

Performance Guarantees in Learning and Robust Control

Ross Boczar



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2019-32

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-32.html>

May 13, 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

I am deeply indebted to many people for the completion of this dissertation: Eric Jonas, Evan Sparks, Stephen Tu, and the whole Modest Yachts group; Laurent Lessard and other mentors; my family and friends; and my advisor Ben Recht, for his continued support over the last five years. Finally, I'd like to thank my loving partner Maggie—the brains of this outfit—without whom this would not have been possible.

Performance Guarantees in Learning and Robust Control

by

Ross J Boczar

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Benjamin Recht, Chair

Professor Moritz Hardt

Professor Francesco Borrelli

Spring 2019

Performance Guarantees in Learning and Robust Control

Copyright 2019
by
Ross J Boczar

Abstract

Performance Guarantees in Learning and Robust Control

by

Ross J Boczar

Doctor of Philosophy in Engineering – Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Benjamin Recht, Chair

As the systems we control become more complex, first-principle modeling becomes either impossible or intractable, motivating the use of machine learning techniques for the control of systems with continuous action spaces. As impressive as the empirical success of these methods have been, strong theoretical guarantees of performance, safety, or robustness are few and far between. This manuscript takes a step towards such providing such guarantees by establishing finite-data performance guarantees for identifying and controlling fully- or partially-unknown dynamical systems.

In this manuscript, we explore three different viewpoints that each provide different quantitative guarantees of performance. First, we present a generalization of the classical theory of *integral quadratic constraints*. This generalization leads to a tractable computational procedure for finding exponential stability certificates for partially-unknown feedback systems. Second, we present non-asymptotic lower and upper bounds for core problems in the field of system identification. Finally, using the recently developed *system-level synthesis* framework and tools from high-dimensional statistics, we establish finite-sample performance guarantees for robust output-feedback control of an unknown dynamical system.

To Maggie (we did it!)

Contents

Contents	ii
List of Figures	iv
List of Tables	vi
1 Introduction	1
1.1 Overview	1
1.2 Generating Work	2
1.3 Organization	4
2 Notation	5
3 Exponential Stability Analysis using Integral Quadratic Constraints	7
3.1 Introduction	7
3.2 Preliminaries	9
3.3 A Frequency-Domain Condition	10
3.4 Computation	13
3.5 Exponential Rates from Gain Bounds	16
3.6 IQC Library	19
3.7 Examples	25
3.8 An Aside: Application to Optimization Systems	30
3.9 Conclusion	33
4 \mathcal{H}_∞ Bounds for System Estimation	34
4.1 Introduction	34
4.2 Related Work	37
4.3 System Identification of Finite Impulse Responses	40
4.4 Finite Truncation Error Analysis for Stable Systems	46
4.5 Robust Controller Design	47
4.6 Conclusion	50

5	\mathcal{H}_∞ Bounds for Gain Estimation	53
5.1	Introduction	53
5.2	Related Work	54
5.3	Problem Setup and Main Results	54
5.4	Proof of Main Results	57
5.5	Experiments	62
5.6	Conclusion	63
6	Finite-Data Performance Guarantees for the Output-Feedback Control of an Unknown System	66
6.1	Introduction	66
6.2	Preliminaries	67
6.3	System-Level Synthesis	72
6.4	Sample Complexity Bounds	73
6.5	Experiments	80
6.6	Conclusion	84
	Bibliography	85
	Appendix A Supplementary Material for Chapter 3	93
A.1	Proof of Proposition 3.3.1	93
A.2	Proof of Theorem 3.6.1 and Related Extensions	94
A.3	Proof of Theorem 3.6.2	98
A.4	Computational Considerations	101
A.5	Proof of Theorem 3.8.1	103
	Appendix B Supplementary Material for Chapter 4	106
B.1	Details for Monte–Carlo Simulations	106
B.2	Inverting the Chernoff Bound	107
	Appendix C Supplementary Material for Chapter 5	109
	Appendix D Supplementary Material for Chapter 6	111
D.1	Proof of Lemma 6.4.1	111

List of Figures

3.1	Linear time-invariant system G in feedback with a nonlinearity Δ	9
3.2	Illustration of Remark 1.	11
3.3	Modified feedback diagram with additional multipliers and inputs.	11
3.4	Illustration of Proposition 3.3.2.	12
3.5	Augmented LTI system G in feedback with a nonlinearity Δ	16
3.6	Example stiction nonlinearity.	22
3.7	Monotone and odd bounds for unknown nonlinearities.	23
3.8	LTI system G in feedback with the static sigmoidal nonlinearity $\Delta(x)$	25
3.9	Upper bounds on the exponential convergence rate ρ for the system $G_1(z)$ in feedback as in Figure 3.8.	26
3.10	State decay over time of the system $G_1(z)$ in feedback as in Figure 3.8 with $b = 1$ for various initial conditions.	27
3.11	Upper bounds on the exponential convergence rate ρ for the system $G_2(z)$ in feedback as in Figure 3.8.	28
3.12	Plot of the monotone and quasi-odd asymmetric nonlinearity $\phi(x)$ with its associated bounds.	29
3.13	Comparison of monotone and quasi-odd Zames–Falb IQC rate certificates.	29
3.14	Feedback interconnection between a system G and a nonlinearity ϕ	31
3.15	The perturbed system, including additive input noise w	32
4.1	Closed-loop experimental setup.	48
4.2	Loop-shaping curves from the experimental setup of Figure 4.1.	50
4.3	The pointwise frequency μ value for both reference tracking $T_{r \rightarrow e}$ and noise insensitivity $T_{n \rightarrow e}$	51
4.4	Reference tracking behavior of the closed loop with the model G_{fir} and the actual plant G	52
4.5	Reference tracking behavior as the FIR truncation length is varied.	52
5.1	Performance profiles for the plugin, power method A, power method B, and weighted Thompson Sampling estimators (with coefficient decay).	64
5.2	Performance profiles for the plugin, power method A, power method B, and weighted Thompson Sampling estimators (without coefficient decay).	65

6.1	The standard optimal control problem (6.3) for the plant \mathbf{P}	69
6.2	The disturbance rejection problem for a SISO plant \mathbf{G}	70
6.3	The reference tracking problem for a SISO plant \mathbf{G}	71
6.4	FIR approximation length T required to achieve small relative error in the robust performance objective.	81
6.5	“Swarm” plot of relative improvement δJ from using the approximate SLS procedure across multiple random instances of plants and output noise, for number of experiments m	83
6.6	Comparison on upper bound ΔJ and actual suboptimality gap $\widehat{\Delta J}$	84

List of Tables

5.1 Experiment Parameters	63
A.1 Variable Choices for Specific Zames–Falb Proofs	98

Acknowledgments

I am deeply indebted to many people for the completion of this dissertation; almost all of them I will forget to mention, but the ones I managed to remember to thank are listed here:

- Eric Jonas, Evan Sparks, Stephen Tu, and the whole Modest Yachts group, as well as Ahmed El Alaoui and Henry Milner—for being amazing labmates, collaborators, drinking buddies, and friends, and for being the nicest, smartest group of people I have ever met;
- Ali Rahimi, Anders Rantzer, Andy Packard, Ani Adhikari, James Hing, John Platt, Kevin Jamieson, Larry Venetsky, Laurent Lessard, Mahdi Soltanolkotabi, Michael Jordan, Nikolai Matni, Sam Burden, and Shivaram Venkataraman—for generously extending their mentorship and advice, all of which was needed to navigate the past five years (and beyond);
- Dad, Dale, Christian, Jay, Lane, Mom, Nancy, Steve, Timber, and the rest of my family—for their unconditional love and support throughout this process;
- Deremy, Gabe, James T., Prince, RLO, and Weaz—for being there since Day 1;
- Francesco Borrelli, Ken Goldberg, Moritz Hardt, and Murat Arcaç—for graciously serving on my Qualifying Exam and Dissertation Committees;
- Boban Zarkovich, Kattt Atchley, Ria Briggs, and Shirley Salanio—for being wonderful staff who made my job easy;
- Eric, Jon Kuroda, and Vaishaal Shankar—for helping me with computers;
- Alex, Story, and the fine folks at coffeebar—for giving me energy when I needed it;
- and my advisor Ben Recht—for his continued support (emotional and/or financial) over the last five years, for teaching me a thing or two, and for allowing me to have a fun job.

Finally, I'd like to thank my loving partner—the brains of this outfit—without whom this would not have been possible.

Chapter 1

Introduction

1.1 Overview

This manuscript is about dealing with uncertainty. Specifically, it is about analyzing uncertainty in the context of two fields: control theory, and what we deem “learning theory”—a catch-all for machine learning, statistics, and related areas. These fields have both been established for decades (if not centuries), and while they share many techniques, they often prefer to use different tools to describe and deal with uncertainty.

However, these tools are somewhat at odds with each other. There are two points on which they differ; these are not absolute distinctions, but they are commonly seen in the literature. The first point is one of modeling. On the robust control side, we are often controlling a system based on a model derived from first-principles physics or dynamics. This model will naturally be imperfect, but many times it will be sufficient to drive control synthesis. On the other hand, it is much more common—in the learning theory paradigm—to estimate a parametric or black-box model from a large amount of data.

The second point is how model imperfection is dealt with. In control theory, we are frequently concerned about worst-case performance. This concern is usually motivated by safety-critical applications, such as robotics or autonomous vehicles, where we have to be absolutely sure that we can mitigate the worst-case scenarios. Now, contrast this with the learning theory view. In terms of guarantees, while there is a large body of work on adversarial learning, many results are concerned with performance in expectation, or making probabilistic statements rather than absolute ones (such as results holding with exponentially high probability).

Nonetheless, for complex systems we want to control, we naturally want to use viewpoints and tools from both of these communities. In the age of “big data,” physical systems are equipped with a multitude of sensors with the ability to capture copious amounts of data; it would benefit us to run many trials (if we can afford it) to estimate the properties of the system in question. Furthermore, concerning the uncertainty in our estimation, one would hope that we can deal with our modeling errors and noise processes in a graceful way—one

that is quantitative and can give us a guarantee of safety or performance that is not overly conservative. At the very least, we would expect to be able to discern how sensitive our estimation problem is to different system and environment parameters.

Thus, our general model is as follows. We acquire a dataset of trajectories $\mathcal{S} := \{u^{(i)}, y^{(i)}\}$, where the i -th input and output signals (which are, in general, time histories of vectors) are denoted $u^{(i)}$, $y^{(i)}$ respectively. Now, based on this dataset, we want a quantitative guarantee about what the system looks like and/or how the system performs when controlled. However, this general model immediately raises issues. When we only collect a finite amount of data, there are multiple sources of uncertainty that arise. Noise processes (also present in the infinite-data regime) and truncation effects lead to modeling errors which then can propagate through to error for a control task, if one is not too careful. Therefore, what we truly want to accomplish is to calmly take these sources of uncertainty into account during modeling and control design—while also minimizing the uncertainty as much as possible. However, many statistical results for estimation and control design are asymptotic and only say that (say) an estimator is consistent in the regime of infinite data. Since that is quite a long time to wait, we also want to quantify what we lose going from the infinite to the finite—which is generally difficult!

In any case, our finite-data model naturally leads to a two-stage setup, if we want to estimate a model and then robustly control it. First, we “coarsely” identify a model from data, where our model estimate comes with a tractable description of the model uncertainty. Second, we perform robust control synthesis using this nominal model and uncertainty description. As an added bonus, when we want to robustly control these systems in feedback, whether to regulate the state to zero or have an output track a trajectory, we know from both theory and practice that feedback is somewhat robustifying against many different types of uncertainty.

The central questions addressed by this manuscript are therefore as follows:

- Can we get quantitative performance guarantees for control analysis and synthesis based on uncertain models?
- Can we establish any of these guarantees in the environment where we are estimating a model from data?

1.2 Generating Work

As the systems we control become more complex, first-principle modeling becomes either difficult or intractable, motivating the use of machine learning techniques for the control of systems with continuous state and action spaces. As impressive as the empirical success of these methods have been, strong theoretical guarantees of performance, safety, or robustness are often elusive. This manuscript takes steps towards providing such guarantees by establishing finite-data performance guarantees for identifying and controlling fully- or partially-unknown dynamical systems. In this section, we summarize the works that make

up the majority of this manuscript.

1.2.1 Robust Control

Since the 1960s, the field of robust control—broadly, the field of making actions under uncertainty—has been concerned with obtaining quantitative performance guarantees in the presence of bounded uncertainty; this has included breakthroughs such as the small-gain theorem, passivity theory, dissipativity theory, the structured singular value, and Integral Quadratic Constraints (IQCs) [Megretski and Rantzer, 1997]. The most general of these, IQCs, gives a computational procedure to produce a certificate of stability for a linear, time-invariant (LTI) dynamical system placed in feedback with an unknown nonlinearity—a general and useful paradigm often seen in control theory. In Boczar et al. [2015, 2017], we give a computational procedure to now derive a certificate of *exponential* stability for such a feedback system. This tractable procedure—a small semidefinite program (SDP)—allows one to make guarantees about the size of the internal state of the system when only coarse information (say, a norm bound) about the nonlinearity is known. We take advantage of the existing literature by providing an updated library of “exponential IQCs” satisfied by classes of nonlinearities commonly seen in engineering and computer science. This improvement of now being able to compute a certificate for a specific exponential rate is critical: it is analogous to being able to guarantee a rate of convergence for an optimization algorithm. Indeed, these papers were inspired by a recent line of work, starting with Lessard et al. [2016], which uses the theory of IQCs to prove such rates for the most popular optimization methods.

1.2.2 System Identification and Learning Theory

Though the feedback paradigm described in the previous section is often useful, there are often times when we do not even have a model of the system to start with, yet we still ask how to control it—we usually seek an end-to-end identification and control algorithm. How to design such an algorithm has been the question at the heart of the system identification literature. Partially motivated by advances in deep learning, there has also been a deluge of recent work in deriving end-to-end algorithms for various tasks in reinforcement learning. However, sample complexity bounds for these problems are difficult to derive, and it is not clear which quantities accurately measure the “hardness” of a learning problem. In Boczar et al. [2018], we show finite-sample performance guarantees for a specific, simplified end-to-end problem: robust output feedback control of an LTI system. This case is interesting as it captures many common motifs found in practice: a partially-observed state, a rough description of system uncertainty, and a desire for performance guarantees under these constraints.

Achieving these guarantees requires two components. First, it requires tools from non-asymptotic statistics to analyze a “coarse identification” step based on least-squares fitting of input-output data. This step, which we previously developed in Tu et al. [2017] and showed

was optimal (in a sense) in Tu et al. [2018a], returns a simple description of the system under test and a probabilistic quantification of the uncertainty around it. Our second component is a modification of the recently-developed *system-level synthesis* (SLS) framework [Wang et al., 2019]. This framework lifts a nonconvex feedback design problem into a convex one by optimizing over the end-to-end system response. In our analysis, we modify the SLS framework to take into account our uncertainty description from the first step, which allows us to derive non-asymptotic learning rates for an end-to-end identification and control procedure. Specifically, we show how many input/output sequences of a certain length need to be recorded in order to build a descriptive enough model such that the robust SLS procedure can handle the remaining uncertainty in standard feedback controller design tasks.

1.3 Organization

This thesis is organized as follows. Preliminaries and notation are covered in Chapter 2. Chapter 3 (based on Boczar et al. [2017], a superset of Boczar et al. [2015]) establishes the framework of exponential IQCs, including an application for analyzing stochastic optimization algorithms. Chapters 4 and 5 provide upper and lower bounds in the \mathcal{H}_∞ -norm for coarse system identification and gain estimation; these chapters are based on condensed versions of Tu et al. [2017] and Tu et al. [2018a], respectively. Finally, Chapter 6 uses ideas developed in Chapter 4 and the SLS framework to provide end-to-end control guarantees, as seen in Boczar et al. [2018]. Proofs and other supplementary material for these chapters are given in Appendices A, B, C, and D, respectively. Experimental code can be found at <https://laurentlessard.com/public/code> and <https://github.com/rjbczar>.

Chapter 2

Notation

Here we give some of the more common notation in this manuscript.

Linear Algebra and Analysis

\mathbb{C}^d	canonical complex d -dimensional space
\mathbb{R}^d	canonical real d -dimensional space
v^*, A^*	complex conjugate
v^\top, A^\top	transpose
$(g \circ f)(x)$	function composition $g(f(x))$
$\kappa(A)$	condition number, i.e. $\sigma_1(A)/\sigma_n(A)$
$v^{[i]}$	i -th coordinate of the vector v (to occasionally avoid confusion with v_i)
\mathbb{D}	closed unit disk in complex plane $\{z \in \mathbb{C} \mid z \leq 1\}$
e_i	i -th standard basis vector in \mathbb{R}^d
$\mathbf{1}_{\mathcal{A}}$	indicator function on event \mathcal{A}
$\langle A, B \rangle$	inner product $\text{tr}(A^*B)$ (a^*b for vectors)
$\ A\ , \ v\ $	spectral norm (resp. Euclidean norm for vectors, occasionally made explicit as $\ \cdot\ _2$)
$\ f(z)\ _\infty$	supremum norm taken over \mathbb{T} , i.e. $\sup_{z \in \mathbb{T}} f(z) $
\succ, \succeq	positive (resp. semi-)definite matrix inequality
$\rho(A)$	spectral radius, i.e. $\max_i \lambda_i(A) $

$\text{Toep}(u)$	the lower-triangular Toeplitz matrix where the first column is equal to u
\mathbb{T}	unit circle in complex plane $\{z \in \mathbb{C} \mid z = 1\}$ (also parameterized by $z = e^{j\omega}$ with $\omega \in [0, 2\pi)$)
$\text{tr}(A)$	trace, i.e. $\sum_i A_{ii} = \sum_i \lambda_i(A)$

Signals and Systems

$*$	discrete-time convolution operator
$\ G\ _{\mathcal{H}_\infty}$	\mathcal{H}_∞ -norm, the induced $\ell_2 \rightarrow \ell_2$ norm (equal to $\sup_{z \in \mathbb{T}} \ G(z)\ _2$)
ℓ_p	space of p -norm-bounded signals ($\sum_{k=0}^{\infty} x_k ^p < \infty$, $\ x\ _{\ell_\infty} := \max_k x_k $)
ℓ_2	space of square-summable signals ($\sum_{k=0}^{\infty} x_k ^2 < \infty$)
ℓ_2^ρ	space of ρ -square-summable signals ($\sum_{k=0}^{\infty} \rho^{-2k} x_k ^2 < \infty$, $\rho \in (0, 1)$)
$\ x\ _{\ell_2}$	ℓ_2 -norm, i.e. $\sum_{k=0}^{\infty} x_k ^2$ ($\sum_{k=0}^{\infty} \ x_k\ ^2$ in the vector case)
$\mathcal{RH}_\infty^{m \times n}$	set of $m \times n$ matrices whose elements are proper rational functions with real coefficients analytic outside the closed unit disk (superscript usually suppressed for brevity)
(A, B, C, D)	discrete-time linear time-invariant system ($D = 0$ occasionally omitted in the former)

$$\left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right]$$

G	discrete-time linear time-invariant system, frequency-domain representation
x	ℓ_2 signal, frequency-domain representation
x_t	time domain signal (x_0, x_1, \dots) , indexed at time t
$\hat{x}(z)$	z -transform of the time-domain signal x , $\hat{x}(z) := \sum_{k=0}^{\infty} x_k z^{-k}$

Probability

$\mathbb{E}[\cdot]$	expectation
$\mathcal{N}(\mu, \Sigma)$	normally distributed random variable with mean μ and covariance matrix Σ
$\mathbb{P}(\mathcal{E})$	probability of event \mathcal{E}

Chapter 3

Exponential Stability Analysis using Integral Quadratic Constraints

3.1 Introduction

Analysis in the context of robust control is generally concerned with obtaining absolute performance guarantees about a system in the presence of bounded uncertainty. Examples of such results include the small gain theorem and passivity theory [Zames, 1966], dissipativity theory [Willems, 1972], the structured singular value μ [Doyle, 1982], and integral quadratic constraints (IQCs) [Megretski and Rantzer, 1997].

