Stability and Regret bounds on Distributed Truncated Predictive Control for Networked Dynamical Systems

Eric Xu*† Guannan Qu*‡

October 11, 2023

Abstract

This work is primarily concerned about the distributed control of networked linear time-invariant (LTI) systems. In particular, we propose a truncated predictive control algorithm based on κ-hop neighbourhoods of the agents of the networked system. We establish stability and regret bounds for the proposed algorithm, which shows that the regret decays exponentially when the temporal prediction horizon k and the spatial radius κ increases.

1 Introduction

The control of networked systems has retained great popularity for the past few decades because of its myriad of applications [1–4]. A critical challenge in the control of networked is the distributed nature of the decision making problem: the system is partitioned into nodes where each node may only have access to local state information. Another difficulty is the time-varying and online nature of many decision-making problems. Thus, research on efficiently controlling networked systems in a distributed and online manner has become especially prevalent. One potential solution to both problems that has had great traction over the years is distributed Predictive Control (PC) [5–7]. However, these methods mainly show asymptotic guarantees in the sense of stability and robustness whereas performance guarantees, such as the regret and optimality of the algorithm, are lacking. Recently, progress towards showing such performance guarantees has been made in the distributed control and PC literatures in parallel.

Take the distributed control literature as an example: much progress has been made in the synthesis of distributed controllers that are optimal relative to their distributed structure [8–11]. One such promising distributed control method for networked systems is the utilization of the *spatial decay* property in the optimal centralized controller, i.e., the control gain between the control action of agent i and the state of agent j decays based on the (graph) distance between the two agents. The exact decay rate varies depending on the problem setting [12–16]. Particularly of interest is the decay rate for the problem described in [15]: they show that the truncation of the centralized solution to a finite-horizon linear quadratic cost (LQC) problem to a *κ-hop distributed controller*

---

*Eric Xu and Guannan Qu contributed equally to this work.

†Carnegie Mellon University; e-mail: `ex@andrew.cmu.edu`

‡Carnegie Mellon University; e-mail: `gqu@andrew.cmu.edu`

(i.e. a controller whose gains vanish between nodes that are greater than a distance κ away) has near-optimal performance relative to the centralized solution.

On a different note, the PC literature is mature in its stability and robustness guarantees [17], [18] and has recently developed online performance guarantees in both LTI and LTV (Linear Time-Varying) settings [19–21]. An interesting property established in [21] is *temporal decay* between the predictive states generated by PC controllers that have different initial conditions; they further use this property to determine ISS (Input-to-State Stability) and regret properties of PC in the LTV setting. Bolstered by the success of PC and spatial decay-based distributed control, a natural question to ask is: how can we combine the two to attain efficient distributed and online control for networked systems? Directly combining the works in [15] and [21] is nontrivial because [21]'s proof of ISS and regret relies on its centralized setting: every node observes the global state and based on that, can make fully-accurate predictions about the whole system. In the distributed setting however, each node would only have access to its own state information and state information of nodes nearby it, and thus, the node can no longer make fully-accurate predictions about the whole system. In summary, the problem is: *how can nodes make accurate predictions about the whole system if they only have access to local information? Can we design a distributed PC algorithm that adheres to this information constraint and has comparable performance guarantees to centralized PC?* We deem this problem the *localized information-constrained predictive control problem* (LICPC Problem)

Contribution: In this paper, we develop a graph-truncated PC-style algorithm in which each agent solves a PC problem with information limited to its κ-hop neighbourhood and deploys its resulting immediate control input. In particular, we show that this algorithm attains Input to State Stability (ISS) and exponentially decaying dynamic regret in the horizon variable k and the decentralization factor κ, meaning that for sufficient choices of k and κ, we are able to attain controllers in an online manner that enjoy similar stability and regret guarantees of centralized PC. The key technical contribution underlying these guarantees is the spatial decay property exhibited by the centralized PC controller. Thus, predictions that use local information lead to accurate enough predictions of the whole system for large enough κ thereby solving the LICPC problem. Utilizing these accurate enough predictions then leads us to the desired stability and regret guarantees.

Notation: We denote the cardinality of a set S by $|S|$. The set of positive integers is $\mathbb{Z}_+$, the set of reals as $\mathbb{R}$, and the set of $m \times n$ matrices as $\mathbb{R}^{m \times n}$. Furthermore, we denote by $[T] := \{0, 1, ..., T\}$. The operator norm of a matrix A is denoted as $\|A\|$ and its minimum singular value as $\sigma_{min}(A)$. The 2-norm of a vector v is $\|v\|$. We further denote $A \succ (\succeq)0$ to mean that A is a positive (semi)definite matrix and for a matrix B, the notation $A \succ B$ means that $A - B \succ 0$. The Moore-Penrose pseudo-inverse of a matrix M is denoted $M^\dagger$. We define the notation $v_{0:T}$ for v_t indexed by $t \in [T]$ as the set: $\{v_0, v_1, ..., v_T\}$. For a matrix A we denote the sub-block indexed by $i, j \in \mathcal{V}$, where $\mathcal{V}$ is an index set, as $A[i, j]$; hence, A can be written as a sequence of its sub-blocks: $A = (A[i, j])_{i,j \in \mathcal{V}}$ (similar for vectors). We use $O(\cdot)$, $o(\cdot)$, and $\Theta(\cdot)$ as big-O, little-o, and big-theta notations respectively.

2 Preliminaries

2.1 Problem Setting

We consider a graph $\mathcal{G} := \{\mathcal{V}, \mathcal{E}\}$, where $|\mathcal{V}| = N$ is the set of nodes/agents (we will use the terms interchangeably) and $\mathcal{E}$ is the set of edges between them. We denote $\mathcal{N}_{\mathcal{G}}[i] := \{i\} \cup \{j \in \mathcal{V} : (i,j) \in \mathcal{E}\}$ as the neighbourhood of node i and $d_{\mathcal{G}}(i,j)$ to be the shortest path distance between nodes i and j based on the graph $\mathcal{G}$. We also denote the κ-hop neighbourhood of a node i as $\mathcal{N}_{\mathcal{G}}^{\kappa}[i] := \{j \in \mathcal{V} : d_{\mathcal{G}}(i,j) \leq \kappa\}$ and its boundary as $\partial \mathcal{N}_{\mathcal{G}}^{\kappa}[i] := \{j \in \mathcal{V} : d_{\mathcal{G}}(i,j) = \kappa\}$.

Given the graph $\mathcal{G}$ we can write the LTI dynamics with disturbance w of a node i as

$$x_{t+1}[i] = \sum_{j \in \mathcal{N}_{\mathcal{G}}[i]} A[i,j]x_t[j] + B[i,j]u_t[j] + w_t[i]. \tag{1}$$

Here, we take $x_t[i] \in \mathbb{R}^{n_{x_i}}$, $u_t[i] \in \mathbb{R}^{n_{u_i}}$, and $w_t[i] \in \mathbb{R}^{n_{x_i}}$ to be the state, control action, and disturbance at node i and time $t \in \mathbb{Z}_+$ respectively; furthermore, $A[i,j] \in \mathbb{R}^{n_{x_i} \times n_{x_j}}$ and $B[i,j] \in \mathbb{R}^{n_{x_i} \times n_{u_j}}$. The total number of states and control inputs are then $n_x := \sum_{i \in \mathcal{V}} n_{x_i}$ and $n_u := \sum_{i \in \mathcal{V}} n_{u_i}$. Naturally, we can rewrite (1) as a centralized system:

$$x_{t+1} = Ax_t + Bu_t + w_t, \tag{2}$$

where $x_t = (x_t[i])_{i \in \mathcal{V}}$, $u_t = (u_t[i])_{i \in \mathcal{V}}$, and $w_t = (w_t[i])_{i \in \mathcal{V}}$ are the centralized state, control input, and disturbance and similarly, $A = (A[i,j])_{i,j \in \mathcal{V}}$ and $B = (B[i,j])_{i,j \in \mathcal{V}}$ where these matrices are networked in the sense that for any $(i,j) \notin \mathcal{N}_{\mathcal{G}}[i]$, the submatrices $A[i,j]$ and $B[i,j]$ are 0.

Thus, we define the networked control problem:

$$\min_{x_{0:T}, u_{0:T-1}} \sum_{i \in \mathcal{V}} \left(\sum_{t=0}^{T} f_t[i](x_t[i]) + \sum_{t=0}^{T-1} c_{t+1}[i](u_t[i]) \right),$$
$$\text{s.t. } x_{t+1} = Ax_t + Bu_t + w_t, \quad \forall t \in [T-1], \tag{3}$$
$$x_0 = \bar{x},$$

for which the time varying costs $f_t[i]$ and $c_t[i]$ are decentralized such that they only depend on the state and control of node i at time t. We can thus aggregate the costs into centralized versions of themselves via summation, i.e. $f_t(x_t) := \sum_{i \in \mathcal{V}} f_t[i](x_t[i])$ and similarly for $c_t(u_{t-1})$. Given the dynamical system from before, we introduce two standard concepts:

Definition 1. *The system (A, B) is (L, γ)-stabilizable if there exists a controller K such that for $L > 1$, $\gamma \in (0,1)$, and $\|K\| \leq L$, we have that $\|(A - BK)^t\| \leq L\gamma^t$ for all $t \in \mathbb{Z}_+$.*

We also require the stronger concept of the controllability index of a system (A, B):

Definition 2. *There exists a positive integer $d \leq n$ such that the reduced controllability matrix $\mathcal{C}^d := \begin{bmatrix} B & AB & \cdots & A^{d-1}B \end{bmatrix}$ is full row rank.*

2.2 Predictive Control

Predictive Control (PC) is an algorithm that can be used to solve the finite horizon problem in (3) in an online fashion. In particular, we have that at time t, the controller observes k (the prediction horizon) information tuples $I_{t:t+k-1}$ where each information tuple is defined as $I_t := (A, B, w_t, f_t, c_{t+1})$ and solves the following finite-horizon problem for each $t < T - k$:

$$\tilde{\Psi}_t^k(x, \zeta; F) := \operatorname*{arg\,min}_{y_{0:k}, v_{0:k-1}} \sum_{\tau=0}^{k-1} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}(v_{\tau-1}) + F(y_k), \qquad (4)$$

$$\text{s.t. } y_{\tau+1} = Ay_\tau + Bv_\tau + \zeta_\tau, \quad \forall \tau \in [k-1],$$

$$y_0 = x,$$

where $\zeta \in (\mathbb{R}^n)^k$ is a sequence of k disturbances indexed from 0 to $k - 1$ and $F : \mathbb{R}^n \mapsto \mathbb{R}$ is a terminal cost regularizing the final predictive state. At each predictive time step τ, the predictive state and control are $y_\tau \in \mathbb{R}^n$ and $v_\tau \in \mathbb{R}^m$ respectively. Abusing notation, we will often write $\tilde{\Psi}_t^p(x, \zeta)$ for when the terminal cost can be either $F(\cdot)$ or f_T and $\tilde{\Psi}_t^p(x; \cdot) = \tilde{\Psi}_t^p(x, \zeta; \cdot)$ when the disturbances are unambiguous. Thus, the overall algorithm can be written as:

Algorithm 1 Centralized PC (PC_k)

---
    for $t = 0, 1, ..., T - k - 1$ **do**
        Observe x_t and information tuples $I_{t:t+k-1}$.
        Solve $\tilde{\Psi}_t^k(x_t, w_{t:t+k-1}; F)$ and apply $u_t = v_0$.
    At $t = T - k$, observe x_t and information tuple $I_{t:T-1}$.
    Solve $\tilde{\Psi}_t^k(x_t, w_{t:T-1}; f_T)$ and apply $u_{t:T-1} = v_{0:k-1}$.

