Digital Object Identifier: 10.1109/TPWRS.2017.2704524

Demand Response of Ancillary Service from Industrial Loads Coordinated with Energy Storage

Xiao Zhang, *Student Member, IEEE*, Gabriela Hug, *Senior Member, IEEE*, J. Zico Kolter, Iiro Harjunkoski

***Abstract*—As one of the featured initiatives in smart grids, demand response is enabling active participation of electricity consumers in the supply/demand balancing process, thereby enhancing the power system's operational flexibility in a cost-effective way. Industrial load plays an important role in demand response because of its intense power consumption, already existing advanced monitoring and control infrastructure, and its strong economic incentive due to the high energy costs. As typical industrial loads, cement plants are able to quickly adjust their power consumption rate by switching on/off the crushers. However, in the cement plant as well as other industrial loads, switching on/off the loading units only achieves discrete power changes, which restricts the load from offering valuable ancillary services such as regulation and load following, as continuous power changes are required for these services. In this paper, we overcome this restriction of poor granularity by proposing methods that enable these loads to provide regulation or load following with the support of an on-site energy storage system.**

***Index Terms*—Demand response, industrial load, energy storage system (ESS), model predictive control, regulation, load following.**

I. Introduction

Countries and states all over the world are making efforts to secure a sustainable energy future. To achieve this, renewable generation resources such as wind turbines and solar panels are being deployed more and more widely. For example, in 2015, the U.S. saw the installation of 7.3 GW solar power capacity which contributed 29.5 percent of all newly-installed electric generation capacity and beat out the increase of natural gas capacity for the first time [1]; more than 3 GW of offshore wind power has been connected to the European grid in 2015 - twice as much as in 2014 [2]. However, the power output from these renewable resources is intermittent and uncertain which imposes challenges on power system operation.

The integration of renewable generation requires large amounts of balancing resources to enhance the operational flexibility of the grid. Traditionally, the power system relies on generators to provide such flexibility. However, it is necessary to seek for additional balancing resources to strengthen the balancing capability of the grid. Meanwhile, demand response has demonstrated potentials to enhance the power system's operational flexibility in a cost-effective way [3]–[6]. Demand response can be provided by residential, commercial, or industrial loads, and its benefits have been discussed in many studies, e.g. demand response provided by electric vehicles [7]–[9], residential areas [10]–[12], buildings [13]–[16], data centers [17], aluminum smelters [18], [19], air separation units [20], and steel plants [21]–[23].

Within the realm of demand response, industrial loads have the following advantages [24]–[26]: most industrial loads are already equipped with control, measurement, and communication infrastructures which are necessary for demand response participation; many industrial loads are able to adjust their energy consumption profile and provide power change in a large, fast, and accurate way; industrial loads are also willing to increase their operational complexity in order to participate in demand response programs, which will reduce their electricity cost - usually a significant portion of total cost of operations for many industrial plants. Industrial loads that are ideal candidates for demand response include aluminum smelting pots, furnaces, fans, freezers, pumps, mills, crushers, etc. These industrial loads are not only participating in the energy markets through programs like load shifting, they are also actively providing ancillary services [14], [27], [28] such as spinning reserve, load following, and regulation. In this project, we study approaches to enable ancillary services provision from industrial loads to encourage the industrial loads to actively participate in the power system.

Among the ancillary services, regulation and load following require a very fast response of power change, both up and down, and are very valuable in the markets. Meanwhile, the electricity market has critical tests for market participants to be qualified as resources for regulation or load following. Failing to consistently follow the setpoint instructions sent from the power system control center may lead to penalty and even disqualification. For example, the rules of the Midcontinent Independent System Operator (MISO) state that if the mismatch exceeds the tolerance band for 3 consecutive 5-min intervals, a Failure to Follow Dispatch Flag is set, and the resource needs to pay back the regulation awards plus a penalty to MISO [29], [30]; in PJM Interconnection (PJM), the regulating resources that have not met performance thresholds over a specified time period will be disqualified, where the performance score takes account of the delay, correlation, and precision scores over the last 100 hours. The score calculation procedures can be found in [31].

Lots of industrial loads are able to provide very fast changes in power consumption in both directions. For example, the crushers or mills in the cement industry can be switched on and off very rapidly [32]. However, most of these industrial loads can only provide power changes in a discrete manner, e.g. the power change is several MWs at a time. This poor granularity restricts them from offering regulation or load following, which requires a continuous change of power. Therefore, these demand response resources with fast power changing ability have not been utilized to their fullest potential.

This project was supported by ABB.

X. Zhang is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, USA (e-mail: xiaozhang@cmu.edu).

G. Hug is with the Department of Information Technology and Electrical Engineering, ETH Zurich, Zurich 8092, Switzerland.

J. Z. Kolter is with the Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

I. Harjunkoski is with ABB Corporate Research, Ladenburg, Baden-Wurttemberg 68526, Germany.

To overcome the restriction of poor granularity and deal with the discreteness in its power change, we proposed a coordination framework in which the industrial load provides regulation or load following with the support from an on-site energy storage system (ESS) [33]: the industrial load provides a large but discrete power change which constitutes the main contribution of the service, while the energy storage provides a fine and continuous power change which ensures that the combination of the two accurately follows the desired power signal. These two parts are coordinated by a model predictive control (MPC) approach which utilizes the prediction of the upcoming signals from the power system. As demonstrated in [33], the combination of the industrial loads and energy storage is able to accurately follow the regulation or load following command in a very wide range. That previous work only focuses on the hourly operation and assumes the regulation capacity is given. In this paper, we extend our previous work to study the day-ahead scheduling and provide the tool to determine the optimal quantity of regulation provision for each hour.

The key contributions of this work are the proposed methods for enabling the loads to fully achieve their potentials as demand response resources with the support of an on-site energy storage system. We demonstrate through case studies that the cooperation of the industrial machines and the energy storage system can produce a combined contribution that is greater than the sum of their separate effects. The remainder of the paper is organized as follows. The considered problem is introduced in Section II. We will then describe in Section III how the regulation can be provided by the industrial machines with the support of an energy storage system. The real-time coordination method is based on an MPC for hourly operation. Next, we investigate in Section III how to optimally decide its regulation contributions for the next day, where the day-ahead optimal scheduling method is presented. A case study is discussed in Section V to demonstrate the effectiveness of the proposed approaches. Section VI concludes the paper.

II. Problem Statement

A variety of industrial loads can be switched on and off very rapidly, which enables them to change their power consumption rate fast and frequently, e.g. the crushers in the cement crushing industry [32] and the mills in the thermo-mechanical pulp and paper industry [34]. In this paper, we investigate and provide methods to utilize these industrial loads for the provision of regulation service. Note that the proposed method can also be employed to enable load following.

PJM, the largest competitive wholesale electricity market in the U.S., distinguishes between two different regulation signals: RegD and RegA. RegD is designed for fast reacting resources and we assume that the manufacturing plant participates as such fast acting resource, which is true for industrial demand response resources like aluminum smelters. The regulation signal is in per unit value, i.e. it ranges between -1.0 and 1.0. Suppose the load has committed to provide R MW regulation with a baseline power of B MW, then the regulation command is the RegD signal scaled by R plus B, i.e. the targeted power consumption rate ranges between $B-R$ MW and $B+R$ MW. Hence, the problem to be solved is twofold. First, given a particular B and R, how can the manufacturing plant ensure that it closely follows the given signal? Second, the plant needs to determine the optimal B and R values, such that it can fulfill its regulation commitment at any point in time while not negatively impacting its production. Though previous work has demonstrated the advancement of short-term forecasting in power systems [35]–[37], it is still impossible to accurately forecast the regulation signal over several minutes. However, forecasting its trend with reasonable accuracy over a horizon less than 1 minute is possible, e.g. by using autoregressive-moving-average (ARMA) models. As demonstrated later, this prediction is good enough to coordinate the industrial loads and the energy storage, where the energy storage provides some buffer for prediction errors.