In this chapter, we present a modification of IQC theory, the most general of the aforementioned tools, that allows one to certify *exponential stability* rather than just bounded-input bounded-output (BIBO) stability. Moreover, we can compute numerical bounds on the exponential decay rate of the state. Even when BIBO stable systems are exponentially stable, estimates of the exponential decay rates provided by standard IQC theory are typically very conservative. We will show that this conservatism can be greatly reduced if we directly certify exponential stability and use the method presented herein to compute the associated decay rate.

Our modified IQC analysis was successfully applied in Lessard et al. [2016] to analyze convergence properties of commonly-used optimization algorithms such as the gradient descent method. These algorithms converge at an exponential rate when applied to strongly convex functions, and the modified IQC analysis automatically produces very tight bounds on the convergence rates. Another potential application is in time-critical systems. In embedded model predictive control, for example, it is vital to have robust guarantees that desired error bounds will be met in the allotted time without overflow errors and in spite of fixed-point arithmetic (see Jerez et al. [2014] and references therein).

3.1.1 A Special Case

While a general treatment of exponential bounds is provided in the sequel, it is worth noting that exponential stability can be proven directly for some special cases. To illustrate this fact, consider a linear time-invariant (LTI) discrete-time plant G with state-space realization (A, B, C, D) . Suppose G is connected in feedback with a *strictly-input passive* nonlinearity Δ . A sufficient condition for BIBO stability is that there exists a positive definite matrix $P \succ 0$ and a scalar $\lambda \geq 0$ satisfying the linear matrix inequality (LMI)

$$\begin{bmatrix} A & B \\ I & 0 \end{bmatrix}^\top \begin{bmatrix} P & 0 \\ 0 & -P \end{bmatrix} \begin{bmatrix} A & B \\ I & 0 \end{bmatrix} + \lambda \begin{bmatrix} 0 & C^\top \\ C & D+D^\top \end{bmatrix} \prec 0. \quad (3.1)$$

This result is also related to the Positive Real Lemma (see [Kottenstette and Antsaklis \[2010\]](#) and references therein). If we define $V(x) := x^\top P x$, then (3.1) implies that V decreases along trajectories: $V(x_{k+1}) \leq V(x_k)$ for all k . BIBO stability then follows from positivity and boundedness of V . Observe that when (3.1) holds, we may replace the right-hand side by $-\epsilon P$ for some sufficiently small $\epsilon > 0$. We then conclude that $V(x_{k+1}) \leq (1 - \epsilon)V(x_k)$ for all k and exponential stability follows. We may then maximize ϵ subject to feasibility of (3.1) to further improve the rate bound.

Unfortunately, the approach outlined above of including $-\epsilon P$ fails in the general IQC setting due to the different role played by P in the associated LMI. In IQC theory, the LMI comes from the Kalman–Yakubovich–Popov (KYP) lemma, and although it is structurally similar to (3.1), P is not positive definite in general, and V may not decrease along trajectories. However, our key insight is that by suitably modifying both the LMI *and* the IQC definition, we obtain a more broadly applicable condition for certifying exponential stability.

This chapter is organized as follows. We cover some related work in the remainder of the introduction, we explain our notation and some basic results in Section 3.2, we develop and present our main result in Section 3.3, and we discuss computational considerations in Section 3.4. An explicit construction of the (conservative) rate guarantees implied by finite L_2 gain is given in Section 3.5. In Section 3.6 we provide a library of applicable IQCs. Finally, we present illustrative examples demonstrating the usefulness of our result in Section 3.7 and give an application to stochastic optimization algorithms in Section 3.8. We make some concluding remarks in Section 3.9.

3.1.2 Related Work

It is noted in [Megretski and Rantzer \[1997\]](#) and [Rantzer and Megretski \[1997\]](#) that BIBO stability often implies exponential stability. In particular, exponential stability follows if the nonlinearity satisfies an additional *fading memory* property. So, under mild assumptions, the robust stability guarantee from IQC theory automatically implies exponential stability as well. The proof of this result uses the L_2 gain from the stability analysis to construct an exponential rate bound. However, we will see in Section 3.7 that bounds computed in this way can be very conservative.

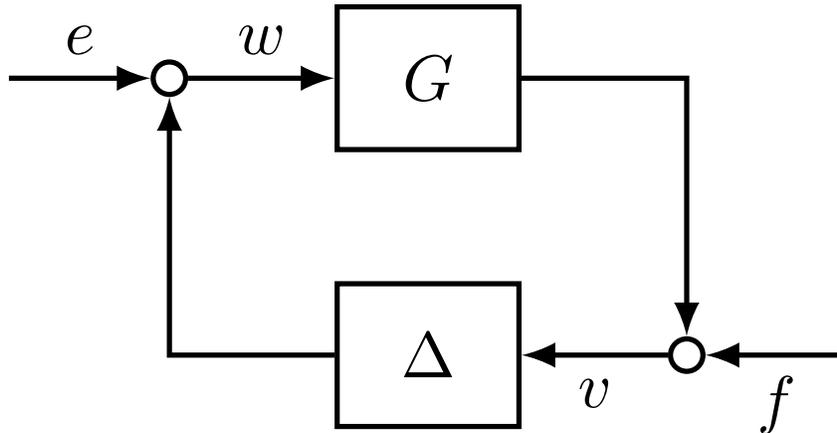


Figure 3.1: Linear time-invariant system G in feedback with a nonlinearity Δ .

Other proofs of exponential stability have appeared in the literature for specific classes of nonlinearities. Some examples include sector-bounded nonlinearities [Corless and Leitmann, 1993; Konishi and Kokame, 1999] and nonlinearities satisfying a Popov IQC [Jönsson, 1997]. These works exploit LMI modifications akin to the one shown in (3.1) earlier in this section.

The analysis in this chapter mirrors the analysis given in Lessard et al. [2016], which presents an approach for proving the robust exponential stability of optimization algorithms. The approach of Lessard et al. [2016] uses a time-domain formulation of IQCs modified to handle exponential stability. In contrast, the present work develops the aforementioned exponential stability analysis entirely in the frequency domain and its applicability is not restricted to the analysis of iterative optimization algorithms. Moreover, we clarify the connection to the seminal IQC results in Megretski and Rantzer [1997]. Parts of this work first appeared in the conference paper Boczar et al. [2015]; since then, an analogous continuous-time formulation with alternative techniques and motivations also appeared in Hu and Seiler [2016].

3.2 Preliminaries

We adopt a setup analogous to the one used in Megretski and Rantzer [1997], with the exception that we will work in discrete time rather than continuous time. Along with the standard notation given in Chapter 2, recall that a sequence $u = (u_0, u_1, \dots)$ is said to be in ℓ_2 if $\sum_{k=0}^{\infty} |u_k|^2 < \infty$. A sequence u_k is said to be in ℓ_2^ρ for some $\rho \in (0, 1)$ if the sequence $(\rho^{-k} u_k)$ is in ℓ_2 , i.e. $\sum_{k=0}^{\infty} \rho^{-2k} |u_k|^2 < \infty$. Note that $\ell_2^\rho \subset \ell_2$. Furthermore, \mathcal{X} -bounded is shorthand for a bounded operator from \mathcal{X} to \mathcal{X} .

Consider the standard setup of Figure 3.1 (known as the *Lur'e system*). The block G contains the known LTI part of the system while Δ contains the part that is uncertain, unknown, nonlinear, or otherwise troublesome. The interconnection is said to be *well-posed*

if the map $(v, w) \mapsto (e, f)$ has a causal inverse. The interconnection is said to be bounded-input bounded-output (BIBO) stable if, in addition, there exists some $\gamma > 0$ such that when G is initialized with zero state,

$$\|v\|_{\ell_2}^2 + \|w\|_{\ell_2}^2 \leq \gamma(\|e\|_{\ell_2}^2 + \|f\|_{\ell_2}^2)$$

for all square-summable inputs f and e . Finally, the interconnection is (internally) *exponentially stable* if there exists some $\rho \in (0, 1)$ and $c > 0$ such that if $f = 0$ and $e = 0$, the state x_k of G will decay exponentially with rate ρ . That is,

$$\|x_k\| \leq c\rho^k \|x_0\| \quad \text{for all } k.$$

We now present the classical IQC definition and stability result, which will be modified in the sequel to guarantee exponential convergence. These results are discrete-time analogs of the main IQC results of Megretski and Rantzer [Megretski and Rantzer \[1997\]](#).

Definition 3.2.1 (IQC). *Signals $y \in \ell_2$ and $u \in \ell_2$ with associated z -transforms $\hat{y}(z)$ and $\hat{u}(z)$ satisfy the IQC defined by a Hermitian complex-valued function Π if*

$$\int_{\mathbb{T}} \begin{bmatrix} \hat{y}(z) \\ \hat{u}(z) \end{bmatrix}^* \Pi(z) \begin{bmatrix} \hat{y}(z) \\ \hat{u}(z) \end{bmatrix} dz \geq 0. \quad (3.2)$$

An ℓ_2 -bounded causal operator Δ satisfies the IQC defined by Π if (3.2) holds for all $y \in \ell_2$ with $u = \Delta(y)$. We also define $\text{IQC}(\Pi(z))$ to be the set of all Δ that satisfy the IQC defined by Π .

Theorem 3.2.1 (Stability result). *Let $G(z) \in \mathcal{RH}_{\infty}^{m \times n}$ and let Δ be a bounded causal operator. Suppose that:*

- i) for every $\tau \in [0, 1]$, the interconnection of G and $\tau\Delta$ is well-posed.*
- ii) for every $\tau \in [0, 1]$, we have $\tau\Delta \in \text{IQC}(\Pi(z))$.*
- iii) there exists $\epsilon > 0$ such that*

$$\begin{bmatrix} G(z) \\ I \end{bmatrix}^* \Pi(z) \begin{bmatrix} G(z) \\ I \end{bmatrix} \preceq -\epsilon I, \quad \forall z \in \mathbb{T}. \quad (3.3)$$

Then, the feedback interconnection of G and Δ is BIBO stable.

3.3 A Frequency-Domain Condition

In this section, we augment Definition 3.2.1 and the classical result of Theorem 3.2.1 to derive a frequency-domain condition that certifies exponential stability.

Definition 3.3.1. The operators ρ_+ , ρ_- are defined as the time-domain, time-dependent multipliers ρ^k, ρ^{-k} , respectively, where $\rho \in (0, 1)$ is a defined constant.

Remark 1. The operator $\rho_- \circ (G(z) \circ \rho_+)$ is equivalent to the operator $G(\rho z)$. This follows from the fact that, for any constant $a > 0$ and signal u_k , the z -transform of $a^{-k}u_k$ is given by $\hat{u}(az)$. See Figure 3.2 for an illustration.

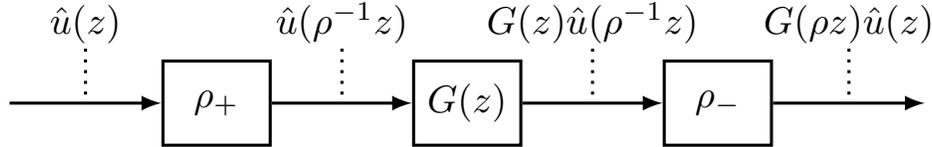


Figure 3.2: Illustration of Remark 1.

In order to show exponential stability of the system in Figure 3.1, we will relate it to BIBO stability of the modified system shown in Figure 3.3. This equivalence is closely related to the theory of stability multipliers [Safonov and Kulkarni, 2000].

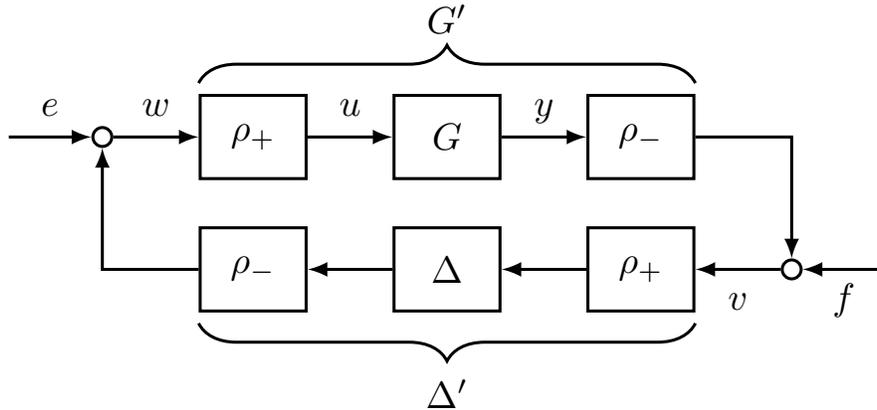


Figure 3.3: Modified feedback diagram with additional multipliers and inputs. For appropriately chosen e and f and with zero initial condition, we show how this diagram is equivalent to that of Figure 3.1.

Proposition 3.3.1. Suppose $G(z)$ has a minimal realization (A, B, C, D) . If the interconnection in Figure 3.3 is BIBO stable, then the interconnection in Figure 3.1 with initial state x_0 is exponentially stable.

Proof. Intuitively, if v and w are small in the BIBO sense compared to e and f , then y must be even smaller. See Appendix A.1 for a detailed proof. ■

In an effort to define IQCs for the transformed system shown in Figure 3.3, we introduce the concept of the ρ -IQC.

Definition 3.3.2 (ρ -IQC). Signals $y \in \ell_2^\rho$ and $u \in \ell_2^\rho$ with associated z -transforms $\hat{y}(z)$ and $\hat{u}(z)$ satisfy the ρ -IQC defined by a Hermitian complex-valued function Π if

$$\int_{\mathbb{T}} \begin{bmatrix} \hat{y}(\rho z) \\ \hat{u}(\rho z) \end{bmatrix}^* \Pi(\rho z) \begin{bmatrix} \hat{y}(\rho z) \\ \hat{u}(\rho z) \end{bmatrix} dz \geq 0. \quad (3.4)$$

An ℓ_2^ρ -bounded causal operator Δ satisfies the ρ -IQC defined by Π if (3.4) holds for all $y \in \ell_2^\rho$ with $u = \Delta(y)$. We also define $\text{IQC}(\Pi(z), \rho)$ to be the set of all Δ that satisfy the ρ -IQC defined by Π .

Note that the concept of a ρ -IQC generalizes that of a regular IQC. Indeed, we have $\text{IQC}(\Pi(z), 1) = \text{IQC}(\Pi(z))$. The restriction of $u \in \ell_2^\rho$ and $y \in \ell_2^\rho$ corresponds to the restriction of $u \in \ell_2$ and $y \in \ell_2$ in the classical definition of IQC [Megretski and Rantzer, 1997]. Now equipped with ρ -IQCs, we can relate Δ' in Figure 3.3 to Δ in Figure 3.1.

Proposition 3.3.2. Let Δ be an ℓ_2^ρ -bounded causal operator, and let Π be a Hermitian complex-valued function. As in Figure 3.3, define $\Delta' := \rho_- \circ (\Delta \circ \rho_+)$, which is equivalently ℓ_2 -bounded. Then, the following statements are equivalent.

- (i) $\Delta \in \text{IQC}(\Pi(z), \rho)$
- (ii) $\Delta' \in \text{IQC}(\Pi(\rho z))$

Proof. We define the discrete Fourier transform of the input and output of Δ as $\hat{y}(z)$ and $\hat{u}(z)$, respectively. Then, from the definition of ρ_+ and ρ_- , we have that $\hat{w}(z) = \hat{u}(\rho z)$ and $\hat{v}(z) = \hat{y}(\rho z)$. Substituting into the IQC definition (3.2), we obtain (3.4) as required. ■

Proposition 3.3.2 is illustrated in Figure 3.4.

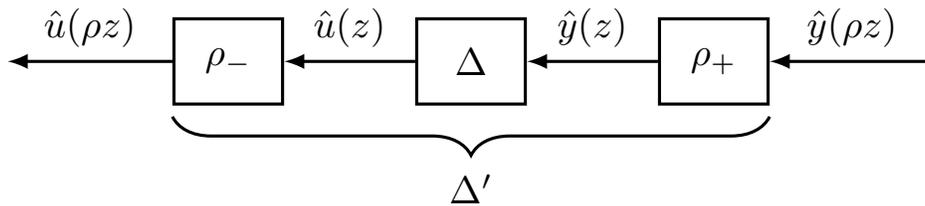


Figure 3.4: Illustration of Proposition 3.3.2.

We now state our main result, an exponential stability theorem analogous to the classical result in Theorem 3.2.1.