In [21], they establish ISS and regret bounds for PC_k. In particular, these bounds depend on the *temporal decay constant* $\delta_{\mathcal{T}}$ (established in their lipschitz result for the optimization problem $\tilde{\Psi}_t^k$ (Theorem 3.3 in [21])): with large enough prediction horizon, k, the dynamic regret will decay exponentially as $O(\delta_{\mathcal{T}}^k)$.

Despite such nice properties, it is quite cumbersome to perform PC_k in the networked setting because every node would require access to the global state and information tuples at every time step t. Hence, we would ideally like to design a similar predictive controller for each node that only has access to limited information; this naturally leads to the aforementioned LICPC problem which we address in the folllowing section.

3 Main Results

The main idea behind the algorithm we will design is that at each time step, every node solves a "κ-hop" version of (4) such that they only have access to state information and future information tuples of other nodes that are within a κ-hop distance of themselves.

3.1 Algorithm

Before explicitly describing the algorithm, we require a way to formalize the idea of some node i having access to only local information. This idea is encapsulated in the following definition:

Definition 3. *For a matrix $M = (M[j,k])_{j,k\in\mathcal{V}}$ we can define its (i,κ)-truncation $M^{(i,\kappa)}$ such that*

$$M^{(i,\kappa)}[j,k] = \begin{cases} M[j,k], & j \in \mathcal{N}_{\mathcal{G}}^{\kappa}[i], \\ 0, & \text{else,} \end{cases} \tag{5}$$

Note that this definition can easily be extended to vectors (i.e. set the entries belonging to j that are outside of $\mathcal{N}_{\mathcal{G}}^{\kappa}[i]$ equal to 0). Thus, we can define the distributed counterpart of (4):

$$\tilde{\psi}_t^k(x, \zeta, \mathcal{N}_{\mathcal{G}}^{\kappa}[i]; F) := \underset{y_{0:k}, v_{0:k-1}}{\arg\min} \sum_{t=0}^{k-1} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}(v_{\tau-1}) + F(y_k), \tag{6}$$

$$\text{s.t. } y_{\tau+1} = A^{(i,\kappa)}y_\tau + B^{(i,\kappa)}v_\tau + \zeta_\tau^{(i,\kappa)}, \quad \tau \in [k-1],$$

$$y_0 = x^{(i,\kappa)}.$$

In (6), while we wrote $\tilde{\psi}_t^k$ as a function of x, ζ, it only uses local information in the sense that only the (i,κ)-truncated version of A, B, ζ, and x are needed for solving (6). Further, for the costs $f_{t+\tau}(y_\tau) = \sum_j f_{t+\tau}[j](y_\tau[j])$, only the costs in the κ-hop neighbourhood $\mathcal{N}_{\mathcal{G}}^{\kappa}[i]$ are needed since $y_\tau[j] = 0$ for j outside of it. Similarly, $c_{t+\tau}(y_\tau) = \sum_j c_{t+\tau}[j](v_{\tau-1}[j])$ only requires costs within the κ-hop neighbourhood $\mathcal{N}_{\mathcal{G}}^{\kappa+1}[i]$ since $B^{(i,\kappa)}[j,k]$ is nonzero for $j \in \mathcal{N}_{\mathcal{G}}^{\kappa}[i]$ and $k \in \mathcal{N}_{\mathcal{G}}[j]$. In other words, to solve (6), agent i only needs to observe $x^{(i,\kappa)}$ and have access to the localized info tuples: $I_{t:t+k-1}^{\kappa}$ where each localized info tuple is $I_t^{\kappa} := (A^{(i,\kappa)}, B^{(i,\kappa)}, w_t^{(i,\kappa)}, f_t^{(i,\kappa)}, c_{t+1}^{(i,\kappa+1)})$ where $f_t^{(i,\kappa)}[j] = f_t[j]$ for $j \in \mathcal{N}_{\mathcal{G}}^{\kappa}[i]$ and 0 for j outside of $\mathcal{N}_{\mathcal{G}}^{\kappa}[i]$ and $c_{t+1}^{(i,\kappa+1)}$ is defined similarly. The full algorithm can be described as in Algorithm 2:

Algorithm 2 Distributed-Truncated PC ($DTPC_k$)
---
for $t = 0, 1, ..., T-1$ **do**
    for $i = 1, ..., N$ **do**
        Agent i observes $x_t^{(i,\kappa)}$ and info tuples $I_{t:t+k-1}^{\kappa}$
        if $t < T-k$ **then**
            Solve $\tilde{\psi}_t^k(x_t, w_{t:t+k-1}, \mathcal{N}_{\mathcal{G}}^{\kappa}[i]; F)$
        else
            Solve $\tilde{\psi}_t^{T-t}(x_t, w_{t:t+k-1}, \mathcal{N}_{\mathcal{G}}^{\kappa}[i]; f_T)$
        Collect $v_0[i]$ from the solution and set $u_t[i] = v_0[i]$
    Apply u_t to the system.

$DTPC_k$ answers the first question in the LICPC problem by representing the information constraint via (i,κ)-truncated versions of the information tuples and state; each agent solves the optimization problem in (6) in which they only have access to (i,κ)-truncated versions of A, B,

ζ, and x. $DTPC_k$ also solves the second question asked by the LICPC problem since the error incurred from truncating the information at each node is negligible for large enough κ due to a spatial decay property (cf. Theorem 2). This leads to ISS and regret guarantees, which we present now.

3.2   Stability and Regret Guarantees

Before presenting our main result, we state our assumptions.

Assumption 1. *The system matrices in (2) satisfy the following:*

1. *$\|A\| \leq L$, $\|B\| \leq L$, and $\|B^\dagger\| \leq L$.*

2. *The reduced controllability matrix has minimum singular value $\sigma_{min}(\mathcal{C}^d) \geq \sigma$ for $\sigma > 0$.*

3. *There exists $\kappa_0 < \text{diam}(\mathcal{G})$ such that for all $i \in \mathcal{V}$ and $\kappa \geq \kappa_0$, the system $(A^{(i,\kappa)}, B^{(i,\kappa)})$ is (L, γ)-stabilizable.*

Where $\text{diam}(\mathcal{G})$ is the diameter of the graph $\mathcal{G}$. Note that the second assumption guarantees existence of an (L, γ)-stable K in the LTI setting. The final assumption essentially says that if we isolate large enough κ-hop neighbourhoods of nodes from a distributed system, then we expect that these isolated κ-hop subsystems to be stabilizable if the original distributed system was stabilizable.

We also make the following assumption on the costs:

Assumption 2. *The costs are well-conditioned such that:*

1. *$f_t(\cdot)$ and $c_t(\cdot)$ are μ-strongly convex, L-smooth, and twice continuously differentiable for all t.*

2. *$F(\cdot)$ is a μ-strongly convex and L-smooth K-function (i.e., $F(x) = \beta(\|x\|)$ where $\beta(\cdot)$ is strictly increasing and $\beta(0) = 0$) and twice continuously differentiable*

3. *$f_t(\cdot)$ and $c_t(\cdot)$ are non-negative and $f_t(0) = c_t(0) = 0$ for all $t = 1, ..., T$. The same goes for $f_0(\cdot)$.*

Finally, we make one last assumption on the τ-hop neighbourhoods of the graph

Assumption 3. *There exists a subexponential function $p(\cdot)$ such that for some distance d, we have that*

$$|j \in \mathcal{V} : d_{\mathcal{G}}(i, j) = d| \leq p(d). \tag{7}$$

Our main contribution is the following ISS and regret bound that mirrors those found in centralized PC. The result is stated in the following theorem:

Theorem 1. *Let the disturbance w_t be uniformly bounded such that $\max_{t \in [T-k]} \sum_{\tau=0}^{k-1} \|w_{t+\tau}\| \leq D_k$, and let $C := \max(\Omega, \Gamma)$ and $\delta := \max(\delta_{\mathcal{S}}, \delta_{\mathcal{T}})$ where Ω and $\delta_{\mathcal{T}}$ are as in Lemma 3 and Γ and $\delta_{\mathcal{S}}$ are as in Theorem 2. Under Assumptions 1, 2, and 3 with constants $\xi = 1 - \sqrt{\delta} > 0$ and L from the first two Assumptions, we take k and κ such that*

$$\kappa \geq \max\left(\kappa_0, \frac{\log \frac{(1-\sqrt{\delta})(1-\delta)}{2C^2 LN}}{\log \delta}\right), \quad k \geq \frac{2\log \frac{\delta^{5/2}(1-\delta)}{4C^3}}{\log \delta},$$

which gives the following ISS bound for the system:

$$\|x_t\| \leq \begin{cases} \frac{C}{\xi}(1-\xi)^{\max(0,t-k)}\|x_0\| + \frac{W}{\xi}, & t \leq T - k, \\ \frac{C^2}{\xi^2}(1-\xi)^{T-2k}\delta^{t+k-T}\|x_0\| + \frac{2CW}{\xi^2}, & \text{else,} \end{cases} \tag{8}$$

where $W := \frac{10C^3 LN}{(1-\delta)^2}D_k$. Further, we have the regret bound:

$$\text{cost}(DTPC_k) - \text{cost}(OPT) =$$
$$O\left(\left(\left(D_k + \frac{\|x_0\| + D_k}{\xi^2}\right)^2 \delta^\kappa + \left(D_k + \frac{\delta^k(\|x_0\| + D_k)}{\xi}\right)^2 \delta^k\right)T + \eta\|x_0\|^2\right) \tag{9}$$

where $\eta = \Theta(\max(\delta^k, \delta^\kappa))$.