As mentioned, the loading units (machines) are switched on/off to follow the regulation command with the support of an on-site energy storage system. For simplicity, we assume that there are M machines which can be switched on and off rapidly, and each machine has the same power consumption rate of ρ MW. Note that in practice the power consumption rates for different machines may not be the same, yet the proposed method can be easily extended to consider this deviation from our assumption. It has been demonstrated that stand-alone storage has significant potentials to support the power system operation [38], [39], whereas in our method the storage helps the industrial load to overcome the restriction of poor granularity. We assume that the storage has a maximum energy capacity of E_s MWh and its charging power is bounded by $-P_s$ and P_s MW. To simplify the problem, we further assume that there is no energy loss associated with the charging and discharging processes. Note that the energy loss can be considered easily by extending the formulations.

In most electricity markets, the market participants bid for their market share for each hour in the following day, and the market share together with the final market price are settled after the bidding process is completed. Hence, in each operating hour, the regulation capacity R and the energy baseline B are pre-determined by the market. The demand response provider is obliged to follow the resulting regulation signal, otherwise, it will be penalized according to market rules. This leads to the proposed **Real-time MPC Coordination** of the industrial plant, as presented in Section III. With the settled R and B, the industrial load and the energy storage are coordinated by the MPC method to optimally follow the regulation command with consideration of the machine switching cost, the regulation command violation, and the storage energy level.

Before the actual demand response participation, the industrial plant needs to determine the optimal R and B for each operating hour to maximize its daily revenue. Since market bidding is not the focus in this paper, we assume the energy price and regulation price are known. This leads to the proposed **Day-ahead Optimal Scheduling** of the industrial plant, in which R and B are determined for every hour of the next day, as presented in Section IV. The Day-ahead Optimal Scheduling aims to decide the R and B for each hour in the next day, while the Real-time MPC Coordination aims to control the devices to optimally follow the regulation commands at real-time for each operating hour with a pre-settled R and B. The scheduling approach takes into account

the revenues from market participation, the cost of regulation provision, as well as the coupling of the crushing process with other processing stages within the industrial plant. Note that the proposed optimization model can be extended to consider the market bidding problem, e.g. through stochastic programming with possible price curves as scenarios [19].

As just mentioned, the coupling between processes needs to be considered in scheduling industrial loads. Usually, there are multiple processing stages in an industrial plant, and successive stages interact with each other through the generation and consumption of intermediate products. Taking the cement plant in Fig. 1 as an example, the crushing machines belong to the first stage which breaks the raw material (e.g. limestone, clay) into finer particles. In the second stage, the kiln heats up these finer particles to a higher temperature for further processing in its following stages. In other words, the crushing machines generate the finer particles and the kiln consumes these finer particles, therefore these two stages are coupled through the intermediate product. In scheduling the cement plant, we want to keep the kiln operating at a constant consumption rate, because it is very expensive to turn on/off the kiln due to its large thermal capacity. As the intermediate product generation rate is proportional to the energy consumption rate of the crushing machines, i.e. the energy baseline B, we need to wisely decide the values of B for each hour in the operating day so that the following processing stages are not interrupted.

III. Real-time MPC Coordination

The real-time coordination framework for hourly operation is illustrated in Fig. 2. At each step t, based on historical regulation commands, the predictor outputs the regulation prediction for the next L steps; then the optimal controller optimizes over the number of active machines x_i and the storage charging power y_i for each time step $i = t, \cdots, t+L$ in the MPC horizon, based on the regulation prediction and previous operation records of the machines; after obtaining the optimization results, only the control decision for the current time step t is applied to the industrial load and the energy storage. Then, the horizon is shifted forward by one time step and the optimization is carried out anew.

A. Prediction

The prediction of the regulation signal is achieved using an ARMA model. We have trained different ARMA models by the Python Time Series Analysis package. For different training data sets [40], the ARMA(2,1) model achieved the best performance, in terms of the Akaike information criterion (AIC) scores. The ARMA (2,1) model is described by:

$$\omega_t = \phi_1 \omega_{t-1} + \phi_2 \omega_{t-2} + \theta_1 \epsilon_{t-1} + \epsilon_t$$

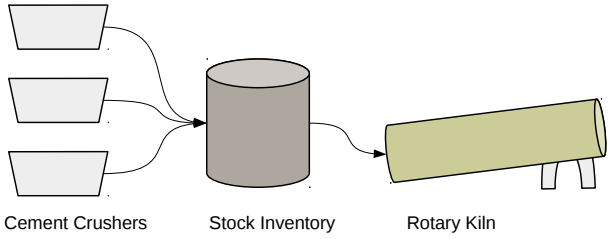

Fig. 1: Cement plant illustration.

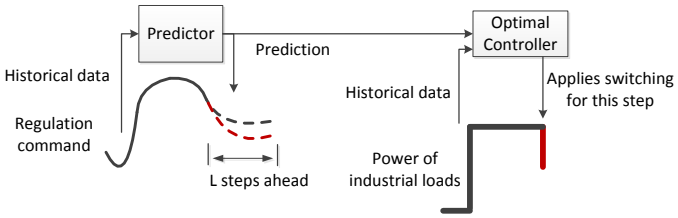

Fig. 2: MPC coordination framework.

in which ω_t stands for the regulation signal and ϵ_t stands for the white noise. The auto-regressive parameters ϕ_1, ϕ_2 and moving-average parameters θ_1 are trained and obtained. The regulation prediction mean squared errors by ARMA(2,1) for different prediction horizons are plotted in Fig. 3 with comparison to the Persistence Prediction and the Mean Prediction approach. The Persistence Prediction uses the latest available observation as prediction and the Mean Prediction uses the average from all available observations as prediction. According to Fig. 3, the ARMA(2,1) model results in a good performance up to horizons of around 1 minute.

B. Optimal Control

The objective of the optimal control is to provide high quality regulation service at low cost. The decision variables for the optimal control are the number of active machines and the charging power for the storage. The regulation capacity and the regulation baseline, denoted as R and B (MW) respectively, are determined in advance by the schedule optimization approach presented in Section IV. Note, as the regulation signal is assumed to be well balanced, i.e. its integral over time is zero, which is the case for RegD signal, the average power consumption of the industrial machines is B MW. This indicates that the throughput from these crushing machines, which is proportional to the energy (MWh) it consumes, is decided by baseline B. The formulations for the optimal control are stated as follows.

1) Objective: We penalize the regulation violation v_i, the amount of switch actions s_i, and the deviation d of the final storage energy level from the targeted level, as in:

$$\text{minimize} \quad \sum_{i \in \mathbb{L}} (\alpha v_i + \beta s_i) + \gamma d \tag{1}$$

in which $\mathbb{L} = \{0, 1, ..., L\}$ is the set of time steps in the current MPC horizon and $\alpha, \beta, \gamma \geq 0$ are the penalty parameters. Different values of the penalty parameters indicate different

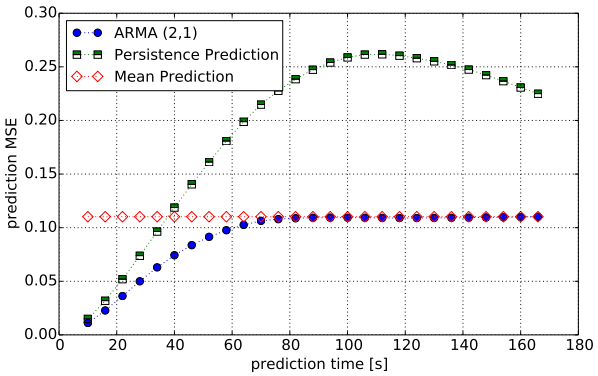

Fig. 3: Prediction mean squared errors.

preferences for the regulation provision. Details of the impact of these parameters are discussed in the case study.

