

Safely Learning to Control the Linear Quadratic Regulator

*Sarah Dean
Stephen Tu
Nikolai Matni
Benjamin Recht*

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2019-41

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-41.html>

May 15, 2019



Copyright © 2019, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

We thank Claire Tomlin, Francesco Borrelli, and the members of the MPC Lab at UC Berkeley for their helpful comments and feedback.
SD is supported by an NSF Graduate Research Fellowship under Grant No. DGE 1752814.
ST is supported by a Google PhD fellowship.
BR is generously supported in part by ONR awards N00014-17-1-2191, N00014-17-1-2401, and N00014-18-1-2833, the DARPA Assured Autonomy (FA8750-18-C-0101) and Lagrange (W911NF-16-1-0552) programs, and an Amazon AWS AI Research Award.

Safely Learning to Control the Linear Quadratic Regulator

by Sarah Dean

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

Professor Ben Recht
Research Advisor

(Date)

* * * * *

Professor Claire Tomlin
Second Reader

(Date)

Safely Learning to Control the Constrained Linear Quadratic Regulator

Sarah Dean, Stephen Tu, Nikolai Matni, and Benjamin Recht*

May 3, 2019

Abstract

We study the constrained linear quadratic regulator with unknown dynamics, addressing the tension between safety and exploration in data-driven control techniques. We present a framework which allows for system identification through persistent excitation, while maintaining safety by guaranteeing the satisfaction of state and input constraints. This framework involves a novel method for synthesizing robust constraint-satisfying feedback controllers, leveraging newly developed tools from *system level synthesis*. We connect statistical results with cost sub-optimality bounds to give non-asymptotic guarantees on both estimation and controller performance.

1 Introduction

Data-driven design has considerable potential in contemporary control systems where precise modeling of the dynamics is intractable, whether due to complex large-scale interactions or nonlinearities resulting from contact forces. However, a large hurdle in the way of practical deployment is the question of maintaining safe operating conditions during the learning process, and furthermore guaranteeing safe execution using learned components.

Motivated by this issue, we study the data-driven design of a controller for the *constrained* Linear Quadratic Regulator (LQR) problem. In constrained LQR, we design a controller for a potentially unknown linear dynamical system that minimizes a given quadratic cost, subject to the additional requirement that both the state and input stay within a specified safe region. This is a problem that has received much attention within the model predictive control (MPC) community.

For the LQR problem with no constraints, a natural method of exploration for learning the dynamics is to excite the system by injecting white noise. When safety is not an issue, this method is effective and recently Dean et al. [1] provide an end-to-end sample complexity

*S. Dean, S. Tu, N. Matni, and B. Recht are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, 94709 USA (email: dean.sarah@berkeley.edu, stephent@berkeley.edu, nmatni@berkeley.edu, brecht@berkeley.edu)

on this “identify-then-control” scheme. However, this method of fails to consider safety, and the injected white noise may lead to constraint violation.

In this paper, we directly address the tension between exploration for learning and safety, which are fundamentally at odds. We do this by synthesizing a controller which simultaneously excites and regulates the system; we propose to learn by additively injecting bounded noise to the control inputs computed by a safe controller. By leveraging the recently developed system level synthesis (SLS) framework for control design [2], we give a computationally tractable algorithm which returns a controller that (a) guarantees the closed loop system remains within the specified constraint set and (b) ensures that enough noise can be injected into the system to obtain a statistical guarantee on learning. To the best of our knowledge, our algorithm is the first to simultaneously achieve both objectives. Furthermore, the controller synthesis is solved by a convex optimization problem whose feasibility is a certificate of safety, and no considerations of robust invariant sets are required.

Our second contribution is to provide a sub-optimality bound on control performance for constrained LQR. Using the same SLS framework, we quantify the excess cost incurred by playing a controller designed on the uncertain dynamics obtained from learning. The sub-optimality depends on both the size of the uncertainty sets and a type of constraint robustness cost gap of the optimal constrained controller for the true system. This allows us to provide the first end-to-end sample complexity guarantee for the control of constrained systems.

Estimation and control of the unconstrained LQR problem has been studied in the non-asymptotic setting [3, 1]. However, the identification schemes rely on pure excitation and system restarts, which is not suitable for the constrained setting. On the other hand, the online learning literature simultaneously considers learning and control, where strategies are based on *optimism in the face of uncertainty* [4] or *Thompson Sampling* [5]. These approaches guarantee system estimation only up to optimal closed-loop equivalence, and do not consider safety. Building on a statistical result by Simchowitz et al. [6] which allows for non-asymptotic guarantees on parameter estimation from a single trajectory of a linear system, Dean et al. [7] provide a robust online method that guarantees parameter estimation and stability throughout.

The design of controllers that guarantee robust constraint satisfaction has long been considered in the context of model predictive control [8], including methods that model uncertainty in the dynamics directly [9], or model it as a bounded state disturbance for computational efficiency [10, 11]. Strategies for incorporating estimation of the dynamics include experiment-design inspired costs [12], decoupling learning from constraint satisfaction [13], and set-membership methods rather than parameter estimation [14, 15]. Due to the receding horizon nature of model predictive controllers, this literature relies on set invariance theory for infinite horizon guarantees [16]. Our framework considers the infinite horizon problem directly, and therefore we do not require computation of invariant sets.

Finally, the machine-learning community has begun to consider safety in reinforcement learning, where much work positions itself as being for general dynamical systems in lieu of providing statistical guarantees [17, 18, 19, 20]. Some works assume the existence of

an initial safe controller for learning [21], and robust MPC methods have been proposed to modify potentially unsafe learning inputs [22]. Our framework gives an alternative procedure for designing such a controller using coarse system estimates. Most similar to this work is that of Lu et al. [23], who propose a method to allow excitation on top of a safe controller, but consider only finite-time safety and require non-convex optimization to obtain formal guarantees. We fix an underlying linear dynamical system with full state observation,

$$x_{k+1} = A_\star x_k + B_\star u_k + w_k, \quad (1.1)$$

where we have initial condition $x_0 \in \mathbb{R}^n$, sequence of inputs $\{u_0, u_1, \dots\} \subseteq \mathbb{R}^d$, and disturbance process $\{w_0, w_1, \dots\} \subseteq \mathbb{R}^n$. Later, we will specify assumptions on the disturbances, including independence over time, bounded second moment, and bounded ℓ_∞ norm. The dynamics matrices (A_\star, B_\star) are unknown. We denote the errors on estimates of the system (\hat{A}, \hat{B}) in terms of the $\ell_p \rightarrow \ell_p$ matrix operator norm as

$$\|A_\star - \hat{A}\|_p := \|\Delta_A\|_p := \varepsilon_{A,p}, \quad \|B_\star - \hat{B}\|_p := \|\Delta_B\|_p := \varepsilon_{B,p},$$

In this work, we will focus on the errors for $p = 2$ and $p = \infty$.

Furthermore, we assume some prior knowledge on the system dynamics, in the form of initial estimates (\hat{A}_0, \hat{B}_0) and uncertainty measures $(\varepsilon_{A,p}^0, \varepsilon_{B,p}^0)$. Without such knowledge, guaranteeing safety in any form would be impossible. We note that the initial estimates may be coarse grained, and the goal of the learning procedure will be to refine this uncertainty prior to optimal control design.

The constrained LQR optimal control problem seeks to minimize the expected quadratic cost, subject to constraints on the state and input. Before further detailing this objective, we develop notation and background on system level synthesis.

1.1 System Level Synthesis

Many approaches to optimal control for systems with constraints involve receding horizon control, where an *open loop* finite-time trajectory is computed at each timestep; indeed, parameterizing optimal control problems by a state feedback controller generally leads to nonconvex optimization. As a motivating example, consider the static feedback law $u_k = Kx_k$ applied to the linear system (1.1). Then we can write

$$x_k = \sum_{t=1}^k (A_\star + B_\star K)^t w_{k-t} + (A_\star + B_\star K)^k x_0,$$

and a similar expression for u_k . Then it is clear that convex constraints on x_k and u_k will be non-convex in K . Instead, we can parametrize the problem in terms of convolution with the closed-loop system response,

$$x_k = \sum_{t=1}^{k+1} \Phi_x(t) w_{k-t}, \quad u_k = \sum_{t=1}^{k+1} \Phi_u(t) w_{k-t} \quad (1.2)$$

where we defined $w_{-1} = x_0$ the fixed initial condition. We note that the expression above is valid for any linear dynamic controller, i.e. any controller which is a linear function of the state and its history. Since the relation in (1.2) is linear, convex constraints on state and input translate to convex constraints on the system response elements. The *system level synthesis* (SLS) framework shows that for any elements $\{\Phi_x(t), \Phi_u(t)\}$ constrained to obey, for all $k \geq 1$,

$$\Phi_x(k+1) = A_\star \Phi_x(k) + B_\star \Phi_u(k), \quad \Phi_x(1) = I,$$

there exists a feedback controller that achieves the desired system responses (1.2). The state-feedback parameterization result in Theorem 1 of Wang et al. [2] formalizes this principle, and therefore any optimal control problem over linear systems can be cast as a constrained optimization problem over system response elements. We remark that similar observations have been used for constrained state feedback control in the finite horizon setting [11].

In what follows, we use boldface letters to denote transfer functions and signals, e.g. $\Phi_x(z) = \sum_{k=1}^{\infty} \Phi_x(k) z^{-k}$ and $\mathbf{x}(z) = \sum_{k=0}^{\infty} x_k z^{-k}$. Under this notation, the affine constraints can be rewritten as

$$\begin{bmatrix} zI - A_\star & -B_\star \end{bmatrix} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = I,$$

and the corresponding control law $\mathbf{u} = \mathbf{K}\mathbf{x}$ is given by $\mathbf{K} = \Phi_u \Phi_x^{-1}$.

1.2 Notation

In this paper, we restrict our attention to the function space \mathcal{RH}_∞ , consisting of discrete-time stable matrix-valued transfer functions. We use $\frac{1}{z}\mathcal{RH}_\infty$ to denote the set of transfer functions \mathbf{G} such that $z\mathbf{G} \in \mathcal{RH}_\infty$. We further define

$$\mathcal{RH}_\infty(C, \rho) := \left\{ \mathbf{M} = \sum_{k=0}^{\infty} M(k) z^{-k} \mid \|M(k)\|_2 \leq C\rho^k, \quad k = 0, 1, 2, \dots \right\}$$

for transfer functions that satisfy a certain decay rate in the spectral norm of their impulse response elements.

When working with transfer functions and signals, we will denote the coefficient of the term of degree k as $\mathbf{G}[k] = G(k)$ and $\mathbf{x}[k] = x_k$. We will also denote $G[k:1]$ as the block row vector of system response elements of \mathbf{G}

$$G[k:1] = [G(k) \quad \dots \quad G(1)].$$

As is standard, we let $\|x\|_p$ denote the ℓ_p -norm of a vector x . For a matrix M , we let $\|M\|_p$ denote its $\ell_p \rightarrow \ell_p$ operator norm. We will consider the \mathcal{H}_2 , \mathcal{H}_∞ , and \mathcal{L}_1 norms, which are infinite horizon analogs of the Frobenius, spectral, and $\ell_\infty \rightarrow \ell_\infty$ operator norms of a matrix, respectively:

$$\|\mathbf{M}\|_{\mathcal{H}_2} = \sqrt{\sum_{k=0}^{\infty} \|M(k)\|_F^2}, \quad \|\mathbf{M}\|_{\mathcal{H}_\infty} = \sup_{\|\mathbf{w}\|_2=1} \|\mathbf{M}\mathbf{w}\|_2, \quad \|\mathbf{M}\|_{\mathcal{L}_1} = \sup_{\|\mathbf{w}\|_\infty=1} \|\mathbf{M}\mathbf{w}\|_\infty$$

Finally, for two numbers a, b , we let $a \lesssim b$ (resp. $a \gtrsim b$) denote that there exists an absolute constant $C > 0$ such that $a \leq Cb$ (resp. $a \geq Cb$).

1.3 Optimal Control Problem

We now describe the constrained optimal control problem that we would want to solve given perfect knowledge of the system dynamics. First, we consider the expected infinite horizon quadratic cost for the system (A_\star, B_\star) in feedback with \mathbf{K} :

$$J(A_\star, B_\star, \mathbf{K})^2 := \frac{1}{\sigma^2} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=0}^{T-1} \mathbb{E}_w [x_k^\top Q x_k + u_k^\top R u_k].$$

We assume that $\{w_k\}$ is any distribution that satisfies $\mathbb{E}[w_k] = 0$ and $\mathbb{E}[w_k w_k^\top] = \sigma^2 I$ and is independent across time, i.e., $w_k \perp w_\ell$ for $\ell \neq k$. Note that any distribution satisfying these constraints induces the same expected quadratic cost, so it is unnecessary to specify a specific distribution.

