

Distributed control of a fleet of batteries

Ana Bušić¹, Md Umar Hashmi¹, and Sean Meyn²

Abstract—Battery storage is increasingly important for grid-level services such as frequency regulation, load following, and peak-shaving. The management of a large number of batteries presents a control challenge: How can we solve the apparently combinatorial problem of coordinating a large number of batteries with discrete, and possibly slow rates of charge/discharge? The control solution must respect battery constraints, and ensure that the aggregate power output tracks the desired grid-level signal.

A distributed stochastic control architecture is introduced as a potential solution. Extending prior research on distributed control of flexible loads, a randomized decision rule is defined for each battery of the same type. The power mode at each time-slot is a randomized function of the grid-signal and its internal state. The randomized decision rule is designed to maximize idle time of each battery, and keep the state-of-charge near its optimal level, while ensuring that the aggregate power output can be continuously controlled by a grid operator or aggregator. Numerical results show excellent tracking, and low stress to individual batteries.

I. INTRODUCTION

Future power grids will have greater volatility due to the higher percentage of weather dependent renewable energy sources connected to the grid. These sources are intermittent, and their power generation has sharp peaks and valleys. This makes load balancing and frequency regulation challenging.

Storage systems have enormous potential value for a range of services. Individual consumers can use storage to enable more effective use of residential solar power, or to reduce electricity costs in regions with time-of-use pricing. Battery energy storage systems (BESS) are increasingly important for grid-level services such as frequency regulation, load following, peak-shaving, and deferral of investment in transmission and generation resources [1].

The value of responsive regulation is now recognized and incentivized through FERC Order 755, issued in 2011. FERC 755 report [2, pages 23-24] contains a survey of experiments conducted by Beacon Power and Primus Power to determine the value of highly responsive resources for ancillary service. Primus claims that batteries provide approximately 76% more “area control error” (ACE) correction when compared with traditional service from generation.

Among the goals of the present paper are:

- Address the apparent combinatorial problem of coordinating a large number of geographically distributed batteries to provide grid services.

Research supported by ANR grant ANR-16-CE05-0008, NSF grant CPS-1259040, and by PGM0 (Gaspard Monge Program for Optimization and operations research).

¹A.B and Md U.H. are with INRIA and the Computer Science Dept. of Ecole Normale Supérieure, Paris, France.

²S.M. is with the Dept. ECE at the University of Florida, Gainesville.

- Ensure that each battery within the fleet does not violate local constraints, such as bounds on the state of charge (SoC).

One challenge is that the grid level signals exhibit ramps and volatility that may be beyond reach for any individual battery. Another is that many battery technologies exhibit strong asymmetry with respect to charging/discharging rates; this is particularly true for lithium-ion technology.

Whether the batteries are distributed across a region or in a central location, techniques are required to ensure that the fleet of batteries can perform desired grid services reliably and accurately, while respecting the inherent constraints of individual batteries [3]–[5]. For most cell chemistries, the effective capacity of a battery is increased with slow discharge rates [6]. It is argued in [7] that the battery lifetime is reduced when discharge current is subject to high variance.

Distributed control is adopted so that the complexity of the control solution does not grow with the number of batteries. It is useful to employ randomization to introduce additional degrees of freedom for control. Each battery in the fleet is modeled as a controlled Markov chain, in which the controlled transition probabilities are common among each battery of the same type; this defines a local randomized control architecture. The randomized decision rule is designed to encourage idle time for each battery, and keep the state-of-charge near its optimal level, while ensuring that the aggregate power output can be continuously controlled by a grid operator or aggregator.

The aggregate behavior is approximated by a mean field model, which is a nonlinear input-output system, as in [8], [9]. This concept was brought to the power systems community in [16]. In numerical results it is found that linearized dynamics are nearly flat over a large bandwidth. An additional macro feedback loop is used to ensure accurate tracking; integral control is justified because of the simple aggregate dynamics.

It is found in numerical experiments that the goals can be achieved with this hierarchical control architecture: the aggregate of batteries accurately tracks the grid signal, and constraints on the battery SoC are strictly enforced. It is also found that individual batteries behave in a nearly deterministic manner – randomization leads to smooth input-output dynamics without causing “chattering” of individual batteries.

