
Perturbation-based Regret Analysis of Predictive Control in Linear Time Varying Systems

Yiheng Lin

California Institute of Technology
Pasadena, CA, USA
yihengl@caltech.edu

Yang Hu

Tsinghua University
Beijing, China
huy18@mails.tsinghua.edu.cn

Guanya Shi

California Institute of Technology
Pasadena, CA, USA
gshi@caltech.edu

Haoyuan Sun

California Institute of Technology
Pasadena, CA, USA
hsun2@caltech.edu

Guannan Qu

California Institute of Technology
Pasadena, CA, USA
gqu@caltech.edu

Adam Wierman

California Institute of Technology
Pasadena, CA, USA
adamw@caltech.edu

Abstract

We study predictive control in a setting where the dynamics are time-varying and linear, and the costs are time-varying and well-conditioned. At each time step, the controller receives the exact predictions of costs, dynamics, and disturbances for the future k time steps. We show that when the prediction window k is sufficiently large, predictive control is input-to-state stable and achieves a dynamic regret of $O(\lambda^k T)$, where $\lambda < 1$ is a positive constant. This is the first dynamic regret bound on the predictive control of linear time-varying systems. We also show a variation of predictive control obtains the first competitive bound for the control of linear time-varying systems: $1 + O(\lambda^k)$. Our results are derived using a novel proof framework based on a perturbation bound that characterizes how a small change to the system parameters impacts the optimal trajectory.

1 Introduction

We study the problem of predictive control in a linear time-varying (LTV) system, where the dynamics is given by $x_{t+1} = A_t x_t + B_t u_t + w_t$. Here, x_t is the state, u_t is the control action, and w_t is the disturbance or exogenous input. At each time step t , the online controller incurs a time-varying state cost $f_t(x_t)$ and control cost $c_t(u_{t-1})$, and then decides its next control action u_t . In deciding u_t the controller makes use of predictions of the next k future disturbances, cost functions, and dynamical matrices, and seeks to minimize its total cost on a finite horizon T . Our main results bound the dynamic regret and competitive ratio of predictive controllers in this LTV setting.

Recently, a growing literature has sought to design controllers that achieve learning guarantees such as static regret [1, 2], dynamic regret [3, 4], and competitive ratio [5]. The most relevant

[†]Yiheng Lin, Yang Hu, Haoyuan Sun, Guanya Shi, and Guannan Qu contributed equally to this work.

[‡]This work was supported by NSF grants CNS-2106403, NGSDI-2105648, and AitF-1637598, with additional support from Amazon AWS, PIMCO, and the Resnick Sustainability Institute. Yiheng Lin was supported by Kortschak Scholars program.

line of work concerns predictive control with learning guarantees, which studies how to leverage the prediction window k to reduce the regret and competitive ratio. This line of work has focused on linear time-invariant (LTI) systems [3, 4, 6, 7]. However, linear time-varying (LTV) systems have received increasing attention in recent years due to their importance in a variety of emerging applications, despite the challenges associated with analysis. For example, in the problem of power grid frequency regulation, the dynamics is determined by the proportion of renewable energy in total power generation, which is time-varying [8, 9]. It is also common to use the LTV systems as an approximation of nonlinear dynamics in predictive control and planning [10–13].

The current lack of progress toward understanding measures like regret and competitive ratio in LTV settings is due to the need for new techniques to generalize the dynamics from LTI to LTV and the costs from quadratic to well-conditioned functions. Specifically, the proof approaches used in previous studies on regret and competitive ratio of predictive control in LTI dynamics with quadratic costs, e.g., [4, 6, 7], require explicitly writing down the cost-to-go function, optimal control actions, and algorithm’s actions as functions of the system parameters. This is very difficult, if not impossible, for general cost functions that do not have a quadratic form. A promising approach that does not require such explicit characterizations is to derive results via reductions from optimal control to online convex optimization with multi-step memory, e.g., [1, 3, 5, 14, 15]. However, such reductions usually do not work well for LTV systems due to the need to represent the problem in control canonical form [3, 5], or due to limitations on the policy class and comparisons to static benchmarks [1, 15].

Perhaps the most prominent approach for controlling LTV systems is Model Predictive Control (MPC), also known as Receding Horizon Control [16]. Generally speaking, at each time step, an MPC-style algorithm solves a predictive trajectory for the future k time steps and commit the first control action in this trajectory. MPC-style algorithms are known to work well in practice, even when the dynamics are non-linear and time-varying, e.g., [13, 17–19]. On theoretical side, the asymptotic behaviors of MPC such as stability and convergence have been studied intensively under general assumptions on dynamics and costs [20–23]. However, non-asymptotic guarantees such as regret and competitive ratio of MPC-style policies have been limited. Despite recent work providing such guarantees in the context of LTI systems with quadratic costs, e.g., [4, 6, 7], the derivation of regret and competitive ratio results for MPC in LTV systems remains open.

Contributions. We provide the first regret and competitive ratio results for a controller in LTV systems with time-varying costs. Specifically, we show that an MPC-style predictive control algorithm (Algorithm 1) achieves a dynamic regret that decays exponentially with respect to the length of prediction window k in the LTV system (Theorem 4.2): $O(\lambda^k T)$, where the decay rate λ is a positive constant less than 1. This almost matches the exponential lower bound for improvement from predictions in the LTI setting shown in [3] in the sense that, to achieve any target regret level, the required length of prediction k shown by our bound differs from the theoretical lower bound by at most a constant factor. With a variation of predictive control (Algorithm 2), we also show the first competitive bound in LTV systems with time-varying well-conditioned costs (Theorem 4.3): $1 + O(\lambda^k)$, where the decay rate λ is identical with the one in the regret bound.

We develop a novel analysis framework based on a perturbation approach. Specifically, instead of solving for the optimal states and control actions like previous analyses in the LTI setting with quadratic costs, e.g., [4, 7], we bound how much impact an perturbation to the system parameters can have on the optimal solution. This type of perturbation bound (Theorem 3.3) can be shown even when the optimal trajectory cannot be written down explicitly, which allows it to be applied in LTV systems with well-conditioned costs. Then, we utilize this perturbation bound to establish results on dynamic regret and the competitive ratio. In addition, we want to emphasize that the perturbation approach we develop is highly modular and extendable. For instance, if a stronger perturbation bound for some specific class of dynamics and/or cost functions can be shown, the dynamic regret of the predictive controller will improve. Similarly, to further generalize the problem setting (e.g., to include additional constraints), one only needs to establish the corresponding perturbation bounds and the regret result will follow.

Another important component of the proof is a novel reduction between LTV control and online optimization. Connections between online optimization and control have received increasing attention in recent years, e.g., [1, 3, 5, 14, 15]. Existing reductions rely on the canonical form, which does not apply to LTV systems, and/or formulations of online optimization with memory of multiple prior time steps, which makes the online problem more challenging. The reduction we present here relies

on neither, and is thus a fundamentally different approach to connect control and online optimization. Further, this reduction is not specific to the predictive control algorithm we study, and we expect it to prove useful for other controllers in future work. A limitation of our reduction framework is that it cannot handle state/control constraints. This limitation is shared by previous works [3, 4, 6, 7], and represents a challenging open question in the literature.

2 Background and Setting

We consider a finite-horizon discrete-time online control problem with linear time-varying (LTV) dynamics, time-varying costs, and disturbances, namely

$$\begin{aligned} \min_{x_{0:T}, u_{0:T-1}} \quad & \sum_{t=1}^T (f_t(x_t) + c_t(u_{t-1})) \\ \text{s.t.} \quad & x_t = A_{t-1}x_{t-1} + B_{t-1}u_{t-1} + w_{t-1}, t = 1, \dots, T, \\ & x_0 = x(0), \end{aligned} \tag{1}$$

where $x_t \in \mathbb{R}^n$, $u_t \in \mathbb{R}^m$, and $w_t \in \mathbb{R}^n$ respectively denote the state, the control action, and the disturbance of the system at time steps $t = 1, \dots, T$, and $x(0) \in \mathbb{R}^n$ is a given initial state. By convention, the hitting cost function $f_t : \mathbb{R}^n \rightarrow \mathbb{R}_+$ and control cost function $c_t : \mathbb{R}^m \rightarrow \mathbb{R}_+$ are assumed to be time-varying and well-conditioned. Define the tuple $\vartheta_t := (A_t, B_t, w_t, f_{t+1}, c_{t+1})$.

In the classical setting where no predictions are available, after observing state x_t at time step t , the algorithm needs to decide the control action u_t before observing ϑ_t , which is an unknown random disturbance input. We use the following event sequence to describe this ordering:

$$x_0, u_0, \vartheta_0, x_1, u_1, \vartheta_1, x_2, \dots, x_{T-1}, u_{T-1}, \vartheta_{T-1}, x_T.$$

We assume that the algorithm has access to the exact predictions of disturbances, cost functions and dynamical matrices in the future k time steps (which are time-varying); i.e., the event sequence is

$$x_0, \vartheta_0, \vartheta_1, \dots, \vartheta_{k-1}, u_0, \vartheta_k, u_1, \vartheta_{k+1}, \dots, u_{T-k-1}, \vartheta_{T-1}, u_{T-k}, u_{T-k+1}, \dots, u_{T-1}.$$

Here we assume all predictions are *exact*, and leave the case of inexact predictions for future work. This prediction model has been used in previous works like [4, 24–26], and is available in many real-world applications such as disturbance estimation in robotics and frequency regulation in power grids. The availability is due to the fact that, in such scenarios as mentioned above, experiments or observations on the dynamics can be conducted repeatedly and consistently, which makes it feasible to train a good predictor based on the data collected from repeated trials.

2.1 Assumptions

As is standard in studies of regret and competitive ratio in linear control problems, we assume the cost functions are well-conditioned.

Assumption 2.1 (Well-conditioned Costs). *The cost functions satisfy the following constraints:*

1. $f_t(\cdot)$ is m_f -strongly convex for $t = 1, \dots, T$, and ℓ_f -strongly smooth for $t = 1, \dots, T - 1$.
2. $c_t(\cdot)$ is both m_c -strongly convex and ℓ_c -strongly smooth for $t = 1, \dots, T$.
3. $f_t(\cdot)$ and $c_t(\cdot)$ are twice continuously differentiable for $t = 1, \dots, T$.
4. $f_t(\cdot)$ and $c_t(\cdot)$ are non-negative, and $f_t(0) = c_t(0) = 0$ for $t = 1, \dots, T$.

Note that assumptions (1) through (3) are quite common [3, 5, 14, 25, 27]. Assumption (4) is less common, but can be satisfied via re-parameterization without loss of generality. Specifically, when the minimizers of state cost f_t and control cost c_t are nonzero, we perform the transformation

$$\begin{aligned} x'_t &\leftarrow x_t - \arg \min_x f_t(x), \quad u'_t \leftarrow u_t - \arg \min_u c_{t+1}(u), \\ w'_t &\leftarrow w_t + A_t \arg \min_x f_t(x) + B_t \arg \min_u c_{t+1}(u). \end{aligned}$$

The intuition of this transformation is that, when the minimizer of the cost function for the next step is known, we can always perform a translation in the state and control space to align the minimizer

with the origin. We refer the interested readers to Example 2.1 for a more intuitive explanation of the above transformation.

Additionally, we need to assume the dynamics are *controllable*. It is crucial that the dynamical system can be steered from an arbitrary initial state to an arbitrary final state via a finite sequence of admissible control actions. For linear time-invariant (LTI) systems, the full-rankness of the *controllability matrix* completely characterizes the reachability of the state space, which is generally used as a standard assumption for analysis [7, 28, 29]. This can be generalized to parallel assumptions for LTV systems as follows. We begin with a definition.

Definition 2.1. For a dynamical system with linear time-varying dynamics $x_t = A_{t-1}x_{t-1} + B_{t-1}u_{t-1} + w_{t-1}$, $t = 1, \dots, T$, the transition matrix $\Phi(t_2, t_1) \in \mathbb{R}^{n \times n}$ (from time step t_1 to t_2) is defined as

$$\Phi(t_2, t_1) := \begin{cases} A_{t_2-1}A_{t_2-2} \cdots A_{t_1} & \text{if } t_2 > t_1 \\ I & \text{if } t_2 \leq t_1 \end{cases}$$

and the controllability matrix $M(t, p) \in \mathbb{R}^{n \times (mp)}$ is defined as

$$M(t, p) := [\Phi(t+p, t+1)B_t, \Phi(t+p, t+2)B_{t+1}, \dots, \Phi(t+p, t+p)B_{t+p}].$$

The dynamical system is called *controllable* if there exists a constant $d \in \mathbb{Z}_+$, such that the controllability matrix $M(t, d)$ is of full row rank for any $t = 1, \dots, T-d$. The smallest constant d with such property is called the *controllability index* of the system.

Given the above definition, we can state the key assumption necessary for the analysis of LTV systems. We use a slightly stronger assumption than being merely controllable, which we refer to as (d, σ) -uniform controllability. It is a natural generalization of its counterpart for LTI systems (see Assumption 2 in [28], where (d, σ) is instead named as (ℓ, ν)).

Assumption 2.2. There exists positive constants a , b , and b' , such that

$$\|A_t\| \leq a, \|B_t\| \leq b, \text{ and } \|B_t^\dagger\| \leq b'$$

hold for all time steps $t = 0, \dots, T-1$, where B_t^\dagger denotes the Moore–Penrose inverse of matrix B_t . Furthermore, there exists a positive constant σ such that

$$\sigma_{\min}(M(t, d)) \geq \sigma$$

holds for all time steps $t = 0, \dots, T-d$, where d denotes the controllability index.

Note that Assumption 2.2 implies $\sigma_{\min}(M(t, p)) \geq \sigma$ for all $p \geq d$ because appending more columns to a matrix with full row rank will not reduce its minimum singular value.

The LTV setting we consider is more general than the settings which existing results on regret and competitive ratio have assumed [1, 3, 4, 7]. We highlight the implications of this general setting for enabling applications in the following examples.

Example 2.1 (Trajectory tracking in LTV systems with well-conditioned costs). Consider a trajectory tracking problem with LTV dynamics and well-conditioned costs, which generalizes the standard linear quadratic tracking problem in [4, 30] with LTI dynamics and quadratic costs. We adopt LTV dynamics $x_{t+1} = A_t x_t + B_t u_t + w_t$ and general well-conditioned cost functions $f_t(\cdot)$, $c_t(\cdot)$ (see Assumption 2.1). With the desired trajectory $d_{1:T}$, we consider a new state $\tilde{x}_t := x_t - d_t$ and a new disturbance $\tilde{w}_t := w_t + A_t d_t - d_{t+1}$. Thus, using the new state and disturbance, the problem naturally fits into our problem setting with k future predictions of $(A_t, B_t, w_t, f_t, c_t, d_{t+1})$. Note that predictive control with LTV dynamics is practical in nonlinear systems [13, 31] because the nonlinearity could be well approximated by LTV models [31].

