Optimal Control of Residential Energy Storage Under Price Fluctuations

Peter van de ven
Department of Mathematics & Computer Science
Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
p.m.v.d.ven@tue.nl

Nidhi Hegde, Laurent Massoulié and Theodoros Salonidis
Technicolor Paris Research Lab
Technicolor
75015 Paris, France
{nidhi.hegde,laurent.massoulie,theodoros.salonidis}@technicolor.com

***Abstract*—An increasing number of retail energy markets exhibit price fluctuations and provide home users the opportunity to buy energy at lower than average prices. However, such cost savings are hard to realize in practice because they require human users to observe the price fluctuations and shift their electricity demand to low price periods. We propose to temporarily store energy of low price periods in a home battery and use it later to satisfy user demand when energy prices are high. This enables home users to save on their electricity bill by exploiting price variability without changing their consumption habits. We formulate the problem of minimizing the cost of energy storage purchases subject to both user demands and prices as a Markov Decision Process and show that the optimal policy has a threshold structure. We also use a numerical example to show that this policy can lead to significant cost savings, and we offer various directions for future research.**

***Index Terms*—Battery storage, dynamic pricing, dynamic programming, energy storage, threshold policy.**

I. INTRODUCTION

Wholesale energy prices exhibit significant fluctuations during each day due to variations in demand and generator capacity. Home users are traditionally not exposed to these fluctuations but pay a fixed retail energy price, as shown in Figure 1(a). Economists have long argued to remove the fixed retail prices in favor of prices that change during the day. Such *dynamic pricing* reflects the prices of the wholesale market and has been predicted to lead to lower demand peaks and lower average level and volatility of the wholesale price [5].

Dynamic pricing has been enabled by recent smart-grid technologies such as smart meters. A first example of dynamic pricing that is being increasingly adopted is time-of-use pricing (Figure 1(b)). Such schemes typically provide two or three price levels (e.g., ‘off-peak’, ‘mid-peak’ and ‘on-peak’) where the level is determined by the time of day. The price levels are determined well in advance and are typically not changed more than once or twice per year. A second example of dynamic pricing is real-time pricing (Figure 1(c)) where the retail energy price changes hourly or half-hourly to reflect the price on the wholesale energy market.

Dynamic pricing creates an opportunity for users to reduce energy costs by exploiting the price fluctuations. However, in practice users show only a minor shift in their demand to match the energy prices [2]–[4], [8]. A possible remedy is to

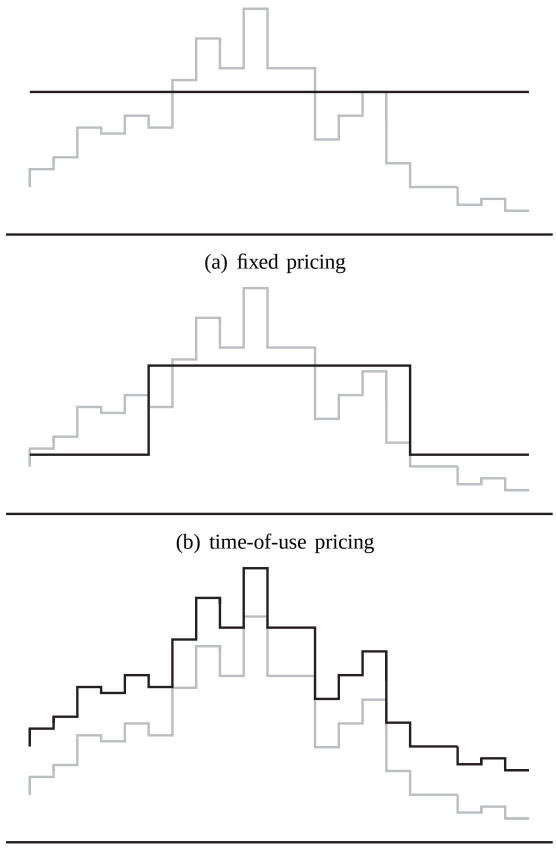

(a) fixed pricing

(b) time-of-use pricing

(c) real-time pricing

Figure 1. The wholesale energy price (gray) and various approaches to retail pricing (black).

equip homes with a battery that can be used for home energy storage. This battery can be charged when the energy price is low and the stored energy can then be used to protect against high prices. This allows users to benefit from the varying energy price without having to adjust their usage patterns accordingly. Energy can be stored both by a dedicated battery, or by using the battery pack of an electric car [9]. In the past such setup was not economically viable due to the high cost of batteries, but current developments have brought storage applications within reach.