2) Regulation Violation: Within the MPC horizon at time t, the regulation signal prediction for step i is denoted by $\widehat{\omega}_{t+i}$. According to the regulation prediction, the regulation violation v_{t+i} at the i-th step is defined as:

$$v_{t+i} \geq |B + R\widehat{\omega}_{t+i} - \rho x_{t+i} - y_{t+i}| \quad \forall i \in \mathbb{L} \tag{2}$$

in which the first two terms on the right side correspond to the regulation command and the last two terms correspond to the plant power consumption rate; ρ is the power consumption rate of the machine, x_{t+i} represents the number of active machine, and y_{t+i} denotes the charging power of the storage system. Since we penalize v_{t+i} in the objective function, i.e. α is positive, the above constraint can be formulated as two linear inequality constraints, i.e.

$$v_{t+i} \geq B + R\widehat{\omega}_{t+i} - \rho x_{t+i} - y_{t+i}$$
$$v_{t+i} \geq -B - R\widehat{\omega}_{t+i} + \rho x_{t+i} + y_{t+i}$$

Similar formulations apply to the following two constraints.

3) Machine Switching: Too much switching of the machines potentially increases degradation and may even damage the machines; that is why we penalize the amount of switch actions in the objective function. The amount of switch actions s_{t+i} at the i-th step is given by:

$$s_i \geq |x_{t+i} - x_{t+i-1}| \quad \forall i \in \mathbb{L} \tag{3}$$

in which the right side represents the change in the number of active machines between time steps.

4) Storage Level Deviation: Another objective is to control the final energy level in the storage by the end of each MPC horizon. Otherwise, if the energy level is near to its full capacity, then there is little room for the storage to contribute to the provision of regulation for the following MPC horizons. This deviation is defined as

$$d \geq |e_{t+L} - \overline{e}| \tag{4}$$

in which $\overline{e}$ is the targeted storage level. We usually set $\overline{e}$ equal to 50% of its energy capacity.

5) Storage Energy Balance: The energy balance for the storage describes the dynamic relationship between stored energy and its charging power, as given by:

$$e_{t+i} - e_{t+i-1} = y_{t+i}\delta \quad \forall i \in \mathbb{L} \tag{5}$$

where δ is the length of one time step. In addition, the energy in the storage is constrained by the storage capacity.

6) Switching Limitation: In practice, the industrial machines cannot be switched on/off without any limitation as the machines could get damaged by too much switching. Hence, we restrict the number of switch actions to be no more than $\bar{s}$ for every successive K steps (typically, $K > L$) for each MPC step t and each time i in the MPC horizon, as given by:

$$\sum_{j=t+i-K}^{t-1} \widetilde{s}_j + \sum_{j=t}^{t+i} s_j \leq \bar{s} \quad \forall i \in \mathbb{L} \tag{6}$$

The first term corresponds to the summation of switch actions that already took place before t, and the second term stands for the possible number of switch actions that may take place from t to the i-th step $t+i$. Note that the above constraint applies to each step i within the MPC time horizon, as we require the switching to not violate the bound on the number of switchings for every successive K steps. Consider an example where L is 20 steps and K corresponds to 100 steps, if there are already $\bar{s}$ times of switching between t-90 and t, then for the MPC horizon starting at t, its first 10 steps cannot allow for any switching. Other constraints on switching limitation can be considered in a similar way, e.g. the requirement for the machines to consume a certain minimum amount of energy for a specific number of successive time steps.

7) Variable Range: The decision variables can take values within the following bounds:

$$x_{t+i} \in \{0, 1, ..., M\} \text{ and } -P_s \leq y_{t+i} \leq P_s \quad \forall i \in \mathbb{L} \tag{7}$$

in which x_{t+i} is an integer variable while y_{t+i} is continuous.

To sum up, the MPC recedes forward and at each time step t, it first predicts the upcoming regulation signals, then optimizes (1) subject to constraints (2)-(7), but only applies the optimal decisions at time step t. The resulting optimization problem is a mixed-integer linear programming (MIP) problem, which can be solved by CPLEX very quickly as the problem size is small. The complexity of solving the MIP problem grows exponentially with the number of integer variables. The integer variables on the other hand, which model the switch statuses of the machines, linearly depends on the number of machines in the plant. In practice, the number of machines in a plant is usually quite small.

Note that the above constraints (2) and (3) are based on our assumption that each machine has the same power consumption rate ρ and we use one variable x to denote the number of active machines. In case the power consumption rates of different machines are not the same, we can extend the above constraints and model the on/off status for each individual machine by using a vector of binary variables. For example, let $\mathbb{M}$ be the set of machines and x_t^m be a binary variable representing the on/off status of machine m with $m \in \mathbb{M}$ at time t, then constraints (2) and (3) need to be extended to

$$v_{t+i} \geq |B + R\widehat{\omega}_{t+i} - \sum_{m \in \mathbb{M}} \rho^m x_{t+i}^m - y_{t+i}| \quad \forall i \in \mathbb{L}$$

where ρ^m is the power consumption rate for machine m, and

$$s_i \geq \sum_{m \in \mathbb{M}} |x_{t+i}^m - x_{t+i-1}^m| \quad \forall i \in \mathbb{L}$$

respectively.

IV. Day-ahead Optimal Scheduling

In Section III, the power baseline B and regulation capacity R are pre-specified for the hourly operation. Here, we now consider the day-ahead scheduling and the goal is to optimally determine the hourly B and R for the load. Specifically for a cement plant, our scheduling objective is to maximize its daily profit which consists of the revenue from industrial production and the revenue from regulation provision minus the cost of energy consumption and the cost of providing regulation. We also need to ensure that the kiln keeps on running at a constant rate continuously and prevent it from turning off, as restarting the kiln is very expensive.

A. Constraints

1) Power Baseline: The power baseline B is the base for regulation provision, which equals the load's power consumption rate if no regulation is provided, i.e. without charging/discharging of energy storage nor switching of machines. Hence, B equals the sum of power from a subset of available machines. Consequently, we have a limited choice of values for B. Note that when the load is providing regulation, B is very close to the average power consumption rate, because the regulation signal is assumed to be well balanced (e.g. RegD signal) and therefore its hourly integral is almost zero.

The baseline B is determined by how many crushers are turned on; in case that the machines are not identical with each other, it is also decided by which machines are turned on. For a fixed number of machines, the combinations of their on/off statuses are limited, hence the choices of B are limited.