Next, we consider polytopic constraint sets of the form

$$\mathcal{X} := \{x : F_x x \leq b_x\}, \quad \mathcal{U} := \{u : F_u u \leq b_u\}.$$

We will constrain the state and input trajectories to lie within these sets for any possible disturbance sequence $\{w_k\}$ satisfying $\|w_k\|_\infty \leq \sigma_w$ for all $k \geq 0$. Putting the cost and the constraints together, the optimal control problem that acts as our baseline is:

$$\begin{aligned} \min_{\mathbf{u} \in \mathcal{K}} \quad & J(A_\star, B_\star, \mathbf{K}) \\ \text{s.t. } \quad & x_0 \text{ fixed}; \quad F_x x_k \leq b_x, \quad F_u u_k \leq b_u \\ & \forall k, \forall \{w_k : \|w_k\|_\infty \leq \sigma_w\}. \end{aligned} \tag{1.3}$$

Above, we let the set \mathcal{K} enumerate all inputs that result from linear dynamic stabilizing feedback controllers for (A_\star, B_\star) of the form $\mathbf{u} = \mathbf{K}\mathbf{x}$. This is made possible by the system level synthesis framework described in the previous section.

As we show in the following section, the optimal control problem given in (1.3) is a convex, but infinite-dimensional problem. It is an idealized baseline to compare our actual solutions to; our sub-optimality guarantees will be with respect to the optimal cost achieved by this problem. It is a relevant baseline, since it optimizes for average case performance but ensures safety for the worst-case behavior, consistent with MPC literature [10, 24]. We remark that an alternative to (1.3) is to replace the worst case constraint behavior with probabilistic chance constraints [25]. We do not work with chance constraints because they are generally difficult to directly enforce on an infinite horizon; arguments around recursive feasibility using robust invariant sets are common in the MPC literature to deal with this issue.

2 Constraint-Satisfying Control

We begin by formulating a method for robustly operating a system while maintaining safety. First we describe a system level synthesis approach to the constrained LQR problem, and then we propose a modification which makes it robust to uncertainties in system dynamics. Finally, we outline a reduction to tractable synthesis via a finite-dimensional optimization problem over a finite impulse response (FIR) approximation.

2.1 A System Level Approach

Using the SLS formulation, we define an optimization problem that solves the optimal control problem presented in the section above.

Proposition 2.1. *The following convex optimization problem solves the optimal control problem (1.3).*

$$\begin{aligned} \min_{\Phi_x, \Phi_u \in \frac{1}{z}\mathcal{RH}_\infty} \quad & \left\| \begin{bmatrix} Q^{1/2} & \\ & R^{1/2} \end{bmatrix} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} \right\|_{\mathcal{H}_2} \\ \text{s.t.} \quad & [zI - A_\star \quad -B_\star] \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = I, \\ & G_x(\Phi_x; k) \leq b_x, \quad G_u(\Phi_u; k) \leq b_u \quad \forall k \geq 1. \end{aligned} \quad (2.1)$$

where the elements of the constraint functions are defined as

$$\begin{aligned} G_x(\Phi_x; k)_j &:= F_{x,j}^\top \Phi_x(k+1)x_0 + \sigma_w \|F_{x,j}^\top \Phi_x[k:1]\|_1, \\ G_u(\Phi_u; k)_j &:= F_{u,j}^\top \Phi_u(k+1)x_0 + \sigma_w \|F_{u,j}^\top \Phi_u[k:1]\|_1, \end{aligned}$$

with j indexing the rows of F_x and F_u and entries of G_x and G_u .

Examining the form of the optimization problem above, we see that the expected quadratic cost transforms into a system \mathcal{H}_2 norm, while the worst-case polytopic constraints on state and input become closed-form polytopic constraints on the system response. For a controller $\mathbf{K} = \Phi_u(\Phi_x)^{-1}$, note that the LQR cost is equivalent to

$$J(A_\star, B_\star, \mathbf{K}) = \left\| \begin{bmatrix} Q^{1/2} & \\ & R^{1/2} \end{bmatrix} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} \right\|_{\mathcal{H}_2}.$$

Lastly, we remark here that the feasibility of the convex synthesis problem in (2.1) for an initial condition x_0 implies that x_0 is a member of a robust control invariant set.

Proof. By the state-feedback parameterization result in Theorem 1 of [2], the SLS parametrization encompasses all internally stabilizing state-feedback controllers acting on the true system (A_\star, B_\star) . Thus, it is necessary only to show that the optimization problem in (2.1) is consistent with that of (1.3) under the system level parametrization. The equivalence between the

LQR cost and the \mathcal{H}_2 system norm is standard and we refer to the Appendix of [1], which presents this reformulation in terms of system responses for Gaussian process noise. The argument applies unchanged to any distribution satisfying our assumptions.

Therefore, it remains to consider the inequality constraints. Because the constraints must be satisfied robustly, it is equivalent to consider

$$\max_{\{w_k\}} F_x \sum_{t=1}^{k+1} \Phi_x(t) w_{k-t} \leq b_x \iff F_x \Phi_x(k+1) x_0 + \max_{\{w_k\}} F_x \sum_{t=1}^k \Phi_x(t) w_{k-t} \leq b_x .$$

Then considering elements in the second term with each j indexing the rows of F_x ,

$$\max_{\{w_k\}} F_{x,j}^\top \sum_{t=1}^k \Phi_x(t) w_{k-t} = \sum_{t=1}^k \max_{\|w\|_\infty \leq \sigma_w} F_{x,j}^\top \Phi_x(t) w = \sum_{t=1}^k \sigma_w \|F_{x,j}^\top \Phi_x(t)\|_1 = \sigma_w \|F_{x,j}^\top \Phi_x[k:1]\|_1 .$$

Thus the inequality constraint on the function $G_x(\Phi_x; k)$ is an equivalent condition. A similar computation holds for the input constraint. \square

We note the appearance of the row-wise ℓ_1 norm over the multiplication of the the system response elements with the constraint matrix. The expression can be understood as an analog to the $\ell_\infty \rightarrow \ell_\infty$ operator norm, mediated by the shape of the polytope. In the next section when we introduce robustness to system dynamics, the \mathcal{L}_1 system norm will come into play for similar reasons.

2.2 Robust Control

We are further motivated to reformulate the optimal control problem in terms of system responses so that we can transparently consider uncertainties in the dynamics. We are interested in controller synthesis under model errors, where only nominal estimates of the system are known. Suppose that the system responses $\hat{\Phi}_x, \hat{\Phi}_u$ are designed on the estimated system. Then

$$\begin{bmatrix} zI - \hat{A} & -\hat{B} \end{bmatrix} \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u \end{bmatrix} = I, \text{ and define } \hat{\Delta} := \begin{bmatrix} \Delta_A & \Delta_B \end{bmatrix} \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u \end{bmatrix} .$$

The model mismatch impacts the closed-loop behavior of the true system in feedback with the controller $\hat{\mathbf{K}} = \hat{\Phi}_u(\hat{\Phi}_x)^{-1}$ in a transparent way. Supposing that $\|\hat{\Delta}\|_{\mathcal{H}_\infty} < 1$, it achieves the system response and bounded cost:

$$\begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u \end{bmatrix} (I + \hat{\Delta})^{-1}, \quad J(A_\star, B_\star, \hat{\mathbf{K}}) \leq \frac{1}{1 - \|\hat{\Delta}\|_{\mathcal{H}_\infty}} J(\hat{A}, \hat{B}, \hat{\mathbf{K}}) . \quad (2.2)$$

This follows from the robust stability result in Theorem 2 of Matni et al. [26] and the submultiplicativity of the \mathcal{H}_2 and \mathcal{H}_∞ norms. Note that $\|\hat{\Delta}\| < 1$ is a sufficient condition for the existence of the inverse $(I + \hat{\Delta})^{-1}$ for any induced norm $\|\cdot\|$ by the small gain theorem.

This expression provides an upper bound on the cost achieved for a controller designed system estimates which depends only on the system estimates and the size of the error system $\hat{\Delta}$. Motivated by this bound, consider the following robust optimization problem:

$$\begin{aligned} \hat{J}(\gamma, \tau) := & \min_{\Phi_x, \Phi_u \in \frac{1}{z} \mathcal{RH}_\infty} \frac{1}{1 - \gamma} J(\hat{A}, \hat{B}, \mathbf{K}) \\ \text{s.t. } & \begin{bmatrix} zI - \hat{A} & -\hat{B} \end{bmatrix} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = I, \\ & \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x \\ \varepsilon_{B,2} \Phi_u \end{bmatrix} \right\|_{\mathcal{H}_\infty} \leq \gamma, \quad \left\| \begin{bmatrix} \varepsilon_{A,\infty} \Phi_x \\ \varepsilon_{B,\infty} \Phi_u \end{bmatrix} \right\|_{\mathcal{L}_1} \leq \tau, \\ & G_x^\tau(\Phi_x; k) \leq b_x, \quad G_u^\tau(\Phi_u; k) \leq b_u \quad \forall k \geq 1. \end{aligned} \quad (2.3)$$

where $0 \leq \gamma, \tau \leq 1$ are fixed parameters and for each j and k ,

$$G_x^\tau(\Phi_x; k)_j := G_x(\Phi_x; k)_j + \frac{\tau}{1 - \tau} \max(\sigma_w, \|x_0\|_\infty) \|F_{x,j}^\top \Phi_x[k+1:1]\|_1,$$

and similarly for $G_u^\tau(\Phi_u; k)_j$.

In this problem, γ bounds the increase in the \mathcal{H}_2 cost due to the dynamics uncertainty, while τ determines the increase in the state and input values with respect to the constraints. In fact, both values can be related to bounding an enlarge noise process $\tilde{\mathbf{w}} = (I + \hat{\Delta})^{-1} \mathbf{w}$ driving the system.

Theorem 2.2. *Any controller designed from a feasible solution to the robust control problem (2.3) for any $0 \leq \gamma, \tau < 1$ will stabilize the true system. Furthermore, the state and input constraints will be satisfied.*

Proof. First, note that by the system norm constraints,

$$\|\hat{\Delta}\|_{\mathcal{H}_\infty} \leq \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x \\ \varepsilon_{B,2} \Phi_u \end{bmatrix} \right\|_{\mathcal{H}_\infty} < 1, \quad \|\hat{\Delta}\|_{\mathcal{L}_1} \leq \left\| \begin{bmatrix} \varepsilon_{A,\infty} \Phi_x \\ \varepsilon_{B,\infty} \Phi_u \end{bmatrix} \right\|_{\mathcal{L}_1} < 1. \quad (2.4)$$

Then by (2.2), the true system trajectory will be given by the stable system responses

$$\mathbf{x} = \Phi_x(I + \hat{\Delta})^{-1} \mathbf{w}, \quad \mathbf{u} = \Phi_u(I + \hat{\Delta})^{-1} \mathbf{w}.$$

Therefore, the state constraints are satisfied as long as

$$\begin{aligned} b_x & \geq \max_{\mathbf{w}} F_x \mathbf{x}[k] = \max_{\mathbf{w}} F_x(\Phi_x(I + \hat{\Delta})^{-1} \mathbf{w})[k] \\ & = \max_{\mathbf{w}} F_x(\Phi_x \mathbf{w})[k] - F_x(\Phi_x \hat{\Delta}(I + \hat{\Delta})^{-1} \mathbf{w})[k]. \end{aligned}$$

The first term reduces to $G_x(\Phi_x; k)$ as in the non-robust case. Because information about $\hat{\Delta}$ is not known, we resort to a sufficient condition to bound the second term, letting $\tilde{\mathbf{w}} =$

$$\widehat{\Delta}(I + \widehat{\Delta})^{-1}\mathbf{w},$$

$$\begin{aligned} |F_{x,j}^\top(\Phi_x \tilde{\mathbf{w}})[k]| &\leq \|F_{x,j}^\top \Phi_x[k+1:1]\|_1 \|\tilde{\mathbf{w}}\|_\infty \\ &\leq \|F_{x,j}^\top \Phi_x[k+1:1]\|_1 \frac{\|\widehat{\Delta}\|_{\mathcal{L}_1}}{1 - \|\widehat{\Delta}\|_{\mathcal{L}_1}} \max(\sigma_w, \|x_0\|_\infty) \\ &\leq \|F_{x,j}^\top \Phi_x[k+1:1]\|_1 \frac{\tau}{1 - \tau} \max(\sigma_w, \|x_0\|_\infty). \end{aligned}$$

Consequently, a sufficient condition for satisfying state constraints is to have for j indexing rows of F_x and elements of b_x ,

$$G_x(\Phi_x; k)_j + \frac{\tau}{1 - \tau} \max(\sigma_w, \|x_0\|_\infty) \|F_{x,j}^\top \Phi_x[k+1:1]\|_1 \leq b_{x,j}.$$

Therefore, the constraints on Φ_x imply that the state constraints are satisfied. Similar logic shows that the constraints on Φ_u imply that the input constraints are satisfied. \square

We remark on differences between the presented constraint-tightening and approaches common in the MPC literature. In both cases, the uncertainty in the dynamics induces an enlarged disturbance process. Common in the MPC literature, the additive disturbance approximation assigns $\tilde{w}_k = w_k + \Delta_A x_k + \Delta_B u_k$, or equivalently $\tilde{\mathbf{w}} = \mathbf{w} + \Delta_A \mathbf{x} + \Delta_B \mathbf{u}$. Then $\|\tilde{\mathbf{w}}\|_\infty$ can be bounded using the state and input constraint sets. Because this bound degrades as the constraint sets increase in size, this strategy can lead to counterintuitive behavior. On the other hand, our approach writes $\tilde{\mathbf{w}} = \mathbf{w} + \widehat{\Delta}(I + \widehat{\Delta})^{-1}\mathbf{w}$. While bounding $\|\tilde{\mathbf{w}}\|_\infty$ in this setting does not depend on the constraint set, it is affected by the size of the initial condition. Further comparison between the two approaches is presented in Appendix B.