 ISBN: 978-1-61208-136-6

In this paper, we address the problem of organizing home 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 difficult to handle due to the stochastic nature of price and demand and due to the fact that we aim to minimize the long-term costs. In our approach, we model the problem as a Markov Decision Process and we show that there exists a threshold-based stationary cost-minimizing policy. When the battery level is below this threshold, the battery is charged up to it, while the battery is discharged when above the threshold. By comparing the costs incurred under this policy with the cost of satisfying the demand directly from the grid, we can show that energy storage may lead to significant cost savings. In the current paper we provide an outline of this approach, which will be discussed in more detail in a follow-up paper.

To the best of our knowledge, previous work on home energy storage is limited and does not propose optimal control solutions subject to stochastic price and demand fluctuations. Home 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 [6]. This problem has been studied assuming no demands and 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 [1], [7]. However, these formulations do not take into account the stochasticity in prices and demands and do not allow for long-term cost optimization. Our approach can be readily adapted to an arbitrage problem in an infinite-horizon setting, where the behavior of the price process may be stochastic. A similar threshold-based optimal policy can be shown to hold in this case.

We also recently became aware of a parallel work that uses a model similar to ours to investigate control of energy storage in the context of *data centers* [10]. The model in [10] 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 contrast, our model incorporates battery inefficiencies and we show that the optimal scheduling policy is threshold-based.

The rest of the paper is structured as follows. In Section II we introduce the model and describe the various decision variables. In Section III we propose an optimal policy and apply this policy to a numerical example. Section IV gives some concluding remarks and directions for future research.

II. Model outline

Consider a residential 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 buffer level (i.e., the state of charge) of the battery at time t, $t = 0, 1, \dots$. Let $\bar{B}$ represent the maximum buffer level, so $B(t) \in [0, \bar{B}]$. 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 finite support $D(t) \in [0, \bar{D}]$, as does the price $P(t) \in [0, \bar{P}]$.

Denote by $\Omega = [0, \bar{D}] \times [0, \bar{P}]$ the set of possible realizations of demand and price, and for any $\boldsymbol{x} \in \Omega$, denote by $d(\boldsymbol{x})$ and $p(\boldsymbol{x})$ the corresponding price and demand, respectively. We assume that the demand and price level may be correlated, and evolve according to some stationary process. Specifically, we denote by $f_{\boldsymbol{x}}(\boldsymbol{y})$ the probability density function of moving from state $\boldsymbol{x}$ to state $\boldsymbol{y}$ in the next slot, for any $\boldsymbol{x}, \boldsymbol{y} \in \Omega$.

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. 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 from the battery used towards satisfying demand , see Figure 2.

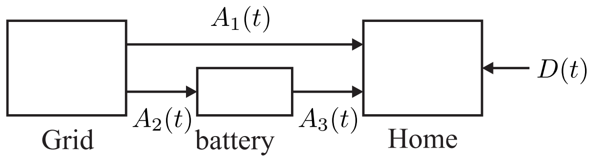

Figure 2. A graphical representation of the model.

We assume $A_1(t), A_2(t), A_3(t) \geq 0$, and since all demand must be met, we require that

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

The battery has a finite charging rate, so the amount of energy that can be used for charging the battery is bounded as $A_2(t) \leq \bar{A}$, for some $\bar{A} \in [0, \bar{B}]$.