Theorem 3.3.1 (Exponential stability). Fix $\rho \in (0, 1)$. Let $G(\rho z) \in \mathcal{RH}_\infty^{m \times n}$ and Δ be an ℓ_2^ρ -bounded causal operator. Furthermore, suppose that:

- i) for every $\tau \in [0, 1]$, the interconnection of G and $\tau\Delta$ is well-posed.

ii) for every $\tau \in [0, 1]$, we have $\tau\Delta \in \text{IQC}(\Pi(z), \rho)$.

iii) there exists $\epsilon > 0$ such that

$$\begin{bmatrix} G(\rho z) \\ I \end{bmatrix}^* \Pi(\rho z) \begin{bmatrix} G(\rho z) \\ I \end{bmatrix} \preceq -\epsilon I, \quad \forall z \in \mathbb{T}. \quad (3.5)$$

Then, the interconnection of G and Δ shown in Figure 3.1 is exponentially stable with rate ρ .

Proof. We apply Theorem 3.2.1 to the interconnection in Figure 3.3 with operators $G(\rho z)$ and Δ' and the IQC $\Pi(\rho z)$.

- (a) Since Figure 3.1 and Figure 3.3 have the same interconnection structure, well-posedness is equivalent.
- (b) Due to the equivalence of IQCs in Proposition 3.3.2,

$$\begin{aligned} \tau\Delta \in \text{IQC}(\Pi(z), \rho) &\iff \rho_- \circ ((\tau\Delta) \circ \rho_+) \in \text{IQC}(\Pi(\rho z)) \\ &\iff \tau(\rho_- \circ (\Delta \circ \rho_+)) \in \text{IQC}(\Pi(\rho z)) \\ &\iff \tau\Delta' \in \text{IQC}(\Pi(\rho z)). \end{aligned}$$

- (c) This is condition iii) of Theorem 3.2.1 using $G(\rho z)$ and Δ' .

Thus, these three conditions ensure BIBO stability of the system in Figure 3.3. We then apply Proposition 3.3.1 to arrive at exponential stability of Figure 3.1.

Note that the assumption $G(\rho z) \in \mathcal{RH}_\infty^{m \times n}$ restricts us to verifying rates that are no faster than the rate of convergence of the plant G operating in open-loop, which corresponds to the largest (in magnitude) pole of $G(z)$. Assuming WLOG that $\Delta(0) = 0$, this is clear as $\Delta \equiv 0$ (corresponding to open-loop G) satisfies any ρ -IQC. ■

3.4 Computation

As in the classical IQC setting, to guarantee stability, the frequency-domain inequality (FDI) (3.5) must be verified for every $\omega \in [0, 2\pi)$. However, if the IQC in question exhibits a particular factorization, then the discrete-time KYP Lemma can be applied to convert the infinite-dimensional FDI to a finite-dimensional LMI. We now review these results.

Definition 3.4.1. We say Π has a factorization (Ψ, M) if

$$\Pi(z) = \Psi(z)^* M \Psi(z),$$

where Ψ is a stable linear time-invariant system, M is a constant Hermitian matrix, and $\Psi(z)^*$ denotes the conjugate transpose of $\Psi(z)$.

Remark 2. Definition 3.4.1 is similar to J -spectral factorization (see Zhang et al. [2001] and references therein), except we require them to hold for arbitrary $z \in \mathbb{C}$. Spectral factorizations are commonly evaluated on the unit circle for discrete systems (c.f. the imaginary axis for continuous-time systems). In such cases, we have $z^* = z^{-1}$ for all $z \in \mathbb{T}$ and $s^* = -s$ for all $s \in j\mathbb{R}$. For this reason, factorizations are conventionally written using the para-Hermitian conjugate defined as $\Psi^\sim(z) := \Psi^\top(z^{-1})$ (c.f. $\Psi^\sim(s) := \Psi^\top(-s)$ for continuous time). Although these definitions are equivalent to $\Psi(z)^*$ (c.f. $\Psi(s)^*$) in general, we cannot use the para-Hermitian conjugate for our factorization because we require it to hold for all $z \in \mathbb{C}$.

Remark 3. If $\Pi(z)$ has a factorization (Ψ, M) and $\Psi(\rho z)$ is stable, then by Parseval's Theorem, (3.4) is equivalent to

$$\sum_{k=0}^{\infty} \rho^{-2k} z_k^\top M z_k \geq 0, \quad \text{where } z := \Psi \begin{pmatrix} y \\ u \end{pmatrix}.$$

The KYP lemma, stated below, is attributed to Kalman, Yakubovich, and Popov. A short proof and further references can be found in Rantzer [1996].

Lemma 3.4.1 (Discrete-time KYP Lemma). *Suppose A, B, M are given matrices where M is Hermitian and A has no eigenvalues on the unit circle. Then the following FDI:*

$$\begin{bmatrix} (zI - A)^{-1}B \\ I \end{bmatrix}^* M \begin{bmatrix} (zI - A)^{-1}B \\ I \end{bmatrix} \prec 0$$

holds for all $z \in \mathbb{T}$ if and only if there exists a $P = P^\top$ and $\lambda \geq 0$ satisfying the LMI

$$\begin{bmatrix} A & B \\ I & 0 \end{bmatrix}^\top \begin{bmatrix} P & 0 \\ 0 & -P \end{bmatrix} \begin{bmatrix} A & B \\ I & 0 \end{bmatrix} + \lambda M \prec 0.$$

An algebraic transformation of the discrete-time KYP lemma then gives a computational certificate for verifying exponential stability rates.

Corollary 3.4.1. *Suppose the realization of G is given by (A, B, C, D) and assume Π has a factorization (Ψ, M) , where the realization of Ψ is given by*

$$\Psi = \left[\begin{array}{c|cc} A_\Psi & B_{\Psi_1} & B_{\Psi_2} \\ \hline C_\Psi & D_{\Psi_1} & D_{\Psi_2} \end{array} \right].$$

Then (3.5) is equivalent to the existence of $P = P^\top$ and $\lambda \geq 0$ such that

$$\begin{bmatrix} \hat{A}^\top P \hat{A} - \rho^2 P & \hat{A}^\top P \hat{B} \\ \hat{B}^\top P \hat{A} & \hat{B}^\top P \hat{B} \end{bmatrix} + \lambda \begin{bmatrix} \hat{C}^\top \\ \hat{D}^\top \end{bmatrix} M \begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix} \prec 0 \quad (3.6)$$

where $(\hat{A}, \hat{B}, \hat{C}, \hat{D})$ are defined as

$$\left[\begin{array}{c|c} \hat{A} & \hat{B} \\ \hline \hat{C} & \hat{D} \end{array} \right] := \left[\begin{array}{cc|c} A & 0 & B \\ B_{\Psi_1}C & A_{\Psi} & B_{\Psi_2} + B_{\Psi_1}D \\ \hline D_{\Psi_1}C & C_{\Psi} & D_{\Psi_2} + D_{\Psi_1}D \end{array} \right].$$

Proof. Related to a derivation in Seiler [2015], one can show that

$$\begin{bmatrix} G(z) \\ I \end{bmatrix}^* \Pi(z) \begin{bmatrix} G(z) \\ I \end{bmatrix} = [\star]^* M \begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix} \begin{bmatrix} (zI - \hat{A})^{-1} \hat{B} \\ I \end{bmatrix}$$

where \star denotes the repeated part of the quadratic form surrounding M . Similarly, we have

$$\begin{aligned} \rho^{-2} \begin{bmatrix} G(\rho z) \\ I \end{bmatrix}^* \Pi(\rho z) \begin{bmatrix} G(\rho z) \\ I \end{bmatrix} &= [\star]^* \rho^{-2} M \begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix} \begin{bmatrix} (\rho z I - \hat{A})^{-1} \hat{B} \\ I \end{bmatrix} \\ &= [\star]^* \rho^{-2} M \begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix} \begin{bmatrix} (zI - \rho^{-1} \hat{A})^{-1} \rho^{-1} \hat{B} \\ I \end{bmatrix}. \end{aligned}$$

If $\rho^{-1} \hat{A}$ has no eigenvalues on the unit circle, we may then invoke Lemma 3.4.1 (applied to $\rho^{-1} \hat{A}$, $\rho^{-1} \hat{B}$, and the appropriate M term) and multiply through by ρ^2 to show that (3.5) is equivalent to the existence of $P = P^\top$ and $\lambda \geq 0$ such that (3.6) holds, as required. ■

With the advent of fast interior-point methods to solve LMIs, the feasibility of the LMI (3.6) can often be quickly ascertained for any fixed ρ^2 . Since the size of the LMI is often on the order of the size of the system G and the IQC Π , many practical linear systems lead to LMIs of relatively moderate size.

Finding the best upper bound amounts to minimizing ρ^2 subject to (3.6) being feasible. This type of problem occurs frequently in robust control and is known as a *generalized eigenvalue optimization problem* (GEVP) [Boyd et al., 1994]. The GEVP is not an LMI because (3.6) is not jointly linear in ρ^2 and P . One simple approach to solving the GEVP is to perform a bisection search on ρ^2 , but there are more sophisticated methods available; see for example Boyd and El Ghaoui [1993]. However, to be precise, applying a bisection search on ρ^2 requires the ρ -IQC to obey a certain monotonicity property, which we now define.

Definition 3.4.2 (Monotonicity). *We say an IQC $\Pi(z)$ satisfies the monotonicity property if for all $0 < \rho \leq \rho' < 1$, we have:*

$$\Delta \in \text{IQC}(\Pi(z), \rho) \implies \Delta \in \text{IQC}(\Pi(z), \rho').$$

All of the ρ -IQCs discussed herein satisfy the monotonicity property. If an IQC does not satisfy this property, then a grid search may be used instead of bisection.

Remark 4. *The results above may also be carried through in continuous time. In that case,*

an equation analogous to (3.5) must be satisfied for $G(s - a)$ for all $\omega \in [0, \infty)$, and can be verified by finding $P = P^\top$ and $\lambda \geq 0$ such that

$$\begin{bmatrix} \hat{A}^\top P + P\hat{A} - 2aP & P\hat{B} \\ \hat{B}^\top P & 0 \end{bmatrix} + \lambda \begin{bmatrix} \hat{C}^\top \\ \hat{D}^\top \end{bmatrix} M \begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix} \prec 0.$$

An alternative continuous-time formulation is detailed in *Hu and Seiler [2016]*.

3.5 Exponential Rates from Gain Bounds

In *Megretski and Rantzer [1997]*, IQC analysis is used to certify L_2 stability¹ of interconnected systems. As noted in *Megretski and Rantzer [1997]*: “for general classes of ordinary differential equations, exponential stability is equivalent to the input/output stability...”. However, while input/output stability often implies exponential stability, we will show that rates constructed from ℓ_2 bounds can be conservative. This fact justifies the use of a dedicated technique for certifying exponential rates rather than using an ℓ_2 analysis.

To make explicit this suboptimal rate, we will need two results. First, a generalization of Theorem 3.2.1 that allows us to optimize the ℓ_2 gains over any pair of signals. We will consider the scenario of Figure 3.5, which is slightly more general than the setup in Figure 3.1.

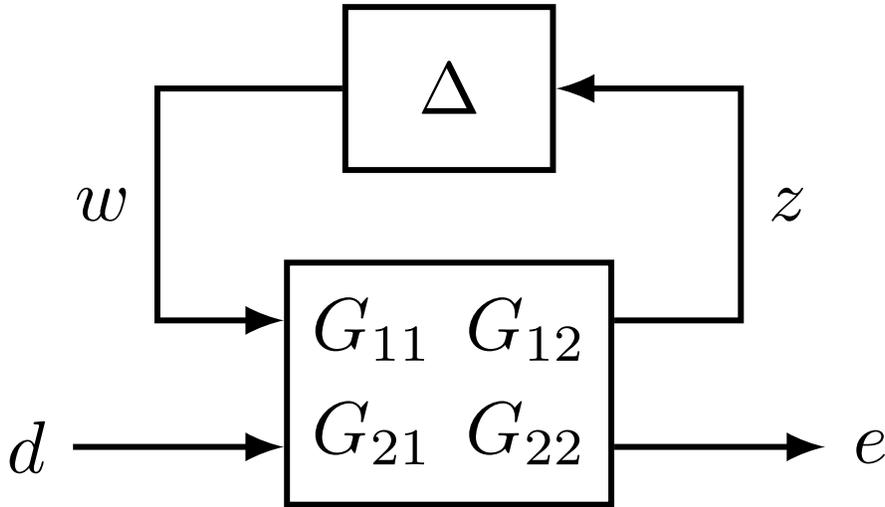


Figure 3.5: Augmented LTI system G in feedback with a nonlinearity Δ .

We would like to show that the input d and output e satisfy an IQC of the form

$$\int_{\mathbb{T}} \begin{bmatrix} \hat{d}(z) \\ \hat{e}(z) \end{bmatrix}^* \Pi_p(z) \begin{bmatrix} \hat{d}(z) \\ \hat{e}(z) \end{bmatrix} dz \geq 0. \quad (3.7)$$

¹ L_2 being the continuous-time analog of ℓ_2 .

The following result appears (for example) in [Apkarian and Noll \[2006\]](#), and a complete proof is given in [Summers \[2012\]](#).

Theorem 3.5.1. *Let $G(z) \in \mathcal{RH}_\infty^{m \times n}$ and let Δ be a bounded causal operator. Suppose G is partitioned according to the dimensions of the input and output channels in [Figure 3.5](#). Suppose the interconnection of G_{11} and Δ is well-posed and stable and $\Delta \in \text{IQC}(\Pi(z))$. If there exists $\epsilon > 0$ such that*

$$[\star]^* \begin{bmatrix} \Pi(z) & 0 \\ 0 & -\Pi_p(z) \end{bmatrix} \begin{bmatrix} G_{11}(z) & G_{12}(z) \\ I & 0 \\ 0 & I \\ G_{21}(z) & G_{22}(z) \end{bmatrix} \preceq -\epsilon I \quad \forall z \in \mathbb{T} \quad (3.8)$$

then for all $d \in \ell_2$ and $e \in \ell_2$, [Equation \(3.7\)](#) is satisfied.

Remark 5 (see [Summers \[2012\]](#)). *In [Theorem 3.5.1](#), if $\Pi_{p,22} \preceq 0$ and (G_{11}, Δ) satisfies assumptions (i) and (ii) of [Theorem 3.2.1](#), then stability of the (G_{11}, Δ) interconnection is immediate as the $(1, 1)$ block of the FDI provides the remaining requirement for stability in [Theorem 3.5.1](#).*

Next, we will need a way to convert an ℓ_2 gain into an exponential rate bound. The sequel is similar to [\[Megretski and Rantzer, 1997, Prop. 1\]](#), but presented here with an explicit rate construction and adapted for discrete-time systems.

Lemma 3.5.1. *Define the recursion with $x_0 = 0$ by:*

$$x_{k+1} = \phi(x_k) + g_k \quad k = 0, 1, 2, \dots \quad (3.9)$$

where $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ satisfies $\phi(0) = 0$. Suppose that there exists a constant $c > 0$ such that whenever $g \in \ell_2$, and (g, x) is a valid trajectory of [\(3.9\)](#), then

$$\|x\|_{\ell_2}^2 \leq c \|g\|_{\ell_2}^2. \quad (3.10)$$

Then, we also have the bound

$$\|x_{k+1}\|_2^2 \leq \sum_{i=0}^k c \left(1 - \frac{1}{c}\right)^{k-i} \|g_i\|_2^2.$$

Proof. We write $(x, g) \in \mathcal{S}$ to denote a valid trajectory of [\(3.9\)](#). Define the function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ as follows:

$$V(\xi) := \sup_{\substack{x_1 = \xi \\ g \in \ell_2 \\ (g, x) \in \mathcal{S}}} (\|x\|_{\ell_2}^2 - c \|g\|_{\ell_2}^2 + c \|\xi\|_2^2).$$

The first step is to bound $V(\xi)$. Note that $x_0 = 0$ and $\phi(0) = 0$, implies $\xi = x_1 = g_0$. A simple lower bound is found by specializing to $g_1 = g_2 = \dots = 0$, and an upper bound is found by using (3.10). Combining, we have

$$\|\xi\|_2^2 \leq V(\xi) \leq c \|\xi\|_2^2. \quad (3.11)$$

Now, fix $(\bar{g}, \bar{x}) \in \mathcal{S}$ to be any feasible trajectory of (3.9). We may lower-bound $V(\bar{x}_1)$ by setting $g_1 = \bar{g}_1$ and shifting the entire x and g vectors forward one timestep:

$$\begin{aligned} V(\bar{x}_1) &\geq \sup_{\substack{x_1 = \bar{x}_1 \\ g \in \ell_2, g_1 = \bar{g}_1 \\ (g, x) \in \mathcal{S}}} (\|x\|_{\ell_2}^2 - c \|g\|_{\ell_2}^2 + c \|\bar{x}_1\|_2^2) \\ &= V(\bar{x}_2) + \|\bar{x}_1\|_2^2 - c \|\bar{g}_1\|_2^2 \\ &\geq V(\bar{x}_2) + \frac{1}{c} V(\bar{x}_1) - c \|\bar{g}_1\|_2^2, \end{aligned}$$

where the final step follows from (3.11). Rearranging, we obtain

$$V(\bar{x}_2) \leq \left(1 - \frac{1}{c}\right) V(\bar{x}_1) + c \|\bar{g}_1\|_2^2.$$

We may lower-bound $V(\bar{x}_3)$ by setting $g_1 = \bar{g}_2$ and using a similar argument. Continuing in this fashion, we see that

$$V(\bar{x}_{k+1}) \leq \left(1 - \frac{1}{c}\right) V(\bar{x}_k) + c \|\bar{g}_k\|^2 \quad \text{for } k = 0, 1, 2, \dots$$

It then follows that for all k , we have

$$V(\bar{x}_{k+1}) \leq \left(1 - \frac{1}{c}\right)^k V(\bar{x}_1) + \sum_{i=1}^k c \left(1 - \frac{1}{c}\right)^{k-i} \|\bar{g}_i\|^2.$$

Applying the bound (3.11) one final time, we conclude that

$$\begin{aligned} \|\bar{x}_{k+1}\|^2 &\leq V(\bar{x}_{k+1}) \\ &\leq \left(1 - \frac{1}{c}\right)^k V(\bar{x}_1) + \sum_{i=1}^k c \left(1 - \frac{1}{c}\right)^{k-i} \|\bar{g}_i\|^2 \\ &\leq c \left(1 - \frac{1}{c}\right)^k \|\bar{x}_1\|^2 + \sum_{i=1}^k c \left(1 - \frac{1}{c}\right)^{k-i} \|\bar{g}_i\|^2 \\ &= \sum_{i=0}^k c \left(1 - \frac{1}{c}\right)^{k-i} \|\bar{g}_i\|^2, \end{aligned}$$

where we used $\bar{g}_0 = \bar{x}_1$ in the final step. ■

By combining Theorem 3.5.1 and Lemma 3.5.1, we can find exponential rate bounds for LTI systems in feedback with nonlinearities that satisfy IQCs. First, use the setup of Figure 3.5 with $d = g$ and $e = x$. Next, view the ℓ_2 bound in (3.10) is an IQC as in (3.7), with

$$\Pi_p = \begin{bmatrix} c & 0 \\ 0 & -1 \end{bmatrix}.$$

Then, transform Figure 3.1 into augmented form by setting

$$\begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix} = \left[\begin{array}{c|cc} A & B & B \\ \hline C & D & D \\ I & 0 & 0 \end{array} \right].$$

Finally, the appropriate initial condition can be set by using $g = d = [x_0^\top \ 0 \ 0 \ \dots]^\top$. Applying Lemma 3.5.1 leads to a bound of the form $\|x_{k+1}\|_2^2 \leq c \left(1 - \frac{1}{c}\right)^k \|x_0\|_2^2$; equivalently, an exponential rate of $\rho = \sqrt{1 - \frac{1}{c}}$.