2.3 Finite Dimensional Reduction

To make controller synthesis tractable, we can solve a finite approximation to optimization problem (2.3) wherein we only optimize over the first L impulse response elements of Φ_x and Φ_u , treating them as finite impulse response (FIR) filters. We show that in this setting, the optimization variables and constraints admit finite-dimensional representations. We first reformulate the constraints. Starting with the affine constraint, we have for $k = 1, \dots, L - 1$

$$\begin{aligned} \Phi_x(1) &= I, \quad \Phi_x(k+1) = \widehat{A}\Phi_x(k) + \widehat{B}\Phi_u(k) \\ V &= \widehat{A}\Phi_x(L) + \widehat{B}\Phi_u(L), \end{aligned} \tag{2.5}$$

where we will also optimize over V , a term which captures the tail of the system responses that we ignore in the synthesis.

Next, considering the system norm constraints, the \mathcal{H}_∞ norm can be reduced to a compact SDP over Φ_x, Φ_u, γ as in Theorem 5.8 of Dumitrescu [27], described explicitly for this setting in Appendix G.3 of Dean et al. [7]. For the \mathcal{L}_1 norm, the constraint becomes an $\ell_\infty \rightarrow \ell_\infty$ operator norm bound,

$$\left\| \begin{bmatrix} \varepsilon_{A,\infty} \Phi_x \\ \varepsilon_{B,\infty} \Phi_u \end{bmatrix} \right\|_{\mathcal{L}_1} = \left\| \begin{bmatrix} \varepsilon_{A,\infty} \Phi_x[L:1] \\ \varepsilon_{B,\infty} \Phi_u[L:1] \end{bmatrix} \right\|_\infty + \|V\|_\infty \leq \tau, \tag{2.6}$$

where the tail variable V enters transparently. For $k = 1, \dots, L-1$, the inequality constraints on $G_x^\tau(\Phi_x; k)$ and $G_u^\tau(\Phi_u; k)$ remain. For any $k \geq L$, the expression reduces to, for each j ,

$$\begin{aligned} \left(\sigma_w + \frac{\tau \max(\sigma_w, \|x_0\|_\infty)}{1 - \tau} \right) \|F_{x,j}^\top \Phi_x[L : 1]\|_1 &\leq b_{x,j}, \\ \left(\sigma_w + \frac{\tau \max(\sigma_w, \|x_0\|_\infty)}{1 - \tau} \right) \|F_{u,j}^\top \Phi_u[L : 1]\|_1 &\leq b_{u,j}. \end{aligned} \quad (2.7)$$

Therefore, the synthesis problem becomes, for any fixed $0 \leq \gamma, \tau < 1$,

$$\begin{aligned} \min_{\Phi_x, \Phi_u, V} \quad & \frac{1}{1 - \gamma} \sum_{t=0}^L \text{Tr}(\Phi_x(t)^\top Q \Phi_x(t) + \Phi_u(t)^\top R \Phi_u(t)) \\ \text{s.t.} \quad & (2.5), \text{SDP}(\Phi_x, \Phi_u, \gamma - \|V\|_2), (2.6), (2.7), \\ & 1 \leq k \leq L - 1 : G_x^\tau(\Phi_x; k) \leq b_x, \quad G_u^\tau(\Phi_u; k) \leq b_u. \end{aligned} \quad (2.8)$$

This is a finite dimensional SDP. The controller given by $\mathbf{K} = \Phi_u \Phi_x^{-1}$ can be written in an equivalent state-space realization (A_K, B_K, C_K, D_K) via Theorem 2 of Anderson et al. [28].

3 Suboptimality Guarantees

How much is control performance degraded by uncertainties about the dynamics? In this section, we derive a sub-optimality bound which answers this question for the constrained LQR problem. First, consider the addition of an outer minimization over γ and τ :¹

$$\min_{\gamma, \tau} \hat{J}(\gamma, \tau). \quad (3.1)$$

Denote the solution to the true optimal control problem as (Φ_x^*, Φ_u^*) , then define $\mathbf{K}_* = \Phi_u^* \Phi_x^{*-1}$ and $J_* = J(A_*, B_*, \mathbf{K}_*)$. Additionally, define constants related to the optimal system norm and the dynamics uncertainties:

$$\zeta_2 = \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x^* \\ \varepsilon_{B,2} \Phi_u^* \end{bmatrix} \right\|_{\mathcal{H}_\infty}, \quad \zeta_\infty = \left\| \begin{bmatrix} \varepsilon_{A,\infty} \Phi_x^* \\ \varepsilon_{B,\infty} \Phi_u^* \end{bmatrix} \right\|_{\mathcal{L}_1}.$$

Before stating the suboptimality result, we define a robust version of the optimal constrained controller. Define the doubly robust constraint sets

$$\begin{aligned} \bar{G}_x^\zeta(\Phi_x; k)_j &= F_{x,j}^\top \Phi(k+1)x_0 + \frac{\sigma_w}{1 - \zeta_\infty} \|F_{x,j}^\top \Phi_x[k : 1]\|_1 \\ &\quad + \frac{2\zeta_\infty \max(\sigma_w, \|x_0\|_\infty)}{1 - 2\zeta_\infty} \|F_{x,j}^\top \Phi_x[k+1 : 1]\|_1, \end{aligned}$$

¹ The objective (2.3) is unimodal in γ, τ individually, and therefore this outer minimization can be achieved by searching over the box $[0, 1) \times [0, 1)$. For less computational complexity, the minimization need only be over a single outer variable: $\max(\gamma, \tau)$. In this case, the sub-optimality bound will retain the same flavor, but the norm distinctions between cost and constraints will be less clear.

and similarly for $\bar{G}_u^\zeta(\Phi_u; k)_j$. Note that these constraint sets resemble the robust constraint sets with $\tau = 2\zeta_\infty$ and an enlarged noise process with $\sigma_w \leftarrow \frac{\sigma_w}{1-\zeta_\infty}$. Then the optimal robustly constrained controller is defined as

$$\begin{aligned}
(\Phi_x^c, \Phi_u^c) \in \underset{\Phi_x, \Phi_u \in \frac{1}{z}\mathcal{RH}_\infty}{\operatorname{argmin}} \quad & \left\| \begin{bmatrix} Q^{1/2} & \\ & R^{1/2} \end{bmatrix} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} \right\|_{\mathcal{H}_2} \\
\text{s.t.} \quad & [zI - A_\star \quad -B_\star] \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = I, \\
& \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x \\ \varepsilon_{B,2} \Phi_u \end{bmatrix} \right\|_{\mathcal{H}_\infty} \leq \zeta_2, \quad \left\| \begin{bmatrix} \varepsilon_{A,\infty} \Phi_x \\ \varepsilon_{B,\infty} \Phi_u \end{bmatrix} \right\|_{\mathcal{L}_1} \leq \zeta_\infty, \\
& \bar{G}_x^\zeta(\Phi_x; k) \leq b_x, \quad \bar{G}_u^\zeta(\Phi_u; k) \leq b_u \quad \forall k \geq 1.
\end{aligned} \tag{3.2}$$

with $\mathbf{K}_c = \Phi_u(\Phi_x^c)^{-1}$. Note that this optimization problem designs a controller which has similar closed-loop behavior to the true optimal controller (i.e. system norms of similar size), but satisfies more stringent state and input constraints. The relative *robustness cost gap* is defined as

$$M_\zeta = \frac{J(A_\star, B_\star, \mathbf{K}_c) - J(A_\star, B_\star, \mathbf{K}_\star)}{J(A_\star, B_\star, \mathbf{K}_\star)}. \tag{3.3}$$

Now we are ready to state a data independent bound on the sub-optimality of robust controllers synthesized on estimated dynamics.

Theorem 3.1. *Suppose that the robust optimal constrained controller problem is feasible. As long as $\zeta_2 \leq \frac{1}{4\sqrt{2}}$ and $\zeta_\infty \leq \frac{1}{2}$, we have that the cost achieved by $\hat{\mathbf{K}} = \hat{\Phi}_u \hat{\Phi}_x^{-1}$ synthesized from the minimizers of (3.1) satisfies*

$$\frac{J(A_\star, B_\star, \hat{\mathbf{K}}) - J_\star}{J_\star} \leq 4\sqrt{2}(1 + M_\zeta)(\varepsilon_{A,2} + \varepsilon_{B,2}\|\mathbf{K}_\star\|_{\mathcal{H}_\infty})\|\Phi_x^\star\|_{\mathcal{H}_\infty} + M_\zeta.$$

This result is stated in terms of quantities related to the unknown true system, with the goal of highlighting properties of systems that make them harder or easier to robustly control. We see that the bound grows with the \mathcal{H}_∞ norm of the optimal controller and closed-loop response. It also grows with the robustness cost gap M_ζ .

Remark that M_ζ is difficult to characterize analytically; in general, it requires checking the boundaries of the robust polytopic constraints. Once errors ε_A and ε_B are small enough that the optimal system response $(\Phi_x^\star, \Phi_u^\star)$ satisfies the doubly robust constraints, we will have $M_\zeta = 0$. If the optimal controller saturates its constraints, then M_ζ will be nonzero for any nonzero estimation errors. In Section 5 (Figure 2), we numerically characterize this optimality gap for a double integrator example.

Proof. Using (2.2) along with the norm bounds (2.4) and the constraints in optimization problem (3.1),

$$J(A_\star, B_\star, \hat{\mathbf{K}}) \leq \frac{1}{1-\hat{\gamma}} J(\hat{A}, \hat{B}, \hat{\mathbf{K}}).$$

Next, we will connect the optimal system response to the estimated system by constructing a feasible solution to the robust optimization problem. We will use the following lemma

Lemma 3.2. *Under the conditions of Theorem 3.1, we have that the following is a feasible solution to (3.1)*

$$\tilde{\Phi}_x = \Phi_x^c(I - \Delta)^{-1}, \quad \tilde{\Phi}_u = \Phi_u^c(I - \Delta)^{-1}, \quad \tilde{\gamma} = \frac{\sqrt{2}\zeta_2}{1 - \sqrt{2}\zeta_2}, \quad \tilde{\tau} = \frac{\zeta_\infty}{1 - \zeta_\infty}.$$

where we define $\Delta := - \begin{bmatrix} \Delta_A & \Delta_B \end{bmatrix} \begin{bmatrix} \Phi_x^c \\ \Phi_u^c \end{bmatrix}$.

The proof of Lemma 3.2 follows by checking that the proposed solution satisfies all the constraints and is presented in Appendix A. Applying Lemma 3.2,

$$\frac{1}{1 - \tilde{\gamma}} J(\hat{A}, \hat{B}, \hat{\mathbf{K}}) \leq \frac{1}{1 - \tilde{\gamma}} J(\hat{A}, \hat{B}, \mathbf{K}_c).$$

This is true because $(\hat{\mathbf{K}}, \hat{\gamma})$ is the optimal solution to (3.1), so objective function with feasible $(\tilde{\Phi}_u \tilde{\Phi}_x^{-1} = \mathbf{K}_c, \tilde{\gamma})$ is an upper bound. Then we have

$$\begin{aligned} J(A_\star, B_\star, \hat{\mathbf{K}}) &\leq \frac{1}{1 - \tilde{\gamma}} J(\hat{A}, \hat{B}, \mathbf{K}_c) \\ &\leq \frac{1}{1 - \tilde{\gamma}} \frac{1}{1 - \|\Delta\|_2} J(A_\star, B_\star, \mathbf{K}_c) \\ &\leq \frac{1}{1 - \frac{\sqrt{2}\zeta}{1 - \sqrt{2}\zeta}} \frac{1}{1 - \sqrt{2}\zeta} (1 + M) J(A_\star, B_\star, \mathbf{K}_\star) \\ &\leq \left(1 + 4\sqrt{2}\zeta_2(1 + M) + M\right) J(A_\star, B_\star, \mathbf{K}_\star). \end{aligned}$$

The second inequality follows from an application of (2.2) with the roles of the nominal and true systems switched. The final follows from bounding $\|\Delta\|_2$ by $\sqrt{2}\zeta$ and noticing that $\frac{x}{1-x} \leq 2x$ for $0 \leq x \leq \frac{1}{2}$, where we set $x = 2\sqrt{2}\zeta_2$. Finally, we note that

$$\zeta_2 = \left\| \begin{bmatrix} \varepsilon_{A,2} I \\ \varepsilon_{B,2} \mathbf{K}_\star \end{bmatrix} \Phi_x^\star \right\|_{\mathcal{H}_\infty} \leq (\varepsilon_{A,2} + \varepsilon_{B,2} \|\mathbf{K}_\star\|_{\mathcal{H}_\infty}) \|\Phi_x^\star\|_{\mathcal{H}_\infty}.$$

□

Here, we briefly remark that a similar sub-optimality bound can be derived for the finite problem in (2.8). In short, controllers synthesized by minimizing over γ and τ will satisfy a sub-optimality bound of the form in Theorem 3.1 with an additional factor due to the FIR truncation. The formal statement and proof of this result are deferred to Appendix A, but we highlight here that the cost penalty incurred due to FIR approximation decays exponentially in the horizon L over which the approximation is taken.

