Automatica 107 (2019) 306–316

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Hierarchical MPC schemes for periodic systems using stochastic programming ☆

Ranjeet Kumar [a], Michael J. Wenzel [b], Matthew J. Ellis [b], Mohammad N. ElBsat [b], Kirk H. Drees [b], Victor M. Zavala [a,*]

[a] Department of Chemical and Biological Engineering, University of Wisconsin—Madison 1415 Engineering Dr, Madison, WI 53706, United States
[b] Johnson Controls International, 507 E. Michigan St., Milwaukee, WI, United States

ARTICLE INFO

Article history:
Received 28 April 2018
Received in revised form 27 March 2019
Accepted 14 May 2019
Available online 7 June 2019

Keywords:
Model predictive control
Multiscale
Hierarchical
Cutting planes

ABSTRACT

We show that stochastic programming provides a framework to design hierarchical model predictive control (MPC) schemes for periodic systems. This is based on the observation that, if the state policy of an infinite-horizon problem is periodic, the problem can be cast as a stochastic program (SP). This reveals that it is possible to update periodic state targets by solving a retroactive optimization problem that progressively accumulates historical data. Moreover, we show that the retroactive problem is a statistical approximation of the SP and thus delivers optimal targets in the long run. Notably, the computation of the optimal targets can be achieved without data forecasts. The SP setting also reveals that the retroactive problem can be seen as a high-level hierarchical layer that provides targets to guide a low-level MPC controller that operates over a short period at high time resolution. We derive a retroactive scheme tailored to linear systems by using cutting plane techniques and suggest strategies to handle nonlinear systems and to analyze stability properties.

1. Introduction

A well-known challenge arising in model predictive control (MPC) is the computational complexity associated with the length of the planning horizon and with the time resolution of the state and control policies (Rawlings & Mayne, 2009). These issues are often encountered in energy system applications that exhibit phenomena and disturbances emanating at multiple timescales. For instance, in energy systems, long horizons are often required to respond to low-frequency (e.g., seasonal) supply/demand variations and peak electricity costs (e.g., demand charges) while fine time resolutions are needed to modulate high-frequency variations (e.g., from wind/solar supply) and to participate in real-time markets (Braun, 1990; Dowling, Kumar, & Zavala, 2017). Computational complexity issues are often handled using receding horizon (RH) approximations, which are practical but do not provide optimality guarantees (Risbeck, Maravelias, Rawlings, & Turney, 2017).

Hierarchical MPC schemes (Scattolini & Colaneri, 2007, Picasso, De Vito, Scattolini, & Colaneri, 2010) have been recently proposed to handle multiple scales (and associated computational complexity issues). These schemes, however, do not provide optimality guarantees in the sense that the computed policies match those of the long-horizon problem of interest. The hierarchical scheme proposed in Zavala (2016) uses adjoint information obtained from a long-term but coarse controller to guide a short-term controller operating at fine time resolutions. Computational experiments are provided to demonstrate that this approach can achieve optimality but no guarantees are given. Moreover, such an approach requires smoothness and continuity of the adjoint profiles, which is not guaranteed in general applications.

The hierarchical scheme proposed in this work relies on the observation that, if the optimal policy of an infinite horizon problem is periodic (or can be approximated with a periodic policy), the problem can be cast as a stochastic programming (SP) problem. Periodicity is a property that is commonly observed in systems driven by exogenous factors with strong periodic components (e.g., energy demands and prices) (Huang, Harinath, & Biegler, 2011; Risbeck, Maravelias, Rawlings, & Turney, 2015). Under the SP abstraction, the inter-period trajectory of the exogenous factors is interpreted as a realization of a random variable that triggers a periodic trajectory of the system states (the states at the beginning and end of the period are the same). Moreover, the periodic states are interpreted as design variables and

☆ The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Riccardo Scattolini under the direction of Editor Ian R. Petersen.

* Corresponding author.
E-mail addresses: rkumar32@wisc.edu (R. Kumar), mike.wenzel@jci.com (M.J. Wenzel), matthew.j.ellis@jci.com (M.J. Ellis), mohammad.elbsat@jci.com (M.N. ElBsat), kirk.h.drees@jci.com (K.H. Drees), victor.zavala@wisc.edu (V.M. Zavala).

https://doi.org/10.1016/j.automatica.2019.05.054

operational policies over the periods are interpreted as recourse variables. We have recently observed that the SP representation provides a mechanism to construct hierarchical MPC schemes in which a *long-term* (supervisory) MPC controller provides periodic targets to guide a short-term MPC controller (Kumar et al., 2018a). Under nominal conditions with *perfect forecasts*, we have shown that the hierarchical scheme delivers an optimal policy for the infinite horizon problem. For the more relevant case of *imperfect forecasts*, the hierarchical scheme needs to re-compute periodic targets. While this can certainly be done using an RH scheme (e.g., computes targets by anticipating multiple future periods), such an approach would not provide optimality guarantees. In fact, to the best of our knowledge, no RH scheme currently exists that can provide optimal policies in the presence of imperfect forecast information. Specifically, standard proactive RH schemes use historical data to compute forecasts and associated control actions. A fundamental issue with proactive approaches is that no optimality guarantees can be provided unless the forecast is perfect.

The contribution of this work is the observation that, under a periodic setting, one can derive *retroactive* hierarchical MPC schemes that accumulate real historical data to *asymptotically deliver optimal targets*. The retroactive design principle thus offers optimality guarantees and, notably, does not require data forecasts. The retroactive approach thus provides key advantages over proactive RH schemes. The targets obtained with the retroactive scheme are used to guide a low-level controller operating at fine time resolutions within the periods. In the case of linear systems, one can derive a specialized retroactive scheme by using *incremental cutting-plane (CP) techniques* (Higle & Sen, 1991). The SP setting also reveals strategies to construct retroactive schemes for nonlinear systems and to obtain the desired stability properties. We demonstrate the concepts using a battery application and compare the proposed retroactive hierarchical MPC scheme with a proactive MPC approach for periodic systems.

The paper is structured as follows. In Section 2, we provide basic definitions and describe the problem setting. In Section 3, we introduce the concept of retroactive optimization, derive an incremental CP scheme for linear systems, discuss implementation details, and discuss extensions for nonlinear systems. Computational experiments are presented in Section 4.

2. Basic definitions and setting

In this work, we derive schemes to compute approximate solutions for the long-horizon problem $\mathbf{O}_m$:

$$\min_{u_\xi,x_\xi,\eta,x_0} \frac{1}{m}\sum_{\xi\in\Xi}\sum_{t\in\mathcal{T}}\varphi_1(x_{\xi,t},u_{\xi,t},d_{\xi,t})+\eta$$

$$\text{s.t. } \varphi_2(x_{\xi,t},u_{\xi,t},d_{\xi,t})\le\eta,\ \xi\in\Xi,t\in\mathcal{T} \tag{1a}$$

$$x_{\xi,t+1}=f(x_{\xi,t},u_{\xi,t},d_{\xi,t}),\ \xi\in\Xi,t\in\bar{\mathcal{T}} \tag{1b}$$

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

$$x_{1,0}=x_0 \tag{1d}$$

$$x_{\xi,t}\in\mathcal{X},u_{\xi,t}\in\mathcal{U}. \tag{1e}$$

Here, the horizon with $p:=(n+1)\cdot m$ time steps is partitioned into a set of *periods* $\Xi:=\{1,\dots,m\}$ with intra-period times $\mathcal{T}:=\{0,\dots,n\}$. For convenience, we define the sets $\bar{\mathcal{T}}:=\mathcal{T}\setminus\{n\}$ and $\bar{\Xi}:=\Xi\setminus\{m\}$. The controls, states, and data trajectories at period $\xi\in\Xi$ and intra-period time $t\in\mathcal{T}$ are denoted as $u_{\xi,t}\in\mathbb{R}^{n_u}, x_{\xi,t}\in\mathbb{R}^{n_x}$, and $d_{\xi,t}\in\mathbb{R}^{n_d}$, respectively. We define the notation $u_\xi:=\{u_{\xi,t}\}_{t\in\mathcal{T}}$, $x_\xi:=\{x_{\xi,t}\}_{t\in\mathcal{T}}$, and $d_\xi:=\{d_{\xi,t}\}_{t\in\mathcal{T}}$ to denote the inner period trajectories. The problem *data* is interpreted as exogenous disturbances or factors driving the system (e.g., market prices, demands, weather, model errors).

The mapping $\varphi_1:\mathbb{R}^{n_u}\times\mathbb{R}^{n_x}\times\mathbb{R}^{n_d}\to\mathbb{R}$ is a time-additive cost function and the mapping $\varphi_2:\mathbb{R}^{n_u}\times\mathbb{R}^{n_x}\times\mathbb{R}^{n_d}\to\mathbb{R}$ is a time-max (peak) cost function. We assume both of these functions to be bounded in their domains. Minimizing the variable $\eta\in\mathbb{R}$ subject to the constraints (1a) is equivalent to minimize the peak cost $\max_{\xi\in\Xi}\max_{t\in\mathcal{T}}\varphi_2(x_{\xi,t},u_{\xi,t},d_{\xi,t})$. Consequently, since $\varphi_2(\cdot)$ is assumed to be bounded, we have that η is bounded as well. The mapping $f:\mathbb{R}^{n_u}\times\mathbb{R}^{n_x}\times\mathbb{R}^{n_d}\to\mathbb{R}^{n_x}$ describes the system dynamics and $\mathcal{X}$ and $\mathcal{U}$ are non-empty feasible sets for states and controls, respectively. For reasons that will become apparent, the initial state $x_0\in\mathbb{R}^{n_x}$ in the above formulation is treated as a decision variable.

