Optimizing the Level of Commitment in Demand Response

Joshua Comden, Zhenhua Liu, and Yue Zhao
Stony Brook University
{joshua.comden, zhenhua.liu, yue.zhao.2}@stonybrook.edu

ABSTRACT

Demand response (DR) is a cost-effective and environmentally friendly approach for mitigating the uncertainties in renewable energy integration by taking advantage of the flexibility of the customers' demand. Existing DR programs, however, suffer from the inflexibility of commitment levels. In particular, these programs can be split into two classes depending on whether customers are fully committed or fully voluntary to provide demand response. Full commitment makes customers reluctant to participate, while the load serving entity (LSE) cannot rely on voluntary participation for reliability and dispatchability considerations. This paper proposes a generalized DR framework called Flexible Commitment Demand Response (FCDR) to allow for explicit choices of the level of commitment. We perform numerical simulations to demonstrate that the optimal level of commitment in FCDR brings in significant (around 50%) social cost reductions, consistently under various settings. This benefits both the LSE and customers simultaneously. Further, lower cost and higher levels of commitment can be simultaneously achieved with the optimal level of DR commitment.

1. INTRODUCTION

One of the major issues with the integration of renewable energy sources into the power grid is the increased variability that they bring [3]. Additionally, the limited capability to accurately predict this variability makes it challenging for the load serving entities (LSEs) to respond to it [6]. If this variability is not sufficiently addressed, it will limit the further penetration of renewables into the grid and even result in blackouts [5].

Various approaches have been implemented or proposed to address this issue. These include improving renewable generation forecast [18], aggregating diverse renewable sources [23], fast-responding reserve generators, energy storage [9, 11], and demand response (DR) [21], among others. In particular, in 2013, the California state legislature enforced a solution by passing a bill that requires 1,325 MW of grid energy storage by 2020 [17, 20]. In order for this solution to be cost-effective, the price of storage needs to be within the range of \$700-750/kWh. However, in 2013 when the law was passed, prices were about three times that amount [19].

Compared to energy storage, demand response has advantages to provide reserves to the LSEs in a cost-effective and environmentally friendly way [13, 21]. Despite the great potential, the increase in the amount of DR is much slower than that of renewable integration [22], partially evidenced by the California's move towards more grid-level energy storage. There are multiple reasons about this, but the level of DR commitment is an important factor.

Roughly speaking, there are two types of DR programs based on how much commitment customers need to make in the electric load reduction. In the first type, customers are required to make *full* commitment in load reduction, e.g., regulations service [15], capacity bidding [10]. Accordingly, a customer takes all the responsibilities of managing the risks/uncertainties in meeting its own commitment. Since such a full commitment is often costly, a higher payment is necessary to incentivize a customer to participate in DR programs than what are normally used in practice. In the latter type, customers do not need to make any commitments, and therefore are willing to participate in DR programs. Examples include emergency demand response programs [16] and coincident peak pricing [1]. The drawback, however, is that from LSE's perspective this sort of "voluntary" demand response is not reliable or sufficiently dispatchable.

This actually raises the following fundamental question: **what should be the optimal level of DR commitment?** This paper serves as a first step towards answering this question by making the following two main contributions:

(i) The design of the Flexible Commitment Demand Response framework: We propose a generalized framework for demand response programs called *Flexible Commitment Demand Response* (FCDR). Under FCDR, we explicitly allow an LSE to specify her requirement on the level of commitment from the customers by setting the value of a parameter ρ. The two types of DR programs mentioned above are special cases of the proposed FCDR with $\rho = 1$ and $\rho = 0$, respectively. We include the details of FCDR in Section 2.2.

(ii) Performance evaluations of FCDR: We conduct numerical experiments to illustrate the benefits of allowing the possibility of different levels of DR commitment. Our study yields the following key insights:

- The optimal level of commitment brings in a significant social cost reduction. For instance in our case studies, we achieve a consistent decrease of around 50%. Moreover, this benefits LSE and customers at the same time, so that both sides are incentivized to stay in the program.

- While lower price and higher levels of DR commitment are usually considered to be conflicting with each other, our results highlight the (somewhat surprising) opportunity to achieve both with the optimal level of commitment.

- We provide some insights about who should decide the level of commitment ρ. The result is particularly exciting as the optimal decision of LSE is near the social optimum. This means that we can rely on the LSE's (selfish/rational) decision to reach the social optimum.