For each such choice, we term the corresponding machines' statuses as profile $p \in \mathbb{P}$, where $\mathbb{P}$ is the set of all possible profiles, and denote the baseline power as B_p MW. Note that B_p and $\mathbb{P}$ can be obtained as parameters once we know the plant's configuration. We then use binary variable $z_{p,h}$ to denote whether the plant chooses a baseline with profile p, where $z_{p,h} \in \{0, 1\}$ and $h \in \mathbb{H}$ with $\mathbb{H}$ as the set of hours in the scheduling horizon; this is equivalent to using a binary vector of size m to indicate the on/off statuses of m machines, as there is a bijective mapping between the profile p and the binary vector. Since the load chooses only one baseline profile for each hour, we have the following constraint regarding the choice of baseline:

$$\sum_{p\in\mathbb{P}} z_{p,h} = 1 \quad \forall h \in \mathbb{H} \tag{8}$$

2) Regulation Capacity: Unlike the power baseline, the regulation capacity R is continuous; meanwhile, the maximum capacity of R depends on the baseline power B, as B determines the available machines for switching. We use continuous variable $r_{p,h}$ to denote the regulation capacity the load provides given that machine switching profile p is chosen during hour h; for each hour h, only the chosen power baseline's correspondent $r_{p,h}$ is nonzero. Since the energy storage itself is able to provide a regulation of P_s MW, here we consider the case when $r_{p,h}$ is greater than P_s, i.e., we switch machines to provide a larger amount of regulation than the possible amount by merely using the storage. We have the following constraint on the bounds for $r_{p,h}$

$$R_p^{lo} z_{p,h} \le r_{p,h} \le R_p^{up} z_{p,h} \quad \forall p \in \mathbb{P} \tag{9}$$

where R_p^{lo} and R_p^{up} are bounds associated with baseline profile p. Note that if a baseline profile is not chosen, then its correspondent regulation $r_{p,h}$ is zero because $z_{p,h}$ is zero.

3) Inventory Stock: Since the production rate of the machines is proportional to their energy consumption rate, the choice of p impacts the intermediate product generation rate. We assume that the following stage consumes the intermediate product at a constant rate τ, and any interruption to the next stage should be avoided as this could be costly. Therefore, there always needs to be enough intermediate product in stock. We use variable q_h to represent the quantity of intermediate product in the inventory after hour h. With the assumption that the production rate is proportional to the power consumption rate, unit conversion can be done and we denote q_h with the unit of MWh for simplicity; similarly, the consumption rate τ of the next stage is also in the unit of MW, and the intermediate production generation rate is B_p MW. This leads to the following constraint regarding the dynamic balance for the inventory stock:

$$q_h + \sum_{p\in\mathbb{P}} B_p z_{p,h} - \tau = q_{h+1} \tag{10}$$

The above constraint should always hold in order to keep the next stage running. Besides, the inventory stock is subject to the following bounds:

$$Q^{lo} \le q_h \le Q^{up} \tag{11}$$

where the parameters Q^{lo} and Q^{up} (also converted to the unit of MWh) correspond to the minimum and maximum amount of intermediate stock allowed in the inventory.

4) Regulation Cost: The provision of regulation does not come for free as it leads to more switching of the machines and therefore increased equipment degradation. Hence, we need to quantify the hourly regulation cost before optimizing the day-ahead scheduling. In other words, we want to choose the optimal regulation capacity R for each hour, and before that, we need to know the cost corresponding to each possible R. As discussed in Section III, the regulation provision cost depends on the actual AGC signal trace. However, the relationship between the cost and the AGC signal is nonlinear and complex; besides, the AGC signal itself is uncertain and impossible to predict over a long interval (e.g. more than 5 minutes). To simplify the problem and focus on the day-ahead operation of the industrial plant, we approximate the hourly regulation cost by only considering the cost associated with the switching of machines.

In order to quantify the cost associated with the machine switching, we need to know the switching amount for each hour, but that amount is unknown before the actual hourly operation. However, since the AGC signal is approximately normally distributed and its trace is similar from hour to hour, the average hourly switching amount from historical operations can serve as an approximation of future hourly regulation cost in our day-ahead scheduling. As the average hourly switching also depends on the values of B and R, we run numerical simulations with various settings of R and B using historical AGC traces (published by PJM) and record the number of switching for these cases. This gives an approximate average number of switching for each B and R. We then use that average as the approximated switching cost (as a function of B and R) in the day-ahead scheduling problem to determine the optimal B and R for each operating hour. From simulations, we observe that the switching cost linearly depends on R for a given B, as seen in Fig. 10. Hence, we model the switching cost to be $C_p^0 + C_p^1 R$ for a given power baseline profile p; in other words, we assume the switching cost increases linearly with regulation capacity R, while the coefficients C_p^0 and C_p^1 are determined by the power baseline. The coefficients C_p^0 and C_p^1 can be obtained by regression of the historical operation records, which is presented in detail in Section V-C.

To sum up, in consideration of all possible baselines, the hourly switching cost is expressed as follows:

$$C_h = \sum_{p\in\mathbb{P}} (C_p^0 z_{p,h} + C_p^1 r_{p,h}) \tag{12}$$

Note that $r_{p,h}$ is zero when $z_{p,h}$ is zero according to Eq. (9). The coefficients C_p^0 and C_p^1 can take into account the monetary cost of the switch actions, hence the hourly switching cost C_h is in the unit of \$.

B. Objective Function

Suppose the hourly energy price, regulation price, and industrial product market price are known, then the regulation revenue is proportional to R, and both the energy cost and industrial product revenue are proportional to B. We use $\lambda_{R,h}$ to denote the hourly regulation price, and use $\lambda_{E,h}$ to denote the hourly net energy profit price which equals the unit revenue in industrial production minus the unit cost in energy consumption. Since we assume the kiln operates at a constant speed and hence the final product yield does not change, we only maximize the net revenue from the crushing stage. Then our final optimization problem is:

$$\begin{aligned} \underset{z_{p,h},\ r_{p,h}}{\text{maximize}} \quad & \sum_{h\in\mathbb{H}} (\lambda_{E,h} \sum_{p\in\mathbb{P}} B_p z_{p,h} + \lambda_{R,h} \sum_{p\in\mathbb{P}} r_{p,h} - C_h) \\ \text{subject to} \quad & (8)-(12) \end{aligned}$$

which is a mixed-integer linear programming problem, where the number of integer variables linearly depends on the number of machines in the plant, as these variables indicate the statuses of the crushers. And again we use CPLEX to solve the problem.

V. Case Study

A. Industrial Plant Parameters and ESS Cost Estimation

For the case study, we consider a cement plant with $M = 4$ crushing machines. These machines can be switched on and off rapidly. The power consumption rate of each machine is either $\rho = 2$ MW (when it is on) or zero (when it is off). The plant has an on-site energy storage system whose maximum charging/discharging power is $P_s = 3$ MW and its energy capacity is $E_s = 9$ MWh. The crushing machines constitute the first stage of the cement production process, followed by a cement kiln burner which heats the intermediate products to a certain temperature. The kiln burner cannot be interrupted and consumes the intermediate product at a constant rate. The constant consumption rate is equivalent to $\tau = 4$ MW. The intermediate product between the crushing and burning is stored in an inventory, which has a maximum capacity that is equivalent to Q^{up} = 20 MWh and its initial stock is 10 MWh.

In order to estimate the benefit and cost from providing regulation by installing an ESS for such a cement plant, we use \$1,000/kW to estimate the cost of a Li-ion energy storage system whose capacity (MWh) / power (MW) ratio is 3 [41], [42]. Hence, a 3 MW (9 MWh) Li-ion ESS roughly costs \$3M. The cement industry has been one of the least profitable businesses, whose net profit margin is only 2.03% [43]. The traditional scheduling strategy, in such a capital-intensive industry, is to produce as much as possible by running all machines at their full capacities, in order to fully utilize the invested equipment and resources. With the falling prices of the cement products, it could be profitable to reduce the production rate, and utilize its capabilities in power response to earn revenues in the electricity markets. For the above cement plant, its nominal power consumption rate is 8 MW and its nominal production rate is 480 tons of cement per hour [44].