The infinite-horizon problem $\mathbf{O}_\infty$ is obtained by setting $\lim_{m\to\infty}$. If the data d_ξ, $\xi\in\Xi$ is known, the solution of $\mathbf{O}_\infty$ provides a policy with the *best possible performance*. Unfortunately, problem $\mathbf{O}_m$ becomes computationally difficult to solve for large m (long horizons) and/or for large n (fine time resolutions). An approximate solution for $\mathbf{O}_\infty$ is often computed by using a *proactive RH scheme*, which approximates the policy by moving forward in time and by planning over short time horizons (e.g., that span a few periods). Our goal is to derive an *alternative* strategy that uses hierarchical MPC schemes that approximate the solution of $\mathbf{O}_\infty$ by using a *periodic policy*. We define a periodic policy as follows:

Definition 2.1. A *policy is said to be periodic if it satisfies the* periodicity constraints $x_{\xi,0}=x_{\xi,n}$ for all $\xi\in\Xi$.

To obtain a periodic policy, we enforce periodicity constraints at every period (after every n steps). These constraints, together with the continuity constraints (1c), can be expressed as $x_{\xi+1,0}=x_{\xi,0},\ \xi\in\bar{\Xi}$. This set of constraints, in turn, can be reformulated as $x_{\xi,0}=x_0,\ \xi\in\Xi$. These modifications give the periodic problem $\mathbf{P}_m$:

$$\min_{u_\xi,x_\xi,\eta,x_0} \frac{1}{m}\sum_{\xi\in\Xi}\sum_{t\in\mathcal{T}}\varphi_1(x_{\xi,t},u_{\xi,t},d_{\xi,t})+\eta \tag{2a}$$

$$\text{s.t. } \varphi_2(x_{\xi,t},u_{\xi,t},d_{\xi,t})\le\eta,\ \xi\in\Xi,t\in\mathcal{T} \tag{2b}$$

$$x_{\xi,t+1}=f(x_{\xi,t},u_{\xi,t},d_{\xi,t}),\ \xi\in\Xi,t\in\bar{\mathcal{T}} \tag{2c}$$

$$x_{\xi,0}=x_0,\ \xi\in\Xi \tag{2d}$$

$$x_{\xi,n}=x_0,\ \xi\in\Xi \tag{2e}$$

$$x_{\xi,t}\in\mathcal{X},u_{\xi,t}\in\mathcal{U}. \tag{2f}$$

Remark. The feasible region of $\mathbf{P}_m$ is smaller than that of $\mathbf{O}_m$ (since the latter does not enforce periodicity every n steps). Consequently, the performance of $\mathbf{P}_m$ is expected to be inferior to that of $\mathbf{O}_m$. In some applications, however, the deterioration of performance might not be significant. For instance, in Kumar et al. (2018a), it is shown that enforcing state periodicity constraints for a battery system (obtained from $\mathbf{P}_m$) results in a policy with a cost that is 0.2% larger than the cost achieved without periodicity constraints (obtained from $\mathbf{O}_m$).

The solution of $\mathbf{P}_m$ provides a better approximation to the solution of $\mathbf{O}_m$ as we increase n (increasing the size of the period). This is because periodicity constraints are enforced less often. For a fixed value horizon p, we have that in the limit when $n=p$ and $m=1$, we obtain the best possible periodic policy (enforcing periodicity at the beginning and at the end of the horizon). On the other hand, in the limit when $n=1$ and $m=p$, we have the worst possible periodic policy (a steady-state policy). Setting $n=p$ and $m=1$ and eliminating all the periodicity constraints make $\mathbf{P}_m$ and $\mathbf{O}_m$ equivalent. From these observations, we note that the length of the period n can be used as a design parameter to find a periodic policy of $\mathbf{P}_m$ that properly approximates the policy of $\mathbf{O}_m$.

The goal of $\mathbf{P}_m$ is to find a periodic state x_0, peak cost η, as well as intra-period policies u_ξ, x_ξ $\xi \in \Xi$ that minimize the time-additive and peak costs. We note that, by construction, $x_{\xi+1,0} = x_{\xi,0} = x_0$ holds for all $\xi \in \Xi$ at any solution of $\mathbf{P}_m$. If we have perfect knowledge of the data $d_\xi, \xi \in \Xi$, the infinite-horizon problem $\mathbf{P}_\infty$ identifies a periodic policy with the best possible performance. We denote a solution of $\mathbf{P}_\infty$ as $w^* = (x_0^*, \eta^*)$, and u_ξ^*, x_ξ^*. Unfortunately, problem $\mathbf{P}_m$ also becomes computationally difficult to solve for large m and/or n. Here, one can address this issue by using a proactive RH scheme with periodicity constraints, as proposed in Huang et al. (2011). We will see, however, that periodicity results in a structure that enables the derivation of *retroactive schemes* that offer key advantages over proactive schemes. In particular, we will see that retroactive schemes can deliver a solution of $\mathbf{O}_\infty$ (while proactive schemes cannot).

By analyzing the structure of $\mathbf{P}_m$, we can see that the only coupling between periods arises from the variables x_0 and η. Consequently, by fixing these variables, we can decompose (2) into individual *period subproblems* of the form:

$$\min_{u_\xi, x_\xi} \sum_{t\in\mathcal{T}} \varphi_1(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) \tag{3a}$$

$$\text{s.t. } \varphi_2(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) \le \eta,\ t \in \mathcal{T} \tag{3b}$$

$$x_{\xi,t+1} = f(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}),\ t \in \bar{\mathcal{T}} \tag{3c}$$

$$x_{\xi,0} = x_0 \tag{3d}$$

$$x_{\xi,n} = x_0 \tag{3e}$$

$$x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \tag{3f}$$

We define this problem as $\mathbf{S}_\xi$ and we define its optimal cost as $h(w, d_\xi)$, where $w := (x_0, \eta) \in \mathcal{W} := \mathcal{X} \times \mathbb{R}$. In the following discussion, we use the short-hand notation $h_\xi := h(w, d_\xi)$.

We now make the following key assumption regarding the nature of the data and of the subproblem cost $h(w, d_\xi)$.

Assumption 2.1. Assume that $\{d_\xi\}_{\xi=1}^\infty$ is a sequence of independent and identically distributed (i.i.d.) realizations of a continuous random variable D with associated probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with associated codomain $\Omega' \subseteq \mathbb{R}^{n_d}$ and σ-algebra $\mathcal{F}'$. Moreover, assume that the cost function $h : \mathcal{W} \times \Omega' \to \mathbb{R}$ is continuous and bounded in its domain.

Under the assumption of i.i.d. realizations and a bounded function $h(\cdot)$, we can establish that the sample average $\frac{1}{m}\sum_{\xi=1}^m h(w, d_\xi)$ converges pointwise with probability one to $\mathbb{E}[h(w, D)]$ as $m \to \infty$. Here, $\mathbb{E}[\cdot]$ is the expectation operator. This property is the strong law of large numbers (LLN) and is key because it reveals that the infinite horizon problem $\mathbf{P}_\infty$ can be interpreted as a *stochastic programming (SP) problem* of the form:

$$\min_{w\in\mathcal{W}} \ \phi(w) := g(w) + \mathbb{E}\left[h(w, D)\right]. \tag{4}$$

Under this representation, periods are realizations d_ξ of D, $w = (x_0, \eta)$ are *design* (target) variables, and x_ξ, u_ξ are (recourse) policies associated with realization d_ξ. Here, $g : \mathcal{W} \to \mathbb{R}$ is a cost function for the design variables. In the context of $\mathbf{P}_\infty$, we have $g(w) = \eta$, but this function can be generalized to enforce a cost also on x_0. Function $g(\cdot)$ is bounded because $\phi_2(\cdot)$ (and thus η) are assumed to be bounded. Since (4) is equivalent to $\mathbf{P}_\infty$, it delivers optimal targets w^*. We define the solution set of (4) as $S \subseteq \mathcal{W}$.

Remark. The i.i.d. requirement on $\{d_\xi\}$ ensures that LLN holds. This requirement can be enforced by defining a sufficiently long period duration n (to eliminate autocorrelation between periods). This approach is commonly used in statistical extrapolation and time series analysis (Box, Jenkins, Reinsel, & Ljung, 2015; Ragan & Manuel, 2008). The independence requirement can also be relaxed by allowing for cost sequences that have bounded correlation (Hu, Rosalsky, & Volodin, 2008). To give an idea of why this is the case, we provide the following result.

Property 2.1. *Assume that $\{h_\xi\}_{-\infty}^\infty$ is a sequence of identically distributed random variables with expected value $\mu = \mathbb{E}[h(x, D)]$. Assume also that there exists $0 \le c < \infty$ such that $\sum_{\xi=-\infty}^\infty \mathbb{E}[(h_{\xi'} - \mu)(h_\xi - \mu)] \le c$ holds for all $\xi' = 1, \dots, \infty$. Then, $\lim_{m\to\infty} \frac{1}{m}\sum_{\xi=1}^m h_\xi = \mathbb{E}[h(x, D)]$.*

Proof. Define $S_m := \sum_{\xi=1}^m h_\xi$, $\bar{h}_m := S_m/m$, and $\mu := \mathbb{E}[h(x, D)]$. The variance of S_m (denoted as $V[S_m]$) is:

$$\begin{aligned}\mathbb{E}[(S_m - m\,\mu)^2] &= \sum_{\xi'=1}^m \sum_{\xi=1}^m \mathbb{E}[(h_{\xi'} - \mu)(h_\xi - \mu)] \\ &\le \sum_{\xi'=1}^m \sum_{\xi=-\infty}^\infty \mathbb{E}[(h_{\xi'} - \mu)(h_\xi - \mu)] \\ &\le m\,c.\end{aligned}$$

We thus have $V[\bar{h}_m] \le c/m$ and, from Chebyshev's inequality (with parameter $\kappa > 0$),

$$\begin{aligned}\mathbb{P}(|\bar{h}_m - \mu| > \kappa) &\le \frac{1}{\kappa^2}\mathbb{E}[(\bar{h}_m - \mu)^2] \\ &= \frac{1}{\kappa^2 m^2}\mathbb{E}[(S_m - m\mu)^2] \\ &\le \frac{m\,c}{m^2\kappa^2} = \frac{c}{m\,\kappa^2}.\end{aligned}$$

We thus have $\lim_{m\to\infty} \bar{h}_m = \mu$ with probability one. □

Remark. We can guarantee that a solution for $\mathbf{P}_m$ exists if the period length n is long enough such that one can find a control policy u_ξ that delivers a feasible solution for $\mathbf{S}_\xi$ for all fixed $w \in \mathcal{W}$ and $d_\xi, \xi \in \Xi$. Here, a feasible solution of $\mathbf{S}_\xi$ is one that satisfies the control and state constraints as well as the periodicity and peak constraints. This assumption is compatible with that used in the standard MPC literature (a sufficiently long horizon is assumed such that one can find a feasible control policy that satisfies the state and terminal constraints) (Subramanian, Rawlings, & Maravelias, 2014; Zanon, Grüne, & Diehl, 2017). In the SP literature, such an assumption is equivalent to assuming that $\mathbf{S}_\xi$ has *relatively complete recourse* (i.e., $\mathbf{S}_\xi$ has a feasible solution for any $w \in \mathcal{W}$ and $d_\xi \in \Omega'$).

