A Stochastic MPC Framework for Stationary Battery Systems

Ranjeet Kumar, Michael J. Wenzel, Matthew J. Ellis, Mohammad N. ElBsat
Kirk H. Drees, and Victor M. Zavala

***Abstract*—We present a stochastic model predictive control (MPC) framework to determine real-time commitments in energy and frequency regulation markets for a stationary battery system while simultaneously mitigating long-term demand charges for an attached load. The framework solves a two-stage stochastic program over a receding horizon that maximizes the expected profit and that factors in uncertainty of the load, energy prices, regulation prices and dispatch signals. We use a Ledoit-Wolf covariance estimator to generate load and price scenario profiles from limited historical data. We benchmark the performance of stochastic MPC against that of perfect information MPC and deterministic MPC for different prediction horizon lengths and demand charge discounting strategies. We use real load data for a typical university campus and price and regulation data from PJM. We find that stochastic MPC can recover 83% of the ideal value of the battery, which we define as the expected savings obtained by installing the battery and operating it under perfect information MPC. In contrast, deterministic MPC can only recover 73% of this ideal value. We also find that operating the battery under stochastic MPC improves the battery payback period by 12.1% while operating it under perfect information improves it by 27.9%.**

***Index Terms*—frequency regulation; energy markets; demand charges; stochastic systems; predictive control; battery storage**

I. Introduction

BATTERY systems are flexible assets that can provide energy and frequency regulation (FR) capacity for independent system operators (ISOs) and that can aid utility companies by providing demand-side management capabilities [1], [2]. This flexibility is becoming increasingly valuable as more intermittent renewable power is injected into the grid [3], [4]. In particular, ISOs have reported an increased demand for fast dispatchable resources and FR services to mitigate high-frequency power fluctuations and stranded power issues [5].

The economic benefits of using stationary battery systems to provide FR have been studied in [6]–[10]. The focus of [6]–[8] is on exploring the economic opportunities of using battery systems to provide FR capacity in day-ahead markets. The studies in [9] and [10] propose real-time control strategies for providing FR services to the grid. The potential of harnessing the flexibility of batteries from collections of electric vehicles has also been studied in [11]–[14]. The economic benefits of using battery systems to reduce demand charges in buildings and microgrids (via peak shaving) have been studied in [15]–[20]. A common limitation of these studies is that they assume that the battery either provides FR or demand-side management capabilities. The use of batteries for *simultaneous* FR and demand charge mitigation has been studied in [2], [14], [21]–[24]. These studies assume perfect knowledge of all the market and load parameters. Consequently, while these approaches are valuable for design and planning tasks, they can be limited in real-time operational settings where diverse uncertainties can hinder performance.

Stochastic optimization models have been proposed to maximize the performance of batteries from electric vehicles while capturing uncertainty in market prices and FR signals [3], [11], [25]–[27]. Some of these studies use stochastic dynamic programming methods to optimize charging policies and FR commitments [3], [11], [26]. In [3], model uncertainties in prices and FR signals as Markov random variables and demonstrate that their proposed method results in lower charging costs than a deterministic model predictive control (MPC) strategy. In [25] a strategy for computing optimal bidding strategies for an electric vehicle aggregator participating in day-ahead energy and FR markets is proposed. All of these studies consider mobile battery systems. Demand-side management in a commercial building using both mobile and stationary batteries was studied in [27].

We propose a stochastic MPC framework to obtain optimal participation strategies in energy and FR markets for stationary batteries while mitigating demand charges. We propose to use the Ledoit-Wolf covariance estimator to generate realizations from historical data. We use the framework to study the flexibility and economic benefits provided by a battery system attached to an aggregated load from a collection of buildings. We compare the benefits of stochastic MPC policies over those obtained with standard deterministic MPC and perfect information MPC strategies. We also study the effect of the prediction horizon length and demand charge discounting on the performance of the MPC policies. Using real load data for a typical university campus, we find that stochastic MPC can recover 83% of the ideal value of the battery (obtained by operating the battery under perfect information) while deterministic MPC can only recover 73%. We also find that operating the battery under stochastic MPC improves the battery payback period by 12.1% while operating it under perfect information improves it by 27.9%.

II. Problem Formulation

We begin by describing the decision-making setting under which the battery is operated. We also describe a deterministic MPC model that will serve as a basis for comparison.

A. Decision-Making Setting

The goal is to determine the optimal participation strategies for the battery in energy and FR markets operated by ISO (here we consider PJM) while simultaneously mitigating demand

Ranjeet Kumar and Victor M. Zavala are with the Department of Chemical and Biological Engineering,University of Wisconsin-Madison, 1415 Engineering Dr, Madison, WI 53706, USA (e-mail: victor.zavala@wisc.edu).

Michael J. Wenzel, Matthew J. Ellis, Mohammad N. ElBsat, and Kirk H. Drees are with Johnson Controls International, 507 E. Michigan St., Milwaukee, WI 53202, USA.

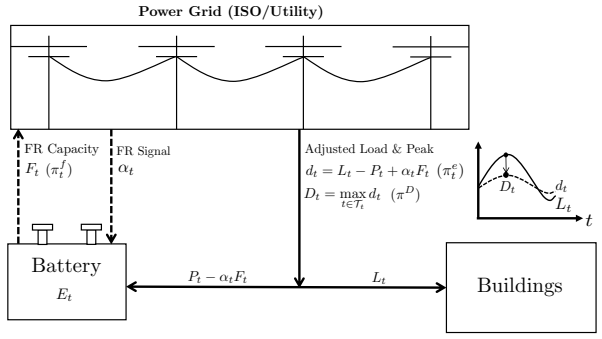

Fig. 1: Interactions between battery system, ISO, and utility.

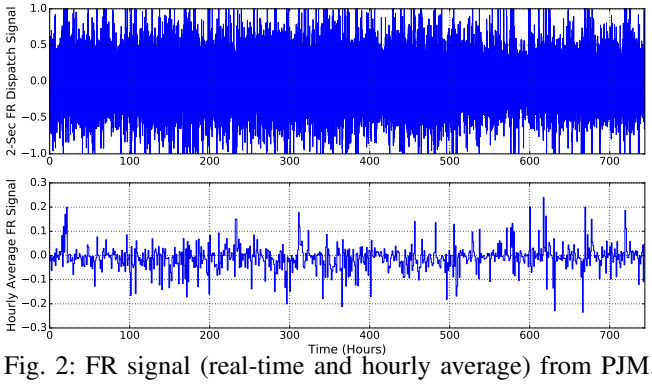

Fig. 2: FR signal (real-time and hourly average) from PJM.

charges from a utility company. The various cost and revenue components that are considered are:

- *Energy Transactions (hourly)*: The battery system purchases energy to recharge and discharges to provide energy for the building and for the FR signal. The energy transactions are charged at the real-time energy price.
- *Frequency Regulation Capacity (hourly)*: The ISO compensates the battery for providing an operational band (compensated based on FR capacity prices) around a charge/discharge level (charged at real-time energy prices). The ISO can request the battery to dispatch a fraction of the committed capacity based on the grid requirements (e.g., every 2 seconds in PJM). The real-time FR signal from the ISO is zero mean with a bounded range of [-1,+1]. In our planning setting we consider an hourly average FR signal. The hourly-averaged dispatch signal is also a zero-mean signal with a smaller variance (the variance is reduced due to averaging).
- *Demand Charges (monthly)*: The attached load is charged for the peak demand (at a fixed demand charge price) incurred over a month by the utility company.