Theorem 1 says that the regret is exponentially decaying in κ and k, meaning that if we pick κ and k equal to $\Theta(\log T)$, we can attain $o(1)$ dynamic regret. This result parallels that of [21] with differences. One difference is that in contrast to a single decay constant in [21], our bound involves two decay constants $\delta_{\mathcal{S}}$ and $\delta_{\mathcal{T}}$, which correspond to the spatial and temporal decay constants respectively. The spatial decay constant $\delta_{\mathcal{S}}$ depends on the graph and networked structure of the dynamics in (1). The temporal decay constant $\delta_{\mathcal{T}}$ is attributed to lipschitz property of (4) and its costs' conditioning. Aside from differences in the constants, the ISS result attains an extra factor of ξ in the last k time steps due to having to explicitly solve (6) at every time step. Furthermore, the regret bound incurs extra error (the κ term) from the truncation as expected.

4 Proof

We begin the proof of Theorem 1 with a roadmap,

Step 1: Prove exponential decay between the solution to the centralized problem in (4) versus the solution to (6). The main idea is that we can express the KKT conditions of both (4) and (6) as the matrix equation $H(z)z = b$ for which we can directly apply the exponential decay of the inverse of H from [14] and [15] to show exponential decay.

Step 2: Show the ISS bound on $DTPC_k$. Here, we will utilize the exponential decay from Step 1 to establish a bound on the difference between the next state produced by $DTPC_k$ and the state that would have been produced by PC_k. Then, since PC_k is already ISS, we expect $DTPC_k$ also to be ISS.

Step 3: We finally prove the regret bound of $DTPC_k$. In order to prove the regret, we will use assumption 2 and relate the difference between the trajectory from $DTPC_k$ and the offline optimal trajectory $\tilde{\Psi}_0^T(x_0, w_{0:T-1}; f_T)$ via the ISS bound.

We note that due to space limit, the proofs for some auxiliary results are omitted, and its explicit proof can be found in the Appendix of [22].

4.1 Proof of Step 1

The main goal of Step 1 is to prove the following result:

Theorem 2. *Let q^c be the solution vector containing the primal and dual variables for the centralized problem in (4) and let q^d be that of the distributed problem in (6) for some $i \in \mathcal{V}$. Let the prediction horizon be $\ell \leq k$ and the truncation factor be $\kappa \geq \kappa_0$ as in Assumption 1. Under the Assumptions 1, 2, and 3, we have the following decay result:*

$$\|q^c[i] - q^d[i]\| \leq \Gamma(\|x\| + D_\ell)\delta_{\mathcal{S}}^\kappa \tag{10}$$

where the closed brackets $[\cdot]$ will henceforth denote the spatial indexing of some vector or matrix by the network graph $\mathcal{G}$ underlying the dynamics, i.e. $q^c[i]$ means all the entries in q^c corresponding to node i (including all time steps); $D_\ell := \max_{t \in [T-1]} \sum_{\tau=0}^{\ell-1} \|w_{t+\tau}\|$. The constants Γ and $\delta_{\mathcal{S}}$ are defined as

$$\delta_{\mathcal{S}} := \frac{\rho+1}{2}, \quad \Gamma := \frac{4\alpha^2 \delta_{\mathcal{S}} L}{(1-\delta_{\mathcal{S}})^2} \left(\sup_{d \in \mathbb{Z}_+} (\rho/\delta_{\mathcal{S}})^d p(d) \right)^2 p(1),$$

where the function $p(\cdot)$ is from Assumption 3 and constant L is from Assumptions 1 and 2. The constants ρ and α are as in Theorem 2.

Before proving the result above, we first revisit the KKT conditions of the optimization problems in (4) and (6) and then we introduce two key auxiliary results to proving Theorem 2.

Let $t \in [T]$ be arbitrary and $\ell \leq k$ be the horizon variable. Then we define the total cost as

$$\hat{f}(z) := \sum_{\tau=0}^{\ell-1} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{\ell} c_{t+\tau}(v_{\tau-1}) + g(y_\ell), \tag{11}$$

where $g(\cdot)$ is either $F(\cdot)$ or $f_T(\cdot)$ and $z := (y_0.v_0, ..., y_{\ell-1}, v_{\ell-1}, y_\ell)$ is the trajectory.

Let J be the constraint jacobian

$$J(A, B) := \begin{bmatrix} I & & & \\ -A & -B & I & \\ & & \ddots & \\ & & & -A & -B & I \end{bmatrix}, \tag{12}$$

which has $\ell + 1$ rows. Let, λ be the dual variables. The KKT conditions of (6) are then

$$\begin{bmatrix} \nabla\hat{f}(z^d) + (J^d)^\top \lambda^d \\ J^d z^d \end{bmatrix} = \begin{bmatrix} 0 \\ \begin{bmatrix} x^{(i,\kappa)} \\ \zeta^{(i,\kappa)} \end{bmatrix} \end{bmatrix}, \tag{13}$$

where $J^d := J(A^{(i,\kappa)}, B^{(i,\kappa)})$ and (z^d, λ^d) is primal-dual solution of $\tilde{\psi}_t^\ell$. Further, we have the following lemma:

Lemma 1 (Lemma 1 in [23])**.** *For μ-strongly convex, L-smooth, and twice continuously-differentiable $\hat{f}$, For each z and z' there exists symmetric $G(z, z')$ such that $\mu I \preceq G(z) \preceq LI$ and $\nabla\hat{f}(z) - \nabla\hat{f}(z') = G(z, z')(z - z')$.*

Using Lemma 1, we can write (13) as $\tilde{H}^d q^d = b^d$ where

$$\tilde{H}^d := \begin{bmatrix} G(z^d, 0) & (J^d)^\top \\ J^d & \end{bmatrix}, \quad b^d := \begin{bmatrix} 0 \\ \begin{bmatrix} x^{(i,\kappa)} \\ \zeta^{(i,\kappa)} \end{bmatrix} \end{bmatrix}, \tag{14}$$

and q^d is short hand notation for (z^d, λ^d). Furthermore, when taking the difference of the KKT conditions for equations (4) and (6), we get,

$$\begin{bmatrix} \nabla \hat{f}(z^c) - \nabla \hat{f}(z^d) + (J^c)^\top \lambda^c - (J^d)^\top \lambda^d \\ J^c z^c - J^d z^d \end{bmatrix} = \begin{bmatrix} 0 \\ \begin{bmatrix} x^\perp \\ \zeta^\perp \end{bmatrix} \end{bmatrix},$$

where the notation z^c and $J^c := J(A, B)$ denote the solution and constraint jacobian of $\tilde{\Psi}_t^\ell$ respectively and the vectors on the RHS are $x^\perp := x - x^{(i,\kappa)}$ and $\zeta^\perp := \zeta - \zeta^{(i,\kappa)}$. Applying Lemma 1 to $\nabla \hat{f}(z^c) - \nabla \hat{f}(z^d)$, we get

$$\underbrace{\begin{bmatrix} G(z^c, z^d) & (J^c)^\top \\ J^c & \end{bmatrix}}_{:=H^c} q^c - \underbrace{\begin{bmatrix} G(z^c, z^d) & (J^d)^\top \\ J^d & \end{bmatrix}}_{:=H^d} q^d = \underbrace{\begin{bmatrix} 0 \\ \begin{bmatrix} x^\perp \\ \zeta^\perp \end{bmatrix} \end{bmatrix}}_{:=b^\perp}, \tag{15}$$

where q^c is short hand for $(z^c, \lambda^c))$.

The purpose of rewriting the KKT conditions using the H matrices defined above is because we have the following result from [14] and [15] on such H matrices:

Lemma 2 (Theorem 3.6 in [14] and Theorems A.3 and A.4 in [15]). *Consider the following matrix H*

$$H = \begin{bmatrix} G & J(A, B)^\top \\ J(A, B) & 0 \end{bmatrix},$$

such that G is a block diagonal matrix with singular values $\mu I \preceq G \preceq LI$, and the system (A, B) is (L, γ)-stabilizable. Further, let $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ be any graph whose set of nodes partitions H such that for any $i, j \in \mathcal{V}$, if $d_\mathcal{G}(i, j) > 1$, then $H[i, j] = 0$. Then, the KKT matrix H has singular values bounded as $\mu_H \leq \sigma(H) \leq L_H$ and its inverse H^{-1} is spatially exponentially decaying with respect to the graph $\mathcal{G}$

$$\|H^{-1}[i, j]\| \leq \alpha \rho^{d_\mathcal{G}(i,j)}, \tag{16}$$

where the constants are defined as:

$$\mu_J := \frac{(1-\gamma)^2}{L^2(1-L^2)}, \quad L_H := 2L + 1$$

$$\mu_H := \left(\frac{1}{\mu} + \left(1 + \frac{2L_H}{\mu} + \frac{L_H^2}{\mu^2} \right) \frac{L_H}{\mu_J} \right)^{-1},$$

$$\rho := \left(\frac{L_H^2 - \mu_H^2}{L_H^2 + \mu_H^2} \right)^{\frac{1}{2}}, \quad \alpha := \frac{L_H}{\mu_H^2 \rho}.$$

With the results above, we are ready to prove Theorem 2.