- The optimal level of commitment increases with larger DR amount, and varies based on the customers' demand characteristics.

The most related work in literature is probably [14]. [14] studies the program where the LSE generates random numbers to decide whether a particular customer is required to provide the peak reduction.

2. FLEXIBLE COMMITMENT DEMAND RESPONSE FRAMEWORK

We start by modeling the LSE and customers, and then propose the flexible commitment demand response program.

2.1 Modeling the system

We consider a system with one LSE and a set of customers denoted by $\mathcal{N}$.

Customers

For customer $i \in \mathcal{N}$, her electricity demand is within a range $[l_i, h_i]$. We employ a random variable Y_i to represent her demand at a future time since she may not know it exactly beforehand due to uncertainties. The probability density function of Y_i is denoted by $f_i(y)$. Although future demand may depend on other factors (e.g. time of day, weather), we leave it as general as possible to allow the framework to be tailored for various statistical models.

In the current operation of the power grid, it is possible that future demand may be correlated with the LSE's need for DR but our own data analysis shows this correlation to be weak at best (cf. Figure 2(c)). Additionally as renewable energy penetration increases, it will add more uncertainty into the future supply. This will result in the correlation being further weakened, e.g. make DR necessary at non-peak hours [12]. For simplicity, we model the future demand as an independent random variable instead of conditioning on the LSE's DR demand. However, our model can be adapted to incorporate this dependency.

For customer i, we define its demand reduction as $x_i = y_i - y_i'$, where y_i is what her demand would realize without DR, and y_i' is her actual demand after performing DR. The internal cost for customer i to reduce its demand by $x_i \geq 0$ is modeled by a function $C_i(x_i)$ with a nondecreasing nonnegative marginal cost. Specifically in our numerical simulations, we consider the one-sided quadratic customer cost function:

$$C_i(x_i) = \begin{cases} c_i x_i^2, & \text{if } x_i \geq 0 \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

where the constant $c_i > 0$ represents the customer's increasing rate of disutility from load reduction. We leave it up for future work to study the effects of different customer cost function structures.

The customer's total cost is the difference between the internal cost and payment px_i received from the LSE (as explained in the following subsection):

$$C_i(x_i) - px_i \tag{2}$$

The load serving entity

We assume that the LSE has a targeted expected aggregate demand reduction D in total from all the customers. In order to reach this target, the LSE incentivizes the customers with payment. Specifically, the LSE sets a price p and if customer i provides x_i amount of demand response, she gets paid px_i. Note here x_i is normally a non-decreasing function of p.

In many cases, the realized DR amount $\sum_i x_i$ is different from D. The LSE incurs a cost $G\left(D - \sum_i x_i\right)$ for managing the difference. In our numerical simulations we use a one-sided quadratic cost function:

$$G\left(D - \sum_i x_i\right) = \begin{cases} a(D - \sum_i x_i)^2, & \text{if } D \geq \sum_i x_i \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

where $a > 0$ is a constant that represents the increasing cost rate of coping with the mismatch. Our model is general enough to incorporate other cost functions.

The LSE's total cost is the summation of the mismatch cost and the payment given to all the customers for their individual reductions:

$$G\left(D - \sum_i x_i\right) + p \sum_i x_i. \tag{4}$$

In order to model the total cost to the society, we simply sum the total cost of the individual customers and the LSE's total cost and thus the payments cancel:

$$G\left(D - \sum_i x_i\right) + \sum_i C_i(x_i). \tag{5}$$

2.2 The FCDR framework

The Flexible Commitment Demand Response framework consists of two key components: one parameter ρ (level of commitment), and one structural requirement for customers specified by a threshold parameter, y_i^l, for customer i. *The proposed FCDR requires customer i to reduce her demand to a level no higher than y_i^l with probability no less than $1 - \rho$.* In other words, the LSE allows the customers to violate the commitment of demand reduction to y_i^l with probability up to $1 - \rho$ when reductions may be too severe for the customer (see the peak load levels in Figure 1(b)). This explains the reason why ρ is called the level of commitment. Intuitively, as we increase ρ, the customer has less leeway to miss her commitment. $\rho = 1$ implies a full commitment and $\rho = 0$ implies complete voluntariness. We assume that, after agreeing to some ρ, the customers must meet this probabilistic commitment level. This can be enforced with a very high penalty otherwise.

