On the Sample Complexity of Stabilizing LTI Systems on a Single Trajectory

Yang Hu
SEAS, Harvard University
Massachusetts, USA
yanghu@g.harvard.edu

Adam Wierman
CMS, California Institute of Technology
California, USA
adamw@caltech.edu

Guannan Qu
ECE, Carnegie Mellon University
Pennsylvania, USA
gqu@andrew.cmu.edu

Abstract

Stabilizing an unknown dynamical system is one of the central problems in control theory. In this paper, we study the sample complexity of the learn-to-stabilize problem in Linear Time-Invariant (LTI) systems on a single trajectory. Current state-of-the-art approaches require a sample complexity linear in n, the state dimension, which incurs a state norm that blows up exponentially in n. We propose a novel algorithm based on spectral decomposition that only needs to learn "a small part" of the dynamical matrix acting on its unstable subspace. We show that, under proper assumptions, our algorithm stabilizes an LTI system on a single trajectory with $O(k \log n)$ samples, where k is the instability index of the system. This represents the first sub-linear sample complexity result for the stabilization of LTI systems under the regime when $k = o(n)$.

1 Introduction

Linear Time-Invariant (LTI) systems, namely $x_{t+1} = Ax_t + Bu_t$, where $x_t \in \mathbb{R}^n$ is the state and $u_t \in \mathbb{R}^m$ is the control input, are one of the most fundamental dynamical systems in control theory, and have wide applications across engineering, economics and societal domains. For systems with known dynamical matrices (A, B), there is a well-developed theory for designing feedback controllers with guaranteed stability, robustness, and performance [1, 2]. However, these tools cannot be directly applied when (A, B) is unknown.

Driven by the success of machine learning [3, 4], there has been significant interest in learning-based (adaptive) control, where the learner does not know the underlying system dynamics and learns to control the system in an online manner, usually with the goal of achieving low regret [5–13].

Despite the progress, an important limitation in this line of work is a common assumption that the learner has a priori access to a known *stabilizing* controller. This assumption simplifies the learning task, since it ensures a bounded state trajectory in the learning stage, and thus enables the learner to learn with low regret. However, assuming a known stabilizing controller is not practical, as *stabilization* itself is nontrivial and considered equally important as any other performance guarantee.

To overcome this limitation, in this paper we consider the *learn-to-stabilize* problem, i.e., learning to stabilize an unknown dynamical system without prior knowledge of any stabilizing controller.

[†]This work is supported by NSF Grants CNS-2146814, CPS-2136197, CNS-2106403, NGSDI-2105648, EPCN-2154171, with additional support from Amazon AWS.

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

Understanding the learn-to-stabilize problem is of great importance to the learning-based control literature, as it serves as a precursor to any learning-based control algorithms that assume knowledge of a stabilizing controller.

The learn-to-stabilize problem has attracted extensive attention recently. For example, [14] and [15] adopt a model-based approach that first excites the open-loop system to learn dynamical matrices (A, B), and then designs a stabilizing controller, with a sample complexity scaling linearly in n, the state dimension. However, a linearly-scaling sample complexity could be unsatisfactory for some specific instances, since the state trajectory still blows up exponentially when the open-loop system is unstable, incurring a $2^{\tilde{\Theta}(n)}$ state norm, and hence a $2^{\tilde{\Theta}(n)}$ regret (in LQR settings, for example). Another recent work [16] proposes a policy-gradient-based discount annealing method that solves a series of discounted LQR problems with increasing discount factors, and shows that the control policy converges to a near-optimal policy. However, this model-free approach only guarantees a $\mathrm{poly}(n)$ sample complexity. In fact, to the best of our knowledge, state-of-the-art learn-to-stabilize algorithms with theoretical guarantees always incur state norms exponential in n.

It has been shown in [15] that all *general-purpose* control algorithms are doomed to suffer a *worst-case* regret of $2^{\Omega(n)}$. This result is intuitive, since from an information-theoretic perspective, a complete recovery of A should take $\Theta(n)$ samples since A itself involves n^2 parameters. However, this does not rule out the possibility that we can achieve better regret in *specific* systems. Our work is motivated by the observation that it is not always necessary to learn the whole matrix A to stabilize an LTI system. For example, if the system is open-loop stable, we do not need to learn anything to stabilize it. For general LTI systems, it is still intuitive that open-loop *stable "modes"* exist and need not be learned for the learn-to-stabilize problem. So, we focus on learning a controller that stabilizes only the *unstable "modes"*, making it possible to learn a stabilizing controller without exponentially exploding state norms. The central question of this paper is:

Can we exploit instance-specific properties of an LTI system to learn to stabilize it on a single trajectory, without incurring a state norm exponentially large in n?

Contribution. In this paper, we answer the above question by designing an algorithm that stabilizes an LTI system with only $O(k \log n)$ state samples along a single trajectory, where k is the *instability index* of the open-loop system and is defined as the number of unstable "modes" (i.e., eigenvalues with moduli larger than 1) of matrix A. Our result is significant in the sense that k can be considerably smaller than n for practical systems and, in such cases, our algorithm stabilizes the system using asymptotically fewer samples than prior work; specifically, it only incurs a state norm (and regret) in the order of $2^{O(k \log n)}$, much smaller than $2^{O(n)}$ of prior state of the art when $k \ll n$.

To formalize the concept of unstable "modes" for the presentation of our algorithm and analysis, we formulate a novel framework based on the spectral decomposition of dynamical matrix A. More specifically, we focus on the *unstable subspace* E_{u} spanned by the eigenvectors corresponding to unstable eigenvalues, and consider the system dynamics "restricted" to it — states are orthogonally projected onto E_{u}, and we only have to learn the effective part of A within subspace E_{u}, which takes only $O(k)$ samples. The formulation is explained in detail in Section 3.1 and Appendix A. We comment that this idea of decomposition is in stark contrast to prior work, which in one way or another seeks to learn the entire A (or other similar quantities).

Related work. Our work contributes to and builds upon related works described below.

Learning for control assuming known stabilizing controllers. There has been a large literature on learning-based control with known stabilizing controllers. For example, one line of research utilizes model-free policy optimization approaches to learn the optimal controller for LTI systems [5–7, 17–30]. All of these works require a known stabilizing controller as an initializer for the policy search method. Another line of research uses model-based methods, i.e., learning dynamical matrices (A, B) first before designing a controller, which also require a known stabilizing controller (e.g., [31–39]). Compared to these works, we focus on the learn-to-stabilize problem without knowledge of an initial stabilizing controller, which can serve as a precursor to existing learning-for-control works that require a known stabilizing controller.

Learning to stabilize on a single trajectory. Stabilizing linear systems over *infinite* horizons with asymptotic convergence guarantees is a classical problem that has been studied extensively in a wide range of papers such as [40–42]. On the other hand, the problem of system stabilization over *finite* horizons remains partially open and has not seen significant progresses. Algorithms incurring a $2^{O(n)}O(\sqrt{T})$ regret have been proposed in settings that rely on relatively strong assumptions of controllability and strictly stable transition matrices [13, 43], which has recently been improved to $2^{\tilde{O}(n)}+\tilde{O}(\text{poly}(n)\sqrt{T})$ [14, 15]. Another model-based approach that merely assumes stabilizability is introduced in [44], though it does not provide guarantees on regret or sample complexity. A more recent model-free approach based on policy gradient [16] provides a novel perspective into this problem, yet it can only guarantee a $\text{poly}(n)$ sample complexity. Compared to these previous works, our approach requires only $O(k \log n)$ samples, incurring a sub-exponential state norm.

Another recent work [45] proposes to do partial system identification via projecting the state onto a lower-dimensional subspace, which is similar in intuition with our work. However, the problem considered there is system stabilization with a fixed initial state x_0, and their approach only eliminates the unstable component along that specific trajectory in k steps when x_0 lies in a k-dimensional subspace. In contrast, our approach finds a stabilizing controller for the system with sub-linear sample complexity along an arbitrary trajectory regardless of the initial state.

Learning to stabilize on multiple trajectories. There are also works [12, 46] that do not assume known stabilizing controllers and learn the full dynamics before designing an optimal stabilizing controller. While requiring $\tilde{\Theta}(n)$ samples which is larger than $\tilde{O}(k)$ of our work, those approaches do not have the exponentially large state norm issue as they allow *multiple trajectories*; i.e., the state can be "reset" to 0 so that it won't get too large. In contrast, we focus on the more challenging single-trajectory scenario where the state cannot be reset.

System Identification. Our work is also related to the system identification literature, which focuses on learning the system parameters of dynamical systems, with early works like [47] focusing on asymptotic guarantees, and more recent works such as [48–53] focusing on finite-time guarantees. Our approach also identifies the system (partially) before constructing a stabilizing controller, but we only identify a part of A rather than the entire A.

2 Problem Formulation

We consider a noiseless LTI system $x_{t+1} = Ax_t + Bu_t$, where $x_t \in \mathbb{R}^n$ and $u_t \in \mathbb{R}^m$ are the *state* and *control input* at time step t, respectively. The dynamical matrices $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$ are unknown to the learner. The learner is allowed to learn about the system by interacting with it on a *single trajectory* — the initial state is sampled uniformly at random from the unit hyper-sphere surface in $\mathbb{R}^n$, and then, at each time step t, the learner is allowed to observe x_t and freely determine u_t. The goal of the learner is to learn a stabilizing controller, which is defined as follows.

Definition 2.1 (Stabilizing Controller)**.** *Control rule $u_t = f_t(x_t, x_{t-1}, \cdots, x_0)$ is called a **stabilizing controller** if and only if the closed-loop system $x_{t+1} = Ax_t + Bu_t$ is asymptotically stable; i.e., for any $x_0 \in \mathbb{R}^n$, $\lim_{t\to\infty} \|x_t\| = 0$ is guaranteed in the closed-loop system.*

To achieve this goal, a simple strategy is to let the system run in open loop to learn (A, B) via least squares, and then design a stabilizing controller based on the learned dynamical matrices. However, as has been discussed in the introduction, such a simple strategy inevitably induces an exponentially large stage norm that is potentially improvable.[1] A possible remedy for this is to learn "a small part" of (A, B) that is crucial for stabilization. Driven by such intuition, the core problem of this paper is to characterize what is the "small part" and design an algorithm to learn it.

Note that, although it is common to include an additive disturbance term w_t in the LTI dynamics, the introduction of stochasticity does not provide additional insights into our decomposition-based algorithm, but rather, merely makes the analysis more technically challenging. Therefore, here we simply omit the noise in theoretical results for the clarity of exposition, and will show by numerical experiments that our algorithm can also handle disturbances (see Appendix H).

Notation. For $z \in \mathbb{C}$, $|z|$ is the modulus of z. For a matrix $A \in \mathbb{R}^{p\times q}$, $A^\top$ denotes the transpose of A; $\|A\|$ is the induced 2-norm of A (equal to its largest singular value), and $\sigma_{\min}(A)$ is the smallest singular value of A; when A is square, $\rho(A)$ denotes the spectral radius of A, and $\kappa_{\mathrm{e}}(A)$ denotes the condition number of the matrix consisting of A's eigenvectors as columns. The space spanned by

[1]More sophisticated exploration strategies might be adopted to learn (A, B) [13, 15, 44], but as long as the control inputs do not completely cancel out the "dominant part" of the states, the above intuition still holds to a large extent as the 'dominant part" of the state is still blowing up exponentially.

$\{v_1, \cdots, v_p\}$ is denoted by $\mathrm{span}(v_1, \cdots, v_p)$, and the column space of A is denoted by $\mathrm{col}(A)$. For two subspaces U, V of $\mathbb{R}^n$, $U^\perp$ is the orthogonal complement of U, and $U \oplus V$ is the direct sum of U and V. The zero matrix and identity matrix are denoted by $\mathbf{0}$, I, respectively.

3 Learning to Stabilize from Zero (LTS$_0$)

The core of this paper is a novel algorithm, Learning to Stabilize from Zero (LTS$_0$), that utilizes a decomposition of the state space based on a characterization of the notion of unstable "modes". The decomposition and other preliminaries for the algorithm are first introduced in Section 3.1, and then we proceed to describe LTS$_0$ in Section 3.2.

3.1 Algorithm Preliminaries

We first introduce the decomposition of the state space in Section 3.1.1, which formally defines the "small part" of A mentioned in the introduction. Then, we introduce τ-hop control in Section 3.1.2, so that we can construct a stabilizing controller based only on the "small part" of A (as opposed to the entire A). Together, these two ideas form the basis of LTS$_0$.

3.1.1 Decomposition of the State Space

Consider the open-loop system $x_{t+1} = Ax_t$. Suppose that A is diagonalizable, and let $\lambda_1, \cdots, \lambda_n$ denote the eigenvalues of A, which are assumed to be distinct and satisfy

$$|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_k| > 1 > |\lambda_{k+1}| \geq \cdots \geq |\lambda_n|.$$

We define the *eigenspaces* associated to these eigenvalues: for a real eigenvalue $\lambda_i \in \mathbb{R}$ corresponding to eigenvector $v_i \in \mathbb{R}^n$, we associate with it a 1-dimensional space $E_i = \mathrm{span}(v_i)$; for a complex eigenvalue $\lambda_i \in \mathbb{C} \setminus \mathbb{R}$ corresponding to eigenvector $v_i \in \mathbb{C}^n$, there must exist some j such that $\lambda_j = \bar{\lambda}_i$ (corresponding to eigenvector $v_j = \bar{v}_i$), and we associate with them a 2-dimensional space $E_i = E_j = \mathrm{span}((v_i + \bar{v}_i), \mathrm{i}(v_i - \bar{v}_i))$. Further, define the *unstable subspace* $E_\mathrm{u} := \bigoplus_{i \leq k} E_i$ and *stable subspace* $E_\mathrm{s} := \bigoplus_{i > k} E_i$.

As discussed earlier, we only need to learn "a small effective part" of A associated with the unstable "modes", or the unstable eigenvectors of A. For this purpose, in the following we formally define a decomposition based on the orthogonal projection onto the unstable subspace E_u. This decomposition forms the foundation of our algorithm.

The $\boldsymbol{E}_\mathbf{u} \oplus \boldsymbol{E}_\mathbf{u}^\perp$-decomposition. Suppose the unstable subspace E_u and its orthogonal complement $E_\mathrm{u}^\perp$ are given by *orthonormal* bases $P_1 \in \mathbb{R}^{n \times k}$ and $P_2 \in \mathbb{R}^{n \times (n-k)}$, respectively, namely

$$E_\mathrm{u} = \mathrm{col}(P_1),\ E_\mathrm{u}^\perp = \mathrm{col}(P_2).$$

Let $P = [P_1\ P_2]$, which is also orthonormal and thus $P^{-1} = P^\top = [P_1\ P_2]^\top$. For convenience, let $\mathit{\Pi}_1 := P_1 P_1^\top$ and $\mathit{\Pi}_2 = P_2 P_2^\top$ be the *orthogonal* projectors onto E_u and $E_\mathrm{u}^\perp$, respectively. With the state space decomposition, we proceed to decompose matrix A. Note that E_u is an invariant subspace with regard to A (but $E_\mathrm{u}^\perp$ not necessarily is), there exists $M_1 \in \mathbb{R}^{k \times k}$, $\mathit{\Delta} \in \mathbb{R}^{k \times (n-k)}$ and $M_2 \in \mathbb{R}^{(n-k) \times (n-k)}$, such that

$$AP = P \begin{bmatrix} M_1 & \mathit{\Delta} \\ & M_2 \end{bmatrix} \Leftrightarrow M := \begin{bmatrix} M_1 & \mathit{\Delta} \\ & M_2 \end{bmatrix} = P^{-1}AP.$$

In the decomposition, the top-left block $M_1 \in \mathbb{R}^{k \times k}$ represents the action of A on the unstable subspace. Matrix M_1, together with P_1, is the "small part" we discussed in the introduction. Note that M_1 (P_1) is only k-by-k (n-by-k) and thus takes much fewer samples to learn compared to the entire A. It is also evident that M_1 inherits all unstable eigenvalues of A, while M_2 inherits all stable eigenvalues. Finally, we provide the system dynamics in the transformed coordinates. Let $y = [y_1^\top\ y_2^\top]^\top$ be the coordinate representation of x in the basis formed by column vectors of P (i.e., $x = Py$). The system dynamics in y-coordinates is

$$\begin{bmatrix} y_{1,t+1} \\ y_{2,t+1} \end{bmatrix} = P^{-1}AP \begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} + P^{-1}Bu_t = \begin{bmatrix} M_1 & \mathit{\Delta} \\ & M_2 \end{bmatrix} \begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} + \begin{bmatrix} P_1^\top B \\ P_2^\top B \end{bmatrix} u_t. \tag{1}$$

The $\boldsymbol{E}_\mathbf{u} \oplus \boldsymbol{E}_\mathbf{s}$-decomposition. In the above $E_\mathrm{u} \oplus E_\mathrm{u}^\perp$-decomposition, $E_\mathrm{u}^\perp$ is in general *not* an invariant subspace with respect to A. This can be seen from the top-right $\mathit{\Delta}$ block in M, which represents how much of the state is "moved" by A from $E_{\mathrm{u}}^{\perp}$ into E_{u} in one step. The absence of invariant properties in $E_{\mathrm{u}}^{\perp}$ is sometimes inconvenient in the analysis. Hence, we introduce another invariant decomposition that is used in the proof as follows. Specifically, $\mathbb{R}^n$ can be naturally decomposed into $E_{\mathrm{u}} \oplus E_{\mathrm{s}}$, and further both E_{u} and E_{s} are invariant with respect to A. We also represent $E_{\mathrm{u}} = \mathrm{col}(Q_1)$ and $E_{\mathrm{s}} = \mathrm{col}(Q_2)$ by their *orthonormal* bases, and define $Q = [Q_1\ Q_2]$. Note that, these two subspaces are generally not orthogonal, so we additionally define $Q^{-1} =: [R_1^\top R_2^\top]^\top$. Details are deferred to Appendix A.1.

Lastly, we comment that when A is symmetric, the $E_{\mathrm{u}} \oplus E_{\mathrm{u}}^{\perp}$- and $E_{\mathrm{u}} \oplus E_{\mathrm{s}}$-decompositions are identical because $E_{\mathrm{u}}^{\perp} = E_{\mathrm{s}}$ in such symmetric cases. While $E_{\mathrm{u}}^{\perp} \neq E_{\mathrm{s}}$ in general cases, the "closeness" between $E_{\mathrm{u}}^{\perp}$ and E_{s} also contributes to the sample complexity bound in Section 4. For that reason, we formally define such "closeness" between subspaces in Definition 3.1. We point out that the definition has clear geometric interpretations and leads to connections between the bases of E_{s} and $E_{\mathrm{u}}^{\perp}$, which is technical and thus deferred to Appendix A.2.

Definition 3.1 (ξ-Close Subspaces). *For $\xi \in (0, 1]$, the subspaces $E_{\mathrm{u}}^{\perp} = \mathrm{col}(P_2), E_{\mathrm{s}} = \mathrm{col}(Q_2)$ are called **ξ-close** to each other, if and only if $\sigma_{\min}(P_2^\top Q_2) > 1 - \xi$.*

3.1.2 τ-hop Control

This section discusses the design of controller based only on the "small part" of A, i.e., the P_1 and M_1 matrices discussed in Section 3.1.1, as opposed to the entire A matrix. Note that the main objective of this subsection is to introduce the idea of our controller design when M_1 and P_1 are known without errors, whereas in Section 3.2 we fully introduce Algorithm 1 that learns M_1 and P_1 before constructing the stabilizing controller.

As discussed in Section 3.1.1, we can view M_1 as the "restriction" of A onto the unstable subspace E_{u} (spanned by the basis in P_1) and it captures all the unstable eigenvalues of A. Since only M_1 and P_1 are known while M_2 and P_2 are unknown, a simple idea is to "restrict" the system trajectory entirely to E_{u} such that the effect of A is fully captured by M_1, the part of A that is known. However, such a restriction is not possible because, even if the current state x_t is in E_{u} (so Ax_t is also in E_{u}), $x_{t+1} = Ax_t + Bu_t$ is generally not in E_{u} with non-zero u_t. To address this issue, recall that a desirable property of the stable component is that it spontaneously dies out in open loop. Therefore, we propose the following *τ-hop controller* design, where the control input is only injected every τ steps — in this way, we let the stable component die out exponentially between two consecutive control injections. Consequently, when we examine the states every τ steps, we could expect that the trajectory appears approximately "restricted to" the unstable subspace E_{u}.

More formally, a τ-hop controller only injects non-zero u_t for $t = s\tau$, $s \in \mathbb{N}$. Let $\tilde{x}_s := x_{s\tau}$ and $\tilde{u}_s := u_{s\tau}$ to be the state and input every τ time steps. We can write the dynamics of the τ-hop control system as $\tilde{x}_{s+1} = A^\tau \tilde{x}_s + A^{\tau-1} B \tilde{u}_s$. We also let $\tilde{y}_s$ to denote the state under $E_{\mathrm{u}} \oplus E_{\mathrm{u}}^{\perp}$-decomposition, i.e. $\tilde{y}_s = P^\top \tilde{x}_s$. Then the state evolution can be written as

$$\begin{bmatrix} \tilde{y}_{1,s+1} \\ \tilde{y}_{2,s+1} \end{bmatrix} = P^{-1} A^\tau P \begin{bmatrix} \tilde{y}_{1,s} \\ \tilde{y}_{2,s} \end{bmatrix} + P^{-1} A^{\tau-1} B \tilde{u}_s = M^\tau \begin{bmatrix} \tilde{y}_{1,s} \\ \tilde{y}_{2,s} \end{bmatrix} + \begin{bmatrix} P_1^\top A^{\tau-1} B \\ P_2^\top A^{\tau-1} B \end{bmatrix} \tilde{u}_s, \tag{2}$$

where we define $B_\tau := P_1^\top A^{\tau-1} B$ for simplicity, and

$$M^\tau = \left(\begin{bmatrix} M_1 & \\ & M_2 \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \Delta \\ & \mathbf{0} \end{bmatrix} \right)^\tau = \begin{bmatrix} M_1^\tau & \sum_{i=0}^{\tau-1} M_1^i \Delta M_2^{\tau-1-i} \\ & M_2^\tau \end{bmatrix} =: \begin{bmatrix} M_1^\tau & \Delta_\tau \\ & M_2^\tau \end{bmatrix}.$$

Now we consider a state feedback controller $\tilde{u}_s = K_1 \tilde{y}_{1,s}$ in the τ-hop control system that only acts on the unstable component $\tilde{y}_{1,s}$, the closed-loop dynamics of which can then be written as

$$\tilde{y}_{s+1} = \begin{bmatrix} M_1^\tau + P_1^\top A^{\tau-1} B K_1 & \Delta_\tau \\ P_2^\top A^{\tau-1} B K_1 & M_2^\tau \end{bmatrix} \tilde{y}_s. \tag{3}$$

In (3), the bottom-left block becomes $P_2^\top A^{\tau-1} B K_1$, which is exponentially small in τ. Therefore, with a properly chosen τ, the closed-loop dynamical matrix in (3) is almost block-upper-triangular with the bottom-right block very close to $\mathbf{0}$ (recall that M_2 is a stable matrix). As a result, if we select K_1 such that $M_1^\tau + P_1^\top A^{\tau-1} B K_1$ is stable, then (3) will become stable as well. There are different ways to select such K_1, and in this paper, we focus on the simple case that B is an n-by-k matrix and $P_1^\top A^{\tau-1} B$ is an invertible square matrix (see Assumption 4.3′), in which case selecting

$$K_1 = -(P_1^\top A^{\tau-1} B)^{-1} M_1^\tau \tag{4}$$

will suffice. Note that such a controller design will also need the knowledge of $P_1^\top A^{\tau-1}B$, which has the same dimension as M_1 (a k-by-k matrix) and takes only $O(k)$ additional samples to learn. For the case that B is not n-by-k, similar controller design can be done (but in a slightly more involved way), and we defer the discussion to Appendix C.

