Handling Long Horizons in MPC: A Stochastic Programming Approach

Ranjeet Kumar¶, Michael J. Wenzel‡, Matthew J. Ellis‡,
Mohammad N. ElBsat‡, Kirk H. Drees‡, Victor M. Zavala¶
¶Department of Chemical and Biological Engineering
University of Wisconsin-Madison, 1415 Engineering Dr, Madison, WI 53706, USA
‡Johnson Controls International Plc
507 E. Michigan St., Milwaukee, WI 53202, USA

***Abstract*— We propose an approach to handle long horizons in model predictive control (MPC). The approach is based on the observation that, if periodicity constraints are enforced over short-term stages, the long horizon MPC problem can be cast as a stochastic programming (SP) problem. The SP representation reveals a mechanism to construct a hierarchical MPC scheme under which a high-level (long-horizon) MPC controller provides periodic state targets to guide a low-level (short-term) MPC controller. We show that this hierarchical scheme is optimal under nominal (perfect forecast) conditions and can be extended to handle imperfect forecasts by correcting the targets in real-time. We demonstrate our concepts using a building system with stationary battery storage, where the goal is to use the battery to mitigate monthly demand charges while collecting revenue from hourly frequency regulation markets. We demonstrate that the hierarchical MPC scheme yields improved performance over standard MPC schemes because it can systematically capture long-term effects.**

I. Introduction

MPC provides a powerful planning framework for industrial applications but it is well-known that its scope is hindered by the computational complexity associated to the length of the planning horizon. This limitation is relevant in applications such as energy systems, inventory management, and scheduling. In the context of energy systems, long planning horizons are often needed to capture monthly peak electricity costs (demand charges) [1]. Peak electricity demands remain a serious concern for utilities, because they pose capacity challenges to the power grid. This situation has created incentives for energy storage [2], [3]. In this context, MPC has been used in [4] to mitigate demand charges of a building HVAC system. Due to the inability to handle long horizons, the authors penalize the peak electricity demand over the short-term receding horizon. Unfortunately, this approach results in overly conservative MPC policies and deteriorate economic performance. Long-term discounting cost factors are used in [5] to reduce the conservatism of short-term MPC, but this approach is ad-hoc and does not have optimality guarantees.

The work in [6] proposes a hierarchical MPC scheme to handle long time horizons. Here, adjoint information obtained from a long-term and coarse MPC controller is used to guide the control policy of a short-term MPC controller. It is demonstrated that this approach can achieve optimality. Unfortunately, such an approach requires continuity of the adjoint profiles, which is not guaranteed in general formulations with peak costs (such as those arising from demand charges) and integer decisions (as those arising in scheduling formulations). Alternative hierarchical MPC schemes [7], [8] available in the literature provide feasibility but not optimality guarantees.

In this work we show that, when periodicity constraints are enforced over short-term stages, the long-horizon MPC problem can be cast as a stochastic programming (SP) problem. Under this setting, a short-term planning stage represents a scenario, the design variables are the initial/terminal periodic states, and the recourse decisions are the intra-stage operational policies. The SP representation suggests a hierarchical MPC scheme in which optimal design decisions are used as targets to guide a short-term MPC controller. We show that, under nominal conditions with perfect forecasts, the hierarchical scheme yields an optimal policy of the long-term MPC problem. The SP setting also reveals strategies to correct targets and short-term policies when forecasts are imperfect. The proposed approach can be easily generalized to handle peak costs, by incorporating an additional design variable that imposes a demand budget to the short-term MPC controller. The approach can also handle general nonlinear models and integer decisions and can enable the use of parallel decomposition schemes.

We demonstrate the performance of the hierarchical MPC scheme using an application in buildings with electricity storage. Here, the goal is to determine optimal short-term (hourly) participation strategies in frequency regulation (FR) and energy markets while simultaneously mitigating long-term (monthly) demand charges from utilities. We compare the solution of the hierarchical MPC scheme with that from a standard MPC scheme that uses a discounting (weighting) factor to account for the inability of solving the long horizon problem. We confirm that the proposed hierarchical MPC scheme achieves optimal policies under perfect forecasts of the electrical load. Under imperfect forecasts, we find that the cost from the hierarchical MPC scheme provides superior performance over standard MPC because it can capture long-

Corresponding Author: victor.zavala@wisc.edu

term variability. The proposed developments can also be used in other applications that need to handle long horizons, such as inventory management and scheduling formulations.

II. Hierarchical MPC Scheme

We consider an MPC planning problem of the form:

$$\min_{u_t} \sum_{t\in\mathscr{T}} \varphi_1(x_t,u_t,d_t) + \max_{t\in\mathscr{T}} \varphi_2(x_t,u_t,d_t) \tag{1a}$$

$$\text{s.t. } x_{t+1} = f(x_t,u_t,d_t),\ t\in\bar{\mathscr{T}} \tag{1b}$$

$$x_0 = \bar{x}_0,\ x_t\in\mathscr{X}, u_t\in\mathscr{U}. \tag{1c}$$

Here, $\mathscr{T} := \{0,...,N\}$ is the set of time steps with N being the final time in the horizon (with $\bar{\mathscr{T}} = \mathscr{T}\setminus\{N\}$), $\varphi_1(\cdot)$ is a time-additive cost function, and $\varphi_2(\cdot)$ is a time-max (peak) cost function. The controls, states, and true (realized) disturbances at time t are denoted as u_t, x_t, and d_t, respectively. The forecasted disturbance at time t is $\hat{d}_t$ and initial state is $\bar{x}_0$.