We will compare three scheduling options and focus on its net cost in the electricity markets per ton of cement produced. The cement plant is assumed to operate 24 hours a day and 7 days a week. The market prices are taken from a typical day (Jan 30, 2017) in MISO, where the average prices over the day are \$29.87/MWh for energy, \$11.36/MW for regulation capacity, \$0.72/MW for regulation mileage, and \$2.45/MW for spinning reserve capacity. In option (1), the plant operates at its full capacity with a power consumption rate of 8 MW and does not provide any ancillary service. Its daily cement throughput is 11520 tons and its daily energy cost is \$5735.04. In option (2), the plant operates at its full capacity and also provides 8 MW spinning reserve. Since the spinning reserve dispatch rate is very rare (less than 0.5% according to [30]), we assume there is no actual deployment. Its daily cement throughput is 11520 tons and its daily net cost is \$5264.64 (energy \$5735.04, spinning reserve \$470.40) in the electricity markets. In option (3), the plant operates at its half capacity and provides 7 MW regulation with the help of the ESS, which corresponds to the simulation discussed later and shown in Fig. 6. In calculating the regulation mileage payment, we assume a mild mileage as 20 p.u. per hour, which corresponds to an AGC signal as shown in Fig. 4. Its daily cement throughput is 5760 tons and its daily net cost is \$-1460.16 (energy \$2867.52, regulation capacity \$1908.48, regulation mileage \$2419.20). Hence, the net cost in the electricity market per ton of cement processed (\$/ton) are 0.50, 0.46, and -0.25, respectively.

By comparing option (2) and (3), we observe that the saving increase in the electricity markets is \$0.71/ton on the typical day, if providing regulation instead of spinning reserve. Therefore, the investment of the 3 MW ESS can be covered by producing around 4.23 millions (3M/0.71) of tons of cement, with the assumption that the prices on the typical day can approximate the average prices. In other words, if we choose option (3) and let the plant process 240 tons of cement per hour, the ESS can be covered in two years. Of course, the analysis here is not comprehensive enough and it is only meant to illustrate the idea that lowering down production rate and providing regulation could be profitable for the cement plant.

B. Simulations of Real-time MPC Coordination

The simulations of real-time operations are studied to demonstrate the benefits of the proposed MPC coordination approach. Since the market cycle is an hour, the simulations are performed hourly. As R and B are given values in the hourly operation, here we consider the cement plant providing $R = 5$ MW regulation at a baseline of $B = 2$ MW; as seen later in Section V-D, this setting of R and B corresponds to the scheduling result in Fig. 11, e.g. hour 8 and hour 20. The regulation command ranges between -3 MW and 7 MW. Note that this range of the regulation command is 10 MW,

which is much higher than that of the energy storage (6 MW). The 20 minutes regulation signal for the simulations is plotted in Fig 4, together with the ARMA(2,1) prediction at a few distinct time instances. The length of the time step is $\delta = 2$ seconds, and the prediction horizon is $L = 15$ steps. The penalties α, β, γ in Section III are set to 10, the targeted final energy is $\bar{e} = 0.5$ MWh, and we require the maximum number of switch actions to be $\bar{s} = 10$ times for every successive 5 minutes, i.e. $K = 150$ steps.

The simulations of hourly operation, i.e. the real-time following of the AGC signal, the switch actions, and the energy storage operation, are plotted in Fig. 5. The dashed lines in the middle plot are the bounds for the charging/discharging power of the storage. According to the simulation, the integral of regulation violation over the hour is 0 MWh, i.e. there is no violation at all, the total number of switch actions is 15, and the storage energy level at the end of the hour is 0.64 MWh, i.e the energy deviation is 0.14 MWh. These results demonstrate that the coordination method proposed for real-time operation is able to utilize the advantages of both the industrial loads and the energy storage, and provides high-quality regulation service to support the power system operation.

We also study the case of providing an even larger amount of regulation. We consider the cement providing $R = 7$ MW regulation at a baseline of $B = 4$ MW, which corresponds to the scheduling result in Fig. 12, e.g. from hour 4 to hour 22. The simulation results are displayed in Fig. 6 and all the configurations are exactly the same as the simulation in Fig. 5. Over the hour, there is only a slight violation - the integral of regulation violation over the hour is 0.001 MWh, the total number of switch actions is 27, and the storage energy level at the end of the hour is 0.60 MWh, i.e the energy deviation is 0.10 MWh. Compared with Fig. 5, the machines are switched more frequently and there is some violation of regulation following, which is the cost of increasing the regulation capacity.

The sensitivity with respect to changes in parameter settings is investigated by simulations with different penalty values. If we impose a stronger switch limitation constraint, e.g. requiring the maximum number of switches to be 3 for every successive 5 minutes, then the switching frequency will decrease, as demonstrated by the simulation results in Fig. 7. The total number of switches decreases to 21 times, but the regulation violation increases to 0.25 MWh. If we also increase the penalty on switch actions β, the total number of switch actions is expected to further decrease. For example, increasing β to 100 while keeping all other parameters the same as Fig. 7 yields the simulation results in Fig. 8, in which the total

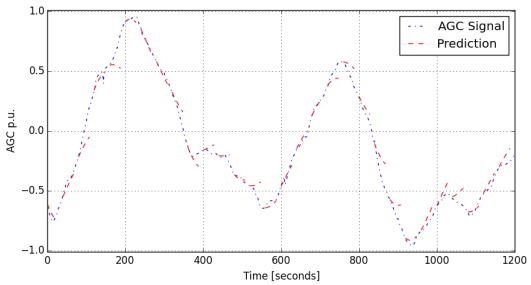

Fig. 4: Regulation signal (AGC) over 20 minutes and its prediction.

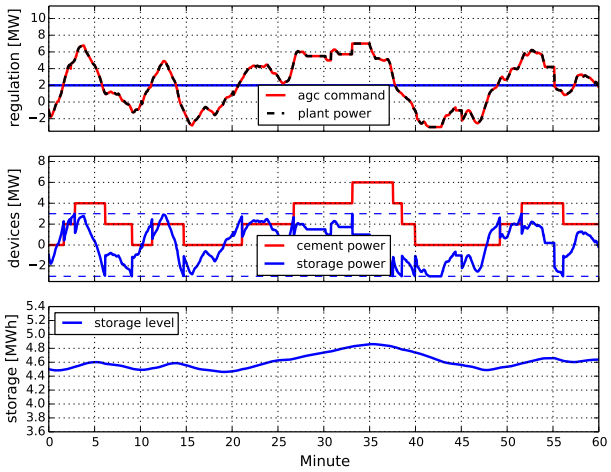

Fig. 5: Hourly real-time simulations with R = 5 MW, B = 2 MW.

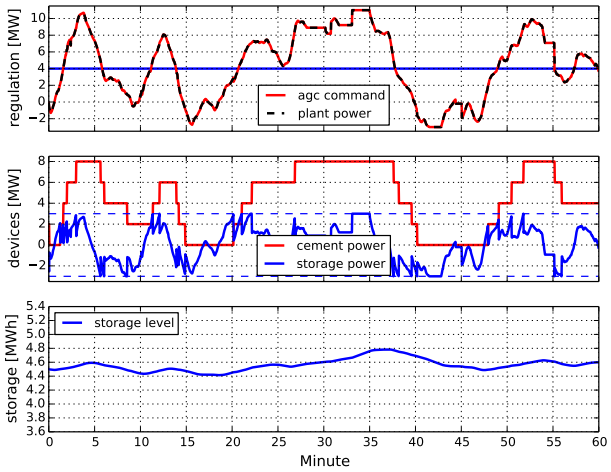

Fig. 6: Hourly real-time simulations with R = 7 MW, B = 4 MW.

number of switches further decreases to 19 times, however, there is even more regulation violation, i.e. 0.27 MWh.