Under the FCDR framework, the average amount of DR is calculated as the *effective demand reduction*. Specifically, it is the expected difference between unaffected desired demand (baseline) Y_i and demand after reduction Y_i', calculated as follows:

$$\mathbb{E}[x_i] = \mathbb{E}[Y_i - Y_i'] \tag{6}$$

Note that in practice the LSE can combine the demand responses from multiple (heterogenous) customers to mitigate the randomness of each individual's flexible commitment demand response. Our data analysis shows that the load demands of individual customers are only weakly correlated (cf. Figure 1(c)).

2.3 Decision making of LSE and customers

Under the FCDR framework, we consider that the LSE determines the price p and the level of commitment ρ, and each customer i decides its y_i^l [1]. This way, customer i is given the freedom to *not* participate in DR by choosing its $y_i^l = h_i$.

Customers

Each customer makes the load level decision Y_i' as well as y_i^l to minimize her expected total cost (2) while fulfilling the

[1]Whether the LSE or the customers should decide y_i^l is a future direction.

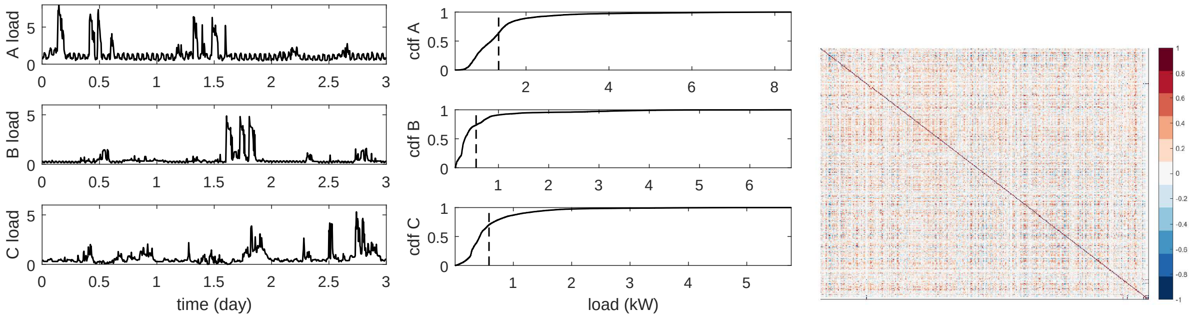

(a) Load (kW) (b) Cumulative distribution functions (c) 395 Homes

Figure 1: For Homes A, B, and C: (a) Load (kW) trace in five-minute intervals from May 6-8, 2012, (b) Cumulative distribution of the five-minute loads (kW) of the 33 days along with the means (dashed lines). (c) Heatmap showing the correlation coefficient matrix for one day of loads from 395 buildings.

level of commitment ρ:

$$\min_{y_i^l, Y_i'} \quad \mathbb{E}_{Y_i}[C_i(Y_i - Y_i') - p(Y_i - Y_i')] \tag{7}$$

$$\text{s.t.} \qquad \Pr(Y_i' \leq y_i^l) \geq \rho. \tag{8}$$

Under full commitment ($\rho = 1$), if the customer's realized desired demand is $Y_i > y_i^l$, then the customer must reduce her demand by $Y_i - y_i^l$. But **what should the customer decide if she is given the chance to miss the committed demand response by a probability** $1 - \rho$**?**

Customer i's optimal choice of y_i^l: The customer minimizes its cost by avoiding events that would achieve the highest costs. Since the customer i's cost of reducing demand is a function with a nondecreasing nonnegative marginal cost in the amount demand reduction, it would be best to *not reduce* load demand when $y_i - y_i^l$ is large and only reduce demand when $y_i - y_i^l$ is low. Specifically, we proved that (the details are omitted due to space limit) there exists a *threshold* y_i^u such that when $Y_i > y_i^u$, then the customer would decide not to reduce her demand and otherwise reduce to y_i^l. The customer must also make sure that the chosen threshold y_i^u does not violate her commitment obligation; thus, the requirement of ρ (8) simply becomes a lower bound on y_i^u:

$$y_i^u \geq F_i^{-1}(\rho) \tag{9}$$

where $F_i^{-1}(\rho)$ is the inverse cumulative distribution function of Y_i. This redefines problem (7), (8) to be based on only choosing its lower and upper thresholds. Since the load is reduced only when $Y_i \in (y_i^l, y_i^u)$, the customer optimization problem becomes:

$$\min_{y_i^l, y_i^u} \quad \int_{y_i^l}^{y_i^u} \left(C_i(t - y_i^l) - p(t - y_i^l)\right) f_i(t)dt \tag{10}$$

such that (9) is satisfied.

