Sizing and management of storage and demand response in the renewables-rich smart power grid

Millen Kanabar
Indian Institute of Technology, Bombay
Mumbai, India

Jayakrishnan Nair
Indian Institute of Technology, Bombay
Mumbai, India

ABSTRACT

The intermittency and unpredictability of solar and wind generation remains a key challenge as we attempt to transition to a predominantly renewables-powered electricity grid. Two key mechanisms will play a key role in addressing this challenge: dynamic operation of large-scale energy storage, and demand response. In this paper, we analyse the joint management of storage and demand response from the standpoint of a utility. Specifically, we consider a contract-based demand response (DR) model, whereby the utility is allowed curtail the electricity consumption of participating customers (industrial or retail) by at most a certain prescribed amount, subject to a further constraint on how often this curtailment can be applied. Under these constraints, we consider a storage management mechanism, which triggers DR when the charge level on the battery drops below a certain threshold. We derive large buffer asymptotics for this model, which provides tractable approximations of the loss of load probability and the frequency of demand curtailment, as a function of the battery size and the DR parameters.

CCS CONCEPTS

• **Hardware** → **Smart grid**; • **Mathematics of computing** → *Queueing theory*; • **Networks** → *Network performance analysis*; • **Theory of computation** → Random walks and Markov chains.

KEYWORDS

renewables, demand response, energy storage, Markov models

ACM Reference Format:
Millen Kanabar and Jayakrishnan Nair. 2021. Sizing and management of storage and demand response in the renewables-rich smart power grid. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Our power grid has been designed, over the past century, to operate with small supply side uncertainty (conventional generation being extremely reliable), and small demand side uncertainty (at a reasonable level of aggregation, demand can be predicted quite accurately ahead of time). This enables us to match supply and demand at all times via small control interventions in real time. However, as the world transitions to a power grid that is powered predominantly by renewables, the inherent intermittency of solar and wind generation, the predominant drivers of this transition, will result in a considerable surge in the supply side uncertainty. Effective and economical management of this supply side uncertainty is the key challenge, as we seek to increase the penetration of renewable generation worldwide.

The two predominant mechanisms that are available to address the supply side uncertainty induced by renewable generation are: (i) dynamic operation of large scale battery storage, and (ii) demand response. The former allows us store energy in times of excess renewable generation, which can be later be utilized at times of deficit. On the other hand, demand response enables electricity consumption to adapt to the state of the grid, so that part of the demand can be curtailed, or deferred, at times when the grid is supply constrained.

In this paper, we analyse the joint operation of battery storage and demand response from the standpoint of a utility. Specifically, we consider a utility that operates a grid-scale battery, which can be charged/discharged dynamically based on the state of renewable generation. Additionally, the utility has demand response (DR) contracts with some of its retail/industrial customers, that allow the utility to curtail their demand by at most a certain prescribed amount, subject to a constraint on how often such curtailment can be applied. (In return, the utility compensates the customers who participate in the DR program via a suitable financial incentive.) DR arrangements of this form have been implemented by several utilities, including at conEdison, where various programs for load curtailment are available for stay-at-home consumers [10]. Alternatively, the responsibility of demand curtailment can be outsourced to third party aggregators, as is also common practice in several jurisdictions.

Our goal is to propose and analyse a mathematical model for joint operation of storage and DR by the utility, to guide both operational aspects (i.e., the scheduling of battery charging/discharging, and demand curtailment) as well as planning aspects (the sizing of the battery, and the number of customers to enter into DR contracts with) of the utility's decision making. The idea behind the model is that the state of charge of the battery storage is modulated by the instantaneous surplus/deficit in generation, with demand curtailment kicking in once the state of charge drops below a threshold. To the best of our knowledge, a tractable mathematical model that captures these dynamics is missing in the literature.

We model supply and demand uncertainty via a Markovian model. Specifically, we model the instantaneous net generation (supply minus demand) as a function of a background Markov chain, the state of which includes all those factors that influence generation and demand. The battery level is then modulated by the state of the background Markov process. Further, once the battery level drops below a threshold, demand curtailment kicks in (in one shot, or in a phased manner), which decreases the rate at which the battery drains. In summary, we model the evolution of the state

Conference'17, July 2017, Washington, DC, USA
2021. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

of charge of the battery as a Markov modulated fluid queue (see [1, 14]), with state dependent drift to account for the demand response (see [3]). We further analyse the large buffer asymptotics for this system, which provides tractable approximations of performance metrics of interest, including the long run fraction of time demand curtailment is applied, the loss of load probability (the long run fraction of time the utility needs to procure additional power in real time to meet its demand), and the loss of load rate (the rate of real time energy procurement).