Example 2.2 (Power grid frequency regulation). Consider the frequency regulation problem in [8], where state $x = [\theta^\top, \omega^\top]^\top$ represent the status of a power plant, and power generation $p_{in} \in \mathbb{R}^n$ is the control action. The continuous-time dynamics is given by

$$\underbrace{\begin{bmatrix} \dot{\theta} \\ \dot{\omega} \end{bmatrix}}_{\hat{x}(t)} = \underbrace{\begin{bmatrix} 0 & I \\ -M(t)^{-1}L & -M(t)^{-1}D \end{bmatrix}}_{\hat{A}(t)} \underbrace{\begin{bmatrix} \theta \\ \omega \end{bmatrix}}_{x(t)} + \underbrace{\begin{bmatrix} 0 \\ M(t)^{-1} \end{bmatrix}}_{\hat{B}(t)} \underbrace{p_{in}}_{u(t)}.$$

Algorithm 1 Predictive Control (PC_k)

- 1: **for** $t = 0, 1, \dots, T - k - 1$ **do**
 - 2: Observe current state x_t and receive predictions $\vartheta_{t:t+k-1}$.
 - 3: Solve and commit control actions $u_t := \tilde{\psi}_t^k(x_t, w_{t:t+k-1}; F)_{v_0}$.
 - 4: At time step $t = T - k$, observe current state x_t and receive predictions $\vartheta_{t:T-1}$.
 - 5: Solve and commit control actions $u_{t:T-1} := \tilde{\psi}_t^k(x_t, w_{t:T-1}; 0)_{v_0:k-1}$.
-

Here $M(t)$ denotes the rotational inertia matrix, which is time-varying and is determined by the proportion of renewable power in total power generation at time t , and can be accurately predicted in a certain time horizon [32, 33]; L and D are known system parameters. Using standard discretization techniques, we can formulate a discrete-time linear time-varying system $x_{t+1} = A_t x_t + B_t u_t + w_t$, where A_t and B_t are determined by $\hat{A}(t)$ and $\hat{B}(t)$. The cost functions are quadratic costs which penalizes frequency deviation [8]. This setting fits into our predictive control algorithm, since the controllers have accurate predictions of A_t and B_t in the near future due to predictability of $M(t)$.

2.2 Predictive Control

We study a classical predictive control (PC) algorithm inspired by model predictive control, which solves the optimization problem of k future time steps (where k is called the *prediction window*). Specifically, the algorithm receives the dynamics and disturbances of the next k time steps, calculates the optimal solution, and then applies the first control action of the optimal solution. The PC algorithm with prediction window k is denoted as PC_k .

More formally, At time step $t < T - k$, PC_k solves the optimization problem $\tilde{\psi}_t^k(x_t, w_{t:t+k-1}; F)$. Since we need to consider horizon lengths other than k , for arbitrary $p \geq 1$ and time step t , we define the optimization problem $\tilde{\psi}_t^p(x, \zeta; F)$ as

$$\begin{aligned} \tilde{\psi}_t^p(x, \zeta; F) := \arg \min_{y_0:p, v_0:p-1} & \sum_{\tau=1}^p f_{t+\tau}(y_\tau) + \sum_{\tau=1}^p c_{t+\tau}(v_{\tau-1}) + F(y_k) \\ \text{s.t. } & y_\tau = A_{t+\tau-1} y_{\tau-1} + B_{t+\tau-1} v_{\tau-1} + \zeta_{\tau-1}, \tau = 1, \dots, p, \\ & y_0 = x, \end{aligned} \quad (2)$$

where $x \in \mathbb{R}^n$ is the initial state, $\zeta \in (\mathbb{R}^n)^p$ (indexed by $0, \dots, p-1$) is a sequence of disturbances, and $F: \mathbb{R}^n \rightarrow \mathbb{R}$ is a standard terminal cost function regularizing the final state. Here we additionally require that the terminal cost F has the form $F(x) = \alpha(\|x\|)$, where $\alpha: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a convex K-function (i.e. continuous increasing function with 0 at the origin, see Appendix A for definition) that is twice continuously differentiable. For each time step $\tau = 1, \dots, k$, $y_\tau \in \mathbb{R}^n$ is the predictive state, and $v_\tau \in \mathbb{R}^m$ is the predictive control action. To make the algorithm well-defined, at time step $t = T - k$, PC_k can finish the rest of the trajectory optimally by committing $u_{T-k:T-1} = \tilde{\psi}(x_{T-k}, w_{T-k:T-1}; 0)$. The pseudo-code of predictive control is given in Algorithm 1.

It is also important in our framework to study the behavior of predictive control under some fixed terminal point. So, for prediction length $p \geq 1$ and time step t , we define an auxiliary optimization problem with a strict terminal constraint $y_p = z$ as follows:

$$\begin{aligned} \psi_t^p(x, \zeta, z) := \arg \min_{y_0:p, v_0:p-1} & \sum_{\tau=1}^p f_{t+\tau}(y_\tau) + \sum_{\tau=1}^p c_{t+\tau}(v_{\tau-1}) \\ \text{s.t. } & y_\tau = A_{t+\tau-1} y_{\tau-1} + B_{t+\tau-1} v_{\tau-1} + \zeta_{\tau-1}, \tau = 1, \dots, p, \\ & y_0 = x, y_p = z, \end{aligned} \quad (3)$$

where the optimal value is denoted by $\psi_t^p(x, \zeta, z)$. We define this auxiliary optimization problem besides $\tilde{\psi}_t^p(x, \zeta; F)$ because we need to fix both the initial state and the terminal state, for example, when expressing a sub-trajectory of the offline optimal trajectory as the solution of an optimization problem. ψ also allows us to study the impact of the perturbation at the terminal state on the optimal trajectory directly, which will be useful in the proof of dynamic regret (Theorem 4.2) and competitive ratio (Theorem 4.3).

Throughout the paper, we use $\{(x_t, u_t)\}_{t=1}^T$ to denote the trajectory of predictive control, and use $\{(x_t^*, u_t^*)\}_{t=1}^T$ to denote the offline optimal trajectory (i.e., the optimal solution of (1)). We also use several standard definitions and notations in linear algebra and optimization, which we detail in Appendix A for clarity. In particular, we use vector 2-norms and induced matrix 2-norms throughout this paper unless otherwise specified.

Two standard metrics will be utilized to assess the performance of our predictive control algorithms, namely *dynamic regret* $\sup_{x(0), w} (\text{cost}(\text{ALG}) - \text{cost}(\text{OPT}))$ (the additional cost of our algorithm against the optimal algorithm) and *competitive ratio* $\sup_{x(0), w} \frac{\text{cost}(\text{ALG})}{\text{cost}(\text{OPT})}$ (the worst-case ratio of the cost of our algorithm over that of the optimal algorithm).

3 A Perturbation Approach

In order to study the regret and competitive ratio of controllers in LTV systems, we develop a new analysis based on a perturbation approach, which we introduce in this section. This approach is based on developing bounds on how much the solutions to (2) and (3) change with respect to perturbations to the initial/terminal states and the disturbance sequence. Our perturbation bounds are related to the concept of *incremental stability* defined in [34], but not exactly the same because we consider the optimal trajectory in a finite horizon whereas the incremental stability focuses on asymptotic behavior over an infinite horizon. Simply stated, the key to our approach is to derive the perturbation bound in Theorem 3.3, which states that if the target variable we are concerned with is the h -th predictive state/control input, while the perturbation occurs at the τ -th time step, then the impact on the target variable is exponentially small with respect to the time difference $|h - \tau|$.

Proving such a result directly is challenging because of the complexity of the LTV dynamical constraints in (2) and (3). Thus, we develop a novel reduction from LTV systems to fully-actuated systems, i.e., systems where the controller can steer the system to any state in the whole space \mathbb{R}^n freely at every time step. This special case is a form of online optimization called *smoothed online convex optimization* (SOCO), and has received considerable attention recently, e.g., [24, 27, 35]. We exploit the controllability of the dynamics to analyze the LTV system in chunks of d time steps. A sequence of d time steps combined together can be thought as a fully-actuated system and thus we can formulate a SOCO problem, which is $(1/d)$ -times as long as the original LTV system. In this section, we first show the perturbation bound for SOCO in Section 3.1, and then we leverage our reduction to derive a result for general LTV systems in Section 3.2.

3.1 Smoothed Online Convex Optimization

The classic setting of SOCO is an online game played by an agent against an adversary: at each time step t , the adversary reveals a hitting cost function \hat{f}_t , a switching cost function \hat{c}_t , and a disturbance (or exogenous input) \hat{w}_t . The agent picks a decision point $\hat{x}_t \in \mathbb{R}^n$, and incurs a stage cost of $\hat{f}_t(\hat{x}_t) + \hat{c}_t(\hat{x}_t, \hat{x}_{t-1}, \hat{w}_{t-1})$. The agent seeks to minimize the total cost it incurs throughout the game. The offline optimal cost is defined as the minimum cost if the agent has full knowledge of the costs and disturbances at the start of the game. Instead of analyzing the performance of an online algorithm directly, our focus is on studying how the perturbations of the system parameters (initial state, terminal state, and disturbances) impact the offline optimal solution. These results are critical for deriving the guarantees for predictive control in the online setting in Section 4.

To begin, observe that when the initial state \hat{x}_0 , terminal state \hat{x}_p , and the disturbances \hat{w} are given, the optimal p -step trajectory of SOCO can be obtained from the unconstrained optimization problem

$$\hat{\psi}(\hat{x}_0, \hat{w}, \hat{x}_p) := \arg \min_{\hat{x}_{1:p-1}} \sum_{\tau=1}^{p-1} \hat{f}_\tau(\hat{x}_\tau) + \sum_{\tau=1}^p \hat{c}_\tau(\hat{x}_\tau, \hat{x}_{\tau-1}, \hat{w}_{\tau-1}), \quad (4)$$

where the objective is a convex function of the decision variables $\hat{x}_{1:p-1}$. Since (4) is an unconstrained optimization problem, the gradient of its objective equals zero at $\hat{\psi}(\hat{x}_0, \hat{w}, \hat{x}_p)$. Using this, we can further show that the directional derivative of $\hat{\psi}(\hat{x}_0, \hat{w}, \hat{x}_p)$ along some direction e , denoted by χ , satisfies the linear equation $M\chi = \delta$, where symmetric matrix M is the Hessian of the objective and vector δ is determined by the direction e . A special structure of the objective of (4) is that the correlations only occur in two consecutive time steps. This implies that its Hessian M is block

tri-diagonal. Such tri-diagonal structure of M has been noted by previous work, e.g. [36], and have been leveraged to solve the linear equation $M\chi = \delta$ quickly. In contrast, we focus on the exponential decay phenomena M^{-1} exhibits, i.e., the magnitudes of entries decay exponentially with respect to their distances to the main diagonal [37]. Bounding each entry of $\chi = M^{-1}\delta$ separately gives us the following perturbation bound. We state this result formally in Theorem 3.1, and its proof can be found in Appendix B.

Theorem 3.1. *Given a tuple $(\hat{x}_0, \hat{w}, \hat{x}_p)$ that contains the initial state, the disturbances, and the terminal state in this order, we consider the optimal solution of the SOCO problem*

$$\hat{\psi}(\hat{x}_0, \hat{w}, \hat{x}_p) := \arg \min_{\hat{x}_{1:p-1}} \sum_{\tau=1}^{p-1} \hat{f}_\tau(\hat{x}_\tau) + \sum_{\tau=1}^p \hat{c}_\tau(\hat{x}_\tau, \hat{x}_{\tau-1}, \hat{w}_{\tau-1})$$

indexed by $1, \dots, p-1$. Assume $\hat{f}_\tau : \mathbb{R}^n \rightarrow \mathbb{R}$ is μ -strongly convex, $\hat{c}_\tau : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^r \rightarrow \mathbb{R}$ is convex and ℓ -strongly smooth, and both are twice continuously differentiable for $\tau = 1, \dots, p$, then

$$\|\hat{\psi}(\hat{x}_0, \hat{w}, \hat{x}_p)_h - \hat{\psi}(\hat{x}'_0, \hat{w}', \hat{x}'_p)_h\| \leq C_0 \left(\lambda_0^{h-1} \|\hat{x}_0 - \hat{x}'_0\| + \sum_{\tau=0}^{p-1} \lambda_0^{|h-\tau|-1} \|\hat{w}_\tau - \hat{w}'_\tau\| + \lambda_0^{p-h-1} \|\hat{x}_p - \hat{x}'_p\| \right)$$

for all $1 \leq h \leq p-1$, where $C_0 = (2\ell)/\mu$ and $\lambda_0 = 1 - 2 \cdot \left(\sqrt{1 + (2\ell/\mu)} + 1 \right)^{-1}$.

As a remark, we do not require the hitting cost \hat{f}_τ to be strongly smooth, or the switching cost \hat{c}_τ to be strongly convex in Theorem 3.1. This makes the assumptions on the SOCO costs $\hat{f}_\tau, \hat{c}_\tau$ weaker than the assumptions on the LTV costs f_τ, c_τ defined in (1).

3.2 Linear Time-Varying System

We now build upon the SOCO perturbation result to derive a perturbation result for LTV systems. In particular, we show an exponentially-decaying perturbation bound for our LTV system by reducing it to SOCO and apply Theorem 3.1. As we have discussed, LTV systems are more difficult than SOCO because the dynamics prevent the online agent from picking the next state x_{t+1} freely at a given state x_t . We overcome this obstacle by redefining the decision points as illustrated in Figure 1. Specifically, given state x_t at time step t as the last decision point, we then ask the online agent to decide state x_{t+d} at time step $(t+d)$ rather than x_{t+1} at time step $(t+1)$.

Since d is the controllability index, x_{t+d} can be picked freely from the whole space \mathbb{R}^n regardless of x_t . We also utilize the *principle of optimality*, e.g. if $y_{0:k}, v_{0:k-1}$ is the optimal solution to $\psi_t^k(x, \xi, z)$, then $y_{i:j}, v_{i:j-1}$ is the optimal solution to $\psi_{t+i}^{j-i}(y_i, \xi_{i:j-1}, y_j)$ for any $0 \leq i < j \leq k$. Therefore, the trajectory between time t and $(t+d)$ can be recovered by solving $\psi_t^d(x_t, w_{t:t+d-1}, x_{t+d})$. So we are able to formulate a valid SOCO problem on the sequence of time steps $t, t+d, t+2d, \dots$.

Naturally, the hitting cost at time step $(t+d)$ remains the same, while the switching cost becomes $\xi_t^d(x_t, w_{t:t+d-1}, x_{t+d})$, where the function ξ_t^d is defined as

$$\xi_t^d(x, \zeta, z) := \psi_t^d(x, \zeta, z) - f_{t+p}(z). \quad (5)$$

An illustration of the reduction can be found in Figure 1. We would like to point out that our reduction from optimal control to SOCO is novel in that it leverages the principle of optimality to apply to more general LTV settings, as opposed to the reduction via control canonical forms in [3] that is specific to LTI systems. Unlike the switching costs in [14, 27, 35, 38] which are explicitly defined as the ℓ_2 -distance or squared ℓ_2 -distance, the switching cost ξ_t^d here is defined implicitly as the optimal value of an optimization problem. Lemma 3.2 shows that the switching cost defined in (5) satisfies the requirements of Theorem 3.1, which allows us to obtain the desired perturbation bound.

Lemma 3.2. *Under Assumption 2.1 and 2.2, for integer $p \geq d$, we have*

1. $\psi_t^d(x, \zeta, z)$ is $L_1(p)$ -Lipschitz in (x, ζ, z) ;
2. $\xi_t^d(x, \zeta, z)$ is convex and $L_2(p)$ -strongly smooth in (x, ζ, z) .