It will be shown later that (10) is not convex, and the customer can use an exhaustive search on deciding y_u and y_l. Although this naive algorithm has a quadratic time complexity, one thousand points in each decision variable for a residential customer who has a 0-10kW load can give 0.01kW precision. This is good enough for control purpose.

The load serving entity

The LSE decides the price p and the commitment level ρ to minimize its expected total cost (4):

$$\min_{p,\rho} \quad \mathbb{E}_{\mathbf{Y}}\left[G(D - \textstyle\sum_i x_i(p,\rho,Y_i)) + p\sum_i x_i(p,\rho,Y_i)\right] \tag{11}$$

where $\mathbf{Y} = \{Y_i\}_{i \in \mathcal{N}}$, and $x_i(p, \rho, Y_i)$ are functions that depend on how each customer behaves toward p, ρ and their desired demand Y_i, as fully characterized above. Note that although the uncertainty of Y_i is not explicitly stated in the objective function, it is implicitly taken into account when taking the expectation of $G(\cdot)$.

The LSE can decide the price p and commitment level ρ by exhaustive search since price has limited granularity and the shape of the cost function (11) is quite smooth around the optimal commitment level (see Fig. 3(b)), which allows moderate precision to be sufficient.

3. DATA ANALYSIS

3.1 Customer Loads

We use the Smart* Data Set obtained from the University of Massachusetts Trace Repository [8]. The specific data we use is the load data from three different homes located in Western Massachusetts given in one-second intervals from 33 days between May 1, 2012 through June 11, 2012.We average them into five-minute intervals and use these for our trace-based simulations. A concurrent three day sample of the three homes is given in Figure 1(a), which shows peak loads are non-overlapping in many cases.

Figure 1(b) displays the empirical cumulative distribution function of the three homes over the 33 days along with their means. The large peak-to-mean ratios indicate the potential benefit that a reduced commitment level for DR would bring for a customer.

Customers' loads are only weakly correlated, if at all. The northwest corner of Figure 2(c) shows a heatmap of the correlation matrix between the three homes which shows virtually little correlation. In addition to the three homes, the University of Massachusetts Trace Repository [8] provides the loads of 443 buildings for one complete day in one-second intervals which we averaged into five-minute intervals. A heatmap of the correlation matrix between 395 of the buildings is shown in Figure 1(c) which gives evidence that most customers have loads which are only weakly correlated.

3.2 Load Serving Entity

We use the following data from the ISO New England for the Western Massachusetts load zone [7]. The specific data is given in one-hour intervals of the real-time load, day-ahead market load demand, and the real-time price for the same days as the previously described customer data (Homes A, B, and C). We calculate the real-time load mismatch of supply and demand as the difference between the real-time load and day-ahead market provisioned load. Figure 2(a) displays a

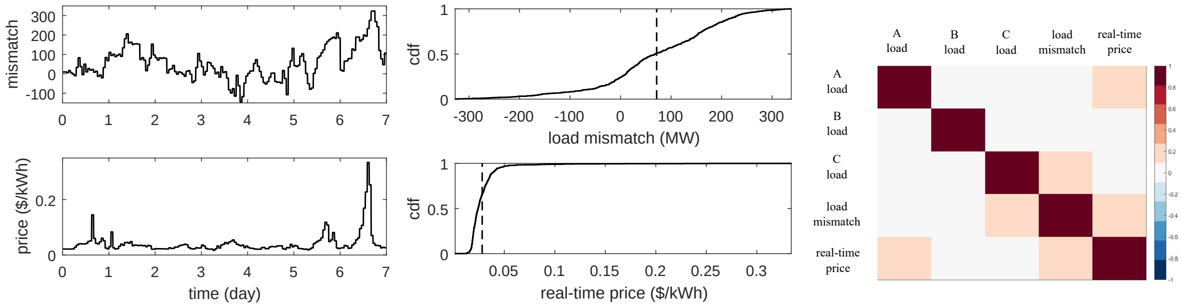

(a) Load mismatch (MW) and price (b) Cumulative distribution functions (c) Homes A, B, and C

Figure 2: (a) The ISO-NE load mismatch (MW) between real-time load and day-ahead provisioned load for the Western Massachusetts zone, and the real-time market price from May 23-29, 2012. For the same data during a period of 33 days: (b) Cumulative distribution functions for the load mismatch of ISO-NE, and real-time price, (c) Heatmap showing the correlation coefficient matrix for of loads from homes A, B, C, the load mismatch of ISO-NE, and real-time price.

one week sample of the load mismatch and its corresponding real-time price. The southeast corner of Figure 2(c) shows that they are weakly correlated over the sampled 33 days. The empirical cumulative distribution functions of the load mismatch and the real-time price for the 33 days are shown in Figure 2(b). In particular, the peak-to-mean ratio of the real-time price is 10 which indicates that DR can be extremely valuable for an LSE.