We also point out that, for the case where A is symmetric, selecting $\tau = 1$ should work well. This is because $\varDelta_\tau = \mathbf{0}$ in (3) for the symmetric case, and therefore, the matrix in (3) will be triangular even for $\tau = 1$. This will result in a simpler algorithm and controller design, and hence a better sample complexity bound, which we will present as Theorem 4.2 in Section 4.

We end this subsection with some comments on the role of τ-hop stabilizing controllers. One may wonder if the controller design proposed here would be compatible with many downstream tasks, since the closed-loop system stabilized by a τ-hop controller will still experience periodical fluctuations in state norms (although in a bounded manner). However, we want to emphasize again that the τ-hop controller can serve as a precursor to any online control algorithm that assumes a known stabilizing controller, which includes system identification from stable trajectories (see, e.g., [48, 50]) and controller designs using the identified system. In this way the state norm fluctuation is only transient, and does not harm to the overall performance significantly.

3.2 Algorithm

Our algorithm, LTS_0, is divided into 4 stages: (i) learn an orthonormal basis P_1 of the unstable subspace E_u (Stage 1); (ii) learn M_1, the restriction of A onto the subspace E_u (Stage 2); (iii) learn $B_\tau = P_1^\top A^{\tau-1}B$ (Stage 3); and (iv) design a controller that seeks to cancel out the "unstable" M_1 matrix (Stage 4). This is formally described as Algorithm 1 below.

Algorithm 1 LTS_0: Learning a τ-hop Stabilizing Controller

Stage 1: learn the unstable subspace of $\boldsymbol{A}$.
Run the system in open loop for t_0 steps for initialization.
Run the system in open loop for k more steps and let $D \leftarrow [x_{t_0+1} \ \cdots \ x_{t_0+k}]$.
Calculate $\hat{\varPi}_1 \leftarrow D(D^\top D)^{-1}D^\top$.
Calculate the top k (normalized) eigenvectors $\hat{v}_1, \cdots \hat{v}_k$ of $\hat{\varPi}_1$, and let $\hat{P}_1 \leftarrow [\hat{v}_1 \ \cdots \ \hat{v}_k]$.
Stage 2: approximate $\boldsymbol{M_1}$ on the unstable subspace.
Solve the least squares $\hat{M}_1 \leftarrow \arg\min_{M_1 \in \mathbb{R}^{k\times k}} \mathcal{L}(M_1) := \sum_{t=t_0+1}^{t_0+k} \|\hat{P}_1^\top x_{t+1} - \hat{M}_1 \hat{P}_1^\top x_t\|^2$.
Stage 3: restore $\boldsymbol{B_\tau}$ for $\boldsymbol{\tau}$-hop control.
for $i \leftarrow 1, \cdots, k$ **do**
    Let the system run in open loop for ω time steps.
    Run for τ more steps with initial $u_{t_i} = \alpha\|x_{t_i}\|e_i$, where $t_i = t_0 + k + i\omega + (i-1)\tau$.
Let $\hat{B}_\tau \leftarrow [\hat{b}_1 \cdots \hat{b}_k]$, where the i^{th} column $\hat{b}_i \leftarrow \frac{1}{\alpha\|x_{t_i}\|}\left(\hat{P}_1^\top x_{t_i+\tau} - \hat{M}_1^\tau \hat{P}_1^\top x_{t_i}\right)$.
Stage 4: construct a $\boldsymbol{\tau}$-hop stabilizing controller $\boldsymbol{K}$.
Construct the τ-hop stabilizing controller $\hat{K} \leftarrow -\hat{B}_\tau^{-1}\hat{M}_1^\tau \hat{P}_1^\top$.

In the remainder of this section we provide detailed descriptions of the four stages in LTS_0.

Stage 1: Learn the unstable subspace of $\boldsymbol{A}$. It suffices to learn an orthonormal basis of E_u. We notice that, when A is applied recursively, it will push the state closer to E_u. Therefore, when we let the system run in open loop (with control input $u_t \equiv 0$) for t_0 time steps, the ratio between the norms of unstable and stable components will be magnified exponentially, and the state lies "almost" in E_u. As a result, the subspace spanned by the next k states, i.e. the column space of $D := [x_{t_0+1} \ \cdots \ x_{t_0+k}]$, is very close to E_u. This motivates us to use the orthogonal projector onto $\mathrm{col}(D)$, namely $\hat{\varPi}_1 = D(D^\top D)^{-1}D^\top$, as an estimation of the projector $\varPi_1 = P_1 P_1^\top$ onto E_u. Finally, the columns of $\hat{P}_1$ are restored by taking the top k eigenvectors of $\hat{\varPi}_1$ with largest eigenvalues (they should be very close to 1), which form a basis of the estimated unstable subspace.

Stage 2: Learn $\boldsymbol{M_1}$ on the unstable subspace. Recall that M_1 is the "dynamical matrix" for the E_u-component under the $E_\mathrm{u} \oplus E_\mathrm{u}^\perp$-decomposition. Therefore, to estimate M_1, we first calculate the coordinates of the states $x_{t_0+1:t_0+k}$ under basis P_1; that is, $\hat{y}_{1,t} = \hat{P}_1^\top x_t$, for $t = t_0 + 1, \ldots, t_0 + k$.

Then, we use least squares to estimate M_1, which minimizes the square loss over $\hat{M}_1$

$$\mathcal{L}(\hat{M}_1) := \sum_{t=t_0+1}^{t_0+k} \|\hat{y}_{1,t+1} - \hat{M}_1 \hat{y}_{1,t}\|^2 = \sum_{t=t_0+1}^{t_0+k} \|\hat{P}_1^\top x_{t+1} - \hat{M}_1 \hat{P}_1^\top x_t\|^2. \tag{5}$$

It can be shown that the unique solution to (5) is $\hat{M}_1 = \hat{P}_1^\top A \hat{P}_1$ (see Appendix B).

Stage 3: Restore B_τ for τ-hop control. In this step, we restore the B_τ that quantifies the "effective component" of control inputs restricted to E_u (see Section 3.1.2 for detailed discussion). Note that equation (2) can be rewritten in terms of $y_{1,t}$ as

$$y_{1,t_i+\tau} = M^\tau y_{1,t_i} + \Delta_\tau y_{2,t_i} + B_\tau u_{t_i}.$$

Hence, for the purpose of estimation, we simply ignore the Δ_τ term, and take the i^{th} column as

$$\hat{b}_i \leftarrow \frac{1}{\|u_{t_i}\|} \left(\hat{P}_1^\top x_{t_i+\tau} - \hat{M}_1^\tau \hat{P}_1^\top x_{t_i}\right),$$

where u_{t_i} is parallel to e_i, and the magnitude of u_{t_i} is set to be large enough as $\alpha\|x_{t_i}\|$ to amplify its effect so that the estimation error of A is comparatively negligible. Here we introduce an adjustable constant α to guarantee that the E_u-component still constitutes a non-negligible proportion of the state after injecting u_{t_i}, so that the iterative restoration of columns could continue.

It is evident that the ignored $\Delta_\tau P_2^\top x_{t_i}$ term will introduce an extra estimation error. Since Δ_τ contains a factor of $M_1^{\tau-1}\Delta$ that explodes with respect to τ, this part can only be bounded if $\frac{\|P_2^\top x_{t_i}\|}{\|x_{t_i}\|}$ is sufficiently small. For this purpose, we introduce ω heat-up steps (running in open loop with 0 control input) to reduce the ratio to an acceptable level, during which time the projection of state onto $E_\mathrm{u}^\perp$ automatically diminishes over time since $\rho(M_2) = |\lambda_{k+1}| < 1$.

Stage 4: Construct a τ-hop stabilizing controller K. Finally, we can design a controller that cancels out M_1^τ in the τ-hop system. As mentioned in Section 3.1.2, we shall focus on the case where B is an n-by-k matrix for the sake of exposition (the case for general B will be discussed in Appendix C). The invertibility of B_τ can be guaranteed under certain conditions (Assumption 4.3$'$); further, $\hat{B}_\tau$ is also invertible as long as it is close enough to B_τ. In this case, the τ-hop stabilizing controller can be simpliy designed as $\hat{K}_1 = -\hat{B}_\tau^{-1}\hat{M}_1^\tau$ in y-coordinates where we replace B_τ and M_1 in (4) with their estimates. When we return to the original x-coordinates, the controller becomes $\hat{K} = -\hat{B}_\tau^{-1}\hat{M}_1^\tau \hat{P}_1^\top$. Note that $\hat{K}$ (and $\hat{K}_1$) appears with a hat to emphasize the use of estimated projector $\hat{P}_1$, which introduces an extra estimation error to the final closed-loop dynamics.

It is evident that the algorithm terminates in $t_0 + k(1 + \omega + \tau)$ time steps. In the next section, we show how to choose the parameters to guarantee both stability and sub-linear sample complexity.

Finally, we remark that, although for the ease of exposition we have assumed here the instability index k is known, it is fine to use an estimate of k that is larger than its true value in practice — i.e., the algorithm still outputs a stabilizing controller since the performance analysis only relies on the ratio between eigenvalues and the stability of λ_{k+1}, and the complexity only suffers little if the guess of k is close to its true value.

4 Stability Guarantee

In this section, we formally state the assumptions and show the sample complexity for the proposed algorithm to find a stabilizing controller. Our first assumption is regarding the spectral properties of A, where we require all eigenvalues to appear without multiplicity (so that we can learn a complete basis of each eigenspace), and marginally stable eigenvalues (i.e., those with moduli 1) are eliminated (so that eigenspaces are either stable or unstable). We would like to point out that it is common practice (e.g., [50]) to discuss marginally stable eigenvalues separately, since it obscures the distinction between stable and unstable components and is thus technically challenging.

Assumption 4.1 (Spectral Property)**.** *A is diagonalizable with instability index k, with distinct eigenvalues $\lambda_1, \cdots, \lambda_n$ satisfying $|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_k| > 1 > |\lambda_{k+1}| \geq \cdots \geq |\lambda_n|$.*

The assumption is mild in the sense that matrices satisfying Assumption 4.1 are dense in $\mathbb{R}^{n\times n}$, and our final complexity bound only depends logarithmically on the condition number of eigenvectors

$\kappa_{\mathrm{e}}(A)$ and the eigen-gap λ_k/λ_{k+1} (see Theorem 4.1 and the discussion below). Thus any matrix A that violates Assumption 4.1 can be handled via small perturbations.

Our second assumption is regarding how to choose the initial state, which again is standard. The initialization must be randomized to eliminate the coincidence where x_0 has zero (oblique) projection onto some eigenvector v_i, in which case we cannot learn about v_i and thus D is not invertible.

Assumption 4.2 (Initialization)**.** *The initial state of the system is sampled uniformly at random on the unit hyper-sphere surface in* $\mathbb{R}^n$*.*

Lastly, we assume the system to be (d, σ)-strongly controllable, which is standard in literature.

Assumption 4.3 ((ν, σ)-Strong Controllability)**.** *The system is* (ν, σ)*-strongly controllable; i.e.,* $\sigma_{\min}(C_\nu) > \sigma$*, where* $C_\nu := [A^{\nu-1}B \;\; A^{\nu-2}B \;\; \cdots \;\; AB \;\; B]$ *is the* ν*-step controllability matrix.*

Above are all the assumptions we need. However, we remind the readers that, when we introduce the τ-hop controller design in Section 3.1.2, B is assumed to have k columns and certain assumptions are needed to guarantee the invertibility of B_1. Indeed, for the ease of exposition, we first consider this special case in presenting our main result (Theorem 4.1) below, where we impose the following Assumption 4.3′ regarding the controllability within the unstable subspace E_{u} instead of the more general Assumption 4.3 (recall that R_1 is defined in the $E_{\mathrm{u}} \oplus E_{\mathrm{s}}$-decomposition in Section 3.1.1). Discussions on how to handle the more general Assumption 4.3 via a transformation to the special case (where Assumption 4.3′ holds) are deferred to Appendix C.

Assumption 4.3′ (c-Effective Control in Unstable Subspace)**.** $B \in \mathbb{R}^{n\times k}$, $\sigma_{\min}(R_1 B) > c\|B\|$.

Note that Assumption 4.3′ has a clear intuition — every direction in the unstable subspace receives at least a proportion of c from the influence of any control input. This assumption is reasonable in that, if $\sigma_{\min}(R_1 B) \approx 0$, the control input u has to be very large to push the state along the direction corresponding to the smallest singular value, which could induce excessively large control cost. We can also interpret the lower bound on $\sigma_{\min}(R_1 B)$ as a special case of Assumption 4.3 (i.e., $(1, c\|B\|)$-strong controllablility). Details can be found in Appendix C.

In the following we present the main performance guarantees for our algorithm.

Theorem 4.1 (Main Theorem)**.** *Given a noiseless LTI system* $x_{t+1} = Ax_t + Bu_t$ *subject to Assumptions 4.1, 4.2 and 4.3′, and additionally* $|\lambda_1|^2|\lambda_{k+1}| < |\lambda_k|$*, by running LTS*$_0$ *with parameters*

$$\tau = O(1),\; \omega = O(\ell \log k),\; \alpha = O(1),\; t_0 = O(k \log n)$$

that terminates within $t_0 + k(1+\omega+\tau) = O(k \log n)$ *time steps, the closed-loop system is exponentially stable with probability* $1 - O(k^{-\ell})$ *over the initialization of* x_0 *for any* $\ell \in \mathbb{N}$*. Here the big-O notation only shows dependence on* k *and* n*, while hiding parameters like* $|\lambda_1|$*,* $|\lambda_k|$*,* $|\lambda_{k+1}|$*,* $\|A\|$*,* $\|B\|$*,* c*,* α*,* ξ *(recall that* $E_{\mathrm{u}}^\perp$ *and* E_{s} *are* ξ*-close),* $\chi(\hat{L}_\tau)$ *(see Lemma D.1), and* $\zeta_\varepsilon(\cdot)$ *(see Lemma G.1), and details can be found in equations (41) through (46).*

Theorem 4.1 shows the proposed LTS$_0$ algorithm can find a stabilizing controller in $\tilde{O}(k)$ steps, which incurs a state norm of $2^{\tilde{O}(k)}$, significantly smaller than the state-of-the-art $2^{\Theta(n)}$ in the $k \ll n$ regime. We would like to point out that this does not violate the lower bound shown in [15], since the state norm degenerates to $2^{\Theta(n)}$ when $k = \Theta(n)$, and might degrade arbitrarily for systems with adversarially designed parameters. Still, for a large proportion of systems with $k \ll n$ and favorable constants, our algorithm achieves better performance than the naive ones. The theoretical result is also verified by numerical experiments, the details of which can be found in Appendix H.

Discussion on constants. Curious readers can refer to Appendix G (equations (41) through (46)) for detailed expressions of the constants hidden behind the big-O notation in the theorem; Table 1 also summarizes all instance-specific constants appearing in the bound. Here we provide a brief overview how the bound depends on the system parameters. It is evident that, for a system with larger ξ (i.e., when E_{u} and E_{s} are "less orthogonal" to each other) or smaller c (i.e., when it costs more to control the unstable subspace), we see a larger τ in (41), a smaller α in (43), and larger t_0 and ω in (45) and (46), respectively, which altogether incur a larger constant term in the sample complexity. This is in accordance with our intuition of the state space decomposition and Assumption 4.3′, respectively.

The bound also relies heavily on the spectral properties of A. The constraint $|\lambda_1|^2|\lambda_{k+1}| < |\lambda_k|$ ensures validity of (41), which is necessary for cancelling out the combined effect of non-orthogonal subspaces E_u and E_s (resulting in $\varDelta_\tau$ in the top-right block) and inaccurate basis $\hat{P}_1$ (resulting in projection error in the bottom-left block) — a system with larger ratio $|\lambda_1|^2|\lambda_{k+1}|/|\lambda_k|$ suffers from more severe side-effects, and thus requires a larger τ and a higher sample complexity. Nevertheless, we believe that this assumption is not essential, and we leave it as future work to relax it.

Another important parameter is the eigen-gap $|\lambda_k|/|\lambda_{k+1}|$ around 1 that determines how fast the stable and unstable components become separable in magnitude when the system runs in open loop, which is utilized in the t_0 initialization steps of Stage 1 and ω heat-up steps of Stage 3. Consequently, a system with smaller eigen-gap $|\lambda_k|/|\lambda_{k+1}|$ requires a larger t_0 (see (10)) and ω (see (46)) and therefore a higher sample complexity.

The condition number of eigenvectors $\kappa_\mathrm{e}(A)$ also contributes to the bound of t_0, the number of initialization steps. It is intuitive that, a large $\kappa_\mathrm{e}(A)$ indicates less orthogonal eigenspaces, which in turn requires a more distinct separation among the magnitudes of different eigen-components of x_{t_0}, so that the stable components interfere less with the unstable ones.

Finally, we would like to point out that all these quantities appear in the bound as *logarithmic* terms, indicating that the sample complexity only degrades mildly when the constants become worse.

A warm-up case. Despite the generality of Theorem 4.1, its proof involves technical difficulties. In Theorem 4.2, we include results for the special case where A is real symmetric, which leads to a simpler choice of algorithm parameters and a cleaner sample complexity bound.

Theorem 4.2. *Given a noiseless LTI system $x_{t+1} = Ax_t + Bu_t$ subject to Assumptions 4.1, 4.2 and 4.3′ with symmetric A, by running LTS$_0$ with parameters $\tau = 1$, $\omega = 0$, $\alpha = 1$, $t_0 = O(k \log n)$ that terminates within $t_0 + k(1+\omega+\tau) = O(k \log n)$ time steps, the closed-loop system is exponentially stable with probability 1 over the initialization of x_0. Here the big-O notation only shows dependence on k and n, while hiding parameters like $|\lambda_1|$, $|\lambda_k|$, $|\lambda_{k+1}|$, $\|A\|$, $\|B\|$, c, and $\chi(\hat{L}_1)$ (see Lemma D.1), and details can be found in equation (18).*

Although Theorem 4.2 takes a simpler form, its proof still captures the main insight of our analysis. For this reason, we use the proof of Theorem 4.2 as a warm-up example in Appendix F before we present the proof ideas of the main Theorem 4.1.

5 Proof Outline

In this section we will give a high-level overview of the key proof ideas for the main theorems. The full proof details can be found in Appendices E, F and G as indicated below.

Proof Structure. The proof is largely divided into two steps. In Step 1, we examine how accurate the learner estimates the unstable subspace E_u in Stage 1 and 2. We will show that $\varPi_1$, P_1 and M_1 can be estimated up to an error of δ within $t_0 = O(k \log n - \log \delta)$ steps. In Step 2, we examine the estimation error of M_1 and B_τ in Stage 2 and 3 (and thus $\hat{K}_1$), based on which we will eventually show that the τ-hop controller output by Algorithm 1 makes the system asymptotically stable. The proof is based on a detailed spectral analysis of the closed-loop dynamical matrix.

Overview of Step 1. To upper bound the estimation errors in Stage 1 and 2, we only have to notice that the estimation error of $\varPi_1$ completely captures how well the unstable subspace is estimated, and all other bounds should follow directly from it. The bound on $\|\varPi_1 - \hat{\varPi}_1\|$ is shown in Theorem 5.1, together with a bound on $\|P_1 - \hat{P}_1\|$ presented in Corollary 5.2.

Theorem 5.1. *For a noiseless linear dynamical system $x_{t+1} = Ax_t$, let E_u be the unstable subspace of A, $k = \dim E_\mathrm{u}$ be the instability index of the system, and $\varPi_1$ be the orthogonal projector onto subspace E_u. Then for any $\varepsilon > 0$, by running Stage 1 of Algorithm 1 with an arbitrary initial state that terminates in $(t_0 + k)$ time steps, where*

$$t_0 = O\left(\frac{k \log n - \log \varepsilon + \log \kappa_\mathrm{e}(A)}{2 \log \frac{|\lambda_k|}{|\lambda_{k+1}|}}\right),$$

the matrix $D^\top D$ is invertible with probability 1 (where $D = [x_{t_0+1} \ \cdots \ x_{t_0+k}]$), and in such cases we shall obtain an estimated $\hat{\varPi}_1 = D(D^\top D)^{-1} D^\top$ with error $\|\hat{\varPi}_1 - \varPi_1\| < \varepsilon$.

Corollary 5.2. *Under the premises of Theorem 5.1, for any orthonormal basis* $\hat{P}_1$ *of* $\mathrm{col}(\hat{\Pi}_1)$ *(where* $\hat{\Pi}_1$ *is obtained by Algorithm 1), there exists a corresponding orthonormal basis* P_1 *of* $\mathrm{col}(\Pi_1)$*, such that* $\|\hat{P}_1 - P_1\| < \sqrt{2k}\varepsilon =: \delta$, $\|\hat{M}_1 - M_1\| < 2\|A\|\delta$.

The proofs are deferred to Appendix E due to limited length.

Overview of Step 2. To analyze the stability of the closed-loop system, we shall first write out the closed-loop dynamics under the τ-hop controller. Recall in Section 3.1.2 we have defined $\tilde{u}_s, \tilde{x}_s, \tilde{y}_s$ to be the control input, state in x-coordinates, and state in y-coordinates in the τ-hop control system, respectively. Using these notations, the learned controller can be written as

$$\tilde{u}_s = \hat{K}\tilde{x}_s = \hat{K}_1\hat{P}_1^\top P\tilde{y}_s = \begin{bmatrix} \hat{K}_1\hat{P}_1^\top P_1 \\ \hat{K}_1\hat{P}_1^\top P_2 \end{bmatrix} \tilde{y}_s$$

in y-coordinates (as opposed to $\hat{K}_1\tilde{y}_s$). Therefore, the closed-loop τ-hop dynamics should be

$$\tilde{y}_{s+1} = \begin{bmatrix} M_1^\tau + P_1^\top A^{\tau-1}B\hat{K}_1\hat{P}_1^\top P_1 & \Delta_\tau + P_1^\top A^{\tau-1}B\hat{K}_1\hat{P}_1^\top P_2 \\ P_2^\top A^{\tau-1}B\hat{K}_1\hat{P}_1^\top P_1 & M_2^\tau + P_2^\top A^{\tau-1}B\hat{K}_1\hat{P}_1^\top P_2 \end{bmatrix} \begin{bmatrix} \tilde{y}_{1,s} \\ \tilde{y}_{2,s} \end{bmatrix} =: \hat{L}_\tau\tilde{y}_s, \qquad (6)$$

and we will show it to be asymptotically stable (i.e., $\rho(\hat{L}_\tau) < 1$). Note that $\hat{L}_\tau$ is given by a 2-by-2 block form, we can utilize the following lemma to assist the spectral analysis of block matrices, the proof of which is deferred to Appendix D.