4 Learning with Control

Finally, we connect the previous results on robust control with system estimation. We adopt control actions that both keep the system safe and provide excitation,

$$\mathbf{u} = \mathbf{K}\mathbf{x} + \eta \quad (4.1)$$

where each $\eta = (\eta_0, \eta_1, \dots)$ is stochastic and ℓ_∞ -bounded, i.e. $\|\eta_k\|_\infty \leq \sigma_\eta$. Given a trajectory sequence $\{(x_k, u_k)\}_{k=0}^T$, we propose to learn the dynamics (A_\star, B_\star) via least-squares regression on a trajectory of length T :

$$(\hat{A}, \hat{B}) \in \operatorname{argmin}_{(A, B)} \sum_{k=0}^{T-1} \frac{1}{2} \|Ax_k + Bu_k - x_{k+1}\|_2^2. \quad (4.2)$$

In what follows, we will prove a statistical rate on the least-squares estimate (\hat{A}, \hat{B}) in terms of the system response and the trajectory length.

The bulk of the proof for the statistical rate comes from a general theorem regarding linear-response time series data from Simchowitz et al. [6]. Recently, this proof was adopted by Dean et al. [7] to show a rate of estimation in the setting given by (4.1) when both η and the disturbance w are Gaussian distributed. We modify the reduction given by Dean et al. to the case when the excitation and disturbance are no longer Gaussian, but instead zero-mean and bounded. We assume that w_k and η_k are both zero-mean sequences with independent coordinates and finite fourth moments. In particular, we assume $\mathbb{E}_{w_k}[w_k(i)^2] \leq \sigma_w^2$, $\mathbb{E}_{w_k}[w_k(i)^4] \lesssim \sigma_w^4$, $\mathbb{E}_{\eta_k}[\eta_k(i)^2] \leq \sigma_\eta^2$, $\mathbb{E}_{\eta_k}[\eta_k(i)^4] \lesssim \sigma_\eta^4$. These assumptions are quickly verified for common distributions such as uniform on a compact interval or over a discrete set of points. The main estimation result is the following.

Theorem 4.1. *Fix a failure probability $\delta \in (0, 1)$. Suppose the stochastic disturbance $\{w_k\}$ and the input disturbance $\{\eta_k\}$ satisfy the assumptions above. Assume for simplicity that $\sigma_\eta \leq \sigma_w$, and that the stabilizing controller \mathbf{K} achieves a SLS response $\Phi_x \in \frac{1}{z}\mathcal{RH}_\infty(C_x, \rho)$, $\Phi_u \in \frac{1}{z}\mathcal{RH}_\infty(C_u, \rho)$. Let $C_K^2 := nC_x^2 + dC_u^2$. Then as long as the trajectory length T satisfies the condition:*

$$T \gtrsim T_0 := (n + d) \log \left(\frac{dC_u^2}{\delta} + \frac{\sigma_w^2}{\sigma_\eta^2} \frac{\rho^2 C_u^2 C_K^2}{\delta(1 - \rho^2)} \left(1 + \|B_\star\|_2^2 + \frac{\|x_0\|_2^2}{\sigma_w^2 T} \right) \right), \quad (4.3)$$

we have the following bound on the least-squares estimation errors that holds with probability at least $1 - \delta$,

$$\max\{\|\Delta_A\|_2, \|\Delta_B\|_2\} \lesssim \frac{\sigma_w C_u}{\sigma_\eta} \sqrt{\frac{n + d}{T}} \sqrt{\log \left(\frac{dC_u}{\delta} + \frac{\sigma_w}{\sigma_\eta} \frac{\rho C_u C_K}{\delta(1 - \rho^2)} \left(1 + \|B_\star\|_2 + \frac{\|x_0\|_2}{\sigma_w \sqrt{T}} \right) \right)}.$$

The proof of this result is presented in Appendix C. We remark on the interpretation of statistical learning bounds. A priori guarantees, like the one presented here, depend on

quantities related to the underlying true system. As such, they are not directly useful when the system is unknown,² but rather they indicate qualities of systems that make them easier or harder to estimate. For example, we see that this bound on estimation errors decreases with the size of the excitation, and increases with the size of the process noise. The bound also increases with C_u , a constant which bounds the gain from disturbance to control inputs.

Applying our robustness results to the learning controller (4.1), it is possible to guarantee safety *during* the estimation process. We can define an expanded noise process $\tilde{w}_k = B_\star \eta_k + w_k$ to synthesize \mathbf{K}_0 using the robust control synthesis problem (2.3). As long as it is feasible for initial system estimates (A_0, B_0) , initial dynamics uncertainties $(\varepsilon_A^0, \varepsilon_B^0)$, σ_w replaced with³ $\sigma_\eta \|B_\star\|_\infty + \sigma_w$, and $b_{u,j}$ replaced with $b_{u,j} - \sigma_\eta \|F_{u,j}\|_1$, then the control law $\mathbf{u} = \mathbf{K}_0 \mathbf{x} + \eta$ stabilizes the true system, satisfies state and input constraints, and allows for learning at the rate given in Theorem 4.1.

Finally, we connect the sub-optimality result to the statistical learning bound for an end-to-end sample complexity bound on the constrained LQR problem. Define M_T to be the value of M_ζ when the definition determined by (3.2) has values set as

$$\zeta_\infty = \varepsilon \sqrt{n+d} \left\| \begin{bmatrix} \Phi_x^\star \\ \Phi_u^\star \end{bmatrix} \right\|_{\mathcal{L}_1}, \quad \zeta_2 = \varepsilon \left\| \begin{bmatrix} \Phi_x^\star \\ \Phi_u^\star \end{bmatrix} \right\|_{\mathcal{H}_\infty}, \quad \varepsilon \gtrsim \frac{\sigma_w C_u}{\sigma_\eta} \sqrt{\frac{n+d}{T}} \mathcal{O}(\sqrt{\log(d/\delta)}).$$

Corollary 4.2. *Assume initial feasibility of the learning problem. Then for a trajectory of length*

$$T \gtrsim T_0 \frac{\sigma_w^2 C_u^2}{\sigma_\eta^2} \max \left\{ (n+d)(1 + \|\mathbf{K}_\star\|_{\mathcal{L}_1})^2 \|\Phi_x^\star\|_{\mathcal{L}_1}^2, (1 + \|\mathbf{K}_\star\|_{\mathcal{H}_\infty})^2 \|\Phi_x^\star\|_{\mathcal{H}_\infty}^2 \right\}, \quad (4.4)$$

the cost achieved by $\hat{\mathbf{K}} = \hat{\Phi}_u (\hat{\Phi}_x)^{-1}$ synthesized from (2.3) on the least-squares estimates \hat{A}, \hat{B} satisfies with probability at least $1 - \delta$,

$$\frac{J(A_\star, B_\star, \hat{\mathbf{K}}) - J_\star}{J_\star} \lesssim \frac{\sigma_w C_u}{\sigma_\eta} \sqrt{\frac{n+d}{T}} (1 + M_T) (1 + \|\mathbf{K}_\star\|_{\mathcal{H}_\infty}) \|\Phi_x^\star\|_{\mathcal{H}_\infty} \mathcal{O}(\sqrt{\log(d/\delta)}) + M_T.$$

Proof. (Sketch) This result follows by combining the statistical guarantee in Theorem 4.1 with the sub-optimality bound in Theorem 3.1. Note that we use the naïve bound $\varepsilon_{A,\infty} \leq \sqrt{n} \varepsilon_{A,2}$ and similarly $\varepsilon_{B,\infty} \leq \sqrt{d} \varepsilon_{B,2}$; this results in an extra factor of $(n+d)$ appearing in (4.4) and in the definition of M_T . \square

Our final result depends both on the true system and the initial system estimates by way of the learning controller, which affects T_0 and constants C_x , C_u , and ρ . The system constraints enter through their effect on M_T .

² Statistical bounds in terms of data-dependent quantities can also be worked out; however, modern methods like bootstrapping generally provide tighter statistical guarantees [29].

³ Note that since the quantity $\|B_\star\|_\infty$ would not generally be known, it can be bounded by $\|B_0\|_\infty + \varepsilon_{B,\infty}^0$.

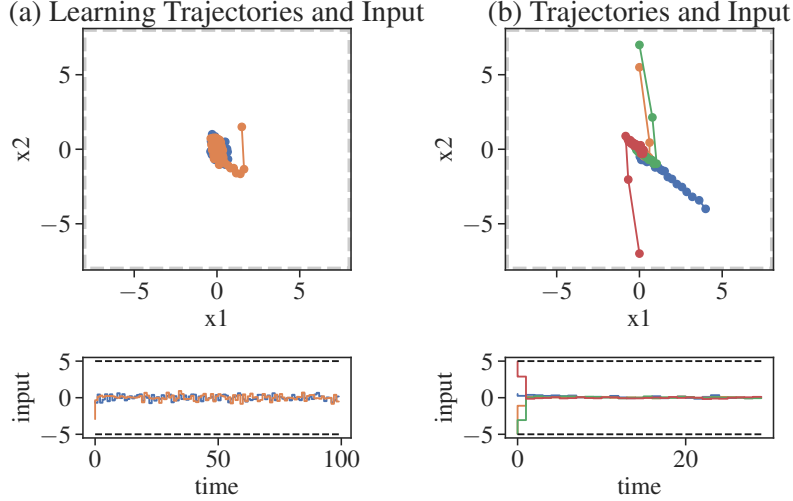


Figure 1: Safe learning trajectories synthesized with coarse initial estimates (a), then robust execution with reduced model errors (b).

5 Numerical Experiments

We demonstrate the utility of this framework on a double integrator example. In this case, the true dynamics are given by

$$x_{k+1} = \begin{bmatrix} 1 & 0.1 \\ 0 & 1 \end{bmatrix} x_k + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u_k + w_k$$

with the constraints as states bounded between -8 and 8 , and inputs bounded in between -4 and 4 . We have $\sigma_w = 0.1$. Our initial estimate comes from a randomly generated initial perturbation of the true system with $\varepsilon_{A,\infty} = \varepsilon_{B,\infty} = 0.1$. Safe controllers are generated with finite truncation length $L = 15$, and for larger initial conditions, the system is warm-started with a finite-time robust controller with horizon 20 to reduce the initial condition.

Figure 1 displays safe trajectories and input sequences for several example initial conditions. In 1a, the plotted trajectories are used for learning: the controller both regulates and excites the system ($\sigma_\eta = 0.5$), and is robust to initial uncertainties. Figure 1b demonstrates an ability to operate closer to the constraints when there is less uncertainty: in this case, there is no added excitation ($\sigma_\eta = 0$) and the system estimates are better specified ($\varepsilon_\infty = 0.001$), so larger initial conditions are feasible.

Figure 2a displays the decreasing estimation errors over time, demonstrating learning. Shaded areas represent quartiles over 400 trials. Figure 2b displays the trade-off between safety and exploration by showing the largest value of σ_η for which the robust synthesis is feasible, given a size for the state constraint set r_x . Here, we leave $x_0 = 0$, and examine a variety of errors in the dynamics estimates. As the uncertainties in the dynamics decrease, higher levels of both safety and exploration are achievable. Finally, Figure 2c shows the value of the relative robust sub-optimality gap M_ζ for the given example. We see a sharp




Figure 2: Over time, estimation errors decrease (a). As safety requirements increase, the maximum feasible excitation decreases (b). The robust cost suboptimality gap M_ζ displays an abrupt transition from feasibility to near-optimality with error size (c).

transition from infeasibility for $\varepsilon \geq 0.05$ to near-optimality for $\varepsilon \leq 0.03$. This indicates that the gap may be most significant as a feasibility condition for our sub-optimality guarantees to hold.

6 Discussion

In this paper, we propose a method for learning unknown linear systems while ensuring that they satisfy state and input constraints. By synthesizing a controller that both excites and regulates the system, we address the trade-off between safety and exploration directly. We further derive an end-to-end finite sample bound on the performance of constrained LQR controllers synthesized from collected data.