In practice, we suggest that the plant operators choose their own penalties according to their preferences. For example, in an electricity market where the regulation quality is highly valued, a higher regulation violation penalty α is recommended; meanwhile, if switching the machines is very expensive, then the operator should use a large switch action penalty β.

C. Quantifying Hourly Regulation Cost

In order to optimize the scheduling for the day-ahead operation as in Section IV, we need to quantify the hourly cost of regulation provision. As discussed in Section IV-A4, we approximate the hourly regulation cost by using the average hourly switching quantities from the records of historical operation.

To obtain the historical regulation cost, we simulate the MPC coordination for each hour over three months by using the historical AGC signal published by PJM. All the three penalties α, β, γ are still chosen as 10. Simulations with different choices of baseline power B and varying regulation capacity R are studied. For the considered plant, we have three baseline powers to choose from, i.e. 2, 4, or 6 MW, leaving at least one machine for switching. The regulation cost given different pairs of (R, B) are obtained for each hour in the

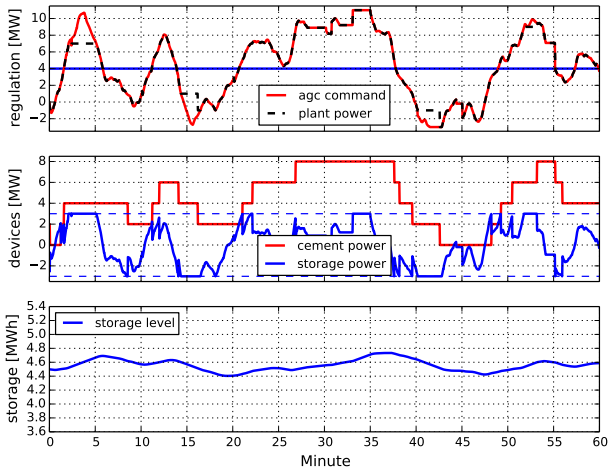

Fig. 7: Stronger switching limitation with R = 7 MW, B = 4 MW.

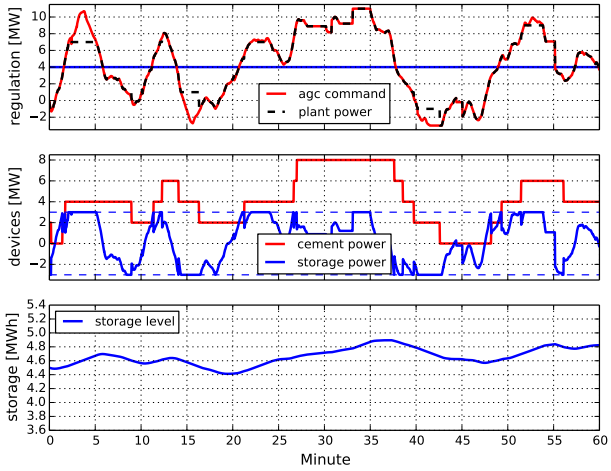

Fig. 8: Increased penalty on switching with R = 7 MW, B = 4 MW.

historical data. We present the trace of each regulation cost component over two days in Fig 9(a), 9(b), and 9(c); we only present the plots with B = 4 MW here, while the observations are similar when B is 2 or 6 MW. From these plots, we observe that: (1) a larger regulation capacity generally leads to a larger amount of switch actions; (2) the regulation violation is zero for most of the hours, and when it is not zero, the violation increases with the regulation capacity; (3) the energy storage deviation fluctuates around zero as it depends on the integral of the hourly AGC signal. The hourly energy consumption by the machines over these two days is plotted in Fig. 9(d), which demonstrates that the hourly energy consumption rate is very close to the baseline power.

In the day-ahead scheduling, we only consider the switching cost. With the simulated hourly switching cost over these three months, we can obtain the average amount of switch actions $\bar{C}_{B,R}$ for any given pair of R and B. This average switch action is displayed in Fig 10, where all three possible baseline power are considered. For each baseline power B, we apply linear regression to fit the output C (response/dependent variable) to the input R (explanatory/independent variable). The fitted relationships are also plotted as dashed lines in Fig 10. These fitted linear relationships which map R to C under a chosen B provide the regulation cost coefficients C_b^0

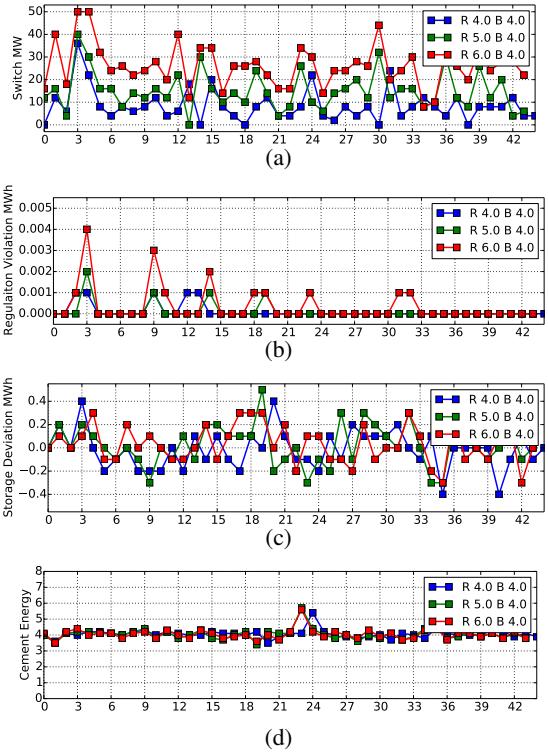

Fig. 9: Hourly simulation results over 48 hours.

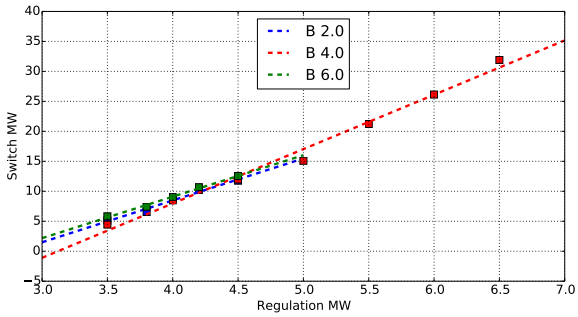

Fig. 10: Average hourly switch MW and its linear fitting.

and C_b^1 for the day-ahead scheduling.

D. Simulations of Day-ahead Optimal Scheduling

Here we apply the optimal scheduling proposed in Section IV to the industrial plant. The hourly electricity energy price and regulation price are taken from historical records from MISO. The profit from cement production is assumed to be \$30/MWh, i.e., for every 1 MWh energy consumption, the machines generate products worth of \$30. The coefficients of hourly regulation cost, C_b^0 and C_b^1, are taken from the regression result in Fig. 10, with the assumption that the monetary cost of switching is \$0.5/MW. All these scheduling optimizations can be solved by CPLEX within minutes as the problem size is small. In practice, its computation does not need to be very fast as the problem is intended for day-ahead scheduling which only needs to be solved a few times every day. The day-ahead scheduling is presented in Fig. 11. From the result we see that when the energy price is lower, the baseline power is higher, i.e. the cement plant takes advantage of the lower energy price and consumes more energy by

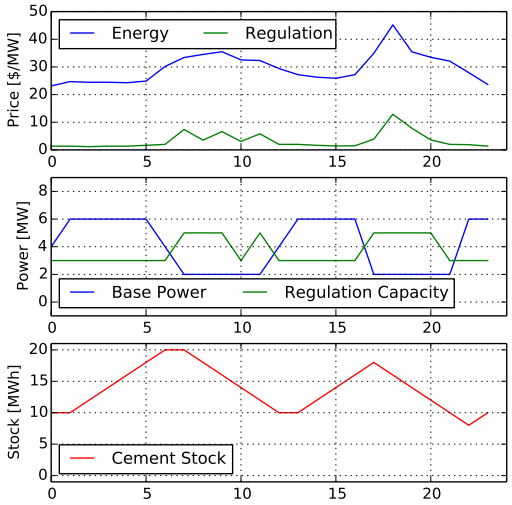

Fig. 11: Day-ahead scheduling result.