3. Hierarchical MPC schemes

The SP representation opens a number of interesting directions. In particular, it provides a mechanism to derive *hierarchical MPC schemes*. For instance, one can use a statistical approximation of (4) (equivalently of $\mathbf{P}_\infty$) to provide targets x_0, η that guide a short-term MPC controller of the form (3). As we discuss next, approximations of $\mathbf{P}_\infty$ can be constructed and solved in a tractable manner by using well-established SP techniques (CarøE & Schultz, 1999; Geoffrion, 1972; Zavala, Laird, & Biegler, 2008).

3.1. Retroactive optimization

A key observation that arises from the SP representation is that one can derive *retroactive optimization* schemes that accumulate data over time to refine targets. To explain how such a scheme would work, assume that the system is currently at the beginning of period $m+1$ and that the data history $\{d_\xi\}_{\xi=1}^m$ is known. We use this information to solve the problem:

$$\min_{w\in\mathcal{W}} \ \phi_m(w) := g(w) + \frac{1}{m}\sum_{\xi=1}^{m} h(w,d_\xi). \tag{5}$$

This problem is equivalent to $\mathbf{P}_m$, and because $\{d_\xi\}_{\xi=1}^m$ is i.i.d., $\mathbf{P}_m$ is a statistical approximation of $\mathbf{P}_\infty$. In the SP literature, problem (5) is known as a sample average approximation (SAA). We define the solution set of $\mathbf{P}_m$ as $S_m \subseteq \mathcal{W}$. A solution of $\mathbf{P}_m$ is used to update the targets for the next period $w_{m+1} = (x_{m+1,0}, \eta_{m+1})$. A solution of $\mathbf{P}_m$ also implicitly contains optimal (retroactive) policies $u_\xi, x_\xi,\ \xi = 1,\ldots,m$ associated with the historical sequence $\{d_\xi\}_{\xi=1}^m$. These retroactive policies are interpreted as policies that the system *would have taken* given knowledge of the data.

Given that the system is at the current state target $x_{m,0}$ (obtained from the solution of $\mathbf{P}_{m-1}$), we use the targets w_{m+1} to guide a *short-term MPC controller* over period $m+1$. At this point, however, the data d_{m+1} is not known, so we use a *forecast* $\hat{d}_{m+1}$ to find policies that optimize the system over the next period $m+1$ while satisfying the targets w_{m+1}. This can be interpreted as solving $\mathbf{S}_\xi$ for w_{m+1} given $\hat{d}_{m+1}$. The forecast $\hat{d}_{m+1}$ is typically obtained by using forecasting techniques such as AR, ARMA, or ARIMA time series models, or covariance estimators (Kumar et al., 2018b). At the beginning of the next period $m+2$, the actual data realization d_{m+1} reveals itself and we use this to solve the approximation $\mathbf{P}_{m+1}$ to obtain new targets w_{m+2}.

The retroactive scheme is consistent because, from the law of large numbers, we know that accumulating data over time will yield an asymptotically exact statistical approximation $\lim_{m\to\infty} \phi_m(w) = \phi(w)$ and thus the targets obtained with $\mathbf{P}_m$ will provide a solution to $\mathbf{P}_\infty$ as $m\to\infty$. This asymptotic convergence result is formally stated in the following theorem (see Theorem 5.3 in Shapiro, Dentcheva, and Ruszczyński (2009)).

Theorem 3.1. *Suppose that there exists a compact set $C \subset \mathbb{R}^{n_w}$ such that: (i) the solution set S of $\mathbf{P}_\infty$ is nonempty and contained in C, (ii) the function $\phi(w)$ is finite-valued and continuous on C, (iii) $\phi_m(w)$ converges to $\phi(w)$ with probability one as $m\to\infty$ uniformly in $w\in C$, and with probability one for large enough m the solution set S_m of $\mathbf{P}_m$ is nonempty and contained in C. Then, the Hausdorff distance between the solution sets $\mathbb{D}(S_m, S)$ converges to zero with probability one as $m\to\infty$.*

The above result implies that the distance of any solution of $\mathbf{P}_m$ to the solution set of $\mathbf{P}_\infty$ converges to zero as $m\to\infty$. Statistical approximation results for SPs also indicate that the probability that a solution of $\mathbf{P}_m$ is in the solution set of $\mathbf{P}_\infty$ increases exponentially with m (Theorem 5.16 in Shapiro et al. (2009)). In other words, the *probability of finding better targets than those obtained with $\mathbf{P}_m$ decays exponentially as information is accumulated over time.*

These asymptotic optimality results provide a key advantage of the retroactive scheme over traditional proactive RH schemes. This is based on a fundamental design difference: the retroactive scheme uses past (but real) data while proactive RH schemes use future (but approximate) data forecasts. Moreover, proactive schemes discard historical data when computing new targets. The fact that historical data is discarded prevents RH schemes from offering asymptotic optimality guarantees.

3.2. Incremental cutting plane scheme

The structure of $\mathbf{P}_m$ can be exploited using decomposition strategies and this enables scalability to large values of m. In this section, we provide a decomposition scheme for linear systems, and we then discuss potential extensions to nonlinear systems.

The retroactive scheme for linear systems proposed is based on an *incremental cutting plane* (CP) scheme. Our approach is an adaptation of the stochastic decomposition scheme proposed in Higle and Sen (1991) to tackle linear SPs. To derive this linear setting, we assume that $g(w) = c_w^T w$ where $c_w \in \mathbb{R}^{n_w}$ is a cost vector, we assume that the set $\mathcal{W} \subseteq \mathbb{R}^{n_w}$ is polyhedral, and we assume that $\mathbf{S}_\xi$ has the form:

$$h(w,d_\xi) := \min_{y\in\mathbb{R}_+^{n_y}} \ c_\xi^T y \quad \text{s.t.} \quad Wy = r_\xi - Tw. \tag{6}$$

Here, the data realization is given by $d_\xi = (c_\xi, r_\xi)$. We use $y(w,d_\xi)$ to denote the primal solution vector containing the intra-period trajectories (x_ξ, u_ξ) and some additional dummy variables. The intra-period dynamics are captured using W, T that are coefficient matrices. The structure of the recourse problem is used to simplify algebraic manipulations and is done without loss of generality (the results that we present hold provided that the recourse problem is a linear program). The representation of $\mathbf{S}_\xi$ allows us to express its dual form in the following compact form:

$$\max_\pi \ \pi^T(r_\xi - Tw) \quad \text{s.t.} \quad W^T\pi \le c_\xi. \tag{7}$$

Here, $\pi(w,d_\xi)$ is a dual solution vector (dual variables of $\mathbf{S}_\xi$) and we recall that $\pi(w,d_\xi)^T(r_\xi - Tw) = h(w,d_\xi)$ holds for $w\in\mathcal{W}$ (from strong duality). We assume that the feasible set of the dual subproblem is a non-empty, compact, and convex polyhedral set, and therefore the polyhedron represented by the set $\mathcal{P} := \{\pi \mid W^T\pi \le c_\xi\}$ is a pointed polyhedron for all $(w,d_\xi)\in\mathcal{W}\times\Omega'$. As a result, $\mathbf{S}_\xi$ has a finite number of dual vertices or extreme points (where a vertex of the polyhedron $\mathcal{P}$ is a vector $\pi\in\mathcal{P}$ such that we cannot find two distinct vectors $\pi_1,\pi_2\in\mathcal{P}$ and a scalar $\lambda\in[0,1]$, such that $\pi = \lambda\pi_1 + (1-\lambda)\pi_2$ (see Chapter 2 of Bertsimas and Tsitsiklis (1997)). Moreover, because the support Ω is finite, the set of dual vertices for all subproblems $\mathbf{S}_\xi$ is finite (see Theorem 2.9 in Chapter 2 of Bertsimas and Tsitsiklis (1997)). We use $\mathbb{V}$ to denote the set of all these dual vertices. Consequently, by definition, $\pi(w,d_\xi)\in\mathbb{V}$ for all $(w,d_\xi)\in\mathcal{W}\times\Omega'$.