Proof of Theorem 2. By (15), we have,

$$b^\perp = H^c q^c - H^d q^d = H^c(q^c - q^d) - (H^d - H^c)q^d$$
$$\implies q^c - q^d = (H^c)^{-1}\left(b^\perp - H^\perp q^d\right)$$

where $H^\perp = H^c - H^d$. Then we obtain the upper bound:

$$\|q^c[i] - q^d[i]\| = \left\|\sum_{j\in\mathcal{V}}(H^c)^{-1}[i,j](b^\perp[j] - \sum_{k\in\mathcal{V}}H^\perp[j,k]q^d[k])\right\|$$

$$\leq \alpha \sum_{j\in\mathcal{V}\backslash\mathcal{N}_\mathcal{G}^\kappa[i]}\rho^{d_\mathcal{G}(i,j)}\left(\|b^\perp\| + \sum_{k\in\partial\mathcal{N}_\mathcal{G}^\kappa[i]\cap\mathcal{N}_\mathcal{G}[j]}\sum_{m\in\mathcal{N}_\mathcal{G}^\kappa[i]}\|H^\perp[j,k]\|\|(\tilde{H}^d)^{-1}[k,m]\|\|b^d[m]\|\right)$$

$$\leq \alpha \sum_{j\in\mathcal{V}\backslash\mathcal{N}_\mathcal{G}^\kappa[i]}\rho^{d_\mathcal{G}(i,j)}\left(\|b^\perp\| + 2\alpha L\sum_{k\in\partial\mathcal{N}_\mathcal{G}^\kappa[i]\cap\mathcal{N}_\mathcal{G}[j]}\sum_{m\in\mathcal{N}_\mathcal{G}^\kappa[i]}\rho^{d_\mathcal{G}(k,m)}\|b^d\|\right)$$

$$\leq \alpha \sum_{d=\kappa+1}^{\infty}(\rho/\delta_\mathcal{S})^d p(d)\delta_\mathcal{S}^d\left(\|b^\perp\| + 2\alpha Lp(1)\sum_{s=0}^{2\kappa}(\rho/\delta_\mathcal{S})^s p(s)\delta_\mathcal{S}^s\|b^d\|\right)$$

$$\leq \frac{2\alpha^2\delta_\mathcal{S}Lp(1)}{(1-\delta_\mathcal{S})^2}\left(\sup_{d\in\mathbb{Z}^+}\left(\frac{\rho}{\delta_\mathcal{S}}\right)^d p(d)\right)^2 (\|b^\perp\| + \|b^d\|)\delta_\mathcal{S}^\kappa$$

In the first inequality, we have used (i) Lemma 2 applied to H^c; (ii) for $j \in \mathcal{N}_\mathcal{G}^\kappa[i]$, the j'th entry/row of $b^\perp$ and $H^\perp$ are zero (hence the outer sum is over $j \in \mathcal{V}\setminus\mathcal{N}_\mathcal{G}^\kappa[i]$); (iii) $H^\perp[j,k]$ is nonzero only when $k \in \mathcal{N}_\mathcal{G}[j]$ since $H^\perp$ consists of networked matrices, and $q^d[k]$ is only nonzero in $\mathcal{N}_\mathcal{G}^\kappa[i]$ (hence the inner sum is over $k \in \partial\mathcal{N}_\mathcal{G}^\kappa[i] \cap \mathcal{N}_\mathcal{G}[j]$); (iv) $\tilde{H}^d q^d = b^d$. The second inequality uses Lemma 2 applied to $\tilde{H}^d$ and the fact that $\|H^\perp[j,k]\| \leq 2L$. The third inequality follows from Assumption 3. Finally, the inequality in (10) follows from $\|b\| \leq \|x\| + \sum_{\tau=0}^{\ell-1}\|\zeta_\tau\|$ for either $b^\perp$ or b^d. $\qquad\square$

4.2 Proof of Step 2 (ISS)

To show ISS, we first show a recursive upper bound on the states generated by $DTPC_k$, denoted as $x_{0:T}$.

Theorem 3. *Let $C := \max(\Gamma, \Omega)$ and we denote $\delta := \max(\delta_\mathcal{S}, \delta_\mathcal{T})$ where Γ and $\delta_\mathcal{S}$ are as in Theorem 2 and Ω and $\delta_\mathcal{T}$ are from Lemma 3. Under the Assumptions of Theorem 1, we have the following upper bounds on the states generated by $DTPC_k$. For $1 \leq t+1 \leq k$:*

$$\|x_{t+1}\| \leq C\sum_{m=0}^{t}(L_N\delta^{\kappa+m} + 2C^2\delta^{2k-m-3})\|x_{t-m}\| + C\|x_0\| + W, \tag{17}$$

for $k \leq t+1 \leq T-k$:

$$\|x_{t+1}\| \leq C\sum_{m=0}^{k-1}(L_N\delta^{\kappa+m} + 2C^2\delta^{2k-m-3})\|x_{t-m}\| + W, \tag{18}$$

and finally for $t + 1 > T - k$:

$$\|x_{t+1}\| \leq \sum_{m=0}^{t+k-T} 2CL_N \|x_{t-m}\| \delta^{\kappa+m} + C\delta^{t+k-T+1} \|x_{T-k}\| + W. \tag{19}$$

where $L_N := CLN$ and W is as in Theorem 1.

Theorem 3 directly leads to the ISS in (8) by induction.

Proof of (8). First, it is easy to verify that the ISS bound holds for $t = 0$. For the induction step, we show the case $2k \leq t \leq T - k - 1$, and the other cases are similar. In other words, we assume the ISS bound holds for all $t_0 \leq t$ for some $t \in [2k, T - k - 1]$, and now show the ISS bound also holds for $t + 1$. Since $k < t + 1 \leq T - k$, we have from Theorem 3 that

$$\|x_{t+1}\| \leq C \sum_{m=1}^{k} (L_N \delta^{\kappa+m-1} + 2C^2 \delta^{2k-m-2}) \|x_{t-m+1}\| + W. \tag{20}$$

By the induction hypothesis, the upper bound in (20) becomes,

$$C \sum_{m=1}^{k} \overbrace{(L_N \delta^{\kappa+m-1} + 2C^2 \delta^{2k-m-2})}^{:=\Lambda_m} \left(\frac{C}{\xi}(1-\xi)^{t-m-k+1} \|x_0\| + \frac{W}{\xi} \right) + W, \tag{21}$$

where the max disappears since $t \geq 2k - 1$. Then, note the choice of k and κ in Theorem 1 guarantees that $C \sum_m \Lambda_m \leq 1 - \xi$, which allows us to upper bound (21) as

$$\left(\frac{C}{\xi}(1-\xi)^{t-k+1} \|x_0\| \right) C \sum_{m=1}^{k} \Lambda_m (1-\xi)^{-m} + \frac{W}{\xi}. \tag{22}$$

Finally, the selection of $1 - \xi = \sqrt{\delta}$, k, and κ in Theorem 1 further guarantees that $C \sum_m \Lambda_m (1 - \xi)^{-m} \leq 1$, which obtains the desired bound in (8) and concludes the induction. $\square$

The remainder of this subsection will be devoted to proving Theorem 3. We first require a couple definitions and supplementary results. First, we define the following terminal state predictive control problem similar to (4):

$$\Psi_t^k(x, \zeta, \bar{x}) := \underset{y_{0:k}, v_{0:k-1}}{\arg\min} \sum_{\tau=0}^{k} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}(v_{\tau-1}), \tag{23}$$

$$\text{s.t. } y_{\tau+1} = Ay_\tau + Bv_\tau + \zeta_\tau, \quad \forall \tau \in [k-1],$$

$$y_0 = x, \quad y_k = \bar{x},$$

By [21], the following lipschitz property of Ψ_t^k holds.

Lemma 3. *(Theorem 3.3 in [21]) Under assumptions 1 and 2 and for horizon length $k \geq d$, the controllability index, given any $(x, \zeta, \bar{x})$ and $(x', \zeta', \bar{x}')$,*

$$\|\Psi_t^k(x, \zeta, \bar{x})_{y_m} - \Psi_t^k(x', \zeta', \bar{x}')_{y_m}\| \leq \Omega \left(\delta_{\mathcal{T}}^m \|x - x'\| + \delta_{\mathcal{T}}^{k-m} \|\bar{x} - \bar{x}'\| + \sum_{l=0}^{k-1} \delta_{\mathcal{T}}^{|m-l|} \|\zeta_l - \zeta_l'\| \right) \quad (24)$$

where $D := \sup_{t \in [T-k]} \|w_t\|$ and the constants Ω and $\delta_{\mathcal{T}}$ are defined in Theorem 3.3 in [21]

Furthermore, as a Corollary of Lemma 2, we have the following lipschitz result on $\tilde{\Psi}_t^\ell$

Corollary 1. *Let $\ell \in [T]$ and q denote the optimal vector of primal and dual variables to $\tilde{\Psi}_t^\ell(x, \zeta; g(\cdot))$ and q' be that of $\tilde{\Psi}_t^\ell(x', \zeta'; g(\cdot))$. We partition q and q' by the the temporal graph $\mathcal{G}_\ell := \{\{[\ell]\}, \{(0, 1), ..., (\ell-1, \ell)\}\}$ such that we can denote q_m and q_m' indexed by $m \in [\ell]$. The lipschitz property then follows from Lemma 2:*

$$\|q_m - q_m'\| \leq \alpha \left(\rho^m \|x - x'\| + \sum_{l=0}^{\ell-1} \rho^{|m-l|} \|\zeta_l - \zeta_l'\| \right) \quad (25)$$

$$\leq \alpha \left(\rho^m \|x - x'\| + \frac{2D_k}{1-\rho} \right), \quad (26)$$

Utilizing the above two results, we attain the following:

Lemma 4. *Let $y_{m+1} := \tilde{\Psi}_{t-m}^k(x_{t-m})_{y_{m+1}}$ and $y_{m+1}' := \tilde{\Psi}_{t-m-1}^k(x_{t-m-1})_{y_{m+2}}$, then we have the following upper bound on the norm of their difference:*

$$\|y_{m+1} - y_{m+1}'\| \leq CL_N(\|x_{t-m-1}\| + D_k)\delta^{\kappa+m+1} + C^2\delta^{k-m-2}\left(\delta^{k-1}(\|x_{t-m}\| + C\delta\|x_{t-m-1}\|) + \frac{6C}{1-\delta}D_k \right), \quad (27)$$

The proof is in Appendix-B of [22], from which we acquire a few auxiliary results. First, by the principle of optimality, we have

$$\tilde{\Psi}_{t-m-1}^k(x_{t-m-1})_{y_{m+2}} = \tilde{\Psi}_{t-m}^{k-1}(x_{t-m}^c)_{y_{m+1}}, \quad (28)$$

where $x_{t-m}^c := \tilde{\Psi}_{t-m-1}^k(x_{t-m-1})_{y_1}$. Next, we also have

$$\|x_{t-m} - x_{t-m}^c\| \leq L_N(\|x_{t-m-1}\| + D_k)\delta^\kappa, \quad (29)$$

which follows directly from rewriting the norm as

$$\|x_{t-m} - x_{t-m}^c\| = \|B(u_{t-m-1} - \tilde{\Psi}_{t-m-1}^k(x_{t-m-1})_{v_0})\|, \quad (30)$$

where $u_{t-m-1} := (\tilde{\psi}_{t-m-1}^k(x_{t-m-1}, \mathcal{N}_{\mathcal{G}}^\kappa[i])_{v_0[i]})_{i \in \mathcal{V}}$, and then applying Theorem 2 to get the upper bound in (29). Now we may prove Theorem 3.