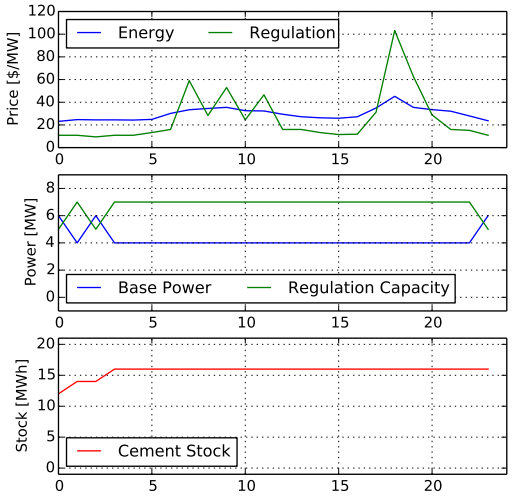

Fig. 12: Day-ahead scheduling with higher regulation price.

speeding up production; while when the energy price is higher, the baseline power is lower and the cement plant consumes less energy. We also observe that the regulation capacity increases when the regulation price is higher, e.g. around hour 8 and hour 18. If we manually increase the regulation price, e.g. artificially multiply the regulation price by 8, the industrial plant will concentrate on the regulation provision, as in Fig. 12. During most of the hours, the baseline power is 4 MW, which guarantees the largest regulation availability.

VI. Conclusion

The key contribution of this work is the proposed approaches for providing the most valuable ancillary services such as regulation and load following by the combination of industrial loads, which can adjust their power consumption only in large discrete steps, and an on-site energy storage device, which provides the more granular power adjustments. Both real-time and day-ahead operations are considered: an MPC approach determines the hourly operation and an optimal scheduling approach handles the daily operation.

Given these proposed approaches, the loads are enabled with more options in supporting the power system operation. The loads are able to overcome the restriction of poor granularity and provide regulation or load following ancillary services. The daily scheduling method provides a tool for plant operators to optimally arrange their production activities with demand response provision; it also helps plant operators to better understand how much profit can be earned from demand response participation, which can encourage more industrial loads to actively contribute to power system operation. Note that the approaches proposed in this paper apply to a variety of loads, e.g. cement crushing, paper milling, and can enable both load following and regulation service; the proposed framework can also help coordinate the industrial loads with other power system components such as commercial buildings, electric vehicles, etc.

For future work, it is necessary to extend the proposed scheduling method to develop a bidding tool that considers the uncertainty of market prices; it is also crucial to address some practical concerns such as the charging loss in the storage.

References

[1] "U.S. solar market sets new record, installing 7.3 gw of solar pv in 2015," Posted: Feb 22, 2016. [Online]. Available: http://www.seia.org/news/us-solar-market-sets-new-record-installing-73-gw-solar-pv-2015

[2] "Germany leads europe in offshore wind energy growth," Posted: February 02, 2016. [Online]. Available: http://www.theguardian.com/environment/2016/feb/02/germany-leads-europe-in-offshore-wind-energy-growth

[3] M. G. Vayá and G. Andersson, "Self scheduling of plug-in electric vehicle aggregator to provide balancing services for wind power," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 2, pp. 886–899, 2016.

[4] H. Marzooghi, G. Verbič, and D. J. Hill, "Aggregated demand response modelling for future grid scenarios," *Sustainable Energy, Grids and Networks*, vol. 5, pp. 94–104, 2016.

[5] J. M. Morales, A. J. Conejo, H. Madsen, P. Pinson, and M. Zugno, "Facilitating renewable integration by demand response," in *Integrating Renewables in Electricity Markets*. Springer US, 2014, pp. 289–329.

[6] M. D. Ilić, L. Xie, and J.-Y. Joo, "Efficient coordination of wind power and price-responsive demand - Part I: Theoretical foundations," *IEEE Trans. on Power Systems*, vol. 26, no. 4, pp. 1875–1884, 2011.

[7] J. Wang, C. Liu, D. Ton, Y. Zhou, J. Kim, and A. Vyas, "Impact of plug-in hybrid electric vehicles on power systems with demand response and wind power," *Energy Policy*, vol. 39, no. 7, pp. 4016–4021, 2011.

[8] W. Wei, F. Liu, and S. Mei, "Charging strategies of EV aggregator under renewable generation and congestion: A normalized nash equilibrium approach," *To appear in IEEE Transactions on Smart Grid*, 2016.

[9] Y. Yao, W. Gao, J. Momoh, and E. Muljadi, "Economic dispatch for microgrid containing electric vehicles via probabilistic modelling," in *IEEE North American Power Symposium*, 2015.

[10] Z. Chen, L. Wu, and Y. Fu, "Real-time price-based demand response management for residential appliances via stochastic optimization and robust optimization," *IEEE Transactions on Smart Grid*, vol. 3, no. 4, pp. 1822–1831, Dec 2012.

[11] K. Margellos and S. Oren, "Capacity controlled demand side management: A stochastic pricing analysis," *IEEE Transactions on Power Systems*, vol. 31, no. 1, pp. 706–717, 2016.

[12] E. Mocanu, P. H. Nguyen, M. Gibescu, and W. L. Kling, "Deep learning for estimating building energy consumption," *Sustainable Energy, Grids and Networks*, vol. 6, pp. 91–99, 2016.

[13] I. Beil, I. Hiskens, and S. Backhaus, "Frequency regulation from commercial building HVAC demand response," *Proceedings of the IEEE*, vol. 104, no. 4, pp. 745–757, 2016.

[14] E. Vrettos and G. Andersson, "Scheduling and provision of secondary frequency reserves by aggregations of commercial buildings," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 2, pp. 850–864, 2016.

[15] F. Pallonetto, S. Oxizidis, F. Milano, and D. Finn, "The effect of time-of-use tariffs on the demand response flexibility of an all-electric smart-grid-ready dwelling," *Energy and Buildings*, vol. 128, pp. 56–67, 2016.

[16] C. Wu, W. Tang, K. Poolla, and R. Rajagopal, "Predictability, constancy and contingency in electric load profiles," in *IEEE Smart Grid Communications*, 2016.

[17] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing data center demand response," in *ACM International Conference on Measurement and Modeling of Computer Systems*, 2014, pp. 111–123.

[18] X. Zhang and G. Hug, "Optimal regulation provision by aluminum smelters," in *IEEE Power and Energy Society General Meeting*, 2014.

[19] ——, "Bidding strategy in energy and spinning reserve markets for aluminum smelters' demand response," in *IEEE Innovative Smart Grid Technologies Conference*, 2015.

[20] Q. Zhang, C. F. Heuberger, I. E. Grossmann, A. Sundaramoorthy, and J. M. Pinto, "Air separation with cryogenic energy storage: optimal scheduling considering electric energy and reserve markets," *AIChE Journal*, 2015.