Related literature

The basic distributed control architecture developed in this paper is based upon prior work on “demand dispatch” – the

creation of virtual energy storage from flexible loads [9]–[11]. For loads whose power consumption cannot be varied continuously, a distributed randomized control architecture provides the needed degrees of freedom to track a smooth power reference signal, despite the discrete nature of each load. The use of randomization has been adopted in other work, such as [12]–[14]. A significant difference is the design approach: it is argued here and in our prior research that with appropriate local control at the load or individual battery, the aggregate of resources will appear to the grid operator as an input-output system that is easily controlled. Local control also ensures that strict bounds on quality of service to the load are maintained.

While there are many papers on BESS, we are not aware of any work with a comparable distributed control architecture, or comparable results. The prior work [8], [15] considers low-frequency services such as “valley filling”, based on BESS in which each battery is a residential electric vehicle. There are many recent papers on managing generation along with grid level storage, such as [17], [18].

The recent work [5] surveys mathematical models for estimating the cost of running LiIon batteries, emphasizing the importance of maintaining the SoC near its optimal level. This paper and [19] focus on control loops to regulate the SoC level, and the latter focuses on estimating potential revenue from BESS in typical ancillary service markets. Reference [20] considers a nonlinear model for a single battery, and compares control techniques to address the conflicting goals of managing SoC and tracking a regulation signal. These papers provided part of the motivation for the present work.

While the setting is entirely different, [7] provides insights and control techniques for managing batteries in portable electronic devices. An MDP (Markov Decision Process) model is proposed to obtain good performance for the device, while ensuring that the battery is not subject to stress. It is claimed that the battery control system will extend the battery service lifetime by more than 20%. It is not clear if the methodology of [7] can be extended to grid applications.

II. DISTRIBUTED CONTROL DESIGN

Consider a large collection of batteries, potentially distributed across a large geographical region within a single balancing authority (BA). In numerical experiments we will consider data from PJM; the largest BA in the U.S. [21]. A BA maintains supply-demand balancing within its operating region by regulating frequency to its nominal value (60 Hz in the U.S.), and regulating power flow between adjacent BAs. This requires resources to provide ancillary services – traditionally, power output from generators is ramped up and down to track a signal broadcast from the BA. At PJM, for frequency regulation there are two signals: RegA and RegD. The latter is distinguished by its higher frequency content, and it is also conditioned so that it is approximately zero energy over long time-horizons. Fig. 1 shows the cumulative sum of the RegD signal over a typical day in 2015. While

the sum eventually returns to zero, on this day it took a long positive excursion before returning to zero at midnight.




Fig. 1. Cumulative sum of PJM RegD – data from January 4, 2015.

The distributed control approach of [9], [22] is based on a Markovian nominal model for an individual resource (flexible loads in these papers). The main difficulty of extending this approach to batteries is that the nominal behavior for batteries is highly deterministic: if there is no need for storage, the batteries should stay idle at the mid-range SoC value. This will extend battery life, and also ensure that there is ancillary service capacity in each direction (both charging and discharging). The first step in our design is to construct a stochastic nominal model that is close to this deterministic nominal behavior.

A. Nominal model design

Fig. 2 illustrates the overall control architecture in which design is based on two components: there is a local decision rule at each battery based on a common signal $\{\zeta_t\}$ that is broadcast to all batteries. Error feedback based on the regulation signal $\{r_t\}$ and measurements of aggregate power is used to ensure accurate tracking.

Time is discrete with time steps of h seconds. In the next section, for PJM RegD signal tracking we will use $h = 2$. We assume that an individual battery can be in three different modes of operation: charging, discharging and idle. The state of a battery has two components, the mode and the state of charge (SoC):

$$x = (m, s),$$

where $m \in \{\text{ch}, \text{dis}, \text{id}\}$ and $s \in [0, 1]$, with 1 corresponding to a fully charged battery.

While a battery is in a charging mode ($m = \text{ch}$), it is charging with a constant rate denoted δ_{ch} . Its new state after h seconds is $(\text{ch}, s + h\delta_{\text{ch}})$. Similarly, if a battery is discharging, then its new state after h seconds is $(\text{dis}, s - h\delta_{\text{dis}})$. A battery that is idle does not change its SoC, so the new state after h seconds is (id, s) .

