbf{x}}^- \leq \beta_{2,\mathbf{x}}^- \leq \beta_{1,\mathbf{x}}^+ \leq \beta_{2,\mathbf{x}}^+$.

Then, for any $\beta_{\mathbf{x}}^- \in [\beta_{1,\mathbf{x}}^-, \beta_{2,\mathbf{x}}^-]$ and $\beta_{\mathbf{x}}^+ \in [\beta_{1,\mathbf{x}}^+, \beta_{2,\mathbf{x}}^+]$, the policy (9) is optimal. \blacksquare

Remark 1: Note that the proof of Theorem 1 describes a continuum of optimal policies, since any choice for $\beta_{\mathbf{x}}^- \in [\beta_{1,\mathbf{x}}^-, \beta_{2,\mathbf{x}}^-]$ and $\beta_{\mathbf{x}}^+ \in [\beta_{1,\mathbf{x}}^+, \beta_{2,\mathbf{x}}^+]$ defines a solution to (7).

C. Proof of Corollary 1

Proof: In case $\eta_c = \eta_d = 1$, we see that $[\beta_{1,x}^-, \beta_{2,x}^-] = [\beta_{1,x}^+, \beta_{2,x}^+]$, and the result readily follows. ■

D. Proof of Proposition 1

Proof: We have that $\beta_x = 0$ if for all $\beta \in U_x(0)$,

$$\begin{aligned} H_x(0, 0) &\leq H_x(\beta, 0) \\ \Leftrightarrow d(\mathbf{x})p(\mathbf{x}) - (d(\mathbf{x}) + \beta)p(\mathbf{x}) \\ &+ \alpha \int_{\mathbf{y} \in \Omega} f_x(\mathbf{y})(J_{\mathbf{y}}(0) - J_{\mathbf{y}}(\beta)) \leq 0 \\ \Leftrightarrow \alpha \int_{\mathbf{y} \in \Omega} f_x(\mathbf{y})(J_{\mathbf{y}}(0) - J_{\mathbf{y}}(\beta)) &\leq \beta p(\mathbf{x}). \end{aligned}$$

This holds if $J_{\mathbf{y}}(b_1) - J_{\mathbf{y}}(b_2) \leq (b_2 - b_1)\frac{p(\mathbf{x})}{\alpha}$ for all $b_1 < b_2$, $\mathbf{y} \in \Omega$.

In order to verify that $\beta_x = \bar{B}$, we need to show that for all $\beta \in U_x(\bar{B})$,

$$\begin{aligned} H_x(\bar{B}, \bar{B}) &\leq H_x(\beta, \bar{B}) \\ \Leftrightarrow d(\mathbf{x})p(\mathbf{x}) - (d(\mathbf{x}) + (\bar{B} - \beta))p(\mathbf{x}) \\ &+ \alpha \int_{\mathbf{y} \in \Omega} f_x(\mathbf{y})(J_{\mathbf{y}}(\bar{B}) - J_{\mathbf{y}}(\beta)) \leq 0 \\ \Leftrightarrow \alpha \int_{\mathbf{y} \in \Omega} f_x(\mathbf{y})(J_{\mathbf{y}}(\beta) - J_{\mathbf{y}}(\bar{B})) &\geq (\bar{B} - \beta)p(\mathbf{x}). \end{aligned}$$

This holds if $J_{\mathbf{y}}(b_1) - J_{\mathbf{y}}(b_2) \geq (b_2 - b_1)\frac{p(\mathbf{x})}{\alpha}$ for all $b_1 < b_2$, $\mathbf{y} \in \Omega$. ■

E. Proof of Proposition 2

Proof: We will show that for all $b_1, b_2 \in \mathcal{B}$, $b_1 < b_2$ and $\mathbf{x} \in \Omega$,

$$J_{\mathbf{x}}(b_1) - J_{\mathbf{x}}(b_2) \leq (b_2 - b_1)\bar{P}. \quad (12)$$

It then readily follows that in both cases (i) and (ii), the first condition of Proposition 1 is satisfied.