We consider a partition (in lexicographic order) of the time horizon set $\mathscr{T}$ (and of $\bar{\mathscr{T}}$) into M *stages* with shorter horizons. We denote each stage by $\xi\in\Xi$ with $\mathscr{T}_\xi := \{0,...,N_\xi\}$ being the inner stage time horizon and $\Xi := \{0,...,M\}$ being the set of all stages, satisfying $\mathscr{T} = \cup_{\xi\in\Xi}\mathscr{T}_\xi$ and $\sum_{\xi\in\Xi} N_\xi = N$. For convenience, we define the set $\bar{\Xi} := \Xi\setminus\{M\}$. Note that each stage Ξ can have a different time horizon. We partition states, controls, and disturbance policies into stages and denote the stage policies as $u_{\xi,t}$, $x_{\xi,t}$, and $d_{\xi,t}$ for $\xi\in\Xi$ and $t\in\mathscr{T}_\xi$. We use these partitions to reformulate the MPC problem in the following *equivalent* form:

$$\min_{u_{\xi,t}} \sum_{\xi\in\Xi}\sum_{t\in\mathscr{T}_\xi} \varphi_1(x_{\xi,t},u_{\xi,t},d_{\xi,t}) + \max_{\xi\in\Xi}\max_{t\in\mathscr{T}_\xi} \varphi_2(x_{\xi,t},u_{\xi,t},d_{\xi,t}) \tag{2a}$$

$$\text{s.t. } x_{\xi,t+1} = f(x_{\xi,t},u_{\xi,t},d_{\xi,t}),\ \xi\in\Xi, t\in\bar{\mathscr{T}}_\xi \tag{2b}$$

$$x_{\xi+1,0} = x_{\xi,N_\xi}, \xi\in\bar{\Xi} \tag{2c}$$

$$x_{0,0} = \bar{x}_0,\ x_{\xi,t}\in\mathscr{X}, u_{\xi,t}\in\mathscr{U}. \tag{2d}$$

Here, the constraint (2c) enforces continuity between stages. We modify the MPC planning problem by assuming that periodicity is enforced at the end of every stage:

$$\min_{u_{\xi,t},x_{\xi,0}} \sum_{\xi\in\Xi}\sum_{t\in\mathscr{T}_\xi} \varphi_1(x_{\xi,t},u_{\xi,t},d_{\xi,t}) + \max_{\xi\in\Xi}\max_{t\in\mathscr{T}_\xi} \varphi_2(x_{\xi,t},u_{\xi,t},d_{\xi,t}) \tag{3a}$$

$$\text{s.t. } x_{\xi,t+1} = f(x_{\xi,t},u_{\xi,t},d_{\xi,t}),\ \xi\in\Xi, t\in\bar{\mathscr{T}}_\xi \tag{3b}$$

$$x_{\xi+1,0} = x_{\xi,N_\xi}, \xi\in\bar{\Xi} \tag{3c}$$

$$x_{\xi,N_\xi} = x_{\xi,0}, \xi\in\Xi \tag{3d}$$

$$x_{\xi,t}\in\mathscr{X}, u_{\xi,t}\in\mathscr{U}. \tag{3e}$$

In this formulation (that we call *long-term MPC*), $x_{0,0}$ is a free variable that we seek to optimize. We also note that the periodicity constraint (3d) together with the stage continuity constraints (3c) can be expressed as $x_{\xi+1,0} = x_{\xi,0}$, $\xi\in\bar{\Xi}$. These constraints, in turn, can be reformulated as $x_{\xi,0} = x_0$, $\xi\in\Xi$ by introducing an additional *lifting* variable x_0. Consequently, the goal of the long-term MPC formulation (3) is to find the optimal periodic state x_0 and control policies $u_{\xi,t}$, $\xi\in\Xi, t\in\mathscr{T}_\xi$ that minimize the time-additive and peak costs. The assumption on periodicity is reasonable in applications such as energy and inventory management systems, in which disturbance profiles (e.g, demands and prices) have a strong periodic component [9]–[11].