We describe next the change of mode of operation of a battery. Denote by $p_{\text{ch}}, p_{\text{dis}}, p_{\text{id}} : [0, 1] \rightarrow [0, 1]$ the functions that model the probability to stay in the charging, discharging or idle mode respectively. Bounds on SoC impose the constraint $p_{\text{ch}}(s) = 0$, for all $s > 1 - h\delta_{\text{ch}}$ and $p_{\text{dis}}(s) = 0$ for all $s < h\delta_{\text{dis}}$.

In a state (m, s) , a first decision (first coin flip) is made to decide if there is a change of mode of the battery. If so then the second decision is made to decide which of the remaining two modes the battery is going to switch to. This choice is done with the probabilities proportional to the values of the p functions of the remaining two modes.




Fig. 2. BESS control architecture: Local control at each battery combined with error feedback based on aggregate power is used to track the regulation signal. The power output of an individual battery Y_t^i takes on a finite number of values.

As an example, suppose the current state is $x = (\text{ch}, s)$. The new state, after h seconds is then:

$$\begin{aligned} x^+ &:= (\text{ch}, s + h\delta_{\text{ch}}), & \text{with probability } & p_0^+(x) \\ x^0 &:= (\text{id}, s), & \text{with probability } & p_0^0(x) \\ x^- &:= (\text{dis}, s - h\delta_{\text{dis}}), & \text{with probability } & p_0^-(x) \end{aligned} \quad (1)$$

where $p_0^+(x) = p_{\text{ch}}(s)$,

$$p_0^0(x) = (1 - p_{\text{ch}}(s)) \frac{p_{\text{id}}(s)}{p_{\text{id}}(s) + p_{\text{dis}}(s)}$$

$$\text{and } p_0^-(x) = (1 - p_{\text{ch}}(s)) \frac{p_{\text{dis}}(s)}{p_{\text{id}}(s) + p_{\text{dis}}(s)}$$

The design of $p_{\text{ch}}, p_{\text{dis}}, p_{\text{id}}$ used in the numerical results of this paper is given in Fig. 3. In this design, the target SoC interval was set to 40 - 80% SoC. If the battery is charging, it will remain charging with probability almost 1 until it reaches 40% SoC. The probability to keep charging then decreases and reaches almost 0 at 80% SoC. The design of p_{dis} is symmetrical. The function p_{id} has values almost 1 for 50-70% SoC values and it is almost 0 outside the target interval.




Fig. 3. Design of switching probability functions $p_{\text{ch}}, p_{\text{dis}}, p_{\text{id}}$.

The nominal Markov model for an individual battery is then given by a transition kernel P_0 . Its state space \mathbf{X} is the union of three intervals; let \mathcal{B} denotes its Borel sigma field. For any state x , the probability measure $P_0(x, \cdot)$ is supported on at most three states. From (1) we see that for the state $x = (\text{ch}, s)$, and for any $A \in \mathcal{B}$, $P_0(x, A) =$

$$p^+(x)\mathbb{I}\{x^+ \in A\} + p^0(x)\mathbb{I}\{x^0 \in A\} + p^-(x)\mathbb{I}\{x^- \in A\}$$

B. Controlled Markov model for an individual battery

Based on this nominal model, a controlled Markov model for an individual battery can be obtained following similar steps as in [22]. The main difference is that our model evolves on a continuous state space compared to the discrete state space design in this previous work.

Let $\mathcal{U}(x)$ denote the power output at state x ; it is assumed that this takes on just three values: $\mathcal{U}(x) = 0$ for $x = (\text{id}, s)$,

and the remaining values are consistent with the Markov model:

$$\frac{|\mathcal{U}(\text{ch}, s)|}{\mathcal{U}(\text{dis}, s)} = \frac{\delta_{\text{ch}}}{\delta_{\text{dis}}}, \quad 0 \leq s \leq 1.$$

For any fixed value of ζ , the probability measure $P_\zeta(x, \cdot)$ is also supported on at most three states. For the family of transition matrices, the three states (x^+, x^0, x^-) depend on x but not on ζ : $P_\zeta(x, A) = p_\zeta^+(x)\mathbb{I}\{x^+ \in A\} + p_\zeta^0(x)\mathbb{I}\{x^0 \in A\} + p_\zeta^-(x)\mathbb{I}\{x^- \in A\}$, $x \in \mathbf{X}$, $A \in \mathcal{B}$.