Here $L_1(p) = C(p)(1 + \ell \cdot C(p)/m_c)$, $L_2(p) = \ell \cdot C(p)^2 + \ell^2 \cdot C(p)^4/m_c$, where $\ell = \max(\ell_f, \ell_c)$,

$$C(p) = \begin{cases} O(a^{3p}) & \text{if } a > 1; \\ O(p^2) & \text{if } a = 1; \\ O(1) & \text{if } a < 1. \end{cases}$$




Figure 1: Illustration of the reduction from LTV to SOCO. Here we consider a simple example where $t = 0$ and $p = vd$. At time step 0, the agent cannot steer the system to an arbitrary target state at the next time step due to dynamical constraints. However, given (d, σ) -uniform controllability, the controller is able to enforce an arbitrary target state after d time steps, which prompts the transformation to a SOCO problem with a decision point in every d time steps.

In Lemma 3.2, we use $O(\cdot)$ to hide quantities a, b , and $1/\sigma$; the precise expression of $C(p)$ and the proof of Lemma 3.2 can be found in Appendix C. Using the reduction from LTV to SOCO, we obtain a perturbation bound for the LTV systems (2) and (3) in Theorem 3.3, the proof of which is deferred to Appendix D.

Theorem 3.3. *Consider the optimization problem defined in (2) and (3) and with a horizon length $p \geq d$. Under Assumptions 2.1 and 2.2, given any (x, ζ, z) and (x', ζ', z') ,*

$$\begin{aligned} \|\tilde{\psi}_t^p(x, \zeta; F)_{y_h} - \tilde{\psi}_t^p(x', \zeta'; F)_{y_h}\| &\leq C \left(\lambda^h \|x - x'\| + \sum_{\tau=0}^{p-1} \lambda^{h-\tau} \|\zeta_\tau - \zeta'_\tau\| \right) \\ \|\psi_t^p(x, \zeta, z)_{y_h} - \psi_t^p(x', \zeta', z')_{y_h}\| &\leq C \left(\lambda^h \|x - x'\| + \sum_{\tau=0}^{p-1} \lambda^{h-\tau} \|\zeta_\tau - \zeta'_\tau\| + \lambda^{p-h} \|z - z'\| \right) \end{aligned}$$

hold for all time steps t . Here we define $L_0 = \max_{d \leq p \leq 2d-1} L_2(p)$, and the constants are given by

$$\lambda = \left(1 - 2 \left(\sqrt{1 + (2L_0/m_c)} + 1 \right)^{-1} \right)^{\frac{1}{2d-1}}, C = \frac{2L_0}{m_c} \cdot \left(1 - 2 \left(\sqrt{1 + (2L_0/m_c)} + 1 \right)^{-1} \right)^{-1}.$$

Theorem 3.3 allows us to bound the distance between any two trajectories so long as they can be expressed as the optimal solutions of the optimization problem (2) or (3). For example, to bound the norm of each state in the predictive trajectory $\tilde{\psi}_t^p(x, \zeta; F)$, we only need to set $x' = 0, \zeta' = 0$ in the first inequality because an all zero trajectory can be expressed as $\tilde{\psi}_t^p(0, 0; F)$. The formal statement of this result can be found in Appendix E.

4 Performance Guarantees for Predictive Control

We now demonstrate the power of the perturbation approach in Section 3.2 by obtaining bounds on regret and competitive ratio. The key intuition behind our analysis is the following: at time step t , if the predictive controller with prediction window k is given the knowledge of x_t^* and x_{t+k}^* , it can fully recover the offline optimal states and control inputs for the future k time steps, x_{t+1}^* and $u_{t:t+k-1}^*$, from $\psi_t^k(x_t^*, w_{t:t+k-1}, x_{t+k}^*)$. However, without the knowledge of the offline optimal states, the predictive controller solves $\psi_t^k(x_t, w_{t:t+k-1}, x_{t+k})$ instead, where x_{t+k} is implicitly determined by the k -th predictive state of $\tilde{\psi}_t^k(x_t, w_{t:t+k-1}; F)$. We overcome this gap with our perturbation approach (specifically, Theorem 3.3 and its corollaries), which allows us to bound the distance between the controller's trajectory and the offline optimal trajectory.

4.1 Dynamic Regret

We first bound the dynamic regret of predictive control. For this analysis, a key observation is that the offline optimal trajectory is given by $x^* = \tilde{\psi}_0^T(x_0, w_{0:T-1}; 0)_{y_{1:T}}$. Furthermore, the optimal trajectory starting at time step t with state x_t is equivalent to the trajectory of predictive control with prediction window $(T-t)$ and no terminal cost, i.e. $\tilde{\psi}_t^{T-t}(x_t, w_{t:T-1}; 0)_{y_{1:T-t}}$. Using Theorem 3.3, we can bound the change in decision points against the change in prediction window k . Lemma 4.1 formalizes this:

Lemma 4.1. *For any integers p, h such that $p \geq h \geq 1$ and time step $t < T - p$, we have*

$$\left\| \tilde{\psi}_t^p(x_t, w_{t:t+p-1}; F)_{y_h} - \tilde{\psi}_t^{p+1}(x_t, w_{t:t+p}; F)_{y_h} \right\| \leq 2C\lambda^{p-h} \left(C\lambda^p \|x_t\| + \frac{2C}{1-\lambda} \sup_{0 \leq \tau \leq T-1} \|w_\tau\| \right).$$

Then, we can follow the road map below to bound the dynamic regret of predictive control PC_k :

- Given the well-conditioned state/control costs, it suffices to bound the distance between PC_k 's trajectory and the offline optimal trajectory (i.e., $\|x_t - x_t^*\|$) to show the dynamic regret result. See inequalities (22) and (23) in Appendix H for technical details.
- At each time step t , the optimal next state (under an imaginary terminal cost F) from the current state x_t is given by $\tilde{\psi}_t^{T-t}(x_t, w_{t:T-1}; F)_{y_1}$. However, reaching the optimal next state from x_t requires full knowledge of the future costs, dynamics, and disturbances. Although PC_k cannot reach the optimal next state due to incomplete knowledge of the future, it can leverage the predictions of future k steps to decide a near-optimal control action from state x_t . By cumulatively summing up the bounded difference in Lemma 4.1 and applying Theorem 3.3, one can show the suboptimality, measured by the distance $\|x_{t+1} - \tilde{\psi}_t^{T-t}(x_t, w_{t:T-1}; F)_{y_1}\|$, is in the order of $O(\lambda^k)$. See inequality (19) in Appendix H for technical details.
- Using the exponentially-decaying LTV perturbation bound in Theorem 3.3, we can convert the per-step suboptimality bounds to a global suboptimality bound on $\|x_t - x_t^*\|$ that is also in the order of $O(\lambda^k)$. See inequalities (20) and (21) in Appendix H for technical details.

The $O(\lambda^k)$ upper bound on the distance between the algorithm's trajectory and the offline optimal trajectory leads to the regret bound in Theorem 4.2.

Theorem 4.2. *Suppose $\|w_t\| \leq D$ for some constant D at each time step t . Let λ, C, L_0 be the decay rate and constants defined in Theorem 3.3. If prediction window $k \geq d$ is sufficiently large, such that*

$$k \geq 1 + \log \left(\frac{1}{1-\delta} \cdot C \left(\frac{2C}{1-\lambda} + \lambda \right) \right) / \log \left(\frac{1}{\lambda} \right) \quad (6)$$

for some positive constant $\delta \in (0, 1)$, then the trajectory of PC_k satisfies:

- (Input-to-state Stability) The norm of each state x_t is upper bounded by

$$\|x_t\| \leq \begin{cases} \frac{C}{\delta} \cdot (1-\delta)^{\max(0, t-k)} \|x_0\| + \frac{2C}{\delta(1-\lambda)} \left(1 + \frac{2C}{1-\lambda} \right) D & \text{if } 0 < t \leq T-k \\ \frac{C^2}{\delta} \cdot (1-\delta)^{T-2k} \lambda^{t+k-T} \|x_0\| + \left(\frac{2C^2}{\delta(1-\lambda)} \left(1 + \frac{2C}{1-\lambda} \right) + \frac{2C}{1-\lambda} \right) D & \text{if } T-k < t \leq T. \end{cases}$$

- (Dynamic Regret) The dynamic regret of PC_k is upper bounded by

$$\text{cost}(PC_k) - \text{cost}(OPT) = O \left(\left(D + \frac{\lambda^k (\|x_0\| + D)}{\delta} \right)^2 \lambda^k T + \lambda^k \|x_0\|^2 \right),$$

where the notation hides quantities $a, b', \ell_f, \ell_c, C, 1/(1-\lambda)$ and L_0 .

An implication of Theorem 4.2 is that to obtain $o(1)$ dynamic regret when the norm of disturbances are uniformly upper bounded, it suffices to use a prediction window of length $\Theta(\log T)$. This parallels the result shown in [4], although in a more general setting.

4.2 Competitive Ratio

We now focus on bounding the competitive ratio of predictive control. Here, we study a modification of the predictive control algorithm we have considered to this point. In particular, we introduce a

Algorithm 2 Predictive Control with Replan Window h ($PC_{(k,h)}$)

- 1: Suppose $T = n_0h + m_0$, where integers $n_0 \geq 0, k - h + 1 \leq m_0 \leq k$.
 - 2: **for** $t = 0, h, \dots, n_0(h - 1)$ **do**
 - 3: Observe current state x_t and receive predictions $\vartheta_{t:t+k-1}$.
 - 4: Solve and commit control actions $u_{t:t+h-1} := \tilde{\psi}_t^k(x_t, w_{t:t+k-1}; F)_{v_{0:h-1}}$.
 - 5: At time step $t = n_0h$, observe current state x_t and receive predictions $\vartheta_{t:T-1}$.
 - 6: Solve and commit control actions $u_{t:T-1} := \tilde{\psi}_t^{m_0}(x_t, w_{t:T-1}; 0)_{v_{0:m_0-1}}$.
-

replan window h , as defined in Algorithm 2 which we denote as $PC_{(k,h)}$. This style of algorithm has been considered previously in the SOCO literature, where it has been shown to obtain a constant competitive ratio in some settings where MPC does not [39].

Our analysis approach highlights why this modification is beneficial for competitive ratio. Specifically, we obtain the competitive ratio bound by applying a potential method building on [40]. We define the potential function as the squared distance between the algorithm’s trajectory and the offline optimal trajectory, i.e., $\phi_t(x_t, x_t^*) = \|x_t - x_t^*\|^2$, which is standard in the literature [5, 14, 27]. We study how this potential function changes over time. Intuitively, we need to upper bound the increment of this potential function by the offline optimal cost to obtain a competitive ratio result. To achieve this, the algorithm needs to “move closer” to the offline optimal trajectory rather than “moving further away” from it. Recall that Theorem 3.3 gives that

$$\|\psi_t^k(x_t, w_{t:t+k-1}; F)_{y_h} - \psi_t^k(x_t^*, w_{t:t+k-1}; F)_{y_h}\| \leq C\lambda^h \|x_t - x_t^*\|. \quad (7)$$

When the algorithm commits the first predictive state ($h = 1$), the left hand side of (7) might be larger than $\|x_t - x_t^*\|$ when $C\lambda > 1$. Thus, the algorithm must “wait” until the right hand side of (7) becomes smaller than $\|x_t - x_t^*\|$. This is accomplished in Algorithm 2 via the replan window h .

Our main result for this section is the following competitive ratio bound for $PC_{(k,h)}$.

Theorem 4.3. *Let λ, C, L_0 be the decay rate and constants defined in Theorem 3.3. In Algorithm 2, if the replan window h satisfies $h \geq \max\{\log((1 + \varepsilon)C)/\log(1/\lambda), d\}$ for some positive constant ε , and the prediction window k satisfies $k \geq h + d$, then it has competitive ratio*

$$\sup_{x(0), w} \frac{\text{cost}(PC_{(k,h)})}{\text{cost}(OPT)} = 1 + O\left(\varepsilon^{-1} \left(\frac{L_0 + \ell_f}{m_f}\right)^{1/2} \cdot C\lambda^{k-1-h}\right),$$

where the notation only hides a small numerical constant.

Note that when the constant ε and the replan window h are fixed, the competitive ratio is on the order of $1 + O(\rho^k)$ as the length of prediction k tends to infinity. One potential line of future work is to understand if the replan window is necessary. It may be possible to either strengthen the constants given in Theorem 3.3 or improve our proof approach so as to eliminate the requirement on h .

References

- [1] Naman Agarwal, Elad Hazan, and Karan Singh. Logarithmic regret for online control. *Advances in Neural Information Processing Systems*, 32, 2019.
- [2] Max Simchowitz and Dylan Foster. Naive exploration is optimal for online LQR. In *International Conference on Machine Learning*, pages 8937–8948. PMLR, 2020.
- [3] Yingying Li, Xin Chen, and Na Li. Online optimal control with linear dynamics and predictions: Algorithms and regret analysis. *Advances in Neural Information Processing Systems*, 32, 2019.
- [4] Chenkai Yu, Guanya Shi, Soon-Jo Chung, Yisong Yue, and Adam Wierman. The power of predictions in online control. *Advances in Neural Information Processing Systems*, 33, 2020.
- [5] Guanya Shi, Yiheng Lin, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Online optimization with memory and competitive control. In *Advances in Neural Information Processing Systems*, pages 20636–20647, 2020.

- [6] Chenkai Yu, Guanya Shi, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Competitive control with delayed imperfect information. *arXiv preprint arXiv:2010.11637*, 2020.
- [7] Runyu Zhang, Yingying Li, and Na Li. On the regret analysis of online LQR control with predictions. *arXiv preprint arXiv:2102.01309*, 2021.
- [8] Patricia Hidalgo-Gonzalez, Rodrigo Henriquez-Auba, Duncan Callaway, and Claire Tomlin. Frequency regulation using data-driven controllers in power grids with variable inertia due to renewable energy. In *2019 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5, 2019.
- [9] Guannan Qu, Yuanyuan Shi, Sahin Lale, Anima Anandkumar, and Adam Wierman. Stable online control of linear time-varying systems. *arXiv preprint arXiv:2102.01309*, 2021.
- [10] Jorge Nocedal and Stephen J Wright. Sequential quadratic programming. *Numerical optimization*, pages 529–562, 2006.
- [11] Daniel Morgan, Soon-Jo Chung, and Fred Y Hadaegh. Model predictive control of swarms of spacecraft using sequential convex programming. *Journal of Guidance, Control, and Dynamics*, 37(6):1725–1740, 2014.
- [12] Guanya Shi, Wolfgang Hönig, Xichen Shi, Yisong Yue, and Soon-Jo Chung. Neural-swarm2: Planning and control of heterogeneous multirotor swarms using learned interactions. *arXiv preprint arXiv:2012.05457*, 2020.
- [13] Paolo Falcone, Manuela Tufo, Francesco Borrelli, Jahan Asgari, and H Eric Tseng. A linear time varying model predictive control approach to the integrated vehicle dynamics control problem in autonomous systems. In *2007 46th IEEE Conference on Decision and Control*, pages 2980–2985. IEEE, 2007.
- [14] Gautam Goel and Adam Wierman. An online algorithm for smoothed regression and LQR control. In *Proceedings of the Machine Learning Research*, volume 89, pages 2504–2513, 2019.
- [15] Naman Agarwal, Brian Bullins, Elad Hazan, Sham Kakade, and Karan Singh. Online control with adversarial disturbances. In *International Conference on Machine Learning*, pages 111–119. PMLR, 2019.
- [16] Carlos E Garcia, David M Prett, and Manfred Morari. Model predictive control: Theory and practice—a survey. *Automatica*, 25(3):335–348, 1989.
- [17] Ugo Rosolia and Francesco Borrelli. Learning model predictive control for iterative tasks. a data-driven control framework. *IEEE Transactions on Automatic Control*, 63(7):1883–1896, 2017.
- [18] Milan Korda and Igor Mezić. Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica*, 93:149–160, 2018.
- [19] Frank Allgöwer and Alex Zheng. *Nonlinear model predictive control*, volume 26. Birkhäuser, 2012.
- [20] Moritz Diehl, Rishi Amrit, and James B Rawlings. A lyapunov function for economic optimizing model predictive control. *IEEE Transactions on Automatic Control*, 56(3):703–707, 2010.
- [21] David Angeli, Rishi Amrit, and James B Rawlings. On average performance and stability of economic model predictive control. *IEEE transactions on automatic control*, 57(7):1615–1626, 2011.
- [22] David Angeli, Alessandro Casavola, and Francesco Tedesco. Theoretical advances on economic model predictive control with time-varying costs. *Annual Reviews in Control*, 41:218–224, 2016.
- [23] Lars Grüne and Simon Pirkelmann. Economic model predictive control for time-varying system: Performance and stability results. *Optimal Control Applications and Methods*, 41(1):42–64, 2020.