The FDI of Theorem 3.5.1 can be transformed into an LMI in a manner similar to that described in Section 3.4. This LMI is linear in P and c , so it can be efficiently solved to find the minimal c and in turn the smallest exponential rate ρ . However, we will show in Section 3.7 that this “classical” rate ρ can be conservative.

3.6 IQC Library

In this section, we show classes of nonlinearities describable by ρ -IQCs and therefore applicable to Theorem 3.3.1 to prove robust exponential stability of an interconnected system. In the case where $\rho = 1$, these ρ -IQCs reduce to standard IQCs as given in Megretski and Rantzer [1997]. This class of IQCs will be constructed for single-input single-output systems, but they may be adapted for square multi-input multi-output systems where the nonlinearity is of the form $\text{diag}(\{\Delta_i\})$ for a scalar Δ .

3.6.1 Noisy Multiplication

As noted for continuous time in Hu and Seiler [2016], nonlinearities of the form $\Delta(y_k) \equiv \delta_k y_k$ for some unknown and/or time-varying δ_k may satisfy ρ -IQCs. As Δ and ρ_\pm commute, in the parlance of Proposition 3.3.2 we have that $\Delta = \Delta'$, so $\Delta \in \text{IQC}(\Pi, 1)$ implies $\Delta \in \text{IQC}(\Pi, \rho)$. See Megretski and Rantzer [1997] for examples of noisy multiplication.

3.6.2 Uncertain Time Delay

The following is a discrete-time analog of the ρ -IQC first developed in [Hu and Seiler \[2016\]](#). Let Δ be the operator defined by

$$\Delta(y_k) = \begin{cases} 0, & k < \tau \\ y_{k-\tau}, & k \geq \tau \end{cases},$$

for some unknown τ in $[0, \tau_0]$, where τ_0 is known. Now, observe that

$$\begin{aligned} \Delta'(y_k) &= \rho^{-k} \Delta(\rho^{-k} y_k) = \rho^{-k} \cdot \begin{cases} 0, & k < \tau \\ \rho^{-(k-\tau)} y_{k-\tau}, & k \geq \tau \end{cases} \\ &= \rho^{-\tau} \Delta(y_k). \end{aligned}$$

Since $\Delta'(y_k) = \rho^{-\tau} \Delta(y_k)$, we may transform the system into one with a block diagonal nonlinearity $\text{diag}\{\Delta, \rho^{-\tau}\}$. We can then use existing IQCs for noisy multiplication and time delays, always using $\Pi(\rho z)$ instead of $\Pi(z)$.

Alternatively, with any bounded Hermitian function $X(\rho z) = X(\rho z)^* \succeq 0$, we see that

$$\begin{aligned} &\begin{bmatrix} \hat{y}(\rho z) \\ \hat{u}(\rho z) \end{bmatrix}^* \begin{bmatrix} \rho^{-2\tau_0} X(\rho z) & 0 \\ 0 & -X(\rho z) \end{bmatrix} \begin{bmatrix} \hat{y}(\rho z) \\ \hat{u}(\rho z) \end{bmatrix} \\ &= \begin{bmatrix} \hat{y}(\rho z) \\ \rho^{-\tau} \hat{y}(\rho z) \end{bmatrix}^* \begin{bmatrix} \rho^{-2\tau_0} X(\rho z) & 0 \\ 0 & -X(\rho z) \end{bmatrix} \begin{bmatrix} \hat{y}(\rho z) \\ \rho^{-\tau} \hat{y}(\rho z) \end{bmatrix} \\ &= (\rho^{-2\tau_0} - \rho^{-2\tau}) \hat{y}(\rho z)^* X(\rho z) \hat{y}(\rho z) \geq 0. \end{aligned}$$

Thus, $\Delta \in \text{IQC}(\text{diag}\{\rho^{-2\tau_0} X(z), -X(z)\}, \rho)$.

3.6.3 Pointwise IQCs

A nonlinearity Δ satisfies a pointwise IQC with a factorization (Ψ, M) if $z_k^\top M z_k \geq 0$ for each k . In other words, the IQC holds pointwise in time. In this case, Δ also satisfies the associated ρ -IQC for all $\rho < 1$. Examples of pointwise IQCs include the γ *norm-bounded IQC*

$$\Pi = \begin{bmatrix} \gamma^2 & 0 \\ 0 & -1 \end{bmatrix},$$

and the $[\alpha, \beta]$ *sector-bounded IQC*, given by

$$\Pi = \begin{bmatrix} -2\alpha\beta & \alpha + \beta \\ \alpha + \beta & -2 \end{bmatrix}.$$

The latter corresponds to a nonlinearity Δ that satisfies

$$(\Delta(x) - \beta x)^\top (\Delta(x) - \alpha x) \leq 0 \quad \forall x.$$

Note that the norm-bounded IQC is a special case of the sector IQC with the sector $[-\gamma, \gamma]$. These IQCs hold even if Δ is time-varying, if Δ satisfies the IQC at each k .

3.6.4 Zames–Falb IQCs

A nonlinearity Δ is *slope-restricted on* $[\alpha, \beta]$ where $0 \leq \alpha \leq \beta \leq \infty$ if the following relation holds for all x, y .

$$(\Delta(x) - \Delta(y) - \alpha(x - y))^\top (\Delta(x) - \Delta(y) - \beta(x - y)) \leq 0.$$

This relation states that the chord joining input-output pairs of Δ has a slope that is bounded between α and β . This class of functions satisfies the Zames–Falb family of IQCs (see [Zames and Falb \[1968\]](#) and [Heath and Wills \[2005\]](#)). We give the definition below.²

Proposition 3.6.1. *A nonlinearity Δ that is static and slope-restricted on $[\alpha, \beta]$ satisfies the Zames–Falb IQC*

$$\Pi = \begin{bmatrix} -\alpha\beta(2-\hat{h}-\hat{h}^*) & \alpha(1-\hat{h})+\beta(1-\hat{h}^*) \\ \alpha(1-\hat{h}^*)+\beta(1-\hat{h}) & -(2-\hat{h}-\hat{h}^*) \end{bmatrix} \quad (3.12)$$

where $\hat{h}(z)$ is any proper transfer function with impulse response $h := (h_0, h_1, \dots)$ that satisfies $\|h\|_1 \leq 1$ and $h_k \geq 0$ for all k . Moreover, it admits the factorization

$$\Psi = \begin{bmatrix} \beta(1-\hat{h}) & -(1-\hat{h}) \\ -\alpha & 1 \end{bmatrix} \quad \text{and} \quad M = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

If Δ is odd ($\Delta(-x) = -\Delta(x)$), then we may remove the constraint that $h_k \geq 0$ for all k .

In general, for a fixed ρ , only a subset of Zames–Falb IQCs will be ρ -IQCs. We now give a characterization of this subset.

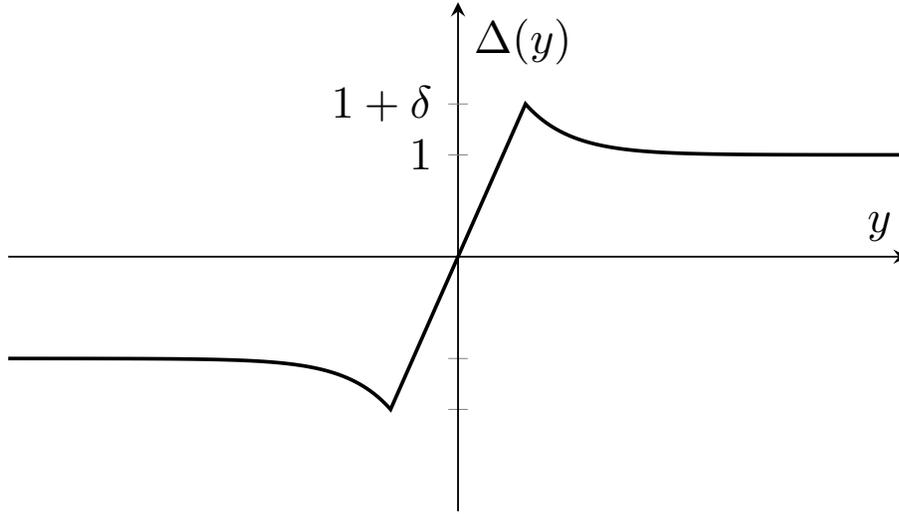
Theorem 3.6.1 (Zames–Falb ρ -IQC). *Suppose Δ is static and slope-restricted on $[\alpha, \beta]$. Then $\Delta \in \text{IQC}(\Pi(z), \rho)$ where Π is the Zames–Falb IQC (3.12) and \hat{h} also satisfies the additional constraint*

$$\sum_{k=0}^{\infty} \rho^{-2k} |h_k| \leq 1.$$

Proof. The proof involves rewriting the IQC as a discrete-time sum which can be split into parts that can separately be shown to be nonnegative. See [Appendix A.2](#) for the full proof of [Theorem 3.6.1](#) and related extensions. ■

The concept of a ρ -IQC can also be extended to handle noncausal Zames–Falb multipliers as in [Freeman \[2018\]](#). Sector-bounded and/or slope-restricted functions show up in various specialized contexts. We will derive ρ -IQCs for two such cases: stiction nonlinearities and quasi-monotone/quasi-odd nonlinearities.

²The $\beta = \infty$ case for this and similar IQCs considers only the β terms, i.e. $\Pi_{[\alpha, \infty]} = \lim_{\beta \rightarrow \infty} \beta^{-1} \Pi$.

Figure 3.6: Example stiction nonlinearity (taken from [Rantzer \[2001\]](#).)

3.6.4.1 Stiction Nonlinearities

Stiction nonlinearities (shown in Figure 3.6) satisfy Zames–Falb ρ -IQCs with additional constraints on the coefficients h_k .

Corollary 3.6.1 (Stiction ρ -IQC). *Suppose Δ is a stiction nonlinearity with slope $1/\epsilon$ and overshoot δ as defined in [Rantzer \[2001\]](#). Then $\Delta \in \text{IQC}(\Pi(z), \rho)$ where Π is the $[0, 1/\epsilon]$ Zames–Falb IQC (3.12) and H satisfies the additional constraint*

$$\sum_{k=0}^{\infty} \rho^{-2k} |h_k| \leq \frac{1 - \delta}{1 + \delta}.$$

3.6.4.2 Quasi-monotone and Quasi-odd Nonlinearities

Following the definition in [Heath et al. \[2015\]](#) (shown in Figure 3.7), quasi-monotone and quasi-odd nonlinearities also satisfy Zames–Falb ρ -IQCs under additional constraints on the h_k .

Corollary 3.6.2 (Quasi-monotone/odd ρ -IQC). *Suppose Δ is static and is quasi-monotone or quasi-odd as defined in [Heath et al. \[2015\]](#). Then $\Delta \in \text{IQC}(\Pi(z), \rho)$ where Π is the Zames–Falb IQC (3.12) and H satisfies the additional constraint*

$$\sum_{k=0}^{\infty} \gamma_k^{-1} \rho^{-2k} |h_k| \leq 1, \quad \text{where} \quad \gamma_k^{-1} := \begin{cases} R_m, & h_k \geq 0 \\ R_o, & h_k < 0 \end{cases}.$$

Given a fixed ρ , searching over (finite) h_k when solving the feasibility LMI using this IQC is still a convex problem. To see this, observe that we can equivalently write this constraint

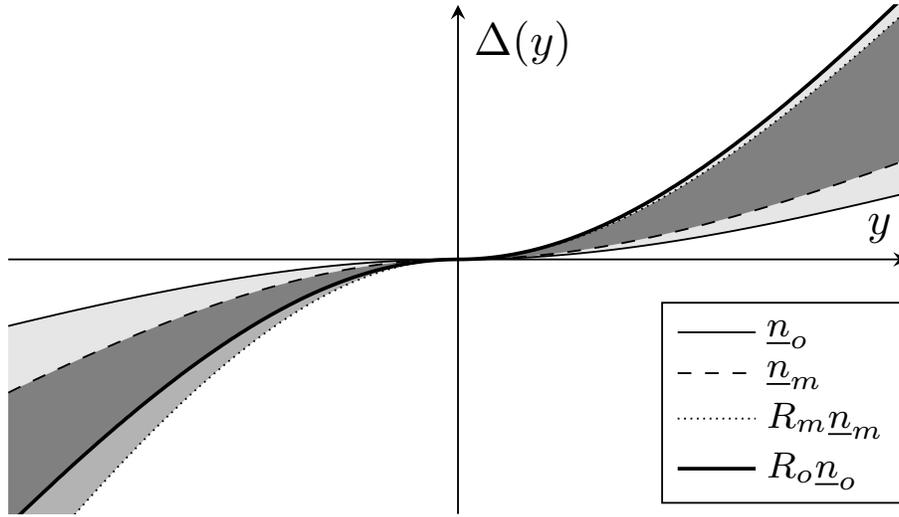


Figure 3.7: Monotone and odd bounds for unknown nonlinearities (modified from [Heath et al. \[2015\]](#)). The nonlinearity must lie within envelopes generated by multiplicative perturbations of a known monotone linearity \underline{n}_m (perturbation between 1 and $R_m \geq 1$) and a known monotone odd nonlinearity \underline{n}_o (perturbation between 1 and $R_o \geq 1$). In this example, the nonlinearities of interest lie in the darkest region, the intersection of both envelopes.

on the h_k (assuming $R_m \geq R_o$, the other case is similar) as

$$R_o \sum_{k=0}^K \rho^{-2k} |h_k| + (R_m - R_o) \sum_{k=0}^K \rho^{-2k} \cdot \max(h_k, 0) \leq 1.$$

However, the proof of Corollary 3.6.2 will show that these general Zames–Falb ρ -IQCs can be written as a nonnegative linear combination of “off-by- j ” ρ -IQCs. Thus, when solving (3.6) it is sufficient to search over all nonnegative linear combinations of simpler ρ -IQC atoms, rather than formulating the constraint on the h_k explicitly. Whether this is more efficient depends on the specific problem dimensions.

It is a useful exercise to derive the correct chain of implications for this constraint (and others), which is as follows:

- Compared to the odd Zames–Falb IQC, a quasi-odd IQC as defined in Corollary 3.6.2 gives *less information* about the nonlinearity ϕ , i.e. we must provide a certificate of stability for every nonlinearity in a *larger class*.
- Since $R_m, R_o \geq 1$, the weights satisfy $\gamma_k^{-1} \geq 1$, so there is *less freedom* in choosing the h_k .
- This restriction in choosing h_k leads to a *smaller* feasible set for the LMI. Thus, the upper bound we find for the convergence rate will be *larger*.

3.6.5 Repeated Sector Nonlinearities

We say a real symmetric matrix Γ is (ρ, H) -*diagonally dominant* if, for a symmetric matrix of nonnegative proper transfer functions \hat{H} with impulse responses $H_{ij,k}$, we have that $\Gamma_{ii} \geq 0$, $\Gamma_{ij} \leq 0$ (for $i \neq j$), $H_{ij,k} \geq 0$, $\sum_{k=0}^{\infty} \rho^{-2k} |H_{ij,k}| \leq 1 \forall (i, j)$ and

$$\Gamma_{ii} \geq \sum_{j=1, j \neq i}^n |\Gamma_{ij}| + \sum_{j=1}^n \sum_{k=0}^{\infty} \rho^{-2k} |H_{ij,k}| \quad \forall i.$$

We call Γ just *diagonally dominant*³ if the above holds with $H = 0$ and $\rho = 1$. Now, let Δ be a repeated monotone scalar nonlinearity in some sector; that is, $\Delta(y) = \text{diag}\{\phi(y_i)\}$.

Proposition 3.6.2. Δ satisfies the pointwise ρ -IQC

$$\Pi = \begin{bmatrix} 0 & \Gamma \\ \Gamma & 0 \end{bmatrix}$$

for any symmetric diagonally dominant matrix Γ .

Proof. This proposition is analogous to Theorem 1 in the Appendix of D'Amato et al. [2001] with $H = 0$. ■

Theorem 3.6.2. Assume Γ is (ρ, H) -diagonally dominant. Then, if a static nonlinearity ϕ is $[\alpha, \beta]$ slope-restricted, the repeated nonlinearity $\Delta(y) = \text{diag}\{\phi(y_i)\}$ satisfies the ρ -IQC

$$\Pi = \begin{bmatrix} -\alpha\beta(2\Gamma - \hat{H} - \hat{H}^*) & \alpha(\Gamma - \hat{H}) + \beta(\Gamma - \hat{H}^*) \\ \alpha(\Gamma - \hat{H}^*) + \beta(\Gamma - \hat{H}) & -2\Gamma + \hat{H} + \hat{H}^* \end{bmatrix}. \quad (3.13)$$

Moreover, it admits the factorization

$$\Psi = \begin{bmatrix} \beta(\Gamma - \hat{H}) & -(\Gamma - \hat{H}) \\ -\alpha I & I \end{bmatrix}, \quad M = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}.$$

Proof. The proof is similar in spirit to that of Theorem 3.6.1 but more involved; see Appendix A.3. ■

See Appendix A.4 for a note on how to search over general nonnegative combinations of ρ -IQCs of the form (3.13), which is not immediately apparent.