Proof of Theorem 3. We only prove the case for $k \leq t \leq T - k - 1$, and the proofs for $t \leq k - 1$ and $t \geq T + k$ are similar. We start by comparing the norm of $\|x_{t+1}\|$ as

$$\|x_{t+1}\| = \left\| \tilde{\Psi}_t^k(x_t)_{y_1} - (x_{t+1}^c - x_{t+1}) \right\|$$

$$\leq \sum_{m=0}^{k-2} \|\tilde{\Psi}_{t-m}^k(x_{t-m})_{y_{m+1}} - \tilde{\Psi}_{t-m-1}^k(x_{t-m-1})_{y_{m+2}}\| + \|\tilde{\Psi}_{t-k+1}^k(x_{t-k+1})_{y_k}\| + L_N(\|x_t\| + D_k)\delta^\kappa$$

$$\leq \sum_{m=0}^{k-2} \left(CL_N(\|x_{t-m-1}\| + D_k)\delta^{\kappa+m+1} + C^2\delta^{k-m-2}\left(\delta^{k-1}\left(\|x_{t-m}\| + C\delta\|x_{t-m-1}\| \right) + \frac{6C}{1-\delta}D_k \right) \right)$$

$$+ C\delta^k\|x_{t-k+1}\| + \frac{2C}{1-\delta}D_k + L_N(\|x_t\| + D_k)\delta^\kappa$$

$$\leq C\sum_{m=0}^{k-1} \left(L_N\delta^{\kappa+m} + 2C^2\delta^{2k-m-3} \right)\|x_{t-m}\| + W.$$

The first equality uses the definition $x_{t+1}^c = \tilde{\Psi}_t^k(x_t)_{y_1}$. In the first inequality, we telescope $\tilde{\Psi}_t^k(x_t)_{y_1}$ and apply (29). In the second inequality, we apply Lemma 4 and Corollary 1 with $x' = 0$ and $\zeta' = 0$. In the final inequality, we combine terms and upper bound by the geometric sum to attain the desired bound. □

4.3 Proof of Step 3 (Regret)

To prove the regret result in Theorem 1 we first prove the following critical result which bounds the difference between $DTPC_k$'s trajectory and the optimal offline trajectory:

Theorem 4. *Let $x_{0:T}$ be the trajectory generated by $DTPC_k$ and $x_{0:T}^*$ is the optimal offline trajectory. For $t+1 \leq T-k$, we have that*

$$\|x_{t+1} - x_{t+1}^*\| \leq C\delta^k\left(2C\delta^T\|x_0\| + \frac{4C}{1-\delta}D_k \right) + \sum_{m=0}^t \frac{4C^2L_N}{\delta(1-\delta)^2}\delta^m\left((\delta^\kappa + \delta^{2k})\|x_{t-m}\| + (\delta^\kappa + \delta^k)D_k \right),$$

(31)

and for $t+1 \geq T-k+1$, we have that

$$\|x_{t+1} - x_{t+1}^*\| \leq \sum_{m=0}^t CL_N(\|x_{t-m}\| + D_k)\delta^{\kappa+m} + \sum_{m=1}^{T-k} \frac{4C^3}{\delta(1-\delta)^2}\delta^{t+k-T+m}\left(\delta^{2k}\|x_{T-k-m}\| + \delta^k D_k \right), \quad (32)$$

Proof. For $t+1 \leq T-k$: Let $\hat{x}_{0:T}$ denote the state trajectory of $\tilde{\Psi}_0^T(x_0; F)$. Then we can upper bound the norm of the difference as

$$\|x_{t+1} - x_{t+1}^*\| \leq \|x_{t+1} - \hat{x}_{t+1}\| + \|\hat{x}_{t+1} - x_{t+1}^*\|.$$

By Lemma 3 and Corollary 1 and for $t+1 \leq T-k$, we have the upper bound from equation (21) in [21]

$$\|\hat{x}_{t+1} - x_{t+1}^*\| \leq C\delta^k\left(2C\delta^T\|x_0\| + \frac{4C}{1-\delta}D_k \right). \quad (33)$$

Then we make the observation that

$$\tilde{\Psi}_{t-m-1}^{T-t+m+1}(x_{t-m-1})_{y_{m+2}} = \tilde{\Psi}_{t-m}^{T-t+m}(x_{t-m}^{c'})_{y_{m+1}}, \quad (34)$$

where $x^{c'}_{t-m} := \tilde{\Psi}^{T-t+m+1}_{t-m-1}(x_{t-m-1})_{y_1}$. Thus we have,

$$\|x_{t+1} - \hat{x}_{t+1}\| = \|x_{t+1} - \tilde{\Psi}^T_0(x_0; F)_{y_{t+1}}\|$$

$$\leq \|x_{t+1} - \tilde{\Psi}^{T-t}_t(x_t)_{y_1}\| + \sum_{m=0}^{t-1} \|\tilde{\Psi}^{T-t+m}_{t-m}(x_{t-m})_{y_{m+1}} - \tilde{\Psi}^{T-t+m+1}_{t-m-1}(x_{t-m-1})_{y_{m+2}}\|$$

$$\leq \|x_{t+1} - x^{c'}_{t+1}\| + \sum_{m=0}^{t-1} C\delta^{m+1}\|x_{t-m} - x^{c'}_{t-m}\|$$

$$\leq \|x_{t+1} - x^c_{t+1}\| + \|x^c_{t+1} - x^{c'}_{t+1}\| + \sum_{m=0}^{t-1} C\delta^{m+1}\left(\|x_{t-m} - x^c_{t-m}\| + \|x^c_{t-m} - x^{c'}_{t-m}\|\right),$$

where we've applied the observation from above (34) and Corollary 1 in the second inequality. We then apply the bound in (29) and the bound

$$\|\tilde{\Psi}^k_t(x_t)_{y_1} - \tilde{\Psi}^{T-t}_t(x_t)_{y_1}\| \leq \frac{4C^2}{\delta(1-\delta)^2}(\delta^{2k}\|x_t\| + \delta^k D_k), \tag{35}$$

(Equation 19 in [21]) which holds for $t+1 \leq T-k$. This completes the proof for $t+1 \leq T-k$. The case for $t+1 \geq T-k+1$ is similar so its proof is omitted here. □

Before finishing the proof of Theorem 1, we require the following two Lemmas. First, we have from [21]:

Lemma 5 (Lemma F.2 in [21]). *Suppose $h : \mathbb{R}^n \mapsto \mathbb{R}_+$ is a convex and L-smooth continuously differentiable function. Then for $x, x' \in \mathbb{R}^n$ and $\eta > 0$, we have that*

$$h(x) - (1+\eta)h(x') \leq \frac{L}{2}\left(1 + \frac{1}{\eta}\right)\|x - x'\|^2.$$

The above Lemma 5 will be applied to both f_t and c_t.

In addition, we require the following bound on the one step terminal state problem ($k = 1$) in (23), the proof is in Appendix-C of [22]

Lemma 6. *Let v and v' be the two "one-step" terminal state problems: $\Psi^1_t(x_t, x_{t+1})_{v_0}$ and $\Psi^1_t(x'_t, x'_{t+1})_{v_0}$. Under Assumptions 1 and 2, we have that*

$$\|v - v'\|^2 \leq C^2\left(\|x_t - x'_t\|^2 + \|x_{t+1} - x'_{t+1}\|^2\right), \tag{36}$$

where C is as in Theorem 3.

We are now ready to prove the regret in Theorem 1.

Proof of Theorem 1. First, denote $\bar{u}_t := \Psi^1_t(x_t, x_{t+1})_{v_0}$. Then consider the difference of the costs c_{t+1} which can be written as:

$$c_{t+1}(u_t) - (1+\eta)c_{t+1}(\bar{u}_t) + (1+\eta)\left(c_{t+1}(\bar{u}_t) - (1+\eta)c_{t+1}(u^*_t)\right)$$

$$\leq \frac{L}{2}\left(1+\frac{1}{\eta}\right)\left(\|u_t - \bar{u}_t\|^2 + (1+\eta)\|\bar{u}_t - u_t^*)\|^2\right)$$

$$\leq C^3\left(1+\frac{1}{\eta}\right)\left(L_N^2(\|x_t\| + D_k)^2\delta^{2\kappa} + \|x_{t+1} - x_{t+1}^c\|^2 + (1+\eta)\left(\|x_t - x_t^*\|^2 + \|x_{t+1} - x_{t+1}^*\|^2\right)\right)$$

$$\leq 2C^3L_N^2\left(1+\frac{1}{\eta}\right)\left((\|x_t\| + D_k)^2\delta^{2\kappa} + (1+\eta)(\|x_t - x_t^*\|^2 + \|x_{t+1} - x_{t+1}^*\|^2)\right). \tag{37}$$

Where in the first inequality, we have applied Lemma 5. To obtain the second inequality, by the parallelogram identity, we have the upper bound

$$\|u_t - \bar{u}_t\|^2 \leq 2\left(\|u_t - \tilde{\Psi}_t^{k'}(x_t)_{v_0}\|^2 + \|\tilde{\Psi}_t^{k'}(x_t)_{v_0} - \bar{u}_t\|^2\right), \tag{38}$$

where $k' := \min(k, T-t)$. Applying Theorem 2 to $\|u_t - \tilde{\Psi}_t^{k'}(x_t)_{v_0}\|$ and Lemma 6 to $\|\tilde{\Psi}_t^{k'}(x_t)_{v_0} - \bar{u}_t\|$ gives us the following upper bound of (38):

$$2\left(L_N^2(\|x_t\| + D_k)^2\delta^{2\kappa} + C^2(\|x_{t+1} - x_{t+1}^c\|^2)\right). \tag{39}$$

Observe that

$$\|\bar{u}_t - u_t^*\| = \|\Psi_t^1(x_t, x_{t+1})_{v_0} - \Psi_t^1(x_t^*, x_{t+1}^*)_{v_0}\|,$$

for which we can apply Lemma 6. This, together with (39) gives us the second inequality in (37). Finally, we attain the last inequality by applying (29) again. We take $\eta = \Theta(\max(\delta^k, \delta^\kappa))$. Denoting $1 + \eta' = (1+\eta)^2$, we have

$$\text{cost}(DTPC_k) - (1+\eta')\text{cost}(OPT) = \sum_{t=1}^{T}(f_t(x_t) - (1+\eta')f_t(x_t^*)) + (c_t(u_{t-1}) - (1+\eta')c_t(u_{t-1}^*))$$

$$\leq 2C^3L_N^2\left(1+\frac{1}{\eta}\right)\sum_{t=0}^{T-1}\left(\|x_{t+1} - x_{t+1}^*\|^2 + (\|x_t\| + D_k)^2\delta^{2\kappa} + (1+\eta)\left(\|x_t - x_t^*\|^2 + \|x_{t+1} - x_{t+1}^*\|^2\right)\right)$$

$$\leq 6C^3L_N^2\left(1+\frac{1}{\eta}\right)(1+\eta)\sum_{t=0}^{T}((\|x_t\| + D_k)^2\delta^{2\kappa} + \|x_t - x_t^*\|^2)$$

$$= \frac{1}{\eta}O\left(\left[\left(D_k + \frac{\|x_0\| + D_k}{\xi^2}\right)^2\delta^{2\kappa} + \left(D_k + \frac{\delta^k(\|x_0\| + D_k)}{\xi}\right)^2\delta^{2k}\right]T\right). \tag{40}$$

Where in the first inequality we apply (37), Lemma 5, and that $\eta \leq \eta'$. In the second inequality, we merge all of the norms of differences. In the final line, we apply the ISS bound and Theorem 4 to attain the final expression in (40).