In order to show that (12) holds, we apply the finite horizon framework presented in the proof of Lemma 1, and use induction on the horizon n . First, it is readily seen that

$$J_{\mathbf{x},0}(b_1) - J_{\mathbf{x},0}(b_2) = 0 \leq (b_2 - b_1)\bar{P}.$$

Now assume that

$$J_{\mathbf{x},n-1}(b_1) - J_{\mathbf{x},n-1}(b_2) \leq (b_2 - b_1)\bar{P}.$$

Then

$$\begin{aligned} &J_{\mathbf{x},n}(b_1) - J_{\mathbf{x},n}(b_2) \\ &= ((\beta_{\mathbf{x}}^*(b_1) - b_1) - (\beta_{\mathbf{x}}^*(b_2) - b_2))p(\mathbf{x}) \\ &+ \alpha \int_{\mathbf{y} \in \Omega} f_x(\mathbf{y})(J_{\mathbf{y},n-1}(\beta_{\mathbf{x}}^*(b_1)) - J_{\mathbf{y},n-1}(\beta_{\mathbf{x}}^*(b_2)))d\mathbf{y} \\ &\leq ((\beta_{\mathbf{x}}^*(b_1) - b_1) - (\beta_{\mathbf{x}}^*(b_2) - b_2))\bar{P} \\ &+ \alpha(\beta_{\mathbf{x}}^*(b_2) - \beta_{\mathbf{x}}^*(b_1))\bar{P} = (b_2 - b_1)\bar{P}, \end{aligned}$$

completing the proof. ■

F. Proof of Proposition 3

Proof: For i.i.d. prices we see that the future prices do not depend on the current state, i.e., $G_{\mathbf{x}} \equiv G$. Consequently, we have that the subdifferential of G is also independent of \mathbf{x} , and we write $\partial G(\beta) = [\sigma^-(\beta), \sigma^+(\beta)]$. Now, the sets $B_{1,\mathbf{x}}^-, B_{1,\mathbf{x}}^+, B_{2,\mathbf{x}}^-$ and $B_{2,\mathbf{x}}^+$ in (11) only depend on \mathbf{x} through $p(\mathbf{x})$, and are non-increasing in $p(\mathbf{x})$ due to the non-decreasingness of σ^- and σ^+ . Thus, the thresholds are non-increasing in $p(\mathbf{x})$ as well. ■

G. Details of Example 1

We denote by \mathbf{x}_i the state such that $p(\mathbf{x}_i) = i$, $i = 1, \dots, 4$. First, observe that in this example the control set does not depend on the battery level, and $\beta_{\mathbf{x}}^*(b) \equiv \beta_{\mathbf{x}}^*$. Consequently, for all $\mathbf{y} \in \Omega$, $0 \leq b_1 \leq b_2 \leq \bar{B}$,

$$J_{\mathbf{y}}(b_1) - J_{\mathbf{y}}(b_2) = (b_2 - b_1)p(\mathbf{y}). \quad (13)$$

By Proposition 2 we know that $\beta_4 = 0$. In order to show that $\beta_1 = 1$, $\beta_2 = 0$ and $\beta_3 = 1$, we need to demonstrate that (cf. Proposition 1):

$$\begin{aligned} &\frac{1}{2}(J_{\mathbf{x}_1}(\beta) - J_{\mathbf{x}_1}(1)) + \frac{1}{2}(J_{\mathbf{x}_3}(\beta) - J_{\mathbf{x}_3}(1)) \\ &\geq p_1(1 - \beta)p_1/\alpha, \end{aligned} \quad (14)$$

$$J_{\mathbf{x}_1}(0) - J_{\mathbf{x}_1}(\beta) \leq p_2\beta/\alpha, \quad (15)$$

$$J_{\mathbf{x}_4}(\beta) - J_{\mathbf{x}_4}(1) \geq p_3(1 - \beta)/\alpha, \quad (16)$$

respectively.

By (13),

$$\begin{aligned} &\frac{1}{2}(J_{\mathbf{x}_1}(\beta) - J_{\mathbf{x}_1}(1)) + \frac{1}{2}(J_{\mathbf{x}_3}(\beta) - J_{\mathbf{x}_3}(1)) \\ &= 2(1 - \beta) \geq (1 - \beta)/\alpha, \end{aligned}$$

which corresponds to (14). For β_2 we have

$$J_{\mathbf{x}_1}(0) - J_{\mathbf{x}_1}(\beta) = \beta \leq \beta p_2/\alpha,$$

which satisfies (15). Finally, we see that also (16) is satisfied since

$$J_{\mathbf{x}_4}(\beta) - J_{\mathbf{x}_4}(1) = (1 - \beta)p_4 \geq (1 - \beta)p_3/\alpha.$$