- [24] Yiheng Lin, Gautam Goel, and Adam Wierman. Online optimization with predictions and non-convex losses. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(1):1–32, 2020.
- [25] Yingying Li, Guannan Qu, and Na Li. Online optimization with predictions and switching costs: Fast algorithms and the fundamental limit. *IEEE Transactions on Automatic Control*, 2020.
- [26] Minghong Lin, Zhenhua Liu, Adam Wierman, and Lachlan LH Andrew. Online algorithms for geographical load balancing. In *Proceedings of the International Green Computing Conference (IGCC)*, pages 1–10, 2012.
- [27] Gautam Goel, Yiheng Lin, Haoyuan Sun, and Adam Wierman. Beyond online balanced descent: An optimal algorithm for smoothed online optimization. *Advances in Neural Information Processing Systems*, 32:1875–1885, 2019.
- [28] Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalence is efficient for linear quadratic control. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [29] Karl Johan Åström and Richard M Murray. *Feedback systems*. Princeton university press, 2010.
- [30] Brian DO Anderson and John B Moore. *Optimal Control: Linear Quadratic Methods*. Courier Corporation, 01 2007.
- [31] Paolo Falcone, Francesco Borrelli, H Eric Tseng, Jahan Asgari, and Davor Hrovat. Linear time-varying model predictive control and its application to active steering systems: Stability analysis and experimental validation. *International Journal of Robust and Nonlinear Control: IFAC-Affiliated Journal*, 18(8):862–875, 2008.
- [32] Madasthu Santhosh, Chintham Venkaiah, and D. M. Vinod Kumar. Current advances and approaches in wind speed and wind power forecasting for improved renewable energy integration: A review. *Engineering Reports*, 2(6):e12178, 2020.
- [33] Ruby Nageem and Jayabarathi R. Predicting the power output of a grid-connected solar panel using multi-input support vector regression. *Procedia Computer Science*, 115:723–730, 2017.
- [34] Duc N Tran, Björn S Rüffer, and Christopher M Kellett. Incremental stability properties for discrete-time systems. In *Proceedings of Conference on Decision and Control (CDC)*, pages 477–482. IEEE, 2016.
- [35] CJ Argue, Anupam Gupta, and Guru Guruganesh. Dimension-free bounds for chasing convex functions. In *Conference on Learning Theory*, pages 219–241. PMLR, 2020.
- [36] Brandon Amos, Ivan Jimenez, Jacob Sacks, Byron Boots, and J. Zico Kolter. Differentiable MPC for End-to-end Planning and Control. In *Advances in Neural Information Processing Systems*, 2018.
- [37] Stephen Demko, William F Moss, and Philip W Smith. Decay rates for inverses of band matrices. *Mathematics of computation*, 43(168):491–499, 1984.
- [38] Niangjun Chen, Gautam Goel, and Adam Wierman. Smoothed online convex optimization in high dimensions via online balanced descent. In *Proceedings of Conference On Learning Theory (COLT)*, pages 1574–1594, 2018.
- [39] Niangjun Chen, Joshua Comden, Zhenhua Liu, Anshul Gandhi, and Adam Wierman. Using predictions in online optimization: Looking forward with an eye on the past. *ACM SIGMETRICS Performance Evaluation Review*, 44(1):193–206, 2016.
- [40] Nikhil Bansal and Anupam Gupta. Potential-function proofs for gradient methods. *Theory of Computing*, 15(1):1–32, 2019.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] Our main results in Sections 3 and 4 are consistent with our abstract and introduction.
 - (b) Did you describe the limitations of your work? [Yes] We described some potential improvements to our analysis at the end of our introduction and our description of Theorem 4.3.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] The goal of this paper is to advance the theoretical understanding of meta-adaptive control. We believe such fundamental work will not have any direct societal impacts.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We did not find any aspect of our work that may lead to ethics violations.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 2.1 for our assumptions and the motivations behind them.
 - (b) Did you include complete proofs of all theoretical results? [Yes] See Appendices for complete proofs of our results. We also referenced the relevant sections of the Appendix when we introduced the Theorems.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Definitions and Notations

Definition A.1. We use the follow convention on linear algebra:

1. $\|\cdot\|$ denotes the (Euclidean) 2-norm for vectors and the induced 2-norm for matrices:

$$\|v\| = \sqrt{v_1^2 + v_2^2 + \cdots + v_n^2}, \quad v \in \mathbb{R}^n$$

$$\|A\| = \sup_{v \in \mathbb{R}^n \setminus \{0\}} \frac{\|Ax\|}{\|x\|}, \quad A \in \mathbb{R}^{m \times n};$$

2. $\sigma(A)$ is the collection of singular values of a matrix A , also known as the singular spectrum;
3. $\sigma_{\min}(A)$ denotes the smallest singular value of a matrix A ;
4. $A \geq 0$ indicates that a matrix A is positive semi-definite.

The notions of strong-convexity and smoothness are used throughout this paper:

Definition A.2. A real-valued function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is called ℓ -strongly smooth if

$$g(y) \leq g(x) + \langle \nabla g(x), y - x \rangle + \frac{\ell}{2} \|y - x\|_2^2$$

and is called m -strongly convex if

$$g(y) \geq g(x) + \langle \nabla g(x), y - x \rangle + \frac{m}{2} \|y - x\|_2^2$$

for any $x, y \in \mathbb{R}^n$. Here $\langle \cdot, \cdot \rangle$ denotes the standard inner product of vectors.

We also require the terminal cost to be a K-function, the definition of which is given below.

Definition A.3. A function $g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is said to be a K-function (or belongs to class K), if it is continuous, strictly increasing, and satisfies $g(0) = 0$.

For ease of reference, we summarize in the following table all the notation used in the paper.

Notation	Meaning
A_t, B_t	dynamical matrices of the system at time step t
x_t, u_t, w_t	state, control action, and disturbance at time step t
f_t, c_t	hitting cost function and control cost function at time step t
m_f, m_c	strong convexity parameters of cost functions
ℓ_f, ℓ_c	strong smoothness parameters of cost functions
$\tilde{\psi}_t^p(x, \zeta; F)$	optimal trajectory [†] from step t in the future p steps (free terminal state)
$\tilde{v}_t^p(x, \zeta; F)$	optimal value from step t in the future p steps (free terminal state)
$\psi_t^p(x, \zeta, z)$	optimal trajectory [†] from step t in the future p steps (fixed terminal state z)
$v_t^p(x, \zeta, z)$	optimal value from step t in the future p steps (fixed terminal state z)
$y_\tau, v_\tau, \zeta_\tau$	(in ψ and $\tilde{\psi}$) predictive state, control action, and disturbance within the optimization
F, z	terminal constraint function F (in ψ) or fixed terminal state z (in $\tilde{\psi}$)
$\{(x_t, u_t)\}_{t=1}^T$	the trajectory of our predictive control algorithm
$\{(x_t^*, u_t^*)\}_{t=1}^T$	the offline optimal trajectory (i.e., $\tilde{\psi}_0^T(x(0), w; 0)$)
$\hat{\psi}(\hat{x}_0, \hat{w}, \hat{x}_p)$	the converted SOCO trajectory with initial state \hat{x}_0 , disturbances \hat{w} , and terminal state \hat{x}_p
$\hat{f}_\tau, \hat{c}_\tau, \hat{x}_{1:p-1}$	hitting cost, transition cost, and optimal trajectory of the converted SOCO problem

[†] Trajectory contains both states and control inputs, which are referred to by subscripts y_τ and u_τ , respectively.

B Proof of Theorem 3.1

In the next lemma we will use the notation A_{S_R, S_C} to denote the submatrix obtained by selecting the blocks indexed by some set $S_R \times S_C$ while preserving their relative order. Specifically, consider a

matrix $A \in \mathbb{R}^{\omega n \times \omega n}$ formed by $\omega \times \omega$ blocks $A_{i,j} \in \mathbb{R}^{\omega \times \omega}$. Let $i_1 < \dots < i_{|S_R|}$ be the elements in $S_R \subseteq \{1, \dots, \omega\}$, and $j_1 < \dots < j_{|S_C|}$ be the elements in $S_C \subseteq \{1, \dots, \omega\}$, both in ascending order. Then $A_{S_R, S_C} \in \mathbb{R}^{|S_R| \times |S_C|}$ is defined as a block matrix

$$A_{S_R, S_C} := \begin{bmatrix} A_{i_1, j_1} & A_{i_1, j_2} & \cdots & A_{i_1, j_{|S_C|}} \\ A_{i_2, j_1} & A_{i_2, j_2} & \cdots & A_{i_2, j_{|S_C|}} \\ \vdots & \vdots & \ddots & \vdots \\ A_{i_{|S_R|}, j_1} & A_{i_{|S_R|}, j_2} & \cdots & A_{i_{|S_R|}, j_{|S_C|}} \end{bmatrix}.$$

For a diagonal block matrix $D = \text{diag}(D_1, \dots, D_\omega)$ and a set $S \subseteq \{1, \dots, \omega\}$, we use the shorthand notation $D_S := \text{diag}(D_{i_1}, D_{i_2}, \dots, D_{i_{|S|}})$, where $i_1 < \dots < i_{|S|}$ are the elements in S .

Lemma B.1. *Suppose A is a positive definite matrix in $\mathbb{S}^{\omega n}$ formed by $\omega \times \omega$ blocks $A_{i,j} \in \mathbb{R}^{\omega \times \omega}$. Assume that A is q -banded for an even positive integer q ; that is*

$$A_{i,j} = 0, \forall |i - j| > q/2.$$

Let $[a_0, b_0]$ ($b_0 > a_0 > 0$) be the smallest interval containing the spectrum $\sigma(A)$. Suppose $D = \text{diag}(D_1, \dots, D_\omega)$, where $D_i \in \mathbb{S}^{\omega}$ is positive semi-definite. Let $M = ((A + D)^{-1})_{S_R, S_C}$ as defined above, where $S_R, S_C \subseteq \{1, \dots, \omega\}$. Then we have $\|M\| \leq C\gamma^{\hat{d}}$, where

$$C = \frac{2}{a_0}, \gamma = \left(\frac{\sqrt{\text{cond}(A)} - 1}{\sqrt{\text{cond}(A)} + 1} \right)^{2/q}, \hat{d} = \min_{i \in S_R, j \in S_C} |i - j|.$$

Here $\text{cond}(A) = b_0/a_0$ denotes the condition number of matrix A .

Proof of Lemma B.1. We first prove the lemma for the special case where $D = 0$.

For the case $\hat{d} \neq 0$, write $\hat{d} = vq/2 + \kappa$ for integers v, κ satisfying $v \geq 0, 1 \leq \kappa \leq q/2$. Following the same approach as the proof of Proposition 2.2 in [37], we see that there exists a polynomial p_v of degree v , where

$$\|A^{-1} - p_v(A)\| \leq \frac{1}{a_0} \cdot \frac{(1 + \sqrt{\text{cond}(A)})^2}{2 \text{cond}(A)} \gamma^{\hat{d}} \leq C\gamma^{\hat{d}},$$

where the last inequality holds because $\text{cond}(A) \geq 1$.

Since p_v has degree $v < \frac{2\hat{d}}{q}$ and A is q -banded, the matrix $p_v(A)$ satisfies $(p_v(A))_{i,j} = 0$ for any $i \in S_R$ and $j \in S_C$. We then obtain

$$\|P\| = \left\| (A^{-1})_{S_R, S_C} \right\| = \left\| (A^{-1} - p_v(A))_{S_R, S_C} \right\| \leq \|A^{-1} - p_v(A)\| \leq C\gamma^{\hat{d}},$$

because 2-norm of a submatrix cannot be larger than that of the original matrix.

For the case $\hat{d} = 0$, as $\|P\| = \left\| (A^{-1})_{S_R, S_C} \right\| \leq \|A^{-1}\| = \frac{1}{a_0} \leq C$, the result trivially holds.

Now we show the general case (where $D_i \geq 0$ for $1 \leq i \leq n$) through a reduction to the special case. Define a positive definite matrix $N := (a_0 I + D) \in \mathbb{S}^{n\omega}$, and then define matrix $H \in \mathbb{S}^{n\omega}$ as follows,

$$H = N^{-\frac{1}{2}}(A + D)N^{-\frac{1}{2}}.$$

We start by showing that $I \leq H \leq \frac{b_0}{a_0} \cdot I$. For any $x \in \mathbb{R}^{n\omega}$, we observe

$$\begin{aligned} x^\top H x &= x^\top N^{-\frac{1}{2}} A N^{-\frac{1}{2}} x + x^\top N^{-\frac{1}{2}} D N^{-\frac{1}{2}} x \\ &\geq x^\top N^{-\frac{1}{2}} a_0 I N^{-\frac{1}{2}} x + x^\top N^{-\frac{1}{2}} D N^{-\frac{1}{2}} x \\ &= x^\top N^{-\frac{1}{2}} (a_0 I + D) N^{-\frac{1}{2}} x \\ &= \|x\|^2. \end{aligned}$$

For the other inequality, we have

$$\begin{aligned}
x^\top Hx &= x^\top N^{-\frac{1}{2}} AN^{-\frac{1}{2}} x + x^\top N^{-\frac{1}{2}} DN^{-\frac{1}{2}} x \\
&\leq x^\top N^{-\frac{1}{2}} b_0 I N^{-\frac{1}{2}} x + x^\top N^{-\frac{1}{2}} DN^{-\frac{1}{2}} x \\
&= x^\top N^{-\frac{1}{2}} (a_0 I + D) N^{-\frac{1}{2}} x + (b_0 - a_0) x^\top N^{-1} x \\
&\leq \|x\|^2 + \frac{b_0 - a_0}{a_0} \cdot \|x\|^2 \\
&= \frac{b_0}{a_0} \cdot \|x\|^2.
\end{aligned}$$