According to PJM rules, the decision of allocating commitments can be changed 90 minutes before the start of that trading hour. Consequently, the battery has the opportunity to reschedule commitments. Such commitments can be computed under a receding horizon (RH) by using an MPC scheme.

III. Stochastic MPC Framework

The combined battery and load system is illustrated in Figure 1. This illustrates the interaction between the system, the ISO, and the utility. The proposed stochastic MPC scheme solves a two-stage stochastic programming problem at every time t over the prediction horizon $\mathcal{N}_t := \{t+1, t+2, ..., t+N\}$. Here, N is the length of the prediction horizon and index $t \in \mathbb{Z}_+$ corresponds to the time instant $t \cdot h$, where $h = 1$ hour. The optimization problem at time t uses realizations for prices, loads, and FR dispatch signals over $\mathcal{N}_t$. The MPC scheme is run over a month $\mathcal{M} := \{0, ..., M\}$, with $M = 720$.

The following are the model parameters, data, and variables involved used in the formulation of the MPC scheme:

1) Model Parameters and Data:

- $L_t \in \mathbb{R}_+$: Attached load [kW].
- $\pi_t^e \in \mathbb{R}$: Electricity price [$/kWh].
- $\hat{\pi}_t^e \in \mathbb{R}$: Forecast of electricity price [$/kWh].
- $\pi_t^f \in \mathbb{R}_+$: FR capacity price [$/kW].
- $\pi^D \in \mathbb{R}_+$: Demand charge [$/kW].
- $\alpha_t \in [-1, 1]$: Hourly average fraction of FR capacity requested by ISO [-]. If $\alpha_t > 0$, the ISO sends power to the battery while if $\alpha_t < 0$ the ISO withdraws power.
- $\overline{E} \in \mathbb{R}_+$: Battery capacity [kWh].
- $\overline{P} \in \mathbb{R}_+$: Maximum discharging rate [kW].
- $\underline{P} \in \mathbb{R}_+$: Maximum charging rate [kW].
- $\rho \in \mathbb{R}_+$: Minimum fraction of battery capacity reserved as a buffer for a unit FR capacity [kWh/kW].
- $\overline{\Delta P} \in \mathbb{R}_+$: Maximum ramping limit [kW/h].
- $\widehat{D}_t \in \mathbb{R}_+$: Peak load observed up to time $t \in \mathcal{M}$ [kW].

2) Model Variables:

- $P_t \in \mathbb{R}$: Net battery charge/discharge rate [kW]. If $P_t > 0$, the battery is being discharged and if $P_t < 0$ the battery is being charged.
- $F_t \in \mathbb{R}_+$: FR capacity provided to ISO [kW].
- $E_t \in \mathbb{R}_+$: Battery state of charge (SOC) [kWh].
- $d_t \in \mathbb{R}_+$: Load requested from utility [kW].
- $D_t = \max_{k \in \mathcal{N}_t} d_k$: Peak load over horizon $\mathcal{N}_t$ [kW]

All quantities with subindex t are *held constant* over the time interval $[(t-1) \cdot h, t \cdot h]$. The FR capacity F_t represents a *symmetric band* around the middle point P_t (the net charge/discharge rate). The battery thus commits to a discharge level P_t and an FR band of length $2 \cdot F_t$. The actual power requested by the ISO (averaged over an hour) is $\alpha_t F_t$. A value of $\alpha_t = -1$ indicates that the ISO requests all the FR capacity by charging the battery while $\alpha_t = 1$ indicates that the ISO requests all the capacity by discharging. When convenient, we use the notation $\pi_t = (\pi_t^e, \pi_t^f)$. We also use $\pi_{\mathcal{N}_t}$ to denote the price trajectory over $\mathcal{N}_t$. We use a similar notation to denote trajectories for load and FR signals.

The *uncertain quantities* (π_{t+1}^e, π_{t+1}^f, L_{t+1}, α_{t+1}) can only be forecasted an hour in advance (i.e., at time t). To represent uncertainty in these forecasts (and over any future times) we model them as random variables with realizations denoted by $\xi \in \Xi$ (where Ξ is the scenario set). At time t, a realization of the uncertain profiles for load, prices, and FR signals over $\mathcal{N}_t$ are denoted as $L_{\mathcal{N}_t}(\xi), \pi_{\mathcal{N}_t}(\xi)$ and $\alpha_{\mathcal{N}_t}(\xi)$, respectively. To account for uncertainty in the decision-making process, we define the net power P_{t+1} and FR commitments F_{t+1} over the next immediate trading hour as *first-stage decisions* (here-and-now) that need to be made prior to observing uncertainty (i.e., when the market is settled). The battery power and FR capacity trajectories $P_k(\xi)$ and $F_k(\xi)$ for all $k \in \mathcal{N}_t \setminus \{t+1\}$ are modeled as second-stage or *recourse decisions* that can

be corrected once uncertainty reveals. The residual and peak demands are also recourse variables that we express as $d_{\mathcal{N}_t}(\xi)$ and $D_t(\xi) = \max_{k\in\mathcal{N}_t} d_k(\xi)$. The SOC at time $t+1$ only depends on the previous storage E_t and onP_{t+1}. Consequently, E_{t+1} is also a first-stage variable. The rest of the trajectory $E_k(\xi)$ for $k \in \mathcal{N}_t \setminus \{t+1\}$ are second-stage variable because $P_k(\xi)$ are second-stage variables for $k \in \mathcal{N}_t \setminus \{t+1\}$.

A. Objective Function

We maximize the expected profit over horizon $\mathcal{N}_t$:

$$\mathbb{E}\left[\sum_{k\in\mathcal{N}_t}(\pi^f_k(\xi)F_k(\xi) - \pi^e_k(\xi)d_k(\xi)) - \frac{\pi^D}{\sigma_t}\max_{k\in\mathcal{N}_t} d_k(\xi)\right]. \tag{III.1}$$

Here, the expectation is computed using a scenario subset $\bar{\Xi} \subseteq \Xi$. The *net residual demand* is $d_k(\xi) = L_k(\xi) - P_k(\xi) + \alpha_k(\xi)F_k(\xi)$. The first term is the revenue obtained from the provision of FR *capacity*, the second term is the cost of purchasing power from the ISO, and the third term is the demand charge. The parameter σ_t is a *discounting factor* for the monthly demand charge price π^D, which can be used to adjust the demand charge when a prediction horizon of less than a month is used (i.e., $N \leq M$). When such a discounting factor is not used, the MPC scheme can operate conservatively because it will try to prevent the peak demand charge over the next immediate horizon $\mathcal{N}_t$. We note that the terms $\pi^e_k(\xi)L_k(\xi)$ in the total net profit are constant quantities and thus do not affect the solution of the optimization problem. Consequently, the term $\pi^f_k(\xi)F_k(\xi) - \pi^e_k(\xi)d_k(\xi)$ can be simplified as $\pi^e_k(\xi)(P_k(\xi) - \alpha_k(\xi)F_k(\xi)) + \pi^f_k(\xi)F_k(\xi))$.