The values of $p_\zeta^+(x)$, $p_\zeta^0(x)$ and $p_\zeta^-(x)$ are set as follows:

$$\begin{aligned} p_\zeta^+(x) &= p_0^+(x) \exp(\zeta \mathcal{U}(x^+) - \Lambda_\zeta(x)) \\ p_\zeta^0(x) &= p_0^0(x) \exp(\zeta \mathcal{U}(x^0) - \Lambda_\zeta(x)) \\ p_\zeta^-(x) &= p_0^-(x) \exp(\zeta \mathcal{U}(x^-) - \Lambda_\zeta(x)) \end{aligned} \quad (2)$$

with $\Lambda_\zeta(x) = \log\left(\sum_{m \in \{-, 0, +\}} p_0^m(x) \exp(\zeta \mathcal{U}(x^m))\right)$.

This is similar to the *myopic design* introduced in [23]. It can be interpreted as an instance of the optimal design introduced in [9], but with a time-horizon of *one* rather than infinite. This transformation of the nominal model will encourage discharge when $\zeta > 0$, and charging when $\zeta < 0$.

C. Mean Field Model

The family of transition kernels $\{P_\zeta : \zeta \in \mathbb{R}\}$ is constructed to define local decision making: Each battery evolves as a controlled Markov chain on a finite state space, with common input $\zeta = (\zeta_0, \zeta_1, \dots)$. It is assumed that the scalar signal ζ is broadcast to each battery. If a battery is in state x at time t , and the value ζ_t is broadcast, then the battery transitions to the state x' according to the probability measure $P_{\zeta_t}(x, \cdot)$.

Let $X_t^i = (M_t^i, S_t^i)$ denote the state of the i th battery at time t . The dynamics of the first component are governed by the randomized policy; the second component denotes SoC, which evolves as a controlled random walk:

$$\begin{aligned} S_{t+1}^i &= S_t^i + h\delta_{\text{ch}} & \text{charge mode, } M_t^i &= 1 \\ S_{t+1}^i &= S_t^i - h\delta_{\text{dis}} & \text{discharge mode, } M_t^i &= -1 \\ S_{t+1}^i &= S_t^i & \text{idle, } M_t^i &= 0 \end{aligned} \quad (3)$$

The empirical probability measure is defined as the average,

$$\mu_t^N(A) = \frac{1}{N} \sum_{i=1}^N \mathbb{I}\{X_t^i \in A\}, \quad A \in \mathcal{B}.$$

As in prior work for the discrete state space [9], [10], [22], we assume that average power output is obtained through




Fig. 4. Bode plots for linear models obtained using Least Squares.

measurements or state estimation. At time t , this is expressed in terms of the empirical probability measure:

$$y_t^N = \frac{1}{N} \sum_{i=1}^N Y_t^i = \int \mathcal{U}(x) \mu_t^N(dx), \quad t \geq 0, \quad x \in \mathcal{X},$$

where $Y_t^i = \mathcal{U}(X_t^i)$ as shown in Fig. 2.

The mean field model is defined by the recursion

$$\mu_{t+1}(A) = \int P_{\zeta_t}(x, A) \mu_t(dx) \quad (4)$$

Under general conditions on the model and on μ_0 it can be shown that μ_t^N is approximated by μ_t . One condition is that the input ζ_t is a continuous function of $(y_0^N, \dots, y_{t-1}^N, \dots)$ that depends upon t , but does not depend upon N .

Under these conditions, the average power output is approximated using the mean-field model:

$$y_t = \int \mathcal{U}(x) \mu_t(dx), \quad t \geq 0. \quad (5)$$

In prior work on demand dispatch it is found that although the input-output mean field dynamics are nonlinear, the dynamics are accurately approximated by a linear model, and often the linearized dynamics are minimum phase [9]. Desirable aggregate behavior can also be obtained through design [22], [23].

III. NUMERICAL RESULTS

The following three scenarios were considered, differentiated by their time-to-charge T_{ch} , time-to-discharge T_{dis} , and power ratings:

Scenario	T_{ch}	$\mathcal{U}(\text{ch}, s)$	T_{dis}	$\mathcal{U}(\text{dis}, s)$
S1:	2hr	-0.5kW	2hr	0.5kW
A1:	2hr	-0.5kW	1hr	1kW
A2:	2hr	-0.5kW	30min	2kW

In each case, a collection of 1,000 batteries can deliver at most 500kW of power. In the symmetric model the lower limit is -500kW. The regulation signal was scaled to respect these constraints.