We can reformulate the peak cost function to obtain the following equivalent form of (3):

$$\min_{u_{\xi,t},x_0,\eta} \sum_{\xi\in\Xi}\sum_{t\in\mathscr{T}_\xi} \varphi_1(x_{\xi,t},u_{\xi,t},d_{\xi,t}) + \eta \tag{4a}$$

$$\text{s.t. } \varphi_2(x_{\xi,t},u_{\xi,t},d_{\xi,t}) \le \eta, \xi\in\Xi, t\in\mathscr{T}_\xi \tag{4b}$$

$$x_{\xi,t+1} = f(x_{\xi,t},u_{\xi,t},d_{\xi,t}),\ \xi\in\Xi, t\in\bar{\mathscr{T}}_\xi \tag{4c}$$

$$x_{\xi,0} = x_0, \xi\in\Xi \tag{4d}$$

$$x_{\xi,N_\xi} = x_0, \xi\in\Xi \tag{4e}$$

$$x_{\xi,t}\in\mathscr{X}, u_{\xi,t}\in\mathscr{U}. \tag{4f}$$

We denote the solution of (4) as $x^*_{\xi,t}, u^*_{\xi,t}, \eta^*$ and note that $\eta^* = \max_{\xi\in\Xi}\max_{t\in\mathscr{T}_\xi} \varphi_2(x^*_{\xi,t},u^*_{\xi,t},d_{\xi,t})$ holds. Consequently, η^* is the peak cost over the entire planning horizon. We also note that $x^*_{\xi+1,0} = x^*_{\xi,0} = x^*_0$. From the structure of (4) we note that the only coupling between stages arises from the variables x_0 and η. Consequently, (3) can be seen as a *stochastic programming problem* in which stages are operational scenarios, x_0 and η are *design* variables, and $x_{\xi,t}, u_{\xi,t}$ are *scenario* (recourse) policies. When no peak costs are present, the only design variable of the long-term MPC formulation is the periodic state x_0.

By fixing the design variables to their optimal values x^*_0 and η^*, we can decompose problem (4) into M stage subproblems of the form:

$$\min_{u_{\xi,t}} \sum_{t\in\mathscr{T}_\xi} \varphi_1(x_{\xi,t},u_{\xi,t},d_{\xi,t}) + \eta^* \tag{5a}$$

$$\text{s.t. } \varphi_2(x_{\xi,t},u_{\xi,t},d_{\xi,t}) \le \eta^*,\ t\in\mathscr{T}_\xi \tag{5b}$$

$$x_{\xi,t+1} = f(x_{\xi,t},u_{\xi,t},d_{\xi,t}),\ t\in\bar{\mathscr{T}}_\xi \tag{5c}$$

$$x_{\xi,0} = x^*_0 \tag{5d}$$

$$x_{\xi,N_\xi} = x^*_0 \tag{5e}$$

$$x_{\xi,t}\in\mathscr{X}, u_{\xi,t}\in\mathscr{U}. \tag{5f}$$

The subproblem has the structure of a standard optimal control problem with periodicity constraints. The SP formulation (3) thus suggests a *hierarchical MPC scheme*, in which the long-term MPC problem (4) (equivalently (2)) guides a short-term MPC controller (with formulation (5)). The communication between the hierarchical levels arises in the form of targets for the periodic state x^*_0 and the peak cost η^*. Figure 1 provides a sketch of the hierarchical MPC scheme. Because the targets $x^*_{\xi,0}$ and η^* are optimal, the solution of the MPC stage problem yields an optimal stage trajectory $u^*_{\xi,t}$ and $x^*_{\xi,t}$ (or a trajectory that achieves the same optimal stage cost). Consequently, the targets obtained from the SP problem (2) are optimal when the *disturbance forecasts are perfect* ($d_{\xi,t} = \hat{d}_{\xi,t}$) and the solution of the SP problem is equivalent to that of the long-horizon MPC problem (3).

The SP setting also indicates that, when the forecasts are imperfect (i.e., $\hat{d}_{\xi,t} \neq d_{\xi,t}$), the targets x^*_0, η^* will not be optimal and the stage (recourse) problem (5) might not have a feasible solution. This situation can be mitigated by penalizing deviations from the target (which will find the closest feasible point) or by re-optimizing the design targets using the actual realized disturbances when they become available in real-time. The SP setting also reveals that by increasing the number of scenarios in the planning problem (e.g., by using multiple possible realizations of the forecast), it is possible to find targets that remain optimal and feasible for all stage realizations. Consequently, the proposed approach provides a framework to easily construct robust formulations. Another important feature of the hierarchical MPC approach is that it avoids the need to weight or discount the long-term costs in the short-horizon MPC subproblem. Moreover, one can solve the SP problem efficiently by using decomposition schemes based on parallel linear algebra and Benders/Lagrangian decomposition [12], [13].

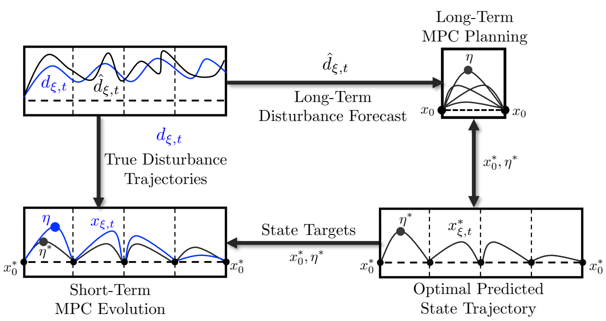

Fig. 1: Hierarchical MPC scheme.