The buffer level of the battery evolves according to

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

and the energy costs in slot t is given by

$$g(t) = (A_1(t) + A_2(t))P(t).$$

Our goal is to choose in each slot $A_1(t)$, $A_2(t)$ and $A_3(t)$ as to minimize the total discounted cost

$$g = \sum_{t=0}^{\infty} g(t)\alpha^t, \tag{1}$$

with $0 < \alpha < 1$ the discount factor. Note that the total discounted cost is finite, since the per-slot costs are bounded.

Before we consider in more detail the infinite-horizon problem (1), we first note that it is never optimal to charge and discharge the battery in the same slot, i.e., we have $A_2(t)A_3(t) = 0$, $t = 0, 1, \dots$. This is intuitively clear, because charging and discharging the battery in the same slot 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.

This observation simplifies the minimization problem significantly, by reducing the number of decision variables. Specifically, denote by $\Delta(t)$ the buffer level difference between slot t and $t+1$, i.e.,

$$B(t+1) = B(t) + \Delta(t).$$

Then, in view of the restriction on simultaneous charging and discharging:

$$\begin{aligned} A_1(t) &= D(t) + \Delta(t)\eta_d \mathbf{1}_{\{\Delta(t)<0\}}, \\ A_2(t) &= \Delta(t)\eta_c^{-1}\mathbf{1}_{\{\Delta(t)>0\}}, \\ A_3(t) &= -\Delta(t)\eta_d \mathbf{1}_{\{\Delta(t)<0\}}. \end{aligned}$$

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

$$g(t) = (D(t) + (\Delta(t))^+\eta_c^{-1} + (\Delta(t))^-\eta_d)P(t),$$

with $x^+ = \max\{x, 0\}$ and $x^- = -\max\{-x, 0\}$.

III. The optimal policy

In this section, we discuss how to choose in each slot the $\Delta(t)$ that minimizes the total discounted costs. To this end, we rewrite our model as a Markov decision process. We denote by $J_{\boldsymbol{x}}(b)$ the minimal total discounted costs, starting from state $\boldsymbol{x} \in \Omega$, and buffer level $b \in [0, \bar{B}]$. The cost function satisfies the Bellman equation

$$J_{\boldsymbol{x}}(b) = \inf_{\delta \in U_{\boldsymbol{x}}(b)} \{\gamma_{\boldsymbol{x}}(\delta) + \alpha G_{\boldsymbol{x}}(b+\delta)\},$$

with $\gamma_{\boldsymbol{x}}(\delta) = (d(\boldsymbol{x}) + \delta^+\eta_c^{-1} + \delta^-\eta_d)p(\boldsymbol{x})$ the immediate costs, $G_{\boldsymbol{x}}(b') = \int_{\boldsymbol{y}\in\Omega} f_{\boldsymbol{x}}(\boldsymbol{y})J_{\boldsymbol{y}}(b')\mathrm{d}\boldsymbol{y}$, and $U_{\boldsymbol{x}}$ the control set that contains all allowed decisions for the difference in buffer level. It is readily seen that

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

where $U_{\boldsymbol{x}}^-(b) = -\min\{b, d(\boldsymbol{x})\}$ and $U_{\boldsymbol{x}}^+(b) = \min\{\bar{B}-b, \bar{A}\}$.

In the remainder we restrict ourselves to the case $\eta_c = \eta_d = 1$, although the optimal policy for the more general case is similar to the policy described below for the completely efficient scenario.

The optimal policy specifies for each price $p = p(\boldsymbol{x})$ a buffer threshold $\beta(p) \in [0, \bar{B}]$. If $b \le \beta(p)$, then the optimal policy is to charge the buffer as close to $\beta(p)$ as the control set allows. If $b \ge \beta(p)$, the battery should be discharged up to $\beta(p)$ within the boundaries of the control set. Formally stated, the cost-minimizing choice for the buffer difference $\Delta_{\boldsymbol{x}}^*(b)$ is given by

$$\Delta_{\boldsymbol{x}}^*(b) = \begin{cases} \min\{\beta(p) - b, U_{\boldsymbol{x}}^+(b)\}, & b \le \beta(p), \\ \max\{\beta(p) - b, U_{\boldsymbol{x}}^-(b)\}, & b \ge \beta(p). \end{cases} \qquad (2)$$