In each of these three scenarios, a linear approximation of the input-output dynamics was obtained using Least Squares. A seventh order model was sufficient in all cases.

Fig. 4 shows the Bode plots obtained for models S1 and A1. In each case, the magnitude plot is nearly constant over the frequency range of interest, $f \in [10^{-4}, 5 \times 10^{-2}]$.




Fig. 5. Bode plot, margins and root locus for the loop transfer function.

While the magnitude plots are nearly perfectly flat, the true dynamics are only approximately linear. To achieve accurate tracking we introduce error feedback, with input ζ_t , error $e_t = r_t - y_t^N$, and “plant transfer function” G_p : the linear model for input-output dynamics.

Consider the symmetric model S1 for which the Bode plot of G_p is shown in Fig. 4. Design of the compensator G_c is simplest in continuous time. The loop transfer function is the product $G = G_c G_p$, and the compensator G_c is constructed so that it has high gain at low frequencies, and the crossover frequency f_c is chosen at the desired closed-loop bandwidth. Since by definition $|G(jf_c)|_{\text{dB}} = 0$, a glance at Fig. 4 implies that $|G_c(jf_c)|_{\text{dB}} \approx 35$ if $f_c \in [10^{-4}, 5 \times 10^{-2}]$.

The flat magnitude plot for G_p motivates the use of pure integral control, $G_c(s) = K_I^\circ s^{-1}$, with $K_I^\circ = 50f_c$ (using $|50|_{\text{dB}} \approx 35$). This is implemented in discrete time:

$$\zeta_t = K_I \sum_{i=0}^t e_i, \quad K_I = hK_I^\circ \quad (7)$$

where $h = 2$ is the sampling time in seconds.

On choosing $f_c = 5 \times 10^{-2}$ we obtain $K_I^\circ = 250 \times 10^{-2}$, and hence $K_I = 5$. Fig. 5 shows a Bode plot of the loop transfer function with this design (represented in continuous time), along with the root locus plot for the discrete-time loop transfer function.

A. Tracking and SoC Performance




Fig. 6. Tracking of PJM’s RegD Normalized Signal Test in Scenario S1 with 1000 batteries. Also shown is the control signal (7)

All experiments use PJM’s Normalized Signal Test for RegD to define the reference signal $\{r_t\}$ [21]. The initial condition for each battery was at rest, with SoC at 50%. The impact on the initial condition was insignificant in these and other experiments.

Consider first the symmetric model, Scenario S1, in which each battery requires 2hr to fully charge (resp. discharge),




Fig. 7. Evolution of SoC for the symmetric model.

and draws 0.5 kW when charging. The tracking results and control signal obtained using the integral gain $K_I = 5$ are displayed in Fig. 6.

It is found that an individual battery operates in a nearly deterministic manner, and the SoC is maintained to the desired limits, even though the local control is randomized. The plots marked “ $\eta = 1.0$ ” in Fig. 7 show a comparison of the mean SoC and the evolution of SoC for an individual battery. All goals have been achieved: grid level tracking is nearly perfect, and battery constraints are maintained.

Tracking results obtained with only 200 batteries are shown in Fig. 8. There is some extra high-frequency error due to quantization, but the performance is well within the requirements of PJM.




Fig. 8. Tracking of PJM's RegD Normalized Signal Test in Scenario S1 with only 200 batteries

The PI design was repeated in the two asymmetric models. Tracking is more challenging in these cases. This is true in part because there is a capacity reduction associated with asymmetry. It is also found that nonlinear effects are more apparent in asymmetric models.

Results for the asymmetric case A2 are displayed in Fig. 9. It is clear that the true system is nonlinear: when the regulation signal is large and positive, the closed loop system appears to be under-damped (oscillation is observed). When the regulation signal is negative then the closed loop system appears over-damped. Similar but less extreme results were observed in the case A1.