B. Constraints

The net charged/discharged battery power plus the FR capacity provided must be within the maximum discharging and charging rates $\overline{P}$ and $\underline{P}$:

$$P_k(\xi) + F_k(\xi) \leq \overline{P},\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.2a}$$
$$P_k(\xi) - F_k(\xi) \geq -\underline{P},\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.2b}$$

These constraints make allowable battery charge/discharge limits a function of the committed FR capacity. The larger the FR commitment, the less capacity available to charge/discharge the battery. The storage dynamics are given by:

$$E_k(\xi) = E_{k-1}(\xi) - P_k(\xi) + \alpha_k(\xi)F_k(\xi),\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.3}$$

The following constraint is used to ensure that a certain amount of energy is reserved for the committed FR capacity over the interval $(k-1, k)$:

$$\rho F_k(\xi) \leq E_{k-1}(\xi) \leq \overline{E} - \rho F_k(\xi),\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.4a}$$
$$\rho F_k(\xi) \leq E_k(\xi) \leq \overline{E} - \rho F_k(\xi),\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.4b}$$

These constraints impose a safety margin to account for the fact that we do not capture the variability of the 2-second FR signal. Capturing such high time resolutions would significantly increase computational complexity. The ramping limits on the net power are given by:

$$-\overline{\Delta P} \leq P_k(\xi) - P_{k-1}(\xi) \leq \overline{\Delta P},\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.5}$$

where the allowable ramp $\overline{\Delta P}$ is often tuned to prevent premature damage to the battery. The residual demand d_k requested from the utility is:

$$d_k(\xi) = L_k(\xi) - P_k(\xi) + \alpha_k(\xi)F_k(\xi),\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.6}$$

We assume that the ISO does not allow the battery to sell back electricity. This is modeled by using the constraint:

$$P_k(\xi) + F_k(\xi) \leq L_k(\xi),\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.7}$$

This constraint can also be written as $L_k(\xi) - P_k(\xi) - F_k(\xi) \geq 0$. Consequently, when the ISO fully uses the FR capacity we have that $\alpha_k(\xi) = -1$ and the constraint implies that $d_k(\xi) \geq 0$ because $d_k(\xi) = L_k(\xi) - P_k(\xi) - \alpha_k(\xi)F_k(\xi)$. The constraint thus guarantees that the battery does not sell excess power. We use the following peak demand *carryover* constraint:

$$\max_{k\in\mathcal{N}_t} d_k(\xi) \geq \widehat{D}_t,\ \xi \in \bar{\Xi} \tag{III.8}$$

The bounds on the variables are given by:

$$0 \leq E_k(\xi) \leq \overline{E},\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.9a}$$
$$-\underline{P} \leq P_k(\xi) \leq \overline{P},\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.9b}$$
$$0 \leq F_k(\xi) \leq \overline{P},\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{III.9c}$$

We enforce the fact that P_{t+1} and F_{t+1} are first-stage variables by using the non-anticipativity constraints:

$$P_{t+1}(\xi) = \mathbb{E}\left[P_{t+1}(\xi)\right],\ \xi \in \bar{\Xi} \tag{III.10a}$$
$$F_{t+1}(\xi) = \mathbb{E}\left[F_{t+1}(\xi)\right],\ \xi \in \bar{\Xi} \tag{III.10b}$$

Figure 3 sketches the implementation of the stochastic RH scheme. The problem at time t uses realizations for the uncertain variables that we denote as $L_{\mathcal{N}_t}(\xi)$, $\alpha_{\mathcal{N}_t}(\xi)$, $\pi^f_{\mathcal{N}_t}(\xi)$, and $\pi^e_{\mathcal{N}_t}(\xi)$. We can denote the problem solved at time t as $\mathcal{P}_t(L_{\mathcal{N}_t}(\xi), \alpha_{\mathcal{N}_t}(\xi), \pi^f_{\mathcal{N}_t}(\xi), \pi^e_{\mathcal{N}_t}(\xi), E_t, P_t, \widehat{D}_t)$. For convenience, we simplify the notation to $\mathcal{P}_t(E_t, P_t, \widehat{D}_t)$. From the solution of this problem we obtain the commitments P_{t+1} and F_{t+1}. The actual (true) residual demand depends on the actual load and is given by $d_{t+1} = L_{t+1} - P_{t+1} - \alpha_{t+1}F_{t+1}$, where L_{t+1} and α_{t+1} are actual values.

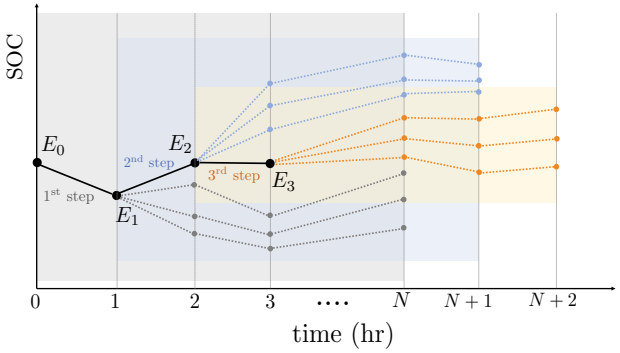

Fig. 3: Sketch of stochastic MPC scheme.

We can now summarize the stochastic MPC scheme run over a month period $\mathcal{M}$ as follows:

- START at $t = 0$ with E_0 and $\widehat{D}_0 = 0$ given. REPEAT for $t \in \mathcal{M}$:

- Solve $\mathcal{P}_t\left(E_t, P_t, \widehat{D}_t\right)$ by using the forecast realizations $\pi_{\mathcal{N}_t}(\xi)$, loads $L_{\mathcal{N}_t}(\xi)$, and FR signals $\alpha_{\mathcal{N}_t}(\xi)$ to obtain commitments P_{t+1} and F_{t+1}. Implement these decisions over $(t, t+1)$ and update SOC $E_{t+1} = E_t - P_{t+1} + \alpha_{t+1}F_{t+1}$ using a random *realized* FR signal $\alpha_{t+1}(\xi)$.
- Using the *actual* realized load L_{t+1} and FR signal α_{t+1} compute the actual net demand $d_{t+1} = L_{t+1} - P_{t+1} + \alpha_{t+1}F_{t+1}$. Update peak carryover demand as $\widehat{D}_{t+1} = \max\{\widehat{D}_t, d_{t+1}\}$.
- Set $t \leftarrow t+1$. If $t = M$ STOP, otherwise RETURN.