REFERENCES

- [1] P. Van de Ven, N. Hegde, L. Massoulié, and T. Salonidis, "Optimal control of residential energy storage under price fluctuations," in *Proc. of IARIA ENERGY*, Venice, Italy, May 22–27 2011.
- [2] S. Borenstein, "The long-run efficiency of real-time electricity pricing," *The Energy Journal*, vol. 26, no. 3, pp. 93–116, 2005.
- [3] M. Albadi and E. El-Saadany, "A summary of demand response in electricity markets," *Electric Power Systems Research*, vol. 78, no. 11, pp. 1989–1996, 2008.
- [4] H. Allcott, "Rethinking real-time electricity pricing," *Resource and Energy Economics*, vol. 33, no. 4, pp. 820–842, 2011.
- [5] G. Barbose and C. Goldman, "A survey of utility experience with real time pricing," Berkeley National Laboratory, Tech. Rep. LBNL-54238, 2004.
- [6] M. Lijesen, "The real-time price elasticity of electricity," *Energy Economics*, vol. 29, no. 2, pp. 249–258, 2007.
- [7] S. Peterson, J. Whitacre, and J. Apt, "The economics of using plug-in hybrid electric battery packs for grid storage," *Journal of Power Sources*, vol. 195, no. 8, pp. 2377–2384, 2010.

- [8] R. Urgaonkar, B. Urgaonkar, M. Neely, and A. Sivasubramanian, "Optimal power cost management using stored energy in data centers," in *Proc. of ACM SIGMETRICS*, San Jose, CA, June 7–11 2011, pp. 221–232.
- [9] F. Graves, T. Jenkins, and D. Murphy, "Opportunities for electricity storage in deregulating markets," *The Electricity Journal*, vol. 12, no. 8, pp. 46–56, 1999.
- [10] K. Ahlert and C. Van Dinter, "Sensitivity analysis of the economic benefits from electricity storage at the end consumer level," in *Proc. of IEEE Bucharest Power Tech Conference*, Bucharest, Romania, June 28–July 2 2009.
- [11] W. Hu, Z. Chen, and B. Bak-Jensen, "Optimal operation strategy of battery energy storage system to real-time electricity price in Denmark," in *Proc. of IEEE Power and Energy Society General Meeting*, Minneapolis, MN, July 25–29 2010.
- [12] I. Koutsopoulos, V. Hatzi, and L. Tassiulas, "Optimal energy storage control policies for the smart power grid," in *Proc. of IEEE SmartGrid-Comm*, Brussels, Belgium, October 17–20 2011, pp. 475–480.
- [13] Y. Guo, Z. Ding, Y. Fang, and D. Wu, "Cutting down electricity cost in internet data centers by using energy storage," in *Proc. of IEEE GLOBECOM*, Houston, TX, December 5–9 2011.
- [14] P. Harsha and M. Dahleh, "Optimal sizing of energy storage for efficient integration of renewable energy," in *Proc. of CDC*, Orlando, FL, December 12–15 2011.
- [15] Y. Zhou, A. Scheller-Wolf, N. Secomandi, and S. Smith, "Managing wind-based electricity generation with storage and transmission capacity," Tepper Working Paper, Tech. Rep. 2011-E36, 2011.
- [16] P. Harsha and M. Dahleh, "Optimal management and sizing of energy storage under dynamic pricing for the efficient integration of renewable energy," 2011, submitted.
- [17] D. Bertsekas and S. Shreve, *Stochastic Optimal Control: The Discrete-Time Case*, 2nd ed. Athena Scientific, 1996.
- [18] Independent Electricity System Operator. Operator. (2012) Electricity prices. [Online]. Available: <http://www.ieso.ca/imoweb/marketdata/hoep.asp>
- [19] I. Richardson, M. Thomson, D. Infield, and C. Clifford, "Domestic electricity use: A high-resolution energy demand model," *Energy and Buildings*, vol. 42, no. 10, pp. 1878–1887, 2010.
- [20] Y. Li and P. Flynn, "Electricity deregulation, spot price patterns and demand-side management," *Energy*, vol. 31, no. 6-7, pp. 908–922, 2006.
- [21] R. Rockafellar, *Convex Analysis*, 2nd ed. Princeton University Press, 1972.