4. NUMERICAL RESULTS

4.1 Setup

We simulate a system with one LSE and 1,000 homogeneous customers. For each customer, her demand has the same probability density function. The empirical distribution functions used for three different simulations are from Homes A, B, and C in the previous section (cf. Figure 1(b)) [8]. The cost of reducing her demand is modeled as a one-sided quadratic function (1) with $c_i = \$1/(\text{kWh})^2$.

The goal of the LSE is to obtain a certain aggregated amount (e.g. 40kWh) of demand reduction from the 1000 customers (resulting in 0.04kWh) for a given five-minute timeslot. Recall the average demand of a customer is 1kW [2], which means the reduction is about 50% given the five-minute time frame. The LSE faces a penalty of not meeting the aggregated DR goal, modeled as another quadratic cost function in (3) with $a = \$0.002/(\text{kWh})^2$.

Since the functions are quite complicated and non-convex as we can see later in the results, both the LSE and customers use exhaustive search to determine their decision variables.

4.2 Economic benefit of FCDR

The LSE performs an exhaustive search over ρ and p to minimize its cost (11). For each pair of ρ and p, each customer responds with an amount of FCDR that minimizes its own cost. To observe the impact of different commitment levels ρ in reducing costs, we plot in Figures 3(a), 3(b) and 3(c) the costs to the society (i.e., the LSE plus the customers), the LSE, and an individual customer, as a function of ρ. These costs are normalized by the cost with a fully committed DR program, namely, $\rho = 1$. To understand the inner working of FCDR, we further plot the total expected DR and the optimal price chosen by the LSE as a function ρ in Figures 3(d) and 3(e). We make the following observations:

- With the optimal level of commitment ρ, the cost to the society is significantly reduced from that with a ρ close to 1. Indeed, from Figure 3(e), when ρ is close to 1, the cost of performing DR by the customers increases significantly, so that a much higher price is needed to incentivize them to supply the desired amount of DR. In the extreme case with $\rho \approx 1$, the price to incentivize the customers can become so high that the LSE would instead resort to other expensive alternatives, characterized by its cost function (3). This is evidenced in Figure 3(d), where with $\rho \approx 1$ the extracted DR falls to zero, meaning that the LSE has turned away from using DR.

- The optimal level of commitment ρ leads to a lower cost to the society than a fully voluntary DR (i.e., $\rho = 0$). This is because the requirement of ρ restricts the action space of the customers in a way that benefits the society. In particular, a higher level of commitment ρ decreases the uncertainty of the extracted DR.

- There is a discontinuity in Figures 3(a), 3(c), 3(d) and 3(e). This is caused by the non-convexity of the customer's optimization problem, as will be shown in more detail in the next subsection.

- There is a flat section towards lower ρ in all of these curves. This is due to the following reason: a fully voluntary DR ($\rho = 0$) can already lead to certain non-zero level of commitment ρ_{vol}, so that setting $\rho \leq \rho_{vol}$ does not change the behavior of the customers at all from that with a fully voluntary DR program.

Interestingly, we further observe that the optimal choice of ρ by the society is located relatively near that by the LSE. To investigate this in more detail, we plot the optimal ρ by the society and that by the LSE with varying levels of D in Figures 4(a), 4(b), and 4(c). We observe that they coincide with each other closely. As a result, the proposed FCDR framework in which LSE optimizes ρ and p in (11) achieves the social optimum. In addition, as expected, we observe that the optimal ρ increases with a higher amount of desired DR. Also note that at higher levels of D, all three homes approach an optimal commitment level around 0.95 which implies that differences in customer load distributions become less important as the amount of desired DR increases.

4.3 Optimal decisions by LSE and customers

We now provide more insights into the decision making processes of the LSE and the customers in the proposed FCDR framework.