We highlight that we impose a lower bound for the peak demand $D_t = \max_{k\in\mathcal{N}_t} d_k$ (given by $\widehat{D}_t$). This arises from carrying over the observed peak until time t in the RH scheme. The need for such a lower bound is observed by noticing that the peak demand over $\mathcal{M}$ can be computed as a recursion $\widehat{D}_{k+1} = \max(\widehat{D}_k, d_{k+1}), k \in \mathcal{M}$ with $\widehat{D}_0 = 0$. Consequently, at time t, the peak demand D_t over the horizon $\mathcal{N}_t$ must be at least $\widehat{D}_t$, and thus $D_t \geq \widehat{D}_t$ and this implies that $\max_{k\in\mathcal{N}_t} d_k \geq \widehat{D}_t$ must hold. The *carryover* constraint also indicates that the peak demand is in fact a state of the system.

C. Deterministic MPC Framework

In traditional deterministic MPC, we solve a single scenario problem at every time t consisting of the most likely realization (i.e., the mean forecast). The mean forecasts for loads, prices, and FR signals are $\hat{L}_{\mathcal{N}_t}, \hat{\alpha}_{\mathcal{N}_t}, \hat{\pi}_{\mathcal{N}_t}$. The optimization problem solved at time t is $\mathcal{P}_t(\hat{L}_{\mathcal{N}_t}, \hat{\alpha}_{\mathcal{N}_t}, \hat{\pi}_{\mathcal{N}_t}, E_t, P_t, \widehat{D}_t)$. When convenient, we use the compact form $\mathcal{P}_t(E_t, P_t, \widehat{D}_t)$. The deterministic RH scheme is:

- START at $t = 0$ with E_0 and $\widehat{D}_0 = 0$ given. REPEAT for $t \in \mathcal{M}$:
- Solve $\mathcal{P}_t(E_t, P_t, \widehat{D}_t)$ by using the forecast prices $\hat{\pi}_{\mathcal{N}_t}$, loads $\hat{L}_{\mathcal{N}_t}$, and FR signals $\hat{\alpha}_{\mathcal{N}_t}$ to obtain commitments P_{t+1} and F_{t+1}. Implement these decisions over $(t, t+1)$ and update the new battery SOC to $E_{t+1} = E_t - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$ using the true FR signal $\alpha_{t+1}(\xi)$.
- Using the true load $L_{t+1}(\xi)$ and FR signal $\alpha_{t+1}(\xi)$ compute the true demand $d_{t+1}(\xi) = L_{t+1} - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$. Update the carryover peak as $\widehat{D}_{t+1}(\xi) = \max\{\widehat{D}_t(\xi), d_{t+1}(\xi)\}$.
- Set $t \leftarrow t+1$. If $t = M$ STOP, otherwise RETURN.

A limitation of deterministic MPC is that it only uses mean forecast information of the loads and prices and neglects their uncertainty when making commitments. This can introduce significant inefficiencies that we seek to quantify.

IV. Benchmarking Stochastic MPC

We compare the performance of the proposed stochastic MPC framework against that of deterministic MPC and of perfect information MPC. We also describe a methodology to generate realizations from historical data.

A. Performance Metrics

To quantify the impact of forecast errors on the performance of *deterministic MPC*, we use the commitment trajectories $P_{\mathcal{M}}$ and $F_{\mathcal{M}}$ collected over the month and different *realizations* of the loads $L_{\mathcal{M}}(\xi)$, prices $\pi_{\mathcal{M}}(\xi)$, and FR signals $\alpha_{\mathcal{M}}(\xi)$, to compute the actual monthly net profit. Since the profit is usually negative, we find it more convenient to use the net cost (the negative net profit). The monthly net cost is a random variable with realizations:

$$\Phi^{det}(\xi) = \sum_{t\in\mathcal{M}} (\pi_t^e(\xi)(\alpha_t(\xi)F_t - P_t) - \pi_t^f(\xi)F_t) + \pi^D \max_{t\in\mathcal{M}} d_t(\xi), \quad \xi \in \Xi \tag{IV.11}$$

where $d_t(\xi) = L_k(\xi) - P_t - \alpha_t(\xi)F_t$. We use the commitment policies $P_{\mathcal{M}}$, $F_{\mathcal{M}}$ obtained under the *stochastic MPC scheme* and the *actual* realizations of the loads, prices, and FR signals to compute the accumulated profit for every realization of the form (IV.11) as well (i.e., we use the same realizations but different commitments). The net profit under the stochastic policy is denoted as $\Phi^{stoch}(\xi)$. The profit of stochastic and deterministic MPC is computed for *every realization* of the uncertain variables $\xi \in \Xi$. This accounts for the fact that, in practice, the stochastic MPC problem cannot be solved for every possible realization of uncertainty (i.e., we often have that $|\bar{\Xi}| \ll |\Xi|$). To evaluate the *ideal* performance of MPC (we call this *perfect information MPC*), we compute commitments for *every realization* $\xi \in \Xi$ of the loads, prices, and FR signals. The monthly cost for this scheme is given by:

$$\Phi^{perf}(\xi) = \sum_{t\in\mathcal{M}} (\pi_t^e(\xi)(\alpha_t(\xi)F_t(\xi) - P_t(\xi)) - \pi_t^f(\xi)F_t(\xi)) + \pi^D \max_{t\in\mathcal{M}} d_t(\xi), \quad \xi \in \Xi. \tag{IV.12}$$

In this case the decisions are a function of the realizations because we assume that we can compute policies by perfectly anticipating uncertainty. These policies, however, are not implementable in practice. One is often interested in evaluating the expected value of the stochastic solution (which is defined as $\mathbb{E}[\Phi^{det}(\xi) - \Phi^{stoch}(\xi)]$ and denoted as VSS). It is also often convenient to assess the value of having perfect information, which we define as $\mathbb{E}[\Phi^{det}(\xi) - \Phi^{perf}(\xi)]$ and which is denoted as VPI [28]. The first metric evaluates the benefit of stochastic over deterministic MPC while the second metric evaluates the impact of not being able to predict perfectly.

Of particular interest in our context is a metric that we call the *expected value of the battery*. To compute this value, we evaluate performance under the assumption that no battery has been installed (FR profit and peak shaving capabilities are not available). The total cost thus reduces to the demand charge:

$$\Phi^{nobat}(\xi) = \pi^D \max_{t\in\mathcal{M}} L_t(\xi). \tag{IV.13}$$

The ideal expected value of the battery is defined as $\mathbb{E}[\Phi^{nobat}(\xi) - \Phi^{perf}(\xi)]$. The value of the battery under stochastic MPC is $\mathbb{E}[\Phi^{nobat}(\xi) - \Phi^{stoch}(\xi)]$ and under deterministic MPC is $\mathbb{E}[\Phi^{nobat}(\xi) - \Phi^{det}(\xi)]$. These metrics are used to isolate the effect of the battery, which is usually significantly smaller than the total cost. Another important metric is the *payback period*, which is the ratio of the investment cost and the value of the battery (the expected savings).