Related Literature

Our work is related to two strands of the literature, (i) on demand response, and (ii) on the sizing/operation of battery storage in the smart power grid.

We begin with a brief survey of the DR literature. Prior work on DR can be broadly classified into two groups: the analysis being focused on the customer's viewpoint, or the utility/ISO viewpoint. Some representative papers in the former category are [12, 17]; the typical model for DR in this line of work is that the utility employs dynamic real time pricing to modulate customer demand. The present paper is more closely related to the latter cluster of the DR literature, where the focus is on the utility/ISO side; see, for example, [4–6, 8, 9, 15].

Turning to storage sizing, the primary approach in the literature is to formulate an optimization problem, whose solution balances the cost of provisioning battery storage with associated reliability improvement; see, for example, [7, 11]. On the other hand, the literature on the scheduling of energy storage typically uses a Markov decision process (MDP) framework; see, for example, [18, 19]. In these papers, structural properties of the optimal policy of the MDP are established, which are exploited to compute the optimal policy more efficiently, or to develop heuristic algorithms.

There is also recent literature on unit commitment (at the ISO level) in the presence of renewables, storage, and/or demand response; see, for example, [13, 16]. These papers focus the formulation and numerical solution of the unit commitment problem.

In contrast to the above literature, the approach in the present paper is to develop a realistic, but also analytically tractable, model for battery storage and DR from the standpoint of a utility, that informs both operational and planning aspects.

2 MODEL AND PRELIMINARIES

Model for supply and demand uncertainty: Consider an electric utility company, which faces uncertain supply and demand processes. While demand uncertainty stems from unpredictable swings in electricity consumption by customers (modest, in practice), supply uncertainty stems from the renewable generation that is linked to the utility. The renewable resources might be operated directly by the utility, or owned by retail/industrial customers whose electricity consumption gets offset by their own renewable generation. We model the overall supply and demand uncertainty via the following Markovian model.

The utility is associated with an electricity generation process $g(t)$ that is a sum total of its procurements from long term contracts and the day ahead market, plus the renewable generation linked to the utility. Further, we denote the aggregate electricity demand seen by the utility by $d(t)$. We model the *net generation* $r(t) = g(t) - d(t)$ (i.e., the difference between the generation and the demand processes) as a function of a background Markov chain $\{X(t)\}_{t \geq 0}$. Formally, this chain is assumed to be an irreducible, time-reversible, Continuous-Time Markov Chain (CTMC) over a finite state space S. With every state $i \in S$, we associate a net generation $r_i = g_i - d_i$, where g_i and d_i denote the generation and demand associated with state i, respectively. Thus, the net generation at time t is given by $r(t) = r_{X(t)} = g_{X(t)} - d_{X(t)}$. For technical reasons, we make the assumption that $r_i \neq 0$ for all $i \in S$.

The background Markov chain $\{X(t)\}_{t \geq 0}$ thus captures, within its state, all those factors that influence supply and demand from the standpoint of the utility. This might include past and present weather conditions, seasonal aspects, the time of the day, and also past values of supply and demand. Note that this model can capture arbitrary dependencies between generation and demand.

Dynamics of battery storage in the absence of DR: We assume that the utility operates a grid-scale battery having capacity $b_{\max}$. We begin by describing the baseline model for the evolution of the state of charge $b(t)$ of this battery in the absence of demand response. $b(t)$ is modulated by the net generation process $r(t)$, subject to boundary conditions:

$$\frac{d}{dt}b(t) = \begin{cases} 0 & b(t) = 0 \textit{ and } r(t) < 0 \\ 0 & b(t) = b_{\max} \textit{ and } r(t) > 0 \\ r(t) & \textit{otherwise} \end{cases} \tag{1}$$

Note that the battery is charged when the net generation is positive (i.e., generation exceeds demand), and discharged when the net generation is negative (i.e., generation is less than demand). The boundary conditions enforce that battery cannot be drained if empty, or charged if full. Note that constraints on charge/discharge rates of the battery can also be incorporated by appropriately restricting the set of values $(r_i,\ i \in S)$ can take.