In the linear setting, the cost function of $\mathbf{P}_\infty$ (given by $\phi(w) = c_w^T w + \mathbb{E}[h(w,D)]$) can be outer-approximated using CPs accumulated over $\xi = 1,\ldots,m$ as:

$$\underline{\phi}_m(w) := \max\{\alpha_\xi^m + (c_w + \beta_\xi^m)^T w \mid \xi = 1,\ldots,m\}, \tag{8}$$

where the coefficients $\alpha_\xi^m, \beta_\xi^m$ are selected to match:

$$\alpha_\xi^m + (\beta_\xi^m)^T w = \frac{1}{m}\sum_{\xi=1}^{m} (\pi_\xi^m)^T (r_\xi - Tw). \tag{9}$$

Here, $\pi_\xi^m \in \operatorname{argmax}\{\pi^T(r_\xi - Tw_m) \mid \pi\in\mathbb{V}_m\}$ for $\xi = 1,\ldots,m$, where $\mathbb{V}_m \subseteq \mathbb{V}$ is the collection of vertices accumulated up to period m and $w_m = (x_{m,0}, \eta_m)$. For convenience, we define the function $h_m(w,d_\xi) := \max\{\pi^T(r_\xi - Tw) \mid \pi\in\mathbb{V}_m\}$ and note that $h_m(w,d_\xi) = (\pi_\xi^m)^T(r_\xi - Tw)$ holds. Moreover, we note that $h(w,d_\xi) = \max\{\pi^T(r_\xi - Tw) \mid \pi\in\mathbb{V}\}$ holds. Since $\mathbb{V}_m \subseteq \mathbb{V}$, we have that $(\pi_\xi^m)^T(r_\xi - Tw) \le h(w,d_\xi)$.

The *running cost* is given by $\phi_m(w) = c_w^T w + \frac{1}{m}\sum_{\xi=1}^m h(w,d_\xi)$. We will prove that $\underline{\phi}_m(w)$ underestimates the running cost $\phi_m(w)$ for all m and converges to the infinite-horizon cost $\phi(w)$ as $m\to\infty$. Consequently, at period m, we update the targets by solving the *master* problem $\mathbf{M}_m$:

$$\min_{w\in\mathcal{W}} \ \underline{\phi}_m(w). \tag{10}$$

This problem is a tractable surrogate of $\min_{w\in\mathcal{W}} \ \phi_m(w)$, because it captures the recourse subproblems by using hyperplanes. This becomes particularly important as information is accumulated over time. The solution of $\mathbf{M}_m$ is used to update the targets w_{m+1}, which in turn are used to solve the recourse subproblem $S_{\xi+1}$ and with this, obtain a new dual vertex to be stored in $\mathbb{V}_{m+1}$. The CP scheme is summarized as:

(1) Initialize $m \leftarrow 1$, $\mathbb{V}_m \leftarrow \emptyset$, and w_m.
(2) At period time $m+1$:
(3) Observe d_m and solve $\mathbf{S_m}$ to obtain $\pi(w_m, d_m)$.
(4) Update $\mathbb{V}_m \leftarrow \mathbb{V}_{m-1} \cup \{\pi(w_m, d_m)\}$.
(5) (a) Obtain $\pi_\xi^m \in \operatorname{argmax}\{\pi^T(r_\xi - Tw_m) | \pi \in \mathbb{V}_m\}$ for all $\xi = 1, \ldots, m$.
 (b) Get α_m^m and β_m^m from (9).
 (c) Update $\alpha_\xi^m \leftarrow \frac{m-1}{m}\alpha_\xi^{m-1}$, $\beta_\xi^m \leftarrow \frac{m-1}{m}\beta_\xi^{m-1}$ for $\xi = 1, \ldots, m-1$.
(6) Solve $\mathbf{M}_m$ and obtain updated targets w_{m+1}.
(7) Shift period time $m \leftarrow m+1$ and return to Step 2.

We now prove that the CP scheme delivers a sequence of targets $\{w_m\}_{m=1}^\infty$ that converges to optimal targets w^* of $\mathbf{P}_\infty$. Our analysis follows along the lines of that presented in Higle and Sen (1991).

Theorem 3.2. *The CP $\alpha_m^m + (c_w + \beta_m^m)^T w_m$ generated in period m provides a statistically valid lower bound (statistically based estimate of a lower bound) for $\phi(w)$ for all $w \in \mathcal{W}$.*

Proof. Because $\mathbb{V}_m \subseteq \mathbb{V}$ we have that,

$$\max\{\pi^T(r_\xi - Tw_m) | \pi \in \mathbb{V}_m\} \le \max\{\pi^T(r_\xi - Tw_m) | \pi \in \mathbb{V}\},$$

and $(\pi_\xi^m)^T(r_\xi - Tw_m) = h(w_m, d_\xi),\ \xi = 1, \ldots, m$. We thus have

$$\frac{1}{m}\sum_{\xi=1}^m (\pi_\xi^m)^T(r_\xi - Tw_m) \le \frac{1}{m}\sum_{\xi=1}^m h(w_m, d_\xi).$$

Furthermore, $\pi^T(r_\xi - Tw) \le h(w, d_\xi)$ for any $\pi \in \mathbb{V}$ and:

$$c_w^T w + \frac{1}{m}\sum_{\xi=1}^m (\pi_\xi^m)^T(r_\xi - Tw) \le c_w^T w + \frac{1}{m}\sum_{\xi=1}^m h(w, d_\xi) = \phi_m(w),\ \ w \in \mathcal{W}.$$

The result follows from (9) and by noticing that $\phi_m(w)$ is a statistical approximation of $\phi(w)$. □

As more observations d_ξ are collected, it is important that all the collected CPs provide a statistically valid lower bound for $\phi(w)$. This is the goal of Step 5b in the CP scheme. In particular, at period $m+p$ with $p > 0$:

$$\frac{1}{m+p}\sum_{\xi=1}^m (\pi_\xi^m)^T(r_\xi - Tw) \le \frac{1}{m+p}\sum_{\xi=1}^m h(w, d_\xi) \le \frac{1}{m+p}\sum_{\xi=1}^{m+p} h(w, d_\xi)$$

Thus, in the $(m+p)^{th}$ period, the CP

$$c_w^T w + \frac{1}{m+p}\sum_{\xi=1}^m (\pi_\xi^m)^T(r_\xi - Tw) = \alpha_m^{(m+p)} + (\beta_m^{(m+p)} + c_w)^T w$$

still provides a statistically valid lower bound for $\phi(w)$.

We now explore the limiting behavior of $h_m(\cdot)$, which embeds all CP information accumulated over time.

Lemma 3.1. *The sequence $\{h_m(\cdot)\}_{m=1}^\infty$ converges uniformly on $\mathcal{W}$.*

Proof. $\mathbb{V}_m \subseteq \mathbb{V}_{m+1} \subseteq \mathbb{V}$ implies that $h_m(w, d_\xi) \le h_{m+1}(w, d_\xi) \le h(w, d_\xi)$ for all $w \in \mathcal{W}$ and d_ξ. Since $\{h_m(\cdot)\}_{k=1}^\infty$ increases monotonically and is bounded by the finite function $h(\cdot)$, it follows that $\{h_m(\cdot)\}_{k=1}^\infty$ converges point-wise to some function $\varphi(\cdot)$ satisfying $\varphi(w, d_\xi) \le h(w, d_\xi)$ for all $w \in \mathcal{W}$, $d_\xi \in \Omega'$. Since $\mathbb{V}_m \subseteq \mathbb{V}_{m+1} \subseteq \mathbb{V}$, we have $\bar{\mathbb{V}} = \lim_{m\to\infty} \mathbb{V}_m \subseteq \mathbb{V}$. Since $\mathbb{V}$ is a finite set, so $\bar{\mathbb{V}}$ is also a finite set, and we have that:

$$\begin{aligned}\varphi(w, d_\xi) &= \lim_{m\to\infty} h_m(w, d_\xi)\\ &= \lim_{m\to\infty} \max\{\pi^T(r_\xi - Tw) | \pi \in \mathbb{V}_m\}\\ &= \max\{\pi^T(r_\xi - Tw) | \pi \in \bar{\mathbb{V}}\},\end{aligned}$$

and it follows that $\varphi(\cdot)$ is a continuous function. As $\mathcal{W} \times \Omega'$ is a compact space, and $\{h_m(\cdot)\}_{m=1}^\infty$ is a monotonic sequence of continuous functions, it follows that it converges uniformly to $\varphi(\cdot)$ (see Theorem 7.13 in Rudin et al. (1964)). □

We now show that the sequence of CPs generated provides support points for $\phi(\cdot)$ in the limit $m \to \infty$.

Theorem 3.3. *Let $\{w_{m_k}\}_{k=1}^\infty$ be an infinite subsequence of $\{w_m\}_{m=1}^\infty$. If $w_{m_k} \to \bar{w}$ then, with probability one,*

$$\lim_{k\to\infty} \frac{1}{m_k}\sum_{\xi=1}^{m_k} (\pi_\xi^{m_k})^T(r_\xi - Tw_{m_k}) = \mathbb{E}[h(\bar{w}, D)].$$

In addition, every limit of $\{\alpha_{m_k}^{m_k}, \beta_{m_k}^{m_k} + c_w\}_{k=1}^\infty$ defines a support of $\phi(w)$ at $\bar{w}$, with probability one.