III. Stationary Battery Study

Batteries are flexible assets that can be used to provide energy and frequency regulation (FR) capacity for independent system operators (ISOs) and that can aid utility companies by providing demand-side management capabilities for buildings or manufacturing facilities [14]. The use of batteries for *simultaneous* FR and demand charge mitigation has been studied in [15]. In such settings, the objective function of the MPC scheme takes the form of the formulation in (1), where the additive costs represent market revenue and peak costs represent demand charges. The proposed battery framework is based on recent work by the authors reported in [16] and is sketched in Figure 2.

A. Long-Term MPC Formulation

The elements of the long-term MPC formulation for the battery planning problem include the following model parameters, data, and variables:

a) Model Parameters and Data:

- $L_{\xi,t} \in \mathbb{R}$: Buildings load [kW].
- $\pi^{e}_{\xi,t} \in \mathbb{R}$: Market price for electricity [\$/kWh].
- $\pi^{f}_{\xi,t} \in \mathbb{R}_{+}$: Market price for regulation capacity [\$/kW].
- $\pi^{D} \in \mathbb{R}_{+}$: Demand charge (monthly) [\$/kW].
- $\alpha_{\xi,t} \in [-1,1]$: Fraction of FR capacity requested by ISO [-]. If $\alpha_t > 0$, the ISO sends a power to the battery while if $\alpha_t < 0$ the ISO withdraws power.
- $\overline{E} \in \mathbb{R}$: Battery storage capacity [kWh].
- $\overline{P} \in \mathbb{R}$: Maximum discharging rate (power) [kW].
- $\underline{P} \in \mathbb{R}$: Maximum charging rate (power) [kW].
- $\rho \in \mathbb{R}$: Minimum fraction of battery capacity reserved for frequency regulation [kWh/kW].
- $\overline{\Delta P} \in \mathbb{R}$: Maximum ramping limit [kW/h].

The forecasted disturbances in the battery problem are loads, prices, and regulation signals.

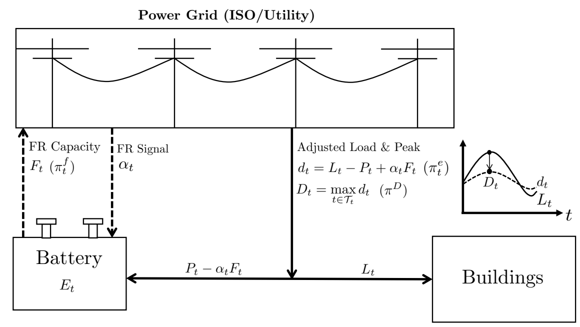

Fig. 2: Interactions battery, buildings, ISO, and utility.

b) Model Variables: We define the following variables for the SP formulation, each replicated for all stages $\xi \in \Xi$:

- $P_{\xi,t} \in \mathbb{R}$: Net battery discharge rate (power) [kW].
- $F_{\xi,t} \in \mathbb{R}_{+}$: FR capacity provided to ISO [kW].
- $E_{\xi,t} \in \mathbb{R}_{+}$: State of charge (SOC) of the battery [kWh].
- $d_{\xi,t} \in \mathbb{R}_{+}$: Load requested from utility [kW].
- $D = \max_{\xi \in \Xi} \max_{t \in \mathcal{T}_{\xi}} d_{\xi,t}$: Peak load over horizon $\mathcal{T}$ [kW]

c) Objective Function: We minimize total cost, which is given by the demand charge and the revenues collected from power and regulation:

$$\sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}_{\xi}} \left(-\pi^{e}_{\xi,t} (P_{\xi,t} - \alpha_{\xi,t} F_{\xi,t}) - \pi^{f}_{\xi,t} F_{\xi,t} \right) + \pi^{D} D. \quad (6)$$

Here, the revenues represent the time-additive cost and the demand charge is the time-max cost.

d) Constraints: The constraints of the SP are replicated for every realization $\xi \in \Xi$. The net charged/discharged battery power plus the FR capacity provided must be within the maximum discharging and charging rates $\overline{P}$ and $\underline{P}$:

$$P_{\xi,t} + F_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}_{\xi}, \xi \in \Xi \quad (7a)$$

$$P_{\xi,t} - F_{\xi,t} \geq -\underline{P},\ t \in \mathcal{T}_{\xi}, \xi \in \Xi \quad (7b)$$

The storage dynamics are given by the difference equation:

$$E_{\xi,t+1} = E_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t} F_{\xi,t},\ t \in \bar{\mathcal{T}}_{\xi}, \xi \in \Xi \quad (8)$$

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

$$\rho F_{\xi,t} \leq E_{\xi,t} \leq \overline{E} - \rho F_{\xi,t},\ t \in \mathcal{T}_{\xi}, \xi \in \Xi \quad (9a)$$

$$\rho F_{\xi,t} \leq E_{\xi,t+1} \leq \overline{E} - \rho F_{\xi,t},\ t \in \bar{\mathcal{T}}_{\xi}, \xi \in \Xi \quad (9b)$$