[21] H. Hadera, I. Harjunkoski, G. Sand, I. E. Grossmann, and S. Engell, "Optimization of steel production scheduling with complex time-sensitive electricity cost," *Computers & Chemical Engineering*, vol. 76, pp. 117–136, 2015.

[22] X. Zhang, G. Hug, Z. Kolter, and I. Harjunkoski, "Computational approaches for efficient scheduling of steel plants as demand response resource," in *IEEE Power Systems Computation Conference*, 2016.

[23] X. Zhang, G. Hug, and I. Harjunkoski, "Cost-effective scheduling of steel plants with flexible EAFs," *IEEE Trans. on Smart Grid*, 2016.

[24] T. Samad and S. Kiliccote, "Smart grid technologies and applications for the industrial sector," *Computers & Chemical Engineering*, vol. 47, pp. 76 – 84, 2012.

[25] L. Merkert, I. Harjunkoski, A. Isaksson, S. Säynevirta, A. Saarela, and G. Sand, "Scheduling and energy–industrial challenges and opportunities," *Computers & Chemical Engineering*, vol. 72, pp. 183–198, 2015.

[26] I. Harjunkoski and H. Hadera, "Industrial tools and needs," in *Alternative Energy Sources and Technologies*. Springer, 2016, pp. 415–438.

[27] X. Zhang, G. Hug, Z. Kolter, and I. Harjunkoski, "Industrial demand response by steel plants with spinning reserve provision," in *47th North American Power Symposium*, 2015.

[28] Q. Zhang, M. F. Morari, I. E. Grossmann, A. Sundaramoorthy, and J. M. Pinto, "An adjustable robust optimization approach to scheduling of continuous industrial processes providing interruptible load," *Computers & Chemical Engineering*, vol. 86, pp. 106–119, 2016.

[29] *Business Practices Manual 02: Energy and Operating Reserve Markets*, MISO, Oct 2016. [Online]. Available: https://www.misoenergy.org/Library/BusinessPracticesManuals/

[30] D. Todd, M. Caufield, B. Helms, A. P. Generating, I. M. Starke, B. Kirby, and J. Kueck, "Providing reliability services through demand response: A preliminary evaluation of the demand response capabilities of Alcoa Inc," *ORNL/TM*, vol. 233, 2008.

[31] *PJM Manual 12: Balancing Operations*, PJM, Feb 2017. [Online]. Available: http://www.pjm.com/ /media/documents/manuals/m12.ashx

[32] R. Vujanic, S. Mariéthoz, P. Goulart, and M. Morari, "Robust integer optimization and scheduling problems for large electricity consumers," in *IEEE American Control Conference*, 2012, pp. 3108–3113.

[33] X. Zhang, G. Hug, Z. Kolter, and I. Harjunkoski, "Model predictive control of industrial loads and energy storage for demand response," in *IEEE PES General Meeting*, 2016.

[34] H. Hadera, P. Wide, I. Harjunkoski, J. Mäntysaari, J. Ekström, G. Sand, and S. Engell, "A mean value cross decomposition strategy for demand-side management of a pulping process," *Computer Aided Chemical Engineering*, vol. 37, pp. 1931–1936, 2015.

[35] F. Golestaneh, P. Pinson, and H. B. Gooi, "Very short-term nonparametric probabilistic forecasting of renewable energy generation with application to solar energy," *IEEE Transactions on Power Systems*, vol. PP, no. 99, pp. 1–14, 2016.

[36] X. Sun, P. B. Luh, K. W. Cheung, W. Guan, L. D. Michel, S. Venkata, and M. T. Miller, "An efficient approach to short-term load forecasting at the distribution level," *IEEE Transactions on Power Systems*, vol. 31, no. 4, pp. 2526–2537, 2016.

[37] M. Wytock and J. Z. Kolter, "Large-scale probabilistic forecasting in energy systems using sparse gaussian conditional random fields," in *52nd IEEE Conference on Decision and Control*, 2013, pp. 1019–1024.

[38] S. Chen, T. Zhang, H. Gooi, R. D. Masiello, and W. Katzenstein, "Penetration rate and effectiveness studies of aggregated BESS for frequency regulation," *IEEE Trans. on Smart Grid*, vol. 7, no. 1, pp. 167–177, 2016.

[39] X. Zhang, F. Gao, X. Lv, H. Lv, Q. Tian, J. Ma, W. Yin, and J. Dong, "Line loss reduction with distributed energy storage systems," in *IEEE Innovative Smart Grid Technologies*, 2012.

[40] "The PJM fast response regulation signal," PJM, Posted: September, 2013. [Online]. Available: http://www.pjm.com/markets-and-operations/ancillary-services/mkt-based-regulation/fast-response-regulation-signal.aspx

[41] N. Ivanova, "Lithium-ion costs to fall by up to 50% within five years," July 2015. [Online]. Available: http://analysis.energystorageupdate.com/lithium-ion-costs-fall-50-within-five-years

[42] N. DiOrio, A. Dobos, and S. Janzou, "Economic analysis case studies of battery energy storage with SAM," *National Renewable Energy Laboratory*, 2015.

[43] "Most and least profitable business types," January 2011. [Online]. Available: https://www.bloomberg.com/news/photo-essays/2011-01-18/most-and-least-profitable-business-types

[44] R. T. Lidbetter and L. Liebenberg, "Load-shifting opportunities for a typical south african cement plant," in *IEEE Industrial and Commercial Use of Energy*, 2011, pp. 17–25.

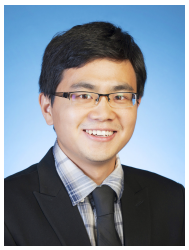

Xiao Zhang (S13) received the B.S. degree in electrical engineering from Tsinghua University in Beijing, China. He is currently pursuing the Ph.D. degree at Carnegie Mellon University in Pittsburgh, USA. His research interests include optimization and algorithms as well as their applications.

Gabriela Hug (S05, M08, SM14) was born in Baden, Switzerland. She received the M.Sc. degree in electrical engineering from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, in 2004 and the Ph.D. degree from the same institution in 2008. After her PhD, she worked in the Special Studies Group of Hydro One in Toronto, Canada and from 2009 - 2015 she was an Assistant Professor at Carnegie Mellon University in Pittsburgh, USA. Currently, she is an Associate Professor at the Power Systems Laboratory at ETH Zurich. Her research is dedicated to control and optimization of electric power systems.

J. Zico Kolter is an Assistant Professor in the School of Computer Science at Carnegie Mellon University, with appointments in the Computer Science Department, the Institute for Software Research (in the Societal Computing program), and affiliated appointments with the Machine Learning Department, the Robotics Institute, and the Electrical and Computer Engineering Department. His work focuses on machine learning and optimization, with a specific focus on applications in smart energy systems. From an algorithmic standpoint, he has worked on fast optimization algorithms for a number of problems and for general convex programs, large-scale probabilistic modeling, stochastic optimization, and deep learning. On the application side, he has worked on energy disaggregation, probabilistic forecasting for energy systems, and model predictive control techniques for industrial control in the electrical grid.

Iiro Harjunkoski is Corporate Research Fellow at ABB Corporate Research Germany. He holds a PhD in Chemical Engineering from Åbo Akademi University, Finland (1997). He is an expert on planning and scheduling, mathematical modeling and optimization of production processes. Further research interests include vertical and horizontal integration of Manufacturing Execution Systems (MES) with the objective to maximize the overall throughput and energy efficiency. He regularly receives requests for consulting or project collaboration both inside and outside of ABB and is actively networking with top universities to advance the development of planning and scheduling for industrial processes.