There are several directions for possible extensions of this work. To mitigate the conservativeness of the robust controller, tighter bounds on the uncertainty in the system response $\hat{\Delta}$ could be derived for structured settings, where more than just the norm of the error is known. To connect this work to experiment design literature, the objective in the synthesis problem (2.3) could be replaced with an exploration inspired cost function for the learning stage.

Alternatively, the constrained LQR problem could be cast in the setting of online learning, where one seeks to minimize cost at all times, including during learning. This would require an analysis of recursive feasibility, to understand the transition that occurs when controllers are updated based on refined system estimates. It would also likely require a finer analysis on the performance loss characterized by M_ζ . Finally, we remark that the exploration vs. safety trade-off is most compelling for nonlinear systems, and a linear-time-varying extension to this work may be applicable to those problems.

ACKNOWLEDGMENT

We thank Francesco Borrelli and the members of the MPC Lab at UC Berkeley for their helpful comments and feedback. SD is supported by an NSF Graduate Research Fellowship under Grant No. DGE 1752814. ST is supported by a Google PhD fellowship. BR is generously supported in part by ONR awards N00014-17-1-2191, N00014-17-1-2401, and N00014-18-1-2833, the DARPA Assured Autonomy (FA8750-18-C-0101) and Lagrange (W911NF-16-1-0552) programs, and an Amazon AWS AI Research Award.

References

- [1] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, “On the sample complexity of the linear quadratic regulator,” *arXiv preprint arXiv:1710.01688*, 2017.
- [2] Y.-S. Wang, N. Matni, and J. C. Doyle, “A system level approach to controller synthesis,” *arXiv preprint arXiv:1610.04815*, 2016.
- [3] C.-N. Fiechter, “PAC adaptive control of linear systems,” in *Conference on Learning Theory*, 1997.
- [4] Y. Abbasi-Yadkori and C. Szepesvári, “Regret bounds for the adaptive control of linear quadratic systems,” in *Conference on Learning Theory*, 2011.
- [5] M. Abeille and A. Lazaric, “Thompson sampling for linear-quadratic control problems,” in *AISTATS*, 2017.
- [6] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht, “Learning without mixing: Towards a sharp analysis of linear system identification,” in *Conference on Learning Theory*, 2018.
- [7] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, “Regret bounds for robust adaptive control of the linear quadratic regulator,” in *Neural Information Processing Systems*, 2018.
- [8] A. Bemporad and M. Morari, “Robust model predictive control: A survey,” in *Robustness in identification and control*. Springer, 1999, pp. 207–226.
- [9] M. V. Kothare, V. Balakrishnan, and M. Morari, “Robust constrained model predictive control using linear matrix inequalities,” *Automatica*, vol. 32, no. 10, pp. 1361–1379, 1996.
- [10] D. Q. Mayne, M. M. Seron, and S. Raković, “Robust model predictive control of constrained linear systems with bounded disturbances,” *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.

- [11] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski, “Optimization over state feedback policies for robust control with constraints,” *Automatica*, vol. 42, no. 4, pp. 523–533, 2006.
- [12] C. A. Larsson, C. R. Rojas, and H. Hjalmarsson, “MPC oriented experiment design,” *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 9966–9971, 2011.
- [13] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin, “Provably safe and robust learning-based model predictive control,” *Automatica*, vol. 49, no. 5, pp. 1216–1226, 2013.
- [14] M. Tanaskovic, L. Fagiano, R. Smith, and M. Morari, “Adaptive receding horizon control for constrained MIMO systems,” *Automatica*, vol. 50, no. 12, pp. 3019–3029, 2014.
- [15] M. Lorenzen, F. Allgöwer, and M. Cannon, “Adaptive model predictive control with robust constraint satisfaction,” *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 3313–3318, 2017.
- [16] F. Blanchini, “Set invariance in control,” *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.
- [17] F. Berkenkamp and A. P. Schoellig, “Safe and robust learning control with gaussian processes,” in *European Control Conference (ECC)*, 2015.
- [18] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, “Safe model-based reinforcement learning with stability guarantees,” in *Neural Information Processing Systems*, 2017.
- [19] G. Dalal, K. Dvijotham, M. Vecerik, T. Hester, C. Paduraru, and Y. Tassa, “Safe exploration in continuous action spaces,” *arXiv preprint arXiv:1801.08757*, 2018.
- [20] Y. Chow, O. Nachum, E. Duenez-Guzman, and M. Ghavamzadeh, “A lyapunov-based approach to safe reinforcement learning,” *arXiv preprint arXiv:1805.07708*, 2018.
- [21] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, “Learning-based model predictive control for safe exploration and reinforcement learning,” *arXiv preprint arXiv:1803.08287*, 2018.
- [22] K. P. Wabersich and M. N. Zeilinger, “Linear model predictive safety certification for learning-based control,” *arXiv preprint arXiv:1803.08552*, 2018.
- [23] T. Lu, M. Zinkevich, C. Boutilier, B. Roy, and D. Schuurmans, “Safe exploration for identifying linear systems via robust optimization,” *arXiv preprint arXiv:1711.11165*, 2017.
- [24] F. Oldewurtel, C. N. Jones, and M. Morari, “A tractable approximation of chance constrained stochastic MPC based on affine disturbance feedback,” in *47th IEEE Annual Conference on Decision and Control*, 2008.

- [25] M. Farina, L. Giulioni, and R. Scattolini, “Stochastic linear model predictive control with chance constraints—a review,” *Journal of Process Control*, vol. 44, pp. 53–67, 2016.
- [26] N. Matni, Y.-S. Wang, and J. Anderson, “Scalable system level synthesis for virtually localizable systems,” in *IEEE 56th Annual Conference on Decision and Control*, 2017.
- [27] B. Dumitrescu, *Positive Trigonometric Polynomials and Signal Processing Applications*. Springer, 2007, vol. 103.
- [28] J. Anderson and N. Matni, “Structured state space realizations for sls distributed controllers,” in *Allerton*, 2017.
- [29] B. Efron, “Bootstrap methods: another look at the jackknife,” in *Breakthroughs in statistics*. Springer, 1992, pp. 569–593.
- [30] R. Ibragimov and S. Sharakhmetov, “The exact constant in the rosenthal inequality for random variables with mean zero,” *Theory of Probability & Its Applications*, vol. 46, no. 1, pp. 127–132, 2002.

A Extended Sub-optimality Results and Proofs

This section contains proofs necessary for the sub-optimality results. It further develops a sub-optimality bound for the case of FIR truncation. We begin by considering the frequency response elements of composed transfer functions. First, define, for $\mathbf{D} = \sum_{k=1}^{\infty} D(k)z^{-k}$,

$$\text{Toep}_k(D) := \begin{bmatrix} D(1) & & \\ \vdots & \ddots & \\ D(k) & \dots & D(1) \end{bmatrix}.$$

Then we have the following lemma.

Lemma A.1. *For $\mathbf{M} = \sum_{k=1}^{\infty} M(k)z^{-k}$ and $\mathbf{D} = \sum_{k=1}^{\infty} D(k)z^{-k}$ the frequency response elements of $\tilde{\mathbf{M}} = \mathbf{M}\mathbf{D}$ are given by*

$$\tilde{M}(k) = M[k : 1]D(1 : k).$$

where we use the notation $D(1 : k)$ for the vertical concatenation of $D(1)$ through $D(k)$. Further, we have

$$\tilde{M}[k : 1] = M[k : 1]\text{Toep}_k(D).$$

Proof. For the first statement, notice that we can write

$$\mathbf{M}\mathbf{D} = \sum_{k=1}^{\infty} \sum_{t=1}^k M(t)D(k-t)z^{-k} = \sum_{k=1}^{\infty} M[k : 1]D(1 : k)z^{-k}.$$

Then for the second, we have

$$\begin{aligned} \tilde{M}[k : 1] &= [M[k : 1]D(1 : k) \quad \dots \quad M[1]D(1)] \\ &= M[k : 1] \underbrace{\begin{bmatrix} D(1) & & \\ \vdots & \ddots & \\ D(k) & \dots & D(1) \end{bmatrix}}_{\text{Toep}_k(D)}. \end{aligned}$$

□

A.1 Proof of Lemma 3.2

Proof of Lemma 3.2. The proposed feasible solution is

$$\tilde{\Phi}_x = \Phi_x^c(I - \Delta)^{-1}, \quad \tilde{\Phi}_u = \Phi_u^c(I - \Delta)^{-1}, \quad \tilde{\gamma} = \frac{\sqrt{2}\zeta_2}{1 - \sqrt{2}\zeta_2}, \quad \tilde{\tau} = \frac{\zeta_{\infty}}{1 - \zeta_{\infty}},$$

where $\Delta = - \begin{bmatrix} \Delta_A & \Delta_B \end{bmatrix} \begin{bmatrix} \Phi_x^c \\ \Phi_u^c \end{bmatrix}$. By construction, $\tilde{\Phi}_x$ and $\tilde{\Phi}_u$ satisfy the equality constraints. Next, notice that

$$\|\Delta\|_{\mathcal{H}_\infty} = \left\| \begin{bmatrix} \frac{1}{\varepsilon_{A,2}} \Delta_A & \frac{1}{\varepsilon_{B,2}} \Delta_B \end{bmatrix} \begin{bmatrix} \varepsilon_{A,2} \Phi_x^c \\ \varepsilon_{B,2} \Phi_u^c \end{bmatrix} \right\|_{\mathcal{H}_\infty} \leq \sqrt{2} \zeta_2,$$

due to the sub-multiplicative property of the \mathcal{H}_∞ norm and the norm constraint in the definition of (Φ_x^c, Φ_u^c) . By similar logic $\|\Delta\|_{\mathcal{L}_1} \leq \zeta_\infty$.

Considering the \mathcal{H}_∞ norm constraint,

$$\sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \tilde{\Phi}_x \\ \varepsilon_{B,2} \tilde{\Phi}_u \end{bmatrix} \right\|_{\mathcal{H}_\infty} = \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x^c \\ \varepsilon_{B,2} \Phi_u^c \end{bmatrix} (I - \Delta)^{-1} \right\|_{\mathcal{H}_\infty} \leq \sqrt{2} \zeta_2 \frac{1}{1 - \|\Delta\|_{\mathcal{H}_\infty}} \leq \frac{\sqrt{2} \zeta_2}{1 - \sqrt{2} \zeta_2} = \tilde{\gamma},$$

where the inequality follows due to the fact that $\sqrt{2} \|\Delta\|_{\mathcal{H}_\infty} \leq 1$. Then consider the \mathcal{L}_1 norm constraint,

$$\left\| \begin{bmatrix} \varepsilon_{A,\infty} \tilde{\Phi}_x \\ \varepsilon_{B,\infty} \tilde{\Phi}_u \end{bmatrix} \right\|_{\mathcal{L}_1} \leq \zeta_\infty \frac{1}{1 - \|\Delta\|_{\mathcal{L}_1}} \leq \frac{\zeta_\infty}{1 - \zeta_\infty} = \tilde{\tau}$$

where the decomposition of the inverse is valid by assumption on the size of ζ_∞ .