The battery ramp discharge rate is constrained as:

$$-\overline{\Delta P} \leq P_{\xi,t+1} - P_{\xi,t} \leq \overline{\Delta P},\ t \in \bar{\mathcal{T}}_{\xi}, \xi \in \Xi \quad (10)$$

The residual demand d_k requested from the utility is:

$$d_{\xi,t} = L_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t} F_{\xi,t},\ t \in \mathcal{T}_{\xi}, \xi \in \Xi \quad (11)$$

The peak demand must satisfy $d_{\xi,t} \leq D$, $t \in \mathcal{T}_\xi, \xi \in \Xi$. We assume that the ISO does not allow the battery to sell back electricity. This is modeled by using the constraint $P_{\xi,t} + F_{\xi,t} \leq L_{\xi,t}$, $t \in \mathcal{T}_\xi, \xi \in \Xi$. We enforce non-anticipativity constraint on the initial SOC $E_{\xi,0} = E_0$. Finally, we enforce periodicity constraints on the SOC $E_{\xi,N_\xi} = E_0$, $\xi \in \Xi$. The bounds on the variables are:

$$0 \leq E_{\xi,t} \leq \overline{E}, t \in \mathcal{T}_\xi, \xi \in \Xi \tag{12a}$$

$$-\underline{P} \leq P_{\xi,t} \leq \overline{P}, t \in \mathcal{T}_\xi, \xi \in \Xi \tag{12b}$$

$$0 \leq F_{\xi,t} \leq \overline{P}, t \in \mathcal{T}_\xi, \xi \in \Xi \tag{12c}$$

The SP is solved to obtain the targets for the periodic SOC E_0^* and for the peak demand D^*. These targets are then used to guide a short-term MPC controller that we describe next.

B. Short-Term MPC Formulation

For simplicity, we assume that the short-term MPC controller only updates its control policies at the beginning of stage $t = t_\xi$ (where $t_\xi = \xi N_\xi, \xi \in \bar{\Xi}$) over horizon $\mathcal{T}_\xi := \{t, t+1, ..., t+N_\xi\}$. The short-term problem at time t_ξ uses forecasts for prices and loads over the horizon $\mathcal{T}_\xi$ (in the perfect information case this matches the scenarios of the long-term MPC formulation). The solution of the problem at time t_ξ is implemented for a block of N_ξ hours (i.e., N_ξ represents the update frequency in the MPC scheme). This approach is different than the traditional MPC approach in which the control policies are updated at every time step *within* the stage ξ. The short-term MPC formulation is:

$$\min_{P_{\xi,t}, F_{\xi,t}} \sum_{t \in \mathcal{T}_\xi} \left(-\pi^e_{\xi,t}(P_{\xi,t} - \alpha_{\xi,t} F_{\xi,t}) - \pi^f_{\xi,t} F_{\xi,t} \right) + \pi^D D^* \tag{13a}$$

$$s.t. \quad P_{\xi,t} + F_{\xi,t} \leq \overline{P}, \; t \in \mathcal{T}_\xi \tag{13b}$$

$$P_{\xi,t} - F_{\xi,t} \geq -\underline{P}, \; t \in \mathcal{T}_\xi \tag{13c}$$

$$E_{\xi,t+1} = E_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t} F_{\xi,t}, \; t \in \bar{\mathcal{T}}_\xi \tag{13d}$$

$$\rho F_{\xi,t} \leq E_{\xi,t} \leq \overline{E} - \rho F_{\xi,t}, \; t \in \mathcal{T}_\xi \tag{13e}$$

$$\rho F_{\xi,t} \leq E_{\xi,t+1} \leq \overline{E} - \rho F_{\xi,t}, \; t \in \bar{\mathcal{T}}_\xi \tag{13f}$$

$$-\overline{\Delta P} \leq P_{\xi,t+1} - P_{\xi,t} \leq \overline{\Delta P}, \; t \in \bar{\mathcal{T}}_\xi \tag{13g}$$

$$d_{\xi,t} = L_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t} F_{\xi,t}, \; t \in \mathcal{T}_\xi \tag{13h}$$

$$P_{\xi,t} + F_{\xi,t} \leq L_{\xi,t}, \; t \in \mathcal{T}_\xi \tag{13i}$$

$$E_{\xi,N} = E_0^* \tag{13j}$$

$$E_{\xi,0} = E_0^* \tag{13k}$$

$$d_{\xi,t} \leq D^* \tag{13l}$$

$$0 \leq E_{\xi,t} \leq \overline{E}, \; t \in \mathcal{T}_\xi \tag{13m}$$

$$-\underline{P} \leq P_{\xi,t} \leq \overline{P}, \; t \in \mathcal{T}_\xi \tag{13n}$$

$$0 \leq F_{\xi,t} \leq \overline{P}, \; t \in \mathcal{T}_\xi \tag{13o}$$

Here, E_0^* and D^* are fixed targets. In the perfect information case, the performance of the short-term MPC controller is optimal (because the true disturbance profiles match the forecasts used in the long-term MPC formulation). We now provide a strategy to update the short-term stage policies for the case of imperfect forecasts.