Proof. By definition, we have that

$$h_{m_k}(w_{m_k}, d_\xi) = (\pi_\xi^{m_k})^T(r_\xi - Tw_{m_k})$$

$$\frac{1}{m_k}\sum_{\xi=1}^{m_k} h_{m_k}(w_{m_k}, d_\xi) = \frac{1}{m_k}\sum_{\xi=1}^{m_k} (\pi_\xi^{m_k})^T(r_\xi - Tw_{m_k}).$$

By Lemma 3.1, there exists $\varphi(\cdot) \le h(\cdot)$ such that $\{h_m\}_{m=0}^\infty$ converges uniformly to $\varphi(\cdot)$. We thus have:

$$\lim_{k\to\infty} \frac{1}{m_k}\sum_{\xi=1}^{m_k} \left[h_{m_k}(w_{m_k}, d_\xi) - \varphi(\bar{w}, d_\xi)\right] = 0,$$

and $\lim_{k\to\infty} \frac{1}{m_k}\sum_{\xi=1}^{m_k} h(w, d_\xi) = \mathbb{E}[h(w, D)]$ with probability one. It is now sufficient to show that $\varphi(\bar{w}, d_\xi) = h(\bar{w}, d_\xi)$ with probability one. Let d_ξ be a given realization and suppose that, for every $\delta > 0$, we have $\mathbb{P}\{|D - d_\xi| < \delta\}$. Then, for every $\delta > 0$, $|d_{m_k} - d_\xi| < \delta$ infinitely often, with probability one. Because $h(\cdot)$ is a continuous function and $\{h_m(\cdot)\}_{m=1}^\infty$ is uniformly convergent, for every $\epsilon > 0$ there exist a $\delta > 0$ and $N < \infty$ such that $|(\bar{w}, d_\xi) - (w, d)| < \delta$ and with:

$$|h(\bar{w}, d_\xi) - h(\bar{w}, d)| < \frac{\epsilon}{3}$$

$$|h_{m_k}(\bar{w}, d_\xi) - h_{m_k}(\bar{w}, d)| < \frac{\epsilon}{3}, \quad k \ge N.$$

Consequently, because $\lim_{k\to\infty} w_{m_k} = \bar{w}$, we have $\mathbb{P}\{|D - d_\xi| < \delta\}$ implies that for every $\epsilon > 0$ there exists a further subsequence $\{(w_{m_k'}, d_{m_k'})\}_{k=1}^\infty$ and $K < \infty$ such that

$$|h(\bar{w}, d_\xi) - h_{m_k'}(\bar{w}, d_{m_k'})| < \frac{\epsilon}{3}$$

$$|h(\bar{w}, d_{m_k'}) - h(w_{m_k'}, d_{m_k'})| < \frac{\epsilon}{3}$$

$$|h_{m_k'}(w_{m_k'}, d_{m_k'}) - h_{m_k'}(w_{m_k'}, d_\xi)| < \frac{\epsilon}{3}$$

for all $m'_k \geq K$. By construction, $h_{m'_k}(w_{m'_k}, d_{m'_k}) = h(w_{m'_k}, d_{m'_k})$. Thus, for every $\epsilon > 0$, there exists a subsequence $\{w_{m'_k}\}_{k=1}^{\infty}$ and $K < \infty$ such that

$$|h(\bar{w}, d_\xi) - h_{m'_k}(w_{m'_k}, d_\xi)| < \epsilon$$

for all $m'_k \geq K$. Consequently, it follows that $\varphi(\bar{w}, d_\xi) = h(\bar{w}, d_\xi)$. Finally, since Ω' is compact, we have that $\mathbb{P}\{|D - d_\xi| < \delta\}$ for some $\delta > 0$ and for only finitely many values of d_ξ, with probability one. Thus, with probability one, $\varphi(\bar{w}, d_\xi) = h(\bar{w}, d_\xi)$ for all but a finite number of realizations. Hence, with probability one,

$$\lim_{k\to\infty} \frac{1}{m_k} \sum_{\xi=1}^{m_k} (\pi_\xi^{m_k})^T (r_\xi - T w_{m_k}) = \mathbb{E}[h(w, D)].$$

Moreover, since $h(w, d_\xi) = \max\{\pi^T(r_\xi - Tw) | \pi \in \mathbb{V}\}$ and $\mathbb{V}_m \subset \mathbb{V}$ for all m, it follows that

$$c_w^T w + \frac{1}{m_k} \sum_{\xi=1}^{m_k} h(w, d_\xi) \geq c_w^T w + \frac{1}{m_k} \sum_{\xi=1}^{m_k} \pi_\xi^{m_k^T}(r_\xi - Tw)$$
$$= \alpha_{m_k}^{m_k} + \left(\beta_{m_k}^{m_k} + c_w\right)^T w.$$

We conclude that, with probability one, $\phi(w)$ is at least as large as any limiting cut that is associated with the subsequence of cuts defined by $\{(\alpha_{m_k}^{m_k}, \beta_{m_k}^{m_k} + c)\}_{k=1}^{\infty}$. Thus, any limiting cut defines a support of $\phi(w)$ at $\bar{w}$. □

Theorem 3.4. *There exists a subsequence of* $\{w_{m_k}\}_{k=1}^{\infty}$, *satisfying* $\lim_{k\to\infty}(\underline{\phi}_{m_k}(w_{m_k}) - \underline{\phi}_{m_k-1}(w_{m_k})) = 0$.

Proof. See proof of Theorem 3 in Higle and Sen (1991).

We now establish the main convergence result.

Theorem 3.5. *There exists a subsequence* $\{w_{m_k}\}_{k=1}^{\infty}$, *such that every accumulation point of* $\{w_{m_k}\}_{k=1}^{\infty}$ *is an optimal solution* w^* *of* $\mathbf{P}_\infty$, *with probability one.*

Proof. Let w^* be an optimal solution of $\mathbf{P}_\infty$. From Theorem 3.4, there exists a subsequence $\{w_{m_k}\}_{k=1}^{\infty}$ such that $\lim_{k\to\infty}(\underline{\phi}_{m_k}(w_{m_k}) - \underline{\phi}_{m_k-1}(w_{m_k})) = 0$. Let $\{w_{m_k}\}_{k\in\mathcal{K}}$ be a further subsequence such that $\lim_{k\in\mathcal{K}} w_{m_k} = \bar{w}$. Compactness of $\mathcal{W}$ ensures that $\bar{w} \in \mathcal{W}$, and thus $\phi(w^*) \leq \phi(\bar{w})$. We know that:

$$\underline{\phi}_m(w) \leq c_w^T w + \frac{1}{m} \sum_{\xi=1}^{m} h(w, d_\xi), \;\; w \in \mathcal{W}, \tag{11}$$

and thus,

$$\limsup_{m\in\mathcal{M}} \underline{\phi}_m(w^*) \leq c_w^T w^* + \mathbb{E}[h(w^*, D)] = \phi(w^*) \tag{12}$$

with probability one for any index set $\mathcal{M}$. Since w_m minimizes $\underline{\phi}_{m-1}(\cdot)$ on $\mathcal{W}$, we have $\underline{\phi}_{m-1}(w_m) \leq \underline{\phi}_{m-1}(w^*)$ for all m. From Theorem 3.3, $\lim_{k\in\mathcal{K}} \phi_{m_k}(w_{m_k}) \leq \phi(\bar{w})$ with probability one and, also by definition, $\lim_{k\in\mathcal{K}} \underline{\phi}_{m_k}(w_{m_k}) - \underline{\phi}_{m_k-1}(w_{m_k}) = 0$. Thus, we have $\lim_{k\in\mathcal{K}} \underline{\phi}_{m_k-1}(w_{m_k}) = \phi(\bar{w})$, with probability one. Combining these results we obtain:

$$\begin{aligned} \phi(w^*) &\leq \phi(\bar{w}) \\ &= \limsup_{k\in\mathcal{K}} \underline{\phi}_{m_k-1}(w_{m_k}) \\ &\leq \underline{\phi}_{m_k-1}(w^*) \leq \phi(w^*), \end{aligned}$$

with probability one. We thus have that $\phi(\bar{w}) = \phi(w^*)$. □

3.3. Short-term MPC controller

The CP scheme is guaranteed to deliver optimal targets as data is accumulated over time. Notably, because the scheme is inherently retroactive, it achieves optimal targets without using any data forecasts. So the question is: How does the proposed scheme accommodate forecast information? From an implementation stand-point, another important question is: What metrics can one use to monitor optimality of the targets?

The proposed scheme offers a couple of mechanisms to embed forecast information. First, initial guesses for the targets $w_m = (x_{m,0}, \eta_m)$ at period $m = 1$, can be obtained by solving $\mathbf{P}_{m'}$ for some $m' \geq 1$ that uses a data forecast $\{\hat{d}_\xi\}_{\xi=1}^{m'}$. Because the forecast is expected to contain errors, the initial targets are expected to be suboptimal but these will be refined as true data is obtained. Also, as discussed in Section 3.1, the proposed scheme also enables the incorporation of forecasts at the beginning of each period to compute the intra-period policies. In particular, given that the system is at x_m, the current guess for the targets w_{m+1}, and a forecast $\hat{d}_{m+1}$ over period $m+1$, one can compute the internal control and state policies $\hat{y}_{m+1} = (\hat{x}_{m+1}, \hat{u}_{m+1})$ that satisfy the targets w_{m+1} using an intra-period MPC controller. In the linear case, this can be done by solving:

$$\min_{y\in\mathbb{R}_+^{n_y}} \;\hat{c}_{m+1}^T y \;\;\text{s.t.}\;\; Wy = \hat{r}_{m+1} - Tw_{m+1}. \tag{13}$$

Clearly, the policies $\hat{y}_{m+1}$ are suboptimal because the forecast $\hat{d}_{m+1}$ will deviate from the true realization d_{m+1} (this becomes known at the end of period $\xi+1$). Because the cost function $h(\cdot,\cdot)$ is continuous, one can use standard perturbation results to show that the optimality error in the intra-period policy is bounded by the forecast error as $|h(w_{m+1}, d_{m+1}) - h(w_{m+1}, \hat{d}_{m+1})| \leq L\|d_{m+1} - \hat{d}_{m+1}\|$ for some Lipschitz constant $L \in \mathbb{R}_+$ (see Theorem 4.156 in Bonnans and Shapiro (2013)). This implies that the quality of the forecast does affect the optimality with respect to the intra-period policies. Interestingly, however, the short-term MPC controller only needs to have a forecast over a period of length n (as opposed to over the entire horizon p). Consequently, the hierarchical MPC controller is less sensitive to forecast errors than standard proactive MPC approaches. In addition, we emphasize that the forecast quality does not affect the optimality of the targets.