Then it remains to show that the robust state and input constraints are satisfied. For compactness we will write $c_0 = \max(1, \frac{1}{\sigma_w} \|x_0\|_\infty)$. Recall that they are given by

$$G_x^\tau(\tilde{\Phi}_x; k)_j = F_{x,j}^\top \tilde{\Phi}_x(k+1)x_0 + \sigma_w \|F_{x,j}^\top \tilde{\Phi}_x[k:1]\|_1 + \frac{\tau \sigma_w c_0}{1 - \tau} \|F_{x,j}^\top \tilde{\Phi}_x[k+1:1]\|_1.$$

Note that $\tilde{\Phi}_x = \Phi_x^c + \Phi_x^c \Delta (I - \Delta)^{-1}$. Define the frequency response elements of $\Delta (I - \Delta)^{-1}$ by $D(k)$. Then, using Lemma A.1 we have

$$\tilde{\Phi}(k) = \Phi_x^c(k) + \Phi_x^c[k:1] D(1:k), \quad \tilde{\Phi}[k:1] = \Phi_x^c[k:1] + \Phi_x^c[k:1] \text{Toep}_k(D).$$

Lemma A.2. Assume that $\zeta_\infty < 1$ and let D represent the frequency response of $\Delta(1 - \Delta)^{-1}$. Then for any vector v ,

$$\|v^\top \text{Toep}_k(D)\|_1 \leq \frac{\zeta_\infty}{1 - \zeta_\infty} \|v\|_1.$$

Proof.

$$\begin{aligned} \|v^\top \text{Toep}_k(D)\|_1 &= \|\text{Toep}_k(D)^\top v\|_1 \leq \|\text{Toep}_k(D)^\top\|_1 \|v\|_1 = \|\text{Toep}_k(D)\|_\infty \|v\|_1 \\ &\leq \|\Delta(1 - \Delta)^{-1}\|_{\mathcal{L}_1} \|v\|_1 \leq \frac{\zeta_\infty}{1 - \zeta_\infty} \|v\|_1. \end{aligned}$$

Above, we make use of the fact that the ℓ_1 and ℓ_∞ norms are duals, and therefore $\|A\|_\infty = \|A^\top\|_1$. The second inequality holds because $\text{Toep}_k(D)$ is a truncation of the semi-infinite Toeplitz matrix associated with the operator $\Delta(1 - \Delta)^{-1}$. The final decomposition is valid because $\|\Delta\|_{\mathcal{L}_1} \leq \zeta_\infty < 1$. \square

Now we are ready to consider the state constraint indexed by j and k ,

$$\begin{aligned} G_x^{\tilde{\tau}}(\tilde{\Phi}_x; k)_j &= F_{x,j}^\top (\Phi_x^c(k+1) + \Phi_x^c[k+1:1]D(1:k+1))x_0 \\ &\quad + \sigma_w \|F_{x,j}^\top \Phi_x^c[k:1](I + \text{Toep}_k(D))\|_1 \\ &\quad + \frac{\tilde{\tau}}{1-\tilde{\tau}} \sigma_w c_0 \|F_{x,j}^\top \Phi_x^c[k+1:1](I + \text{Toep}_{k+1}(D))\|_1 \end{aligned}$$

Considering each term individually,

$$\begin{aligned} F_{x,j}^\top (\Phi_x^c(k+1) + \Phi_x^c[k+1:1]D(1:k+1))x_0 \\ = F_{x,j}^\top \Phi_x^c(k+1)x_0 + F_{x,j}^\top \Phi_x^c[k+1:1]D(1:k+1)x_0, \end{aligned}$$

Then defining E_1 to contain an identity in the first block and zeros elsewhere,

$$\begin{aligned} F_{x,j}^\top \Phi_x^c[k+1:1]D(1:k+1)x_0 &\leq \|F_{x,j}^\top \Phi_x^c[k+1:1]\text{Toep}_{k+1}(D)E_1\|_1 \|x_0\|_\infty \\ &\leq \|F_{x,j}^\top \Phi_x^c[k+1:1]\|_1 \frac{\zeta_\infty}{1-\zeta_\infty} \|x_0\|_\infty, \end{aligned}$$

The first inequality is Hölder's inequality and the second by Lemma A.2. Next, the second term:

$$\sigma_w \|F_{x,j}^\top \Phi_x^c[k:1](I + \text{Toep}_k(D))\|_1 \leq \sigma_w \frac{1}{1-\zeta_\infty} \|F_{x,j}^\top \Phi_x^c[k:1]\|_1,$$

where the inequality holds by Lemma A.2. Finally, the last term,

$$\frac{\tilde{\tau}}{1-\tilde{\tau}} \sigma_w c_0 \|F_{x,j}^\top \Phi_x^c[k+1:1](I + \text{Toep}_{k+1}(D))\|_1 \leq \frac{\zeta_\infty}{1-2\zeta_\infty} \sigma_w c_0 \frac{1}{1-\zeta_\infty} \|F_{x,j}^\top \Phi_x^c[k+1:1]\|_1$$

where we apply Lemma A.2 and plug in the definition of $\tilde{\tau}$:

$$\frac{\tilde{\tau}}{1-\tilde{\tau}} = \frac{\zeta_\infty}{(1-\zeta_\infty)(1-\frac{\zeta_\infty}{1-\zeta_\infty})} = \frac{\zeta_\infty}{1-2\zeta_\infty}.$$

The resulting sum is

$$\begin{aligned} G_x^{\tilde{\tau}}(\tilde{\Phi}_x; k)_j &\leq F_{x,j}^\top \Phi_x^c(k+1)x_0 + \frac{\sigma_w}{1-\zeta_\infty} \|F_{x,j}^\top \Phi_x^c[k:1]\|_1 \\ &\quad + \left(\frac{\zeta_\infty}{1-\zeta_\infty} \|x_0\|_\infty + \frac{\zeta_\infty}{1-2\zeta_\infty} \frac{c_0 \sigma_w}{1-\zeta_\infty} \right) \|F_{x,j}^\top \Phi_x^c[k+1:1]\|_1 \end{aligned}$$

Then considering constants around the final term,

$$\frac{\zeta_\infty}{1-\zeta_\infty} \|x_0\|_\infty + \frac{\zeta_\infty}{1-2\zeta_\infty} \frac{c_0 \sigma_w}{1-\zeta_\infty} \leq \frac{1}{1-\zeta_\infty} \left(\zeta_\infty + \frac{\zeta_\infty}{1-2\zeta_\infty} \right) c_0 \sigma_w = \frac{2\zeta_\infty}{1-2\zeta_\infty} c_0 \sigma_w.$$

Thus, we see that $G_x^{\tilde{\tau}}(\tilde{\Phi}_x; k)_j \leq \bar{G}_x^\zeta(\Phi_x^c; k)_j \leq b_j$ due to the constraints on Φ_x^c . A similar computation with the input constraints shows the same result. Therefore, the proposed solution is feasible. \square

A.2 Finite Dimensional Sub-optimality

We can recover sub-optimality guarantees in the case that we optimize over only a finite set of system response variables. Define the truncated responses $\Phi_x^L = \sum_{k=1}^L \Phi_x(t)z^{-k}$ and similarly for Φ_u^L , and notice that

$$\begin{bmatrix} zI - A & -B \end{bmatrix} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = I \iff \begin{bmatrix} zI - A & -B \end{bmatrix} \begin{bmatrix} \Phi_x^L \\ \Phi_u^L \end{bmatrix} = I + \frac{1}{z^L} \Phi_x(L+1). \quad (\text{A.1})$$

This reformulation allows for the optimization over the FIR filters Φ_x^L and Φ_u^L , plus an additional variable that represents the tail of the true response. Applying this observation to the robust synthesis problem (2.3) yields the following finite dimensional problem:

$$\begin{aligned} \widehat{J}_L(\gamma, \tau) := & \min_{\substack{\Phi_x^L, \Phi_u^L \in \mathcal{RH}_\infty \\ V}} \frac{1}{1-\gamma} J(\widehat{A}, \widehat{B}, \mathbf{K}) \\ \text{s.t. } & \begin{bmatrix} zI - \widehat{A} & -\widehat{B} \end{bmatrix} \begin{bmatrix} \Phi_x^L \\ \Phi_u^L \end{bmatrix} = I + \frac{1}{z^L} V, \\ & \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x^L \\ \varepsilon_{B,2} \Phi_u^L \end{bmatrix} \right\|_{\mathcal{H}_\infty} + \|V\|_2 \leq \gamma, \quad \left\| \begin{bmatrix} \varepsilon_{A,\infty} \Phi_x^L \\ \varepsilon_{B,\infty} \Phi_u^L \end{bmatrix} \right\|_{\mathcal{L}_1} + \|V\|_\infty \leq \tau, \\ & G_x^\tau(\Phi_x^L; k) \leq b_x, \quad G_u^\tau(\Phi_u^L; k) \leq b_u \quad \forall k. \end{aligned} \quad (\text{A.2})$$

We now show that so long as the horizon L is suitably large, all of the properties that hold for the solution of infinite horizon problem (2.3) carry over to the solution of the finite dimensional approximation (A.2).

Proposition A.3. *Any controller designed from a feasible solution to the finite robust control problem (A.2) for any $0 \leq \gamma, \tau < 1$ stabilizes the true system and ensures that state and input constraints will be satisfied.*

Proof. Consider the affine constraint on Φ_x^L and Φ_u^L . We can equivalently write, using the observation in (A.1),

$$\begin{bmatrix} zI - \widehat{A} & -\widehat{B} \end{bmatrix} \begin{bmatrix} \Phi_x^L \\ \Phi_u^L \end{bmatrix} = I + \frac{1}{z^L} V \implies \begin{bmatrix} zI - A_\star & -B_\star \end{bmatrix} \begin{bmatrix} \Phi_x^L \\ \Phi_u^L \end{bmatrix} = I + \underbrace{\frac{1}{z^L} V + \widehat{\Delta}}_{\widehat{\Delta}_L}$$

where we define $\widehat{\Delta} = \Delta_A \Phi_x^L + \Delta_B \Phi_u^L$. The noting that

$$\begin{aligned} \|\widehat{\Delta}_L\|_{\mathcal{H}_\infty} &\leq \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x^L \\ \varepsilon_{B,2} \Phi_u^L \end{bmatrix} \right\|_{\mathcal{H}_\infty} + \|V\|_2, \\ \|\widehat{\Delta}_L\|_{\mathcal{L}_1} &\leq \left\| \begin{bmatrix} \varepsilon_{A,\infty} \Phi_x^L \\ \varepsilon_{B,\infty} \Phi_u^L \end{bmatrix} \right\|_{\mathcal{L}_1} + \|V\|_\infty, \end{aligned}$$

by (2.2), the system is stabilized and the true system trajectory is given by

$$\mathbf{x} = \Phi_x(I + \widehat{\Delta}_L)^{-1} \mathbf{w}, \quad \mathbf{u} = \Phi_u(I + \widehat{\Delta}_L)^{-1} \mathbf{w}.$$

The rest of the proof follows exactly as in the proof of Theorem 2.2 □

To bound the sub-optimality of this robust synthesis, we need to connect the optimal controller to the robust problem. For this feasibility result, it is necessary to understand the system response induced by putting the optimal robustly constrained controller $\mathbf{K}_c = \mathbf{\Phi}_u^c(\mathbf{\Phi}_x^c)^{-1}$, as defined in (3.2), in feedback with the estimated dynamics (\hat{A}, \hat{B}) . We denote this system responses as $\mathbf{\Phi}_x^c(I - \mathbf{\Delta})^{-1}$. As before, $\mathbf{\Delta} = -\Delta_A \mathbf{\Phi}_x^c - \Delta_B \mathbf{\Phi}_u^c$.

We thus define constants related to the decay rate of the system responses. Let C_\star and ρ_\star be any constants such that $\|\Phi_x^c(k)\|_p \leq C_\star \rho_\star^k$ for $p \in \{2, \infty\}$. Note that these constants must exist since the optimal closed loop will be exponentially stable. Let C_Δ and ρ_Δ similarly be any constants that bound the norm of the frequency response elements of $(I - \mathbf{\Delta})^{-1}$. We note that these constants can be bounded by multiples of C_\star and ρ_\star , as worked out in Appendix G of Dean et al. [7].

Lemma A.4. *Suppose that for $p \in \{2, \infty\}$, $\|\Phi_x^c(k)\|_p \leq C_\star \rho_\star^k$ and $\|(I - \mathbf{\Delta})^{-1}[k]\|_p \leq C_\Delta \rho_\Delta^k$. Define $C = \frac{C_\star C_\Delta}{\log(\frac{\max(\rho_\star, \rho_\Delta)+1}{2 \max(\rho_\star, \rho_\Delta)})}$ and $\rho = \frac{\max(\rho_\star, \rho_\Delta)+1}{2}$. Then for $p \in \{2, \infty\}$ and all $t \geq 0$,*

$$\|\Phi_x^c(I - \mathbf{\Delta})^{-1}[k]\|_p \leq C \rho^k.$$

We defer the proof of this lemma to the end of the section. Next, we redefine the optimal robustly constrained responses $(\mathbf{\Phi}_x^c, \mathbf{\Phi}_u^c)$ in (3.2) to satisfy the modified doubly robust constraint sets

$$\bar{G}_x^\zeta(\mathbf{\Phi}_x; k)_j = F_{x,j}^\top \Phi(k+1)x_0 + \frac{\sigma_w}{1 - \zeta_\infty} \|F_{x,j}^\top \Phi_x[k : 1]\|_1 + \frac{3\zeta_\infty}{1 - 3\zeta_\infty} \sigma_w c_0 \|F_{x,j}^\top \Phi_x[k+1 : 1]\|_1,$$

Here, we have changed a factor of 2 to a factor of 3 in the last term to account for the finite approximation. We will define M_ζ as in (3.3) using the new $(\mathbf{\Phi}_x^c, \mathbf{\Phi}_u^c)$. We are now ready to state the sub-optimality result.

Theorem A.5. *Suppose that the truncation length is such that*

$$L \geq \frac{\log \left(C \max\left(\frac{1-\sqrt{2}\zeta_2}{\sqrt{2}\zeta_2}, \frac{1-\zeta_\infty}{\zeta_\infty}\right) \right)}{\log(1/\rho)} - 1.$$

As long as $\zeta_\infty < \frac{1}{3}$ and $\zeta_2 < \frac{1}{3\sqrt{2}}$, we have that the cost achieved by $\hat{\mathbf{K}}(L) = \hat{\mathbf{\Phi}}_u \hat{\mathbf{\Phi}}_x^{-1}$ synthesized from the minimizers of $\min_{\gamma, \tau} \hat{J}_L(\gamma, \tau)$ satisfies

$$\frac{J(A_\star, B_\star, \hat{\mathbf{K}}_L) - J_\star}{J_\star} \leq 6\sqrt{2}(1 + M_\zeta)(\varepsilon_{A,2}\|\mathbf{\Phi}_x^\star\|_{\mathcal{H}_\infty} + \varepsilon_{B,2}\|\mathbf{\Phi}_u^\star\|_{\mathcal{H}_\infty}) + M_\zeta.$$

To prove this result, we first construct a feasible solution to the truncated synthesis problem (A.2).