B. Scenario Generation from Limited Historical Data

One often is interested in using historical data to model the uncertainty of random variables and generate scenarios

(realizations). In our context, we seek to generate scenarios by sampling an empirical multivariable Gaussian distribution that captures short-term and long-term time correlations. Unfortunately, historical data is often insufficient to obtain a sample covariance. For instance, if we seek to capture daily and weekly correlations, we can only obtain 52 samples from a full year of data and the sample covariance has a dimension of $24\times 7 = 168$. Consequently, the covariance matrix does not have full rank. To address this issue, we use the Ledoit-Wolf estimator [29]. We consider the norm of a $p\times p$ matrix to be a scaled version of the Frobenius norm $||A||^2 = \text{Tr}(AA^T)/p$ and the associated inner product to be $\langle A_1, A_2\rangle = \text{Tr}(A_1A_2^T)/p$. The advantage of dividing the trace by the dimension p is that the norm of the identity matrix is one. Assume that we have n independent and identically distributed (i.i.d.) observations of a p-dimensional random variable X with mean 0 and covariance Σ_n, where $p \gg n$ (in our case $n = 52$ and $p = 168$). Let X_n denote a $p\times n$ matrix of the n realizations of the x vector. Our goal is to find the linear combination $\Sigma_n^* = \rho_1 I_n + \rho_2 S_n$ of the identity matrix I_n and the sample covariance matrix $S_n = X_n X_n^T/n$. Here, $\rho_1, \rho_2 \geq 0$ and we define the expected quadratic loss $\mathbb{E}[||\Sigma_n^* - \Sigma_n||]$ where Σ_n is the true (unobservable) covariance matrix. We also define the following scalar quantities, which play a central role in the analysis: $\mu_n = \langle \Sigma_n, I_n\rangle, \alpha_n^2 = ||\Sigma_n - \mu_n I_n||^2, \beta_n^2 = \mathbb{E}[||S_n - \Sigma_n||^2]$, and $\delta_n^2 = \mathbb{E}[||S_n - \mu_n I_n||^2]$. Using these definitions, it can be shown that $\mathbb{E}[S_n] = \Sigma_n$ and that $\alpha_n^2 + \beta_n^2 = \delta_n^2$ hold. We obtain the optimal weighted combination $\Sigma_n^* = \rho_1 I_n + \rho_2 S_n$ by solving the quadratic program:

$$\min_{\rho_1,\rho_2} \mathbb{E}[||\Sigma_n^* - \Sigma_n||^2] \text{ s.t. } \Sigma_n^* = \rho_1 I_n + \rho_2 S_n. \qquad \text{(IV.14)}$$

The solution of this quadratic program gives the optimal weights $\rho_1 = (\beta_n^2/\delta_n^2)\mu_n$ and $\rho_2 = (\alpha_n^2/\delta_n^2)$. The optimal covariance matrix Σ_n^* (obtained by substituting the weights into the constraint of (IV.14)), however, is not a *bona fide* estimator because it requires hindsight knowledge of the scalar functions of the true (but unobservable) covariance Σ_n ($\mu_n, \alpha_n^2, \beta_n^2$, and δ_n^2). This problem is addressed by obtaining an alternative estimator for Σ_n^*. To do this, we note that computing Σ_n^* does not require knowledge of the whole matrix Σ_n directly, but only of the scalars $\mu_n, \alpha_n^2, \beta_n^2$, and δ_n^2. We thus define estimators of $\mu_n, \alpha_n^2, \beta_n^2$ and δ_n^2 that converge to the true ones for large n. We define $m_n = \langle S_n, I_n\rangle$, $d_n^2 = ||S_n - m_n I_n||^2$, $\bar{b}_n^2 = \frac{1}{n^2}\sum_{k=1}^n \left|\left|x_k x_k^T - S_n\right|\right|^2$, $b_n^2 = \min(\bar{b}_n^2, d_n^2)$, and $a_n^2 = d_n^2 - b_n^2$. It is possible to prove that $\mathbb{E}[m_n] = \mu_n$, and that $m_n - \mu_n$, $m_n^2 - \mu_n^2$, $d_n^2 - \delta_n^2$, $\bar{b}_n^2 - \beta_n^2$ and $b_n^2 - \beta_n^2$, $a_n^2 - \alpha_n^2$ all converge to 0 as n increases. This yields the *bona fide* estimator $\tilde{\Sigma}_n = (b_n^2/d_n^2)m_n I_n + (a_n^2/d_n^2)S_n$.

C. Computational Experiments

We consider a stationary battery that has a capacity 0.5 MWh, rated power of 1 MW, and we assume a ramping limit of 0.5 MW/hr. We use historical data for energy prices and FR prices and signals from PJM. Historical load data from a typical university campus is used. We consider cases with MPC prediction horizon lengths of $N = 24$ (1-day horizon) and $N = 168$ (7-day horizon). The horizon lengths that we consider are short compared to the demand charge period (one month). Consequently, we also consider strategies to estimate the *discounting factor* σ_t. In the first case we consider a constant factor of $\sigma_t = 1$, which means that no discounting of the long-term demand charge rate is used. In the second case we use a discounting $\sigma_t = M/N$. We perform MPC simulations for a period of one month. Load realizations for this period are generated from one year of historical data using the Ledoit-Wolf covariance estimator. The one-year load (obtained from a large university campus) and the generated realizations are shown in Figure 4. The prediction errors for the realizations and for the mean profile remain in the range of $[-20\%, 10\%]$. This indicates that the proposed scenario generation approach provides reasonable estimates and preserves the correlation structure of the load.

The energy price profile obtained from PJM along with the generated Ledoit-Wolf realizations are shown in Figure 5. We observe that the profile also exhibits strong periodicity. The FR price data obtained from PJM is shown Figure 6a. We observe that, compared to the energy price, the FR price profile exhibits less correlation structure. We have found that the FR price follows a log-normal distribution and has high variance. In this case, we use the empirical log-normal distribution to generate FR price realizations. We note that the significant conservatism is expected with respect to this variable. To generate realizations of the hourly average FR dispatch signal, we note that the signal approximately follows a Gaussian distribution, which indicates that there is small temporal correlation. Figure 7 shows the realizations generated by sampling from the empirical Gaussian.