C. Analysis for Imperfect Forecasts

Under imperfect forecast case, the disturbance profile observed in real time can be different from the one that was used to for long-term planning and determine the targets. Therefore, the target will not be optimal for the short-term MPC problem and can also be infeasible. To account for infeasibility, we add a penalty term to the cost function of the short-term MPC controller (13a) of the form $\pi^S S$. Here, the slack variable S is penalized by the cost π^S, where π^S must satisfy $\pi^S > \pi^D$ to guarantee that, if the peak target is feasible, then $S = 0$. In the case of an infeasible peak target, the MPC problem returns the smallest value of S, and therefore, the closest feasible operating point.

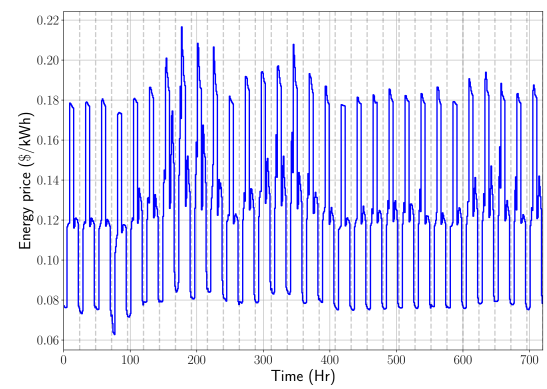

(a) Energy price for a month.

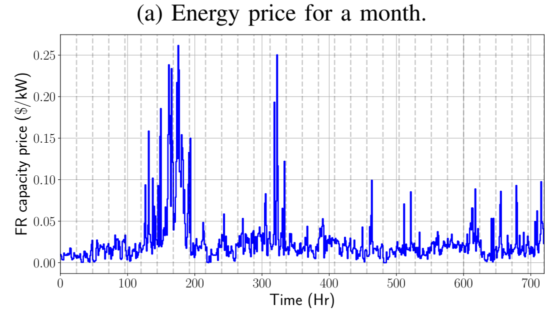

(b) Regulation capacity price for a month.

Fig. 3: Market price data used for the case studies.

IV. Results

We consider a utility-scale stationary battery that has a capacity 0.5 MWh, rated power of 1 MW for both charge and discharge, and we assume a ramping limit of 0.5 MW/hr and at least 10% of battery capacity reserved for FR at every hour. We use historical data for one month for energy prices and FR prices from PJM Interconnection, shown in Figures 3a, 3b. Historical load data from a typical university campus for a month is used as the disturbance profile and is shown in Figure 4. We consider a planning horizon of one month (i.e., $N = 720$) and we use stages of 24 hours to create the SP formulation (i.e., $N_\xi = 24$ and $M = 30$).

We assume perfect forecasts in the first case study considered. For these experiments we compare the cost of the hierarchical MPC scheme with the cost of a long-term MPC formulation *with and without periodicity constraints*. This comparison seeks to evaluate the impact of assuming an optimal periodic policy. In the second study we evaluate the performance of the hierarchical MPC scheme under imperfect forecasts. To do so, we compare performance against a standard MPC scheme that performs hourly updates of the control policy and that uses a prediction horizon of 24 hours. For the standard MPC approach we do not impose periodicity constraints. Instead, we use a discounting (weighting) factor of $1/30$ for the demand charges [16].

A. Perfect Forecasts

Figure 4 compares the policies obtained with the long-term MPC planning problem (labeled as Long-Term MPC) and the hierarchical MPC scheme (labeled as Hierarchical). The grey vertical lines denote 24-hour stages (scenarios). We see that the policies are identical; the equivalence indicates that the solutions of the stage subproblems are unique (for fixed targets D^* and E_0^*). Figure 5 shows the SOC policy obtained from the long-term MPC problem with no periodicity constraints (1). It can be observed that the SOC is close to periodic. The cost items obtained with the different formulations are summarized in Table I. We see that the total cost obtained with no periodicity constraints is only 0.002% lower than that obtained with periodicity constraints (2). We thus conclude that assuming periodicity does not limit performance.

B. Imperfect Forecasts

Figure 6 shows the forecasted load and the true (realized) load used in the case with imperfect forecasts. It can be observed that the peak in the realized load (33,315 kW) is 1800 kW higher than the peak in the forecasted load (31,486 kW). The design peak demand (D^*) obtained for the hierarchical MPC scheme using the forecasted load profile is 31,186 kW. This target value D^* leads to infeasibility of the short-term MPC controller when using the realized load profile. In Table II, we compare the cost items obtained with the different MPC schemes. We observe that the total cost obtained from the hierarchical MPC is 0.9% higher than that obtained with the long-term MPC formulation that uses the realized load to compute policies (perfect information). By comparing the demand charges of these two schemes, we see that hierarchical MPC is able to identify the optimal demand charge of long-term MPC. This is important because the demand charge is a significant component of the total cost. The higher total cost of hierarchical MPC is thus attributed to the suboptimal periodic SOC targets obtained from long-term MPC using the forecasted load profiles. We also evaluate the performance of standard MPC schemes, one that uses the realized load (perfect information) and another that uses the forecasted load (imperfect information) to compute policies. By comparing long-term MPC and standard MPC with realized loads (perfect forecasts), we note that standard MPC yields a suboptimal policy. This highlights that the use of the discount factor only provides and ad-hoc approximation. Under imperfect forecasts, standard MPC results in a higher total cost compared to hierarchical MPC. This result highlights that hierarchical MPC provides a more