Lemma A.6. *Under the conditions of Theorem A.5, the following is a feasible solution to (A.2):*

$$\begin{aligned} \tilde{\mathbf{\Phi}}_x &= \mathbf{\Phi}_x^c(I - \mathbf{\Delta})^{-1}[1 : L], \quad \tilde{\mathbf{\Phi}}_u = \mathbf{\Phi}_u^c(I - \mathbf{\Delta})^{-1}[1 : L], \quad \tilde{V} = \mathbf{\Phi}_x^c(I - \mathbf{\Delta})^{-1}[L+1], \\ \tilde{\gamma} &= \frac{2\sqrt{2}\zeta_2}{1 - \sqrt{2}\zeta_2}, \quad \tilde{\tau} = \frac{2\zeta_\infty}{1 - \zeta_\infty}. \end{aligned}$$

Proof of Lemma A.6. First, we consider the affine constraint,

$$\begin{bmatrix} zI - \hat{A} & -\hat{B} \end{bmatrix} \begin{bmatrix} \tilde{\Phi}_x \\ \tilde{\Phi}_u \end{bmatrix} = I + \frac{1}{z^L} V.$$

Applying the observation in (A.1),

$$\begin{aligned} \begin{bmatrix} zI - \hat{A} & -\hat{B} \end{bmatrix} \begin{bmatrix} \tilde{\Phi}_x \\ \tilde{\Phi}_u \end{bmatrix} = I + \frac{1}{z^L} V &\iff \begin{bmatrix} zI - \hat{A} & -\hat{B} \end{bmatrix} \begin{bmatrix} \Phi_x^c \\ \Phi_u^c \end{bmatrix} (I - \Delta)^{-1} = I \\ &\iff \left(\begin{bmatrix} zI - A_\star & -B_\star \end{bmatrix} + \begin{bmatrix} \Delta_A & \Delta_B \end{bmatrix} \right) \begin{bmatrix} \Phi_x^c \\ \Phi_u^c \end{bmatrix} (I - \Delta)^{-1} = I \end{aligned}$$

The right hand side of this equation is true by the definition of Δ and the affine constraint on Φ_x^c and Φ_u^c , and thus the affine constraint on $\tilde{\Phi}_x$ and $\tilde{\Phi}_u$ is satisfied.

Next, we bound the norm of \tilde{V} . Using Lemma A.4, we have that

$$\|\Phi_x^c(I - \Delta)^{-1}[L+1]\|_p \leq C\rho^{L+1}.$$

Next, we consider the \mathcal{H}_∞ norm constraint

$$\begin{aligned} \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \tilde{\Phi}_x \\ \varepsilon_{B,2} \tilde{\Phi}_u \end{bmatrix} \right\|_{\mathcal{H}_\infty} + \|V\|_2 &\leq \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x^c \\ \varepsilon_{B,2} \Phi_u^c \end{bmatrix} (I - \Delta)^{-1} \right\|_{\mathcal{H}_\infty} + C\rho^{L+1} \\ &\leq \sqrt{2}\zeta_2 \frac{1}{1 - \|\Delta\|_{\mathcal{H}_\infty}} + \frac{\sqrt{2}\zeta_2}{1 - \sqrt{2}\zeta_2} \\ &\leq \frac{2\sqrt{2}\zeta_2}{1 - \sqrt{2}\zeta_2} = \tilde{\gamma}. \end{aligned}$$

The first inequality follows from the bound on $\|\tilde{V}\|_2$ and because the proposed $\tilde{\Phi}_x$ is a truncation of $\Phi_x^c(I - \Delta)^{-1}$ and similarly for $\tilde{\Phi}_u$. The second follows by the sub-multiplicative property, because $\|\Delta\|_{\mathcal{H}_\infty} \leq \zeta_2 < 1$, and due to the assumption on L . Then similarly considering the \mathcal{L}_1 norm constraint,

$$\begin{aligned} \left\| \begin{bmatrix} \varepsilon_{A,\infty} \tilde{\Phi}_x \\ \varepsilon_{B,\infty} \tilde{\Phi}_u \end{bmatrix} \right\|_{\mathcal{L}_1} + \|V\|_\infty &\leq \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x^c \\ \varepsilon_{B,2} \Phi_u^c \end{bmatrix} (I - \Delta)^{-1} \right\|_{\mathcal{L}_1} + C\rho^{L+1} \\ &\leq \zeta_\infty \frac{1}{1 - \|\Delta\|_{\mathcal{L}_1}} + \frac{\zeta_\infty}{1 - \zeta_\infty} \\ &\leq \frac{2\zeta_\infty}{1 - \zeta_\infty} = \tilde{\tau}. \end{aligned}$$

Then it remains only to show that the robust state and input constraints are satisfied. Notice that $\tilde{\Phi}_x$ is a truncation of $\Phi_x^c(1 - \Delta)^{-1}$. For $k > L$, the frequency response elements are

zero, and therefore the constraints are trivially satisfied. For $0 \leq k \leq L$, we have, by Lemma A.1, using block-Toeplitz notation, that

$$\tilde{\Phi}(k) = \Phi_x^c(k) + \Phi_x^c[k : 1]D(1 : k)$$

where we take $D(k)$ to represent the frequency response elements of $\Delta(1 - \Delta)^{-1}$. Then the proof of constraint satisfaction follows as in Section A.1, up to the computations with the definition of $\tilde{\tau}$

$$\frac{\tilde{\tau}}{1 - \tilde{\tau}} = \frac{2\zeta_\infty}{(1 - \zeta_\infty)(1 - \frac{\zeta_\infty}{1 - \zeta_\infty})} = \frac{2\zeta_\infty}{1 - 3\zeta_\infty}.$$

This affects the constant factors around the final term of the sum

$$\frac{\zeta_\infty}{1 - \zeta_\infty} \|x_0\|_\infty + \frac{2\zeta_\infty}{1 - 3\zeta_\infty} \frac{c_0\sigma_w}{1 - \zeta_\infty} \leq \frac{1}{1 - \zeta_\infty} \left(\zeta_\infty + \frac{2\zeta_\infty}{1 - 3\zeta_\infty} \right) c_0\sigma_w = \frac{3\zeta_\infty}{1 - 3\zeta_\infty} c_0\sigma_w$$

Thus, we see that $G_x^{\tilde{\tau}}(\tilde{\Phi}_x; k)_j \leq \bar{G}_x^\zeta(\Phi_x^c; k)_j \leq b_j$ due to the constraints on Φ_x^c . A similar computation with the input constraints shows the same result. Therefore, the proposed solution is feasible. \square

We are now ready to prove the main sub-optimality result.

Proof of Theorem A.5. Recall that we denote the minimizers of $\min_{\gamma, \tau}$ (A.2) as $(\hat{\Phi}_x, \hat{\Phi}_u, \hat{V}, \hat{\gamma}, \hat{\tau})$. Let

$$\hat{\Delta}_L := \hat{\Delta} + \frac{1}{z^L} \hat{V}.$$

Then we have that that

$$\|\hat{\Delta}_L\|_{\mathcal{H}_\infty} \leq \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \hat{\Phi}_x \\ \varepsilon_{B,2} \hat{\Phi}_u \end{bmatrix} \right\|_{\mathcal{H}_\infty} + \|\hat{V}\|_2 \leq \hat{\gamma},$$

By the constraints of the optimization problem in (A.2).

We now apply observation A.1 and (2.2) with $\hat{\Delta}_L$ to characterize the response achieved by the FIR approximate controller $\hat{\mathbf{K}}(L)$ on the true system (A_\star, B_\star) , giving the following:

$$\begin{aligned} J(A_\star, B_\star, \hat{\mathbf{K}}(L)) &= \left\| \begin{bmatrix} Q^{\frac{1}{2}} & \\ & R^{\frac{1}{2}} \end{bmatrix} \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u \end{bmatrix} (I + \hat{\Delta}_L)^{-1} \right\|_{\mathcal{H}_2} \\ &\leq \frac{1}{1 - \hat{\gamma}} \left\| \begin{bmatrix} Q^{\frac{1}{2}} & \\ & R^{\frac{1}{2}} \end{bmatrix} \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u \end{bmatrix} \right\|_{\mathcal{H}_2}. \end{aligned}$$

The inequality follows because $\|\hat{\Delta}\|_{\mathcal{H}_\infty} \leq \hat{\gamma} < 1$.

Denote by $(\tilde{\Phi}_x, \tilde{\Phi}_u, \tilde{V}, \tilde{\gamma}, \tilde{\tau})$ the feasible solution constructed in Lemma A.6. Then,

$$\begin{aligned} \frac{1}{1-\tilde{\gamma}} \left\| \begin{bmatrix} Q^{\frac{1}{2}} & \\ & R^{\frac{1}{2}} \end{bmatrix} \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u \end{bmatrix} \right\|_{\mathcal{H}_2} &\leq \frac{1}{1-\tilde{\gamma}} \left\| \begin{bmatrix} Q^{\frac{1}{2}} & \\ & R^{\frac{1}{2}} \end{bmatrix} \begin{bmatrix} \tilde{\Phi}_x \\ \tilde{\Phi}_u \end{bmatrix} \right\|_{\mathcal{H}_2} = \frac{1}{1-\tilde{\gamma}} J_L(\hat{A}, \hat{B}, \mathbf{K}_c) \\ &\leq \frac{1}{1-\tilde{\gamma}} J(\hat{A}, \hat{B}, \mathbf{K}_c) \leq \frac{1}{1-\tilde{\gamma}} \frac{1}{1-\|\Delta\|_2} J(A_\star, B_\star, \mathbf{K}_c) \\ &\leq \frac{1}{1-\tilde{\gamma}} \frac{1}{1-\sqrt{2}\zeta_2} (1+M_\zeta) J(A_\star, B_\star, \mathbf{K}_\star), \end{aligned}$$

where the first inequality follows from the optimality of $(\hat{\Phi}_x, \hat{\Phi}_u, \hat{V}, \hat{\gamma})$, the equality and second inequality from the fact that $(\tilde{\Phi}_x, \tilde{\Phi}_u)$ are truncations of the response of \mathbf{K}_c on (\hat{A}, \hat{B}) to the first L time steps. The second to last inequality follows from an application of (2.2) with the roles of the nominal and true systems switched. The final inequality follows by definition of the robustness cost gap and bounding $\|\Delta\|_2$ by $\sqrt{2}\zeta_2$. We then simplify

$$\frac{1}{1-\tilde{\gamma}} \frac{1}{1-\sqrt{2}\zeta_2} = \frac{1}{1-\frac{2\sqrt{2}\zeta_2}{1-\sqrt{2}\zeta_2}} \frac{1}{1-\sqrt{2}\zeta_2} = \frac{1}{1-3\sqrt{2}\zeta_2} \leq 6\sqrt{2}\zeta_2$$

which follows for $\zeta_2 < \frac{1}{6\sqrt{2}}$. Then finally,

$$\frac{J(A_\star, B_\star, \hat{\mathbf{K}}(L)) - J_\star}{J_\star} \leq 6\sqrt{2}(1+M_\zeta)\zeta_2 + M_\zeta.$$

□

Proof of Lemma A.4. Using Lemma A.1 with $D(t)$ representing the frequency response elements of $(I - \Delta)^{-1}$, we have that

$$\begin{aligned} \|\hat{\Phi}_x^c(k)\|_p &= \|\Phi_x^c(I - \Delta)^{-1}[k]\|_p \leq \sum_{t=1}^k \|\Phi_x^*(t)\|_p \|D(k-t)\|_p \\ &\leq \sum_{t=1}^k C_\star \rho_\star^t C_\Delta \rho_\Delta^{k-t} \leq C_\star C_\Delta k \max(\rho_\star, \rho_\Delta)^k. \end{aligned}$$

Then noting that for any $x \in (0, 1)$, $\alpha > 0$, and $k \geq 1$,

$$\log(kx^k) = \log(k) + k \log(x) \leq \frac{1}{\alpha} k + \log(\alpha) + k \log(x),$$

which is true because $\log\left(\frac{k}{\alpha}\right) \leq \frac{k}{\alpha}$, we have for $p = 2, \infty$

$$\|\hat{\Phi}_x^c(k)\|_p \leq C_\star C_\Delta \alpha \left(e^{1/\alpha} \max(\rho_\star, \rho_\Delta)\right)^k \leq \frac{C_\star C_\Delta}{\log\left(\frac{\max(\rho_\star, \rho_\Delta)+1}{2 \max(\rho_\star, \rho_\Delta)}\right)} \left(\frac{\max(\rho_\star, \rho_\Delta)+1}{2}\right)^k.$$

The final inequality follows from choosing $\alpha = \log\left(\frac{\max(\rho_\star, \rho_\Delta)+1}{2 \max(\rho_\star, \rho_\Delta)}\right)^{-1}$.