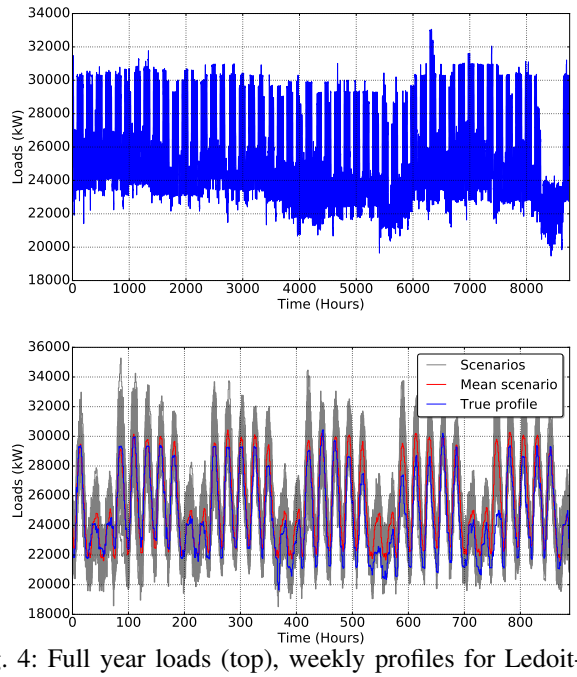

Fig. 4: Full year loads (top), weekly profiles for Ledoit-Wolf (middle), and generated realizations (bottom).

For the simulation of deterministic MPC, we use the mean (most likely) profiles of the random variables as forecasts. We evaluate the policy in $|\Xi| = 1000$ realizations to evaluate the distribution of the monthly cost. In the stochastic MPC

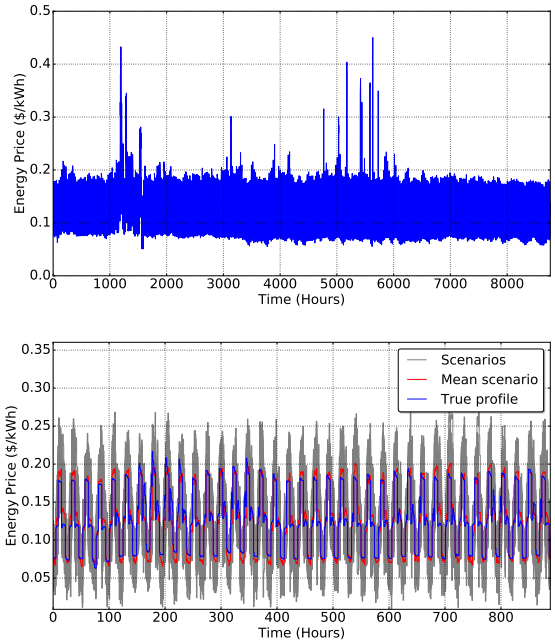

Fig. 5: Full year energy prices (top), weekly profiles for Ledoit-Wolf (middle), and generated realizations (bottom).

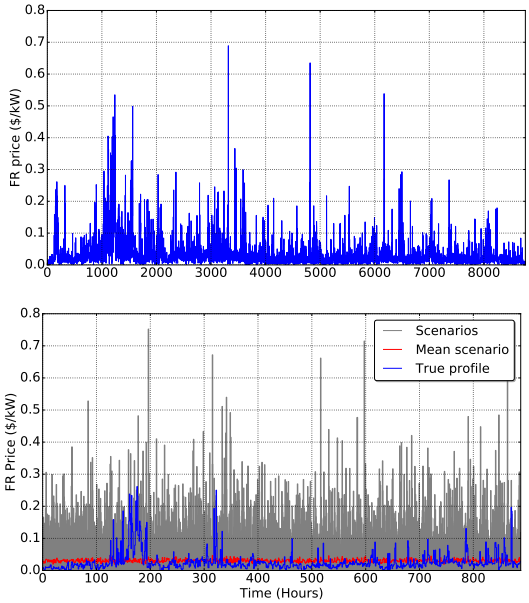

Fig. 6: Historical FR prices and realizations

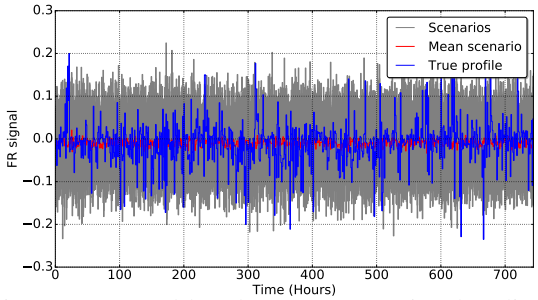

Fig. 7: Generated hourly average FR signal realizations

scheme, we solve a two stochastic stochastic program containing $|\bar{\Xi}| = 50$ realizations to compute the planning policy and evaluate the policy using the same 1000 realizations. To simulate the perfect information MPC scheme, compute policies over individual realization [28]. All optimization problems in the MPC simulations were implemented in the modeling language `JuMP` and solved using `Gurobi` (version 6.5.0).

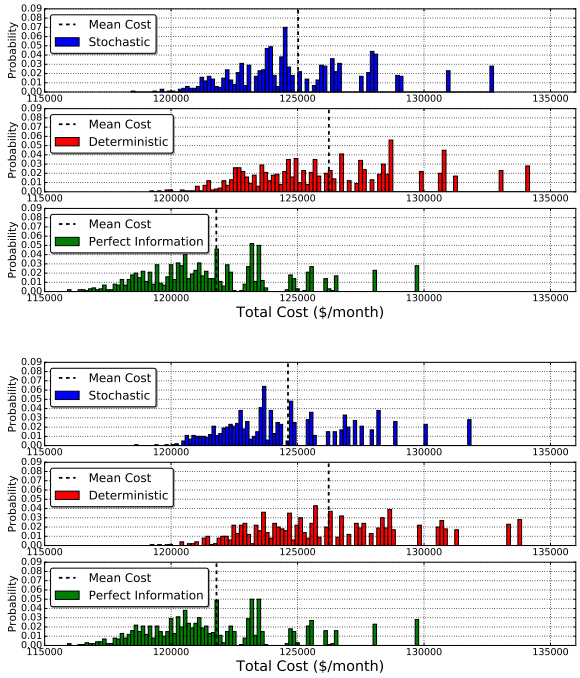

Fig. 8: Distribution of monthly costs for 7-day horizon MPC with undiscounted (top) and discounted demand (bottom).

1) Total Costs and Peak Demands: In Figure 8, we compare the distribution of the total costs (negative net profits) for the 1000 realizations evaluated. The *tail costs* in the deterministic policy are significantly more pronounced, indicating that a larger number of realizations experience high costs. The stochastic MPC scheme mitigates these extreme costs and also has a *mean cost shifted towards the left*. The cost distribution for perfect information MPC is clearly superior but this approach is not implementable. From Table I we observe that average total costs for all MPC schemes decrease as we increase the prediction horizon. This indicates that extending the horizon is beneficial but using a longer horizon does not seem necessary. We can also see that the costs improve when the demand charge is discounted. This is because an *undiscounted demand charge will force the MPC system to operate more conservatively*.

2) FR Commitments and SOC Policies: The mean of the FR participation levels for the stochastic, deterministic, and perfect information MPC schemes are approximately 720, 740, and 780 kW, respectively. We can thus see that the deterministic policy is slightly more aggressive in allocating FR capacities. This, however, comes at the expense of higher peak demand. This highlights that the stochastic policy can better manage stored energy to protect against high peak demands. The perfect information policy can anticipate demand charges and thus be more aggressive in participating in the FR market. Figure 9 compares the *dynamics of the SOC* (for

TABLE I: Expected costs for MPC schemes. TC= Total cost, DC=Demand charge, FR=frequency regulation, and E=Energy.