³Note that the conventional definition of “diagonally dominant” does not restrict the diagonal elements to be nonnegative.

3.7 Examples

3.7.1 Using Multiple IQCs

Using multiple IQCs can lead to more refined ℓ_2 gain bounds. Likewise, using multiple ρ -IQCs can lead to refined exponential rates. In this section, we present numerical examples using both pointwise and dynamic ρ -IQCs. Now, consider a stable discrete-time LTI system $G(z)$ in feedback with the sigmoidal nonlinearity $\Delta(x) = b \arctan(x)$. This interconnection is shown in Figure 3.8.

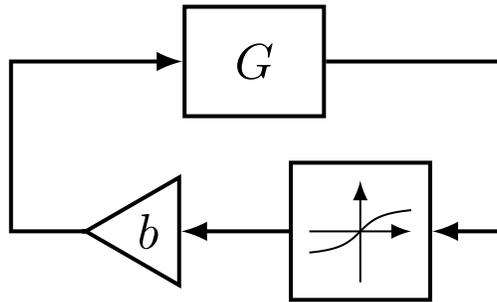


Figure 3.8: LTI system G in feedback with the static sigmoidal nonlinearity $\Delta(x) = b \arctan(x)$.

Since this nonlinearity is static, in the $[0, b]$ sector, and $[0, b]$ slope-restricted, it satisfies the following ρ -IQCs:

$$\Pi_n(z) := \begin{bmatrix} b^2 & 0 \\ 0 & -1 \end{bmatrix} \quad (\text{norm-bounded}) \quad (3.14)$$

$$\Pi_0(z) := \begin{bmatrix} 0 & b \\ b & -2 \end{bmatrix} \quad (\text{sector bounded}) \quad (3.15)$$

$$\Pi_k(z) := \begin{bmatrix} 0 & b(1 - \rho^{2k} \bar{z}^{-k}) \\ b(1 - \rho^{2k} z^{-k}) & -2 + \rho^{2k}(z^{-k} + \bar{z}^{-k}) \end{bmatrix} \quad (\text{off-by-}k \text{ Zames-Falb}) \quad (3.16)$$

where we may choose any $k \geq 1$.

3.7.1.1 A simple bound

For our first case study, we analyzed the interconnection of Figure 3.8 with the LTI system⁴

$$G_1(z) = -\frac{(z+1)(10z+9)}{(2z-1)(5z-1)(10z-1)}.$$

⁴This example was inspired by the continuous-time example given in [Weiland and Scherer \[2015\]](#), which showed that adding more IQCs yields better ℓ_2 gain bounds.

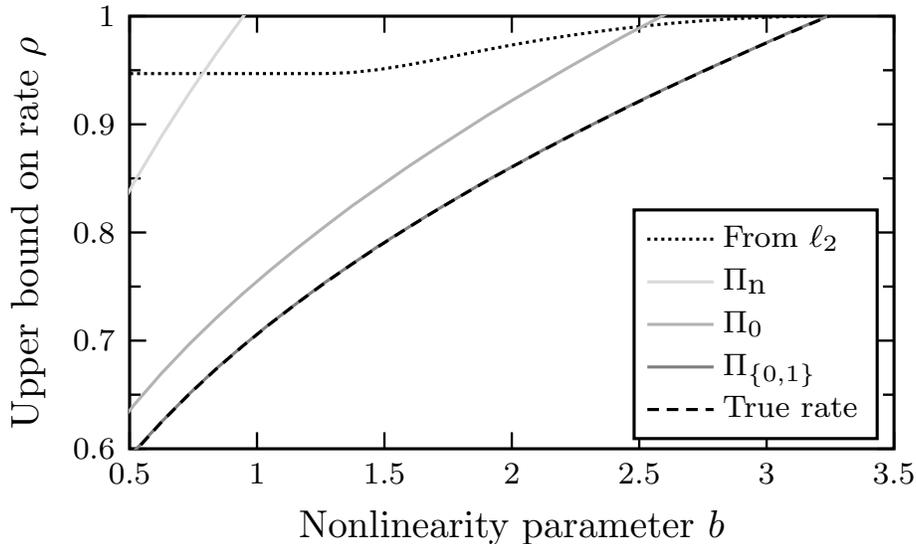


Figure 3.9: Upper bounds on the exponential convergence rate ρ for the system $G_1(z)$ given in (3.7.1.1) in feedback as in Figure 3.8. A tight bound is achieved using two ρ -IQCs. The bound derived from the ℓ_2 gain is very conservative.

We solved the feasibility LMI (3.6) using MATLAB together with CVX [Grant and Boyd, 2008, 2014] to find the fastest guaranteed rate of convergence, and we searched over positive linear combinations of subsets of the IQCs (3.14)–(3.16). Figure 3.9 shows the rate bounds achieved as a function of which IQCs were used. Figure 3.10 shows sample state trajectories for the case $b = 1$.

The true exponential rate can be found by linearizing the system about its equilibrium point; namely, $\Delta(x) \approx bx$. Formally, this is an application of Lyapunov’s indirect method [Khalil, 2002, Thm. 4.13]. The result is that the decay rate should correspond to the maximal pole magnitude of the closed-loop map $G(z)/(1 - bG(z))$. We display the true exponential rate as the dashed black curve in Figure 3.9 and Figure 3.10.

For this example, the ρ -IQC approach yields a tight upper bound to the true exponential rate when we use a combination of the sector and off-by-1 IQCs. We also computed the exponential rate derived from ℓ_2 gain as described in Section 3.5 (dotted line). The ℓ_2 bound is very conservative despite being computed using all available IQCs.

3.7.1.2 A more complex bound

The ρ -IQC approach does not always achieve tight bounds as in the previous example. Consider the same interconnection of Figure 3.8 but this time using

$$G_2(z) = \frac{2z - 1}{10(2z^2 - z + 1)}$$

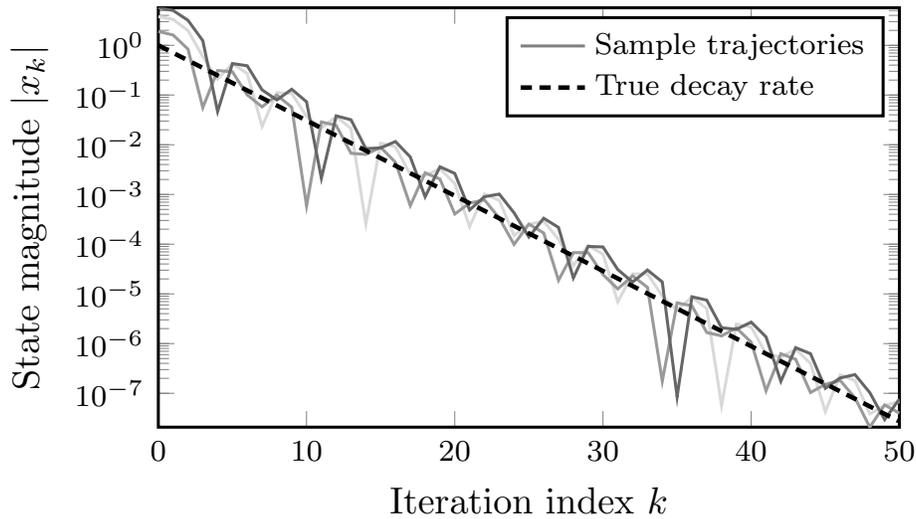


Figure 3.10: State decay over time of the system $G_1(z)$ in feedback as in Figure 3.8 with $b = 1$ for various initial conditions $x_0 \in [-15, 15]$. The dashed black line is ρ^k , where $\rho = .7058$ is the true rate at $b = 1$ in Figure 3.9.

The rate bounds for various ρ -IQCs are shown in Figure 3.11. This time, we again observe that using more IQCs achieves better rate bounds, but the bound is not tight even after using six IQCs. However, if we add the Zames–Falb IQCs corresponding to odd monotone nonlinearities, the rate improves to within a small tolerance of the true rate.

As in the previous example, the best achievable rate derived from an ℓ_2 gain bound as detailed in Section 3.5 is still very conservative when compared to the rates obtained by using the ρ -IQC approach.

3.7.1.3 A quasi-odd nonlinearity

Consider the asymmetric nonlinearity in Figure 3.12, shown with the associated monotone and odd bounds as defined in Heath et al. [2015]. In this example, we have $R_m = 1$ and $R_o = 2$.

Thus, we may invoke Corollary 3.6.2 and use the associated ρ -IQC. Using this system in feedback with the $G(z)$ from the second example, we see in Figure 3.13 that the quasi-odd Zames–Falb IQCs yield better performance than the monotone Zames–Falb IQCs of the same order (which requires all filter coefficients h_k to be positive).

3.7.1.4 Repeated nonlinearities

To illustrate the need for repeated nonlinearity IQCs, first instantiate some stable SISO system G with realization (A, B, C, D) . Now, consider the “extended” 2-input 2-output

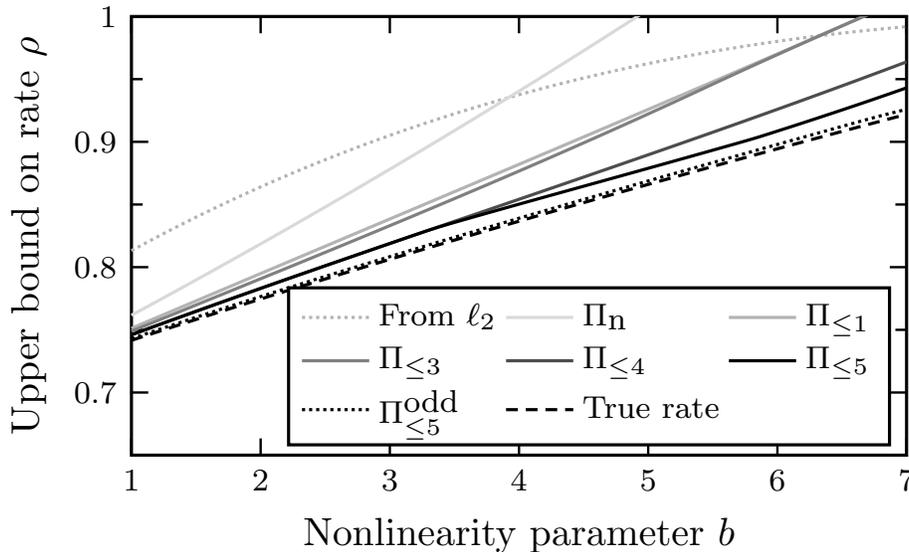


Figure 3.11: Upper bounds on the exponential convergence rate ρ for the system $G_2(z)$ given in (3.7.1.2) in feedback as in Figure 3.8. As we include more ρ -IQCs, we can certify tighter bounds. Once again, the ℓ_2 -derived bound is more conservative.

system

$$G_{\text{ext}} = \left[\begin{array}{c|cc} A & B & -B \\ \hline C & D & 0 \\ C & 0 & D \end{array} \right]$$

and consider G_{ext} in positive feedback with the block-diagonal nonlinearity $\text{diag}\{\Delta_1, \Delta_2\}$. If we constrain $\Delta_1 = \Delta_2$, then the nonlinearities cancel each other out and the system is in open loop. The convergence rate of the state is therefore determined by the largest magnitude eigenvalue of A . However, if our IQC does not capture that the nonlinearity is repeated and instead only assumes each individual nonlinearity is (say) $[0, b]$ slope-restricted, then G must essentially be robust to b -norm bounded nonlinearities in the feedback loop. This will result in a worse rate certificate or even none at all (if G is made unstable by positive feedback).

Indeed, constructing G_{ext} using our previous “tight bound” example with $b = 0.3$ leads to a rate certificate of ≈ 0.825 using only the odd monotone IQC; replacing it with the repeated odd monotone nonlinearity IQC gives a certificate matching the true convergence rate, 0.5.

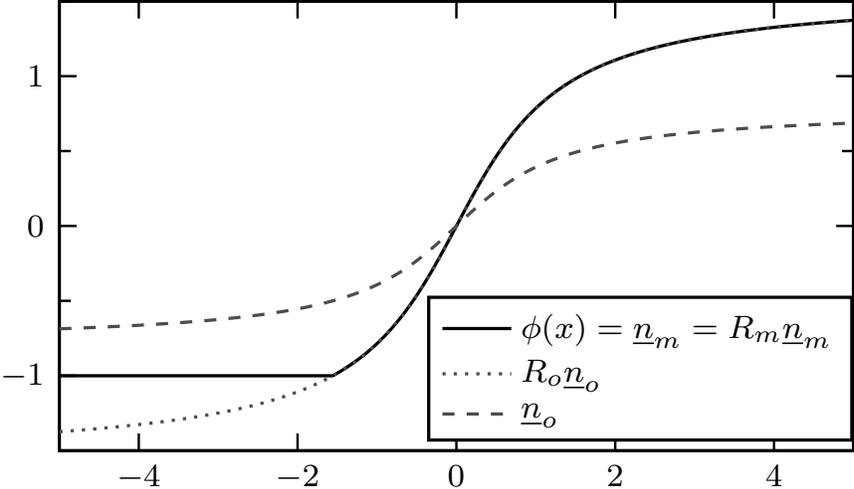


Figure 3.12: Plot of the monotone and quasi-odd asymmetric nonlinearity $\phi(x) = \max\{\arctan(x), -1\}$ with its associated bounds.

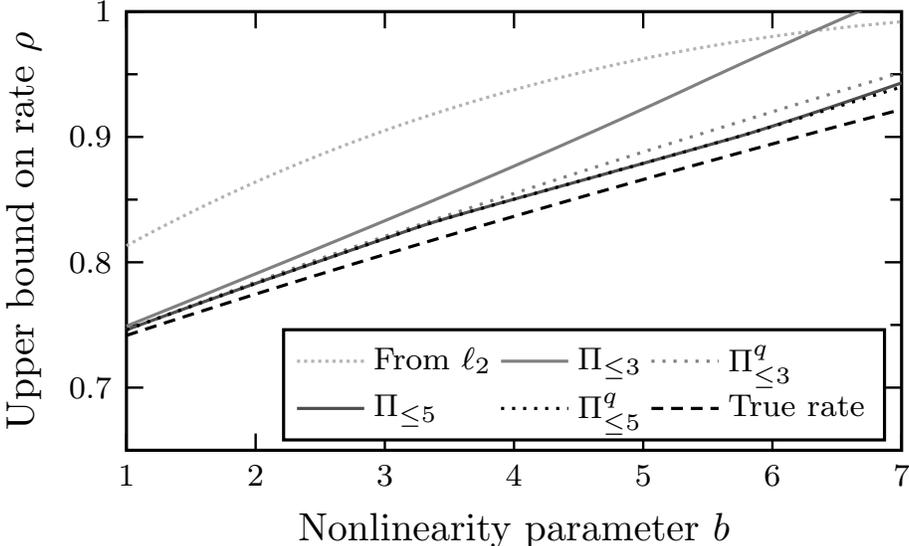


Figure 3.13: Comparison of monotone Zames–Falb and quasi-odd (denoted with superscript q) Zames–Falb IQC rate certificates.

3.8 An Aside: Application to Optimization Systems

3.8.1 Previous Work

As mentioned in the previous sections, LMIs have been a useful tool to verify properties of interconnected dynamical systems [Boyd et al., 1994; Rantzer, 1996]. Recent work by Lessard et al. [2016] looked at an inverse approach, i.e. using tools from control and system theory to verify properties of optimization algorithms. Using techniques similar to ones developed there and in this chapter, we look at generalizing Lessard et al. [2016] to stochastic optimization algorithms.

3.8.2 First-order Optimization Systems

In Lessard et al. [2016], the authors made a fundamental connection between optimization algorithms (specifically, first-order ones) and discrete-time linear dynamical systems. They noted that many commonly-used algorithms be described by a linear discrete-time system G composed with a nonlinear feedback function ϕ —with the nonlinearity being the gradient of the objective function. The feedback in the system is then defined by $u = \nabla f(y)$. For example, the linear system G corresponding to the gradient descent method with step size α can be described by standard state-space equations:

$$\left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right] := \left[\begin{array}{c|c} I & -\alpha I \\ \hline I & 0 \end{array} \right].$$

As seen before, when such a system satisfies an IQC with a factorization (Ψ, M) , analyzing only the linear portions of the constrained system can give stability information about the original system. This decomposition is seen in Figure 3.14 (taken from Lessard et al. [2016]).

We now seek to describe stochastic optimization algorithms, such as stochastic gradient descent, using this framework. Some related results can be found in Hu [2016].

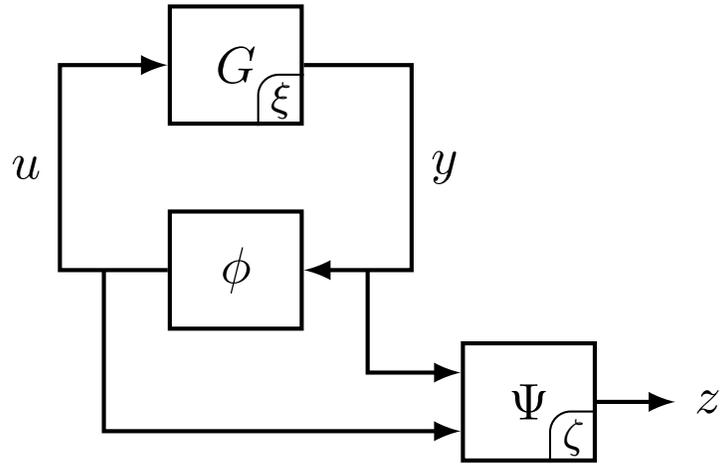
3.8.3 Gradient Descent with Additive Input Noise

We present a preliminary result and contrast it with the main result from Lessard et al. [2016]. Consider an optimization algorithm described by linear system under nonlinear (gradient) feedback, with additive input noise w_k . For example, take the stochastic gradient descent update

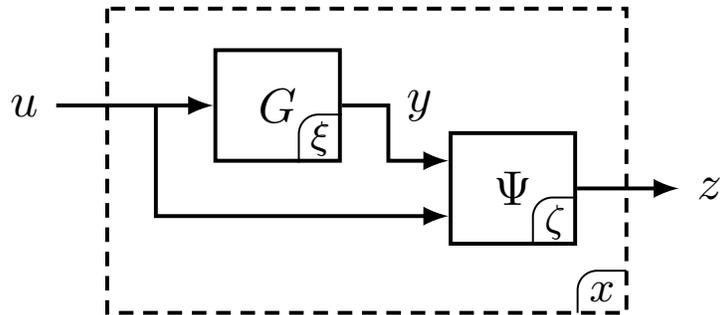
$$x_{k+1} = x_k - \alpha g_k,$$

where g_k is a noisy (but unbiased) estimate of $\nabla f(x_k)$, such as $g_k = \nabla f(x_k) + w_k$ where w_k is independent of the gradient and zero-mean. In general, such an optimization algorithm can be viewed as the dynamical system

$$G : \begin{aligned} \xi_{k+1} &= A\xi_k + B(u_k + w_k) \\ y_k &= C\xi_k \end{aligned} \tag{3.17}$$



(a) An auxiliary system Ψ produces z , a filtered version of the signals y and u .