B. Impact of efficiency loss

In the preceding experiments efficiency losses have been ignored. For example, the round trip efficiency of Tesla's Powerwall is just 92.5%. A linear model was used to test the impact of efficiency losses. For an efficiency parameter $\eta \in (0, 1]$, the evolution of the state of charge of battery i is




Fig. 9. Tracking of PJM's RegD Normalized Signal Test in Scenario A2.

given by

$$\begin{aligned} S_{t+1}^i &= S_t^i + \eta \times h \delta_{\text{ch}} & \text{charge mode, } M_t^i &= 1 \\ S_{t+1}^i &= S_t^i - h \delta_{\text{dis}} / \eta & \text{discharge mode, } M_t^i &= -1 \end{aligned} \quad (8)$$

The power output of the battery remains the same: $Y_t^i = \mathcal{U}(X_t^i)$ for each t and i .

In the results described below we take $\eta = 0.9$. For the 40 minute PJM test signal, this loss presents no impact on tracking performance. There is an observable negative drift in SoC over this time interval – typical results are shown in Fig. 7 in Scenario S1 with 1,000 batteries. The SoC at the end of the experiment is slightly smaller than observed with $\eta = 1$. This would be unacceptable on a 24 hour run: an additional feedback loop is required to draw additional power to make up for losses.




Fig. 10. Mileage nearly reaches the ideal with and without losses. The error at the end of the run is under 3% in case S1.

Excess operation of the fleet of batteries was computed based on *mileage* – a traditional metric of service from batteries [2]. For an individual battery, the cumulative mileage by time T and its average are defined by

$$\mathcal{E}_T^k = \sum_{t=1}^T |\mathcal{U}(X_t^k)|, \quad \mathcal{R}_T = \frac{1}{N} \sum_{k=1}^N \mathcal{R}_T^k \quad T \geq 1. \quad (9)$$

The following sequence is regarded as a metric for *excess operation*:

$$\mathcal{E}_T = \frac{\mathcal{R}_T - \mathcal{R}_T^*}{\mathcal{R}_T^*}, \quad \text{where } \mathcal{R}_T^* = \frac{1}{N} \sum_{t=1}^T |r_t|$$

It can be shown using Jensen's inequality that $\mathcal{E}_T \geq 0$ in the ideal case of perfect tracking. We obtain $\mathcal{E}_T = 0$ if each battery tracks the reference exactly.

With \mathcal{E} defined to be the final value at the end of the 40min PJM test signal, the following are typical values:

Scenario	$\mathcal{E}, \eta = 1$	$\mathcal{E}, \eta = 0.9$
S1	2.51%	2.54%
A1	3.30%	3.62%

with scenario data given in (6). The impact of losses is negligible in these experiments. In all cases, it is surprising to see only about 3% beyond the ideal mileage using a distributed randomized control strategy.

The evolution of mileage in these four cases is shown in Fig. 10. The close match with the ideal means that most batteries in the fleet are cooperating, in the sense that the sign of $\mathcal{U}(X_t^i)$ matches the sign of r_t , or is equal to zero, for the overwhelming majority of indices $i = 1, \dots, 1,000$.

IV. CONCLUSIONS

It is remarkable that a fleet of ‘dumb, slow batteries’ can accurately track a grid level signal with much faster temporal characteristics, while maintaining individual SoC within desired bounds.

There is of course much more work to be done on local control design, and control techniques for the aggregate. What is the best way to resolve the asymmetric behavior observed in Fig. 9? One approach is at the macro-level: gain scheduling could be used, in which the gain K_I is reduced when the regulation signal is positive, and increased when it is negative. An alternative is to refine the local control design to make the aggregate appear more linear, as in [23].

Further research is needed to investigate the heterogeneous setting (e.g. with respect to the capacity, charging and discharging rates), and the robustness to un-modelled dynamics (e.g. charging/discharging rates that depend on the current SoC level). Another direction concerns state estimation both for the aggregate power consumption used as a control feedback and state estimation for the local control.