Thus $I \leq H \leq \frac{b_0}{a_0} \cdot I$, which gives $\text{cond}(H) \leq \frac{b_0}{a_0} = \text{cond}(A)$. Note that H is also q -banded, so we can apply the result of the special case ($D_i = 0, i = 1, \dots, n$) to obtain that

$$\|(H^{-1})_{S_R, S_C}\| \leq 2\gamma_H^{\hat{d}} \leq 2\gamma^{\hat{d}},$$

where $\gamma_H = \left(\frac{\sqrt{\text{cond}(H)-1}}{\sqrt{\text{cond}(H)+1}} \right)^{2/q} \leq \gamma$. Using this inequality, we conclude that

$$\begin{aligned}
\|P\| &= \|((A+D)^{-1})_{S_R, S_C}\| = \left\| \left(N^{-\frac{1}{2}} H^{-1} N^{-\frac{1}{2}} \right)_{S_R, S_C} \right\| \\
&\leq \left\| (a_0 I + D_{S_R})^{-\frac{1}{2}} \right\| \cdot \|(H^{-1})_{S_R, S_C}\| \cdot \left\| (a_0 I + D_{S_C})^{-\frac{1}{2}} \right\| \\
&\leq \frac{1}{a_0} \|(H^{-1})_{S_R, S_C}\| \\
&\leq C\gamma^{\hat{d}}.
\end{aligned}$$

Here we apply the fact that $\|(a_0 I + D_S)^{-\frac{1}{2}}\| \leq \frac{1}{\sqrt{a_0}}$ since $D_S \geq 0$. □

Now we return to the proof of Theorem 3.1

Proof of Theorem 3.1. Let $e = (e_0^\top, \mu^\top, e_p^\top)^\top$ be a vector where $e_0, e_p \in \mathbb{R}^n$ and

$$\mu = [\mu_0, \mu_1, \dots, \mu_{p-1}],$$

for $\mu_i \in \mathbb{R}^r, i = 0, 1, \dots, p-1$. Let θ be an arbitrary real number. Define function $\hat{h} : \mathbb{R}^{(p-1) \times n} \times \mathbb{R}^n \times \mathbb{R}^{p \times r} \times \mathbb{R}^n \rightarrow \mathbb{R}_+$ as

$$\hat{h}(\hat{x}_{1:p-1}, \hat{x}_0, \hat{w}_{0:p-1}, \hat{x}_p) = \sum_{\tau=1}^{p-1} \hat{f}_\tau(\hat{x}_\tau) + \sum_{\tau=1}^p \hat{c}_\tau(\hat{x}_\tau, \hat{x}_{\tau-1}, \hat{w}_{\tau-1}).$$

To simplify the notation, we use $\hat{\zeta}$ to denote the tuple of system parameters, i.e.,

$$\hat{\zeta} := (\hat{x}_0, \hat{w}_{0:p-1}, \hat{x}_p).$$

From our construction, we know that \hat{h} is μ -strongly convex in $\hat{x}_{1:p-1}$, so we use the decomposition $\hat{h} = \hat{h}_a + \hat{h}_b$, where

$$\begin{aligned}
\hat{h}_a(\hat{x}_{1:p-1}, \hat{\zeta}) &= \sum_{\tau=1}^{p-1} \frac{\mu}{2} \|\hat{x}_\tau\|^2 + \sum_{\tau=1}^p \hat{c}_\tau(\hat{x}_\tau, \hat{x}_{\tau-1}, \hat{w}_{\tau-1}), \\
\hat{h}_b(\hat{x}_{1:p-1}, \hat{\zeta}) &= \sum_{\tau=1}^{p-1} \left(\hat{f}_\tau(\hat{x}_\tau) - \frac{\mu}{2} \|\hat{x}_\tau\|^2 \right).
\end{aligned}$$

Since $\hat{\psi}(\hat{\zeta} + \theta e)$ is the minimizer of convex function $\hat{h}(\cdot, \hat{\zeta} + \theta e)$, we see that

$$\nabla_{\hat{x}_{1:p-1}} \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e) = 0.$$

Taking the derivative with respect to θ gives that

$$\begin{aligned} \nabla_{\hat{x}_{1:p-1}}^2 \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e) \frac{d}{d\theta} \hat{\psi}(\hat{\zeta} + \theta e) &= -\nabla_{\hat{x}_0} \nabla_{\hat{x}_{1:p-1}} \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e) e_0 \\ &\quad - \nabla_{\hat{x}_p} \nabla_{\hat{x}_{1:p-1}} \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e) e_p \\ &\quad - \sum_{\tau=0}^{p-1} \nabla_{w_\tau} \nabla_{\hat{x}_{1:p-1}} \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e) \mu_\tau. \end{aligned}$$

To simplify the notation, we define

$$M := \nabla_{\hat{x}_{1:p-1}}^2 \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e), \text{ which is a } (p-1) \times (p-1) \text{ block matrix,}$$

$$R^{(0)} := -\nabla_{\hat{x}_0} \nabla_{\hat{x}_{1:p-1}} \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e), \text{ which is a } (p-1) \times 1 \text{ block matrix,}$$

$$R^{(p)} := -\nabla_{\hat{x}_p} \nabla_{\hat{x}_{1:p-1}} \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e), \text{ which is a } (p-1) \times 1 \text{ block matrix,}$$

$$K^{(\tau)} := -\nabla_{w_\tau} \nabla_{\hat{x}_{1:p-1}} \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e), \forall 0 \leq \tau \leq p-1, \text{ which are } (p-1) \times 1 \text{ block matrices,}$$

where in M , $R^{(0)}$, $R^{(p)}$, the block size is $n \times n$; in $K^{(\tau)}$, the block size is $n \times r$. Hence we can write

$$\frac{d}{d\theta} \hat{\psi}(\hat{\zeta} + \theta e) = M^{-1} \left(R^{(0)} e_0 + R^{(p)} e_p + \sum_{\tau=0}^{p-1} K^{(\tau)} \mu_\tau \right).$$

Recall that $R^{(0)}$, $R^{(p)}$ are $(p-1) \times 1$ block matrices with block size $n \times n$. $\{K^{(\tau)}\}_{0 \leq \tau \leq p-1}$ are $(p-1) \times 1$ block matrices with block size $n \times r$. For $R^{(0)}$ and $K^{(0)}$, only the $(1, 1)$ -th blocks are non-zero. For $R^{(p)}$ and $K^{(p-1)}$, only the $(p-1, 1)$ -th blocks are non-zero. For $K^{(\tau)}$, $\tau = 1, \dots, p-2$, only the $(\tau, 1)$ -th and $(\tau+1, 1)$ -th blocks are non-zero. Hence we see that

$$\begin{aligned} \frac{d}{d\theta} \hat{\psi}(\hat{\zeta} + \theta e)_h &= (M^{-1})_{h,1} R_{1,1}^{(0)} e_0 + (M^{-1})_{h,p-1} R_{p-1,1}^{(p)} e_p \\ &\quad + (M^{-1})_{h,1} K_{1,1}^{(0)} \mu_0 + (M^{-1})_{h,p-1} K_{p-1,1}^{(p-1)} \mu_{p-1} \\ &\quad + \sum_{\tau=1}^{p-2} (M^{-1})_{h,\tau:\tau+1} K_{\tau:\tau+1,1}^{(\tau)} \mu_\tau. \end{aligned}$$

Since the switching costs $c_\tau(\cdot, \cdot, \cdot)$, $\tau = 1, \dots, p$ are ℓ -strongly smooth, we know that the norms of

$$R_{1,1}^{(0)}, R_{p-1,1}^{(p)}, K_{1,1}^{(0)}, K_{p-1,1}^{(p-1)}, \text{ and } \{K_{\tau:\tau+1,1}^{(\tau)}\}_{1 \leq \tau \leq p-2}$$

are all upper bounded by ℓ . Taking norm on both sides gives that

$$\begin{aligned} \left\| \frac{d}{d\theta} \hat{\psi}(\hat{\zeta} + \theta e)_h \right\| &\leq \ell \|(M^{-1})_{h,1}\| \|e_0\| + \ell \|(M^{-1})_{h,p-1}\| \|e_p\| \\ &\quad + \ell \|(M^{-1})_{h,1}\| \|\mu_0\| + \ell \|(M^{-1})_{h,p-1}\| \|\mu_{p-1}\| \\ &\quad + \sum_{\tau=1}^{p-2} \ell \|(M^{-1})_{h,\tau:\tau+1}\| \|\mu_\tau\|. \end{aligned} \tag{8}$$

Note that M can be decomposed as $M = M_a + M_b$, where

$$M_a := \nabla_{\hat{x}_{1:p-1}}^2 \hat{h}_a(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e),$$

$$M_b := \nabla_{\hat{x}_{1:p-1}}^2 \hat{h}_b(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e).$$

Since M_a is block tri-diagonal and satisfies $(\mu + 2\ell)I \geq M_a \geq \mu I$, and M_b is block diagonal and satisfies $M_b \geq 0$, we obtain the following with Lemma B.1:

$$\|(M^{-1})_{h,1}\| \leq \frac{2}{\mu} \lambda_0^{h-1}, \|(M^{-1})_{h,p-1}\| \leq \frac{2}{\mu} \lambda_0^{p-h-1}, \text{ and } \|(M^{-1})_{h,\tau:\tau+1}\| \leq \frac{2}{\mu} \lambda_0^{|h-\tau|-1},$$

where $\lambda_0 := (\sqrt{\text{cond}(M_a)} - 1) / (\sqrt{\text{cond}(M_a)} + 1) = 1 - 2 \cdot \left(\sqrt{1 + (2\ell/\mu)} + 1 \right)^{-1}$.

Substituting this into (8), we see that

$$\left\| \frac{d}{d\theta} \hat{\psi}(\hat{\xi} + \theta e)_h \right\| \leq C_0 \left(\lambda_0^{h-1} \|e_0\| + \sum_{\tau=0}^{p-1} \lambda_0^{|h-\tau|-1} \|\mu_\tau\| + \lambda_0^{p-h-1} \|e_p\| \right),$$

where $C_0 = (2\ell)/\mu$.

Hence we obtain

$$\begin{aligned} \|\hat{\psi}(\hat{\xi})_h - \hat{\psi}(\hat{\xi} + e)_h\| &= \left\| \int_0^1 \frac{d}{d\theta} \hat{\psi}(\hat{\xi} + \theta e)_h d\theta \right\| \\ &\leq \int_0^1 \left\| \frac{d}{d\theta} \hat{\psi}(\hat{\xi} + \theta e)_h \right\| d\theta \\ &\leq C_0 \left(\lambda_0^{h-1} \|e_0\| + \sum_{\tau=0}^{p-1} \lambda_0^{|h-\tau|-1} \|\mu_\tau\| + \lambda_0^{p-h-1} \|e_p\| \right). \end{aligned}$$

This finishes the proof. \square

C Proof of Lemma 3.2

Lemma C.1. *Suppose function $f(x, y)$ is convex and L -strongly smooth in (x, y) , μ -strongly convex in y , and continuously differentiable. Define functions $y^*(x) := \arg \min_y f(x, y)$ and $g(x) := \min_y f(x, y)$. Then, function y^* is $\frac{L}{\mu}$ -Lipschitz and function g is $\left(L + \frac{L^2}{\mu}\right)$ -strongly smooth.*

Proof of Lemma C.1. Let $y^*(x) = \arg \min_y f(x, y)$. This function is well-defined since the strong convexity of $f(x, y)$ in y guarantees that $y^*(x)$ is unique. We see that for all x, x' ,

$$\nabla_y f(x, y^*(x)) = 0 \text{ and } \nabla_y f(x', y^*(x')) = 0.$$

Using these equalities, we obtain

$$\begin{aligned} 0 &= \langle y^*(x) - y^*(x'), \nabla_y f(x, y^*(x)) - \nabla_y f(x', y^*(x')) \rangle \\ &= \langle y^*(x) - y^*(x'), \nabla_y f(x, y^*(x)) - \nabla_y f(x, y^*(x')) \rangle \\ &\quad + \langle y^*(x) - y^*(x'), \nabla_y f(x, y^*(x')) - \nabla_y f(x', y^*(x')) \rangle \\ &\geq \mu \|y^*(x) - y^*(x')\|^2 - \|y^*(x) - y^*(x')\| \cdot \|\nabla_y f(x, y^*(x')) - \nabla_y f(x', y^*(x'))\|, \end{aligned}$$

where we used the fact that a μ -strongly convex function h satisfies

$$\langle a - b, \nabla h(a) - \nabla h(b) \rangle \geq \mu \|a - b\|^2, \forall a, b$$

and the Cauchy-Schwartz inequality in the last inequality. Since f is L -strongly smooth, we see that

$$\|y^*(x) - y^*(x')\| \leq \frac{1}{\mu} \|\nabla_y f(x, y^*(x')) - \nabla_y f(x', y^*(x'))\| \leq \frac{L}{\mu} \|x - x'\|,$$

which implies function y^* is $\frac{L}{\mu}$ -Lipschitz.

Note that the gradient of g is given by

$$\nabla g(x) = \nabla_x f(x, y^*(x)) + \nabla_y f(x, y^*(x)) \frac{\partial y^*(x)}{\partial x} = \nabla_x f(x, y^*(x)),$$

because $\nabla_y f(x, y^*(x)) = 0$. Hence we obtain

$$\begin{aligned} \|\nabla g(x) - \nabla g(x')\| &\leq \|\nabla_x f(x, y^*(x)) - \nabla_x f(x', y^*(x'))\| \\ &\leq \|\nabla_x f(x, y^*(x)) - \nabla_x f(x', y^*(x))\| + \|\nabla_x f(x', y^*(x)) - \nabla_x f(x', y^*(x'))\| \\ &\leq L \|x - x'\| + L \|y^*(x) - y^*(x')\| \\ &\leq \left(L + \frac{L^2}{\mu} \right) \|x - x'\|. \end{aligned}$$

\square

Lemma C.2. Suppose A is a $\omega_1 \times \omega_2$ block matrix. Let A_{ij} denote the (i, j) th block of A , $1 \leq i \leq \omega_1, 1 \leq j \leq \omega_2$. The induced 2-norm of A is upper bounded by

$$\|A\| \leq \left(\sum_{i=1}^{\omega_1} \sum_{j=1}^{\omega_2} \|A_{ij}\|^2 \right)^{\frac{1}{2}}.$$

Proof of Lemma C.2. For unit vector x , we have the following:

$$\begin{aligned} \|Ax\|^2 &= \sum_{i=1}^{\omega_1} \left\| \sum_{j=1}^{\omega_2} A_{ij}x_j \right\|^2 \\ &\leq \sum_{i=1}^{\omega_1} \left(\sum_{j=1}^{\omega_2} \|A_{ij}\| \cdot \|x_j\| \right)^2 \\ &\leq \sum_{i=1}^{\omega_1} \left(\sum_{j=1}^{\omega_2} \|A_{ij}\|^2 \right) \left(\sum_{j=1}^{\omega_2} \|x_j\|^2 \right) \\ &= \sum_{i=1}^{\omega_1} \sum_{j=1}^{\omega_2} \|A_{ij}\|^2. \end{aligned}$$

where we used the definition of the induced 2-norm in the first inequality and the Cauchy-Schwarz inequality in the second inequality. \square

Now we come back to the proof of Lemma 3.2.