Mathematically, the battery evolution model (1) corresponds to a finite buffer Markov modulated fluid queue, a well studied object in the queueing/networking literature (see, for example, [1, 14]). Specifically, the literature characterized the stationary distribution of $\{(b(t), X(t))\}$, which is a Markov process that evolves over the state space $[0, b_{\max}] \times S$.

Demand Response model: We now describe our model for demand response (DR) on part of the utility. DR is captured by the ability to curtail demand to a certain extent when the state of charge of the battery drops below a threshold. This in turn decreases the discharge date of the battery, making loss of load less likely. Formally, we model the DR-adjusted demand process $\hat{d}_i$, when the background chain $X(t) = i$, as follows:

$$\hat{d}_i(b) = \begin{cases} d_i - \left\lceil n \frac{b_{\min} - b}{b_{\min}} \right\rceil \frac{\alpha_i}{n} & b \in [0, b_{\min}) \\ d_i & b \in [b_{\min}, b_{\max}] \end{cases} \tag{2}$$

Here, $b_{\min}$ is the DR-threshold, i.e., DR is activated only then the battery level drops below the level $b_{\min}$. The parameter α_i is the peak demand curtailment when the background process is in state i, and $n \in \mathbb{N}$ represents the number of phases the demand curtailment is rolled out over. If $n = 1$, the utility performs the maximum feasible demand curtailment as soon as the battery level drops below $b_{\min}$. On the other hand, when $n > 1$, the curtailment increases in n phases as the battery level drops progressively lower. (To provide a compact representation of demand curtailment, we have assumed that the demand curtailment increases uniformly as the battery occupancy drops below n uniformly spaced levels between 0 and $b_{\min}$. Our results generalize naturally to a non-uniform, piecewise constant variation of demand curtailment with battery level.)

Thus, the DR-adjusted net generation rate $\hat{r}_i(b)$, is given by

$$\hat{r}_i(b) = \begin{cases} g_i - d_i + \left\lceil n \frac{b_{\min} - b}{b_{\min}} \right\rceil \frac{\alpha_i}{n} & b \in (0, b_{\min}) \\ g_i - d_i & b \in [b_{\min}, b_{\max}). \end{cases} \quad (3)$$

Thus, in the presence of DR, our battery evolution is given by:

$$\frac{d}{dt} b(t) = \begin{cases} 0 & b(t) = 0 \textit{ and } r_{X(t)}(b(t)) < 0 \\ 0 & b(t) = b_{\max} \textit{ and } r_{X(t)}(b(t)) > 0 \\ r_{X(t)}(b(t)) & \textit{otherwise} \end{cases} \quad (4)$$

We note here that mathematically, the model (4) is equivalent to a fluid model for random early dropping (RED) in the networking literature (see [3]). Under our DR model, note that $(b(t), X(t))$ remains a Markov process over the state space $[0, b_{\max}] \times S$. In Section 3, we characterize some steady state properties of this Markov process in the asymptotic regime as $b_{\max} \to \infty$.

Finally, for simplicity, we make the following assumption on the peak demand curtailment:

$$r_i(r_i + \alpha_i) > 0 \quad \forall\, i \in S. \quad (5)$$

This assumption, which is referred to as the absence of a "confluence of drifts" in [3], ensures that the sign or r_i matches that of $r_i + \alpha_i$ for all i. Specifically, this condition implies that for those states i where $r_i < 0$ (i.e., demand exceeds generation), the net generation remains negative even under the maximum demand curtailment. If this condition does not hold, then the drift of the battery level would switch sign at $b_{\min}$ if $X(t) = i$, leading to atoms at that battery level in the stationary distribution. This makes our asymptotic analysis cumbersome. In practice, a confluence of drifts would result in rapidly oscillating ON/OFF switching of demand curtailment, which is undesirable from an engineering perspective.

Performance metrics and preliminary results: We define the Loss of Load Probability (LOLP) as the steady state probability that the battery is empty and the net generation rate is negative. Formally, $\mathsf{LOLP} := \sum_{i \in S} F_i(0) = \sum_{i \in S_-} F_i(0)$. The LOLP is the long run fraction of time that the utility is unable to meet its customer demand using the electricity procured ahead of time, and the energy stored in the battery. When this happens, the utility must buy (typically more expensive) power in the real time market to meet its deficit. Analogously, we define the Loss of Load Rate (LOLR) as defined as the long run time average rate of lost load, i.e, $\mathsf{LOLR} \overset{(a)}{:=} \lim_{t \to \infty} \frac{\int_0^t [\hat{r}(t)]_- \mathbb{1}_{\{b(t)=0\}}}{t} = \sum_{i \in S_-} |r_i| F_i(0)$. Here, $[z]_- := \max(-z, 0)$ and the limit in (a) is to be interpreted in an almost sure sense.