Lemma 5.3 (Block Perturbation Bound)**.** *For 2-by-2 block matrices* $A = \left[\begin{smallmatrix} A_1 & \mathbf{0} \\ \mathbf{0} & A_2 \end{smallmatrix}\right]$, $E = \left[\begin{smallmatrix} \mathbf{0} & E_{12} \\ E_{21} & \mathbf{0} \end{smallmatrix}\right]$, *the spectral radii of* A *and* $A + E$ *differ by at most* $|\rho(A+E) - \rho(A)| \le \chi(A+E)\|E_{12}\|\|E_{21}\|$, *where* $\chi(A+E)$ *is a constant (see Appendix D).*

The above lemma shows a clear roadmap for proving $\rho(\hat{L}_\tau) < 1$. First, we need to guarantee stability of the diagonal blocks — the top-left block is stable because $\hat{K}_1$ is designed to (approximately) eliminate it to zero (which requires the estimation error bound on B_τ), and the bottom-right block is stable because it is almost M_2^τ with a negligible error induced by inaccurate projection. Then, we need to upper-bound the norms of off-diagonal blocks via careful estimation of factors appearing in these blocks. Complete proofs for both cases can be found in Appendices F and G, respectively.

6 Conclusions

This paper provides a new perspective into the learn-to-stabilize problem. We design a novel algorithm that exploits instance-specific properties to learn to stabilize an unknown LTI system on a single trajectory. We show that, under certain assumptions, the sample complexity of the algorithm is upper bounded by $O(k \log n)$, which avoids the $2^{\Theta(n)}$ state norm blow-up of existing methods in the $k \ll n$ regime. This work initiates a new direction in the learn-to-stabilize literature, and many interesting and challenging questions remain open, including handling noises, eliminating the assumptions on spectral properties, and developing better ways to learn the unstable subspace.

References

[1] John C. Doyle, Bruce A. Francis, and Allen R. Tannenbaum. *Feedback Control Theory*. Courier Corporation, 2013.

[2] Geir E. Dullerud and Fernando Paganini. *A Course in Robust Control Theory: A Convex Approach*, volume 36. Springer Science & Business Media, 2013.

[3] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *arXiv preprint arXiv:1504.00702*, 2015.

[4] Yan Duan, Xi Chen, Rein Houthooft, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, pages 1329–1338, 2016.

[5] Maryam Fazel, Rong Ge, Sham M. Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. *arXiv preprint arXiv:1801.05039*, 2018.

[6] Jingjing Bu, Afshin Mesbahi, Maryam Fazel, and Mehran Mesbahi. LQR through the lens of first order methods: Discrete-time case. *arXiv preprint arXiv:1907.08921*, 2019.

[7] Yingying Li, Yujie Tang, Runyu Zhang, and Na Li. Distributed reinforcement learning for decentralized linear quadratic control: A derivative-free policy optimization approach. *arXiv preprint arXiv:1912.09135*, 2019.

[8] Steven J. Bradtke, B. Erik Ydstie, and Andrew G. Barto. Adaptive linear quadratic control using policy iteration. In *Proceedings of 1994 American Control Conference-ACC'94*, volume 3, pages 3475–3479. IEEE, 1994.

[9] Stephen Tu and Benjamin Recht. Least-squares temporal difference learning for the linear quadratic regulator. *arXiv preprint arXiv:1712.08642*, 2017.

[10] Karl Krauth, Stephen Tu, and Benjamin Recht. Finite-time analysis of approximate policy iteration for the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 8512–8522, 2019.

[11] Kemin Zhou, John Comstock Doyle, Keith Glover, et al. *Robust and Optimal Control*, volume 40. Prentice Hall New Jersey, 1996.

[12] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, pages 1–47, 2019.

[13] Yasin Abbasi-Yadkori and Csaba Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 1–26, 2011.

[14] Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Explore more and improve regret in linear quadratic regulators, 2020.

[15] Xinyi Chen and Elad Hazan. Black-box control for linear dynamical systems. *arXiv preprint arXiv:2007.06650*, 2021.

[16] Juan C. Perdomo, Jack Umenberger, and Max Simchowitz. Stabilizing dynamical systems via policy gradient methods. *arXiv preprint arXiv:2110.06418*, 2021.

[17] Tankred Rautert and Ekkehard W. Sachs. Computational design of optimal output feedback controllers. *SIAM Journal on Optimization*, 7(3):837–852, 1997.

[18] Karl Mårtensson and Anders Rantzer. Gradient methods for iterative distributed control synthesis. In *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pages 549–554. IEEE, 2009.

[19] Dhruv Malik, Ashwin Pananjady, Kush Bhatia, Koulik Khamaru, Peter L. Bartlett, and Martin J. Wainwright. Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. *arXiv preprint arXiv:1812.08305*, 2018.

[20] Hesameddin Mohammadi, Armin Zare, Mahdi Soltanolkotabi, and Mihailo R. Jovanović. Convergence and sample complexity of gradient methods for the model-free linear quadratic regulator problem. *arXiv preprint arXiv:1912.11899*, 2019.

[21] Benjamin Gravell, Peyman Mohajerin Esfahani, and Tyler Summers. Learning robust controllers for linear quadratic systems with multiplicative noise via policy gradient. *arXiv preprint arXiv:1905.13547*, 2019.

[22] Zhuoran Yang, Yongxin Chen, Mingyi Hong, and Zhaoran Wang. On the global convergence of actor-critic: A case for linear quadratic regulator with ergodic cost. *arXiv preprint arXiv:1907.06246*, 2019.

[23] Kaiqing Zhang, Zhuoran Yang, and Tamer Basar. Policy optimization provably converges to Nash equilibria in zero-sum linear quadratic games. In *Advances in Neural Information Processing Systems*, pages 11602–11614, 2019.

[24] Kaiqing Zhang, Bin Hu, and Tamer Basar. Policy optimization for H_2 linear control with H_∞ robustness guarantee: Implicit regularization and global convergence. In *Learning for Dynamics and Control*, pages 179–190, 2020.

[25] Luca Furieri, Yang Zheng, and Maryam Kamgarpour. Learning the globally optimal distributed LQ regulator. In *Learning for Dynamics and Control*, pages 287–297, 2020.

[26] Joao Paulo Jansch-Porto, Bin Hu, and Geir Dullerud. Convergence guarantees of policy optimization methods for Markovian jump linear systems. *arXiv preprint arXiv:2002.04090*, 2020.

[27] Joao Paulo Jansch-Porto, Bin Hu, and Geir Dullerud. Policy learning of MDPs with mixed continuous/discrete variables: A case study on model-free control of Markovian jump systems. *arXiv preprint arXiv:2006.03116*, 2020.

[28] Ilyas Fatkhullin and Boris Polyak. Optimizing static linear feedback: Gradient method. *arXiv preprint arXiv:2004.09875*, 2020.

[29] Yujie Tang, Yang Zheng, and Na Li. Analysis of the optimization landscape of linear quadratic Gaussian (LQG) control. In *Learning for Dynamics and Control*, pages 599–610. PMLR, 2021.

[30] Asaf B. Cassel and Tomer Koren. Online policy gradient for model free learning of linear quadratic regulators with $\sqrt{T}$ regret. In *International Conference on Machine Learning*, pages 1304–1313. PMLR, 2021.

[31] Mohamad K. S. Faradonbeh, Ambuj Tewari, and George Michailidis. Finite time analysis of optimal adaptive policies for linear-quadratic systems. *arXiv preprint arXiv:1711.07230*, 2017.

[32] Yi Ouyang, Mukul Gagrani, and Rahul Jain. Learning-based control of unknown linear systems with Thompson sampling. *arXiv preprint arXiv:1709.04047*, 2017.

[33] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 4188–4197, 2018.

[34] Alon Cohen, Tomer Koren, and Yishay Mansour. Learning linear-quadratic regulators efficiently with only $\sqrt{T}$ regret. *arXiv preprint arXiv:1902.06223*, 2019.

[35] Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalent control of LQR is efficient. *arXiv preprint arXiv:1902.07826*, 2019.

[36] Max Simchowitz and Dylan J. Foster. Naive exploration is optimal for online LQR. *arXiv preprint arXiv:2001.09576*, 2020.

[37] Max Simchowitz, Karan Singh, and Elad Hazan. Improper learning for non-stochastic control. *arXiv preprint arXiv:2001.09254*, 2020.

[38] Yang Zheng, Luca Furieri, Maryam Kamgarpour, and Na Li. Sample complexity of linear quadratic Gaussian (LQG) control for output feedback systems. In *Learning for Dynamics and Control*, pages 559–570. PMLR, 2021.

[39] Orestis Plevrakis and Elad Hazan. Geometric exploration for online control. *Advances in Neural Information Processing Systems*, 33:7637–7647, 2020.

[40] Tze Leung Lai. Asymptotically efficient adaptive control in stochastic regression models. *Advances in Applied Mathematics*, 7(1):23–45, 1986.

[41] Han-Fu Chen and Ji-Feng Zhang. Convergence rates in stochastic adaptive tracking. *International Journal of Control*, 49(6):1915–1935, 1989.

[42] Tze Leung Lai and Zhiliang Ying. Parallel recursive algorithms in asymptotically efficient adaptive control of linear stochastic systems. *SIAM Journal on Control and Optimization*, 29(5):1091–1127, 1991.

[43] Morteza Ibrahimi, Adel Javanmard, and Benjamin Van Roy. Efficient reinforcement learning for high dimensional linear quadratic systems. *arXiv preprint arXiv:1303.5984*, 2013.

[44] Mohamad K. S. Faradonbeh, Ambuj Tewari, and George Michailidis. Finite-time adaptive stabilization of linear systems. *IEEE Transactions on Automatic Control*, 64(8):3498–3505, 2019.

[45] Shahriar Talebi, Siavash Alemzadeh, Niyousha Rahim, and Mehran Mesbahi. Online regulation of unstable linear systems from a single trajectory. In *Proceedings of the 59th IEEE Conference on Decision and Control (CDC)*, pages 4784–4789, 2020.

[46] Yang Zheng and Na Li. Non-asymptotic identification of linear dynamical systems using multiple trajectories. *IEEE Control Systems Letters*, 5(5):1693–1698, 2020.

[47] Lennart Ljung. System identification. *Wiley Encyclopedia of Electrical and Electronics Engineering*, pages 1–19, 1999.

[48] Max Simchowitz, Horia Mania, Stephen Tu, Michael I. Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. *arXiv preprint arXiv:1802.08334*, 2018.

[49] Samet Oymak and Necmiye Ozay. Non-asymptotic identification of LTI systems from a single trajectory. In *2019 American Control Conference (ACC)*, pages 5655–5661. IEEE, 2019.

[50] Tuhin Sarkar, Alexander Rakhlin, and Munther A. Dahleh. Finite-time system identification for partially observed LTI systems of unknown order. *arXiv preprint arXiv:1902.01848*, 2019.

[51] Salar Fattahi. Learning partially observed linear dynamical systems from logarithmic number of samples. In *Learning for Dynamics and Control*, pages 60–72. PMLR, 2021.

[52] Han Wang and James Anderson. Large-scale system identification using a randomized svd. *arXiv preprint arXiv:2109.02703*, 2021.

[53] Yu Xing, Benjamin Gravell, Xingkang He, Karl Henrik Johansson, and Tyler Summers. Identification of linear systems with multiplicative noise from multiple trajectory data. *arXiv preprint arXiv:2106.16078*, 2021.

[54] Yuji Nakatsukasa. Off-diagonal perturbation, first-order approximation and quadratic residual bounds for matrix eigenvalue problems. In *Eigenvalue Problems: Algorithms, Software and Applications in Petascale Computing (EPASA)*, Lecture Notes in Computer Science, pages 233–249. Springer, 2015.

[55] E. A. Rawashdeh. A simple method for finding the inverse matrix of Vandermonde matrix. *Matematicki Vesnik: MV19303*, 2019.

[56] F. L. Bauer and C. T. Fike. Norms and exclusion theorems. *Numerische Mathematik*, 2:137–141, 1960.

[57] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 2nd edition, 2013.

[58] Ioannis Chatzigeorgiou. Bounds on the Lambert function and their application to the outage analysis of user cooperation. *IEEE Communications Letters*, 17(8):1505—1508, 2013.

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]
   (b) Did you describe the limitations of your work? [Yes] See the "future work" part of the conclusions in Section 6.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] We don't see any potential societal impacts in such theoretical results.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

(2) If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 4.
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix D, Appendix E, Appendix F and Appendix G.

(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]

Appendix

A Decomposition of the State Space

A.1 The $E_\mathrm{u} \oplus E_\mathrm{s}$-decomposition

It is evident that the following two subspaces of $\mathbb{R}^n$ are invariant with respect to A, namely

$$E_\mathrm{u} := \bigoplus_{i\le k} E_i,\ E_\mathrm{s} := \bigoplus_{i>k} E_i$$

which we refer to as the *unstable subspace* and the *stable subspace* of A, respectively. Since the eigenspaces E_i sum to the whole $\mathbb{R}^n$ space, one natural decomposition is $\mathbb{R}^n = E_\mathrm{u} \oplus E_\mathrm{s}$; accordingly, each state can be uniquely decomposed as $x = x_\mathrm{u} + x_\mathrm{s}$, where $x_\mathrm{u} \in E_\mathrm{u}$ is called the *unstable component*, and $x_\mathrm{s} \in E_\mathrm{s}$ is called the *stable component*.

We also decompose A based on the $E_\mathrm{u} \oplus E_\mathrm{s}$-decomposition. Suppose E_u and E_s are represented by their *orthonormal* bases $Q_1 \in \mathbb{R}^{n\times k}$ and $Q_2 \in \mathbb{R}^{n\times(n-k)}$, respectively, namely

$$E_\mathrm{u} = \mathrm{col}(Q_1),\ E_\mathrm{s} = \mathrm{col}(Q_2).$$

Let $Q = [Q_1\ Q_2]$ (which is invertible as long as A is diagonalizable), and let $R = [R_1^\top\ R_2^\top]^\top := Q^{-1}$. Further, let $\mathit{\Pi}_\mathrm{u} := Q_1R_1$ and $\mathit{\Pi}_\mathrm{s} = Q_2R_2$ be the *oblique* projectors onto E_u and E_s (along the other subspace), respectively. Since E_u and E_s are both invariant with regard to A, we know there exists $N_1 \in \mathbb{R}^{k\times k}$, $N_2 \in \mathbb{R}^{(n-k)\times(n-k)}$, such that

$$AQ = Q\begin{bmatrix} N_1 & \\ & N_2 \end{bmatrix} \Leftrightarrow N := \begin{bmatrix} N_1 & \\ & N_2 \end{bmatrix} = RAQ.$$

Let $z = [z_1^\top\ z_2^\top]^\top$ be the coordinate representation of x in the basis Q (i.e., $x = Qz$). The system dynamics in z-coordinates can be expressed as

$$\begin{bmatrix} z_{1,t+1} \\ z_{2,t+1} \end{bmatrix} = RAQ\begin{bmatrix} z_{1,t} \\ z_{2,t} \end{bmatrix} + RBu_t = \begin{bmatrix} N_1 & \\ & N_2 \end{bmatrix}\begin{bmatrix} z_{1,t} \\ z_{2,t} \end{bmatrix} + \begin{bmatrix} R_1B \\ R_2B \end{bmatrix} u_t.$$

The major advantage of this decomposition is that the dynamical matrix in z-coordinate is block diagonal, so it would be simpler to study the behavior of the open-loop system.

A.2 Geometric Interpretation: Principle Angles

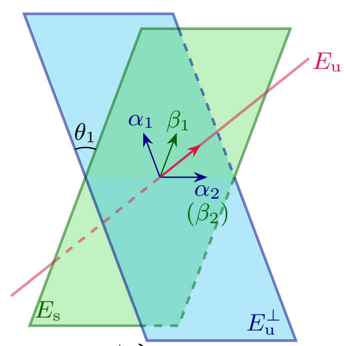

Before going any further, we emphasize that Definition 3.1 is well-defined by itself, since singular values are preserved under orthonormal transformations.

It might seem unintuitive to interpret $\sigma_{\min}(P_2^\top Q_2)$ in Definition 3.1 as a measure of "closeness". However, this is closely related to the *principle angles* between subspaces that generalize the standard angle measures in lower dimensional cases. More specifically, we can recursively define the i^{th} principle angle θ_i $(i = 1, \cdots, n-k)$ as

$$\theta_i := \min\left\{\arccos\left(\frac{\langle x, y\rangle}{\|x\|\|y\|}\right) \;\middle|\; \begin{array}{l} x \in E_\mathrm{u}^\perp,\ x \perp \mathrm{span}(x_1, \cdots, x_{i-1}); \\ y \in E_\mathrm{s},\ y \perp \mathrm{span}(y_1, \cdots, y_{i-1}). \end{array}\right\} =: \angle(x_i, y_i), \tag{7}$$

where x_i and y_i $(i = 1, \cdots, n-k)$are referred to as the i^{th} principle vectors accordingly. Meanwhile, let $P_2^\top Q_2 = U\mathit{\Sigma}V^\top$ be the singular value decomposition (SVD), where $\mathit{\Sigma} = \mathrm{diag}(\sigma_1, \cdots, \sigma_{n-k})$ and $\sigma_1 \ge \cdots \ge \sigma_{n-k}$. Then by an equivalent recursive characterization of singular values, we have

$$\sigma_i = \max_{\substack{\|x\|=\|y\|=1 \\ \forall j<i:\ x\perp x_j,\ y\perp y_j}} x^\top P_2^\top Q_2 y =: \bar{x}_i^\top P_2^\top Q_2 \bar{y}_i.$$

Since P_2 and Q_2 are orthonormal, $\bar{x}_i$ and $\bar{y}_i$ can be regarded as coordinate representations of $x_i = P_2\bar{x}_i$ and $y_i = Q_2\bar{y}_i$, and it can be easily verified that x_i and y_i defined in this way are exactly the minimizers in (7). Hence we conclude that $\sigma_i = \cos\theta_i$. Therefore, $E_\mathrm{u}^\perp$ and E_s are ξ-close if and only if the all principle angles between $E_\mathrm{u}^\perp$ and E_s lie in the interval $[0, \arccos(1-\xi)]$; the above argument also shows that we can find orthonormal bases for $E_\mathrm{u}^\perp$ and E_s so that the angles formed by corresponding vectors are exactly the principle angles.

A.3 Characterization of ξ-close Subspaces

It is naturally expected that the geometric interpretation should inspire more relationships among $P_1 = Q_1, P_2, Q_2, R_1, R_2$ and N_2. We would like to emphasize that P_1, P_2 and Q_1 are not confined to bases consisting of eigenvectors (since they are even not necessarily orthonormal). Meanwhile, since they are only used in the stability guarantee proof, we are granted the freedom to select any orthonormal bases. For simplicity, we will stick to the convention that $P_1 = Q_1$ (and thus $M_1 = N_1$). Further, in Lemma A.1, such freedom is utilized to establish fundamental relationships between the bases in the above two decompositions. The results are concluded as follows.

Lemma A.1. *Suppose $E_{\mathrm{u}}^{\perp}$ and E_{s} are ξ-close. Then we shall select P_2 and Q_2 such that*

(1) $\sigma_{\min}(P_2^\top Q_2) \geq 1-\xi$, $\|P_1^\top Q_2\| \leq \sqrt{2\xi}$, $\|P_2 - Q_2\| \leq \sqrt{2\xi}$.

(2) $\|R_2\| \leq \frac{1}{1-\xi}$, $\|N_2\| \leq \frac{1}{1-\xi}\|A\|$.

(3) $\|P_1^\top - R_1\| \leq \frac{\sqrt{2\xi}}{1-\xi}$, $\|R_1\| \leq \frac{\sqrt{2\xi}}{1-\xi} + 1$.

(4) $\|\mathit{\Delta}\| \leq \frac{2-\xi}{1-\xi}\sqrt{2\xi}\|A\|$.

Proof. (1) Following the above interpretation, take arbitrary orthonormal bases $\bar{P}_2$ and $\bar{Q}_2$ of $E_{\mathrm{u}}^{\perp}$ and E_{s}, respectively, and let $\bar{P}_2^\top \bar{Q}_2 = U\mathit{\Sigma}V^\top$ be the SVD, which translates to

$$(\bar{P}_2U)^\top(\bar{Q}_2V) = \mathit{\Sigma} =: \operatorname{diag}(\sigma_1, \cdots, \sigma_{n-k}).$$

Since U and V are orthonormal matrices, the columns of $\bar{P}_2U$ and $\bar{Q}_2V$ also form orthonormal bases of $E_{\mathrm{u}}^{\perp}$ and E_{s}, respectively. Then ξ-closeness basically says that there exist a basis $\{\alpha_1, \cdots, \alpha_{n-k}\}$ for $E_{\mathrm{u}}^{\perp}$, and a basis $\{\beta_1, \cdots, \beta_{n-k}\}$ for E_{s} (both are assumed to be orthonormal), such that

$$\langle \alpha_i, \beta_j \rangle = \delta_{ij}\sigma_i = \begin{cases} \sigma_i \geq 1-\xi & \text{for any } i = j \\ 0 & \text{for any } i \neq j \end{cases},$$

and we also have $\mathit{\Pi}_2\beta_i = \sigma_i\alpha_i$ and $\mathit{\Pi}_1\alpha_i = \sigma_i\beta_i$ (recall that $\mathit{\Pi}_1, \mathit{\Pi}_2$ are orthogonal projectors onto subspaces $E_{\mathrm{u}}, E_{\mathrm{u}}^{\perp}$, respectively). Therefore, without loss of generality, we shall always select $P_2 = [\alpha_1 \ \cdots \ \alpha_{n-k}]$ and $Q_2 = [\beta_1 \ \cdots \ \beta_{n-k}]$, such that $P_2^\top Q_2 = \operatorname{diag}(\sigma_1, \cdots, \sigma_{n-k})$, and

$$\sigma_{\min}(P_2^\top Q_2) = \min_i |\sigma_i| \geq 1-\xi.$$

Equivalently speaking, for any $\beta = Q_2\eta \in E_{\mathrm{s}}$, we have (note that $\|\eta\| = \|\beta\|$)

$$\|P_2^\top \beta\| = \|P_2^\top Q_2\eta\| \geq \sigma_{\min}(P_2^\top Q_2)\|\eta\| \geq (1-\xi)\|\beta\|,$$

and consequently,

$$\|P_1^\top Q_2\eta\| = \|P_1^\top \beta\| = \sqrt{\|\beta\|^2 - \|P_2^\top \beta\|^2} \leq \sqrt{2\xi}\|\beta\| = \sqrt{2\xi}\|\eta\|,$$

which further shows $\|P_1^\top Q_2\| \leq \sqrt{2\xi}$. To bound $\|P_2 - Q_2\|$, by definition we have

$$\begin{aligned}
\|P_2 - Q_2\| &= \max_{\|\eta\|=1} \|(P_2 - Q_2)\eta\| = \max_{\|\eta\|=1} \left\| \sum_i \eta_i(\alpha_i - \beta_i) \right\| \\
&= \max_{\|\eta\|=1} \sqrt{\sum_{i,j} \eta_i\eta_j(\alpha_i - \beta_i)^\top(\alpha_j - \beta_j)} \\
&= \max_{\|\eta\|=1} \sqrt{\sum_i 2(1-\mu_i)\eta_i^2} \\
&\leq \max_{\|\eta\|=1} \sqrt{2\xi \sum_i \eta_i^2} = \sqrt{2\xi}.
\end{aligned}$$

Here $\eta = [\eta_1, \cdots, \eta_{n-k}]$ is an arbitrary vector in $\mathbb{R}^{n-k}$.