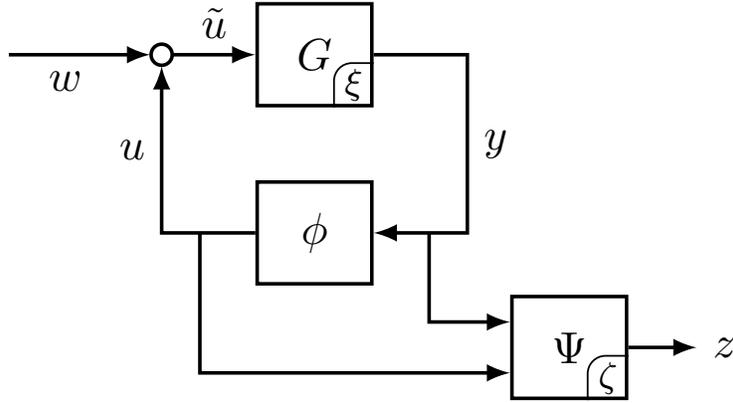
(b) The nonlinearity ϕ is replaced by a constraint on z , so we may remove ϕ entirely.

Figure 3.14: Feedback interconnection between a system G and a nonlinearity ϕ . An IQC is a constraint on (y, u) satisfied by ϕ . We only analyze the constrained system and so we may remove the ϕ block entirely.

where u_k is the gradient of some function f . Now, assume that the gradient satisfies an IQC with a factorization defined by the linear system Ψ , with dynamics

$$\begin{aligned} \zeta_{k+1} &= A_{\Psi} \zeta_k + \begin{bmatrix} B_{\Psi}^y & B_{\Psi}^u \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix} \\ z_k &= C_{\Psi} \zeta_k + \begin{bmatrix} D_{\Psi}^y & D_{\Psi}^u \end{bmatrix} \begin{bmatrix} y_k \\ u_k \end{bmatrix} \end{aligned} \quad (3.18)$$

The block diagram shown in Figure 3.15 shows the overall system with the new noise input w . Consider the dynamics of G and Ψ from (3.17) and (3.18), respectively. As in Corollary 3.4.1,

Figure 3.15: The perturbed system, including additive input noise w .

upon eliminating y the recursions may be combined to obtain another linear system

$$\begin{aligned} x_{k+1} &= \hat{A}x_k + \hat{B}u_k + \begin{bmatrix} B \\ 0 \end{bmatrix} w_k & \text{where } x_k &:= \begin{bmatrix} \xi_k \\ \zeta_k \end{bmatrix}, \\ z_k &= \hat{C}x_k + \hat{D}u_k \end{aligned} \quad (3.19)$$

for some matrices $\hat{A}, \hat{B}, \hat{C}, \hat{D}$.

Recall the definition of a ρ -hard IQC from Lessard et al. [2016, Def. 3], which is similar to the discrete-time ρ -IQC given in Remark 3 but is required to hold for all finite sums.

Definition 3.8.1. *The IQC defined by (Ψ, M) is ρ -hard if*

$$\sum_{t=0}^k \rho^{-2t} (z_t - z_\star)^\top M (z_t - z_\star) \geq 0 \quad \text{for } k = 0, 1, \dots$$

If we now know that the gradient of f satisfies a ρ -hard IQC, we can say something about the convergence of the system—the convergence of the optimization algorithm.

Theorem 3.8.1. *Consider the block interconnection of Figure 3.15. Assume the error signal w_k comes i.i.d. from a distribution with zero mean and covariance matrix Λ_w , and is independent of u_k and x_k . Suppose G is given by (3.17) and ϕ satisfies a ρ -hard IQC defined by (Ψ, M, z_\star) where Ψ is given by (3.18) and $0 \leq \rho \leq 1$. Assume x_\star is a fixed point of the dynamical system given by (3.19) when $w_k = 0$. Consider the LMI*

$$\begin{bmatrix} \hat{A}^\top P \hat{A} - \rho^2 P & \hat{A}^\top P \hat{B} \\ \hat{B}^\top P \hat{A} & \hat{B}^\top P \hat{B} \end{bmatrix} + \lambda \begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix}^\top M \begin{bmatrix} \hat{C} & \hat{D} \end{bmatrix} \preceq 0. \quad (3.20)$$

If (3.20) is feasible for some $P \succ 0$ and $\lambda \geq 0$, then for any x_0 , we have

$$\mathbb{E} \|x_k - x_\star\|^2 \leq \kappa(P) \rho^{2k} \|x_0 - x_\star\|^2 + \frac{\text{tr}(B^\top P_{11} B \Lambda_w)}{\lambda_{\min}(P)} \frac{1 - \rho^{2k}}{1 - \rho^2} \quad \forall k.$$

Proof. The complete proof (given in Appendix A.5) follows in the same vein as the main result in Lessard et al. [2016], which is a direct discrete-time proof in contrast to the frequency-domain results in Section 3.3. ■

We can compare Theorem 3.8.1 to the main result of Lessard et al. [2016, Thm 4], which states that, if the LMI is feasible, then

$$\|x_k - x_\star\|^2 \leq \kappa(P)\rho^{2k}\|x_0 - x_\star\|^2 \quad \text{for all } k.$$

3.9 Conclusion

IQC theory is the most general tool available for certifying robust stability of systems in feedback with unknown, uncertain, or otherwise difficult nonlinearities. As stable systems are often exponentially stable, it is reasonable to want finer control over not only stability, but also exponential decay rate. The generalization presented herein enables the certification of robust *exponential* stability with precise control over the decay rate. Moreover, the library of ρ -IQCs provided shows how this approach can be applied as broadly and efficiently as the classical IQC theory.

Chapter 4

\mathcal{H}_∞ Bounds for System Estimation

4.1 Introduction

Most control design relies on establishing a model of the system to be controlled. For simple physical systems, a model with reasonable fidelity can typically be constructed from knowledge of the physics at hand. However, for complex, uncertain systems, building models from first principles becomes quickly intractable and one usually resorts to fitting models from empirical input/output data. This approach naturally raises an important question: how well must we identify a system in order to control it?

In this chapter, we attempt to answer this question by striking a balance between system identification and robust control. We aim to identify coarse estimates of the true underlying model while coupling our estimation with precise probabilistic bounds on the inaccuracy of our estimates. With a coarse model in hand, we can use standard robust synthesis tools that take into account the derived bounds on the model uncertainty.

More precisely, given an unknown stable discrete-time plant G , we bound the error accrued by fitting a finite impulse response (FIR) approximation to G from noisy output measurements. These bounds balance the sample complexity of estimating an unknown FIR filter against the capability of such a filter to approximate the behavior of G . In particular, we show that notably short FIR filters provide a sufficient approximation to stable systems in order to ensure robust performance for a variety of control design tasks. In particular, we demonstrate considerable savings in experimental measurements as compared to other non-asymptotic schemes that aim to precisely identify G .

In the process of fitting a FIR filter, a natural question arises as to what inputs should be used to excite the unknown system. Of course, due to actuator limitations and other physical constraints, we are not free to choose any arbitrary input. Hence, we model the choice of inputs as an experiment design question, where the practitioner specifies a bounded input set and asks for the best m inputs to use to minimize FIR identification error. We propose a new optimal experiment design procedure for solving this problem, and relate it to the well studied A -optimal experiment design objective from the statistics literature [Pukelsheim,

1993]. This connection is used to study practical cases of input constraints. Specifically, we prove that when the inputs are ℓ_2 -power constrained, then impulse responses are the optimal choice of inputs. However, we show that is not the case when the inputs are ℓ_∞ -constrained. For ℓ_∞ constraints, we construct a deterministic set of inputs which is within a factor of 2 to the optimal solution. Combining these designs with our probabilistic bounds, we show that for estimating a length- r FIR filter G_r , as long as $m \geq 4r$, the residual \mathcal{H}_∞ error $\|G_r - \widehat{G}_r\|_{\mathcal{H}_\infty}$ on the estimate \widehat{G}_r satisfies $\widetilde{O}(1/\sqrt{m})^1$ with high probability. This is a substantial improvement over the $\widetilde{O}(\sqrt{r/m})$ scaling which we show occurs in the ℓ_2 -constrained case. We also prove an information-theoretic lower bound which shows that when the true system happens to be an FIR filter, our bounds are minimax optimal up to constant factors for the given estimation problem.

Experimentally, we show that \mathcal{H}_∞ loop-shaping controller design on the estimated FIR model, using probabilistic bounds, can be used to synthesize controllers with both stability and performance guarantees on the closed loop with the true plant. We also demonstrate that our probabilistic bounds can be estimated directly from data using Monte-Carlo techniques. We hope that our results encourage further investigation into a rigorous foundation for data-driven controller synthesis.

4.1.1 A Sample Complexity Bound for FIR Identification

We now state our main results for this section: upper and lower bounds on the sample complexity of FIR system identification. Let G be a stable, discrete-time SISO LTI system. Suppose we are given query access to G via independent, noisy measurements of the form

$$y_{u,T} := (g * u)_{k=0}^{T-1} + \xi, \quad \xi \sim \mathcal{N}(0, \sigma^2 I_T). \quad (4.1)$$

Above, g denotes the impulse response of G and T is the length of the output we observe (we add the u and t subscripts to reinforce the notions of a generating input and number of samples). We assume that we are allowed to choose any input u contained within a bounded set \mathcal{U} , which is specified beforehand. From these measurements, we can approximate G by a length- r FIR filter $\widehat{G}_r(z)$ as

$$\widehat{G}_r(z) := \sum_{k=0}^{r-1} \widehat{g}_k z^{-k}, \quad (4.2)$$

where $r \leq T$, and the coefficients \widehat{g}_k are estimated from ordinary least-squares (c.f. Section 4.3). We note that the extra degree of freedom in allowing $r \neq T$ reduces the variance of the higher lag terms, which is a standard trick used in the system-identification literature (see e.g. Wahlberg et al. [2010, Section 2]).

The main quantity of interest in this setting is the number of timesteps needed in order to ensure that the \mathcal{H}_∞ -norm of the error $G - \widehat{G}_r$ satisfies the bound $\|G - \widehat{G}_r\|_{\mathcal{H}_\infty} \leq \varepsilon$

¹The notation $\widetilde{O}(\cdot)$ suppresses dependence on polylogarithmic factors.

with probability at least $1 - \delta$ over the randomness of the noise. Here, the total number of timesteps is the product of the number of queries of the form (4.1) times length of the queries. That is, the number of timesteps is $m \times T$, where m is number of measurements taken and T denotes the length of each measurement. This quantity depends on the set \mathcal{U} ; we restrict (for now) to the case where \mathcal{U} is either the unit ℓ_2 -ball or the unit ℓ_∞ -ball. These two sets comprise the most common input constraints found in the controls literature. Nevertheless, our analysis later will cover all ℓ_p -balls for $p \in [1, \infty]$.

In both cases, we must first consider the length of the FIR filter (4.2) required to ensure reasonable approximation in the \mathcal{H}_∞ -norm. We must guarantee that we are able to accurately capture the large components of the impulse response of G . Therefore, we will need some measure of how quickly the impulse response coefficients tend to zero. It turns out that the \mathcal{H}_∞ -norm of G provides a convenient proxy for the decay of the impulse response. Throughout, we will use the following sufficient bound on the truncation length:

Definition 4.1.1 (Sufficient Length Condition). *Let G be stable with stability radius $\rho \in (0, 1)$. Fix a $\varepsilon > 0$. Let $R(\varepsilon)$ be the smallest integer which satisfies*

$$R(\varepsilon) \geq \inf_{\rho < \gamma < 1} \frac{1}{1 - \gamma} \log \left(\frac{\|G(\gamma z)\|_{\mathcal{H}_\infty}}{\varepsilon(1 - \gamma)} \right). \quad (4.3)$$

Equation (4.3) characterizes the approximation error of an FIR filter to G as a balance between the growth of $1/(1 - \gamma)$ versus the decay of the logarithm of $\|G(\gamma z)\|_{\mathcal{H}_\infty}$, as γ varies between $(\rho, 1)$.

We first study the ℓ_2 -ball case. In this case, we will set all m inputs to an impulse; that is, $u^{(i)} = e_1$, where $e_1 \in \mathbb{R}^r$ is the first standard basis vector.

Theorem 4.1.1 (Main result, ℓ_2 -constrained case). *Fix an $\varepsilon > 0$ and $\delta \in (0, 1)$, and suppose that $\mathcal{U} = \{x \in \mathbb{R}^T : \|x\|_2 \leq 1\}$. Let G be stable with stability radius $\rho \in (0, 1)$, and set $r \geq R(\varepsilon/2)$ from (4.3) and $T = 2r$. Set m inputs $u^{(i)} = e_1 \in \mathcal{U}$ for $i = 1, \dots, m$, where $m \geq 1$ satisfies*

$$m \geq C \frac{\sigma^2 r}{\varepsilon^2} \left(\log r + \log \left(\frac{1}{\delta} \right) \right).$$

Then, with probability at least $1 - \delta$, we have $\|G - \widehat{G}_r\|_{\mathcal{H}_\infty} \leq \varepsilon$. Above, C is an absolute positive constant.

Theorem 4.1.1 states that the number of timesteps to achieve identification error ε with ℓ_2 -constrained inputs scales as $\widetilde{O}(\sigma^2 r^2 / \varepsilon^2)$ in the regime when $\sigma / \varepsilon \gg 1$. It also turns out that this input ensemble is optimal for the ℓ_2 -ball case, which we will discuss shortly.

We next turn to the ℓ_∞ -ball case. In this case, we take $m = 2n$ to be an even number, and construct the input ensemble

$$u_{c,t}^{(i)} = \cos \left(\frac{2\pi it}{n} \right) \quad \text{and} \quad u_{s,t}^{(i)} = \sin \left(\frac{2\pi it}{n} \right) \quad \text{for } i = 0, \dots, n - 1. \quad (4.4)$$

With this input ensemble, we prove the following result for ℓ_∞ -constraints.

Theorem 4.1.2 (Main result, ℓ_∞ -constrained case). *Fix an $\varepsilon > 0$ and $\delta \in (0, 1)$, and suppose that $\mathcal{U} = \{x \in \mathbb{R}^T : \|x\|_\infty \leq 1\}$. Let G be stable with stability radius $\rho \in (0, 1)$, and set $r \geq R(\varepsilon/2)$ from (4.3) and $T = 2r$. Set m inputs as described in (4.4), where $m \geq 4r$ satisfies*

$$m \geq C \frac{\sigma^2}{\varepsilon^2} \left(\log r + \log \left(\frac{1}{\delta} \right) \right). \quad (4.5)$$

Then, with probability at least $1 - \delta$, we have $\|G - \widehat{G}_r\|_{\mathcal{H}_\infty} \leq \varepsilon$. Above, C is an absolute positive constant.

In the regime when $\sigma/\varepsilon \gg 1$, Theorem 4.1.2 states that the number of timesteps to achieve identification error ε with ℓ_∞ -constrained inputs scales as $\widetilde{O}(\sigma^2 r / \varepsilon^2)$. This is substantially more efficient than the complexity $\widetilde{O}(\sigma^2 r^2 / \varepsilon^2)$ which arises in the ℓ_2 -constrained case. We conclude by noting that this particular input ensemble is optimal for the ℓ_∞ -ball case up to constants, which we turn our attention to now.

In the case where, G itself is an length- r FIR filter, we also have a lower bound for estimation. We consider the general case when all m inputs are constrained to a unit ℓ_p -ball, where $p \in [1, \infty]$.

Theorem 4.1.3 ([Tu et al., 2017, Thm. 1.3]). *Fix a $p \in [1, \infty]$ and $r \geq 16$. Suppose that $m \geq 1$ measurements $u^{(1)}, \dots, u^{(m)} \in \mathbb{R}^T$ are fixed beforehand, with $T = 2r$ and $\|u^{(i)}\|_p \leq 1$ for all $i = 1, \dots, m$. Let \mathcal{H}_r denote the space of all length- r FIR filters. We have that*

$$\inf_{\widehat{G}} \sup_{G \in \mathcal{H}_r} \mathbb{E} \|\widehat{G} - G\|_{\mathcal{H}_\infty} \geq C \sigma \sqrt{\frac{r^{2/\max(p,2)} \log r}{m}},$$

where the infimum ranges over all measurable functions $\widehat{G} : \otimes_{k=1}^m \mathbb{R}^T \rightarrow \mathcal{H}_r$, and C is an absolute positive constant.

In view of Theorem 4.1.3, we see that the rates prescribed by Theorem 4.1.1 for the ℓ_2 -constrained case and by Theorem 4.1.2 for the ℓ_∞ -constrained case are minimax optimal up to constant factors. We conclude by noting that our choice of $T = 2r$ is arbitrary. Indeed, the same results hold for $T = \lceil (1 + \varepsilon)r \rceil$ for any fixed $\varepsilon > 0$, which only a change in the constant factors.

4.2 Related Work

4.2.1 Transfer Function Identification

Estimating the transfer function of a linear time-invariant system from input/output pairs has been studied in various forms in both the controls literature [Ljung and Wahlberg, 1992; Ljung and Yuan, 1985] and the statistics literature [Gerencsér, 1992; Goldenshluger

and Zeevi, 2001; Shibata, 1980], where it is closely related to estimating the coefficients of a stable autoregressive (AR) process. The main difference between our work and that of autoregressive estimation is that we assume the noise process driving the system is chosen by the practitioner (which we denote as the input to the system), and the stochastic component enters only during the output of the system. This simplifying assumption allows us to provide stronger non-asymptotic guarantees. Also by making prior assumptions on the stability radius of the underlying system, we circumvent the delicate issue of model order selection; a similar assumption is made in Goldenshluger and Zeevi [2001].