Proof of Lemma 3.2. To simplify the notation, we define the stacked state vector y , control vector v , and disturbance vector ζ as

$$y = \begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_p \end{bmatrix}, v = \begin{bmatrix} v_0 \\ v_1 \\ \vdots \\ v_{p-1} \end{bmatrix}, \zeta = \begin{bmatrix} \zeta_0 \\ \zeta_1 \\ \vdots \\ \zeta_{p-1} \end{bmatrix}.$$

Recall that the transition matrix $\Phi(t_2, t_1)$ is defined as

$$\Phi(t_2, t_1) := \begin{cases} A_{t_2-1}A_{t_2-2} \cdots A_{t_1} & \text{if } t_2 > t_1 \\ I & \text{if } t_2 \leq t_1 \end{cases}.$$

Using this, we can express the state vector y as an affine function of initial state x , control v , and disturbance ζ :

$$y = S^x x + S^v v + S^\zeta \zeta, \quad (9)$$

where

$$S^\zeta := \begin{bmatrix} 0 & 0 & \cdots & 0 \\ \Phi(t+1, t+1) & 0 & \cdots & 0 \\ \Phi(t+2, t+1) & \Phi(t+2, t+2) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \Phi(t+p, t+1) & \Phi(t+p, t+2) & \cdots & \Phi(t+p, t+p) \end{bmatrix}, S^x = \begin{bmatrix} \Phi(t, t) \\ \Phi(t+1, t) \\ \Phi(t+2, t) \\ \vdots \\ \Phi(t+p, t) \end{bmatrix},$$

and $S^v = S^\zeta \cdot \text{diag}(B_t, \dots, B_{t+p-1})$.

To simplify the notation, we use the shorthand $M := M(t, p)$ for the controllability matrix and

$$R^\zeta := [\Phi(t+p, t+1), \Phi(t+p, t+2), \dots, \Phi(t+p, t+p)].$$

throughout the proof. Since p is greater than the controllability index d , we know M has full row rank. The dynamical constraints for (5), which is identical to the constraints of (3), can be written as

$$Mv = z - \Phi(t+p, t)x - R^\zeta \zeta.$$

Because M has full row rank, we let $M^\dagger = M^\top (MM^\top)^{-1}$ be the Moore-Penrose pseudo-inverse of M . Let $V \in \mathbb{R}^{(mp) \times (mp-n)}$ be a matrix whose columns constitute an orthonormal basis of $\ker(M)$. Then, we can express all feasible control vector v as

$$v = M^\dagger \left(z - \Phi(t+p, t)x - R^\zeta \zeta \right) + Vr, \quad (10)$$

where r is a free variable that can take any value in \mathbb{R}^{mp-n} .

Let F denote the objective function of ξ_t^P , i.e.,

$$F(y, v) := \left(\sum_{\tau=1}^{p-1} f_{t+\tau}(y_\tau) + c_{t+\tau}(v_{\tau-1}) \right) + c_{t+p}(v_{p-1}).$$

Since we can express the state vector y and control vector v as linear functions of x, z, ζ and r , we can write the switching cost (5) as an unconstrained optimization problem

$$\min_{r \in \mathbb{R}^{mp-n}} F(y(x, z, \zeta, r), v(x, z, \zeta, r)), \quad (11)$$

where functions $y(x, z, \zeta, r)$ and $v(x, z, \zeta, r)$ are determined by

$$\begin{bmatrix} y \\ v \end{bmatrix} = \begin{bmatrix} S^x - S^v M^\dagger \Phi(t+p, t) & S^v M^\dagger & S^\zeta - S^v M^\dagger R^\zeta & S^v V \\ -M^\dagger \Phi(t+p, t) & M^\dagger & -M^\dagger R^\zeta & V \end{bmatrix} \cdot \begin{bmatrix} x \\ z \\ \zeta \\ r \end{bmatrix}. \quad (12)$$

Note that if $a \neq 1$, the following is due to Lemma C.2:

$$\|S^\zeta\| \leq \left(\sum_{i=1}^p \sum_{j=1}^i \|\phi(t+i, t+j)\|^2 \right)^{\frac{1}{2}} \leq \left(\sum_{i=1}^p \sum_{j=1}^i a^{2(i-j)} \right)^{\frac{1}{2}} = \frac{\sqrt{a^{2p+2} - (p+1)a^2 + p}}{|a^2 - 1|}$$

By Lemma C.2, we also have

$$\|S^x\| \leq \sqrt{\frac{a^{2p+2} - 1}{a^2 - 1}}, \|M^\dagger\| \leq \frac{b}{\sigma^2} \cdot \sqrt{\frac{a^{2p} - 1}{a^2 - 1}}, \|S^v\| \leq b \|S^\zeta\|, \|R^\zeta\| \leq \sqrt{\frac{a^{2p} - 1}{a^2 - 1}} \leq \frac{a^p - 1}{a - 1}.$$

Since the norm of a block matrix is upper bounded by the sum of norms of each block, we see that

$$\left\| \begin{bmatrix} S^x - S^v M^\dagger \Phi(t+p, t) & S^v M^\dagger & S^\zeta - S^v M^\dagger R^\zeta & S^v V \\ -M^\dagger \Phi(t+p, t) & M^\dagger & -M^\dagger R^\zeta & V \end{bmatrix} \right\| \leq C(p), \quad (13)$$

where, when $a \neq 1$,

$$C(p) = \left(\frac{b(a^{p+1} + a - 2)}{\sigma^2(a-1)} \cdot \sqrt{\frac{a^{2p} - 1}{a^2 - 1}} + \frac{1+b}{b} \right) \left(\frac{b\sqrt{a^{2p+2} - (p+1)a^2 + p}}{|a^2 - 1|} + 1 \right) + \sqrt{\frac{a^{2p+2} - 1}{a^2 - 1}} - \frac{1}{b}.$$

If $a = 1$, by Lemma C.2, we see that

$$\|S^\zeta\| \leq \left(\sum_{i=1}^p \sum_{j=1}^i \|\phi(t+i, t+j)\|^2 \right)^{\frac{1}{2}} \leq \left(\sum_{i=1}^p \sum_{j=1}^i a^{2(i-j)} \right)^{\frac{1}{2}} = \sqrt{\frac{p(p+1)}{2}}.$$

By Lemma C.2, we also see that

$$\|S^x\| \leq \sqrt{p+1}, \|M^\dagger\| \leq \frac{b}{\sigma^2} \cdot \sqrt{p}, \|S^v\| \leq b \|S^\zeta\|, \|R^\zeta\| \leq \sqrt{p}.$$

Therefore, for (13) to hold when $a = 1$, we need to set

$$C(p) = \left(\frac{b\sqrt{p}}{\sigma^2} (\sqrt{p} + 2) + 1 \right) \left(1 + b\sqrt{\frac{p(p+1)}{2}} \right) + \sqrt{p+1} \cdot \left(1 + \sqrt{\frac{p}{2}} \right).$$

Since F is convex and strongly smooth in (x, u) , and both x, u are affine functions of (y, z, r) , $F(x(y, z, r), u(y, z, r))$ is convex and $\ell \cdot C(p)^2$ -strongly smooth in (y, z, r) . Since $F(x, u)$ is m_c -strongly convex in u , by (10), we have

$$\begin{aligned} \nabla_r^2 F(x(y, z, w, r), u(y, z, w, r)) &\geq V^\top \nabla_u^2 F(x, u) V \\ &\geq m_c I, \end{aligned}$$

where we used that $\|Vv\|_2 = \|v\|_2, \forall v \in \mathbb{R}^{mp-n}$ because the columns of V are orthonormal in the last inequality. Therefore, by Lemma C.1, we know that (11) is convex and $L_2(p)$ -strongly smooth in (y, z) , where

$$L_2(p) := \ell \cdot C(p)^2 + \frac{\ell^2 \cdot C(p)^4}{m_c}.$$

By Lemma C.1, we also know that the optimal solution of (11):

$$r^*(x, z, \zeta) := \arg \min_{r \in \mathbb{R}^{mp-n}} F(y(x, z, \zeta, r), v(x, z, \zeta, r))$$

is $\ell \cdot C(p)^2/m_c$ -Lipschitz. By (12) and (13), we see that

$$\psi_t^p(x, \zeta, z) = \begin{bmatrix} S^x - S^v M^\dagger \Phi(t+p, t) & S^v M^\dagger & S^\zeta - S^v M^\dagger R^\zeta & S^v V \\ -M^\dagger \Phi(t+p, t) & M^\dagger & -M^\dagger R^\zeta & V \end{bmatrix} \cdot \begin{bmatrix} x \\ z \\ \zeta \\ r^*(x, z, \zeta) \end{bmatrix}$$

is $L_1(p)$ -Lipschitz, where

$$L_1(p) = C(p)(1 + \ell \cdot C(p)^2/m_c).$$

D Proof of Theorem 3.3

The proof of Theorem 3.3 is based on the decision-point transformation introduced in Section 3.2.

Recall that we use d to the controllability index as defined in Definition 2.1. To show the perturbation bound of $\psi_t^p(\cdot, \cdot, \cdot)_{y_h}$, suppose h and p satisfy $ud \leq h < (u+1)d$ and $p = vd + r$, where $u, v, r \in \mathbb{N}$ and $0 \leq r < d$. Now we shall select the decision points as

$$y_0, y_d, \dots, y_{(u-1)d}, y_h, y_{(u+2)d}, \dots, y_{(v-1)d}, y_p,$$

which are also denoted by $y_{i_0}, \dots, y_{i_{v-1}}$ for simplicity. Since the distance of any consecutive decision points falls in $[d, 2d)$, we can apply Lemma 3.2 to bound the strong smoothness of switching costs. In the transformed SOCO problem, the disturbance input of the $(\tau-1)$ -th time period is a vector $\hat{w}_{\tau-1} = \zeta_{i_{\tau-1}:i_\tau-1} \in \mathbb{R}^{n \times (i_\tau - i_{\tau-1})}$. Each stage cost $\xi_t^{i_\tau - i_{\tau-1}}(x_{i_{\tau-1}}, \hat{w}_{\tau-1}, x_{i_\tau})$ is convex and $L_2(i_\tau - i_{\tau-1})$ -strongly smooth by Lemma 3.2, and is thus L_0 -strongly smooth by definition. Recall that the solution of the transformed SOCO problem is denoted by $\hat{\psi}(x_t, \zeta, x_{t+p})$. Then by Theorem 3.1 we have

$$\begin{aligned} &\|\psi_t^p(x, \zeta, z)_{y_h} - \psi_t^p(x', \zeta', z')_{y_h}\| \\ &= \|\hat{\psi}(x, \zeta, z)_u - \hat{\psi}(x', \zeta', z')_u\| \\ &\leq C_0 \left(\lambda_0^{u-1} \|x - x'\|_2 + \sum_{\tau=0}^{v-2} \lambda_0^{|u-\tau|-1} \|w_\tau - w'_\tau\|_2 + \lambda_0^{(v-1)-u-1} \|z - z'\|_2 \right) \\ &= C_0 \left(\lambda_0^{u-1} \|x - x'\|_2 + \sum_{\tau=0}^{v-2} \lambda_0^{|u-\tau|-1} \sum_{j=i_\tau}^{i_{\tau+1}-1} \|\zeta_j - \zeta'_j\|_2 + \lambda_0^{(v-1)-u-1} \|z - z'\|_2 \right) \\ &\leq \frac{C_0}{\lambda_0} \left(\lambda^{iu-i_0} \|x - x'\|_2 + \sum_{\tau=0}^{v-2} \sum_{j=i_\tau}^{i_{\tau+1}-1} \lambda^{|j-i_u|} \|\zeta_j - \zeta'_j\|_2 + \lambda^{i_{v-1}-i_u} \|z - z'\|_2 \right) \\ &= C \left(\lambda^h \|x - x'\| + \sum_{\tau=0}^{p-1} \lambda^{|h-\tau|} \|\zeta_\tau - \zeta'_\tau\| + \lambda^{p-h} \|z - z'\| \right). \end{aligned}$$

The last inequality holds because each interval is of length at most $(2d-1)$. Here the constants are

$$C_0 = \frac{2L_0}{m_c}, \lambda_0 = 1 - 2 \cdot \left(\sqrt{1 + (2L_0/m_c)} + 1 \right)^{-1},$$

$$C = C_0/\lambda_0 = \frac{2L_0}{m_c} \left(1 - 2 \cdot \left(\sqrt{1 + (2L_0/m_c)} + 1 \right)^{-1} \right)^{-1}, \lambda = \left(1 - 2 \left(\sqrt{1 + (2L_0/m_c)} + 1 \right)^{-1} \right)^{\frac{1}{2d-1}}.$$

The proof of the perturbation bound of $\psi_t^P(\cdot, \cdot, \cdot)_{y_h}$ is quite similar. The only difference lies in the terminal cost, which can be addressed with the addition of a fixed auxiliary state. Specifically, we append $x_{\text{aux}} = 0$ to the end of the decision point sequence, and define a zero transition cost to the auxiliary state $\hat{c}_v(x_{t+p}, w_{v-1}, x_{\text{aux}}) \equiv 0$ (note that \hat{c}_v is trivially convex and L_0 -strongly smooth). Denote the solution of the modified version of transformed SOCO problem by $\hat{\psi}'(x_t, \zeta, x_{\text{aux}})$, then by the same argument as above, we have

$$\begin{aligned} \|\tilde{\psi}_t^P(x, \zeta)_{y_h} - \tilde{\psi}_t^P(x', \zeta')_{y_h}\| &= \|\hat{\psi}'(x, \zeta, 0)_u - \hat{\psi}'(x', \zeta', 0)_u\| \\ &\leq \dots \leq C \left(\lambda^h \|x - x'\| + \sum_{\tau=0}^{p-1} \lambda^{|h-\tau|} \|\zeta_\tau - \zeta'_\tau\| \right), \end{aligned}$$

where the constants are the same as previously defined. This finishes the proof of Theorem 3.3.

E Stability of the Optimal Trajectory

Corollary E.1 (Stability of the Optimal Trajectory). *For the predicted trajectory from solving (2) with a prediction window $p \geq d$, the norm of the h -th predictive state is bounded above by*

$$\|\tilde{\psi}_t^P(x, \zeta; F)_{y_h}\| \leq C \left(\lambda^h \|x\| + \sum_{\tau=0}^{p-1} \lambda^{|h-\tau|} \|\zeta_\tau\| \right) \leq C \lambda^h \|x\| + \frac{2C}{1-\lambda} \sup_{\tau} \|\zeta_\tau\|.$$

Proof of Corollary E.1. Note that $\tilde{\psi}_t^P(0, 0)_{y_h} = 0$. By Theorem 3.3, we see that

$$\begin{aligned} \|\tilde{\psi}_t^P(x, \zeta)_{y_h}\| &= \|\tilde{\psi}_t^P(x, \zeta)_{y_h} - \tilde{\psi}_t^P(0, 0)_{y_h}\| \\ &\leq C \left(\lambda_1^h \|x\| + \sum_{\tau=0}^{p-1} \lambda_1^{|h-\tau|} \|\zeta_\tau\| \right) \\ &\leq C \lambda_1^h \|x\| + \frac{2C}{1-\lambda_1} \sup_{\tau} \|\zeta_\tau\|, \end{aligned}$$

where the last inequality holds because

$$\sum_{\tau=0}^{p-1} \lambda_1^{|h-\tau|} \leq \frac{2}{1-\lambda_1}.$$

□

F Smoothness of the Optimal Cost

Corollary F.1. *For any time step t and integer p that satisfies $p \geq d$, function $\iota_t^P(\cdot, \zeta, \cdot)$ satisfies that*

$$\iota_t^P(x, \zeta, z) \leq (1 + \eta) \iota_t^P(x', \zeta, z') + \frac{L_0 + \ell_f}{2} \left(1 + \frac{1}{\eta} \right) \left(\|x' - x\|^2 + \|z' - z\|^2 \right), \forall x, x', \zeta, z, z'.$$

Lemma F.2. *Assume a function $g : \mathbb{R}^n \rightarrow \mathbb{R}_+$ is convex, ℓ -strongly smooth and continuously differentiable. For all $x, y \in \mathbb{R}^n$, for all $\eta > 0$, we have*

$$g(x) \leq (1 + \eta)g(y) + \frac{\ell}{2} \left(1 + \frac{1}{\eta} \right) \|x - y\|^2.$$

Proof of Lemma F.2.

$$g(x) - g(y) \leq \langle \nabla g(y), x - y \rangle + \frac{\ell}{2} \|x - y\|^2$$

$$\begin{aligned}
&\leq \frac{\eta}{2\ell} \|\nabla g(y)\|^2 + \frac{\ell}{2\eta} \|x - y\|^2 + \frac{\ell}{2} \|x - y\|^2 \\
&\leq \eta g(y) + \frac{\ell}{2} \left(1 + \frac{1}{\eta}\right) \|x - y\|^2.
\end{aligned}$$

where the second inequality follows from the generalized means inequality and the last inequality holds because

$$0 \leq g\left(y - \frac{1}{\ell} \nabla g(y)\right) \leq g(y) - \left\langle \nabla g(y), \frac{1}{\ell} \nabla g(y) \right\rangle + \frac{\ell}{2} \left\| \frac{\nabla g(y)}{\ell} \right\|^2 = g(y) - \frac{1}{2\ell} \|\nabla g(y)\|^2$$

□

Now we come back to the proof of Corollary F.1.