The period length n introduces an interesting and important trade-off between long-term and short-term performance. As discussed previously, increasing n ensures that the resulting periodic policy better approximates the policy of problem $\mathbf{O}_m$. On the other hand, increasing n indicates that the short-term MPC controller needs to run over a longer period; as a result, it is more computationally expensive and susceptible to forecast errors. For instance, as we have noted, as n is increased the hierarchical MPC scheme resembles a standard proactive MPC scheme (since the periodicity constraints are enforced less often).

The full hierarchical scheme is sketched in Fig. 1 and is summarized as follows:

(1) Initialize $m \leftarrow 1$, $\mathbb{V}_m \leftarrow \emptyset$, and w_m.
(2) Use forecast $\hat{d}_m$ and targets w_m to obtain intra-period policies $\hat{y}_m$ using MPC controller.
(3) Let system transition from $m \to m+1$.
(4) Observe d_m and solve $\mathbf{S_m}$ to obtain $\pi(w_m, d_m)$.
(5) Update $\mathbb{V}_m \leftarrow \mathbb{V}_{m-1} \cup \{\pi(w_m, d_m)\}$.
(6) Obtain α_ξ^m, β_ξ^m, $\xi = 1, \ldots, m$ from cut generator.
(7) Solve $\mathbf{M}_m$ and obtain updated targets w_{m+1}.
(8) Shift period time $m \leftarrow m+1$ and return to Step 2.

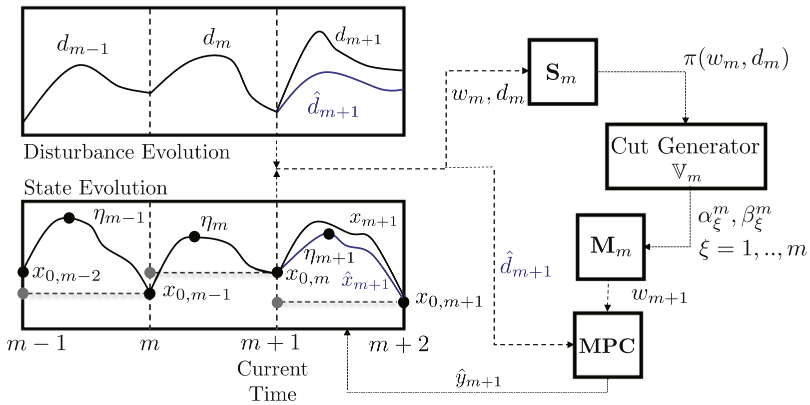

Fig. 1. Sketch of hierarchical scheme using cutting planes.

The SP interpretation allows us to derive metrics and techniques to monitor optimality. We first note that the running cost $\phi_m(w_m)$ evaluated at the current target w_m can be evaluated by solving the sequence of subproblems $\mathbf{S}_\xi$, $\xi = 1, \dots, m$. The running cost is an upper bound of the optimal running cost (obtained by solving $\mathbf{P}_m$). Moreover, the proposed CP scheme offers the guarantee that the cost $\underline{\phi}_m(w_m)$ is a lower bound of the running cost $\phi_m(w_m)$, which is an asymptotically exact statistical approximation of $\phi(w_m)$. The difference between the running cost (upper bound) and the lower bound is known in the SP literature as the optimality gap and is formally defined as $\epsilon_m := (\phi_m(w_m) - \underline{\phi}_m(w_m))/\phi_m(w_m)$. Here, we refer to ϵ_m as the current gap. The convergence of the CP scheme guarantees that $\lim_{m\to\infty} \epsilon_m = 0$. We note that this gap can be used to measure the quality of the CP approximation but should be used with care when interpreting optimality. In particular, the gap can only be used as a measure of optimality in the limit $m \to \infty$ (once the running cost $\phi_m(w_m)$ is close to the actual cost $\phi_\infty(w_m)$). Consequently, one usually resorts to computing confidence intervals for $\phi_m(w_m)$ by using inference (evaluate the cost at w_m but using different combinations of realizations) (Linderoth, Shapiro, & Wright, 2006). Motivated by this, in benchmarking studies, we are also interested in monitoring the overall optimality gap $\bar{\epsilon}_m := (\phi_\infty(w_m) - \underline{\phi}_m(w_m))/\phi_\infty(w_m)$, provided that $\phi_\infty(w_m)$ can be computed.

3.4. Extensions to nonlinear systems

In the case of linear systems, the CP scheme can approximate the running cost $\phi_m(w)$ using a finite number of supporting hyperplanes, which keeps the master problem $\mathbf{M}_m$ tractable. The SP representation opens the door to other schemes such as proximal point methods for nonlinear (but convex) problems (Bertsekas, 2011). Here, the idea is to prevent the accumulation of data over time by summarizing past information in terms of a proximal term. In our context, for instance, the proximal point strategy will result in a problem of the form at stage m:

$$\min_{w\in\mathcal{W}} \ \mu\|w - w_m\|^2 + g(w) + \frac{1}{m'}\sum_{\xi=m'}^{m} h(w, d_\xi). \tag{14}$$

Here, $\mu\|w-w_m\|^2$ is a regularization term with parameter $\mu > 0$. This term summarizes data before time m' and resembles the arrival cost used in moving horizon estimation (Rao, Rawlings, & Mayne, 2003). In the general case of nonconvex problems, one can use a statistical approximation $\min_{w\in\mathcal{W}} \ \phi_m(w)$ at every period m by using linear algebra decomposition schemes. In particular, it is well-known that problems with the structure of $\mathbf{P}_m$ give linear algebra systems that enable parallel decomposition (Zavala et al., 2008). Given that the coupling is only in the space of the periodic targets x_0, this approach can scale to systems with thousands of states and tens of thousands of periods (realizations). This approach, however, exhibits a fundamental limitation in the number of periods that it can handle (because data is accumulated unboundedly over time). This is, in fact, also a limitation of statistical approximation schemes for SP. To circumvent this issue, one can use clustering techniques that seek to compress the realization space to maintain a tractable approximation (Cao, Laird, & Zavala, 2016). Such techniques are based on the observation that data realizations tend to be redundant and only a small subset actually impacts the cost. One can quantify the error incurred in the scenario compression by using inference techniques (Linderoth et al., 2006).

3.5. Stability considerations

The proposed hierarchical schemes provide important stability properties. A formal treatment of such properties is beyond the scope of this work, but here we present some basic arguments. We first consider the *nominal case*, in which the data realization in each period is the same ($d_\xi = d_{\xi'}$ for all ξ, ξ') and it is known. This is equivalent to assuming that the dynamics are of the form $x_{\xi,t+1} = f(x_{\xi,t}, u_{\xi,t})$ (they are time-invariant). We also assume that no peak cost $\phi_2(\cdot)$ is used. This nominal setting is considered in the periodic MPC formulations of Huang et al. (2011) and Zanon et al. (2017). Both of these formulations use a proactive RH strategy and enforce a periodicity constraint at the end of the horizon. In Huang et al. (2011), it is shown that their MPC scheme delivers a periodic state that is a steady-state (the closed-loop system is stable). In Zanon et al. (2017), it is shown that their MPC scheme converges to a periodic state (not necessarily a steady-state), and is thus stable in this sense. Converse to a periodic state that is not a steady-state is desirable because this provides flexibility. Both of these proactive schemes require dissipativity properties. For this same nominal setting, we note that the solution of P_∞ also delivers a periodic state (by construction). Since the proposed hierarchical scheme converges to the periodic state of P_∞, we thus have that our approach is stable in this sense. This result does not require any dissipativity assumptions.

For the more general case in which the data changes in each period, we have that the dynamics are time-variant. Moreover, we have that the data cannot be forecast perfectly. Surprisingly, for this more challenging setting, we have that the retroactive scheme proposed also delivers the optimal periodic state and thus is stable. This is a remarkable result that no other known scheme reported in the literature provides.

The short-term MPC controller admits a standard state-feedback control representation. This is because the control policy inside the period is updated based on the current state and the desired state target. The long-term controller (updating the targets), however, does not admit such a representation because it is retroactive (and not proactive). In particular, it uses the entire history (and not just the current state) to compute the next targets. In future work, we will formalize the stability analysis of the approach and the associated state-feedback representation.

4. Computational experiments

The performance of the proposed scheme is demonstrated using an application in buildings with electricity storage. The goal of the controller is to determine optimal short-term (hourly) participation strategies in frequency regulation (FR) and energy markets that maximize revenue and simultaneously mitigate long-term demand charges. We consider a utility-scale stationary battery with a capacity of 0.5 MWh, rated power of 1 MW, and a ramping limit of 0.5 MW/h. We use real data for energy and FR prices from PJM Interconnection (shown in Fig. 2). We also use real load data for a typical university campus (shown also in Fig. 2). The periodic components in the energy prices and the load profile can be clearly observed, while periodicity of FR prices is not as strong. The MPC problem is formulated using daily periods of $n = 24$ hours and we consider a horizon of $m = 300$ days (nearly a year). The model parameters include the battery storage capacity ($\overline{E}$ kWh), maximum discharging and charging rates (power) ($\overline{P}$, $\underline{P}$ in kW, respectively), minimum fraction of battery capacity reserved for frequency regulation (ρ in kWh/kW), and maximum ramping limit ($\overline{\Delta P}$ in kW/h). The random data are the loads ($L_{\xi,t}$ kW), market prices for electricity and FR capacity ($\pi^e_{\xi,t}$ \$/kWh and $\pi^f_{\xi,t}$ \$/kW respectively), demand charge (monthly) (π^D in \$/kW), and fraction of FR capacity requested by ISO ($\alpha_{\xi,t}$). The decision variables are net battery discharge rate (power) ($P_{\xi,t}$ in kW), FR capacity provided to ISO ($F_{\xi,t}$ in kW), state-of-charge (SOC) of the battery ($E_{\xi,t}$ in kWh), load requested from utility ($d_{\xi,t}$ in kW) and peak load (D in kW). The formulation minimizes the total cost (negative total revenue), which is given by the demand charge (peak cost) and the revenues collected from the market (time-additive cost). Detailed notation and analysis of this formulation is presented in Kumar et al. (2018a). The problem has the form:

$$
\begin{aligned}
\min\ & \frac{1}{m}\sum_{\xi\in\Xi}\sum_{t\in\mathcal{T}}\left(-\pi^e_{\xi,t}(P_{\xi,t}-\alpha_{\xi,t}F_{\xi,t})-\pi^f_{\xi,t}F_{\xi,t}\right)+\pi^D D\\
\text{s.t.}\ & P_{\xi,t}+F_{\xi,t}\le\overline{P},\ t\in\mathcal{T},\xi\in\Xi\\
& P_{\xi,t}-F_{\xi,t}\ge-\underline{P},\ t\in\mathcal{T},\xi\in\Xi\\
& E_{\xi,t+1}=E_{\xi,t}-P_{\xi,t}+\alpha_{\xi,t}F_{\xi,t},\ t\in\bar{\mathcal{T}},\xi\in\Xi\\
& \rho F_{\xi,t}\le E_{\xi,t}\le\overline{E}-\rho F_{\xi,t},\ t\in\mathcal{T},\xi\in\Xi\\
& \rho F_{\xi,t}\le E_{\xi,t+1}\le\overline{E}-\rho F_{\xi,t},\ t\in\bar{\mathcal{T}},\xi\in\Xi\\
& -\overline{\Delta P}\le P_{\xi,t+1}-P_{\xi,t}\le\overline{\Delta P},\ t\in\bar{\mathcal{T}},\xi\in\Xi\\
& d_{\xi,t}=L_{\xi,t}-P_{\xi,t}+\alpha_{\xi,t}F_{\xi,t},\ t\in\mathcal{T},\xi\in\Xi\\
& d_{\xi,t}\le D,\ t\in\mathcal{T},\xi\in\Xi\\
& P_{\xi,t}+F_{\xi,t}\le L_{\xi,t},\ t\in\mathcal{T},\xi\in\Xi\\
& E_{\xi,0}=E_0,\ \xi\in\Xi\\
& E_{\xi,N_\xi}=E_0,\ \xi\in\Xi\\
& 0\le E_{\xi,t}\le\overline{E},\ t\in\mathcal{T},\xi\in\Xi\\
& -\underline{P}\le P_{\xi,t}\le\overline{P},\ t\in\mathcal{T},\xi\in\Xi\\
& 0\le F_{\xi,t}\le\overline{P},\ t\in\mathcal{T},\xi\in\Xi
\end{aligned}
$$

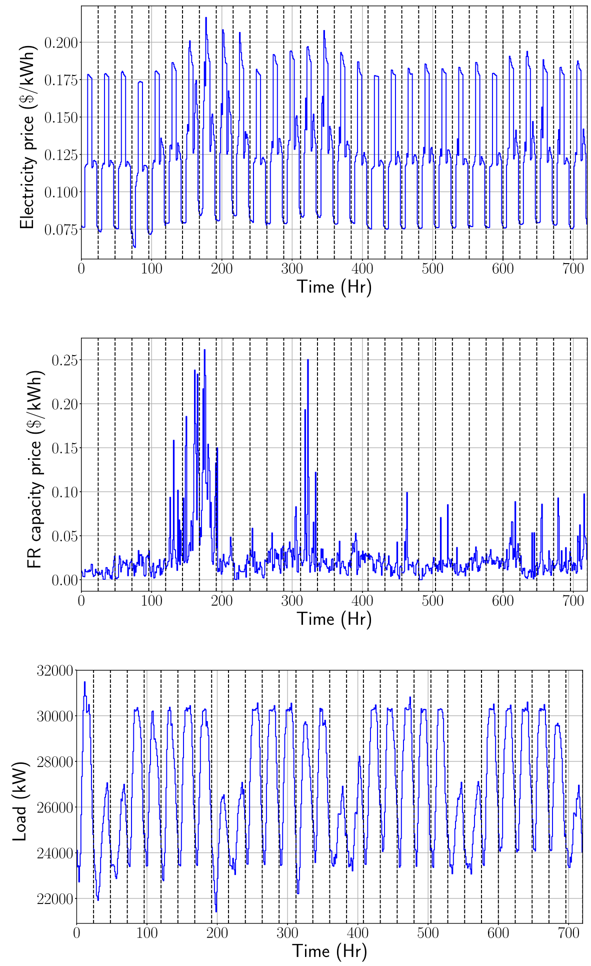

Fig. 2. Energy price (top), FR price (middle), and load (bottom) data.

We first solve the long-horizon MPC problem (1) for a horizon of 300 days assuming perfect knowledge of the data to obtain the optimal policy. We compare this policy against that of a long-horizon MPC formulation that enforces periodicity constraints (2). The optimal and periodic policies over 30 days are presented in Fig. 3. Here, we can see that the policies match. The total cost of the long-horizon MPC problem is \$136,050 while the cost of the long-horizon problem with periodicity constraints is \$136,068. We can see that, in this application, state periodicity arises naturally because the battery needs to maintain a minimum SOC level after each period. We then ran the proposed retroactive CP scheme for 300 periods to progressively update the targets. The evolution of the current gap ϵ_m and overall optimality gap $\bar{\epsilon}_m$ is shown in Fig. 4. We observe that the overall gap eventually vanishes, demonstrating that the CP scheme delivers optimal targets. The overall gap closes to zero close to the end of the horizon, once the peak demand is observed. We also see that the current gap closes to 0% in about 50 periods and stays there for the rest of the horizon. This illustrates that the cutting planes approximate the cost function well, but also that this metric can be misleading. We have also found that the performance of the proposed hierarchical scheme is close to that obtained with the optimal long-term periodic policy. In particular, the cost of the hierarchical MPC is \$139,978 (a difference of 2.89%). Accelerating the convergence of the CP scheme can help decrease this gap. We also compared the performance of the hierarchical scheme with that of a standard (non-periodic) MPC scheme that uses a prediction horizon one and fourteen days. The corresponding costs are

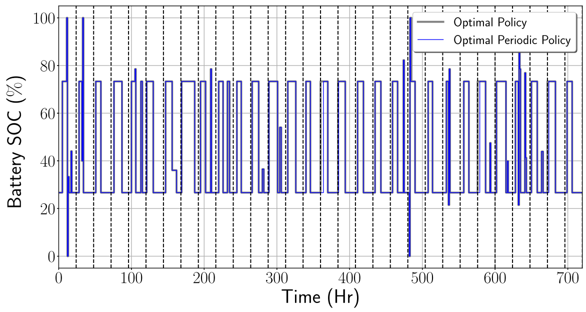

Fig. 3. Optimal and optimal periodic policies.

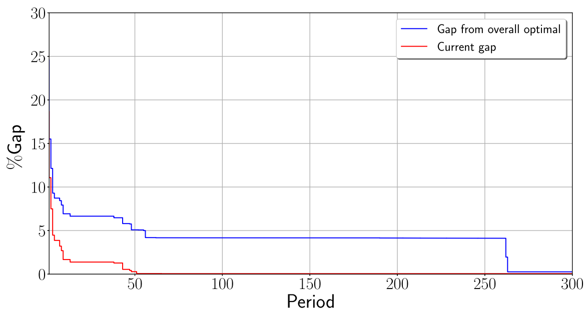

Fig. 4. Evolution of optimality gap.

$139,884 and $138,739. We can see that the cost decreases slowly with the prediction horizon. We also observe that the cost for standard MPC with a one-day prediction horizon is similar to that of the hierarchical scheme. The hierarchical scheme, however, offers optimality guarantees (standard MPC does not).

Fig. 5 shows the evolution of the periodic state and peak targets. We see that the CP scheme adaptively updates the targets as data is accumulated over time. The SOC target settles quickly to the optimal level of 59%. The peak target requires more periods to settle and this behavior is attributed to the fact that the peak load is observed at period 264. After this period the peak settles at its optimal value of 32,935 kW. Fig. 6 shows the intra-period policies for the short-term MPC controller for the first 7 days (periods) of operation. We see that the controller follows the target of the CP scheme.

We also compared the retroactive hierarchical MPC scheme with the proactive MPC approach for periodic systems of Huang et al. (2011). For this scheme, we consider a prediction horizon of 7 days and a period of 1 day. For this comparison, we removed the demand charge (peak cost) from the formulation and assumed that disturbances can be forecast perfectly. In Fig. 7, we present the evolution of the periodic SOC for both approaches. We observe that, with the proactive approach, the periodic state does not converge. This is because the dynamics are time-varying. A similar behavior would be expected from the proactive scheme of Zanon et al. (2017). On the other hand, the retroactive approach converges to a periodic state. Moreover, we found that the cost of the proactive scheme is 5% worse than that of the retroactive scheme. This is because the proactive scheme does not capture the long-term trend of the disturbances. In other words, the retroactive scheme delivers a policy that is close-to-optimal.

5. Conclusions and future work

We showed that stochastic programming provides a framework to design hierarchical MPC schemes for periodic systems.