Most closely related to our work is that of Goldenshluger [1998], where he considers the problem of estimating the impulse response coefficients of a stable SISO LTI system. Goldenshluger provides upper and lower bounds on the ℓ_p -error when the residual between the estimate and the true coefficients is treated as a sequence in ℓ_p for $p \in [1, \infty]$. The main difference between Goldenshluger's setting and ours is that he restricts himself to the case when the input u is ℓ_∞ -constrained, and furthermore assumes only a single realization is available. On the other hand, we make assumption that multiple independent realizations of the system are available, which is reasonable in a controlled laboratory setting. This assumption simplifies the analysis and allows us to study more general ℓ_p -constrained inputs.

4.2.2 System Identification

We now turn our attention to system identification, where the classical results can be found in Ljung [1999]. Sample complexity guarantees in the system identification literature often require strong assumptions, which are difficult to verify. Most analyses are asymptotic and are based on the idea of *persistence of excitation* or *mixing* [McDonald et al., 2017; Vidyasagar and Karandikar, 2008]. There has been some progress in estimating the sample complexity of dynamical system identification using machine learning tools [Campi and Weyer, 2002; Vidyasagar and Karandikar, 2008], but such results typically yield pessimistic sample complexity bounds that are exponential in the degree of the linear system or other relevant quantities.

Two recent results provide polynomial sample complexity for identifying linear dynamical systems. Shah et al. [2012] show that if certain frequency domain measurements are obtained from a linear dynamical system, then the system can be approximately identified by solving a second-order cone programming problem. The degree of the estimated IIR system scales as $(1 - \rho(A))^{-2}$ where $\rho(A)$ denotes the stability radius. Similarly, Hardt et al. [2016] show that one can estimate an IIR system from time domain observations with a number of measurements polynomial in $(1 - \rho(A))^{-2}$, under the assumption that the impulse response coefficients $\{g_k\}_{k \geq 0}$ satisfy the decay law $|g_k| \leq C\rho(A)^k$, where C is considered a constant independent of the degree of the system. In this work, we show that under the same decay assumption, a considerably smaller FIR approximation with degree $\tilde{O}((1 - \rho(A))^{-1})$ suffices to complete many control design tasks. However, there has been a significant amount of recent work from the machine learning community on non-asymptotic rates for estimation in LTI systems; much of it is focused on removing this dependence on the

decay assumption and extending the results to marginally stable systems and/or partially observed systems [Faradonbeh et al., 2017; Hazan et al., 2017, 2018; Oymak and Ozay, 2018; Sarkar and Rakhlin, 2018; Sarkar et al., 2019; Simchowitz et al., 2019].

4.2.3 Robust Control

Classical robust control literature focuses much of its effort on designing a controller while taking into account fixed bounds on the uncertainty in the model. There are numerous algorithms for controller synthesis under various uncertainty specifications, such as coprime factor uncertainty [McFarlane and Glover, 1992] or state-space uncertainty [Packard and Doyle, 1993]. However, there are only a few branches of the robust control literature that couple identification to control design, and the identification procedure best suited for a particular control synthesis scheme is usually not specified.

4.2.4 \mathcal{H}_∞ Identification and Gain Estimation

Most related to our work is the literature on \mathcal{H}_∞ identification. In this literature, noisy input/output data from an unknown stable linear time-invariant (LTI) plant is collected in either the frequency domain [Helmicki et al., 1991; Hindi et al., 2002] or the time domain Chen and Nett [1993]; the goal is often to estimate a model with low \mathcal{H}_∞ error. A comprehensive review of this line of work is given by Chen and Gu [2000].

The main difference between the \mathcal{H}_∞ identification literature and our work is that we assume a probabilistic noise model instead of worst-case (adversarial), and we assume that our identification algorithm is allowed to pick its inputs to the plant G . As we will see, these simplifying assumptions lead to simple algorithms, straightforward analysis, and finite-time sample complexity guarantees.

Another related line of work is the use of the power method [Rojas et al., 2012; Wahlberg et al., 2010] for estimating the \mathcal{H}_∞ -norm of an unknown SISO plant. The key insight is that in the SISO case, a time-reversal trick can be applied to effectively query the system $G^* \circ G$, where G^* denotes the adjoint system. This approach is appealing, since the power method is known to converge exponentially quickly to the leading eigenvector. However, the leading factor in the convergence rate is the ratio of λ_1/λ_2 , and hence providing a finite-time guarantee of this method would require a non-asymptotic analysis of the rate of convergence of the *second* singular value of finite sections of a Toeplitz operator.

4.2.5 Norms of Random Polynomials

Our analysis relies on bounding the norms of random trigonometric polynomials of the form $Q(z) = \sum_{k=0}^{r-1} \varepsilon_k z^k$. The study of the supremum norm of random finite degree polynomials was first initiated by Salem and Zygmund [1954], who studied the setting where the coefficients are drawn from a symmetric Bernoulli distribution supported on $\{\pm 1\}$. Later, Kahane [1994] proved that when the coefficients are distributed as an isotropic Gaussian, then with

probability at least $1 - \delta$, $\|Q\|_\infty \leq O(\sqrt{r \log(r/\delta)})$. More recently, Meckes [2007] extended this result to hold for independent sub-Gaussian random variables by employing standard tools from probability in Banach spaces. In Section 4.3, we extend these results to the case when the coefficients follow a non-isotropic Gaussian distribution. This is important because it allows us to reduce the overall error of our estimate by using non-isotropic covariance matrices from experiment design.

4.3 System Identification of Finite Impulse Responses

Recall from Section 4.1.1 that we are given query access to G via the form

$$y_{u,T} = (g * u)_{k=0}^{T-1} + \xi, \quad \xi \sim \mathcal{N}(0, \sigma^2 I_T),$$

where $u \in \mathcal{U}$ for some fixed set \mathcal{U} . Therefore, the ratio of some measure of the size of \mathcal{U} to σ serves as the signal-to-noise (SNR) ratio for our setting. In what follows, we will always assume \mathcal{U} is a unit ℓ_p -ball for $p \in [1, \infty]$.

Fix a set of m inputs $u^{(1)}, \dots, u^{(m)} \in \mathcal{U}$. Given a realization of $\{y_{u_k, T}\}_{k=1}^m$, we can estimate the first T coefficients $\{g_k\}_{k=0}^{T-1}$ of $G(z) = \sum_{k=0}^{\infty} g_k z^{-k}$ via ordinary least-squares (OLS). Calling the vector $Y := (y_{u^{(1)}, T}, \dots, y_{u^{(m)}, T}) \in \mathbb{R}^{Tm}$, it is straightforward to show that the least squares estimator $\hat{g}_{0:T-1}$ is given by

$$\hat{g}_{0:T-1} := \begin{bmatrix} \hat{g}_0 \\ \hat{g}_1 \\ \vdots \\ \hat{g}_{T-1} \end{bmatrix} = (Z^\top Z)^{-1} Z^\top Y, \quad Z := \begin{bmatrix} \text{Toep}(u^{(1)}) \\ \vdots \\ \text{Toep}(u^{(m)}) \end{bmatrix} \in \mathbb{R}^{Tm \times T}.$$

Let us clarify the $\text{Toep}(u)$ notation. For a vector $u \in \mathbb{R}^T$, $\text{Toep}(u)$ is the $T \times T$ lower-triangular Toeplitz matrix where the first column is equal to u . Later on, we will use the notation $\text{Toep}_{a \times b}(u)$, where a, b are positive integers. This is to be interpreted as the upper left $a \times b$ section of the semi-infinite lower-triangular Toeplitz matrix form by treating u as a zero-padded sequence in $\mathbb{R}^{\mathbb{N}}$.

Above, we assume the matrix $Z^\top Z$ is invertible, which we will ensure in our analysis. From $\hat{g}_{0:T-1}$, we form the estimated finite impulse response \hat{G}_r for any $r \leq T$ as $\hat{G}_r(z) := \sum_{k=0}^{r-1} \hat{g}_k z^{-k}$. The Gaussian output noise assumption means that the error vector $\hat{g}_{0:T-1} - g_{0:T-1}$ is distributed $\mathcal{N}(0, \sigma^2 (Z^\top Z)^{-1})$, and hence $\hat{G}_r - G_r$ is equal in distribution to the random polynomial $Q(z) = \sum_{k=0}^{r-1} \varepsilon_k z^{-k}$ with $\varepsilon \sim \mathcal{N}(0, \sigma^2 E_r (Z^\top Z)^{-1} E_r^\top)$, where $E_r := [I_r \quad 0_{r \times (T-r)}] \in \mathbb{R}^{r \times T}$. Here, $G_r(z) := \sum_{k=0}^{r-1} g_k z^{-k}$ is the length- r FIR truncation of G . Since the covariance matrix will play a critical role in our analysis to follow, we introduce the notation

$$\Sigma(u) := \sum_{k=1}^m \text{Toep}(u_k)^\top \text{Toep}(u_k), \quad (4.6)$$

where m will be clear from context. We will also use the shorthand notation $[M]_{[r]}$, to refer to the $r \times r$ matrix $E_r M E_r^\top$ for any $T \times T$ matrix M .

The roadmap for this section is as follows. In Section 4.3.1, we characterize the behavior of the random quantity $\|Q\|_\infty$ as a function of the covariance matrix $\sigma^2 \Sigma(u)^{-1}$ and the polynomial degree r . Next, we study in Section 4.3.2 the problem of experiment design for choosing the best inputs $u^{(1)}, \dots, u^{(m)}$ to minimize the error $\|Q\|_\infty$. Using these results, we give upper bounds for FIR identification with ℓ_p -constrained inputs in Section 4.3.3. We then combine these results and prove in Section 4.3.4 the main results from Theorem 4.1.1 and Theorem 4.1.2. We will omit some sections and proofs for brevity and refer the reader to Tu et al. [2017].

Process Noise Before we begin our analysis, we note that our upper bounds easily extend to the case where process noise enters the system through the same channel as the input. Specifically, suppose the input signal is corrupted by $\zeta \sim N(0, \sigma_n^2 I_T)$ which is independent of the output noise ξ , and instead of observing $Y_{u,T}$ we observe

$$\tilde{y}_{u,T} = (g * \tilde{u})_{k=0}^{T-1} + \xi, \quad \tilde{u} := u + \zeta.$$

In this setting, the error vector $\hat{g}_{0:T-1} - g_{0:T-1}$ of the least-squares estimator on $\{\tilde{y}_{u_k,T}\}_{k=1}^m$ is distributed

$$\mathcal{N}(0, \Lambda), \quad \Lambda := (Z^\top Z)^{-1} Z^\top (\sigma_n^2 \text{Toep}_{T \times T}(g) \text{Toep}_{T \times T}(g)^\top + \sigma^2 I_T) Z (Z^\top Z)^{-1}.$$

Since g is the impulse response of a stable system, $\|\text{Toep}_{T \times T}(g)\| \leq \|G\|_{\mathcal{H}_\infty}$, and therefore $\Lambda \preceq (\sigma_n^2 \|G\|_{\mathcal{H}_\infty}^2 + \sigma^2) (Z^\top Z)^{-1}$. Thus, the upper bounds carry over to this process noise setting with the variable substitution $\sigma^2 \leftarrow \sigma_n^2 \|G\|_{\mathcal{H}_\infty}^2 + \sigma^2$. The modification to the lower bounds in this setting is more delicate, and we leave this to future work.

4.3.1 A concentration result for the error polynomial

We first address the behavior of the error $\|Q\|_\infty$. Our main tool is a discretization result from Bhaskar et al. [2012]:

Lemma 4.3.1 (Bhaskar et al. [2012]). *Let $Q(z) := \sum_{k=0}^{r-1} \varepsilon_k z^{-k}$, where $\varepsilon_k \in \mathbb{C}$. For any $N \geq 4\pi r$,*

$$\|Q\|_\infty \leq \left(1 + \frac{4\pi r}{N}\right) \max_{k=0, \dots, N-1} |Q(e^{j2\pi k/N})|.$$

Lemma 4.3.1 immediately reduces controlling the \mathcal{H}_∞ -norm² of a system defined by a finite-degree polynomial to controlling the maxima of a finite set of points on the torus.

²The \mathcal{H}_∞ -norm is equivalent to the $\|\cdot\|_\infty$ norm when viewed as a complex polynomial.

Hence, upper bounding the expected value of $\|Q\|_\infty$ and showing concentration is straightforward. Before we state the result, we define some useful notation which we will use throughout this section. For a $z \in \mathbb{C}$, define the vector of monomials $\varphi(z)$ as

$$\varphi(z) := (1, z, z^2, \dots, z^{r-1}) \in \mathbb{C}^r, \quad \varphi_1(z) := \operatorname{Re}\{\varphi(z)\}, \quad \varphi_2(z) := \operatorname{Im}\{\varphi(z)\}, \quad (4.7)$$

where the length r will be implicit from context. Now, using standard concentration results for suprema of Gaussian processes (see e.g. [Boucheron et al. \[2013\]](#)) and [Lemma 4.3.1](#), we have the following error concentration result.

Lemma 4.3.2. *Let $\varepsilon \sim \mathcal{N}(0, V)$ where $\varepsilon \in \mathbb{R}^r$, and put $Q(z) = \sum_{k=0}^{r-1} \varepsilon_k z^{-k}$. Define for $\ell = 1, 2$,*

$$\eta_\ell^2 := \sup_{z \in \mathbb{T}} \varphi_\ell(z)^\top V \varphi_\ell(z).$$

We have that

$$\mathbb{E}\|Q\|_\infty \leq 4\sqrt{2}\eta\sqrt{\log(8\pi r)},$$

where $\eta := \max(\eta_1, \eta_2)$. Furthermore, with probability at least $1 - \delta$, we have

$$\|Q\|_\infty \leq 4\sqrt{2}\eta(\sqrt{\log(8\pi r)} + \sqrt{\log(2/\delta)}).$$

Note that when $V = I$, $\eta^2 \leq r$ which recovers the known results from [Kahane \[1994\]](#) up to constants. Furthermore, when V is diagonal, $\eta^2 \leq \operatorname{tr}(V)$. We will exploit this result in the sequel.

4.3.2 Experiment design

We now consider the problem of choosing a set of inputs $u \in \mathcal{U}$ in order to minimize the expected error of the residual polynomial. Fixing the number of inputs m , the input constraint set \mathcal{U} , and recalling the definition of the covariance $\Sigma(u)$ from [\(4.6\)](#), the optimal experiment design problem is

$$\underset{u^{(1)}, \dots, u^{(m)} \in \mathcal{U}}{\text{minimize}} \quad \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, [\Sigma(u)^{-1}]_{[r]})} \|Q\|_\infty. \quad (4.8)$$

In [\(4.8\)](#) and the sequel, if the covariance matrix $\Sigma(u)$ is not invertible then we assign the function value $+\infty$. Problem [\(4.8\)](#) is difficult to solve as written because the expected value does not have a form which is easy to work with computationally. The following design problem provides a good approximation of [\(4.8\)](#). Let $\{z_1, \dots, z_s\} \subseteq \mathbb{T}$ denote a grid of s points on \mathbb{T} . Consider the problem

$$\underset{u^{(1)}, \dots, u^{(m)} \in \mathcal{U}}{\text{minimize}} \quad \max_{\substack{1 \leq k \leq s \\ \ell=1,2}} \varphi_\ell(z_k)^\top [\Sigma(u)^{-1}]_{[r]} \varphi_\ell(z_k). \quad (4.9)$$

The objective (4.9) minimizes the maximum pointwise variance of $Q(z)$ over all points on the grid $\{z_1, \dots, z_s\}$. If the grid is uniformly spaced and $s \geq 4\pi r$, then by Lemma 4.3.1 we can interpret (4.9) as minimizing an upper bound to the objective function in (4.8), since

$$\begin{aligned} \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, [\Sigma(u)^{-1}]_{[r]})} \|Q\|_\infty &\leq (1 + 4\pi r/s) \mathbb{E} \max_{1 \leq k \leq s} |\langle \varphi(z_k), \varepsilon \rangle| \\ &\leq (1 + 4\pi r/s) \sqrt{2 \log(2s)} \sum_{\ell=1}^2 \max_{1 \leq k \leq s} \sqrt{\varphi_\ell(z_k)^\top [\Sigma(u)^{-1}]_{[r]} \varphi_\ell(z_k)}. \end{aligned} \quad (4.10)$$

However, (4.9) is non-convex in the $u^{(i)}$'s. A convex version of the problem can be written by choosing m_0 inputs $u^{(1)}, \dots, u^{(m)} \in \mathcal{U}$ and solving the semidefinite program (SDP)

$$\begin{aligned} \text{minimize}_{\lambda \in \mathbb{R}^m} \max_{\substack{1 \leq k \leq s \\ \ell=1,2}} \varphi_\ell(z_k)^\top [\Sigma^{-1}]_{[r]} \varphi_\ell(z_k) \quad \text{s.t.} \quad \Sigma = \sum_{i=1}^{m_0} \lambda_i \text{Toep}(u^{(i)})^\top \text{Toep}(u^{(i)}), \quad \lambda^\top \mathbf{1} = 1, \quad \lambda \geq 0. \end{aligned} \quad (4.11)$$

Equation (4.11) is a convex program and can be solved with any off-the-shelf solver such as MOSEK [MOSEK ApS, 2015].

We now study two special cases of \mathcal{U} to show how input constraints can affect design. We first observe that when $\Sigma(u)$ is diagonal, continuing the estimates from (4.10), we have the following upper bound which holds since $\|\varphi_\ell(z)\|_\infty \leq 1$,

$$\mathbb{E}_{\varepsilon \sim \mathcal{N}(0, [\Sigma(u)^{-1}]_{[r]})} \|Q\|_\infty \leq (1 + 4\pi r/s) 2 \sqrt{2 \log(2s) \text{tr}([\Sigma(u)^{-1}]_{[r]})}. \quad (4.12)$$

Even though (4.12) only holds when $\Sigma(u)$ is diagonal, it motivates us to consider the standard A -optimal design problem

$$\text{minimize}_{u^{(1)}, \dots, u^{(m)} \in \mathcal{U}} \text{tr}([\Sigma(u)^{-1}]_{[r]}). \quad (4.13)$$

An advantage of (4.13) versus (4.9) is that the reduced complexity of the objective function allows us to make statements about optimality when \mathcal{U} is an ℓ_p -ball. The analogous SDP relaxation of (4.13), similar to (4.11), is also more efficient to implement in practice for more general \mathcal{U} 's.