We conclude by stating a result from [2] that captures the asymptotic behavior of the LOLP and the LOLR in the absence of DR (i.e., $\alpha_i = 0$ for all i) as $b_{\max} \to \infty$. To state this result, we need to define a few quantities. Let Q denote rate matrix corresponding to the DTMC $\{X(t)\}$, and let $R := diag(r_i,\ i \in S)$. The drift Δ is defined as the long run average net generation in the absence of DR, i.e., $\Delta := \sum_{i \in S} \pi_i r_i$, where $(\pi_i,\ i \in S)$ is the stationary distribution corresponding to the background CTMC $\{X(t)\}$.

Theorem 2.1 (Theorem 2 in [2]). *Suppose that $\Delta > 0$, i.e., the average steady state generation exceeds the average steady state demand. Then*

$$\lim_{b_{\max} \to \infty} \frac{log\ \mathsf{LOLP}}{b_{\max}} = \lim_{b_{\max} \to \infty} \frac{log\ \mathsf{LOLR}}{b_{\max}} = -\lambda_c,$$

where λ_c is the smallest positive eigenvalue of $R^{-1}Q^T$.

Theorem 2.1 states that if $\Delta > 0$, then the LOLP and the LOLR decay exponentially with battery size $b_{\max}$ with decay rate λ_c, implying that these quantities can be made arbitrarily small by provisioning a large enough battery.[1] Indeed, Theorem 2.1 motivates the following approximations for the LOLP and *LOLP* when the battery size is large:

$$\mathsf{LOLP} \approx C e^{-\lambda_c b_{\max}}, \quad \mathsf{LOLR} \approx C' e^{-\lambda_c b_{\max}}, \quad (6)$$

where C, C' are positive constants. These approximations can in fact be used to perform *battery sizing* in practice, as is shown in [2].

3 LARGE BUFFER ASYMPTOTICS

In this section, we present the main technical results of this paper. Specifically, we characterize the asymptotic behavior of the long run fraction of time that DR is active, and the LOLP, as the battery size $b_{\max}$ scales to infinity. These quantities are shown to decay exponentially with the battery size, with decay rates that depend on the parameters of our DR model. These decay rates characterizations highlight the impact of DR, and also guide the configuration of the DR parameters, via crisp approximations of performance measures of interest.

We begin by analysing the asymptotic behavior of the LOLP as $b_{\max}$ scales to infinity, under the proposed DR model. For this, we suppose that

$$b_{\min} = \beta b_{\max},$$

where $\beta \in [0, 1]$, where β is held fixed as $b_{\max}$ is scaled. Note that $\beta = 0$ corresponds to a baseline scenario with no DR.

The asymptotic behavior of the LOLP is defined in terms of the following matrix.

$$R_{DR}^{-1} := diag\left(\frac{\beta}{n} \sum_{m=1}^{n} \frac{1}{g_i - d_i + \alpha_i \frac{m}{n}} + \frac{1 - \beta}{g_i - d_i} \right). \quad (7)$$

Note that if $\beta = 0$, R_{DR}^{-1} matches the matrix R^{-1}.

Theorem 3.1. *Suppose that $\Delta > 0$, i.e., the average steady state generation exceeds the average steady state demand. Under the DR model proposed in Section 2,*

$$\lim_{b_{\max} \to \infty} \frac{log(\mathsf{LOLP})}{b_{\max}} = \lim_{b_{\max} \to \infty} \frac{log(\mathsf{LOLR})}{b_{\max}} = -\lambda_{c,DR},$$

where $\lambda_{c,DR}$ is the smallest positive eigenvalue of $R_{DR}^{-1}Q^T$. Moreover, for $\beta > 0$, $\lambda_{c,DR} > \lambda_c$.