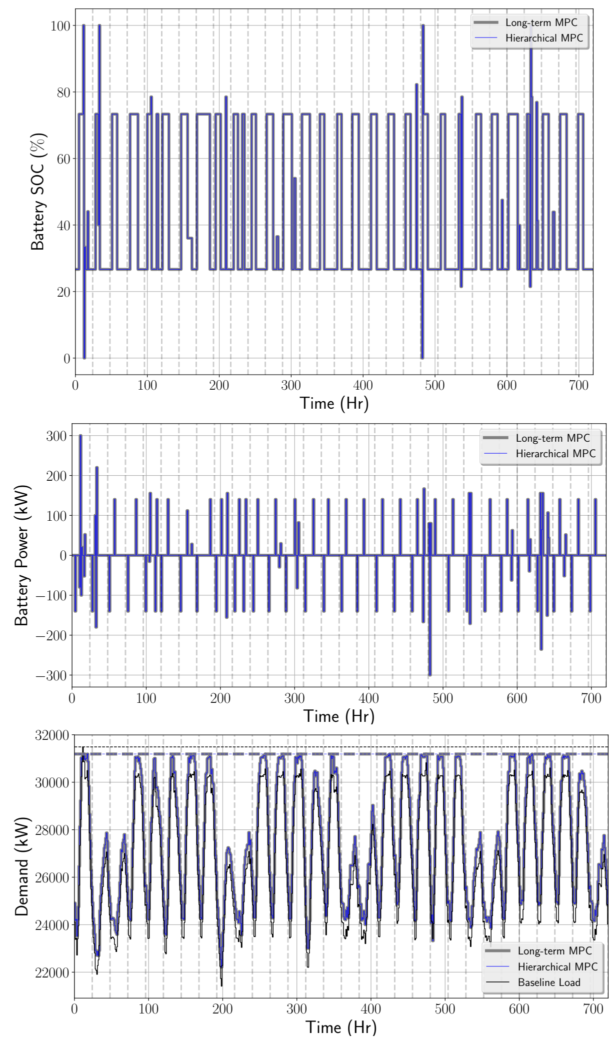

Fig. 4: Comparison of battery SOC (top), battery power (middle) and demand (bottom) policies.

effective approach to handle long-term demand charges because it systematically captures the load variability observed throughout the month.

V. Conclusions and Future Work

We proposed an approach to handle long horizons in MPC based on the observation that, if periodicity constraints are enforced over short-term stages, the long horizon MPC problem can be cast as a stochastic programming (SP) problem. The SP setting reveals a mechanism to construct a hierarchical MPC scheme under which a high-level (long-horizon) MPC controller provides state targets to guide a low-level (short-horizon) MPC controller. We show that this hierarchical MPC architecture is optimal under nominal (perfect forecast) conditions and can be extended to handle imperfect forecasts by correcting short-term policies. Extensions to this work include updates (re-optimization) of the SP formulation to correct the periodic initial/terminal states. The use of a SP setting also opens the door to the use of scalable decomposition methods to tackle the long-term horizon problem. In particular, cutting-plane algorithms can

TABLE I: Comparison of cost items under perfect forecasts.

Cost Item ($/month)	Long-Term MPC (with periodicity)	Hierarhical MPC	Long-Term MPC (without periodicity)
Total cost	114,079.81	114,079.81	114,077.99
Demand charge	129,424.86	129,424.86	129,424.86
FR	-14,861.72	-14,861.72	-14,863.76
Energy	-483.33	-483.33	-483.11

TABLE II: Comparison of cost items under imperfect forecasts.

Cost Item ($/month)	Long-Term MPC (Forecasted Load)	Long-Term MPC (Realized Load)	Hierarchical MPC (Imperfect Forecast)	Standard MPC (Perfect Forecast)	Standard MPC (Imperfect Forecast)
Total cost	114,079.81	122,317.12	123,449.17	123,816.49	126,230.41
Demand charge	129,424.86	137,635.12	137,635.12	138,472.22	141,579.55
FR	-14,861.72	-14,850.11	-13,690.61	-14,178.68	-14,866.15
Energy	-483.33	-467.88	-495.33	-477.04	-482.98