Horizon	**7-day**					
σ_t	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$
Cost Item ($/month)	**Perfect Information**		**Stochastic**		**Deterministic**	
TC	121,788.9	121,794.4	125,022.5	124,621.7	126,237.8	126,229.9
DC	137,077.0	137,090.5	138,572.9	139,153.3	140,865.7	140,920.5
FR	-14,812.9	-14,820.8	-13,073.7	-14,049.7	-14,116.8	-14,180.9
E	-475.1	-475.2	-476.6	-481.9	-511.1	-509.7
Horizon	**1-day**					
σ_t	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$
Cost Item ($/month)	**Perfect Information**		**Stochastic**		**Deterministic**	
TC	121,796.0	122,234.8	125,570.3	125,825.4	126,322.9	125,977.1
DC	136,983.3	137,537.5	138,427.6	141,113.9	140,611.2	141,220.2
FR	-14,715.6	-14,827.8	-12,371.3	-14,811.7	-13,775.2	-14,761.2
E	-471.7	-474.9	-485.9	-476.8	-513.0	-481.8

TABLE II: Value of stochastic solution (VSS) and value of perfect information (VPI).

		7-day		**1-day**	
σ_t		**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$
VSS	**($/month)**	1,215.3	1,608.1	752.6	151.7
	(Approx. $/year)	14,583.6	19,297.2	9,031.2	1,820.4
VPI	**($/month)**	4,448.9	4,435.5	4,526.9	3,742.3
	(Approx. $/year)	53,386.8	53,226.0	54,322.8	44,907.6

TABLE III: Value of the battery under different MPC schemes.

	Without Battery	With Battery			
		Perfect Information MPC	Stochastic MPC	Deterministic MPC	Value of Stochastic MPC
Expected Total Cost ($/month)	138,224	121,794	124,622	126,230	1,608
Expected Savings ($/month)	-	16,430	13,602	11,994	1,608

a random scenario) for the case study with a 7-day horizon and with a discounted demand charge cost. We observe that the deterministic policy exhibits abrupt changes in SOC more frequently. This reflects the fact that the deterministic policy overcommits in FR capacity. The SOC profiles of stochastic and perfect information MPC are similar and smoother.

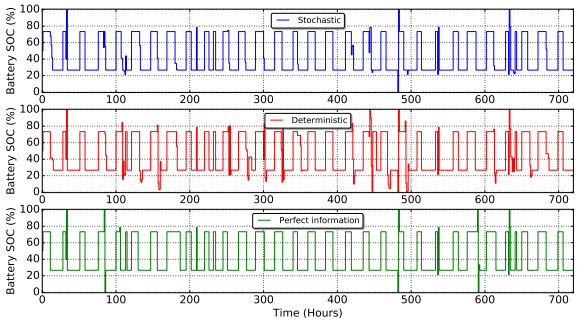

Fig. 9: SOC policy in stochastic, deterministic and perfect information MPC (7-day horizon, discounted demand charge).

3) Expected Costs, Value of the Battery, and Payback: Table I reports the total cost and its components for all the case studies. We observe a general trend of higher expected total costs with shorter horizon lengths and without discounted factors. Moreover, we observe that stochastic MPC outperforms deterministic MPC. In Table II we see that VSS can reach levels of $1,600 per month while VPI can reach up to $4,700 per month. Both of these quantities are small compared to the total cost (which reaches levels of $120,000 per month). This is because the load of the university campus (and thus the demand charge) is significantly larger than the battery capacity. As a result, VSS and VPI are not particularly informative metrics on the added value of MPC.

Table III summarizes the results for the *value of the battery* under different schemes for a 7-day horizon and discounted demand charge. The expected savings obtained by installing the battery and operating it under perfect information (*ideal value of the battery*) reach $16,430 per month (approximately $197,160 by extrapolating over a year). This reveals that the battery is indeed highly valuable. From Table II we observe that FR revenue reaches levels of $14,000 per month (approximately $168,000 per year). From Table III we observe that stochastic MPC can recover *83%* of the ideal value of the battery while deterministic MPC can only recover 73%. Assuming an investment cost of 1.0 M$, the payback period operated under perfect information MPC is 5.0 years, under stochastic MPC is 6.1 years, and under deterministic MPC is 6.94 years. Stochastic MPC achieves a relative improvement in the payback period achieved under deterministic MPC of 12.1%. By comparing the payback period of perfect information MPC and deterministic MPC, we can also conclude that the use of perfect information achieves a relative improvement of 27.9%. We thus see that, by improving the battery management strategy, one can significantly influence its value.

V. Conclusions and Future Work

We present a stochastic MPC framework for evaluating the performance of stationary battery systems that provides FR capacity and mitigates demand charges. We use a Ledoit-Wolf

covariance estimation to generate load and price realizations from limited historical data. Our simulations demonstrate that stochastic MPC can recover 83% of the value of the battery while deterministic MPC can only recover 73% of such value. This indicates that important reductions in the battery payback period can be achieved. Extensions to this work include capturing price and load correlations and high-frequency FR dispatch signals. Such multi-scale implementations will require finer time discretizations and will require alternative algorithms such as stochastic dual dynamic programming.

REFERENCES

[1] D. D. Rastler, "Electricity energy storage technology options: a white paper primer on applications, costs and benefits," *Epri*, p. 170, 2010.

[2] A. Oudalov, D. Chartouni, C. Ohler, and G. Linhofer, "Value Analysis of Battery Energy Storage Applications in Power Systems," in *2006 IEEE PES Power Systems Conference and Exposition*, 2006, pp. 2206–2211.

[3] J. Donadee and M. Ilić, "Stochastic Co - Optimization of Charging and Frequency Regulation by Electric Vehicles," *IEEE*, 2012.

[4] J. Donadee and M. D. Ilić, "Stochastic optimization of grid to vehicle frequency regulation capacity bids," *IEEE Transactions on Smart Grid*, vol. 5, no. 2, pp. 1061–1069, 2014.

[5] K. Kim, F. Yang, V. M. Zavala, and A. A. Chien, "Data Centers as Dispatchable Loads to Harness Stranded Power," *IEEE Transactions on Sustainable Energy*, vol. 3029, no. c, pp. 1–1, 2016.

[6] P. Mercier, R. Cherkaoui, and A. Oudalov, "Optimizing a battery energy storage system for frequency control application in an isolated power system," *IEEE Transactions on Power Systems*, vol. 24, no. 3, pp. 1469–1477, 2009.

[7] G. He, Q. Chen, C. Kang, P. Pinson, and Q. Xia, "Optimal Bidding Strategy of Battery Storage in Power Markets Considering Performance-Based Regulation and Battery Cycle Life," *IEEE Transactions on Smart Grid*, vol. 7, no. 5, pp. 2359–2367, 2016.