Theorem 3.1, which generalizes Theorem 2.1 to our DR model, states that both the LOLP and the LOLR decay exponentially with the battery size $b_{\max}$, with decay rate $\lambda_{c,DR}$. The decay rate $\lambda_{c,DR}$ can in turn be characterized in as the smallest positive eigenvalue of

[1] If $\Delta < 0$, then it is easy to show that the LOLP and LOLR are bounded away from zero for any battery size.

the matrix $R_{DR}^{-1}Q^T$ (that a positive eigenvalue exists for this matrix is proved in [14]).

Importantly, since $\lambda_{c,DR} > \lambda_c$ (when $\beta > 0$), we conclude that the presence of DR causes the LOLP and LOLR to decay *faster* with increasing battery size. This is because demand curtailment provides an additional 'upward' drift to the battery level, making it less likely to drop to zero. Thus, the value of $\lambda_{c,DR}$ provides a quantification of the impact of DR from the standpoint of the utility.

Next, we focus on the long run fraction of time demand curtailment is applied (i.e., DR is active). This is an important performance metric; in practice, one might expect that the DR contract between the utility and customers would specify an upper bound on how often demand curtailment can occur. Under our notation, note that the long run fraction of time demand is curtailed is $\mathbf{1}^T F(b_{\min})$, where $\mathbf{1}$ is a vector of ones having dimension $|S|$.

THEOREM 3.2. *Suppose that* $\Delta > 0$, *i.e., the average steady state generation exceeds the average steady state demand. Under the DR model proposed in Section 2, for* $\beta \in (0, 1)$,

$$\lim_{b_{\max}\to\infty} \frac{log(\mathbf{1}^T F(b_{\min}))}{b_{\max}} = -(1-\beta)\lambda_c,$$

where λ_c *is the smallest positive eigenvalue of the post-threshold matrix* $R^{-1}Q^T$.

Interestingly, the decay rate of the (steady state) probability of demand curtailment is simply $(1-\beta)$ times the decay rate of the LOLP in the absence of DR. Intuitively, this is because the asymptotic behavior of both quantities is dictated by the rare event that a Markov modulated random walk with positive drift, regulated from above at the level $b_{\max}$, drops below its maximum value by a certain large amount, say z. The probability of this rare event, when z is large, is approximately $Ce^{-\lambda_c z}$. Now, in the case of LOLP in the absence of DR, we have $z = b_{\max}$, whereas for the steady state probability of demand curtailment, we have $z = (1-\beta)b_{\max}$.

It follows from the above discussion that one can in fact approximate the steady state probability of demand curtailment as

$$\mathbf{1}^T F(b_{\min}) \approx Ce^{-(1-\beta)\lambda_c b_{\max}}, \tag{8}$$

where $C > 0$ is the same pre-factor that appears in (6). The approximation (8) can be used to determine the values of β and/or $b_{\max}$ in order to meet a prescribed threshold, say δ, on the fraction of time demand curtailment is performed. Specifically, if $b_{\max}$ is large enough, (8) implies that a suitable value for β would be: $\beta \approx 1 - \frac{\log(C/\delta)}{\lambda_c b_{\max}}$.

Finally, we note that $\lambda_{c,DR} > \lambda_c > (1-\beta)\lambda_c$ for $\beta > 0$. In other words, the decay rate associated with the LOLP (in the presence of DR) strictly exceeds that associated with the long run fraction of time that DR is applied. This is to be expected, since loss of load is a *rarer* event as compared to demand curtailment under our model.

4 CONCLUDING REMARKS

A case study demonstrating the applicability of these approximations in practice will appear in the journal version of this work. The primary limitation of the proposed model is that it does not capture battery inefficiency, i.e., it assumes that there are no energy losses associated with the charging/discharing of the battery. Extending the results presented to account for battery inefficiency presents a promising avenue for future work.

One of the motivations behind the present work is that models and tools developed in the networking community to deal with bursty arrival processes, buffer overflows, and capacity provisioning problems can be productively applied to the power grid, the primary source of randomness being the intermittency of renewables. We hope the present work motivates future research along these lines.

REFERENCES

[1] David Anick, Debasis Mitra, and Man M Sondhi. 1982. Stochastic theory of a data-handling system with multiple sources. *Bell System Technical Journal* 61, 8 (1982), 1871–1894.

[2] Vivek Deulkar, Jayakrishnan Nair, and Ankur A Kulkarni. 2020. Sizing Storage for Reliable Renewable Integration: A Large Deviations Approach. *Journal of Energy Storage* 30 (2020), 101443.