(2) By definition, $I = QR = Q_1R_1 + Q_2R_2$. Also recall that $P_1 = Q_1$, so we have $P_1^\top Q_1 = I$ and $P_2^\top Q_1 = \mathbf{0}$. Then by left-multiplying $P_2^\top$ to the equation, we have

$$P_2^\top = P_2^\top Q_1R_1 + P_2^\top Q_2R_2 = P_2^\top Q_2R_2,$$

which further shows

$$\|R_2\| = \|(P_2^\top Q_2)^{-1} P_2^\top\| \le \|(P_2^\top Q_2)^{-1}\| = \frac{1}{\sigma_{\min}(P_2^\top Q_2)} \le \frac{1}{1-\xi}.$$

Therefore, since $N_2 = R_2 A Q_2$, we have

$$\|N_2\| = \|R_2 A Q_2\| \le \|R_2\|\|A\|\|Q_2\| \le \frac{1}{1-\xi}\|A\|.$$

(3) Similarly, by left-multiplying $P_1^\top$ to the equation, we have

$$P_1^\top = P_1^\top Q_1 R_1 + P_1^\top Q_2 R_2 = R_1 + P_1^\top Q_2 R_2,$$

which further shows

$$\|P_1^\top - R_1\| = \|P_1^\top Q_2 R_2\| \le \|P_1^\top Q_2\|\|R_2\| \le \frac{\sqrt{2\xi}}{1-\xi},$$

and therefore $\|R_1\| \le \|P_1^\top - R_1\| + \|P_1^\top\| = 1 + \frac{\sqrt{2\xi}}{1-\xi}$.

(4) A combination of the above results gives

$$\begin{aligned}
\|\varDelta\| &= \|P_1^\top A P_2\| = \|P_1^\top A P_2 - R_1 A Q_2\| \\
&\le \|P_1^\top A (P_2 - Q_2)\| + \|(P_1^\top - R_1) A Q_2\| \\
&\le \|P_1^\top\|\|A\|\|P_2 - Q_2\| + \|P_1^\top - R_1\|\|A\|\|Q_2\| \\
&\le \|A\|\sqrt{2\xi} + \frac{\sqrt{2\xi}}{1-\xi}\|A\| = \frac{2-\xi}{1-\xi}\sqrt{2\xi}\|A\|.
\end{aligned}$$

This completes the proof. □

B Solution to the Least Squares Problem in Stage 2

Lemma B.1 gives the explicit form for the solution to the least squares problem (see Algorithm 1).

Lemma B.1. *Given $D := [x_{t_0+1} \ \cdots \ x_{t_0+k}]$ and $\hat{P}_1 \hat{P}_1^\top = \hat{\varPi}_1 = D(D^\top D)^{-1} D^\top$, the solution*

$$\hat{M}_1 = \arg\min_{M_1} \sum_{t=t_0+1}^{t_0+k} \|\hat{P}_1^\top x_{t+1} - M_1 \hat{P}_1^\top x_t\|^2$$

is uniquely given by $\hat{M}_1 = \hat{P}_1^\top A \hat{P}_1$.

Proof. Here we assume by default that the summation over t sums from t_0+1 to t_0+k. Since M_1 is a stationary point of $\mathcal{L}$, for any $\varDelta$ in the neighbourhood of O, we have

$$\begin{aligned}
0 \le \mathcal{L}(M_1 + \varDelta) - \mathcal{L}(M_1) &= \sum_t \|\hat{y}_{1,t+1} - M_1 \hat{y}_{1,t} - \varDelta \hat{y}_{1,t}\|^2 - \sum_t \|\hat{y}_{1,t+1} - M_1 \hat{y}_{1,t}\|^2 \\
&= \sum_t \langle \varDelta \hat{y}_{1,t}, \hat{y}_{1,t+1} - M_1 \hat{y}_{1,t} \rangle + O(\|\varDelta\|^2) \\
&= \sum_t \operatorname{tr}\left(\hat{y}_{1,t}^\top \varDelta^\top (\hat{y}_{1,t+1} - A \hat{y}_{1,t})\right) + O(\|\varDelta\|^2) \\
&= \sum_t \operatorname{tr}\left(\varDelta^\top (\hat{y}_{1,t+1} - M_1 \hat{y}_{1,t}) \hat{y}_{1,t}^\top\right) + O(\|\varDelta\|^2) \\
&= \operatorname{tr}\left(\varDelta^\top \sum_t (\hat{y}_{1,t+1} - M_1 \hat{y}_{1,t}) \hat{y}_{1,t}^\top\right) + O(\|\varDelta\|^2).
\end{aligned}$$

Since it always holds for any $\varDelta$, we must have

$$\sum_t (\hat{y}_{1,t+1} - M_1 \hat{y}_{1,t}) \hat{y}_{1,t}^\top \ \Leftrightarrow \ M_1 \sum_t \hat{y}_{1,t} \hat{y}_{1,t}^\top = \sum_t \hat{y}_{1,t+1} \hat{y}_{1,t}^\top.$$

Plugging in $\hat{y}_{1,t} = \hat{P}_1^\top x_t$ and $\hat{y}_{1,t+1} = \hat{P}_1^\top A x_t$, we further have

$$M_1 \hat{P}_1^\top X \hat{P}_1 = M_1 \sum_t \hat{P}_1^\top x_t x_t^\top \hat{P}_1 = \sum_t \hat{P}_1^\top A x_t x_t^\top \hat{P}_1 = \hat{P}_1^\top A X \hat{P}_1,$$

where $X := \sum_t x_t x_t^\top = DD^\top$. Since the columns of $\hat{P}_1$ form an orthonormal basis of $\hat{E}_{\mathrm{u}}$, for any $x \in \hat{E}_{\mathrm{u}}$, $\hat{P}_1^\top x$ is the coordinate of x under that basis. The columns of D are linearly independent, so the columns of $\hat{P}_1^\top D$ are also linearly independent, which further shows

$$\mathrm{rank}(\hat{P}_1^\top X \hat{P}_1) = \mathrm{rank}\left((\hat{P}_1^\top D)(\hat{P}_1^\top D)^\top\right) = \mathrm{rank}(\hat{P}_1^\top D) = k.$$

Therefore, $\hat{P}_1^\top X \hat{P}_1$ is invertible, and M_1 is explicitly given by

$$M_1 = (\hat{P}_1^\top A X \hat{P}_1)(\hat{P}_1^\top X \hat{P}_1)^{-1}.$$

Note that $\hat{\Pi}_1 = \hat{P}_1 \hat{P}_1^\top$ is the projector onto subspace $\mathrm{col}(D)$, we must have

$$\hat{P}_1 \hat{P}_1^\top X = (\hat{\Pi}_1 D) D^\top = DD^\top = X,$$

which yields

$$M_1 = (\hat{P}_1^\top A (\hat{P}_1 \hat{P}_1^\top X) \hat{P}_1)(\hat{P}_1^\top X \hat{P}_1)^{-1} = (\hat{P}_1^\top A \hat{P}_1)(\hat{P}_1^\top X \hat{P}_1)(\hat{P}_1^\top X \hat{P}_1)^{-1} = \hat{P}_1^\top A \hat{P}_1.$$

This completes the proof of Lemma B.1. □

It might help understanding to note that, when $\hat{P}_1 = P_1$, for any $x_t, x_{t+1} \in E_{\mathrm{u}}$ we have

$$P_1^\top A x_t = y_{t+1} = M_1 y_t = M_1 P_1^\top x_t,$$

which requires $P_1^\top A = M_1 P_1^\top$, or equivalently, $M_1 = P_1^\top A P_1$ (recall $P_1^\top P_1 = I$).

C Transformation of B with Arbitrary Columns

Throughout the remainder of this paper, we regard B as an n-by-k matrix (i.e., $m = k$). In this section, we show that other cases can be handled in a similar way via proper transformations. Intuitively, since the system is assumed to be strongly controllable, we can "pack" every d consecutive steps together as a new state to obtain a system that is "fully-actuated" in the unstable subspace (i.e., the unstable subspace is controlled by at least k linearly independent inputs).

More specifically, for any integer d, the transformation is given by

$$\tilde{x}_t = \begin{bmatrix} x_{td} \\ x_{td+1} \\ \vdots \\ x_{(t+1)d-1} \end{bmatrix}, \quad \tilde{u}_t = \begin{bmatrix} u_{td-1} \\ u_{td} \\ \vdots \\ u_{(t+1)d-2} \end{bmatrix},$$

$$\tilde{A} = \begin{bmatrix} \mathbf{0} & & A \\ & \ddots & \vdots \\ & \mathbf{0} & A^{d-1} \\ & & A^d \end{bmatrix}, \; \tilde{B} = \begin{bmatrix} B & & & \\ AB & B & & \\ \vdots & \vdots & \ddots & \\ A^{d-1}B & A^{d-2}B & \cdots & B \end{bmatrix},$$

and the dynamics of the transformed system is given by

$$\tilde{x}_{t+1} = \tilde{A}\tilde{x}_t + \tilde{B}\tilde{u}_t.$$

If $\tilde{B}$ contains more than k linearly independent columns, we shall simply select k of them, and pad 0 to the unselected entries of the control input $\tilde{u}_t$.

Assumptions in the transformed system. It is evident that $|\tilde{\lambda}_i| = |\lambda_i|^d$ ($i = 1, \cdots, n$) (denote by $\tilde{\lambda}_i$ the i^{th} eigenvector of $\tilde{A}$). Therefore, the instability index of $\tilde{A}$ is still k, and the transformed system still satisfies Assumption 4.1, as well as $|\lambda_1|^2 |\lambda_{k+1}| < |\lambda_k|$, if the original system do.

The following proposition shows that Assumption 4.3 (i.e., c-effective controllability within unstable subspace) of the original system implies Assumption 4.3′ (i.e., (ν, σ)-strong controllability) of the transformed system, so that the transformation preserves all assumptions. To facilitate the proof, we first introduce a lemma on the smallest singular value of block matrices.

Lemma C.1. *For any row-block matrix $C := [A \;\; B]$ such that $A \in \mathbb{R}^{r\times a}$, $B \in \mathbb{R}^{r\times b}$ with $a+b \le r$, we have $\sigma_{\min}(A) \ge \sigma_{\min}(C)$.*

Proof. Since C is "tall", its smallest (non-trivial) singular value can be equivalently given by

$$\sigma_{\min}(C) = \min_{\substack{v\in\mathbb{R}^{a+b}\\ \|v\|=1}} \|Cv\|.$$

Similar results hold for A and B. Consequently,

$$\sigma_{\min}(C) = \min_{\substack{v\in\mathbb{R}^{a+b}\\ \|v\|=1}} \|Cv\| = \min_{\substack{v_1\in\mathbb{R}^{a}, v_2\in\mathbb{R}^{b}\\ \|v_1\|^2+\|v_2\|^2=1}} \|Av_1 + Bv_2\| \le \|Av_1^*\| = \sigma_{\min}(A),$$

where $v_1^* \in \arg\min_{v_1\in\mathbb{R}^a, \|v_1\|=1} \|Av_1\|$. This completes the proof. $\square$

Proposition C.1. *If the original system is (ν,σ)-strongly controllable, then after the transformation stated above, the transformed system satisfies Assumption 4.3′ with $d \le k$.*

Proof. It can be directly verified that, in the transformed system, matrix R in the $E_{\mathrm{u}} \oplus E_{\mathrm{s}}$ decomposition becomes (possible non-zero blocks are marked by $*$)

$$\tilde{R} = \begin{bmatrix} * & & & \\ & \ddots & & \\ & & * & \\ & & & R \end{bmatrix},$$

and therefore $\tilde{R}_1 = [\mathbf{0} \;\; \cdots \;\; \mathbf{0} \;\; R_1]$. Since $R_1 A = N_1 R_1$, we further have

$$\tilde{R}_1\tilde{B} = [R_1A^{d-1}B \;\; R_1A^{d-2}B \;\; \cdots \;\; R_1B] = [N_1^{d-1}R_1B \;\; N_1^{d-2}R_1B \;\; \cdots \;\; R_1B],$$

which is exactly the d-step controllability matrix of the k-dimensional LTI system (N_1, R_1B). Therefore, it only suffices to establish strong controllability of (N_1, R_1B).

It is evident that similar transform of dynamical matrices preserve strong controllability. Therefore, the system $(N, RB) = (Q^{-1}AQ, Q^{-1}B)$ is also (ν,σ)-strongly controllable by assumption. By definition, this indicates lower-bounded smallest singular value of its ν-step controllability matrix

$$C_\nu = [N^{\nu-1}RB \;\; N^{\nu-2}RB \;\; \cdots \;\; RB] = \begin{bmatrix} N_1^{\nu-1}R_1B & N_1^{\nu-2}R_1B & \cdots & R_1B \\ N_2^{\nu-1}R_2B & N_2^{\nu-2}R_2B & \cdots & R_2B \end{bmatrix}.$$

Consequently, the top k rows of C_ν, which is exactly the ν-step controllability matrix of (N_1, R_1B), also has smallest singular value lower bounded by σ due to Lemma C.1 (note that singular values are preserved by transpose). Therefore, (N_1, R_1B) is also (ν_1,σ)-strongly controllable. Since the sub-system is only k-dimensional, we can always select $\nu_1 \le k$. Taking $d = \nu_1$, and the proof is finished (note that the constant c in Assumption 4.3′ can be simply taken as $\tilde{c} := \sigma/\|\tilde{B}\|$). $\square$

Performance of the transformed system. We also have to show that the transformation only mildly degrades the performance bounds established in Theorems 4.1 and 4.2. Recall that the instance-specific constants involved here are $|\tilde{\lambda}_i|$, $\|\tilde{A}\|$, $\|\tilde{B}\|$, $\tilde{c}$, ξ, $\kappa_{\mathrm{e}}(A)$, and $\xi_\varepsilon(\cdot)$.

It is clear that $|\tilde{\lambda}_i| = |\lambda_i|^d$, and $\|\tilde{A}\|$, $\|\tilde{B}\|$ are upper bounded by

$$\|\tilde{A}\| \le \sqrt{\sum_{i=1}^{d} \|A^i\|^2} = \|A^d\|O(d), \quad \|\tilde{B}\| \le \|B\|\sqrt{\sum_{i=1}^{d}(d-i)\|A^i\|^2} = \|A^d\|\|B\|O(d).$$

Also note that $1/\tilde{c} = \|\tilde{B}\|/\sigma$. Therefore, $|\tilde{\lambda}_i|$, $\|\tilde{A}\|$, $\|\tilde{B}\|$ and $\tilde{c}$ are only lifted to a power of d after transformation. Meanwhile, since $\tilde{A}$ "inherits" all eigenvectors from A (in the last n dimensions), while the eigenvectors corresponding to the first $(d-1)n$ zero eigenvalues can be padded to be orthonormal, ξ and $\kappa_{\mathrm{e}}(A)$ remain unchanged. Finally, since $\tilde{A}^t = \mathrm{diag}(\mathbf{0}, \cdots, \mathbf{0}, A^{td})$ for any $t > 1$, the Gelfand constants $\xi_\varepsilon(\cdot)$ with respect to subspaces of $\mathrm{col}(A)$ also remain unchanged.

As a conclusion, all constants are either unchanged or only lifted to a power of $d \le k \ll n$. Since the constants always appear in logarithmic additive terms (see Appendices F and G for details), the performance bounds will remain the same in terms of orders with regard to k and n.

D Proof of Lemma 5.3

Lemma 5.3 is actually a direct corollary of the following lemma, for which we first need to define $\mathrm{gap}_i(A)$, the *(bipartite) spectral gap around λ_i with respect to A*, namely

$$\mathrm{gap}_i(A) := \begin{cases} \min_{\lambda_j \in \lambda(A_2)} |\lambda_i - \lambda_j| & \lambda_i \in \lambda(A_1) \\ \min_{\lambda_j \in \lambda(A_1)} |\lambda_i - \lambda_j| & \lambda_i \in \lambda(A_2) \end{cases},$$

where $\lambda(A)$ denotes the spectrum of A.

Lemma D.1. *For 2-by-2 block matrices A and E in the form*

$$A = \begin{bmatrix} A_1 & \mathbf{0} \\ \mathbf{0} & A_2 \end{bmatrix},\ E = \begin{bmatrix} \mathbf{0} & E_{12} \\ E_{21} & \mathbf{0} \end{bmatrix},$$

we have

$$|\lambda_i(A+E) - \lambda_i(A)| \le \frac{\kappa_{\mathrm{e}}(A)\kappa_{\mathrm{e}}(A+E)}{\mathrm{gap}_i(A)} \|E_{12}\|\|E_{21}\|.$$

Here $\kappa_{\mathrm{e}}(A)$ is the condition number of the matrix consisting of A's eigenvectors as columns.

Proof. The proof of the lemma can be found in existing literature like [54]. □

Proof of Lemma 5.3. Lemma D.1 basically guarantees that every eigenvalue of $A+E$ is within a distance of $O(\|E_{12}\|\|E_{21}\|)$ from some eigenvalue of A. Hence, by defining $\chi(A+E)$ as the maximum coefficient, namely

$$\chi(A+E) := \frac{\kappa_{\mathrm{e}}(A)\kappa_{\mathrm{e}}(A+E)}{\min_i\{\mathrm{gap}_i(A)\}},$$

we shall guarantee $|\rho(A+E) - \rho(A)| \le \chi(A+E)\|E_{12}\|\|E_{21}\|$. □

E Proof of Theorem 5.1 and its Corollary

The main idea of this proof is to diagonalize A and write the open-loop system dynamics using the basis formed by the eigenvectors of A. Then, we provide an explicit expression for $\hat{\Pi}_1$ and Π_1, based on which we can bound the error. To further derive a bound for $\|\hat{P}_1 - P_1\|$, one only needs to notice that norms are preserved under orthonormal coordinate transformations, so it only suffices to find a specific pair of bases of $E_{\mathrm{u}}^{\perp}$ and E_{s} that are close to each other — and the pair of bases formed by principle vectors (see Appendix A) is exactly what we want. This leads to Corollary 5.2 that is repeatedly used in subsequent proofs.

Without loss of generality, we shall write all matrices in the basis formed by unit eigenvectors $\{w_1, \cdots, w_n\}$ of A. Otherwise, let $W = [w_1 \ \cdots \ w_n]$, and perform change-of-coordinate by setting $\tilde{D} := W^{-1}DW$, $\tilde{\Pi}_1 := W^{-1}\Pi_1 W$, which further gives

$$\tilde{\hat{\Pi}}_1 = \tilde{D}(\tilde{D}^\top \tilde{D})^{-1}\tilde{D}^\top = (W^{-1}DW)(W^{-1}D^\top DW)^{-1}(W^{-1}D^\top W) = W^{-1}\hat{\Pi}_1 W.$$

Note that $\|W^{-1}\hat{\Pi}_1 W - W^{-1}\Pi_1 W\| \le \|W\|\|W^{-1}\|\|\hat{\Pi}_1 - \Pi_1\|$, where the upper bound is only magnified by a constant factor of $\kappa_{\mathrm{e}}(A) = \|W\|\|W^{-1}\|$ that is completely determined by A. Therefore, it is largely equivalent to consider $(\tilde{D}, \tilde{\Pi}_1, \tilde{\hat{\Pi}}_1)$ instead of $(D, \Pi_1, \hat{\Pi}_1)$.

Note that the matrix $D = [x_{t_0+1} \ \cdots \ x_{t_0+k}]$ can be written as

$$D = \begin{bmatrix} d_1 & \lambda_1 d_1 & \cdots & \lambda_1^{k-1} d_1 \\ d_2 & \lambda_2 d_2 & \cdots & \lambda_2^{k-1} d_2 \\ \vdots & \vdots & \ddots & \vdots \\ d_n & \lambda_n d_n & \cdots & \lambda_n^{k-1} d_n \end{bmatrix},$$

where $x_{t_0+1} =: [d_1, \cdots, d_n]^\top$. We first present a lemma characterizing some well-known properties of Vandermonde matrices that we need in the proof.

Lemma E.1. *Given a Vandermonde matrix in variables* $x_1, \cdots, x_n$ *of order* n

$$V := V_n(x_1, \cdots, x_n) = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ x_1 & x_2 & \cdots & x_n \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{n-1} & x_2^{n-1} & \cdots & x_n^{n-1} \end{bmatrix},$$

its determinant is given by

$$\det(V) = \sum_{\pi} (-1)^{\mathrm{sgn}(\pi)} x^0_{\pi(i_1)} x^1_{\pi(i_2)} \cdots x^{n-1}_{\pi(i_n)} = \prod_{j<\ell} (x_\ell - x_j), \tag{8}$$

and its (u, v)*-cofactor is given by*

$$\mathrm{cof}_{u,v}(V) = \begin{vmatrix} 1 & \cdots & 1 & 1 & \cdots & 1 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ x_1^{u-2} & \cdots & x_{v-1}^{u-2} & x_{v+1}^{u-2} & \cdots & x_n^{u-2} \\ x_1^{u} & \cdots & x_{v-1}^{u} & x_{v+1}^{u} & \cdots & x_n^{u} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ x_1^{n-1} & \cdots & x_{v-1}^{n-1} & x_{v+1}^{n-1} & \cdots & x_n^{n-1} \end{vmatrix} = \sigma_{u,v} \prod_{j<\ell\neq v} (x_\ell - x_j). \tag{9}$$

Here coefficients $\sigma_{u,v}$ *are given by* $\sigma_{u,v} := s_{n-u}(x_1, \cdots, x_{v-1}, x_{v+1}, \cdots, x_n)$*, where function* s_m *is defined by* $s_m(y_1, \cdots, y_n) := \sum_{i_1<\cdots<i_m} y_{i_1} \cdots y_{i_m}$.

Proof of Lemma E.1. The proof of (8) can be found in any standard linear algebra textbook, while the proof of (9) can be found in [55]. □

It is evident that the entries in D display a similar pattern as those of a Vandermonde matrix. Based on this observation, we shall further derive the explicit form of $\hat{\Pi}_1$ as in the next lemma.