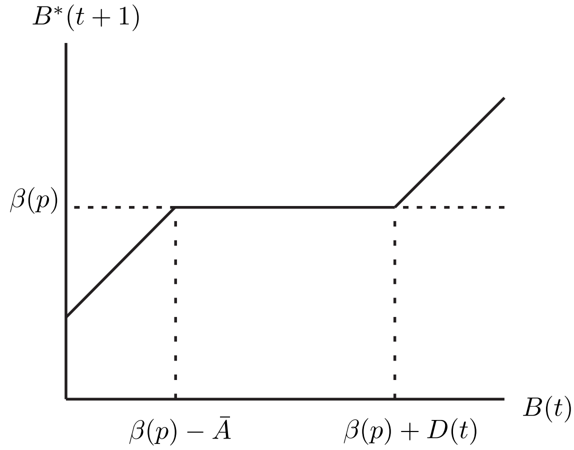

Figure 3. The structure of the optimal policy as a function of $B(t)$.

This policy is illustrated in Figure 3, which shows $B^*(t+1) = B(t) + \Delta_{\boldsymbol{x}}^*(b)$ plotted against $B(t)$, for some price level p.

According to (2), when $B(t) \ge \beta(p)$, the battery will be charged, and all demand will be met from the grid, so $A_1(t) = D(t)$. Conversely, when $B(t) > \beta(p)$ the demand is (partially) met from the battery, and we have $A_1(t) = (D(t) - (B(t) - \beta(p)))^+$ and $A_3(t) = \min\{D(t), B(t) - \beta(p)\}$.

The full analysis of the optimal policy (2) will appear in a subsequent paper.

A. A numerical example

We now present an example of energy storage under time-of-use pricing. We numerically determine the thresholds $\beta(p)$ through policy iteration, and use these to study the cost savings obtained from energy storage. The example below makes various simplifying assumptions on the demand and price processes. However, it clearly demonstrates the functionality of the optimal storage policy and the gains obtained from using energy storage.

Each day is divided into three periods of equal length, with corresponding prices (in Euro/kWh) $p_1 = 0.04$, $p_2 = 0.06$ and $p_3 = 0.07$. The demands (in kWh) are i.i.d. (independent and identically distributed), with the demand in period i, D_i having the distribution $D_i = (d_i + X)^+$, $i = 1, 2, 3$, with $d_1 = 5$, $d_2 = 6$, $d_3 = 7$ and $X \sim \mathcal{N}(0, 1)$. We assume that $\bar{A} = \bar{B}$, so there are no restrictions on charging the battery, and we set $\alpha = 0.99$. We discretize the state space into sections of 0.5 kWh, and use policy iteration to compute the thresholds β_1, β_2 and β_3. Figure 4 shows these thresholds plotted against $\bar{B}$. Note that $\beta_3 = 0$, as is to be expected for the threshold corresponding to the highest price.

We then simulate the demand process for 10^5 slots, and compute the total discounted costs over this period, both with and without battery storage. Figure 5 shows relative cost savings up to 40% obtained from energy storage, plotted against $\bar{B}$. We see that the cost savings initially increase with $\bar{B}$, but converge when the thresholds stabilize.

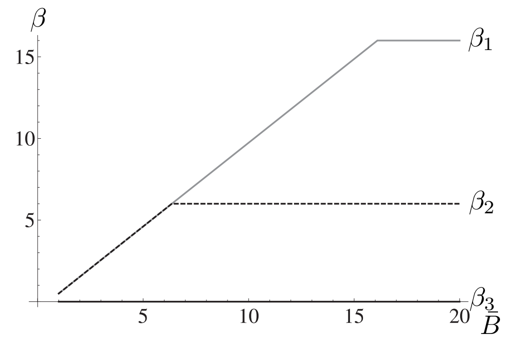

Figure 4. The optimal thresholds as a function of $\bar{B}$.