Since $\eta' = 2\eta + \eta^2$, and hence, $\eta' = \Theta\left(\max(\delta^k, \delta^\kappa)\right)$, we simply add $\eta'\text{cost}(OPT)$ to the RHS of (40) to achieve the bound in (9). This finishes the proof of Theorem 1. □

5 Simulations

We consider the temperature control of a building with HVAC network graph that is a 5×5 mesh grid with diameter 8 and each of its 25 nodes corresponding to a different zone in the building. The states are the temperatures of each zone and their respective temperature integrators and the
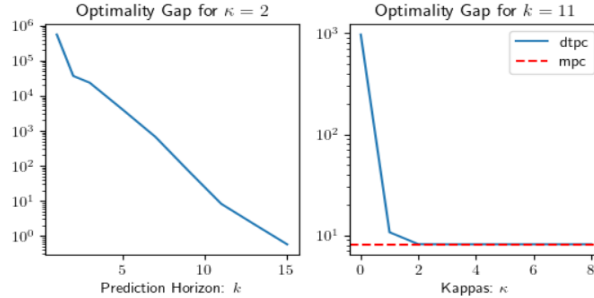
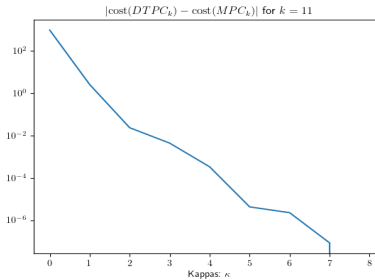Figure 1: $DTPC_k$ versus the OPT for fixed κ (left) and fixed k (right).

Figure 2: $DTPC_k$ versus PC_k for $k = 11$ and varying κ.

control variable at zone i is its manipulated heat generation/absorption. Further details about the system and setup are in Appendix-D in [22]

The results are shown in Figure 1 where we chose a time horizon of $T = 30$. Observe that $DTPC_k$ exhibits the decaying regret behavior as in (9) as κ and k increase. Note that in Figure 1 (right), the regret stops decreasing because the prediction horizon k becomes the bottleneck after κ reaches 2. Figure 2 demonstrates our result in Theorem 2: when κ increases under fixed k, $DTPC_k$ becomes exponentially close to PC_k.

6 Conclusion

In this work we have shown that our algorithm $DTPC_k$ produces trajectories that are similar to those produced by centralized predictive control PC_k. Furthermore, we have shown stability and regret bounds for $DTPC_k$ which guarantee near-optimal performance. As for future work, we would like to extend this analysis to the LTV dynamics case.

References

[1] C. Zhao, U. Topcu, N. Li, and S. Low, "Design and stability of load-side primary frequency control in power systems," *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1177–1189, 2014.

[2] B. Sinopoli, C. Sharp, L. Schenato, S. Schaffert, and S. S. Sastry, "Distributed control applications within sensor networks," *Proceedings of the IEEE*, vol. 91, no. 8, pp. 1235–1246, 2003.

[3] P. Varaiya, "Max pressure control of a network of signalized intersections," *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 177–195, 2013.

[4] S. E. Li, Y. Zheng, K. Li, Y. Wu, J. K. Hedrick, F. Gao, and H. Zhang, "Dynamical modeling and distributed control of connected and automated vehicles: Challenges and opportunities," *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 3, pp. 46–58, 2017.

[5] B. T. Stewart, A. N. Venkat, J. B. Rawlings, S. J. Wright, and G. Pannocchia, "Cooperative distributed model predictive control," *Systems & Control Letters*, vol. 59, pp. 460–469, Aug. 2010.

[6] D. Jia and B. Krogh, "Min-max feedback model predictive control for distributed control with communication," in *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*, vol. 6, pp. 4507–4512 vol.6, 2002.

[7] C. A. Alonso, J. S. Li, J. Anderson, and N. Matni, "Distributed and localized model-predictive control–part i: Synthesis and implementation," *IEEE Transactions on Control of Network Systems*, vol. 10, no. 2, pp. 1058–1068, 2023.

[8] J. Anderson, J. C. Doyle, S. H. Low, and N. Matni, "System level synthesis," *Annual Reviews in Control*, vol. 47, pp. 364–393, 2019.

[9] S. Fattahi, N. Matni, and S. Sojoudi, "Efficient learning of distributed linear-quadratic control policies," *SIAM Journal on Control and Optimization*, vol. 58, no. 5, pp. 2927–2951, 2020.

[10] U. Ghai, U. Madhushani, N. Leonard, and E. Hazan, "A regret minimization approach to multi-agent control," in *Proceedings of the 39th International Conference on Machine Learning* (K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, eds.), vol. 162 of *Proceedings of Machine Learning Research*, pp. 7422–7434, PMLR, 17–23 Jul 2022.

[11] G. Fazelnia, R. Madani, A. Kalbat, and J. Lavaei, "Convex relaxation for optimal distributed control problems," *IEEE Transactions on Automatic Control*, vol. 62, no. 1, pp. 206–221, 2017.

[12] B. Bamieh, F. Paganini, and M. A. Dahleh, "Distributed control of spatially invariant systems," *IEEE Transactions on automatic control*, vol. 47, no. 7, pp. 1091–1107, 2002.

[13] N. Motee and A. Jadbabaie, "Optimal control of spatially distributed systems," *IEEE Transactions on Automatic Control*, vol. 53, no. 7, pp. 1616–1629, 2008.

[14] S. Shin, M. Anitescu, and V. M. Zavala, "Exponential decay of sensitivity in graph-structured nonlinear programs," *SIAM Journal on Optimization*, vol. 32, no. 2, pp. 1156–1183, 2022.

[15] S. Shin, Y. Lin, G. Qu, A. Wierman, and M. Anitescu, "Near-optimal distributed linear-quadratic regulator for networked systems," *SIAM Journal on Control and Optimization*, vol. 61, no. 3, pp. 1113–1135, 2023.

[16] Y. Zhang, G. Qu, P. Xu, Y. Lin, Z. Chen, and A. Wierman, "Global convergence of localized policy iteration in networked multi-agent reinforcement learning," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 7, mar 2023.

[17] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[18] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control* (A. Garulli and A. Tesi, eds.), (London), pp. 207–226, Springer London, 1999.

[19] C. Yu, G. Shi, S.-J. Chung, Y. Yue, and A. Wierman, "The power of predictions in online control," in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20, (Red Hook, NY, USA), Curran Associates Inc., 2020.

[20] R. Zhang, Y. Li, and N. Li, "On the regret analysis of online lqr control with predictions," in *2021 American Control Conference (ACC)*, pp. 697–703, 2021.

[21] Y. Lin, Y. Hu, G. Shi, H. Sun, G. Qu, and A. Wierman, "Perturbation-based regret analysis of predictive control in linear time varying systems," in *Advances in Neural Information Processing Systems* (M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, eds.), vol. 34, pp. 5174–5185, Curran Associates, Inc., 2021.

[22] E. Xu and G. Qu, "Distributed truncated predictive control full paper." https://ex19087.github.io/Distributed_Truncated_Predictive_Control/ACC_Truncated_Predictive_Control_full.pdf, 2023.

[23] G. Qu and N. Li, "On the exponential stability of primal-dual gradient dynamics," *IEEE Control Systems Letters*, vol. 3, no. 1, pp. 43–48, 2019.

7 Appendix

7.1 Proofs of intermediary results in Step 1

7.1.1 Proof of Lemma 2

We note that the proof is a combination of Theorem 3.6 in [14] and Theorems A.3 and A.4 in [15]. For completeness, we provide a proof below with a focus on the differences from the proofs of those Theorems; the steps that are identical are omitted.

The proof for the exponential decay result in equation (16) is available in [14] (Theorem 3.6) which requires upper and lower bounds on the singular values of the H matrix: i.e., $\mu_H \leq \sigma(H) \leq L_H$. So we will prove the bounds on the singular values similar to Theorems A.3 and A.4 in [15].

Let Z be an orthonormal null-space basis matrix for J, i.e. $\text{Col}(Z) = \text{Null}(J)$, then, we shall prove the following uniform regularity conditions:

$$H \preceq L_H I, \quad Z^\top G Z \succeq \mu I, \quad J J^\top \succeq \mu_J I, \tag{41}$$

Proof of $H \preceq L_H I$: Let $H_{\tau,\tau'}$ denote the submatrix of H such that the rows correspond to $(y_\tau, v_\tau, \lambda_\tau)$ and the columns correspond to $(y_{\tau'}, v_{\tau'}, \lambda_{\tau'})$, then we have that

$$H_{\tau,\tau'} = \begin{cases} \begin{bmatrix} G_{y_\tau} & & I \\ & G_{v_\tau} & \\ I & & \end{bmatrix}, & \tau = \tau' \neq \ell, \\[2em] \begin{bmatrix} & & 0 \\ & & 0 \\ -A & -B & \end{bmatrix}, & \tau = \tau' + 1 \neq \ell \\[2em] \begin{bmatrix} G_{y_\ell} & I \\ I & \end{bmatrix}, & \tau = \tau' = \ell \\[1.5em] \begin{bmatrix} & & 0 \\ -A & -B & \end{bmatrix}, & \tau = \tau' + 1 = \ell \\[1.5em] 0, & \tau - \tau' > 1 \end{cases} \tag{42}$$

where each $G_{(\cdot)}$ comes from defining the block diagonal matrices of G such that

$$G := \begin{bmatrix} G_{y_0} & & & & \\ & G_{v_0} & & & \\ & & \ddots & & \\ & & & G_{v_{\ell-1}} & \\ & & & & G_{y_\ell} \end{bmatrix}$$

where each G_{y_τ} and G_{v_τ} corresponds to the predictive state and control y_τ and v_τ respectively. The upper bound $\|H\| \leq 2L + 1$ follows from the following result (Lemma 5.11 from [14]):

$$\|M\| \leq \left(\max_{i \in \mathcal{V}} \sum_{j \in \mathcal{V}} \|M_{ij}\| \right)^{1/2} \left(\max_{j \in \mathcal{V}} \sum_{i \in \mathcal{V}} \|M_{ij}\| \right)^{1/2}. \tag{43}$$

Proof of $Z^\top G Z \succ \mu I$: The bound is immediate from Z being full rank and $G \succ \mu I$.