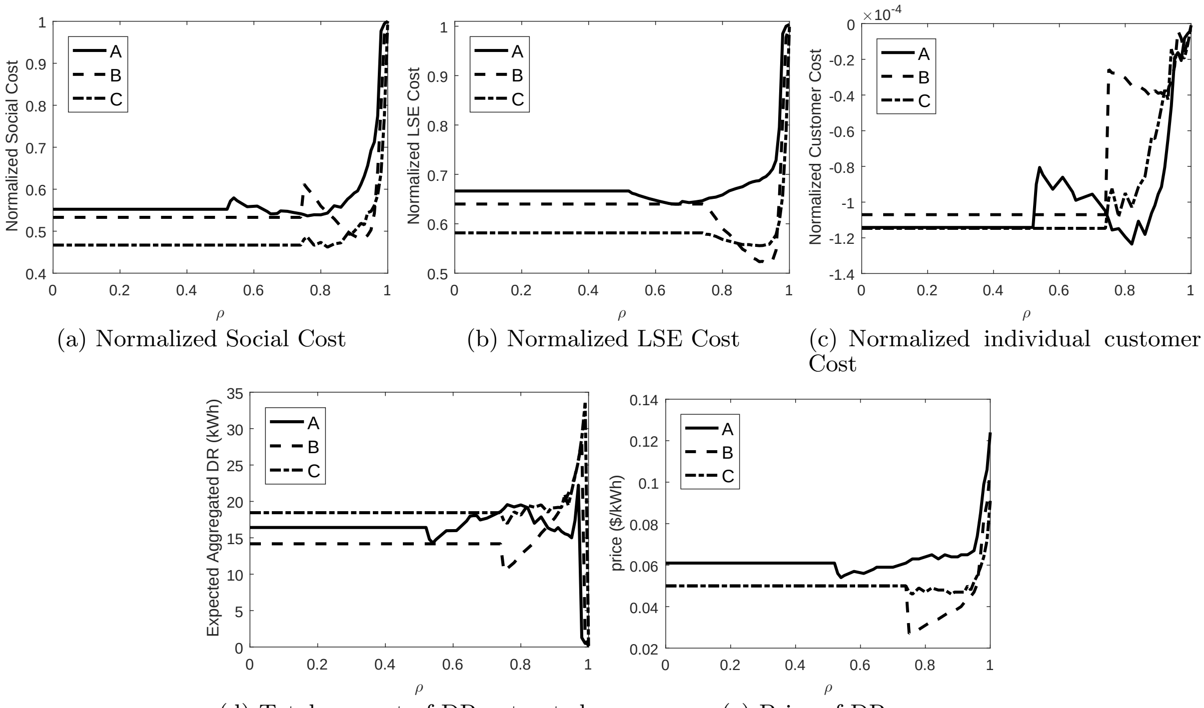

(a) Normalized Social Cost (b) Normalized LSE Cost (c) Normalized individual customer Cost

(d) Total amount of DR extracted (e) Price of DR

Figure 3: Results for empirical distributions from Homes A, B, and C: (a)-(c) illustrate the cost reduction from the perspective of the society, LSE, and a single customer under different commitment levels (ρ). Here, all cost numbers are normalized by the social cost when $\rho = 1$. (d) and (e) show the amounts of demand response extracted, and demand response prices under different ρ's.

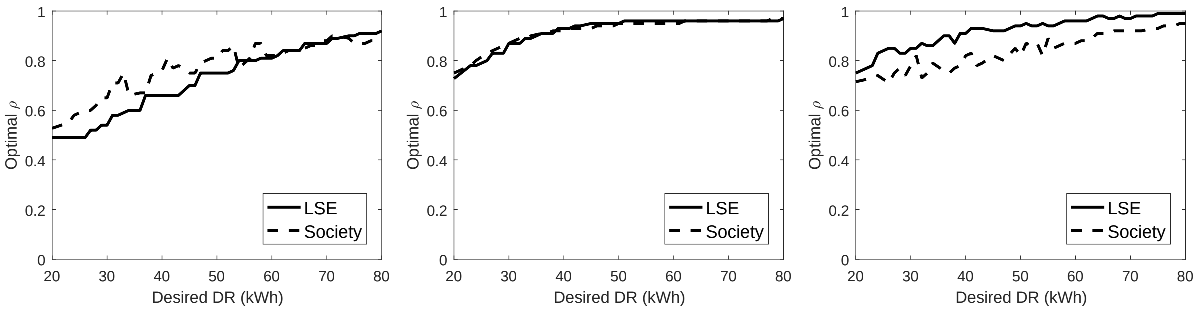

(a) Home A (b) Home B (c) Home C

Figure 4: Optimal ρ for a given desired DR which highlights the LSE's optimal choice of ρ approximates the social optimum.