□

B Additive Disturbance Alternative

In this section we discuss an alternative to the proposed robust control problem in (2.3) in which the constraints are tightened using a traditional additive disturbance approximation. We will contrast this additive reduction with the system-response reduction presented in Section 2.2. Define the additive robust synthesis problem as

$$\begin{aligned} \min_{\Phi_x, \Phi_u} \quad & \frac{1}{1-\gamma} J(\hat{A}, \hat{B}, \mathbf{K}) \\ \text{s.t.} \quad & \begin{bmatrix} zI - \hat{A} & -\hat{B} \end{bmatrix} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = I, \quad \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x \\ \varepsilon_{B,2} \Phi_u \end{bmatrix} \right\|_{\mathcal{H}_\infty} \leq \gamma, \\ & G_x^\varepsilon(\Phi_x; k) \leq b_x, \quad G_u^\varepsilon(\Phi_u; k) \leq b_u \quad \forall k \geq 1. \end{aligned} \tag{B.1}$$

where $0 \leq \gamma < 1$ is a fixed parameter and

$$\begin{aligned} G_x^\varepsilon(\Phi_x; k)_j &= G_x(\Phi_x; k)_j + (\varepsilon_{A,\infty} r_x + \varepsilon_{B,\infty} r_u) \|F_{x,j}^\top \Phi_x[k:1]\|_1, \\ G_u^\varepsilon(\Phi_u; k)_j &= G_u(\Phi_u; k)_j + (\varepsilon_{A,\infty} r_x + \varepsilon_{B,\infty} r_u) \|F_{u,j}^\top \Phi_u[k:1]\|_1, \end{aligned}$$

where we define r_x and r_u as the ℓ_∞ radius of the constraint set, i.e. $r_x = \max_{x \in \mathcal{X}} \|x\|_\infty$ and similarly for r_u .

Theorem B.1. *Any controller designed from a feasible solution to the robust control problem (B.1) for any $0 \leq \gamma < 1$ will stabilize the true system. Furthermore, the state and input constraints will be satisfied.*

Proof. First, note that by the system norm constraint,

$$\|\hat{\Delta}\|_{\mathcal{H}_\infty} \leq \sqrt{2} \left\| \begin{bmatrix} \varepsilon_{A,2} \Phi_x \\ \varepsilon_{B,2} \Phi_u \end{bmatrix} \right\|_{\mathcal{H}_\infty} < 1,$$

so the closed-loop system is stable.

Notice that the dynamics are given by

$$x_{k+1} = \hat{A}x_k + \hat{B}u_k + w_k + \Delta_A x_k + \Delta_B u_k.$$

We will make an additive disturbance approximation by defining $\tilde{w}_k = w_k + \Delta_A x_k + \Delta_B u_k$. Though it is conservative, we can bound

$$\|\tilde{w}_k\|_\infty \leq \sigma_w + \max_{\substack{x \in \mathcal{X} \\ \|\Delta_A\|_\infty \leq \varepsilon_{A,\infty}}} \|\Delta_A x_k\|_\infty + \max_{\substack{u \in \mathcal{U} \\ \|\Delta_B\|_\infty \leq \varepsilon_{B,\infty}}} \|\Delta_B u_k\|_\infty.$$

Now note that if $x_k \in \mathcal{X}$, we have

$$\max_{\substack{x \in \mathcal{X} \\ \|\Delta_A\|_\infty \leq \varepsilon_{A,\infty}}} \|\Delta_A x_k\|_\infty = \varepsilon_{A,\infty} \cdot \max_{\substack{x \in \mathcal{X} \\ \|\Delta\|_\infty \leq 1}} \|\Delta x_k\|_\infty = \varepsilon_{A,\infty} \cdot \max_{x \in \mathcal{X}} \|x_k\|_\infty.$$

Thus as long as state and input constraints are satisfied at time k , we have that $\|\tilde{w}_k\|_\infty \leq \sigma_w + \varepsilon_{A,\infty}r_x + \varepsilon_{B,\infty}r_u := \tilde{\sigma}_w$.

Then, notice that the constraint tightening is equivalent to replacing σ_w in (2.1) with $\sigma_w + \varepsilon_{A,\infty}r_x + \varepsilon_{B,\infty}r_u$. Then by Proposition 2.1, the inequality constraints G_x^ε and G_u^ε imply that the system satisfies the constraints for dynamics given by (\hat{A}, \hat{B}) and any process noise bounded by $\tilde{\sigma}_w$. Thus, we can recursively argue that state and input constraints are satisfied for all k as long as they are satisfied for the initial state x_0 . \square

The next result establishes conditions under which this additive disturbance robust control is better than the system-response robust control.

Proposition B.2. *Let $\hat{\tau}$ be an optimizer of (3.1). The additive disturbance reduction is less conservative than the system-response based reduction if and only if*

$$\max(\sigma_w, \|x_0\|_\infty) \frac{\hat{\tau}}{1 - \hat{\tau}} > \varepsilon_{A,\infty}r_x + \varepsilon_{B,\infty}r_u.$$

From this statement, we can see that the system-response reduction suffers for large initial conditions, while the additive reduction suffers with the size of the constraint sets. For this reason, that additive disturbance reduction would not properly characterize the tradeoff between safety and exploration as plotted in Figure 2b. However, we note that the condition in Proposition B.2 could be used as a switch in a control synthesis problem to decide which reduction to use.

Proof. Recall that $\mathbf{w} = x_0z + \sum_{k=0}^{\infty} w_k z^{-k}$. The system-response reduction writes

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u \end{bmatrix} (I + \hat{\Delta})^{-1} \mathbf{w}.$$

This expression can be viewed as expanding the disturbance by the uncertainty as

$$\tilde{\mathbf{w}}_{SLS} = (I + \hat{\Delta})^{-1} \mathbf{w} = \mathbf{w} + \hat{\Delta}(I + \hat{\Delta})^{-1} \mathbf{w}$$

In the robust constraints (2.3), this additional uncertainty is bounded as

$$\|\hat{\Delta}(I + \hat{\Delta})^{-1} \mathbf{w}\|_\infty \leq \frac{\|\hat{\Delta}\|_{\mathcal{L}_1}}{1 - \|\hat{\Delta}\|_{\mathcal{L}_1}} \max(\sigma_w, \|x_0\|_\infty) \leq \frac{\hat{\tau}}{1 - \hat{\tau}} \max(\sigma_w, \|x_0\|_\infty)$$

On the other hand, the additive disturbance approximation notices that $x_{k+1} = \hat{A}x_k + \hat{B}u_k + w_k + \Delta_A x_k + \Delta_B u_k$, and thus assigns $\tilde{w}_k = w_k + \Delta_A x_k + \Delta_B u_k$, i.e.

$$\tilde{\mathbf{w}}_{add} = \mathbf{w} + \Delta_A \mathbf{x} + \Delta_B \mathbf{u}.$$

Then in the robust constraints (B.1), this additional uncertainty is bounded as

$$\|\Delta_A \mathbf{x} + \Delta_B \mathbf{u}\|_\infty \leq \varepsilon_{A,\infty}r_x + \varepsilon_{B,\infty}r_u.$$

Then the additive bound is worse than the system-response bound as stated. \square

Finally, we note that the sub-optimality result could be extended to controllers synthesized from minimizing (B.1) over γ , and that the main difference would arise from the definition of the robust cost sub-optimality gap M_ζ .

C Learning Results

First, we prove a simple small-ball result for random variables with finite fourth moments.

Proposition C.1. *Let $X \in \mathbb{R}$ be a zero-mean random variable with finite fourth moment, which satisfies the conditions*

$$\mathbb{E}[X^4] \leq C(\mathbb{E}[X^2])^2.$$

Let $a \in \mathbb{R}$ be a fixed scalar and $\theta \in (0, 1)$. We have that

$$\mathbb{P}\{|a + X| \geq \sqrt{\theta \mathbb{E}[X^2]}\} \geq (1 - \theta)^2 / \max\{4, 3C\}.$$

Proof. First, we note that we can assume $a \geq 0$ without loss of generality, since we can perform a change of variables $X \leftarrow -X$. We have that $\mathbb{E}[(a + X)^4] = a^4 + 6a^2\mathbb{E}[X^2] + 4a\mathbb{E}[X^3] + \mathbb{E}[X^4]$. Therefore, by the Paley-Zygmund inequality and Young's inequality, we have that

$$\begin{aligned} \mathbb{P}\{|a + X| \geq \sqrt{\theta \mathbb{E}[X^2]}\} &\geq \mathbb{P}\{|a + X| \geq \sqrt{\theta \mathbb{E}[(a + X)^2]}\} \\ &= \mathbb{P}\{(a + X)^2 \geq \theta \mathbb{E}[(a + X)^2]\} \\ &\geq (1 - \theta)^2 \frac{(\mathbb{E}[(a + X)^2])^2}{\mathbb{E}[(a + X)^4]} \\ &= (1 - \theta)^2 \frac{a^4 + 2a^2\mathbb{E}[X^2] + (\mathbb{E}[X^2])^2}{a^4 + 8a^2\mathbb{E}[X^2] + 3\mathbb{E}[X^4]}. \end{aligned}$$

Now define the function $f(a)$ for $a \geq 0$ as

$$f(a) = \frac{a^4 + 2a^2\mu + \mu^2}{a^4 + 8a^2\mu + 3\beta}.$$

Clearly $f(0) = \mu^2/(3\beta)$ and $f(\infty) = 1$. Since by Jensen's inequality we know that $\mu^2 \leq \beta$, this means $f(0) < f(\infty)$. On the other hand,

$$\frac{d}{da}f(a) = \frac{4a(a^2 + \mu)(\mu(3a^2 - 4\mu) + 3\beta)}{(a^4 + 8a^2\mu + 3\beta)^2}.$$

Assume that $\mu \neq 0$ (otherwise the claim is trivially true). The only critical points of the function f on non-negative reals is at $a = 0$ and $a = \sqrt{\frac{4\mu^2 - 3\beta}{3\mu}}$ if $4\mu^2 \geq 3\beta$. If $4\mu^2 < 3\beta$, then $a = 0$ is the only critical point of f , so $f(a) \geq f(0) = \mu^2/(3\beta) \geq 1/(3C)$. On the other hand, if $4\mu^2 \geq 3\beta$ holds, Some algebra yields that

$$f(\sqrt{(4\mu^2 - 3\beta)/(3\mu)}) = \frac{7\mu^2 - 3\beta}{16\mu^2 - 3\beta} \geq \inf_{\gamma \in [3, 4]} \frac{7 - \gamma}{16 - \gamma} = 1/4.$$

The inequality above holds since $3\beta \in [3\mu^2, 4\mu^2]$. Hence,

$$f(a) \geq \min\{1/4, 1/(3C)\} = 1/\max\{4, 3C\}.$$

□

This gives us the main new technical ingredient needed to prove Theorem 4.1.

Proposition C.2. *Let $\mu \in \mathbb{R}^d$ and $M \in \mathbb{R}^{d \times d}$ be a full rank matrix. Let $w \in \mathbb{R}^d$ be a random vector such that its fourth moment is finite and each coordinate w_i is independent and zero-mean, i.e. $\mathbb{E}[w_i] = 0$. Then, for any fixed $v \in \mathbb{R}^d$,*

$$\mathbb{P}_w \left(|\langle v, \mu + Mw \rangle| \geq \sqrt{\lambda_{\min}(M \Sigma M^\top)/2} \right) \geq \frac{1}{C \cdot C_w},$$

where $\Sigma = \mathbb{E}_w[ww^\top]$, C is an absolute constant, and $C_w = \max_{1 \leq i \leq d} \mathbb{E}[w_i^4]/(\mathbb{E}[w_i^2])^2$. In particular, we can take $C = 192$.

Proof. Fix a $q \in \mathbb{R}^d$. By Rosenthal's inequality, we have that $\mathbb{E}[|\langle q, w \rangle|^4] \leq C \cdot C_w \mathbb{E}[|\langle q, w \rangle|^2]^2$ for an absolute constant—we can take $C = 4$ (see e.g. [30]). We now apply Proposition C.1 with $a = \langle v, \mu \rangle$, $X = \langle M^\top v, w \rangle$, $C = 4C_w$, and $\theta = 1/2$. The claim now follows. \square

With Proposition C.2 in place, the proof of Theorem 4.1 is identical to that of Proposition C.1 in Dean et al. [7], except in the final step of establishing the block martingale condition: instead of using the small-ball calculation for Gaussian distributions we replace it with the estimate provided by Proposition C.2.