Lemma E.2. *The projector* $\hat{\Pi}_1 = D(D^\top D)^{-1} D^\top$ *has explicit form*

$$(\hat{\Pi}_1)_{uv} = \frac{\sum_{\substack{i_2<\cdots<i_k \\ \forall j: i_j \neq u,v}} \alpha_{u,i_2,\cdots,i_k} \alpha_{v,i_2,\cdots,i_k}}{\sum_{i_1<\cdots<i_k} \alpha^2_{i_1,\cdots,i_k}},$$

where the summand $\alpha_{i_1,\cdots,i_k}$ *(with ordered subscript) is defined as*

$$\alpha_{i_1,\cdots,i_k} := \prod_j d_{i_j} \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j}).$$

Proof of Lemma E.2. We start by deriving the explicit form of $(D^\top D)^{-1}$. Note that the determinant (which is also the denominator in the lemma) is given by

$$\begin{aligned} \det(D^\top D) &= \sum_{i_1,\cdots,i_k} \begin{vmatrix} \lambda^0_{i_1} d^2_{i_1} & \lambda^1_{i_2} d^2_{i_2} & \cdots & \lambda^{k-1}_{i_k} d^2_{i_k} \\ \lambda^1_{i_1} d^2_{i_1} & \lambda^2_{i_2} d^2_{i_2} & \cdots & \lambda^{k}_{i_k} d^2_{i_k} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda^{k-1}_{i_1} d^2_{i_1} & \lambda^k_{i_2} d^2_{i_2} & \cdots & \lambda^{2k-2}_{i_k} d^2_{i_k} \end{vmatrix} \\ &= \sum_{i_1,\cdots,i_k} d^2_{i_1} \cdots d^2_{i_k} \lambda^0_{i_1} \lambda^1_{i_2} \cdots \lambda^{k-1}_{i_k} \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j}) \\ &= \sum_{i_1<\cdots<i_k} d^2_{i_1} \cdots d^2_{i_k} \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j}) \sum_{\pi} (-1)^{\mathrm{sgn}(\pi)} \lambda^0_{\pi(j_1)} \lambda^1_{\pi(j_2)} \cdots \lambda^{k-1}_{\pi(j_k)} \\ &= \sum_{i_1<\cdots<i_k} d^2_{i_1} \cdots d^2_{i_k} \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j})^2 \\ &= \sum_{i_1<\cdots<i_k} \alpha^2_{i_1,\cdots,i_k}, \end{aligned}$$

and the (u, v)-cofactor $\mathrm{cof}_{u,v}(D^\top D)$ is given by

$$\begin{aligned}
\mathrm{cof}_{u,v}(D^\top D) &= (-1)^{u+v} \sum_{i_1,\cdots,i_{k-1}} \begin{vmatrix} \lambda_{i_1}^0 d_{i_1}^2 & \cdots & \lambda_{i_{v-1}}^{v-2} d_{i_{v-1}}^2 & \lambda_{i_v}^{v} d_{i_v}^2 & \cdots & \lambda_{i_{k-1}}^{k-1} d_{i_{k-1}}^2 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \lambda_{i_1}^{u-2} d_{i_1}^2 & \cdots & \lambda_{i_{v-1}}^{u+v-4} d_{i_{v-1}}^2 & \lambda_{i_v}^{u+v-2} d_{i_v}^2 & \cdots & \lambda_{i_{k-1}}^{u+k-3} d_{i_{k-1}}^2 \\ \lambda_{i_1}^{u} d_{i_1}^2 & \cdots & \lambda_{i_{v-1}}^{u+v-2} d_{i_{v-1}}^2 & \lambda_{i_v}^{u+v} d_{i_v}^2 & \cdots & \lambda_{i_{k-1}}^{u+k-1} d_{i_{k-1}}^2 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \lambda_{i_1}^{k-1} d_{i_1}^2 & \cdots & \lambda_{i_{u+v-2}}^{k+v-3} d_{i_{v-1}}^2 & \lambda_{i_v}^{k+v-1} d_{i_v}^2 & \cdots & \lambda_{i_{k-1}}^{2k-2} d_{i_{k-1}}^2 \end{vmatrix} \\
&= (-1)^{u+v} \sum_{i_1,\cdots,i_{k-1}} d_{i_1}^2 \cdots d_{i_{k-1}}^2 \lambda_{i_1}^0 \cdots \lambda_{i_{v-1}}^{v-2} \lambda_{i_v}^{v} \cdots \lambda_{i_{k-1}}^{k-1} s_{k-u} \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j}) \\
&= (-1)^{u+v} \sum_{i_1<\cdots<i_{k-1}} s_{k-u} \cdot d_{i_1}^2 \cdots d_{i_{k-1}}^2 \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j}) \cdot \\
&\qquad\qquad \sum_{\pi} (-1)^{\mathrm{sgn}(\pi)} \lambda_{\pi(i_1)}^0 \cdots \lambda_{\pi(i_{v-1})}^{v-2} \lambda_{\pi(i_v)}^{v} \cdots \lambda_{\pi(i_{k-1})}^{k-1} \\
&= (-1)^{u+v} \sum_{i_1<\cdots<i_{k-1}} s_{k-u} s_{k-v} \cdot d_{i_1}^2 \cdots d_{i_{k-1}}^2 \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j})^2,
\end{aligned}$$

where $s_{k-u}(\lambda_{i_1}, \cdots, \lambda_{i_{k-1}})$ is abbreviated to s_{k-u}.

Note that symmetry of $D^\top D$ guarantees $\mathrm{cof}_{v,u}(D^\top D) = \mathrm{cof}_{u,v}(D^\top D)$, so we have

$$(D^\top D)^{-1}_{u,v} = \frac{\mathrm{cof}_{v,u}(D^\top D)}{\det(D^\top D)} = \frac{\mathrm{cof}_{u,v}(D^\top D)}{\det(D^\top D)}.$$

And eventually we shall derive that

$$\begin{aligned}
\hat{P}_{u,v} &= \sum_{p,q} D_{u,p} (D^\top D)^{-1}_{p,q} D^\top_{q,v} \\
&= \frac{1}{\det(D^\top D)} \sum_{p,q} D_{u,p} D_{v,q}\, \mathrm{cof}_{u,v}(D^\top D) \\
&= \frac{1}{\det(D^\top D)} \sum_{p,q} \lambda_u^{p-1} d_u \lambda_v^{q-1} d_v \cdot (-1)^{p+q} \sum_{i_1<\cdots<i_{k-1}} s_{k-p} s_{k-q} \cdot d_{i_1}^2 \cdots d_{i_{k-1}}^2 \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j})^2 \\
&= \frac{1}{\det(D^\top D)} \sum_{i_1<\cdots<i_{k-1}} d_u d_v d_{i_1}^2 \cdots d_{i_{k-1}}^2 \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j})^2 \sum_{p=1}^{k} (-1)^p \lambda_u^{p-1} s_{k-p} \sum_{q=1}^{k} (-1)^q \lambda_v^{q-1} s_{k-q} \\
&= \frac{1}{\det(D^\top D)} \sum_{i_1<\cdots<i_{k-1}} d_u d_{i_1} \cdots d_{i_{k-1}} \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j}) \prod_{\ell} (\lambda_{i_\ell} - \lambda_u) \cdot \\
&\qquad\qquad d_v d_{i_1} \cdots d_{i_{k-1}} \prod_{j<\ell} (\lambda_{i_\ell} - \lambda_{i_j}) \prod_{\ell} (\lambda_{i_\ell} - \lambda_v) \\
&= \frac{1}{\det(D^\top D)} \sum_{\substack{i_2<\cdots<i_k \\ \forall j: i_j \neq u,v}} \alpha_{u,i_2,\cdots,i_k} \alpha_{v,i_2,\cdots,i_k},
\end{aligned}$$

which is in exact the same form as stated in the lemma. □

Now we shall go back to the proof of the main result of this section.

Proof of Theorem 5.1. Recall that $d_i = \lambda_i^{t_0+1} x_{0,i}$. For the clarity of notations, let

$$\theta_{i_1,i_2,\cdots,i_k} := \frac{\alpha_{i_1,i_2,\cdots,i_k}}{\alpha_{1,2,\cdots,k}},$$

and it is evident that $|\theta_{i_1,i_2,\cdots,i_k}| = 1$ only if $(i_1, i_2, \cdots, i_k)$ is a permutation of $(1, 2, \cdots, k)$. For any other $(i_1, i_2, \cdots, i_k)$, by the definition in Lemma E.2 we have

$$|\theta_{i_1,i_2,\cdots,i_k}| \le c_{i_1,i_2,\cdots,i_k} \cdot r^{\sum_j \mathbf{1}_{i_j>k} t_0} \le c \cdot r^{t_0},$$

where $r = \frac{|\lambda_{k+1}|}{|\lambda_k|}$ and $c := \max_{i_1,\cdots,i_k} \{c_{i_1,i_2,\cdots,i_k}\}$. Therefore, since there are $\binom{n}{k}$ different k-tuples $(i_1, \cdots, i_k)$ such that $i_1 < \cdots < i_k$, we have

$$\sum_{i_1<\cdots<i_k} \theta^2_{i_1,\cdots,i_k} - \theta^2_{1,\cdots,k} < c\binom{n}{k} r^{2t_0}.$$

Now we can bound the entries in $\hat{\Pi}_1$. For any $\varepsilon > 0$, we shall select t_0 such that $c\binom{n}{k} r^{2t_0} < \frac{\varepsilon}{n^2}$, where the denominator is always bounded by

$$1 \le \sum_{i_1<\cdots<i_k} \theta^2_{i_1,\cdots,i_k} \le 1 + \frac{\varepsilon}{n^2}.$$

For the nominator, note that for each δ there are fewer entries with exponent δ in the nominator than in the denominator, so we can bound the denominator as

$$\left| \sum_{\substack{i_2<\cdots<i_k \\ \forall j: i_j \ne u,v}} \theta_{u,i_2,\cdots,i_k} \theta_{v,i_2,\cdots,i_k} \right| \le \begin{cases} c\binom{n}{k} r^{2t_0} + 1 & u = v \le k \\ c\binom{n}{k} r^{2t_0} & \text{otherwise} \end{cases}.$$

Therefore, when $u = v \le k$, we have $\sum_{\substack{i_2<\cdots<i_k \\ \forall j: i_j \ne u}} \theta^2_{u,i_2,\cdots,i_k} \ge 1$, which shows

$$\left. \begin{aligned} (\hat{\Pi}_1)_{uv} \ge \left(1 + \frac{\varepsilon}{n^2}\right)^{-1} \ge 1 - \frac{\varepsilon}{n^2} \\ (\hat{\Pi}_1)_{uv} \le 1 + \frac{\varepsilon}{n^2} \end{aligned} \right\} \Rightarrow \left| (\hat{\Pi}_1)_{uv} - (\Pi_1)_{uv} \right| \le \frac{\varepsilon}{n^2};$$

for all other cases, the nominator cannot sum over a permutation of $(1, \cdots, k)$, which gives

$$\left| (\hat{\Pi}_1)_{uv} - (\Pi_1)_{uv} \right| = \left| (\hat{\Pi}_1)_{uv} \right| \le \frac{\varepsilon}{n^2}.$$

Therefore, the overall estimation error is bounded by

$$\|\hat{\Pi}_1 - \Pi_1\| \le \sum_{u,v} \left| (\hat{\Pi}_1)_{uv} - (\Pi_1)_{uv} \right| \le \varepsilon.$$

Recall that the bound is subject to a change-of-basis transformation, and in the general scenario where the eigenvectors of A are not mutually orthogonal, the original prediction error bound should be multiplied by $\kappa_{\mathrm{e}}(A)$. Therefore, to achieve error threshold ε for predictions on Π_i, it is required that $c\binom{n}{k} r^{2t_0} < \frac{\varepsilon}{\kappa_{\mathrm{e}}(A) n^2}$, or equivalently, by *Stirling's Formula*,

$$t_0 > \frac{\log \kappa_{\mathrm{e}}(A) + \log \frac{cn^2}{\varepsilon} + \log \binom{n}{k}}{2 \log \frac{1}{r}} = O\left(\frac{k \log n - \log \varepsilon + \log \kappa_{\mathrm{e}}(A)}{2 \log \frac{|\lambda_k|}{|\lambda_{k+1}|}} \right). \tag{10}$$

This completes the proof. □

Proof of Corollary 5.2. We first construct a specific pair of orthonormal bases $(P_1^*, \hat{P}_1^*)$ that satisfy the corollary. To start with, take an arbitrary initial pair of orthonormal basis $(P_1^\circ, \hat{P}_1^\circ)$, and consider the SVD $(P_1^\circ)^\top \hat{P}_1^\circ = U \Sigma V^\top$, which is equivalent to $(P_1^\circ U)^\top (\hat{P}_1^\circ V) = \Sigma$. Note that the columns of $P_1^\circ U = [w_1 \ \cdots w_k]$ and $\hat{P}_1^\circ V = [\hat{w}_1 \ \cdots \hat{w}_k]$ form orthonormal bases of $\mathrm{col}(\Pi_1)$ and $\mathrm{col}(\hat{\Pi}_1)$, respectively; furthermore, these bases project onto each other accordingly by subscripts, namely

$$\Pi_1 \hat{w}_i = \sigma_i w_i, \ \hat{\Pi}_1 w_i = \sigma_i \hat{w}_i.$$

Now we set $P_1^* := P_1^\circ U$ and $\hat{P}_1^* := \hat{P}_1^\circ V$. Note that

$$|1 - \sigma_i| = \|(\hat{\Pi}_1 - \Pi_1)\hat{w}_i\| < \varepsilon,$$

which shows, by properties of projection matrix Π_1,

$$\|w_i - \hat{w}_i\| = \sqrt{\|w_i - \Pi_1 \hat{w}_i\|^2 + \|\Pi_1 \hat{w}_i - \hat{w}_i\|^2} = \sqrt{|1-\sigma_i|^2 + \|(\hat{\Pi}_1 - \Pi_1)\hat{w}_i\|^2} < \sqrt{2}\varepsilon,$$

and thus

$$\|P_1^* - \hat{P}_1^*\| = \max_{\|z\|=1} \|(P_1^* - \hat{P}_1^*)z\| = \max_{\|z\|=1} \left\| \sum_i z_i (w_i - \hat{w}_i) \right\| \leq \sqrt{k} \cdot \sqrt{2}\varepsilon.$$

To further generalize the proposition to any arbitrary $\hat{P}_1$, we only have to note that there exists an orthonormal matrix T that maps the basis $\hat{P}_1^*$ to $\hat{P}_1 = \hat{P}_1^* T$. Now take $P_1 = P_1^* T$, and we have

$$\|\hat{P}_1 - P_1\| = \|(\hat{P}_1^* - P_1^*)T\| = \|\hat{P}_1^* - P_1^*\| < \sqrt{2k}\varepsilon.$$

As for the estimation error bound for M_1, we can directly write

$$\begin{aligned}
\|P_1^\top A P_1 - \hat{P}_1^\top A \hat{P}_1\| &\leq \|P_1^\top A P_1 - P_1^\top A \hat{P}_1\| + \|P_1^\top A \hat{P}_1 - \hat{P}_1^\top A \hat{P}_1\| \\
&\leq \|A\| \|P_1 - \hat{P}_1\| + \|A\| \|P_1 - \hat{P}_1\| \\
&< 2\|A\|\delta,
\end{aligned}$$

This completes the proof of the corollary. □

Recall that we are allowed to take any orthonormal basis P_1 for E_{u}. Hence we shall always assume by default that P_1 in the proofs are selected as shown in the proof above.

We finish this section with simple but frequently-used bounds on $\|\hat{P}_1^\top P_1\|$ and $\|\hat{P}_1^\top P_2\|$. These factors represent an additional error introduced by using the inaccurate projector $\hat{P}_1$.

Proposition E.1. *Under the premises of Corollary 5.2, $\|I_k - \hat{P}_1^\top P_1\| < \delta$, $\|\hat{P}_1^\top P_2\| < \delta$.*

Proof. Note that $P_1^\top P_1 = I_k$ and $P_1^\top P_2 = O$, it is evident that

$$\begin{aligned}
\|I_k - \hat{P}_1^\top P_1\| &= \|(P_1 - \hat{P}_1)^\top P_1\| < \delta, \\
\|\hat{P}_1^\top P_2\| &= \|(\hat{P}_1 - P_1)^\top P_2\| = \|\hat{P}_1 - P_1\| < \delta.
\end{aligned}$$

This finishes the proof. □

F Proof of Theorem 4.2

We first consider a warm-up case where A is symmetric, which provides some intuition for the general case. In this case, the eigenvectors of A are mutually orthogonal, which guarantees $E_{\mathrm{u}}^\perp = E_{\mathrm{s}}$ (i.e., they are 0-close to each other) and thus $\Delta = \mathbf{0}$. This allows us to select $\tau = 1$, $\omega = 0$ and $\alpha = 1$, and the closed-loop dynamical matrix simplifies to

$$\hat{L}_1 = \begin{bmatrix} M_1 + P_1^\top B \hat{K}_1 \hat{P}_1^\top P_1 & P_1^\top B \hat{K}_1 \hat{P}_1^\top P_2 \\ P_2^\top B \hat{K}_1 \hat{P}_1^\top P_1 & M_2 + P_2^\top B \hat{K}_1 \hat{P}_1^\top P_2 \end{bmatrix}. \tag{11}$$

The norm of the top-left block is in the order of $O(\delta)$ based on the estimation error bound (see Theorem F.1) $\|\hat{B}_1 - B_1\| = O(\sqrt{k}\delta)$, which characterizes how well the controller can eliminate the unstable component. The spectrum of the bottom-right block can be viewed as a perturbation (note that $\|\hat{P}_1^\top P_2\| = O(\delta)$ is small by Proposition E.1) to a stable matrix M_2 (recall $\rho(M_2) = |\lambda_{k+1}|$), which should also be stable as long as δ is small enough. Meanwhile, the top-right block is also approximately zero, since only projection error contributes to the top-right block (again $\|\hat{P}_1^\top P_2\| = O(\delta)$). The above observations together show that $\hat{L}_1$ is in the order of

$$\hat{L}_1 = \begin{bmatrix} O(\delta) & O(\delta) \\ O(1) & |\lambda_{k+1}| + O(\delta) \end{bmatrix}, \tag{12}$$

which is almost lower-triangular. Therefore, we can apply the block perturbation bound to bound the spectrum of $\hat{L}_1$.

We start by showing the estimation error bound for B_1, which is straight-forward since $\Delta = \mathbf{0}$. Note that the upper bound of the norm of our controller $\hat{K}_1$ appears as a natural corollary of it.

Proposition F.1. *Under the premises of Theorem 4.2,* $\|\hat{B}_1 - B_1\| < 4\|A\|\sqrt{k}\delta$.

Proof. Note that the column vector b_i has estimation error bound

$$\begin{aligned}
\|b_i - \hat{b}_i\| &= \frac{1}{\|x_{t_i}\|}\left\|\left(P_1^\top x_{t_i+1} - M_1 P_1^\top x_{t_i}\right) - \left(\hat{P}_1^\top x_{t_i+1} - \hat{M}_1 \hat{P}_1^\top x_{t_i}\right)\right\| \\
&\le \frac{1}{\|x_{t_i}\|}\left(\|(P_1^\top - \hat{P}_1^\top)Ax_{t_i}\| + \|(M_1 P_1^\top - \hat{M}_1 \hat{P}_1^\top)x_{t_i}\|\right) \\
&\le \|P_1^\top - \hat{P}_1^\top\|\|A\| + \|M_1 P_1^\top - M_1 \hat{P}_1^\top\| + \|M_1 \hat{P}_1^\top - \hat{M}_1 \hat{P}_1^\top\| \\
&< \|A\|\delta + \|M_1\|\|P_1^\top - \hat{P}_1^\top\| + \|M_1 - \hat{M}_1\| \\
&< \|A\|\delta + \|A\|\delta + 2\|A\|\delta = 4\|A\|\delta,
\end{aligned}$$

where we repeatedly apply Corollary 5.2 and the fact that $\|M_1\| \le \|A\|$. Then, to bound the error of the whole matrix, we simply apply the definition

$$\|\hat{B}_1 - B_1\| = \max_{\|u\|=1} \|(\hat{B}_1 - B_1)u\| \le \max_{\|u\|=1} \sum_{i=1}^{k} |u_i|\|\hat{b}_i - b_i\| < 4\|A\|\sqrt{k}\delta.$$

This completes the proof. ☐

Corollary F.1. *Under the premises of Theorem 4.2, when (13) holds,* $\|\hat{K}_1\| < \frac{2\|A\|}{c\|B\|}$.

Proof. By Proposition F.1, it is evident that

$$\sigma_{\min}(\hat{B}_1) \ge \sigma_{\min}(B_1) - \|\hat{B}_1 - B_1\| > (c - 4\|A\|\sqrt{k}\delta)\|B\| > \frac{c}{2}\|B\|,$$

where the last inequality requires

$$\delta < \frac{c}{8\|A\|\sqrt{k}}. \tag{13}$$

Recall that $\hat{K}_1 = \hat{B}_1^{-1}\hat{M}_1$, and note that $\|\hat{B}_1^{-1}\| \le \frac{1}{\sigma_{\min}(\hat{B}_1)}$, so we have

$$\|\hat{K}_1\| = \|\hat{B}_1^{-1}\hat{M}_1\| \le \frac{\|\hat{P}_1^\top A \hat{P}_1\|}{\sigma_{\min}(\hat{B}_1)} < \frac{2\|A\|}{c\|B\|}.$$

This completes the proof. ☐

Recall that to apply Lemma 5.3, we need a bound on the spectral radii of diagonal blocks. The top-left block has already been eliminated to approximately $\mathbf{0}$ by the design of $\hat{K}_1$, but the bottom-right block needs some extra work — although M_2 is known to be stable, the inaccurate projection introduces an extra error that perturbs the spectrum. To bound the perturbed spectral radius, we will apply the following perturbation bound known as Bauer-Fike Theorem.

Lemma F.2 (Bauer-Fike)**.** *Suppose* $A \in \mathbb{R}^{n\times n}$ *is diagonalizable, then for any* $E \in \mathbb{R}^{n\times n}$*, we have*

$$|\rho(A) - \rho(A+E)| \le \max_{\hat{\lambda}\in\lambda(A+E)} \min_{\lambda\in\lambda(A)} |\lambda - \hat{\lambda}| \le \kappa_{\mathrm{e}}(A)\|E\|,$$

where $\kappa_{\mathrm{e}}(A)$ *is the condition number of the matrix consisting of* A*'s eigenvectors as columns, and* $\lambda(A)$ *denotes the spectrum of* A*.*

Proof. The proof is well-known and can be found in, e.g., [56]. ☐

Now we are ready to prove the main theorem for any symmetric dynamical matrix A.