REFERENCES

- [1] G. Fitzgerald, J. Mandel, J. Morris, , and H. Touati, “The economics of battery energy storage: How multi-use, customer-sited batteries deliver the most services and value to customers and the grid,” Rocky Mountain Institute, Technical Report, September 2015. [Online]. Available: http://www.rmi.org/electricity_battery_value
- [2] R. Pedroncelli, “Frequency Regulation Compensation in the Organized Wholesale Power Markets – FERC 755,” FERC Docket Nos. RM11-7-000 and AD10-11-000; Order No. 755 – Online, <http://tinyurl.com/FERC755>, October 20 2011.
- [3] E. M. Krieger, “Effects of variability and rate on battery charge storage and lifespan,” Ph.D. dissertation, Princeton University, 2013. [Online]. Available: <http://tinyurl.com/KriegerBatteryThesis13>
- [4] R. Carnegie, D. Gotham, D. Nderitu, and P. Preckel, “Utility scale energy storage systems,” State Utility Forecasting Group, Purdue University, West Lafayette, Indiana, Tech. Rep., 2013.
- [5] S. Karagiannopoulos, A. Rigas, N. Hatziaargyriou, G. Hug, and A. Oudalov, “Battery energy storage capacity fading and control strategies for deterministic and stochastic power profiles,” in *2016 Power Systems Computation Conference (PSCC)*, June 2016, pp. 1–7.
- [6] B. Lawson, “Electropaedia: Energy storage and power generation technologies,” online <http://www.mpoweruk.com/performance.htm>, 2016.
- [7] P. Rong and M. Pedram, “Battery-aware power management based on Markovian decision processes,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 25, no. 7, pp. 1337–1349, July 2006.
- [8] Z. Ma, D. Callaway, and I. Hiskens, “Decentralized charging control for large populations of plug-in electric vehicles,” in *49th IEEE Conference on Decision and Control*, Dec 2010, pp. 206–212.
- [9] S. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren, “Ancillary service to the grid using intelligent deferrable loads,” *IEEE Trans. Automat. Control*, vol. 60, no. 11, pp. 2847–2862, Nov 2015.
- [10] Y. Chen, A. Bušić, and S. Meyn, “Individual risk in mean field control with application to automated demand response,” in *53rd IEEE Conference on Decision and Control*, Dec 2014, pp. 6425–6432.
- [11] —, “State estimation for the individual and the population in mean field control with application to demand dispatch,” *CoRR and to appear, IEEE Transactions on Auto. Control*, 2016. [Online]. Available: <http://arxiv.org/abs/1504.00088v1>
- [12] S. Koch, J. L. Mathieu, and D. S. Callaway, “Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services,” in *Proc. 17th Power Systems Computation Conference*, 2011, pp. 1–7.
- [13] J. Mathieu, S. Koch, and D. Callaway, “State estimation and control of electric loads to manage real-time energy imbalance,” *IEEE Trans. Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.
- [14] C. Ziras, E. Vrettos, and G. Andersson, “Primary frequency control with refrigerators under startup dynamics and lockout constraints,” in *IEEE Power & Energy Society General Meeting*, 2015, pp. 1–5.
- [15] D. Callaway and I. Hiskens, “Achieving controllability of electric loads,” *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, January 2011.
- [16] R. Malhame and C.-Y. Chong, “Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system,” *IEEE Trans. Automat. Control*, vol. 30, no. 9, pp. 854–860, Sep 1985.
- [17] A. I. Bejan, R. J. Gibbens, and F. P. Kelly, “Statistical aspects of storage systems modelling in energy networks,” in *46th Annual Conference on Information Sciences and Systems*, 2012, pp. 1–6.
- [18] I. Prodan and E. Zio, “An optimization-based control approach for reliable microgrid energy management under uncertainties,” in *IEEE Workshop on Integration of Stochastic Energy in Power Systems (ISEPS)*, Nov 2013, pp. 4–7.
- [19] B. Xu, A. Oudalov, J. Poland, A. Ulbig, and G. Andersson, “Bess control strategies for participating in grid frequency regulation,” *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 4024–4029, 2014.
- [20] J. L. Mathieu and J. A. Taylor, “Controlling nonlinear batteries for power systems: Trading off performance and battery life,” in *2016 Power Systems Computation Conference (PSCC)*. IEEE, 2016, pp. 1–7.
- [21] “Markets and operations at PJM,” Online, <http://tinyurl.com/PJMregs>, 2015.
- [22] A. Bušić and S. Meyn, “Distributed randomized control for demand dispatch,” in *IEEE Conference on Decision and Control*, Dec 2016, pp. 6964–6971.
- [23] J. Mathias, A. Bušić, and S. Meyn, “Demand dispatch with heterogeneous intelligent loads,” in *Proc. 50th Annual Hawaii International Conference on System Sciences (HICSS)*, and arXiv 1610.00813, 2017.