When $d \leq p \leq 2d - 1$, since $\xi_t^p(x, \zeta, z)$ is L_0 -strongly smooth by Lemma 3.2, we know

$$l_t^p(x, \zeta, z) = \xi_t^p(x, \zeta, z) + f_{t+p}(z)$$

is $(L_0 + \ell_f)$ -strongly smooth. Therefore, by Lemma F.2, we obtain that

$$l_t^p(x, \zeta, z) \leq (1 + \eta) l_t^p(x', \zeta, z') + \frac{L_0 + \ell_f}{2} \left(1 + \frac{1}{\eta}\right) (\|x' - x\|^2 + \|z' - z\|^2).$$

When $p = 2d$, let $x_1 := \psi_t^p(x', \zeta, z')_{y_d}$. We see that

$$\begin{aligned}
l_t^p(x, \zeta, z) &\leq l_t^d(x, \zeta_{0:d-1}, x_1) + l_{t+d}^d(x_1, \zeta_{d:2d-1}, z) \\
&\leq (1 + \eta) l_t^d(x', \zeta_{0:d-1}, x_1) + \frac{L_0 + \ell_f}{2} \left(1 + \frac{1}{\eta}\right) \|x - x'\|^2 \\
&\quad + (1 + \eta) l_{t+d}^d(x_1, \zeta_{d:2d-1}, z') + \frac{L_0 + \ell_f}{2} \left(1 + \frac{1}{\eta}\right) \|z - z'\|^2 \\
&\leq (1 + \eta) l_t^p(x', \zeta, z') + \frac{L_0 + \ell_f}{2} \left(1 + \frac{1}{\eta}\right) (\|x' - x\|^2 + \|z' - z\|^2).
\end{aligned}$$

When $p > 2d$, let $x_1 := \psi_t^p(x', \zeta, z')_{y_d}$, $x_2 := \psi_t^p(x', \zeta, z')_{y_{p-d}}$. We see that

$$\begin{aligned}
l_t^p(x, \zeta, z) &\leq l_t^d(x, \zeta_{0:d-1}, x_1) + l_{t+d}^{p-2d}(x_1, \zeta_{d:p-d-1}, x_2) + l_{t+p-d}^d(x_2, \zeta_{p-d:p-1}, z) \\
&\leq (1 + \eta) l_t^d(x', \zeta_{0:d-1}, x_1) + \frac{L_0 + \ell_f}{2} \left(1 + \frac{1}{\eta}\right) \|x - x'\|^2 \\
&\quad + l_{t+d}^{p-2d}(x_1, \zeta_{d:p-d-1}, x_2) \\
&\quad + (1 + \eta) l_{t+p-d}^d(x_2, \zeta_{p-d:p-1}, z') + \frac{L_0 + \ell_f}{2} \left(1 + \frac{1}{\eta}\right) \|z - z'\|^2 \\
&\leq (1 + \eta) l_t^p(x', \zeta, z') + \frac{L_0 + \ell_f}{2} \left(1 + \frac{1}{\eta}\right) (\|x' - x\|^2 + \|z' - z\|^2).
\end{aligned}$$

G Proof of Lemma 4.1

Proof of Lemma 4.1. To simplify the notation, we define

$$z := \tilde{\psi}_t^p(x_t; F)_{y_p}, z' := \tilde{\psi}_t^{p+1}(x_t; F)_{y_p}.$$

We see that

$$\left\| \tilde{\psi}_t^p(x_t; F)_{y_h} - \tilde{\psi}_t^{p+1}(x_t; F)_{y_h} \right\| = \left\| \psi_t^p(x_t, z)_{y_h} - \psi_t^p(x_t, z')_{y_h} \right\| \quad (14a)$$

$$\leq C\lambda^{p-h} \|z - z'\| \quad (14b)$$

$$\leq 2C\lambda^{p-h} \left(C\lambda^p \|x_t\| + \frac{2C}{1-\lambda} D \right). \quad (14c)$$

where we used the definition of ψ and $\tilde{\psi}$ in (14a); Theorem 3.3 in (14b); Corollary E.1 in (14c). □

H Proof of Theorem 4.2

Throughout the proof, we will use $\{(\hat{x}_t, \hat{u}_t)\}$ to denote the trajectory of predictive control with prediction window T (PC_T). Recall that $\{(x_t, u_t)\}$ denotes the trajectory of predictive control with prediction window k (PC_k), and $\{(x_t^*, u_t^*)\}$ denotes the offline optimal trajectory (OPT), i.e., the optimal solution of (1).

For simplicity, we will use the shorthand notations

$$\tilde{\psi}_t^P(x; F) := \tilde{\psi}_t^P(x, w_{t:t+p-1}; F) \text{ and } \psi_t^P(x, z) := \psi_t^P(x, w_{t:t+p-1}, z)$$

throughout the proof.

Proof of Theorem 4.2. Since $x_{t+1} = \tilde{\psi}_t^P(x_t; F)_{y_1}$, for all $2 \leq i \leq k$, we have

$$\tilde{\psi}_{t-i}^k(x_{t-i}; F)_{y_i} = \tilde{\psi}_{t-i+1}^{k-1}(x_{t-i+1}; F)_{y_{i-1}}. \quad (15)$$

Therefore, we obtain that for $k \leq t \leq T - k$,

$$\begin{aligned} \|x_t\| &= \|\tilde{\psi}_{t-1}^k(x_{t-1}; F)_{y_1}\| \\ &\leq \sum_{i=1}^{k-1} \|\tilde{\psi}_{t-i}^k(x_{t-i}; F)_{y_i} - \tilde{\psi}_{t-i-1}^k(x_{t-i-1}; F)_{y_{i+1}}\| + \|\tilde{\psi}_{t-k}^k(x_{t-k}; F)_{y_k}\| \end{aligned} \quad (16a)$$

$$= \sum_{i=1}^{k-1} \|\tilde{\psi}_{t-i}^k(x_{t-i}; F)_{y_i} - \tilde{\psi}_{t-i}^{k-1}(x_{t-i}; F)_{y_i}\| + \|\tilde{\psi}_{t-k}^k(x_{t-k}; F)_{y_k}\| \quad (16b)$$

$$\leq \sum_{i=1}^{k-1} 2C\lambda^{k-1-i} \left(C\lambda^{k-1} \|x_{t-i}\| + \frac{2C}{1-\lambda} D \right) + \left(C\lambda^k \|x_{t-k}\| + \frac{2C}{1-\lambda} D \right) \quad (16c)$$

$$\leq C\lambda^{k-1} \left(\lambda \|x_{t-k}\| + 2C \sum_{i=1}^{k-1} \lambda^{k-1-i} \|x_{t-i}\| \right) + \frac{2C}{1-\lambda} \left(1 + \frac{2C}{1-\lambda} \right) D,$$

where we used the triangle inequality in (16a); (15) in (16b); Lemma 4.1 and Corollary E.1 in (16c).

By a similar argument, for $1 \leq t \leq k$, we have

$$\begin{aligned} \|x_t\| &= \|\tilde{\psi}_{t-1}^k(x_{t-1}; F)_{y_1}\| \\ &\leq \sum_{i=1}^{t-1} \|\tilde{\psi}_{t-i}^k(x_{t-i}; F)_{y_i} - \tilde{\psi}_{t-i-1}^k(x_{t-i-1}; F)_{y_{i+1}}\| + \|\tilde{\psi}_0^k(x_0; F)_{y_t}\| \\ &= \sum_{i=1}^{t-1} \|\tilde{\psi}_{t-i}^k(x_{t-i}; F)_{y_i} - \tilde{\psi}_{t-i}^{k-1}(x_{t-i}; F)_{y_i}\| + \|\tilde{\psi}_0^k(x_0; F)_{y_t}\| \\ &\leq \sum_{i=1}^{t-1} 2C\lambda^{k-1-i} \left(C\lambda^{k-1} \|x_{t-i}\| + \frac{2C}{1-\lambda} D \right) + \left(C\lambda^t \|x_0\| + \frac{2C}{1-\lambda} D \right) \\ &\leq C\lambda^{k-1} \left(2C \sum_{i=1}^{t-1} \lambda^{k-1-i} \|x_{t-i}\| \right) + C\|x_0\| + \frac{2C}{1-\lambda} \left(1 + \frac{2C}{1-\lambda} \right) D, \end{aligned} \quad (17)$$

Recall that, under the assumption of (6), the sum of coefficients in (16) and (17) are upper bounded by

$$C\lambda^{k-1} \left(2C \sum_{i=1}^{t-1} \lambda^{k-1-i} \right) \leq 1 - \delta, \quad C\lambda^{k-1} \left(\lambda + 2C \sum_{i=1}^{k-1} \lambda^{k-1-i} \right) < 1 - \delta.$$

Using inequalities (16) and (17), one can show by induction that for $t \leq T - k$

$$\|x_t\| \leq \frac{C}{\delta} \cdot (1 - \delta)^{\max(0, t-k)} \|x_0\| + \frac{2C}{\delta(1-\lambda)} \left(1 + \frac{2C}{1-\lambda} \right) D. \quad (18)$$

For $t \geq T - k + 1$, by Corollary E.1, we see that

$$\begin{aligned} \|x_t\| &= \|\tilde{\psi}_{T-k}^k(x_{T-k}, w_{T-k:T-1}, 0)_{y_{t+k-T}}\| \\ &\leq C\lambda^{t+k-T} \|x_{T-k}\| + \frac{2CD}{1-\lambda} \\ &\leq \frac{C^2}{\delta} \cdot (1-\delta)^{T-2k} \lambda^{t+k-T} \|x_0\| + \left(\frac{2C^2}{\delta(1-\lambda)} \left(1 + \frac{2C}{1-\lambda}\right) + \frac{2C}{1-\lambda} \right) D. \end{aligned}$$

This finishes the proof of ISS of PC_k .

By Lemma 4.1 and (18), we also see that for $t \leq T - k$,

$$\begin{aligned} \|\tilde{\psi}_t^k(x_t; F)_{y_1} - \tilde{\psi}_t^{T-t}(x_t; F)_{y_1}\| &\leq \sum_{p=k}^{T-t} \|\tilde{\psi}_t^p(x_t; F)_{y_1} - \tilde{\psi}_t^{p+1}(x_t; F)_{y_1}\| \\ &\leq \sum_{p=k}^{\infty} 2C\lambda^{p-1} \left(C\lambda^p \|x_t\| + \frac{2C}{1-\lambda} D \right) \\ &= \frac{2C^2}{\lambda(1-\lambda^2)} \cdot (\lambda)^{2k} \|x_t\| + \frac{4C^2}{\lambda(1-\lambda)^2} \cdot \lambda^k D \\ &= O\left(\left(D + \frac{\lambda^k(\|x_0\| + D)}{\delta} \right) \lambda^k \right). \end{aligned} \quad (19)$$

We further obtain that for $t \leq T - k$,

$$\begin{aligned} \|x_t - \hat{x}_t\| &= \|x_t - \tilde{\psi}_0^T(x_0; F)\| \\ &\leq \|x_t - \tilde{\psi}_{t-1}^{T-t+1}(x_{t-1}; F)_{y_1}\| + \sum_{i=1}^{t-1} \|\tilde{\psi}_{t-i}^{T-t+i}(x_{t-i}; F)_{y_i} - \tilde{\psi}_{t-i-1}^{T-t+i+1}(x_{t-i-1}; F)_{y_{i+1}}\| \\ &\leq \|x_t - \tilde{\psi}_{t-1}^{T-t+1}(x_{t-1}; F)_{y_1}\| + \sum_{i=1}^{t-1} C\lambda^i \|x_{t-i} - \tilde{\psi}_{t-i-1}^{T-t+i+1}(x_{t-i-1}; F)_{y_1}\| \\ &= O\left(\left(D + \frac{\lambda^k(\|x_0\| + D)}{\delta} \right) \lambda^k \right), \end{aligned} \quad (20a)$$

$$= O\left(\left(D + \frac{\lambda^k(\|x_0\| + D)}{\delta} \right) \lambda^k \right), \quad (20b)$$

where we used Theorem 3.3 and the fact that $\tilde{\psi}_{t-i-1}^{T-t+i+1}(x_{t-i-1})_{y_{i+1}}$ can be written as

$$\tilde{\psi}_{t-i-1}^{T-t+i+1}(x_{t-i-1}; F)_{y_{i+1}} = \tilde{\psi}_{t-i}^{T-t+i} \left(\tilde{\psi}_{t-i-1}^{T-t+i+1}(x_{t-i-1}; F)_{y_i}; F \right)_{y_i}$$

in (20a); we used we used (19),

$$\|x_{t-i} - \tilde{\psi}_{t-i-1}^{T-t+i+1}(x_{t-i-1}; F)_{y_1}\| = \|\tilde{\psi}_{t-i-1}^k(x_{t-i-1}; F)_{y_1} - \tilde{\psi}_{t-i-1}^{T-t+i+1}(x_{t-i-1}; F)_{y_1}\|$$

and

$$1 + \sum_{i=1}^{t-1} C\lambda^i \leq 1 + \frac{C}{1-\lambda} = O(1)$$

in (20b).