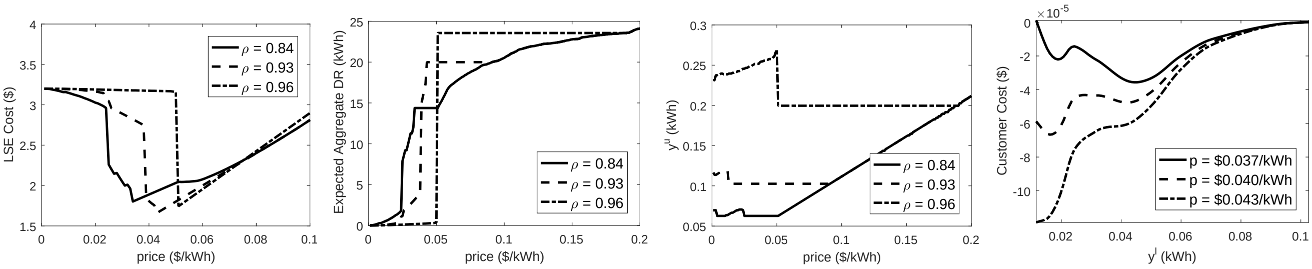

(a) LSE Cost (b) Extracted DR (c) y^u (d) Customer Cost

Figure 5: Home B when $D = 40$ kWh: (a) LSE cost versus price of DR, (b) Extracted DR amount versus price of DR, (c) y^u versus price of DR, (d) Individual customer cost versus y^l when $D = 40$ kWh, $\rho = 0.93$. Different lines correspond to different values of ρ or p.

4.3.0.1 LSE's decision on ρ and p.

In solving (11), given any level of commitment ρ, the LSE needs to decide the optimal DR price p. Figure 5(a) illustrates the complicated tradeoff under the empirical distribution of Home B (see our extended version online for Homes A and C). In general, when the price is too low, there is not enough DR extracted from the customers. This can be seen from Figure 5(b). The LSE therefore suffers from the high penalty of not being able to meet the DR target. On the other hand, if the price is set too high, the LSE pays too much for the extracted DR. As shown in the figures, the tradeoff is complicated and non-convex. The optimal price brings in significant cost reductions. The observations are consistent under all three of these empirical distributions.

4.3.0.2 A customer's decision on y^u and y^l.

Given the level of commitment ρ and price of DR p, the customers need to decide the amount of demand response by setting the values of y^u and y^l. This then leads to different amount of DR provided to the LSE. In particular, Figure 5(b) illustrates the amount of DR provided under the empirical distribution of Home B with different prices of DR. Clearly, the amount of extracted DR is increasing in the price offered. Under the three distributions, there is a series of stages: a convex increasing stage, jump stage(s), smooth increasing stage(s), a flat stage, and a concave increasing stage. In order to understand the reason behind this phenomenon, we provide the optimal values of y^l and y^u selected by the customer under different prices of DR in Figures 5(c), and 5(d). Specifically,

- In the first convex stage, y^u varies (cf. Figure 5(c)), and the probability of responding is $\Pr(Y_i \leq y^u)$ *is greater than* ρ but y^l is high resulting in little effective DR. As p increases, the customer is incentivized to provide more effective DR by adjusting both y^l and y^u since the empirical distribution is not monotonic.
- The jump stage(s) is caused by the *non-convex* cost function when a customer makes its decision on y^l. Specifically, we plot a customer's cost as a function of y^l in Figure 5(d) for different p around the price at which the jump happens in Figure 5(b). It can be seen that: a) finding the optimal y^l involves minimizing a non-convex cost function, and b) as p increases, the optimal y^l can switch from the middle part to either different local minimum or to the lower bound l_i in a discontinuous manner. This thus explains the jump behavior in the amount of DR in Figure 5(b).
- The smooth increasing stage(s) is caused by the migration of the local minimum from increasing the price as can be seen in Figure 5(d).
- In the semi-final flat stage, y^u remains unchanged in order to meet the requirement of response probability ρ, while y^l also stays at the lower bound l_i.
- In the final concave stage, y^l stays at the lower bound l_i, while y^u *increases* with p (cf. Figure 5(d)) because the customer is now willing to provide even higher response probability than ρ due to the sufficiently high price p.