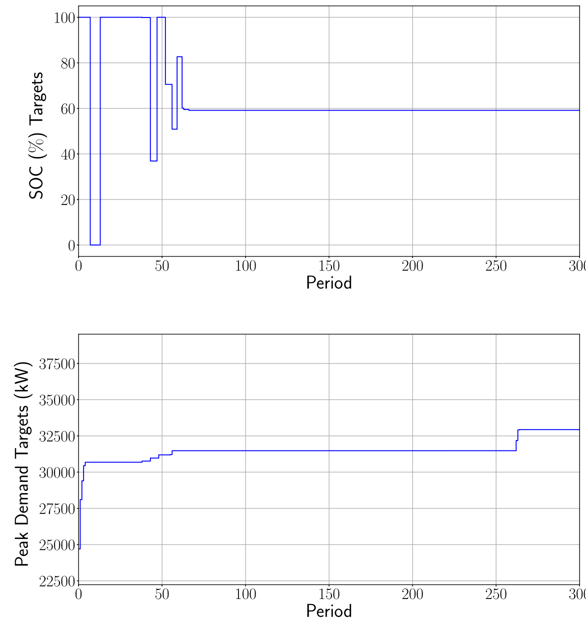

Fig. 5. Evolution of periodic SOC (top) and peak (bottom) targets obtained with cutting-plane scheme.

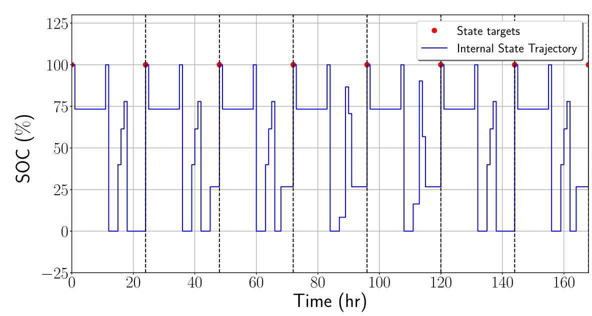

Fig. 6. Evolution of periodic SOC targets and intra-period policies obtained with CP scheme (first seven periods).

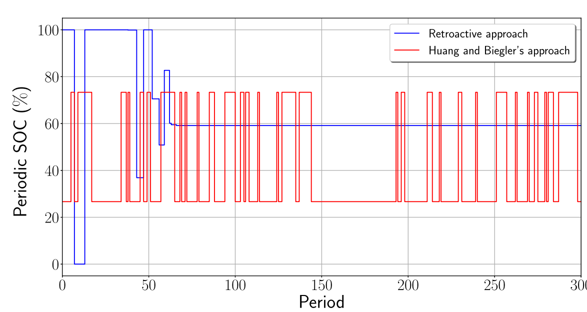

Fig. 7. Comparison of evolution of periodic SOC resulting from the retroactive approach and Huang and Biegler's proactive MPC approach (Huang et al., 2011).

We have shown that, under periodicity, it is possible to compute and refine periodic state targets by solving a retroactive optimization problem that progressively accumulates historical data. The retroactive problem is a statistical approximation of the stochastic program that delivers optimal targets in the long run to guide a short-term MPC controller. The computation of the optimal targets can be achieved without data forecasts. The SP setting opens the door to a number of potential developments in hierarchical MPC scheme. As part of future work, we are interested in exploring schemes for nonlinear systems and to provide optimality and stability results. Moreover, it is necessary to investigate convergent schemes that have faster convergence than cutting planes, that prevent accumulation of large amounts of data over time, and that factor in forecast information in a more effective manner.

Acknowledgment

We acknowledge funding from Johnson Controls International and the National Science Foundation under award NSF-EECS-1609183.

References

Bertsekas, D. P. (2011). Incremental proximal methods for large scale convex optimization. *Mathematical Programming*, *129*(2), 163.

Bertsimas, D., & Tsitsiklis, J. N. (1997). *Introduction to linear optimization, Vol. 6*. Athena Scientific Belmont, MA.

Bonnans, J. F., & Shapiro, A. (2013). Perturbation analysis of optimization problems. Springer Science & Business Media.

Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: forecasting and control*. John Wiley & Sons.

Braun, J. E. (1990). Reducing energy costs and peak electrical demand through optimal control of building thermal storage. In *ASHRAE Transactions, Vol. pt 2* (pp. 876–888).

Cao, Y., Laird, C. D., & Zavala, V. M. (2016). Clustering-based preconditioning for stochastic programs. *Computational Optimization and Applications*, *64*(2), 379–406.

CarøE, C. C., & Schultz, R. (1999). Dual decomposition in stochastic integer programming. *Operations Research Letters*, *24*(1), 37–45.

Dowling, A. W., Kumar, R., & Zavala, V. M. (2017). A multi-scale optimization framework for electricity market participation. *Applied Energy*, *190*, 147–164.

Geoffrion, A. M. (1972). Generalized benders decomposition. *Journal of Optimization Theory and Applications*, *10*(4), 237–260.

Higle, J. L., & Sen, S. (1991). Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research*, *16*(3), 650–669.

Hu, T.-C., Rosalsky, A., & Volodin, A. (2008). On convergence properties of sums of dependent random variables under second moment and covariance restrictions. *Statistics & Probability Letters*, *78*(14), 1999–2005.

Huang, R., Harinath, E., & Biegler, L. T. (2011). Lyapunov stability of economically oriented nmpc for cyclic processes. *Journal of Process Control*, *21*(4), 501–509.

Kumar, R., Wenzel, M. J., Ellis, M. J., ElBsat, M. N., Drees, K. H., & Zavala, V. M. (2018a). Handling long horizons in mpc: A stochastic programming approach. In *2018 annual american control conference (ACC)* (pp. 715–720). IEEE.

Kumar, R., Wenzel, M. J., Ellis, M. J., ElBsat, M. N., Drees, K. H., & Zavala, V. M. (2018b). A stochastic model predictive control framework for stationary battery systems. *IEEE Transactions on Power Systems*, *33*(4), 4397.

Linderoth, J., Shapiro, A., & Wright, S. (2006). The empirical behavior of sampling methods for stochastic programming. *Annals of Operations Research*, *142*(1), 215–241.

Picasso, B., De Vito, D., Scattolini, R., & Colaneri, P. (2010). An mpc approach to the design of two-layer hierarchical control systems. *Automatica*, *46*(5), 823–831.

Ragan, P., & Manuel, L. (2008). Statistical extrapolation methods for estimating wind turbine extreme loads. *Journal of Solar Energy Engineering*, *130*(3), 031011.

Rao, C. V., Rawlings, J. B., & Mayne, D. Q. (2003). Constrained state estimation for nonlinear discrete-time systems: stability and moving horizon approximations. *IEEE Transactions on Automatic Control*, *48*(2), 246–258.

Rawlings, J., & Mayne, D. (2009). *Model predictive control: Theory and design*. Madison, WI: Nob Hill Publishing.

Risbeck, M. J., Maravelias, C. T., Rawlings, J. B., & Turney, R. D. (2015). Cost optimization of combined building heating/cooling equipment via mixed-integer linear programming. In *American control conference (ACC), 2015* (pp. 1689–1694). IEEE.

Risbeck, M. J., Maravelias, C. T., Rawlings, J. B., & Turney, R. D. (2017). A mixed-integer linear programming model for real-time cost optimization of building heating, ventilation, and air conditioning equipment. *Energy and Buildings*, *142*, 220–235.

Rudin, W., et al. (1964). *Principles of mathematical analysis, Vol. 3*. McGraw-hill New York.

Scattolini, R., & Colaneri, P. (2007). Hierarchical model predictive control. In *Decision and control, 2007 46th IEEE conference on* (pp. 4803–4808). IEEE.

Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2009). *Lectures on stochastic programming: modeling and theory*. SIAM.

Subramanian, K., Rawlings, J. B., & Maravelias, C. T. (2014). Economic model predictive control for inventory management in supply chains. *Computers & Chemical Engineering*, *64*, 71–80.

Zanon, M., Grüne, L., & Diehl, M. (2017). Periodic optimal control, dissipativity and mpc. *IEEE Transactions on Automatic Control*, *62*(6), 2943–2949.

Zavala, V. M. (2016). New architectures for hierarchical predictive control. *IFAC-PapersOnLine*, *49*(7), 43–48.

Zavala, V. M., Laird, C. D., & Biegler, L. T. (2008). Interior-point decomposition approaches for parallel solution of large-scale nonlinear parameter estimation problems. *Chemical Engineering Science*, *63*(19), 4834–4845.

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

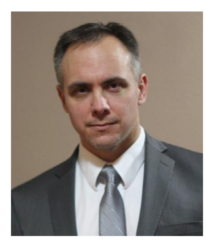

Michael J. Wenzel received his Ph.D. in Electrical Engineering from Marquette University in 2009. He is currently a Senior Staff Engineer at Johnson Controls. His interests include model predictive control, optimization, cost/profit-driven control and probability theory for real-time applications.

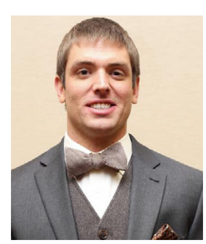

Matthew J. Ellis received his B.S. degree in Chemical and Biological Engineering from the University of Wisconsin—Madison in 2010 and his Ph.D. degree in Chemical and Biomolecular Engineering from the University of California, Los Angeles in 2015. He was the recipient of the UCLA 2016 Outstanding Ph.D. Award in Chemical Engineering. Dr. Ellis joined the advanced development team at Johnson Controls (JCI) in 2015 where he co-designed and developed the core EMPC algorithm used within JCI's central plant optimization offering. Currently, he serves as the technical lead in extending this optimization-based control strategy to building-side HVAC systems. His interests include model predictive control and other advanced control algorithms, mathematical/numerical optimization, system/parameter identification, and cost/profit-driven control and automation algorithms for real-time applications.

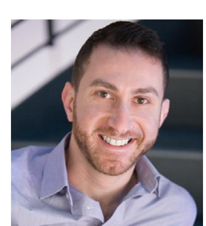

Mohammad N. ElBsat received his B.S. degree in Electrical Engineering from the Lebanese University in 2006 in Hadath, Lebanon. He, then, received his M.S. and Ph.D. also in Electrical Engineering, from Marquette University in 2008 and 2012, respectively. He is currently a Principal Engineer at Johnson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

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

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