By Corollary E.1, we see that

$$\|x_T^* - \hat{x}_T\| \leq 2C\lambda^T \|x_0\| + \frac{4CD}{1-\lambda}.$$

It follows that, by Theorem 3.3, the following holds for all $t \leq T - k$:

$$\|x_t^* - \hat{x}_t\| = \|\psi_0^T(x_0, x_T^*) - \psi_0^T(x_0, \hat{x}_T)\| \leq C\lambda^k \left(2C\lambda^T \|x_0\| + \frac{4CD}{1-\lambda} \right).$$

Combining this inequality with (20) gives

$$\|x_t - x_t^*\| = O\left(\left(D + \frac{\lambda^k(\|x_0\| + D)}{\delta} \right) \lambda^k \right), \quad \forall t \leq T - k. \quad (21)$$

Since

$$(u_t - u_t^*) = B_t^\dagger((x_{t+1} - x_{t+1}^*) - A_t(x_t - x_t^*)),$$

we have

$$\|u_t - u_t^*\| \leq b'(\|x_{t+1} - x_{t+1}^*\| + a\|x_t - x_t^*\|).$$

Therefore, under Corollary F.1, for any $\eta > 0$,

$$\begin{aligned} & \iota_t^1(x_t, x_{t+1}) - (1 + \eta)\iota_t^1(x_t^*, x_{t+1}^*) \\ &= (f_{t+1}(x_{t+1}) - (1 + \eta)f_{t+1}(x_{t+1}^*)) + (c_{t+1}(u_t) - (1 + \eta)c_{t+1}(u_t^*)) \\ &\leq \frac{1}{2}\left(1 + \frac{1}{\eta}\right)\left(\ell_f\|x_{t+1} - x_{t+1}^*\|^2 + \ell_c\|u_t - u_t^*\|^2\right) \\ &\leq \frac{1}{2}\left(1 + \frac{1}{\eta}\right)\left(\ell_f + 2(b')^2\ell_c\right)\|x_{t+1} - x_{t+1}^*\|^2 + \frac{1}{2}\left(1 + \frac{1}{\eta}\right)2a^2(b')^2\ell_c\|x_t - x_t^*\|^2 \\ &\leq \left(1 + \frac{1}{\eta}\right) \cdot \frac{L_4}{2}\left(\|x_t - x_t^*\|^2 + \|x_{t+1} - x_{t+1}^*\|^2\right), \end{aligned} \tag{22}$$

where

$$L_4 := \ell_f + 2(b')^2\ell_c + 2a^2(b')^2\ell_c.$$

Then, for any $\eta > 0$, we obtain the following inequality:

$$\begin{aligned} & \text{cost}(PC_k) - (1 + \eta)\text{cost}(OPT) \\ &= \left(\sum_{t=0}^{T-k-1} \iota_t^1(x_t, x_{t+1}) + \iota_{T-k}^k(x_{T-k}, x_T)\right) - (1 + \eta)\left(\sum_{t=0}^{T-k-1} \iota_t^1(x_t^*, x_{t+1}^*) + \iota_{T-k}^k(x_{T-k}^*, x_T^*)\right) \\ &= \sum_{t=0}^{T-k-1} \left(\iota_t^1(x_t, x_{t+1}) - (1 + \eta)\iota_t^1(x_t^*, x_{t+1}^*)\right) + \left(\iota_{T-k}^k(x_{T-k}, x_T) - (1 + \eta)\iota_{T-k}^k(x_{T-k}^*, x_T^*)\right) \\ &\leq \sum_{t=0}^{T-k-1} \left(\iota_t^1(x_t, x_{t+1}) - (1 + \eta)\iota_t^1(x_t^*, x_{t+1}^*)\right) + \left(\iota_{T-k}^k(x_{T-k}, x_T) - (1 + \eta)\iota_{T-k}^k(x_{T-k}^*, x_T^*)\right) \end{aligned} \tag{23a}$$

$$\leq \left(1 + \frac{1}{\eta}\right) \cdot \frac{L_4}{2} \sum_{t=0}^{T-k-1} \left(\|x_t - x_t^*\|^2 + \|x_{t+1} - x_{t+1}^*\|^2\right) + \left(1 + \frac{1}{\eta}\right) \cdot \frac{L_0 + \ell_f}{2} \|x_{T-k} - x_{T-k}^*\|^2 \tag{23b}$$

$$\begin{aligned} &= \left(1 + \frac{1}{\eta}\right) \cdot L_4 \sum_{t=0}^{T-k-1} \|x_t - x_t^*\|^2 + \left(1 + \frac{1}{\eta}\right) \cdot \frac{L_0 + \ell_f}{2} \|x_{T-k} - x_{T-k}^*\|^2 \\ &\leq \left(1 + \frac{1}{\eta}\right) O\left(\left(D + \frac{\lambda^k(\|x_0\| + D)}{\delta}\right)^2 \lambda^{2k} T\right), \end{aligned} \tag{23c}$$

where we used the fact that the cost of $\tilde{\psi}_{T-k}^k(x_{T-k}; 0)$ is less than or equal to the cost of $\psi_{T-k}^k(x_{T-k}, x_T^*)$ in (23a); we used (22) and Corollary F.1 in (23b); and we used (20) in (23c). \square

To bound the optimal cost, we consider a suboptimal controller inspired by the decision-point transformation, where the controller forces the states $x_d, x_{2d}, \dots, x_{(v-1)d}, x_{vd+r}$ to be 0 (d is the controllability index, and $T = vd + r$). The cost of this suboptimal control is determined by the transformed transition cost $\xi_t^P(\cdot, \cdot, \cdot)$ between each pair of consecutive decision points. By strong smoothness of $\xi_t^P(\cdot, \cdot, \cdot)$ proven in Lemma 3.2, we have

$$\xi_t^P(x, \zeta, 0) \leq \frac{1}{2}L_2(p)\left(\|\zeta\|^2 + \|x\|^2\right) \leq \frac{L_0 D^2}{2} p + \frac{L_0}{2} \|x\|^2,$$

where $L_0 = \max_{d \leq p \leq 2d-1} L_2(p)$. These inequalities add up to

$$\text{cost}(OPT) \leq \xi_0^d(x_0, w_{0:d-1}, 0) + \sum_{\tau=1}^{v-2} \xi_{\tau d}^d(0, w_{\tau d:(\tau+1)d-1}, 0) + \xi_{(v-1)d}^{d+r}(0, w_{(v-1)d:T-1}, 0)$$

$$\begin{aligned} &\leq \frac{L_0 D^2}{2} T + \frac{L_0}{2} \|x_0\|^2 \\ &= O(D^2 T + \|x_0\|^2). \end{aligned}$$

Hence $\text{cost}(OPT) = O(D^2 T + \|x_0\|^2)$. Now we can take $\eta = \Theta(\lambda^k)$ in (23) to get a regret bound of

$$\text{cost}(PC_k) - \text{cost}(OPT) = O\left(\left(D + \frac{\lambda^k(\|x_0\| + D)}{\delta}\right)^2 \lambda^k T + \lambda^k \|x_0\|^2\right).$$

I Proof of Theorem 4.3

Proof of Theorem 4.3. To simplify the notation, we will omit the disturbance sequence $w_{t:t+k-1}$ in $\tilde{\psi}_t^k$ and ψ_t^k throughout the proof. At each time step t , we will use x_t/u_t to denote the state/input of $PC_{(k,h)}$ algorithm and use x_t^*/u_t^* to denote the state/input of the offline optimal. We define

$$\begin{aligned} H_t &:= f_t(x_t), M_t := c_t(u_{t-1}), \\ H_t^* &:= f_t(x_t^*), M_t^* := c_t(u_{t-1}^*). \end{aligned}$$

Let $\tilde{x}_{t+k} := \tilde{\psi}_t^k(x_t^*; F)_{y_k}$, $\bar{x}_{t+k} = \tilde{\psi}_t^k(x_t^*, 0)_{y_k}$.

If $t \leq T - k, t \equiv 0 \pmod{h}$, we have

$$\begin{aligned} &\|x_{t+h} - x_{t+h}^*\|^2 \\ &= \|\tilde{\psi}_t^k(x_t; F)_{y_h} - \psi_t^k(x_t^*, x_{t+k}^*)_{y_h}\|^2 \\ &\leq \left(\|\tilde{\psi}_t^k(x_t; F)_{y_h} - \tilde{\psi}_t^k(x_t^*; F)_{y_h}\| + \|\tilde{\psi}_t^k(x_t^*; F)_{y_h} - \psi_t^k(x_t^*, x_{t+k}^*)_{y_h}\|\right)^2 \end{aligned} \quad (24a)$$

$$\leq (1 + \epsilon) \|\tilde{\psi}_t^k(x_t; F)_{y_h} - \tilde{\psi}_t^k(x_t^*; F)_{y_h}\|^2 + \left(1 + \frac{1}{\epsilon}\right) \|\tilde{\psi}_t^k(x_t^*; F)_{y_h} - \psi_t^k(x_t^*, x_{t+k}^*)_{y_h}\|^2 \quad (24b)$$

$$\leq (1 + \epsilon) \|\tilde{\psi}_t^k(x_t; F)_{y_h} - \tilde{\psi}_t^k(x_t^*; F)_{y_h}\|^2 + \left(1 + \frac{1}{\epsilon}\right) \|\psi_t^k(x_t^*, \tilde{x}_{t+k})_{y_h} - \psi_t^k(x_t^*, x_{t+k}^*)_{y_h}\|^2 \quad (24c)$$

$$\leq (1 + \epsilon) C^2 \lambda^{2h} \|x_t - x_t^*\|^2 + C^2 \lambda^{2(k-1-h)} \cdot \left(1 + \frac{1}{\epsilon}\right) \|x_{t+k}^* - \tilde{x}_{t+k}\|^2 \quad (24d)$$

$$\leq \frac{1}{1 + \epsilon} \|x_t - x_t^*\|^2 + C^2 \lambda^{2(k-1-h)} \cdot \left(1 + \frac{1}{\epsilon}\right) \|x_{t+k}^* - \tilde{x}_{t+k}\|^2, \quad (24e)$$

where we use the triangle inequality in (24a); the AM-GM inequality in (24b); the definition of \tilde{x}_{t+k} in (24c); Theorem 3.3 in (24d); the assumption of Theorem 4.3 on h in (24e).

Since the objective function plus the indicator of the feasible set is m_f -strongly convex in variables $x_{t+1:t+k}$, we see that

$$\|x_{t+k}^* - \bar{x}_{t+k}\|^2 \leq \frac{2}{m_f} \sum_{\tau=1}^k (H_{t+\tau}^* + M_{t+\tau}^*). \quad (25)$$

Since f_{t+k} is m_f -strongly convex, we also see that $\|x_{t+k}^*\|^2 \leq \frac{2}{m_f} H_{t+k}^*$. Recall that the terminal cost $F(x_{t+k}) = \alpha(\|x_{t+k}\|)$, where α is a class K function. By the definition of $\tilde{\psi}$, we see that $\|\tilde{x}_{t+k}\| \leq \|x_{t+k}^*\|$. Therefore, we obtain that

$$\|x_{t+k}^* - \tilde{x}_{t+k}\|^2 \leq 2\|\tilde{x}_{t+k}\|^2 + 2\|x_{t+k}^*\|^2 \quad (26a)$$

$$\leq 2\|\tilde{x}_{t+k}\|^2 + 2\|x_{t+k}^*\|^2 \quad (26b)$$

$$\leq 4\|\tilde{x}_{t+k} - x_{t+k}^*\|^2 + 6\|x_{t+k}^*\|^2 \quad (26c)$$

$$\leq \frac{8}{m_f} \sum_{\tau=1}^k (H_{t+\tau}^* + M_{t+\tau}^*) + \frac{12}{m_f} H_{t+k}^* \quad (26d)$$

$$\leq \frac{20}{m_f} \sum_{\tau=1}^k (H_{t+\tau}^* + M_{t+\tau}^*),$$

where we used Cauchy-Schwarz inequality in (26a) and (26c); we used $\|\tilde{x}_{t+k}\| \leq \|\bar{x}_{t+k}\|$ in (26b); we used (25) in (26d).

Suppose $T = n_0 \cdot h + m_0$, where $n_0 \in \mathbb{Z}_+$ and $k - h + 1 \leq m_0 \leq k$. Summing up the previous inequality for $t = 0, h, 2h, \dots, (n_0 - 1)h$, we obtain that

$$\begin{aligned} \sum_{i=1}^{n_0} \|x_{ih} - x_{ih}^*\|^2 &\leq C^2 \lambda^{2(k-1-h)} \cdot \frac{(1+\epsilon)^2}{\epsilon^2} \cdot \sum_{i=1}^{n_0-1} \|x_{ih+k}^* - \tilde{x}_{ih+k}\|^2 \\ &\leq C^2 \lambda^{2(k-1-h)} \cdot \frac{(1+\epsilon)^2}{\epsilon^2} \cdot \frac{20}{m_f} \cdot \text{cost}(OPT), \end{aligned} \quad (27)$$

where we used (26) in the last inequality.

Therefore, we obtain that for all $\eta > 0$,

$$\begin{aligned} &\text{cost}(PC_{(k,h)}) - (1+\eta) \text{cost}(OPT) \\ &= \left(\sum_{i=0}^{n_0-1} \ell_{ih}^h(x_{ih}, x_{(i+1)h}) + \ell_{n_0h}^{m_0}(x_{n_0h}, x_T) \right) - (1+\eta) \left(\sum_{i=0}^{n_0-1} \ell_{ih}^h(x_{ih}^*, x_{(i+1)h}^*) + \ell_{n_0h}^{m_0}(x_{n_0h}^*, x_T^*) \right) \\ &= \sum_{i=0}^{n_0-1} \left(\ell_{ih}^h(x_{ih}, x_{(i+1)h}) - (1+\eta) \ell_{ih}^h(x_{ih}^*, x_{(i+1)h}^*) \right) + \left(\ell_{n_0h}^{m_0}(x_{n_0h}, x_T) - (1+\eta) \ell_{n_0h}^{m_0}(x_{n_0h}^*, x_T^*) \right) \\ &\leq \sum_{i=0}^{n_0-1} \left(\ell_{ih}^h(x_{ih}, x_{(i+1)h}) - (1+\eta) \ell_{ih}^h(x_{ih}^*, x_{(i+1)h}^*) \right) + \left(\ell_{n_0h}^{m_0}(x_{n_0h}, x_T) - (1+\eta) \ell_{n_0h}^{m_0}(x_{n_0h}^*, x_T^*) \right) \end{aligned} \quad (28a)$$

$$\leq \left(1 + \frac{1}{\eta}\right) \cdot \frac{L + \ell_f}{2} \sum_{i=0}^{n_0-1} \left(\|x_{ih} - x_{ih}^*\|^2 + \|x_{(i+1)h} - x_{(i+1)h}^*\|^2 \right) + \left(1 + \frac{1}{\eta}\right) \cdot \frac{L + \ell_f}{2} \|x_{n_0h} - x_{n_0h}^*\|^2 \quad (28b)$$

$$\begin{aligned} &= \left(1 + \frac{1}{\eta}\right) \cdot (L + \ell_f) \sum_{i=1}^{n_0} \|x_{ih} - x_{ih}^*\|^2 \\ &\leq \left(1 + \frac{1}{\eta}\right) \cdot (L + \ell_f) \cdot C^2 \lambda^{2(k-1-h)} \cdot \frac{20(1+\epsilon)^2}{m\epsilon^2} \cdot \text{cost}(OPT), \end{aligned} \quad (28c)$$

where we use the fact that the PC algorithm (with replan window h) plans optimally after time step nh in (28a); we use Corollary F.1 in (28b); we use (27) in (28c).

By setting $\eta \sim \epsilon^{-1} \left(\frac{L + \ell_f}{m} \right)^{\frac{1}{2}} \cdot C \lambda^{k-1-h}$, we see that the competitive ratio of the $PC_{(k,h)}$ algorithm (with replan window h) is in the order of

$$1 + O\left(\epsilon^{-1} \left(\frac{L + \ell_f}{m} \right)^{\frac{1}{2}} \cdot C \lambda^{k-1-h} \right).$$

□