Proof of Theorem 4.2. With $\tau = 1$, the controlled dynamics under estimated controller $\hat{K}_1$ becomes

$$\hat{L}_1 = \begin{bmatrix} M_1 + P_1^\top B\hat{K}_1\hat{P}_1^\top P_1 & P_1^\top B\hat{K}_1\hat{P}_1^\top P_2 \\ P_2^\top B\hat{K}_1\hat{P}_1^\top P_1 & M_2 + P_2^\top B\hat{K}_1\hat{P}_1^\top P_2 \end{bmatrix}.$$

We first guarantee that the diagonal blocks are stable. For the top-left block,

$$\begin{aligned}
\|M_1 + P_1^\top B\hat{K}_1\| &= \|M_1 - B_1\hat{B}_1^{-1}\hat{M}_1\hat{P}_1^\top P_1\| \\
&\le \|M_1 - \hat{M}_1\| + \|\hat{M}_1 - B_1\hat{B}_1^{-1}\hat{M}_1\| + \|B_1\hat{B}_1^{-1}\hat{M}_1(I_k - \hat{P}_1^\top P_1)\| \\
&\le \|M_1 - \hat{M}_1\| + \|\hat{B}_1 - B_1\|\|\hat{K}_1\| + \|B\|\|\hat{K}_1\|\|I_k - \hat{P}_1^\top P_1\| \\
&< 2\|A\|\delta + \frac{8\|A\|^2\sqrt{k}}{c\|B\|}\delta + \frac{2\|A\|}{c}\delta \qquad (14) \\
&= \frac{2\big(4\sqrt{k}\|A\| + (c+1)\|B\|\big)\|A\|}{c\|B\|}\delta,
\end{aligned}$$

where in (14) we apply Corollary 5.2, Corollary F.1, and Proposition E.1. Meanwhile, for the bottom-right block, note that the norm of the error term is bounded by

$$\|P_2^\top B\hat{K}_1\hat{P}_1^\top P_2\| \le \|B\|\|\hat{B}_1^{-1}\|\|\hat{M}_1\|\|\hat{P}_1^\top P_2\| \le \frac{2\|A\|}{c}\delta.$$

Hence, by Lemma F.2, the spectral radius of the bottom-right block is bounded by

$$\rho(M_2 + P_2^\top B\hat{K}_1\hat{P}_1^\top P_2) \le \rho(M_2) + \tfrac{2}{c}\kappa_{\mathrm{e}}(M_2)\|A\|\delta < 1,$$

where we require (recall that $\rho(M_2) = |\lambda_{k+1}|$)

$$\delta < \frac{c(1 - |\lambda_{k+1}|)}{2\kappa_{\mathrm{e}}(M_2)\|A\|}. \qquad (15)$$

To apply the lemma, it only suffices to bound the spectral norms of off-diagonal blocks. Note that the top-right block is bounded by

$$\|P_1^\top B\hat{K}_1\hat{P}_1^\top P_2\| \le \|B\|\|\hat{K}_1\|\|\hat{P}_1^\top P_2\| < \frac{2\|A\|}{c}\delta,$$

and the bottom-left block is bounded by

$$\|P_2^\top B\hat{K}_1\hat{P}_1^\top P_1\| \le \|B\|\|\hat{K}_1\| \le \frac{2\|A\|}{c}.$$

Now, by Lemma 5.3, we can guarantee that

$$\rho(\hat{L}_1) \le \max\left\{\frac{2\big(4\sqrt{k}\|A\| + 2(c+1)\|B\|\big)\|A\|}{c\|B\|}\delta, |\lambda_{k+1}| + \|B\|\|\hat{K}_1\|\delta\right\} + \frac{4\|A\|^2\chi(\hat{L}_1)}{c^2}\delta < 1,$$

where we require

$$\delta < \min\left\{\frac{1}{\frac{2\left(4\sqrt{k}\|A\|+2(c+1)\|B\|\right)\|A\|}{c\|B\|} + \frac{4\|A\|^2\chi(\hat{L}_1)}{c^2}}, \frac{1 - |\lambda_{k+1}|}{\frac{2\|A\|}{c} + \frac{4\|A\|^2\chi(\hat{L}_1)}{c^2}}\right\}. \qquad (16)$$

So far, we shall recollect all the constraints we need on δ (see (13), (15) and (16)), i.e.,

$$\delta < \min\left\{\frac{c}{8\|A\|\sqrt{k}}, \frac{c(1-|\lambda_{k+1}|)}{2\kappa_{\mathrm{e}}(M_2)\|A\|}, \frac{1 - |\lambda_{k+1}|}{\frac{2\|A\|}{c} + \frac{4\|A\|^2\chi(\hat{L}_1)}{c^2}}, \frac{1}{\frac{2\left(4\sqrt{k}\|A\|+2(c+1)\|B\|\right)\|A\|}{c\|B\|} + \frac{4\|A\|^2\chi(\hat{L}_1)}{c^2}}\right\},$$

which can be simplified (but weakened) to

$$\delta < \frac{c^2(1 - |\lambda_{k+1}|)}{16\sqrt{k}\kappa_{\mathrm{e}}(M_2)\|A\|(\|A\| + \|B\|)\chi(\hat{L}_1)} = O(k^{-1/2}). \qquad (17)$$

We shall rewrite the bound equivalently in terms of t_0 (recall (10) in Appendix E) as

$$t_0 > \frac{\log(cn^2\binom{n}{k}) - \log\frac{c^2(1-|\lambda_{k+1}|)}{16\sqrt{2}k\kappa_{\mathrm{e}}(M_2)\|A\|(\|A\|+\|B\|)\chi(\hat{L}_1)}}{2\log\frac{|\lambda_k|}{|\lambda_{k+1}|}} = O\left(\frac{k\log n}{\log\frac{|\lambda_k|}{|\lambda_{k+1}|}}\right), \qquad (18)$$

since $\kappa_{\mathrm{e}}(A) = 1$. This completes the proof of Theorem 4.2. □

G Proof of the Main Theorem

For the general case, the analysis becomes more challenging for two reasons: on the one hand, we have to apply τ-hop control with τ possibly larger than 1, which potentially increases the norm of B_τ and $\hat{K}_1$; on the other hand, the top-right corner will no longer be $O(\delta)$ with a non-zero Δ (in fact, Δ_τ is in the order of $|\lambda_1|^\tau$ that grows exponentially with respect to τ). To settle these issues, we first introduce two key observations on bounds of major factors:

(1) For an arbitrary matrix X, although $\|X\|$ might be significantly larger than $\rho(X)$, we always have $\|X^t\| = O(\rho(X)^t)$ when t is large enough. This is formally proven as Gelfand's Formula (see Lemma G.1), and helps to establish bounds like $\|M_1\| = O(|\lambda_1|^\tau)$, $\|M_2\| = O(|\lambda_{k+1}|^\tau)$, $\|\Delta_\tau\| = O(|\lambda_1|^\tau)$, $\|P_2^\top A^{\tau-1}\| = O(|\lambda_{k+1}|^\tau)$, and $\|\hat{M}_1^\tau - M_1^\tau\| = O(|\lambda_1|^\tau \delta)$.

(2) When the system runs in open loop for a long period (specifically, for ω time steps), eventually we will see the unstable component expanding and the stable component shrinking, and consequently $\frac{\|P_2^\top A^\omega x\|}{\|A^\omega x\|} = O(|\lambda_k|^{-\omega})$. This cancels out the exponentially exploding $\|\Delta_\tau\|$, and helps to establish the estimation bound $\|\hat{B}_\tau - B_\tau\| = O(|\lambda_1|^\tau \delta)$.

With these in hand, we are ready to upper bound the norms of the blocks in $\hat{L}_\tau$:

(1) *The top-left and bottom-right blocks*: similar to the warm-up case, only to note that dynamical matrices are lifted to their τ^{th} power, and thus $\|\hat{B}_\tau - B_\tau\|$ carries an additional factor of $|\lambda_1|^\tau$.

(2) *The bottom-left block*: $P_2^\top A^{\tau-1}$ contributes an $O(|\lambda_{k+1}|^\tau)$ factor that decays exponentially, while $\hat{K}_1$ contributes an $O(|\lambda_1|^\tau)$ factor that explodes exponentially. The overall bound is in the order of $O(|\lambda_1 \lambda_{k+1}/\lambda_k|^\tau)$, and decays with respect to τ if $|\lambda_1||\lambda_{k+1}|/|\lambda_k| < 1$.

(3) *The top-right block*: the first term is in the order of $O(|\lambda_1|^\tau)$, and the second term is in the order of $O(|\lambda_1 \lambda_{k+1}/\lambda_k|^\tau \delta)$. This block is in the order of $O(|\lambda_1|^\tau)$ when δ is small enough.

Therefore, the closed-loop dynamical matrix is actually in the order of

$$\hat{L}_\tau = \begin{bmatrix} O(|\lambda_1|^{2\tau}\delta) & O\big(|\lambda_1|^\tau + |\lambda_1\lambda_{k+1}/\lambda_k|^\tau \delta\big) \\ O(|\lambda_1\lambda_{k+1}/\lambda_k|^\tau) & O\big(|\lambda_{k+1}|^\tau + |\lambda_1\lambda_{k+1}|^\tau \delta\big) \end{bmatrix}. \tag{19}$$

Finally, by Lemma 5.3, asymptotic stability is guaranteed when $|\lambda_1|^2|\lambda_{k+1}| < |\lambda_k|$ (i.e., the norm of the bottom-left block decays faster than the norm of the top-right block grows), in which case we can set τ to be some constant determined by A and B, and δ in the order of $O(|\lambda_1|^{-2\tau})$.

Technically, we would like to bound the spectral radius of the matrix

$$\hat{L}_\tau = \begin{bmatrix} M_1^\tau + P_1^\top A^{\tau-1} B \hat{K}_1 \hat{P}_1^\top P_1 & \Delta_\tau + P_1^\top A^{\tau-1} B \hat{K}_1 \hat{P}_1^\top P_2 \\ P_2^\top A^{\tau-1} B \hat{K}_1 \hat{P}_1^\top P_1 & M_2^\tau + P_2^\top A^{\tau-1} B \hat{K}_1 \hat{P}_1^\top P_2. \end{bmatrix}$$

using Lemma 5.3. The proof is split into two major building blocks: on the one hand, we introduce the well-known Gelfand's Formula to bound matrices appearing with exponents; on the other hand, we establish the estimation error bound for B_τ (parallel to Lemma F.1) and proceed to bound $\|\hat{K}_1\|$, for which we rely on the instability results shown in Section G.2. Finally, a combination of these building blocks naturally establishes the main theorem.

G.1 Gelfand's Formula

In this section, we will show norm bounds for factors that contain matrix exponents. It is natural to apply the well-known Gelfand's Formula as stated below.

Lemma G.1 (Gelfand's Formula)**.** *For any square matrix X, we have*

$$\rho(X) = \lim_{t\to\infty} \|X^t\|^{1/t}. \tag{20}$$

In other words, for any $\varepsilon > 0$, there exists a constant $\zeta_\varepsilon(X)$ such that

$$\sigma_{\max}(X^t) = \|X^t\| \le \zeta_\varepsilon(X)(\rho(X) + \varepsilon)^t. \tag{21}$$

Further, if X is invertible, let $\lambda_{\min}(X)$ denote the eigenvalue of X with minimum modulus, then

$$\sigma_{\min}(X^t) \ge \frac{1}{\zeta_\varepsilon(X^{-1})} \left(\frac{|\lambda_{\min}(X)|}{1 + \varepsilon|\lambda_{\min}(X)|} \right)^t. \tag{22}$$

Proof. The proof of (20) can be easily found in existing literature (e.g., [57], Corollary 5.6.14), and (21) follows by the definition of limits. For (22), note that

$$\sigma_{\min}(X^t) = \frac{1}{\sigma_{\max}((X^{-1})^t)} \geq \frac{1}{\zeta_\varepsilon(X^{-1})(\rho(X^{-1})+\varepsilon)^t} = \frac{1}{\zeta_\varepsilon(X^{-1})}\left(\frac{|\lambda_{\min}(X)|}{1+\varepsilon|\lambda_{\min}(X)|}\right)^t,$$

where we apply $\sigma_{\min}(X^t) = \sigma_{\max}((X^{-1})^t)^{-1}$ and $\rho(X^{-1}) = |\lambda_{\min}(X)|^{-1}$. □

It is evident that $\rho(A) = \rho(M_1) = \rho(N_1) = |\lambda_1|$, $\lambda_{\min}(M_1) = \lambda_{\min}(N_1) = |\lambda_k|$ and $\rho(M_2) = \rho(N_2) = |\lambda_{k+1}|$ (recall that M_1 and M_2 inherits the unstable and stable eigenvalues, respectively). Therefore, we can use Gelfand's Formula to bound the relevant factors appearing in $\hat{L}_\tau$.

Proposition G.1. *Under the premises of Theorem 4.1, the following results hold for any $t \in \mathbb{N}$:*

(1) $\|B_t\| \leq \zeta_{\varepsilon_1}(A)(|\lambda_1|+\varepsilon_1)^{t-1}\|B\|$*;*

(2) $\|P_2^\top A^t\| \leq \zeta_{\varepsilon_2}(M_2)(|\lambda_{k+1}|+\varepsilon_2)^t$*;*

(3) $\|\Delta_t\| \leq C_\Delta(|\lambda_1|+\varepsilon_1)^t$*, where* $C_\Delta = \zeta_{\varepsilon_1}(M_1)\zeta_{\varepsilon_2}(M_2)\frac{(2-\xi)\sqrt{2\xi}\|A\|}{1-\xi}\frac{2|\lambda_{k+1}|}{|\lambda_1|+\varepsilon_1-|\lambda_{k+1}|-\varepsilon_2}$.

Here (and below) ε_1 and ε_2 are selected to be sufficiently small constants (see (47)).

Proof. (1) This is a direct corollary of Gelfand's Formula, since

$$\|B_t\| = \|P_1^\top A^{t-1}B\| \leq \|A^{t-1}\|\|B\| \leq \zeta_{\varepsilon_1}(A)(|\lambda_1|+\varepsilon_1)^{t-1}\|B\|.$$

(2) It only suffices to recall $\rho(M_2) = |\lambda_{k+1}|$, and note that

$$P_2^\top A^t = P_2^\top PM^tP^{-1} = [\mathbf{0}\ I_{n-k}]M^tP^\top = M_2^tP_2^\top.$$

Hence by Gelfand's Formula we have $\|P_2^\top A^t\| = \|M_2^t\| \leq \zeta_{\varepsilon_2}(M_2)(|\lambda_{k+1}|+\varepsilon_2)^t$.

(3) This is a direct corollary of Lemma A.1(4) and Gelfand's Formula, since

$$\begin{aligned}\|\Delta_t\| &= \left\|\sum_i M_1^i \Delta M_2^{t-1-i}\right\| \leq \|\Delta\|\sum_i \|M_1^i\|\|M_2^{t-1-i}\| \\ &\leq \zeta_{\varepsilon_1}(M_1)\zeta_{\varepsilon_2}(M_2)\frac{(2-\xi)\sqrt{2\xi}\|A\|}{1-\xi}\sum_i(\varepsilon_1+|\lambda_1|)^i(|\lambda_{k+1}|+\varepsilon_2)^{t-1-i} \\ &= C_\Delta(|\lambda_1|+\varepsilon_1)^t.\end{aligned}$$

This finishes the proof of the proposition. □

Proposition G.2. *Under the premises of Theorem 4.1,*

$$\|\hat{M}_1^\tau - M_1^\tau\| < 2\tau\|A\|\zeta_{\varepsilon_1}(A)^2(|\lambda_1|+\varepsilon_1)^{\tau-1}\delta.$$

Proof. Recall that Corollary 5.2 gives $\|M_1 - \hat{M}_1\| < 2\|A\|\delta$. Meanwhile, by Gelfand's Formula,

$$\|M_1^t\| = \|P^\top A^tP\| \leq \|A^t\| \leq \zeta_{\varepsilon_1}(A)(|\lambda_1|+\varepsilon_1)^t,$$
$$\|\hat{M}_1^t\| = \|\hat{P}^\top A^t\hat{P}\| \leq \|A^t\| \leq \zeta_{\varepsilon_1}(A)(|\lambda_1|+\varepsilon_1)^t.$$

Then we have the following bound by telescoping

$$\begin{aligned}\|M_1^\tau - \hat{M}_1^\tau\| &= \left\|\sum_{i=1}^\tau\left(M_1^i\hat{M}_1^{\tau-i} - M_1^{i-1}\hat{M}_1^{\tau-i+1}\right)\right\| \\ &\leq \sum_{i=1}^\tau \|M_1^{i-1}\|\|\hat{M}_1^{\tau-i}\|\|M_1-\hat{M}_1\| \\ &< \tau\cdot\zeta_{\varepsilon_1}(A)^2(|\lambda_1|+\varepsilon_1)^{\tau-1}\cdot 2\|A\|\delta \\ &= 2\tau\|A\|\zeta_{\varepsilon_1}(A)^2(|\lambda_1|+\varepsilon_1)^{\tau-1}\delta.\end{aligned}$$

This finishes the proof. □

Corollary G.2. *Under the premises of Theorem 4.1, when* $\delta < \frac{1}{\tau}$,

$$\|\hat{M}_1^\tau\| < \big(\zeta_{\varepsilon_1}(M_1)(|\lambda_1| + \varepsilon_1) + 2\|A\|\zeta_{\varepsilon_1}(A)\big)(|\lambda_1| + \varepsilon_1)^{\tau-1}.$$

Proof. A combination of Gelfand's Formula and Proposition G.2 yields

$$\begin{aligned}\|\hat{M}_1^\tau\| &\le \|M_1^\tau\| + \|\hat{M}_1^\tau - M_1^\tau\| \\ &\le \zeta_{\varepsilon_1}(M_1)(|\lambda_1| + \varepsilon_1)^\tau + 2\tau\|A\|\zeta_{\varepsilon_1}(A)^2(|\lambda_1| + \varepsilon_1)^{\tau-1}\delta \\ &< \big(\zeta_{\varepsilon_1}(M_1)(|\lambda_1| + \varepsilon_1) + 2\tau\|A\|\zeta_{\varepsilon_1}(A)\delta\big)(|\lambda_1| + \varepsilon_1)^{\tau-1},\end{aligned}$$

where the last inequality requires $\delta < \frac{1}{\tau}$. This completes the proof. □

G.2 Instability of the Unstable Component

We have been referring to E_{s} (and approximately, $E_{\mathrm{u}}^\perp$) as "stable", and E_{u} as "unstable". This leads us to think that the unstable component will constitute an increasing proportion of the state as the system evolves with zero control input. However, in some cases it might happen that the proportion of unstable component does not increase within the first few time steps, although eventually it will explode. This motivates us to formally characterize such instability of the unstable component.

In this section, we aim to establish a fundamental property of A^ω (for large enough ω, of course) that it "almost surely" increases the norm of the state. By "almost surely" we mean that the initial state should have non-negligible unstable component, which happens with probability $1 - \varepsilon$ when we uniformly sample the initial state from the surface of unit hyper-sphere in $\mathbb{R}^n$.

Throughout this section, we use γ to denote the ratio of the unstable component over the stable component within some state x (i.e., $\frac{\|R_1 x\|}{\|R_2 x\|}$). Note that

$$x = \Pi_{\mathrm{u}} x + \Pi_{\mathrm{s}} x = Q_1 R_1 x + Q_2 R_2 x,$$

where Q_1, Q_2 are orthonormal. Hence

$$\|R_1 x\| - \|R_2 x\| \le \|x\| \le \|R_1 x\| + \|R_2 x\|.$$

As a consequence, when $\frac{\|R_1 x\|}{\|R_2 x\|} > \gamma > 1$, we also know that

$$\frac{\|R_1 x\|}{\|x\|} \ge \frac{\|R_1 x\|}{\|R_1 x\| + \|R_2 x\|} > \frac{\gamma}{\gamma + 1}, \quad \frac{\|R_2 x\|}{\|x\|} \le \frac{\|R_2 x\|}{\|R_1 x\| - \|R_2 x\|} < \frac{1}{\gamma - 1}.$$

The following results are presented to fit in the framework of an inductive proof. We first establish the inductive step, where Proposition G.3 shows that the unstable component eventually becomes dominant with a non-negligible initial γ, and Proposition G.4 shows that the unstable component will still constitute a non-negligible part after a control input of mild magnitude is injected. Meanwhile, Proposition G.5 shows that the initial unstable component is non-negligible with large probability.

Proposition G.3. *Given a dynamical matrix A and some constant $\gamma > 0$, for any state x such that $\frac{\|R_1 x\|}{\|R_2 x\|} > \gamma$, for any $\omega \in \mathbb{N}$, we have*

$$\frac{\|R_1 A^\omega x\|}{\|R_2 A^\omega x\|} > \gamma_\omega := C_\gamma \left(\frac{|\lambda_k|}{(1 + \varepsilon_3|\lambda_k|)(|\lambda_{k+1}| + \varepsilon_2)}\right)^\omega,$$

where $C_\gamma := \frac{1}{(1+\frac{1}{\gamma})\zeta_{\varepsilon_3}(N_1^{-1})\zeta_{\varepsilon_2}(N_2)\|R_2\|}$ is a constant related to γ. Specifically, for any $\gamma_+ > 0$, there exists a constant $\omega_0(\gamma, \gamma_+) = O(\log \frac{\gamma_+}{\gamma})$, such that for any $\omega > \omega_0(\gamma, \gamma_+)$, $\frac{\|R_1 x\|}{\|R_2 x\|} > \gamma_+$.

Proof. Recall that $R_1 A^\omega = N_1^\omega R_1$ and $R_2 A^\omega = N_2^\omega R_2$. By Gelfand's Formula we have

$$\begin{aligned}\frac{\|R_1 A^\omega x\|}{\|R_2 A^\omega x\|} &= \frac{\|N_1^\omega R_1 x\|}{\|N_2^\omega R_2 x\|} \ge \frac{\sigma_{\min}(N_1^\omega)\|R_1 x\|}{\|N_2^\omega\|\|R_2\|\|x\|} > \frac{\sigma_{\min}(N_1^\omega)}{(1 + \frac{1}{\gamma})\|N_2^\omega\|\|R_2\|} \\ &\ge \frac{\big(|\lambda_k|/(1 + \varepsilon_3|\lambda_k|)\big)^\omega}{(1 + \frac{1}{\gamma})\zeta_{\varepsilon_3}(N_1^{-1})\zeta_{\varepsilon_2}(N_2)(|\lambda_{k+1}| + \varepsilon_2)^\omega\|R_2\|}\end{aligned}$$

$$= \frac{1}{(1+\frac{1}{\gamma})\zeta_{\varepsilon_3}(N_1^{-1})\zeta_{\varepsilon_2}(N_2)\|R_2\|}\left(\frac{|\lambda_k|}{(1+\varepsilon_3|\lambda_k|)(|\lambda_{k+1}|+\varepsilon_2)}\right)^{\omega}.$$

Therefore, we shall take

$$\omega_0(\gamma,\gamma_+) = \frac{\log \gamma_+/C_\gamma}{\log(|\lambda_k|)/\big((1+\varepsilon_3|\lambda_k|)(|\lambda_{k+1}|+\varepsilon_2)\big)} = O\left(\log\frac{\gamma_+}{\gamma}\right),$$

and the proof is completed. □

Corollary G.3. *Under the premises of Proposition G.3, for any $\omega > \omega_0(\gamma,\gamma_+)$,*

$$\frac{\|P_1^\top A^\omega x\|}{\|A^\omega x\|} > 1 - \frac{2}{\gamma_\omega - 1}, \quad \frac{\|P_2^\top A^\omega x\|}{\|A^\omega x\|} < \frac{1}{\gamma_\omega - 1}.$$

Proof. Note that we have decomposition $x = \Pi_{\mathrm{u}} x + \Pi_1\Pi_{\mathrm{s}} x + \Pi_2\Pi_{\mathrm{s}} x$, where $\|\Pi_{\mathrm{u}} x\| = \|R_1 x\|$ and $\|\Pi_{\mathrm{s}} x\| = \|R_2 x\|$. Hence, for any $\omega > \omega_0(\gamma,\gamma_+)$, we can show that

$$\begin{aligned}\frac{\|P_1^\top A^\omega x\|}{\|A^\omega x\|} &= \frac{\|\Pi_{\mathrm{u}} A^\omega x + \Pi_1\Pi_{\mathrm{s}} A^\omega x\|}{\|A^\omega x\|} \\ &\geq \frac{\|\Pi_{\mathrm{u}} A^\omega x\| - \|\Pi_1\Pi_{\mathrm{s}} A^\omega x\|}{\|A^\omega x\|} \\ &\geq \frac{\|R_1 A^\omega x\| - \|R_2 A^\omega x\|}{\|A^\omega x\|} \\ &> \frac{\gamma_\omega}{\gamma_\omega + 1} - \frac{1}{\gamma_\omega - 1} > 1 - \frac{2}{\gamma_\omega - 1},\end{aligned}$$

and similarly,

$$\frac{\|P_2^\top A^\omega x\|}{\|A^\omega x\|} = \frac{\|\Pi_2\Pi_{\mathrm{s}} A^\omega x\|}{\|A^\omega x\|} \leq \frac{\|\Pi_{\mathrm{s}} A^\omega x\|}{\|A^\omega x\|} < \frac{1}{\gamma_\omega - 1}.$$

The proof is completed. □

Proposition G.4. *Given dynamical matrices A, B and constants $\gamma > 0, \gamma_+ > 1$, for any state x such that $\frac{\|R_1 x\|}{\|R_2 x\|} > \gamma_+$, suppose we feed a control input $\|u\| \leq \alpha\|x\|$ and observe the next state $x' = Ax + Bu$, where α satisfies*

$$\alpha < \frac{\frac{\gamma_+}{\gamma_+ + 1}\sigma_{\min}(M_1) - \frac{\gamma}{\gamma_+ - 1}\frac{1}{1-\xi}\|A\|}{(1 + \frac{\sqrt{2\xi}}{1-\xi} + \frac{\gamma}{1-\xi})\|B\|}. \tag{23}$$

Then we can guarantee that $\frac{\|R_1 x'\|}{\|R_2 x'\|} > \gamma$.