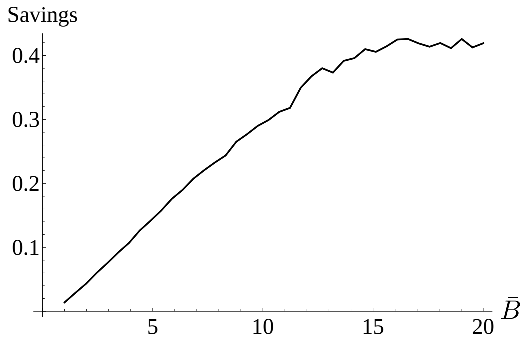

Figure 5. The cost savings as a function of $\bar{B}$.

IV. CONCLUSIONS AND OUTLOOK

We studied the control of residential energy storage under price fluctuations. We introduced a model for the battery operation and argued that the cost-minimizing storage policy is threshold-based. We showed by means of a small numerical example that residential energy storage can lead to significant savings in electricity cost.

We believe that this work opens up several avenues for future research. First, it is necessary to analytically show the existence of optimal threshold policies. Second, the computation of the optimal thresholds for more complex scenarios is challenging. Analytic expressions may be hard to derive and the large state space may make policy iteration computationally infeasible. Thus, it might be necessary to use approximations and bounds for the optimal threshold levels or use simple heuristics that provide reasonable performance.

It would also be interesting to incorporate in the model the battery lifetime and the costs of buying and replacing the battery. Battery lifetime benefits from longer sustained periods of charging and discharging. On one hand, this may complicate analysis because the optimal policy may depend on whether the battery was charged or discharged in previous slots. On the other hand, it may give rise to simple heuristics where the battery is alternatively fully charged and discharged. Taking into account the costs for buying and replacing the battery introduces the problem of battery dimensioning. Smaller batteries are cheaper but may provide less opportunity to exploit price fluctuations.

Finally, we may ask ourselves what will happen to the energy market when a significant fraction of users adopt energy storage. A possible outcome is that the resulting steady demand process will cause convergence of the energy market resulting in smaller price variations. While this is good from the perspective of both energy producers and users without energy storage, a less volatile price process will decrease the possibilities for exploiting price fluctuations for users with storage capacities. Consequently, the cost savings obtained from energy storage may decrease beyond the break-even point. The interplay between energy storage and the energy market is an interesting topic for future research.

REFERENCES

[1] K. Ahlert and C. Van Dinther. 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.

[2] M.H. Albadi and E.F. El-Saadany. A summary of demand response in electricity markets. *Electric Power Systems Research*, 78(11):1989–1996, 2008.

[3] H. Allcott. Rethingking real time electricity pricing. Technical Report http://web.mit.edu/allcott/www/papers.html, MIT, Cambridge, MA, 2010.

[4] G. Barbose and C. Goldman. A survey of utility experience with real time pricing. Technical Report LBNL-54238, Berkeley National Laboratory, 2004.

[5] S. Borenstein. The long-run efficiency of real-time electricity pricing. *The Energy Journal*, 26(3):93–116, 2005.

[6] F. Graves, T. Jenkins, and D. Murphy. Opportunities for electricity storage in deregulating markets. *The Electricity Journal*, 12(8):46–56, 1999.

[7] 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.

[8] M.G. Lijesen. The real-time price elasticity of electricity. *Energy Economics*, 29(2):249–258, 2007.

[9] S.B. Peterson, J.F. Whitacre, and J. Apt. The economics of using plug-in hybrid electric battery packes for grid storage. *Journal of Power Sources*, 195(8):2377–2384, 2010.

[10] R. Urgaonkar, B. Urgaonkar, M.J. Neely, and A. Sivasubramanian. Optimal power cost management using stored energy in data centers. Technical Report https://www.cse.psu.edu/research/publications/tech-reports/2010/CSE-10-008.pdf, Penn State, 2011.