Proof of $JJ^\top \succ \mu_J I$: The proof is the exact same as in Theorem A.4 of [15] so it is omitted here.

Finally, we can use the conditions in (41) to prove the lower bound of H. First we can rewrite the inverse of H by its schur complement

$$H^{-1} = \begin{bmatrix} G & J^\top \\ J & \end{bmatrix}^{-1} = \begin{bmatrix} G^{-1} - G^{-1}J^\top (G_J)^{-1}JG^{-1} & G^{-1}J^\top (G_J)^{-1} \\ (G_J)^{-1}JG^{-1} & -(G_J)^{-1} \end{bmatrix}, \tag{44}$$

where $G_J := JG^{-1}J^\top$. So we may upper bound the norm of H^{-1} as the sum of the norms of its sub-blocks as such:

$$\|H^{-1}\| \leq \|G^{-1}\| + (1 + 2\|JG^{-1}\| + \|JG^{-1}\|^2)\|G_J^{-1}\|, \tag{45}$$

where $\|G^{-1}\| \leq \frac{1}{\mu}$, $\|JG^{-1}\| \leq \frac{L_H}{\mu}$, and we can obtain the upper bound on G_J^{-1} as

$$\|(JG^{-1}J^\top)^{-1}\| = \|(JG^{-1/2}G^{-1/2}J^\top)^{-1}\| = \|(G^{-1/2}(J^\top J)^{1/2}(J^\top J)^{1/2}G^{-1/2})^{-1}\|$$
$$= \|G^{1/2}(J^\top J)^{-1}G^{1/2}\| \leq \|G\|\|(J^\top J)^{-1}\| \leq \frac{L_H}{\mu_J}. \tag{46}$$

Applying the upper bounds to (45), we obtain

$$\|H^{-1}\| \leq \frac{1}{\mu} + \left(1 + \frac{2L_H}{\mu} + \frac{L_H}{\mu^2}\right)\frac{L_H}{\mu_J},$$

as desired. This completes the proof of Lemma 2. $\qquad\square$

7.2 Proofs of intermediary results in Step 2

7.2.1 Complete proof of ISS in Theorem 1

For the case of $t \leq k - 1$, the induction step from (17) is

$$\|x_{t+1}\| \leq C \sum_{m=0}^{t} \Lambda_m \left(\frac{C}{\xi}\|x_0\| + \frac{W}{\xi}\right) + C\|x_0\| + W,$$

for which the choice of k and κ in Theorem 1 guarantees that $C \sum_{m=0}^{t-1} \Lambda_m \leq (1 - \xi)$ and so

$$\|x_{t+1}\| \leq \frac{C}{\xi}\|x_0\| + \frac{W}{\xi}, \tag{47}$$

as desired.

For the case of $k \leq t < 2k - 1$, let $(t - m - k + 1)_+$ denote $\max(0, t - m - k + 1)$, then by Theorem 3,

$$\|x_{t+1}\| \leq C \sum_{m=1}^{k}(L_N \delta^{\kappa+m-1} + 2C^2\delta^{2k-m-2})\|x_{t-m+1}\| + W.$$

And thus, by the induction hypothesis and the condition that $C \sum_m \Lambda_m \leq 1 - \xi$,

$$\|x_{t+1}\| \leq C \sum_{m=1}^{k} \Lambda_m \left(\frac{C}{\xi}(1 - \xi)^{(t-m-k+1)_+}\|x_0\|\right) + \frac{W}{\xi}, \tag{48}$$

for which we can split the sum up as

$$\left(\frac{C}{\xi}(1 - \xi)^{t-k+1}\|x_0\|\right)C\left(\sum_{m=1}^{t-k}\left(\Lambda_m(1 - \xi)^{-m}\right) + \sum_{m=t-k+1}^{k}\Lambda_m(1 - \xi)^{-(t-k+1)}\right) + \frac{W}{\xi}. \tag{49}$$

Equation (49) is clearly upperbounded by

$$\left(\frac{C}{\xi}(1 - \xi)^{t-k+1}\|x_0\|\right)C\sum_{m=1}^{k}\Lambda_m(1 - \xi)^{-m} + \frac{W}{\xi}, \tag{50}$$

which is the same as in equation (22), thus, by the same logic, the desired bound in (8) is obtained for this case.

For the case of $t > T - k$, we have from (19) that

$$\|x_{t+1}\| \leq \sum_{m=0}^{t+k-T} 2CL_N\|x_{t-m}\|\delta^{\kappa+m} + C\delta^{t+k-T+1}\|x_{T-k}\| + W$$

$$\leq \left(\frac{C}{\xi}\delta^{t+k-T+1}\|x_{T-k}\| + \frac{W}{\xi}\right)\left(\sum_{m=0}^{t+k-T} 2CL_N\delta^{\kappa+m} + \xi\right) \leq \frac{C}{\xi}\delta^{t+k-T+1}\|x_{T-k}\| + \frac{W}{\xi}$$

$$\leq \frac{C}{\xi^2}\delta^{t+k-T+1}\left(C(1-\xi)^{T-2k}\|x_0\| + W\right) + \frac{W}{\xi}$$

where the second inequality is obtained through the following fact for $t - m \geq T - k$

$$\|x_{t-m}\| \leq \frac{C}{\xi}\delta^{t+k-T+1}\|x_{T-k}\| + \frac{W}{\xi},$$

which can be easily proved by induction, similar to the previous induction argument for $t \leq k - 1$, cf. eq. (47); the third line comes from applying the condition $\sum_m 2CL_N\delta^{\kappa+m} \leq (1 - \xi)$ (guaranteed by the choice of κ in Theorem 1). The final inequality is obtained by applying the ISS bound for $t = T - k$ in (8). This completes the proof of (8) □

7.2.2 Proof of Corollary 1

As shown in Equation (42), $H_{t,t'} = 0$ for $|t - t'| > 1$ meaning that H satisfies the conditions of Lemma 2 for the graph $\mathcal{G}_\ell$. Thus, we have that

$$\|q_m - q'_m\| = \left\|\sum_{l\in[\ell]} H_{m,l}^{-1}(b_l - b'_l)\right\| \leq \alpha\left(\rho^m\|x - x'\| + \sum_{l=0}^{\ell}\rho^{|m-l|}\|\zeta_l - \zeta'_l\|\right), \tag{51}$$

where we denote

$$b_l := \begin{cases} \begin{bmatrix} 0 \\ x \end{bmatrix}, & l = 0, \\ \begin{bmatrix} 0 \\ \zeta_l \end{bmatrix}, & \text{else,} \end{cases}$$

and similarly for b'_l. The inequality is obtained through application of Lemma 2.

7.2.3 Proof of Lemma 4

First, observe that
$$y'_{m+1} = \tilde{\Psi}^k_{t-m-1}(x_{t-m-1})_{y_{m+2}} = \tilde{\Psi}^{k-1}_{t-m}(x^c_{t-m})_{y_{m+1}}, \tag{52}$$

where $x^c_{t-m} := \tilde{\Psi}^k_{t-m-1}(x_{t-m-1})_{y_1}$. Then, denote

$$\bar{y} := \tilde{\Psi}^k_{t-m}(x_{t-m})_{y_{k-1}}, \quad \bar{y}' := \tilde{\Psi}^{k-1}_{t-m}(x^c_{t-m})_{y_{k-1}},$$

so now we may proceed to prove the Lemma as:

$$\|y_{m+1} - y'_{m+1}\| = \|\Psi^{k-1}_{t-m}(x_{t-m}, \bar{y})_{y_{m+1}} - \Psi^{k-1}_{t-m}(x^c_{t-m}, \bar{y}')_{y_{m+1}}\|$$

$$\leq C\left(\delta^{m+1}\|x_{t-m} - x^c_{t-m}\| + \delta^{k-m-2}\|\bar{y} - \bar{y}'\|\right)$$

$$\leq CL_N(\|x_{t-m-1}\| + D_k)\delta^{\kappa+m+1} + C^2\delta^{k-m-2}\left(\delta^{k-1}\left(\|x_{t-m}\| + \|x^c_{t-m}\|\right) + \frac{4D_k}{1-\delta}\right)$$

$$\leq CL_N(\|x_{t-m-1}\| + D_k)\delta^{\kappa+m+1} + C^2\delta^{k-m-2}\left(\delta^{k-1}\left(\|x_{t-m}\| + C\delta\|x_{t-m-1}\|\right) + \frac{6CD_k}{1-\delta}\right) \quad (53)$$

where in the first equality, we apply the principle of optimality. In the first inequality we apply Lemma 3. In the second inequality, we apply the result in (29) along with the upper bounds

$$\|\bar{y}\| \leq C\left(\delta^{k-1}\|x_{t-m}\| + \frac{2D_k}{1-\delta}\right),$$

and

$$\|\bar{y}'\| \leq C\left(\delta^{k-1}\|x^c_{t-m}\| + \frac{2D_k}{1-\delta}\right),$$

which both come from the Lipschitz bound in Corollary 1. The last inequality then comes from applying Corollary 1 to x^c_{t-m} and combining the D_k terms. $\qquad\square$

7.2.4 Full proof of Theorem 3

For the case $t \leq k - 1$, we have that

$$\|x_{t+1}\| = \|\tilde{\Psi}^k_t(x_t)_{y_1} - (x^c_{t+1} - x_{t+1})\|$$

$$\leq \sum_{m=0}^{t-1} \|\tilde{\Psi}^k_{t-m}(x_{t-m})_{y_{m+1}} - \tilde{\Psi}^k_{t-m-1}(x_{t-m-1})_{y_{m+2}}\| + \|\tilde{\Psi}^k_0(x_0)_{y_{t+1}}\| + L_N(\|x_t\| + D_k)\delta^\kappa$$

$$\leq \sum_{m=0}^{t-1}\left(CL_N\delta^{\kappa+m+1}(\|x_{t-m-1}\| + D_k) + C^2\delta^{k-m-2}\left(\delta^{k-1}\left(\|x_{t-m}\| + C\delta\|x_{t-m-1}\|\right) + \frac{6C}{1-\delta}D_k\right)\right)$$

$$+ C\delta^{t+1}\|x_0\| + \frac{2C}{1-\delta}D_k + L_N(\|x_t\| + D_k)\delta^\kappa$$

$$\leq \sum_{m=0}^{t}(CL_N\delta^{\kappa+m} + 2C^3\delta^{2k-m-3})\|x_{t-m}\| + C\|x_0\| + W. \quad (54)$$

where we use the definition of $\tilde{\Psi}^k_t(x_t)_{y_1}$ in the first equality. We telescope $\tilde{\Psi}^k_t(x_t)_{y-1}$ and the result in (29) in the second line. In the second inequality, we apply Lemma 4 and Corollary 1 with $x' = 0$ and $\zeta' = 0$. Then in the last inequality we combine like terms.