Proof. The proposition can be shown by direct calculation. Let $z = Rx = [z_1^\top, z_2^\top]^\top$. Recall that

$$Rx' = z' = \begin{bmatrix} N_1 z_1 + R_1 B u \\ N_2 z_2 + R_2 B u\end{bmatrix},$$

and note that $\frac{\|z_1\|}{\|x\|} > \frac{\gamma_+}{\gamma_+ + 1}$, $\frac{\|z_2\|}{\|x\|} < \frac{1}{\gamma_+ - 1}$ under the assumptions, so we have

$$\begin{aligned}\frac{\|R_1 x'\|}{\|R_2 x'\|} &= \frac{\|N_1 z_1 + R_1 B u\|}{\|N_2 z_2 + R_2 B u\|} \geq \frac{\|N_1 z_1\| - \|R_1 B u\|}{\|N_2 z_2\| + \|R_2 B u\|} \\ &\geq \frac{\sigma_{\min}(N_1)\|z_1\| - \|R_1 B\|\|u\|}{\|N_2\|\|z_2\| + \|R_2 B\|\|u\|} \\ &\geq \frac{\sigma_{\min}(N_1)\frac{\gamma_+}{\gamma_+ + 1}\|x\| - \alpha\|R_1\|\|B\|\|x\|}{\|N_2\|\frac{1}{\gamma_+ - 1}\|x\| + \alpha\|R_2\|\|B\|\|x\|} \\ &\geq \frac{\sigma_{\min}(M_1)\frac{\gamma_+}{\gamma_+ + 1}\|x\| - \alpha(1 + \frac{\sqrt{2\xi}}{1-\xi})\|B\|\|x\|}{\frac{1}{1-\xi}\|A\|\frac{1}{\gamma_+ - 1}\|x\| + \alpha\frac{1}{1-\xi}\|B\|\|x\|} \\ &> \gamma,\end{aligned}$$

where we apply Lemma A.1 and the convention of taking $N_1 = M_1$. □

Proposition G.5. *Suppose a state x is sampled uniformly randomly from the unit hyper-sphere surface $\mathbb{B}_n \subset \mathbb{R}^n$, then for any constant $\gamma < \min\left\{\frac{1}{2}, \frac{1}{\sqrt{2/(\sigma_{\min}(R_1)k)+1}}\right\}$, we have*

$$\mathrm{Pr}_{x\sim\mathcal{U}(\mathbb{B}_n)}\left[\frac{\|R_1x\|}{\|R_2x\|} > \gamma\right] > 1-\theta(\gamma),$$

where $\theta(\gamma) = \frac{8\sqrt{2}}{\mathrm{B}(\frac{1}{2},\frac{n-1}{2})\sqrt{\sigma_{\min}(R_1)}}\gamma = O(\gamma)$ is a constant bounded linearly by γ.

Proof. Note that

$$\|R_1x\| > \frac{\gamma}{1-\gamma}\|x\| \;\Rightarrow\; \|R_2x\| < \|x\| + \|R_1x\| < \frac{1}{1-\gamma}\|x\| \;\Rightarrow\; \frac{\|R_1x\|}{\|R_2x\|} > \gamma.$$

so we only have to show that $\mathrm{Pr}_{x\sim\mathcal{U}(\mathbb{B}_n)}\left[\|R_1x\| \le \frac{\gamma}{1-\gamma}\right] < \theta(\gamma)$. Now let $R_1^\top R_1 = S^\top DS$ be the eigen-decomposition of $R_1^\top R_1$, where S is selected to be orthonormal such that

$$D = \mathrm{diag}(d_1,\cdots,d_k,0,\cdots,0).$$

Note that the vector $y = Sx =: [y_1,\cdots,y_n]$ also obeys a uniform distribution over $\mathbb{B}_n$, so we have

$$\begin{aligned}\mathrm{Pr}\left[\|R_1x\| \le \tfrac{\gamma}{1-\gamma}\right] &= \mathrm{Pr}\left[x^\top R_1^\top R_1 x \le (\tfrac{\gamma}{1-\gamma})^2\right] = \mathrm{Pr}\left[y^\top D y \le (\tfrac{\gamma}{1-\gamma})^2\right]\\ &\le \mathrm{Pr}\left[d_i y_i^2 \le \tfrac{1}{k}(\tfrac{\gamma}{1-\gamma})^2,\ \forall i = 1,\ldots,k\right]\\ &\le \sum_{i=1}^k \mathrm{Pr}\left[y_i^2 \le \tfrac{1}{d_i k}(\tfrac{\gamma}{1-\gamma})^2\right].\end{aligned}$$

It suffices to bound the probability $\mathrm{Pr}_{y\sim\mathcal{U}(B)}\left[y_i^2 \le \eta\right]$. Note that y can be obtained by first sampling a Gaussian random vector $z \sim \mathcal{N}(0, I_n)$, and then normalize it to get $y = \frac{z}{\|z\|}$. Hence

$$\mathrm{Pr}_{y\sim\mathcal{U}(\mathbb{B}_n)}\left[y_i^2 \le \eta\right] = \mathrm{Pr}_{z\sim\mathcal{N}(0,I_n)}\left[z_i^2 \le \eta\|z\|^2\right] = \mathrm{Pr}_{z\sim\mathcal{N}(0,I_n)}\left[\frac{z_i^2}{\sum_{j\ne i} z_j^2} \le \frac{\eta}{1-\eta}\right],$$

where $w := \frac{z_i^2}{\sum_{j\ne i} z_j^2}$ is known to obey an F-distribution $w \sim \mathcal{F}(1, n-1)$. The c.d.f. of w is known to be $I_{w/(w+n-1)}(\frac{1}{2},\frac{n-1}{2})$, where I denotes the *regularized incomplete Beta function*. Note that

$$I_{w/(w+n-1)}\left(\frac{1}{2},\frac{n-1}{2}\right) = \frac{2w^{1/2}}{(n-1)^{1/2}\mathrm{B}(\frac{1}{2},\frac{n-1}{2})} - \frac{nw^{3/2}}{3(n-1)^{3/2}\mathrm{B}(\frac{1}{2},\frac{n-1}{2})} + O(n^{5/2}),$$

it can be shown that $I_{w/(w+n-1)}\left(\frac{1}{2},\frac{n-1}{2}\right) < \frac{4\sqrt{w}}{\sqrt{n-1}\mathrm{B}(\frac{1}{2},\frac{n-1}{2})}$. Hence

$$\mathrm{Pr}_{y\sim\mathcal{U}(\mathbb{B}_n)}\left[y_i^2 \le \eta\right] = \mathrm{Pr}_{z\sim\mathcal{N}(0,I_n)}\left[\frac{z_i^2}{\sum_{j\ne i} z_j^2} \le \frac{\eta}{1-\eta}\right] < \frac{4\sqrt{\frac{\eta}{1-\eta}}}{\sqrt{n-1}\mathrm{B}(\frac{1}{2},\frac{n-1}{2})},$$

which further gives

$$\mathrm{Pr}\left[\|R_1x\| \le \tfrac{\gamma}{1-\gamma}\right] < \sum_{i=1}^k \frac{4\sqrt{\frac{2}{d_i k}(\frac{\gamma}{1-\gamma})^2}}{\sqrt{n-1}\mathrm{B}(\frac{1}{2},\frac{n-1}{2})} < \frac{8\sqrt{2}}{\mathrm{B}(\frac{1}{2},\frac{n-1}{2})\sqrt{\sigma_{\min}(R_1)}}\gamma = O(\gamma)$$

where we require $\gamma < \min\left\{\frac{1}{2}, \frac{1}{\sqrt{2/(\sigma_{\min}(R_1)k)+1}}\right\}$. □

Combining the previous three propositions, we have shown in an inductive way that the algorithm guarantees $\frac{\|P_2^\top x_{t_i}\|}{\|x_{t_i}\|}$ is constantly upper bounded at each time step t_i ($i = 1,\cdots,k$), which is critical to the estimation error bound of B_τ. This is concluded as the following lemma.

Lemma G.4. *Under the premises of Theorem 4.1, for any constants* ω, γ *such that* $\omega < t_0$ *and* $\gamma < \min\left\{\frac{1}{2}, \frac{1}{\sqrt{2/(\sigma_{\min}(R_1)k)+1}}\right\}$, *the algorithm guarantees*

$$\frac{\|P_2^\top x_{t_i}\|}{\|x_{t_i}\|} < \frac{1}{\gamma_\omega - 1},\ \forall i = 1, \cdots, k$$

with probability $1 - \theta(\gamma)$ *over the initialization of* x_0 *on the unit hyper-sphere surface* $\mathbb{B}_n$, *where*

$$\gamma_\omega := C_\gamma \left(\frac{|\lambda_k|}{(1+\varepsilon_3|\lambda_k|)(|\lambda_{k+1}|+\varepsilon_2)}\right)^\omega.$$

Proof. We proceed by showing that $\frac{\|R_1 x_{t_i}\|}{\|R_2 x_{t_i}\|} > \gamma_\omega$ for $i = 1, \cdots, k$ in an inductive way.

For the base case, it is guaranteed by Proposition G.5 that x_0 satisfies $\frac{\|R_1 x_0\|}{\|R_2 x_0\|} > \gamma$ with probability $1 - \theta(\gamma)$, and Proposition G.3 further guarantees $\frac{\|R_1 x_{t_1}\|}{\|R_2 x_{t_1}\|} > \gamma_\omega$. Here we require $t_0 > \omega$.

For the inductive step, suppose we have shown $\frac{\|R_1 x_{t_i}\|}{\|R_2 x_{t_i}\|} > \gamma_\omega$. Since $\|u_{t_i}\| = \alpha\|x_{t_i}\|$, we have $\frac{\|R_1 x_{t_i+1}\|}{\|R_2 x_{t_i+1}\|} > \gamma$ by Proposition G.4, and again Proposition G.3 guarantees $\frac{\|R_1 x_{t_{i+1}}\|}{\|R_2 x_{t_{i+1}}\|} > \gamma_\omega$.

Now it only suffices to apply Corollary G.3 to complete the proof. □

G.3 Estimation Error of B_τ

Proposition G.6. *Under the premises of Theorem 4.1 and Lemma G.4, when (29) holds,*

$$\|\hat{B}_\tau - B_\tau\| < C_B(|\lambda_1| + \varepsilon_1)^{\tau-1}\delta,$$

where $C_B := \frac{2\sqrt{k}\zeta_{\varepsilon_1}(A)^2\left((2\tau+2)\|A\|+\|B\|\right)}{\alpha}$.

Proof. This is parallel to Lemma F.1. Note that we have to subtract an additional term (induced by non-zero Δ_τ in M^τ) to calculate the actual b_i, so we have

$$\begin{aligned}
\|b_i - \hat{b}_i\| &= \frac{1}{\alpha\|x_{t_i}\|}\left\|\left(P_1^\top x_{t_i+\tau} - M_1^\tau P_1^\top x_{t_i} - \Delta_\tau P_2^\top x_{t_i}\right) - \left(\hat{P}_1^\top x_{t_i+\tau} - \hat{M}_1^\tau \hat{P}_1^\top x_{t_i}\right)\right\| \\
&\le \frac{1}{\alpha\|x_{t_i}\|}\left(\|(P_1 - \hat{P}_1)^\top(A^\tau x_{t_i} + B_\tau u_{t_i})\| + \|M_1^\tau P_1^\top x_{t_i} - \hat{M}_1^\tau \hat{P}_1^\top x_{t_i}\| + \|\Delta_\tau P_2^\top x_{t_i}\|\right) \\
&< \frac{1}{\alpha}\left(\zeta_{\varepsilon_1}(A)^2(|\lambda_1|+\varepsilon_1)^{\tau-1}\left((2\tau+2)\|A\| + \|B\|\right)\delta + \delta\right).
\end{aligned}$$

Here the first term is bounded by

$$\begin{aligned}
\|(P_1 - \hat{P}_1)^\top(A^\tau x_{t_i} + B_\tau u_{t_i})\| &\le \|P_1 - \hat{P}_1\|(\|A^\tau\| + \|A^{\tau-1}B\|)\|x_{t_i}\| \\
&< \|x_{t_i}\|\zeta_{\varepsilon_1}(A)(|\lambda_1|+\varepsilon_1)^{\tau-1}(\|A\| + \|B\|)\delta,
\end{aligned}$$

where in the last inequality we apply Corollary 5.2; the second term is bounded by

$$\begin{aligned}
\|M_1^\tau P_1^\top x_{t_i} - \hat{M}_1^\tau \hat{P}_1^\top x_{t_i}\| &\le (\|M_1^\tau(P_1^\top - \hat{P}_1^\top)\| + \|(M_1^\tau - \hat{M}_1^\tau)\hat{P}_1^\top\|)\|x_{t_i}\| \\
&< \left(\zeta_{\varepsilon_1}(A)(|\lambda_1|+\varepsilon_1)^{\tau-1}\|A\|\delta \right.\\
&\qquad \left. + 2\tau\|A\|\zeta_{\varepsilon_1}(A)^2(|\lambda_1|+\varepsilon_1)^{\tau-1}\delta\right)\|x_{t_i}\| &(24)\\
&\le \|x_{t_i}\|\zeta_{\varepsilon_1}(A)^2(|\lambda_1|+\varepsilon_1)^{\tau-1}(2\tau+1)\|A\|\delta, &(25)
\end{aligned}$$

where in (24) we apply Proposition G.2, and in (25) we apply a simple fact that $\zeta_{\varepsilon_1}(A) \ge 1$; the third term is bounded by

$$\begin{aligned}
\frac{\|\Delta_\tau\|\|P_2^\top x_{t_i}\|}{\|x_{t_i}\|} &\le \frac{C_\Delta(|\lambda_1|+\varepsilon_1)^\tau}{\left[C_\gamma\left(\frac{|\lambda_k|}{(1+\varepsilon_3|\lambda_k|)(|\lambda_{k+1}|+\varepsilon_2)}\right)^\omega - 1\right]} &(26)\\
&< \frac{2C_\Delta(|\lambda_1|+\varepsilon_1)^\tau}{C_\gamma\left(\frac{|\lambda_k|}{(1+\varepsilon_3|\lambda_k|)(|\lambda_{k+1}|+\varepsilon_2)}\right)^\omega} &(27)
\end{aligned}$$

$$< \delta, \tag{28}$$

where in (26) we apply Lemma G.4, while in (27) and (28) we require

$$\omega > \max\left\{ \frac{\log 2/C_\gamma}{\log\big(|\lambda_k|/(1+\varepsilon_3|\lambda_k|)(|\lambda_{k+1}|+\varepsilon_2)\big)}, \frac{\log(2C_\Delta)/(C_\gamma\delta)+\tau\log(|\lambda_1|+\varepsilon_1)}{\log\big(|\lambda_k|/(1+\varepsilon_3|\lambda_k|)(|\lambda_{k+1}|+\varepsilon_2)\big)} \right\}. \tag{29}$$

Finally, to bound the error of the whole matrix, we simply apply the definition

$$\begin{aligned} \|\hat{B}_\tau - B_\tau\| &= \max_{\|u\|=1} \|(\hat{B}_\tau - B_\tau)u\| \le \max_{\|u\|=1} \sum_{i=1}^{k} |u_i| \|\hat{b}_i - b_i\| \\ &< \frac{\sqrt{k}}{\alpha}\left(\zeta_{\varepsilon_1}(A)^2(|\lambda_1|+\varepsilon_1)^{\tau-1}\big((2\tau+2)\|A\|+\|B\|\big)+1\right)\delta \\ &< \frac{2\sqrt{k}\zeta_{\varepsilon_1}(A)^2\big((2\tau+2)\|A\|+\|B\|\big)}{\alpha}(|\lambda_1|+\varepsilon_1)^{\tau-1}\delta. \end{aligned}$$

This completes the proof. □

Corollary G.5. *Under the premises of Theorem 4.1 and Lemma G.4, when (29), (30) and (31) hold,*

$$\sigma_{\min}(\hat{B}_\tau) > \frac{c\|B\|}{4\zeta_{\varepsilon_3}(N_1^{-1})}\left(\frac{|\lambda_k|}{1+\varepsilon_3|\lambda_k|}\right)^{\tau-1}.$$

Proof. We apply the $E_{\mathrm{u}} \oplus E_{\mathrm{s}}$-decomposition. Note that

$$B_\tau = P_1^\top A^{\tau-1}B = P_1^\top(Q_1N_1^{\tau-1}R_1 + Q_2N_2^{\tau-1}R_2)B = N_1^{\tau-1}R_1B + P_1^\top Q_2N_2^{\tau-1}R_2B,$$

so by Gelfand's Formula and Lemma A.1 we have

$$\begin{aligned} \sigma_{\min}(B_\tau) &= \sigma_{\min}(N_1^{\tau-1}R_1B + P_1^\top Q_2N_2^{\tau-1}R_2B) \\ &\ge \sigma_{\min}(N_1^{\tau-1})\sigma_{\min}(R_1B) - \|P_1^\top Q_2\|\|N_2^{\tau-1}\|\|R_2\|\|B\| \\ &\ge \frac{c\|B\|}{\zeta_{\varepsilon_3}(N_1^{-1})}\left(\frac{|\lambda_k|}{1+\varepsilon_3|\lambda_k|}\right)^{\tau-1} - \frac{\sqrt{2\xi}\zeta_{\varepsilon_2}(N_2)\|B\|}{1-\xi}(|\lambda_{k+1}|+\varepsilon_2)^{\tau-1} \\ &> \frac{c\|B\|}{2\zeta_{\varepsilon_3}(N_1^{-1})}\left(\frac{|\lambda_k|}{1+\varepsilon_3|\lambda_k|}\right)^{\tau-1} \end{aligned}$$

where the last inequality requires

$$\frac{\sqrt{2\xi}\zeta_{\varepsilon_2}(N_2)\zeta_{\varepsilon_3}(N_1^{-1})}{c(1-\xi)}\left(\frac{(|\lambda_{k+1}|+\varepsilon_2)(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1} < \frac{1}{2},$$

or equivalently,

$$\tau > \frac{\log\frac{c(1-\xi)}{2\sqrt{2\xi}\zeta_{\varepsilon_2}(N_2)\zeta_{\varepsilon_3}(N_1^{-1})}}{\log\frac{(|\lambda_{k+1}|+\varepsilon_2)(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}} + 1. \tag{30}$$

Therefore, using Proposition G.6, $\sigma_{\min}(\hat{B}_\tau)$ is lower bounded by

$$\begin{aligned} \sigma_{\min}(\hat{B}_\tau) &\ge \sigma_{\min}(B_\tau) - \|\hat{B}_\tau - B_\tau\| \\ &> \frac{c\|B\|}{2\zeta_{\varepsilon_3}(N_1^{-1})}\left(\frac{|\lambda_k|}{1+\varepsilon_3|\lambda_k|}\right)^{\tau-1} - C_B(|\lambda_1|+\varepsilon_1)^{\tau-1}\delta \\ &> \frac{c\|B\|}{4\zeta_{\varepsilon_3}(N_1^{-1})}\left(\frac{|\lambda_k|}{1+\varepsilon_3|\lambda_k|}\right)^{\tau-1}, \end{aligned}$$

where the last inequality requires

$$\delta < \frac{c\|B\|}{4\zeta_{\varepsilon_3}(N_1^{-1})C_B}\left(\frac{|\lambda_k|}{(1+\varepsilon_3|\lambda_k|)(|\lambda_1|+\varepsilon_1)}\right)^{\tau-1}. \tag{31}$$

This completes the proof. □

Finally, using the above bounds, we can easily upper bound the norm of our controller $\hat{K}_1$.

Proposition G.7. *Under the premises of Theorem 4.1, when (29), (30), (31) and $\delta < \frac{1}{\tau}$ hold,*

$$\|\hat{K}_1\| < C_K \left(\frac{(|\lambda_1| + \varepsilon_1)(1 + \varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1},$$

where $C_K := \frac{4\zeta_{\varepsilon_3}(N_1^{-1})\left(\zeta_{\varepsilon_1}(M_1)(|\lambda_1|+\varepsilon_1)+2\|A\|\zeta_{\varepsilon_1}(A)\right)}{c\|B\|}$.

Proof. Recall that the controller is constructed as $\hat{K}_1 = \hat{B}_\tau^{-1}\hat{M}_1^\tau \hat{P}_1^\top$, so we have

$$\|\hat{K}_1\| \le \|\hat{B}_\tau^{-1}\|\|\hat{M}_1^\tau\| = \frac{\|\hat{M}_1^\tau\|}{\sigma_{\min}(\hat{B}_\tau)},$$

and the bound is merely a combination of Corollary G.2 and Corollary G.5 whenever $\delta < \frac{1}{\tau}$. □

G.4 Proof of Theorem 4.1

Now we are ready to combine the above building blocks and present the complete proof of Theorem 4.1. Note that, with all the bounds established above, the proof structure parallels that of Theorem 4.2, the special case with a symmetric dynamical matrix A.