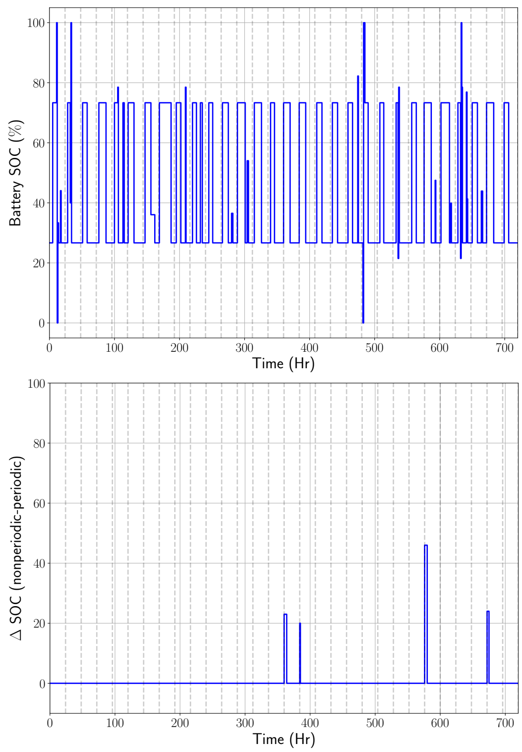

Fig. 5: SOC policy without periodicity (top). Difference of SOC policies with and without periodicity constraints (right).

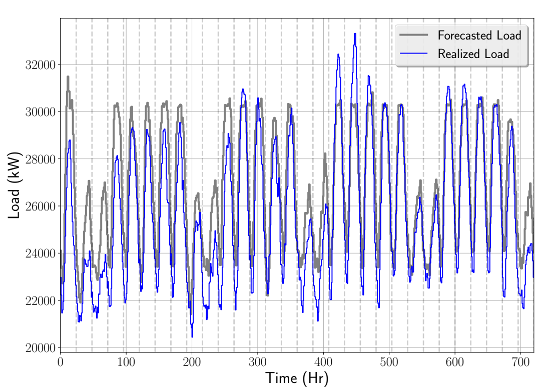

Fig. 6: Forecasted and realized load profiles.

be used to progressively update the high-level MPC layer.

REFERENCES

[1] J. Ma, J. Qin, T. Salsbury, and P. Xu, "Demand reduction in building energy systems based on economic model predictive control," *Chemical Engineering Science*, vol. 67, no. 1, pp. 92–100, 2012.

[2] 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.

[3] K. A. Joshi and N. M. Pindoriya, "Day-ahead dispatch of Battery Energy Storage System for peak load shaving and load leveling in low voltage unbalance distribution networks," in *IEEE Power and Energy Society General Meeting*, vol. 2015-Septe. IEEE, jul 2015, pp. 1–5.

[4] J. E. Braun, "Reducing energy costs and peak electrical demand through optimal control of building thermal storage," in *ASHRAE Transactions*, no. pt 2, 1990, pp. 876–888.

[5] N. R. Patel, M. J. Risbeck, J. B. Rawlings, M. J. Wenzel, and R. D. Turney, "Distributed economic model predictive control for large-scale building temperature regulation," in *American Control Conference (ACC), 2016*. IEEE, 2016, pp. 895–900.

[6] V. M. Zavala, "New architectures for hierarchical predictive control," *IFAC-PapersOnLine*, vol. 49, no. 7, pp. 43–48, 2016.

[7] R. Scattolini and P. Colaneri, "Hierarchical model predictive control," in *Decision and Control, 2007 46th IEEE Conference on*. IEEE, 2007, pp. 4803–4808.

[8] B. Picasso, D. De Vito, R. Scattolini, and P. Colaneri, "An mpc approach to the design of two-layer hierarchical control systems," *Automatica*, vol. 46, no. 5, pp. 823–831, 2010.

[9] R. Huang, E. Harinath, and L. T. Biegler, "Lyapunov stability of economically oriented nmpc for cyclic processes," *Journal of Process Control*, vol. 21, no. 4, pp. 501–509, 2011.

[10] K. Subramanian, J. B. Rawlings, and C. T. Maravelias, "Economic model predictive control for inventory management in supply chains," *Computers & Chemical Engineering*, vol. 64, pp. 71–80, 2014.

[11] M. J. Risbeck, C. T. Maravelias, J. B. Rawlings, and R. D. Turney, "Cost optimization of combined building heating/cooling equipment via mixed-integer linear programming," in *American Control Conference (ACC), 2015*. IEEE, 2015, pp. 1689–1694.

[12] V. M. Zavala, C. D. Laird, and L. T. Biegler, "Interior-point decomposition approaches for parallel solution of large-scale nonlinear parameter estimation problems," *Chemical Engineering Science*, vol. 63, no. 19, pp. 4834–4845, 2008.

[13] A. M. Geoffrion, "Generalized benders decomposition," *Journal of optimization theory and applications*, vol. 10, no. 4, pp. 237–260, 1972.

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

[15] Y. Shi, B. Xu, D. Wang, and B. Zhang, "Using Battery Storage for Peak Shaving and Frequency Regulation: Joint Optimization for Superlinear Gains," *IEEE Transaction on Power Systems*, pp. 1–10, 2017.

[16] R. Kumar, M. J. Wenzel, M. J. Ellis, M. N. ElBsat, K. H. Drees, and V. M. Zavala, "A stochastic model predictive control framework for stationary battery systems," *IEEE Transactions on Power Systems*, 2018.