[8] H. Mohsenian-Rad, "Optimal bidding, scheduling, and deployment of battery systems in California day-ahead energy market," *IEEE Transactions on Power Systems*, vol. 31, no. 1, pp. 442–453, 2016.

[9] D. Kottick, M. Blau, and D. Edelstein, "Battery Energy Storage for Frequency Regulation in an Island Power System," *IEEE Transactions on Energy Conversion*, vol. 8, no. 3, pp. 455–459, 1993.

[10] S. K. S. Singh, S. K. S. Singh, S. Chanana, and Y. Singh, "Frequency Regulation of an Isolated Hybrid Power System with Battery Energy Storage System," *2014 Power and Energy Systems: Towards Sustainable Energy (PESTSE 2014)*, pp. 2–7, 2014.

[11] J. M. Foster and M. C. Caramanis, "Energy reserves and clearing in stochastic power markets: The case of plug-in-hybrid electric vehicle battery charging," *Proceedings of the IEEE Conference on Decision and Control*, pp. 1037–1044, 2010.

[12] S. Han, S. Han, and K. Sezaki, "Development of an optimal vehicle-to-grid aggregator for frequency regulation," *IEEE Transactions on Smart Grid*, vol. 1, no. 1, pp. 65–72, 2010.

[13] J. Lin, K. C. Leung, and V. O. K. Li, "Optimal Scheduling with Vehicle-to-Grid Regulation Service," *IEEE Internet of Things Journal*, vol. 1, no. 6, pp. 556–569, 2014.

[14] C. D. White and K. M. Zhang, "Using vehicle-to-grid technology for frequency regulation and peak-load reduction," *Journal of Power Sources*, vol. 196, no. 8, pp. 3972–3980, 2011.

[15] R. T. de Salis, A. Clarke, Z. Wang, J. Moyne, and D. M. Tilbury, "Energy storage control for peak shaving in a single building," *2014 IEEE PES General Meeting — Conference {&} Exposition*, 2014.

[16] M. P. Johnson, A. Bar-Noy, O. Liu, and Y. Feng, "Energy peak shaving with local storage," in *Sustainable Computing: Informatics and Systems*, vol. 1, no. 3, 2011, pp. 177–188.

[17] J. Leadbetter and L. Swan, "Battery storage system for residential electricity peak demand shaving," *Energy and Buildings*, vol. 55, pp. 685–692, 2012.

[18] A. Rahimi, S. Member, M. Zarghami, M. Vaziri, S. Member, S. Vadhva, and B. Mwh, "A Simple and Effective Approach for Peak Load Shaving Using Battery Storage Systems," in *North American Power Symposium (NAPS), 2013*, 2013, pp. 1–5.

[19] C. Lu, H. Xu, X. Pan, and J. Song, "Optimal sizing and control of battery energy storage system for peak load shaving," *Energies*, vol. 7, no. 12, pp. 8396–8410, 2014.

[20] X. Dong, G. Bao, Z. Lu, Z. Yuan, and C. Lu, "Optimal battery energy storage system charge scheduling for peak shaving application considering battery lifetime," in *Lecture Notes in Electrical Engineering*, vol. 133 LNEE, no. VOL. 2, 2011, pp. 211–218.

[21] J. Alt, M. Anderson, and R. Jungst, "Assessment of utility side cost savings from battery energy storage," *IEEE Transactions on Power Systems*, vol. 12, no. 3, pp. 1112–1120, 1997.

[22] L. Sigrist, E. Lobato, and L. Rouco, "Energy storage systems providing primary reserve and peak shaving in small isolated power systems: An economic assessment," *International Journal of Electrical Power and Energy Systems*, vol. 53, no. 1, pp. 675–683, 2013.

[23] A. Lucas and S. Chondrogiannis, "Smart grid energy storage controller for frequency regulation and peak shaving, using a vanadium redox flow battery," *International Journal of Electrical Power and Energy Systems*, vol. 80, pp. 26–36, 2016.

[24] R. Sebastián, "Application of a battery energy storage for frequency regulation and peak shaving in a wind diesel power system," *IET Generation, Transmission & Distribution*, vol. 10, no. 3, pp. 764–770, 2016.

[25] S. I. Vagropoulos and A. G. Bakirtzis, "Optimal bidding strategy for electric vehicle aggregators in electricity markets," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4031–4041, 2013.

[26] S. J. Moura, D. S. Callaway, H. K. Fathy, and J. L. Stein, "Tradeoffs between battery energy capacity and stochastic optimal power management in plug-in hybrid electric vehicles," *Journal of Power Sources*, vol. 195, no. 9, pp. 2979–2988, 2010.

[27] Y. Wang, B. Wang, C. C. Chu, H. Pota, and R. Gadh, "Energy management for a commercial building microgrid with stationary and mobile battery storage," *Energy and Buildings*, vol. 116, pp. 141–150, 2016.

[28] J. R. Birge and F. Louveaux, *Introduction to stochastic programming*. Springer Science & Business Media, 2011.

[29] O. Ledoit and M. Wolf, "A well-conditioned estimator for large-dimensional covariance matrices," *Journal of Multivariate Analysis*, vol. 88, no. 2, pp. 365–411, 2004.

Ranjeet Kumar received his Bachelor of Technology with Honors from the Indian Institute of Technology Bombay, India in 2014. He is currently a PhD student in the Department of Chemical and Biological Engineering at University of Wisconsin-Madison. His research interests include control, optimization, and electricity markets.

Mohammad N. ElBsat received his B.S. degree in Electrical Engineering from the Lebanese University in 2006 in Hadath, Lebanon. He, then, received his M.S. and Ph.D. also in Electrical Engineering, from Marquette University in 2008 and 2012, respectively. He is currently a Principal Engineer at Johnson Controls and his interests include optimization and control and estimation of nonlinear systems.

Matthew Ellis received his Ph.D. in Chemical and Biomolecular Engineering from the University of California, Los Angeles in 2015. He is currently a Principal Engineer at Johnson Controls. His interests include model predictive control, optimization, and cost/profit-driven control and automation algorithms for real-time applications.

Kirk H. Drees is Director of Energy Services, Research, and Modeling for the Building Efficiency business of Johnson Controls. He holds 40 U.S. patents and has 35 years of experience spanning research, advanced development and productization. Kirk is a registered professional engineer with bachelors and masters degrees in engineering from Iowa State University and Purdue University.

Victor M. Zavala is the Richard H. Soit Assistant Professor in the Department of Chemical and Biological Engineering at the University of Wisconsin-Madison. He holds a B.Sc. degree from Universidad Iberoamericana and a Ph.D. degree from Carnegie Mellon University, both in chemical engineering. He is an associate editor for the Journal of Process Control and a technical editor of Mathematical Programming Computation. His research interests are in the areas of energy systems, high-performance computing, stochastic optimization, and predictive control.