Proof of Theorem 4.1. The proof is again based on Lemma 5.3. We first guarantee that the diagonal blocks are stable. For the top-left block,

$$\begin{aligned}
\|M_1^\tau + P_1^\top A^{\tau-1}B\hat{K}_1\| &= \|M_1^\tau - B_\tau \hat{B}_\tau^{-1}\hat{M}_1^\tau \hat{P}_1^\top P_1\| \\
&\le \|M_1^\tau - \hat{M}_1^\tau\| + \|(B_\tau - \hat{B}_\tau)\hat{B}_\tau^{-1}\hat{M}_1^\tau\| + \|B_\tau \hat{B}_\tau^{-1}\hat{M}_1^\tau(I - \hat{P}_1^\top P_1)\| \\
&\le \|M_1^\tau - \hat{M}_1^\tau\| + \|B_\tau - \hat{B}_\tau\|\|\hat{K}_1\| + \|B_\tau\|\|\hat{K}_1\|\|I - \hat{P}_1^\top P_1\| \\
&\le 2\tau\|A\|\zeta_{\varepsilon_1}(A)^2(|\lambda_1| + \varepsilon_1)^{\tau-1}\delta \\
&\quad + C_B C_K \left(\frac{(|\lambda_1| + \varepsilon_1)^2(1 + \varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1}\delta \qquad (32) \\
&\quad + \zeta_{\varepsilon_1}(A)\|B\|C_K \left(\frac{(|\lambda_1| + \varepsilon_1)^2(1 + \varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1}\delta \\
&< (C_B C_K + \zeta_{\varepsilon_1}(A)\|B\|C_K + 1)\left(\frac{(|\lambda_1| + \varepsilon_1)^2(1 + \varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1}\delta \qquad (33) \\
&< \frac{1}{2}, \qquad (34)
\end{aligned}$$

where in (32) we apply Propositions G.2, G.6, G.7, and E.1; in (33) we require

$$\frac{1}{\tau}\left(\frac{(|\lambda_1| + \varepsilon_1)^2(1 + \varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1} > 2\|A\|\zeta_{\varepsilon_1}(A)^2; \qquad (35)$$

and in (34) we require

$$\delta < \frac{1}{2(C_B C_K + \zeta_{\varepsilon_1}(A)\|B\|C_K + 1)}\left(\frac{(|\lambda_1| + \varepsilon_1)^2(1 + \varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{-(\tau-1)}. \qquad (36)$$

For the bottom-right block, it is straight-forward to see that

$$\begin{aligned}
\|M_2^\tau + P_2^\top A^{\tau-1}B\hat{K}_1\hat{P}_1^\top P_2\| &\le \|M_2^\tau\| + \|P_2^\top A^{\tau-1}\|\|B\|\|\hat{K}_1\|\|\hat{P}_1^\top P_2\| \\
&\le \zeta_{\varepsilon_2}(M_2)(|\lambda_{k+1}| + \varepsilon_2)^\tau \\
&\quad + \zeta_{\varepsilon_2}(M_2)\|B\|C_K \left(\frac{(|\lambda_1| + \varepsilon_1)(|\lambda_{k+1}| + \varepsilon_2)(1 + \varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1}\delta \\
&< 1
\end{aligned}$$

where the last inequality requires

$$\tau > \frac{\log 1/(4\zeta_{\varepsilon_2}(M_2))}{\log(|\lambda_{k+1}| + \varepsilon_2)}, \tag{37}$$

$$\delta < \frac{1}{4\zeta_{\varepsilon_2}(M_2)\|B\|C_K}\left(\frac{(|\lambda_1|+\varepsilon_1)(|\lambda_{k+1}|+\varepsilon_2)(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{-(\tau-1)}. \tag{38}$$

Now it only suffices to bound the spectral norms of off-diagonal blocks. Note that, by applying Proposition G.7 and Proposition G.1, the top-right block is bounded as

$$\begin{aligned}\|\varDelta_\tau + P_1^\top A^{\tau-1}B\hat{K}_1\hat{P}_1^\top P_2\| &\le \|\varDelta_\tau\| + \|B_\tau\|\|\hat{K}_1\|\|\hat{P}_1^\top P_2\| \\ &< C_\varDelta(|\lambda_1|+\varepsilon_1)^\tau \\ &\quad + \zeta_{\varepsilon_1}(A)\|B\|C_K\left(\frac{(|\lambda_1|+\varepsilon_1)^2(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1}\delta \\ &< (C_\varDelta+1)(|\lambda_1|+\varepsilon_1)^\tau\end{aligned}$$

where the last inequality requires

$$\delta < \frac{(|\lambda_1|+\varepsilon_1)^2}{\zeta_{\varepsilon_1}(A)\|B\|C_K}\left(\frac{(|\lambda_1|+\varepsilon_1)^2(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{-\tau}; \tag{39}$$

and the bottom-left block is bounded as

$$\begin{aligned}\|P_2^\top A^{\tau-1}B\hat{K}_1\hat{P}_1^\top P_1\| &\le \|P_2^\top A^{\tau-1}\|\|B\|\|\hat{K}_1\| \\ &< \zeta_{\varepsilon_2}(M_2)\|B\|C_K\left(\frac{(|\lambda_1|+\varepsilon_1)(|\lambda_{k+1}|+\varepsilon_2)(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1}.\end{aligned}$$

Now, by Lemma 5.3, we can guarantee that

$$\rho(\hat{L}_\tau) \le \frac{1}{2} + \chi(\hat{L}_\tau)\frac{(C_\varDelta+1)\zeta_{\varepsilon_2}(M_2)\|B\|C_K}{|\lambda_1|+\varepsilon_1}\left(\frac{(|\lambda_1|+\varepsilon_1)^2(|\lambda_{k+1}|+\varepsilon_2)(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{\tau-1} < 1,$$

which requires

$$\tau > \frac{\log\frac{2(|\lambda_1|+\varepsilon_1)}{\chi(\hat{L}_\tau)(C_\varDelta+1)\zeta_{\varepsilon_2}(M_2)\|B\|C_K}}{\log\frac{(|\lambda_1|+\varepsilon_1)^2(|\lambda_{k+1}|+\varepsilon_2)(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}}. \tag{40}$$

Note that the above constraint makes sense only if $|\lambda_1|^2|\lambda_{k+1}| < |\lambda_k|$.

So far, we shall recollect all the constraints we need on the parameters $\tau, \alpha, \delta, \gamma$ and ω. To start with, all constraints on τ (see (30), (35), (37) and (40)) can be summarized as

$$\begin{aligned}\tau > \max\Bigg\{&\frac{\log\frac{c(1-\xi)}{2\sqrt{2\xi}\zeta_{\varepsilon_2}(N_2)\zeta_{\varepsilon_3}(N_1^{-1})}}{\log\frac{(|\lambda_{k+1}|+\varepsilon_2)(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}} + 1, \frac{\log 1/(4\zeta_{\varepsilon_2}(M_2))}{\log(|\lambda_{k+1}|+\varepsilon_2)}, \frac{\log\frac{2(|\lambda_1|+\varepsilon_1)}{\chi(\hat{L}_\tau)(C_\varDelta+1)\zeta_{\varepsilon_2}(M_2)\|B\|C_K}}{\log\frac{(|\lambda_1|+\varepsilon_1)^2(|\lambda_{k+1}|+\varepsilon_2)(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}}, \\ &-\frac{1}{\log\frac{(|\lambda_1|+\varepsilon_1)^2(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}}W_{-1}\left(-\frac{\log\frac{(|\lambda_1|+\varepsilon_1)^2(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}}{2\|A\|\zeta_{\varepsilon_1}(A)^2\frac{(|\lambda_1|+\varepsilon_1)^2(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}}\right)\Bigg\},\end{aligned}$$

where W_{-1} denotes the non-principle branch of the Lambert-W function. Here we utilize the fact that, for $x > \frac{1}{\log a}$, $y = \frac{a^x}{x}$ is monotone increasing with inverse function $x = -\frac{1}{\log a}W_{-1}\big(-\frac{\log a}{y}\big)$, which can be upper bounded by Theorem 1 in [58] as

$$-\frac{1}{\log a}W_{-1}\left(-\frac{\log a}{y}\right) < \frac{\log y - \log\log a + \sqrt{2(\log y - \log\log a)}}{\log a} < \frac{3(\log y - \log\log a)}{\log a}.$$

By gathering different constants, we have

$$\tau > \frac{\log\frac{\sqrt{\xi}}{1-\xi} + \log\frac{1}{c} + \log\chi(\hat{L}_\tau) + 5\log\bar{\zeta} + \log\frac{\|A\|}{|\lambda_1|-|\lambda_{k+1}|} + C_\tau}{\log\frac{|\lambda_k|}{|\lambda_1|^2|\lambda_{k+1}|}} = O(1), \tag{41}$$

where we define $\bar{\zeta} := \max\{\zeta_{\varepsilon_1}(A), \zeta_{\varepsilon_2}(M_2), \zeta_{\varepsilon_2}(N_2), \zeta_{\varepsilon_3}(N_1^{-1})\}$, and C_τ is a numerical constant. Note that we have to guarantee the denominator to be positive, which gives rise to the additional assumption $|\lambda_1|^2|\lambda_{k+1}| < |\lambda_k|$. Meanwhile, for any $\ell \in \mathbb{N}$, we shall select γ such that

$$\gamma = O(k^{-\ell}), \quad \gamma < \min\left\{\frac{1}{2}, \frac{1}{\sqrt{2/(\sigma_{\min}(R_1)k)} + 1}\right\}, \tag{42}$$

and select α such that (see (23), and we have already guaranteed $\gamma_\omega > 2$ in (29))

$$\alpha < \frac{\frac{2}{3}\sigma_{\min}(M_1) - \frac{\gamma}{1-\xi}\|A\|}{(1 + \frac{\sqrt{2\xi}}{1-\xi} + \frac{\gamma}{1-\xi})\|B\|} = O(1). \tag{43}$$

Now constraints on δ (see (31), (36), (38) and (39)) can be summarized as

$$\delta < \min\left\{\frac{c\|B\|}{4\zeta_{\varepsilon_3}(N_1^{-1})C_B}\left(\frac{|\lambda_k|}{(1+\varepsilon_3|\lambda_k|)(|\lambda_1|+\varepsilon_1)}\right)^{\tau-1},\right.$$
$$\frac{1}{2(C_BC_K + \zeta_{\varepsilon_1}(A)\|B\|C_K + 1)}\left(\frac{(|\lambda_1|+\varepsilon_1)^2(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{-(\tau-1)},$$
$$\frac{1}{4\zeta_{\varepsilon_2}(M_2)\|B\|C_K}\left(\frac{(|\lambda_1|+\varepsilon_1)(|\lambda_{k+1}|+\varepsilon_2)(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{-(\tau-1)},$$
$$\left.\frac{(|\lambda_1|+\varepsilon_1)^2}{\zeta_{\varepsilon_1}(A)\|B\|C_K}\left(\frac{(|\lambda_1|+\varepsilon_1)^2(1+\varepsilon_3|\lambda_k|)}{|\lambda_k|}\right)^{-\tau}\right\},$$

which can be simplified to (C_δ is a constant collecting minor factors)

$$\delta < \frac{C_\delta \alpha c}{\sqrt{k}\bar{\zeta}^3(\|A\| + \|B\|)}|\lambda_1|^{-2\tau} = O(k^{-1/2}|\lambda_1|^{-2\tau}), \tag{44}$$

or we can rewrite the bound equivalently in terms of t_0 (recall (10) in Appendix E) as

$$t_0 > \frac{\log(n^2\binom{n}{k}) + \log k + \log\kappa_{\mathrm{e}}(A) + 2\tau\log|\lambda_1| + 3\log\bar{\zeta} + \log(\|A\| + \|B\|) + \log\frac{\sqrt{2}}{C_\delta\alpha}}{2\log\frac{|\lambda_k|}{|\lambda_{k+1}|}}$$
$$= O\left(\frac{2\tau\log|\lambda_1| + k\log n + \log\kappa_{\mathrm{e}}(A)}{\log\frac{|\lambda_k|}{|\lambda_{k+1}|}}\right), \tag{45}$$

Finally, we select ω such that (see (29), and note that $C_\gamma = O(\gamma) = O(k^{-\ell})$)

$$\omega > \max\left\{\frac{\log\frac{2}{C_\gamma}}{\log\frac{|\lambda_k|}{(1+\varepsilon_3|\lambda_k|)(|\lambda_{k+1}|+\varepsilon_2)}}, \frac{\log\frac{2C_\Delta}{C_\gamma\delta} + \tau\log(|\lambda_1|+\varepsilon_1)}{\log\frac{|\lambda_k|}{(1+\varepsilon_3|\lambda_k|)(|\lambda_{k+1}|+\varepsilon_2)}}\right\},$$

which can be reorganized as

$$\omega > \frac{\log\frac{1}{C_\gamma} + \log\frac{\sqrt{\xi}}{1-\xi} + 2\log\bar{\zeta} + \log\frac{\|A\|}{|\lambda_1| - |\lambda_{k+1}|} + \log\frac{1}{\delta} + C_\omega}{\log\frac{|\lambda_k|}{|\lambda_{k+1}|}} = O(\ell\log k). \tag{46}$$

Note that here $\varepsilon_1, \varepsilon_2, \varepsilon_3$ are taken to be small enough, so that

$$|\lambda_{k+1}| + \varepsilon_2 < 1, \quad (|\lambda_1| + \varepsilon_1)^2(|\lambda_{k+1}| + \varepsilon_2) < \frac{|\lambda_k|}{1 + \varepsilon_3|\lambda_k|}, \quad \varepsilon_3|\lambda_k| < 1. \tag{47}$$

Also, the probability of sampling an admissible x_0 is $1 - \theta(\gamma) = 1 - O(k^{-\ell})$ by the union bound.

Finally, by (41), (45) and (46), we conclude that Algorithm 1 terminates within

$$t_0 + k(1 + \omega + \tau) > \frac{1}{2\log\frac{|\lambda_k|}{|\lambda_{k+1}|}}\left(\underbrace{\log(n^2\tbinom{n}{k})}_{O(k\log n)} + \underbrace{2k\log\frac{1}{C_\gamma}}_{O(k\log k)} + \log k\right) + k$$

$$
\begin{aligned}
&+\frac{\log \kappa_{\mathrm{e}}(A)+2\tau \log|\lambda_1|+3\log\bar{\zeta}+\log(\|A\|+\|B\|)+\log\frac{\sqrt{2}}{C_\delta \alpha}}{2\log\frac{|\lambda_k|}{|\lambda_{k+1}|}}\\
&+\frac{k\left(\log\frac{\sqrt{\xi}}{1-\xi}+2\log\bar{\zeta}+\log\frac{\|A\|}{|\lambda_1|-|\lambda_{k+1}|}+\log\frac{1}{\delta}+C_\omega\right)}{\log\frac{|\lambda_k|}{|\lambda_{k+1}|}}\\
&+\frac{k\left(\log\frac{\sqrt{\xi}}{1-\xi}+\log\frac{1}{c}+\log\chi(\hat{L}_\tau)+5\log\bar{\zeta}+\log\frac{\|A\|}{|\lambda_1|-|\lambda_{k+1}|}+C_\tau\right)}{\log\frac{|\lambda_k|}{|\lambda_1|^2|\lambda_{k+1}|}}\\
=&\ O(k\log n),
\end{aligned}
$$

time steps, which completes the proof. □

For the convenience of readers, we provide a table summarizing all constants appearing in the bound. We would like to point out that, the performance bound here is presented with *instance-specific* parameters to obtain the tightest bound for each system instance. Still, if an *instance-independent* bound is needed for a family of systems, where all parameters are uniformly bounded (from the correct direction), we can simply replace those parameters with their bounds in equations (41) through (46) to obtain an instance-independent bound for that family.

Table 1: Lists of parameters and constants appearing in the bound.

(a) Algorithmic parameters (introduced in Algorithm algorithm 1).

Constant	Appearance	Explanation
t_0	Stage 1	t_0 initialization steps to separate unstable components
ω	Stage 3	ω heat-up steps in each iteration of learning B_τ
α	Stage 3	$\|u_{t_i}\| = \alpha\|x_{t_i}\|$ to keep non-negligible unstable component
τ	Stage 4	τ steps between consecutive control inputs are injected

(b) System parameters (as functions of dynamical matrices).

Constant	Definition	Explanation
λ_i	Assumption 4.1	(complex) eigenvalue of A with i^{th} largest modulus
$\|A\|, \|B\|$	Notation	2-norm of dynamical matrices A and B
c	Assumption 4.3′	c-effective controllability in the unstable subspace E_{u}, i.e., $\sigma_{\min}(R_1 B) > c\|B\|$
ξ	Definition 3.1	$E_{\mathrm{u}}^{\perp}$ and E_{s} are ξ-close subspaces, i.e., $\sigma_{\min}(P_2^\top Q_2) > 1-\xi$
$\chi(\cdot)$	Lemma D.1	perturbation constant for 2-by-2 block diagonal matrices
$\zeta_\varepsilon(\cdot)$	Lemma G.1	Gelfand constant for the norm of matrix exponents
$\kappa_{\mathrm{e}}(\cdot)$	Notation	condition number of the matrix with eigenvectors as columns

(c) Shorthand notations (introduced in proofs).

Constant	Definition	Explanation
C_Δ	Proposition G.1	$C_\Delta := \zeta_{\varepsilon_1}(M_1)\zeta_{\varepsilon_2}(M_2)\frac{(2-\xi)\sqrt{2\xi}\|A\|}{1-\xi}\frac{2\|\lambda_{k+1}\|}{\|\lambda_1\|+\varepsilon_1-\|\lambda_{k+1}\|-\varepsilon_2}$
C_γ	Proposition G.3	$C_\gamma := \frac{1}{(1+\frac{1}{\gamma})\zeta_{\varepsilon_3}(N_1^{-1})\zeta_{\varepsilon_2}(N_2)\|R_2\|}$ (γ is taken according to (42))
C_B	Proposition G.6	$C_B := \frac{2\sqrt{k}\zeta_{\varepsilon_1}(A)^2\left((2\tau+2)\|A\|+\|B\|\right)}{\alpha}$
C_K	Proposition G.7	$C_K := \frac{4\zeta_{\varepsilon_3}(N_1^{-1})\left(\zeta_{\varepsilon_1}(M_1)(\|\lambda_1\|+\varepsilon_1)+2\|A\|\zeta_{\varepsilon_1}(A)\right)}{c\|B\|}$
$\bar{\zeta}$	below (41)	$\bar{\zeta} := \max\{\zeta_{\varepsilon_1}(A), \zeta_{\varepsilon_2}(M_2), \zeta_{\varepsilon_2}(N_2), \zeta_{\varepsilon_3}(N_1^{-1})\}$
$C_\tau, C_\delta, C_\omega$	(41), (44), (46)	collection of numerical constants in (41), (44), (46)

H An Illustrative Example with Additive Noise

Finally, we include an illustrative experiment that shows the performance of our LTS_0 algorithm.

Settings. We evaluate the algorithm in LTI systems with additive noise

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad \text{where } w_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_w^2 I).$$

Here σ_w characterizes the variance (and thus the magnitude) of the noise. The dynamical matrices are randomly generated: A is generated based on its eigen-decomposition $A = V\Lambda V^{-1}$, where the eigenvalues $\Lambda = \operatorname{diag}(\lambda_1, \cdots, \lambda_n)$ are randomly generated by selecting $\lambda_{1:k} \sim \mathcal{U}(1, \lambda_{\max})$ and $\lambda_{k+1:n} \sim \frac{|\lambda_k|}{|\lambda_1|^2} \cdot \mathcal{U}(-1, 1)$ (to ensure $|\lambda_1|^2 |\lambda_{k+1}| < |\lambda_k|$), and the eigenvectors $V = [v_1, \cdots, v_n]$ are generated by random perturbation to a random orthogonal matrix (to avoid tiny ξ); meanwhile, B is generated by random sampling i.i.d. entries from $\mathcal{U}(0, 1)$. For comparability and reproducibility, throughout the experiment we set $k = 3$ and use 0 as the initial random seed.

To compare the performance in different settings, 30 data points are collected for each pair of σ_w and n. It is observed that our algorithm might cause numerical instability issues (e.g., $\operatorname{cond}(D^\top D)$ could be large), so we simply ignore such cases and repeat until 30 data points are collected. The parameters of the algorithm are determined in an adaptive way that minimizes the number of running steps: we search for the minimum t_0 that yields estimation error smaller than δ, search for the minimum τ such that $K = B_\tau^{-1} M_1^\tau P_1^\top$ stabilizes the system, and the ω heat-up steps in Stage 3 could be ended earlier if we already observe $\|\hat{P}_1^\top x\| / \|x\|$ larger than a certain threshold.

Our experimental results are presented in Figure 1 below.

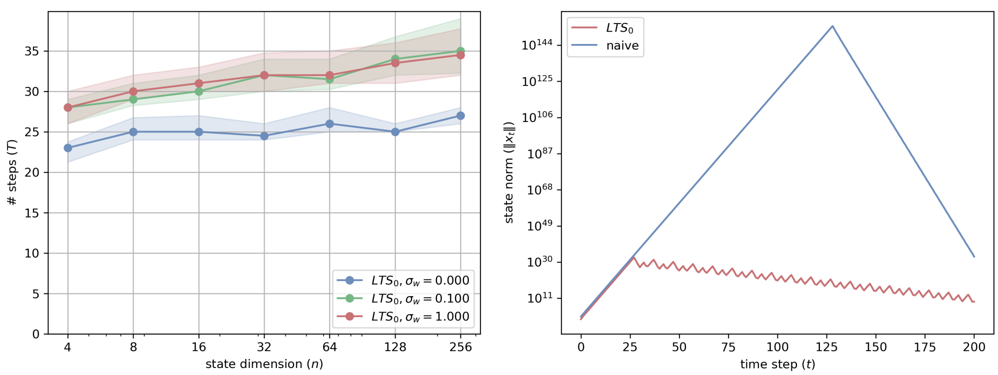

(a) Running steps of LTS$_0$

(b) State norms along one trajectory

Figure 1: Experimental results. *In (a), the line shows the median of running steps, and the shadow marks the range between upper and lower quartiles (the horizontal axis is in log scale). In (b), the trajectories of our algorithm and the naive approach are compared in a randomly-generated system with* $n = 128$ *and* $\sigma_w = 0$ *(the vertical axis is in log scale).*

Performance under different n and σ_w. Figure 1a shows the number of running steps of LTS$_0$ that is needed to learn a stabilizing controller. It is evident that the number of running steps grow almost linearly with regard to $\log n$, which is in accordance with Theorem 4.1.

As for the effect of noise, it is observed that the algorithm needs more steps in systems with noise than in those without noise; nevertheless, the magnitude of noise does not have much influence on the number of running steps. This is also reasonable since the increase is mainly attributed to t_0 — it takes more initial steps to push the state close enough to E_u, such that the estimation error of P_1 drops to acceptable level; however, as the E_u-component grows exponentially fast over time while w_t is i.i.d., the magnitude of noise only plays a minor role in the increase. Noise becomes negligible in later stages due to the disproportionate magnitudes of states and noise.

Analysis on comparison of trajectories. In Figure 1b we study an exemplary trajectory of our LTS$_0$ algorithm, and compare it against that of the naive approach, which first identifies the system and then designs a controller to nullify the unstable eigenvalues by standard pole-placement method. It is evident that our algorithm needs significantly fewer steps, and thus induces far smaller state norms, to learn a controller that effectively stabilizes the system. It is also observed that our controller decreases state norm in a zig-zag manner, which is due to the τ-hop design our algorithm adopts. Nevertheless, a potential drawback of our controller design is that the spectral radius of the controlled system is larger (since we cannot precisely nullify all unstable eigenvalues), resulting in a slower stabilizing rate than the naive approach (compare the decreasing parts of the curves).