For $t \geq T - k$: We continue similarly,

$$\|x_{t+1}\| = \|\tilde{\Psi}^{T-t}_t(x_t)_{y_1} - (x^{c'}_{t+1} - x_{t+1})\|$$

$$\leq \sum_{m=0}^{t+k-T-1} \|\tilde{\Psi}^{T-t+m}_{t-m}(x_{t-m})_{y_{m+1}} - \tilde{\Psi}^{T-t+m+1}_{t-m-1}(x_{t-m-1})_{y_{m+2}}\| + \|\tilde{\Psi}^k_{T-k}(x_{T-k})_{y_{t+k-T+1}}\| + L_N(\|x_t\| + D_k)\delta^\kappa$$

$$= \sum_{m=0}^{t+k-T-1} \|\tilde{\Psi}_{t-m}^{T-t+m}(x_{t-m})_{y_{m+1}} - \tilde{\Psi}_{t-m}^{T-t+m}(x_{t-m}^{c'})_{y_{m+1}}\| + \|\tilde{\Psi}_{T-k}^{k}(x_{T-k})_{y_{t+k-T+1}}\| + L_N(\|x_t\| + D_k)\delta^\kappa$$

$$\leq \sum_{m=0}^{t+k-T-1} C\delta^{m+1}\|x_{t-m} - x_{t-m}^{c'}\| + C\delta^{t+k-T+1}\|x_{T-k}\| + \frac{2C}{1-\delta}D_k + L_N(\|x_t\| + D_k)\delta^\kappa$$

$$\leq \sum_{m=0}^{t+k-T} 2CL_N\|x_{t-m}\|\delta^{\kappa+m} + C\delta^{t+k-T+1}\|x_{T-k}\| + \frac{4CL_N}{1-\delta}D_k, \tag{55}$$

where $x_{t-m}^{c'} := \tilde{\Psi}_{t-m-1}^{T-t+m+1}(x_{t-m-1})_{y_1}$. The first inequality is obtained via telescoping $\tilde{\Psi}_t^{T-t}(x_t)_{y_1}$ and applying the following upper bound for $t - m - 1 \geq T - k$,

$$\|x_{t-m}^{c'} - x_{t-m}\| = \left\| B\left(\tilde{\Psi}_{t-m-1}^{T-t+m+1}(x_{t-m-1})_{v_0} - u_{t-m-1} \right) \right\| \leq L_N(\|x_{t-m-1}\| + D_k)\delta^\kappa \tag{56}$$

where $u_{t-m-1} := \left(\tilde{\psi}_{t-m-1}^{T-t+m+1}(x_{t-m-1}, \mathcal{N}_{\mathcal{G}}^\kappa[i])_{v_0[i]} \right)_{i \in \mathcal{V}}$ and we've applied Theorem 2 to obtain the inequality (56) which resembles the result in (29). The second inequality is obtained through applying Corollary 1 twice; the final inequality is then obtained by applying the result in (56) and combining terms. This completes the proof of Theorem 3 □

7.3 Proofs of intermediary results in Step 3

7.3.1 Complete Proof of Theorem 4

For the case $t \geq T - k$, we have that

$$\|x_{t+1} - x_{t+1}^*\| = \|x_{t+1} - \tilde{\Psi}_0^T(x_0)_{y_{t+1}}\|$$

$$\leq \|x_{t+1} - \tilde{\Psi}_t^{T-t}(x_t)_{y_1}\| + \sum_{m=0}^{t-1} \|\tilde{\Psi}_{t-m}^{T-t+m}(x_{t-m})_{y_{m+1}} - \tilde{\Psi}_{t-m-1}^{T-t+m+1}(x_{t-m-1})_{y_{m+2}}\|$$

$$\leq L_N(\|x_t\| + D_k)\delta^\kappa + \sum_{m=0}^{t-1} C\delta^{m+1}\|x_{t-m} - x_{t-m}^{c'}\|$$

$$\leq L_N(\|x_t\| + D_k)\delta^\kappa + \sum_{m=0}^{t-1} C\delta^{m+1}\left(\|x_{t-m} - x_{t-m}^{\tilde{c}}\| + \|x_{t-m}^{\tilde{c}} - x_{t-m}^{c'}\| \right)$$

$$\leq \sum_{m=0}^{t} CL_N(\|x_{t-m}\| + D_k)\delta^{\kappa+m} + \sum_{m=t+k-T}^{t-1} \frac{4C^3}{\delta(1-\delta)^2}\delta^{m+1}\left(\delta^{2k}\|x_{t-m-1}\| + \delta^k D_k \right) \tag{57}$$

where in the first inequality we telescope $\tilde{\Psi}_0^T(x_0)_{y_{t+1}}$. In the second inequality, we apply the result in (56) and Corollary 1 on the terms in the sum, denoting $x_{t-m}^{c'} := \tilde{\Psi}_{t-m-1}^{T-t+m+1}(x_{t-m-1})_{y_1}$. In the third inequality, we apply the triangle inequality and denote $x_{t-m}^{\tilde{c}} := \tilde{\Psi}_{t-m-1}^{\min(k,T-t+m+1)}(x_{t-m-1})_{y_1}$. Finally, in the last inequality, we apply (29) and (56) to the terms $\|x_{t-m} - x_{t-m}^{\tilde{c}}\|$, and the result in (35) to the terms $\|x_{t-m}^{\tilde{c}} - x_{t-m}^{c'}\|$. This completes the proof of Theorem 4. □

7.3.2 Proof of Lemma 6

Consider the following optimization problem:

$$\arg\min_v c_{t+1}(v)$$
$$\text{s.t. } x_{t+1} = Ax_t + Bv + w_t. \tag{58}$$

Clearly, the v produced from (58) is the same as the one produced from the "one-step" terminal state problem $\Psi_t^1(x_t, x_{t+1})_{v_0}$. We can then split v such that $v = v_y + v_z$ where $v_y \in \text{Col}(B^\top)$ and $v_z \in \text{Null}(B)$. In particular, v_y must be uniquely determined as

$$v_y = B^\dagger(x_{t+1} - Ax_t - w_t), \tag{59}$$

since $x_{t+1} - Ax_t - w_t$ is in the image of B and for any $v \in \mathbb{R}^m$,

$$\|Bv - x_{t+1} + Ax_t + w_t\| \geq \|(BB^\dagger - I)(x_{t+1} - Ax_t - w_t)\| = 0,$$

that is, $B^\dagger(x_{t+1} - Ax_t - w_t)$ is the least-squares solution. Thus, denote $v_z = B_z\omega$ where B_z is the orthonormal basis matrix for $\text{Null}(B)$ and $\omega \in \mathbb{R}^{m-\text{rank}(B)}$. Thus, the optimization problem in (58) becomes the following unconstrained optimization problem

$$\min_\omega c_{t+1}(v_y + B_z\omega) \tag{60}$$

which achieves its optima at ω such that $B_z^\top \nabla c_{t+1}(v_y + B_z\omega) = 0$. Denote $v' = v_y' + B_z\omega'$ to be $\Psi_t^1(x_t', x_{t+1}')_{v_0}$, then comparing the optimality conditions gives us

$$B_z^\top \left(\nabla c_{t+1}(v_y + B_z\omega) - \nabla c_{t+1}(v_y' + B_z\omega')\right) = B_z^\top G_c(v_y - v_y' + B_z(\omega - \omega')) = 0$$
$$\iff B_z^\top G_c B_z(\omega - \omega') = -B_z^\top G_c(v_y - v_y'), \tag{61}$$

where we've applied Lemma 1 in the first equality. Since B_z is full rank and orthonormal, we have that

$$\|\omega - \omega'\| \leq \frac{L}{\mu}\|v_y - v_y'\| \tag{62}$$

by L-smoothness and μ-strong convexity of $c_{t+1}(\cdot)$. Finally, we have that

$$\|v - v'\|^2 = \|v_y - v_y'\|^2 + \|B_z(\omega - \omega')\|^2 \leq \frac{2L^2}{\mu^2}\|v_y - v_y'\|^2$$
$$\leq \frac{4L^6}{\mu^2}\left(\|x_t - x_t'\|^2 + \|x_{t+1} - x_{t+1}'\|^2\right), \tag{63}$$

where the first inequality comes from B_z being orthonormal and (62). Then the second inequality comes from the bound

$$\|v_y - v_y'\| = \|B^\dagger(x_{t+1} - x_{t+1}' - A(x_t - x_t'))\| \leq L^2(\|x_t - x_t'\| + \|x_{t+1} - x_{t+1}'\|) \tag{64}$$

and then applying the parallelogram identity. Γ^2 from Theorem 2 is of the same magnitude as $\frac{4L^6}{\mu^2}$, so we can take $C^2 \approx \frac{4L^6}{\mu^2}$ which finishes the proof. □

7.4 Simulation Details and Setup

Let T_t be the vector containing the temperature of each zone at time t and U_t the respective integrators. Let u_t be the vector containing the manipulated heat generation/absorption of each zone at time t, and let w_t be the disturbances at time t. then the Euler-discretized dynamics of the system for a sampling time $t_s = 1$ (seconds) are

$$\begin{bmatrix} U_{t+1} \\ T_{t+1} \end{bmatrix} = \begin{bmatrix} I & t_s I \\ 0 & I - t_s L \end{bmatrix} \begin{bmatrix} U_t \\ T_t \end{bmatrix} + \begin{bmatrix} 0 \\ 0.5 t_s I \end{bmatrix} u_t + w_t, \tag{65}$$

where $L[i,j] = k_{ij} = 0.05$ for $i,j \in \mathcal{V}$ is the weighted graph Laplacian and each k_{ij} corresponds to the degree of heat exchange between zones i and j. The disturbances are normally distributed as $w_t \sim \mathcal{N}(0, 25I)$, a multivariate Gaussian random variable with mean 0 and covariance matrix $25I$. As for the costs, we set $f_t(x_t) := \frac{1}{2} x_t^\top Q x_t$ for constant $Q = I$ and all t, $F(x_t) = \frac{1}{2} x_t^\top Q_F x_t$ for constant $Q_F = 10I$ and finally we have $c_{t+1}(u_t) = \frac{1}{2} u_t^\top R_t u_t$ for time varying $R_t = \text{diag}(5|\mathbf{Z}|) + I$ where the $\text{diag}(\cdot)$ operator creates a diagonal matrix whose entries correspond to its input vector's, and $\mathbf{Z}$ is the standard multivariate Gaussian with mean 0 and covariance matrix I. The system assumptions of (A, B) in Assumption 1 are further detailed in Section 5 of [15].