5. CONCLUSION

This paper proposes a novel framework for demand response called Flexible Commitment Demand Response (FCDR), which enables explicit choices of the level of DR commitment. Numerical results highlight the great benefits of the optimal level of commitment. In particular, there is roughly a 50% decrease in social cost compared to programs requiring full commitment, and up to 40% of additional demand response extracted with a lower DR price compared to programs with no commitment.

6. ACKNOWLEDGMENTS

We would like to thank Srinivasan Keshav for his insightful comments and suggestions. This work was supported by NSF through CNS-1464388.

7. REFERENCES

[1] Coincident peak. Online. `http://www.fcgov.com/utilities/business/rates/electric/coincident-peak`.
[2] Frequently asked questions. Online. `https://www.eia.gov/tools/faqs/faq.cfm?id=97&t=3`.
[3] The smart grid: an introduction. Online. `http://energy.gov/oe/downloads/smart-grid-introduction-0`.
[4] Quickfacts table. Online, 2010-2014. `http://www.census.gov/quickfacts/table/PST045215/00`.
[5] Grid integration. Online, April 2016. `http://energy.gov/eere/wind/grid-integration`.
[6] Variability of renewable energy sources. Online, April 2016. `http://www.nrel.gov/electricity/transmission/variability.html`.
[7] Zonal information. Online, May 2016. `http://www.iso-ne.com/isoexpress/web/reports/pricing/-/tree/zone-info`.
[8] S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht. Smart*: An open data set and tools for enabling research in sustainable homes. *SustKDD, August*, 111:112, 2012.
[9] E. Bitar, R. Rajagopal, P. Khargonekar, and K. Poolla. The role of co-located storage for wind power producers in conventional electricity markets. In *American Control Conference (ACC)*, pages 3886–3891. IEEE, 2011.
[10] B. K. Cherry. Electric schedule E-CBP. Technical report, Pacific Gas and Electric Company, February 2013.
[11] M. Chowdhury, M. Rao, Y. Zhao, T. Javidi, and A. Goldsmith. Benefits of storage control for wind power producers in power markets. *to appear in IEEE Transactions on Sustainable Energy*, 2016.
[12] P. Denholm and M. Hand. Grid flexibility and storage required to achieve very high penetration of variable renewable electricity. *Energy Policy*, 39(3):1817–1830, 2011.
[13] N. Framework. Roadmap for smart grid interoperability standards, release 1.0 (jan. 2010)(nist special publication 1108).
[14] P. Harsha, M. Sharma, R. Natarajan, and S. Ghosh. A framework for the analysis of probabilistic demand response schemes. *Smart Grid, IEEE Transactions on*, 4(4):2274–2284, 2013.
[15] B. J. Kirby. *Frequency regulation basics and trends*. United States. Department of Energy, 2005.
[16] New York Independent System Operator, Rensselaer, NY. *Emergency Demand Response Program Manual*, 7.1 edition, October 2013.
[17] C. J. Peterman. R. 10-12-007 COM/CAP/jv1. *California Public Utilities Commission*, October 2013.
[18] P. Pinson. Wind energy: Forecasting challenges for its operational management. *Statistical Science*, 28(4):564–585, 2013.
[19] K. Silverstein. California's new energy storage mandate under the microscope. *EnergyBiz*, October 2013. `http://www.energybiz.com/article/13/10/california-s-new-energy-storage-mandate-under-microscope`.
[20] N. Skinner. Assembly bill no. 2514. *California State Assembly*, September 2010.
[21] J. S. Vardakas, N. Zorba, and C. V. Verikoukis. A survey on demand response programs in smart grids: Pricing methods and optimization algorithms. *IEEE Communcation Suverys Tutorials*, 17(1):152–178, 2015.
[22] A. Wierman, Z. Liu, I. Liu, and H. Mohsenian-Rad. Opportunities and challenges for data center demand response. In *Green Computing Conference (IGCC), 2014 International*, pages 1–10. IEEE, 2014.
[23] Y. Zhao, J. Qin, R. Rajagopal, A. Goldsmith, and H. V. Poor. Wind aggregation via risky power markets. *Power Systems, IEEE Transactions on*, 30(3):1571–1581, 2015.