[3] Anwar I Elwalid and Debasis Mitra. 1992. Fluid models for the analysis and design of statistical multiplexing with loss priorities on multiple classes of bursty traffic. In *[Proceedings] IEEE INFOCOM'92: The Conference on Computer Communications*. IEEE, 415–425.

[4] Torgeir Ericson. 2009. Direct load control of residential water heaters. *Energy Policy* 37, 9 (2009), 3502–3512.

[5] A Gomes, CH Antunes, and AG Martins. 2007. A multiple objective approach to direct load control using an interactive evolutionary algorithm. *IEEE Transactions on Power Systems* 22, 3 (2007), 1004–1011.

[6] Nicholas Good, Efthymios Karangelos, Alejandro Navarro-Espinosa, and Pierluigi Mancarella. 2015. Optimization under uncertainty of thermal storage-based flexible demand response with quantification of residential users' discomfort. *IEEE Transactions on Smart Grid* 6, 5 (2015), 2333–2342.

[7] Pavithra Harsha and Munther Dahleh. 2014. Optimal management and sizing of energy storage under dynamic pricing for the efficient integration of renewable energy. *IEEE Transactions on Power Systems* 30, 3 (2014), 1164–1181.

[8] Miao He, Sugumar Murugesan, and Junshan Zhang. 2011. Multiple timescale dispatch and scheduling for stochastic reliability in smart grids with wind generation integration. In *2011 Proceedings IEEE INFOCOM*. IEEE, 461–465.

[9] Libin Jiang and Steven Low. 2011. Multi-period optimal energy procurement and demand response in smart grid with uncertain supply. In *2011 50th IEEE conference on decision and control and European control conference*. IEEE, 4348–4353.

[10] Jeff St. John. 2020. In New York, a New Way for Stay-at-Home Customers to Get Paid for Shaving Peak Energy . https://www.greentechmedia.com/articles/read/in-new-york-a-new-way-for-stay-at-home-customers-to-get-paid-for-shaving-peak-energy/. [Online; accessed 20-February-2021].

[11] Salman Kahrobaee, Sohrab Asgarpoor, and Wei Qiao. 2013. Optimum sizing of distributed generation and storage capacity in smart households. *IEEE Transactions on Smart Grid* 4, 4 (2013), 1791–1801.

[12] Tùng T Kim and H Vincent Poor. 2011. Scheduling power consumption with price uncertainty. *IEEE Transactions on Smart Grid* 2, 3 (2011), 519–527.

[13] Mohammad Amin Mirzaei, Ahmad Sadeghi-Yazdankhah, Behnam Mohammadi-Ivatloo, Mousa Marzband, Miadreza Shafie-khah, and João PS Catalão. 2019. Integration of emerging resources in IGDT-based robust scheduling of combined power and natural gas systems considering flexible ramping products. *Energy* 189 (2019), 116195.

[14] Debasis Mitra. 1988. Stochastic theory of a fluid model of producers and consumers coupled by a buffer. *Advances in Applied Probability* 20, 3 (1988), 646–676.

[15] M. Parsa Moghaddam, A. Abdollahi, and M. Rashidinejad. 2011. Flexible demand response programs modeling in competitive electricity markets. *Applied Energy* 88, 9 (2011), 3257 – 3269. https://doi.org/10.1016/j.apenergy.2011.02.039

[16] Morteza Nazari-Heris, Mohammad Amin Mirzaei, Behnam Mohammadi-Ivatloo, Mousa Marzband, and Somayeh Asadi. 2020. Economic-environmental effect of power to gas technology in coupled electricity and gas systems with price-responsive shiftable loads. *Journal of Cleaner Production* 244 (2020), 118769.

[17] Daniel O'Neill, Marco Levorato, Andrea Goldsmith, and Urbashi Mitra. 2010. Residential demand response using reinforcement learning. In *2010 First IEEE international conference on smart grid communications*. IEEE, 409–414.

[18] Peter M van de Ven, Nidhi Hegde, Laurent Massoulié, and Theodoros Salonidis. 2013. Optimal control of end-user energy storage. *IEEE Transactions on Smart Grid* 4, 2 (2013), 789–797.

[19] Yangfang Helen Zhou, Alan Scheller-Wolf, Nicola Secomandi, and Stephen Smith. 2018. Managing Wind-Based Electricity Generation in the Presence of Storage and Transmission Capacity. *Production and Operations Management* (2018).