Let $F_p^*(T, r)$ denote the optimal value of (4.13) with $\mathcal{U} = B_p^T$ for $p \in [1, \infty]$, where B_p^T denotes the unit ℓ_p -ball in \mathbb{R}^T . We will always assume $T \geq r$. It is not hard to show that $F_p^*(T, r)$ is finite and the value is attained (and hence $\Sigma(u)$ at the optimum is invertible). We claim the following statements about (4.13):

- (a) When $p \in [1, 2]$, the optimal solution is to set $u^{(i)} = e_1$, $i = 1, \dots, m$.
- (b) When $p \in (2, \infty]$, we can solve (4.13) to within a factor of two of optimal by convex programming.

- (c) When $p = \infty$, we can give an exact closed form solution for (4.13) in the special case when $r = 2^n$ with $n \geq 0$, $T = 2^k r$ with $k \geq 1$, and m is a multiple of T .

These statements follow from the following three lemmas, which are proved in full in [Tu et al. \[2017, Section 3.2\]](#).

First, we have an optimization problem which lower bounds the optimal value F_p^* .

Lemma 4.3.3. *For all $p \in [2, \infty]$, we have that*

$$F_p^*(T, r) \geq \frac{1}{m} \inf_{\substack{w \in \mathbb{R}^T \\ \|w\|_{p/2} \leq 1, w \geq 0}} \sum_{i=1}^r \left[\sum_{\ell=1}^{T-r+i} w_\ell \right]^{-1} := \frac{1}{m} D_p(T, r).$$

Lemma 4.3.4. *For $p \in [1, 2]$, we have that $F_p^*(T, r) = \frac{r}{m}$, which is achieved by setting $u^{(1)} = \dots = u^{(m)} = e_1$.*

Lemma 4.3.5. *Fix positive integers m, n, T , and suppose that $m = 2n$ and $n \geq T$. Let $z_i = e^{2\pi j i/n}$ for $i = 0, \dots, n-1$. Define the vectors $u^{(0)}, \dots, u^{(n-1)}$ and $u^{(n)}, \dots, u^{(2n-1)}$ as*

$$u^{(i)} = \operatorname{Re}\{w_p(T, r) \odot \varphi(z_i)\}, \quad u^{(n+i)} = \operatorname{Im}\{w_p(T, r) \odot \varphi(z_i)\}, \quad i = 0, \dots, n-1.$$

We have that the covariance matrix $\Sigma(u)$ satisfies

$$\operatorname{tr}([\Sigma(u)^{-1}]_{[r]}) = \frac{2}{m} D_p(T, r).$$

4.3.3 Upper Bounds on FIR Identification with ℓ_p -constrained Inputs

Combining the results from Section 4.3.1 and Section 4.3.2, we now prove an upper bound on length- r FIR identification when the inputs are ℓ_p -constrained.

Lemma 4.3.6. *Fix positive integers m and r , and set $T = 2r$. Consider the input ensemble $u^{(1)} = \dots = u^{(m)} = e_1$ when $p \in [1, 2]$, or the input ensemble defined in Lemma 4.3.5 when $p \in (2, \infty]$ (with additional restrictions on m, r in this case). Let \widehat{G}_r denote the length- r FIR estimate derived from least-squares, and let G_r denote the length- r FIR truncation of G . With probability at least $1 - \delta$,*

$$\|\widehat{G}_r - G_r\|_{\mathcal{H}_\infty} \leq \begin{cases} 4\sqrt{2}\sigma\sqrt{\frac{r}{m}} \left(\sqrt{\log(8\pi r)} + \sqrt{\log(2/\delta)} \right) & \text{if } p \in [1, 2] \\ 8\sqrt{2}\log 2\sigma\sqrt{\frac{r^{2/p}}{m}} \left(\sqrt{\log(8\pi r)} + \sqrt{\log(2/\delta)} \right) & \text{if } p \in (2, \infty]. \end{cases}$$

Proof. From Lemma 4.3.2, we have that with probability at least $1 - \delta$,

$$\|\widehat{G}_r - G_r\|_{\mathcal{H}_\infty} \leq 4\sqrt{2}\sigma\sqrt{\operatorname{tr}([\Sigma(u)^{-1}]_{[r]})} \left(\sqrt{\log(8\pi r)} + \sqrt{\log(2/\delta)} \right).$$

We only need to upper bound the variance term $\text{tr}([\Sigma(u)^{-1}]_{[r]})$. When $p \in [1, 2]$, by Lemma 4.3.4 the optimal input ensemble is the impulse response $u^{(i)} = e_1$, so we have $\text{tr}([\Sigma(u)^{-1}]_{[r]}) \leq r/m$. On the other hand, when $p \in (2, \infty]$, one can show (see Tu et al. [2017, Lemma 3.5]) that the specified u 's satisfy

$$\text{tr}([\Sigma(u)^{-1}]_{[r]}) \leq \frac{2}{m} D_p(2r, r) \leq \frac{4}{m} r^{2/p} (H_{2r} - H_r) \leq \frac{4 \log 2}{m} r^{2/p},$$

where H_r is the r -th harmonic number. ■

4.3.4 Proof of Main Results

We now prove Theorem 4.1.1 and Theorem 4.1.2. Recall that \widehat{G}_r is the estimated length- r FIR approximation to G , and G_r is the true length- r FIR truncation of G . By the triangle inequality, we have the following error decomposition into an approximation error and an estimation error

$$\|G - \widehat{G}_r\|_{\mathcal{H}_\infty} \leq \underbrace{\|G - G_r\|_{\mathcal{H}_\infty}}_{\text{Approx. error.}} + \underbrace{\|G_r - \widehat{G}_r\|_{\mathcal{H}_\infty}}_{\text{Estimation error.}}.$$

Hence, in order for $\|G - \widehat{G}_r\|_{\mathcal{H}_\infty} \leq \varepsilon$ to hold, it suffices to have both the approximation error $\|G - G_r\|_{\mathcal{H}_\infty} \leq \varepsilon/2$ and the estimation error $\|G_r - \widehat{G}_r\|_{\mathcal{H}_\infty} \leq \varepsilon/2$.

The approximation error is a deterministic quantity, and its behavior is governed by the tail decay of the impulse response coefficients $\{g_k\}_{k \geq 0}$. In Section 4.4, we prove in Lemma 4.4.1 that as long as r satisfies

$$r \geq \inf_{\rho < \gamma < 1} \frac{1}{1 - \gamma} \log \left(\frac{2\|G(\gamma z)\|_{\mathcal{H}_\infty}}{\varepsilon(1 - \gamma)} \right),$$

then we have $\|G - G_r\|_{\mathcal{H}_\infty} \leq \varepsilon/2$.

We now turn our attention to the estimation error. For the case when $p = 2$, Lemma 4.3.6 tells us with probability at least $1 - \delta$, the estimation error satisfies

$$\|G_r - \widehat{G}_r\|_{\mathcal{H}_\infty} \leq 4\sqrt{2}\sigma \sqrt{\frac{r}{m}} \left(\sqrt{\log(8\pi r)} + \sqrt{\log(2/\delta)} \right).$$

Setting the RHS less than $\varepsilon/2$, solving for m , and using the inequality $(a + b)^2 \leq 2(a^2 + b^2)$, we conclude that a sufficient condition on m is

$$m \geq \max \left\{ \frac{256\sigma^2 r}{\varepsilon^2} \left(\log(8\pi r) + \log\left(\frac{2}{\delta}\right) \right), 1 \right\}.$$

This concludes the proof of Theorem 4.1.1.

The proof of Theorem 4.1.2 is nearly identical. For the case when $p = \infty$, Lemma 4.3.6 tells us with probability at least $1 - \delta$, the estimation error satisfies

$$\|G_r - \widehat{G}_r\|_{\mathcal{H}_\infty} \leq 8\sqrt{2 \log 2} \sigma \sqrt{\frac{1}{m}} \left(\sqrt{\log(8\pi r)} + \sqrt{\log(2/\delta)} \right).$$

Setting the RHS less than $\varepsilon/2$ and solving for m , we conclude that the sufficient condition is

$$m \geq \max \left\{ \frac{(1024 \log 2) \sigma^2}{\varepsilon^2} \left(\log(8\pi r) + \log \left(\frac{2}{\delta} \right) \right), 4r \right\}.$$

This concludes the proof of Theorem 4.1.2.

4.4 Finite Truncation Error Analysis for Stable Systems

In Section 4.3, we presented both probabilistic guarantees and experiment design for identification of FIR systems of length r , which were independent of any system specific properties of G . In this section, we analyze how system behavior affects the necessary truncation length needed to reach a desired approximation error tolerance.

In order to provide guarantees, we require that the underlying system G is stable with stability radius $\rho \in (0, 1)$. A standard fact states that stability is equivalent to the existence of a constant $C > 0$ such that the tail decay on the coefficients of the Laurent expansion $G = \sum_{k=0}^{\infty} g_k z^{-k}$ satisfies the following condition

$$|g_k| \leq C \rho^k, \quad k \geq 1. \quad (4.14)$$

Under this assumption, a calculation reveals that as long as

$$r \geq \frac{1}{1 - \rho} \log \left(\frac{C}{\varepsilon(1 - \rho)} \right),$$

then we have that the approximation error $\|G - G_r\|_{\mathcal{H}_\infty}$ satisfies $\|G - G_r\|_{\mathcal{H}_\infty} \leq \varepsilon$.

Unfortunately, without more knowledge of the system at hand, a bound on C in (4.14) is hard to characterize. However, by slightly relaxing the decay condition (4.14), we are able to derive a tail bound using system-theoretic ideas. Intuitively, if a system has long transient behavior, then we expect the constant C in (4.14) to be large, since in order to obtain a small approximation error one needs to capture the transient behavior. The next lemma shows that the \mathcal{H}_∞ -norm provides a sufficient characterization of this transient behavior; we include the proof for completeness.

Lemma 4.4.1 (Goldenshluger and Zeevi [2001, Lemma 1]). *Let $G(z) = \sum_{k=0}^{\infty} g_k z^{-k}$ be a stable SISO LTI system with stability radius $\rho \in (0, 1)$. Fix any γ satisfying $\rho < \gamma < 1$. Then for all $k \geq 1$,*

$$|g_k| \leq \|G(\gamma z)\|_{\mathcal{H}_\infty} \gamma^k.$$

Proof. Define the function $H(z) := G(z^{-1})$, which is analytic for all $|z| \leq 1/\gamma$. It is easy to check that k -th derivative of $H(z)$ evaluated at zero is $H^{(k)}(0) = k!g_k$. Therefore,

$$\begin{aligned} k!|g_k| &= |H^{(k)}(0)| \leq k!\gamma^k \max_{|z| \leq 1/\gamma} |H(z)| = k!\gamma^k \max_{|z| \geq \gamma} |G(z)| \\ &= k!\gamma^k \max_{|z| \geq 1} |G(\gamma z)| = k!\gamma^k \|G(\gamma z)\|_{\mathcal{H}_\infty}. \end{aligned}$$

Above, the first inequality is Cauchy's estimate formula for analytic functions, and the last equality follows from the maximum modulus principle. ■

We note that the technique used in Lemma 4.4.1 of considering the proxy system $G(\gamma z)$ instead of $G(z)$ directly also appears in Boczar et al. [2017] in the context of certifying exponential rates of convergence for linear dynamical systems.

4.5 Robust Controller Design

In Section 4.3, we described how to obtain a FIR system G_{fir} with a probabilistic guarantee that $G = G_{\text{fir}} + \Delta$, where Δ is an LTI system satisfying $\|\Delta\|_{\mathcal{H}_\infty} \leq \varepsilon$. This description of G naturally lends itself to many robust control synthesis methods. In this section, we describe the application of one particular method based on \mathcal{H}_∞ loop-shaping to a particular unknown plant.

Suppose that G is itself an FIR described with z -transform

$$G(z) = |w_0| + \sum_{k=1}^{149} |w_k| \rho^{k-1} z^{-k}, \quad \rho = 0.95, \quad (4.15)$$

where $w_k \sim \mathcal{N}(0, 1)$ are independent Gaussians. In this section, we will detail the design of a reference tracking controller for G using probabilistic guarantees.

4.5.1 Computing Bounds

While the non-asymptotic bounds of Section 4.3 and Section 4.4 give us upper bounds on the error of noisy FIR approximation, the constant factors in the bounds are not optimal. Hence, strictly relying on the bounds will cause oversampling by a constant factor of (say) 10 or more. For real systems, this is extremely undesirable—using the sharpest bound possible is of great practical interest. Fortunately, we can do this via simple Monte–Carlo simulations, which we detail in Section B.1 the appendix. For now, we describe the results of these simulations.

Our first Monte–Carlo simulation establishes that G satisfies the tail decay specified in (4.14) with $C = 3.9703$ and $\rho = 0.95$. If we truncate G with $r = 75$, we see that our worst-case upper bound on $\|G - G_r\|_{\mathcal{H}_\infty}$ is $C \frac{\rho^{r-1}}{1-\rho} = 3.9703 \times \frac{0.95^{74}}{1-0.95} = 1.7840$. In general, assuming we have no other information about G other than the bounds on C and ρ , this is

the sharpest approximation error bound possible, since for any system with real-valued, all non-negative Fourier coefficients, the \mathcal{H}_∞ -norm is simply the sum of the coefficients.

However, if we further assume we know the structure of G as in this case where we know the form of (4.15), but not the values of w_k , we can further sharpen our approximation bound. Specifically, we know that $E_{\text{approx}} := \|G - G_r\|_{\mathcal{H}_\infty} = \sum_{k=75}^{149} |w_k| \rho^{k-1}$, and hence we can perform another Monte–Carlo simulation to estimate the tail probability of this random variable. The result of our simulation is that $\mathbb{P}(E_{\text{approx}} \leq 0.46) \geq 0.99$. This is a substantial improvement over the previous bound of $E_{\text{approx}} \leq 1.7840$ which only uses the information contained in the tail decay.

We can use the same technique to sharpen the estimates from Lemma 4.3.2. We perform our final Monte–Carlo simulation, now on the random variable $E_{\text{noise}} := \frac{\sigma}{\sqrt{N}} \|\sum_{k=0}^{74} \xi_k z^{-k}\|_\infty$, with $\sigma^2 = 1$ and $\xi_k \sim \mathcal{N}(0, 1)$. Note that this corresponds to choosing the inputs to the system as impulse responses, which we recall from Section 4.3.2 is optimal under ℓ_2 -power constraints. With this simulation, we obtain that $\mathbb{P}(E_{\text{noise}} \leq 3.5954) \geq 0.99$.

4.5.2 Controller Design

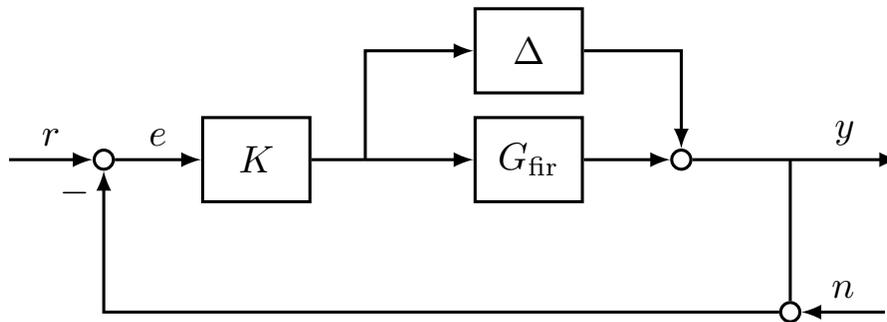


Figure 4.1: Closed-loop experimental setup. The goal is to design the controller K . G_{fir} is estimated from noisy output data, and $\|\Delta\|_{\mathcal{H}_\infty}$ is bounded via Monte–Carlo simulations.

Our goal is to design a controller K in the setup described in Figure 4.1, under the assumption that $\|\Delta\|_{\mathcal{H}_\infty} \leq E_{\text{approx}} + E_{\text{noise}} \leq 4.0554$. This assumption comes from the calculations in Section 4.5.1. We note that $\|\Delta\|_{\mathcal{H}_\infty} / \|G\|_{\mathcal{H}_\infty}$ fluctuates between 10-20%, so G_{fir} is a relatively coarse description of G . We use standard loop-shaping performance goals (see e.g. Doyle et al. [1990]). Let $T_{r \rightarrow e}$ and $T_{n \rightarrow e}$ denote the transfer functions from $r \mapsto e$ and $n \mapsto e$, respectively. At low frequencies, we would like $|T_{r \rightarrow e}|$ to have small gain, and at high frequencies we would like $|T_{r \rightarrow e}| \leq 2$. Similarly, we would like $|T_{n \rightarrow e}| \leq 2$ at low frequencies and $|T_{n \rightarrow e}|$ small at high frequencies. Of course, we would like these goals to be achieved, in addition to closed loop stability, for all $G = G_{\text{fir}} + \Delta$.

We proceed in two steps. We first design a controller with the nominal G_{fir} using \mathcal{H}_∞ loop-shaping synthesis (`mixsyn` in MATLAB). We choose weights to encourage our performance

goals on $T_{r \rightarrow e}$ and $T_{n \rightarrow e}$ to be met. Next, we check that our performance goal is met, in addition to robust stability. To make the computation easier, we check the performance goals separately. First, it is well known (see e.g. Doyle et al. [1990]) that the goal on $T_{r \rightarrow e}$ is met (in addition to robust stability) if the following holds

$$\| |W_1 S| + \gamma |K S| \|_{\mathcal{H}_\infty} < 1, \quad (4.16)$$

where $S = \frac{1}{1+KG_{\text{fir}}}$ and $\gamma = 4.0554$. Specifically, under (4.16), the closed loop with K in feedback with G is stable and achieves the performance guarantee $|T_{r \rightarrow e}(z)| \leq \frac{1}{|W_1(z)|}$ for every frequency $z \in \mathbb{T}$. On the other hand, to the best of our knowledge no simple expression for the performance goal on $T_{n \rightarrow e}$ exists, so we resort to a standard structured singular value (SSV) calculation [Packard and Doyle, 1993].

We generate our controller K via the following MATLAB commands

```
w_c = 0.07; % Cross-over freq
W1 = makeweight(5000, w_c, .5, 1); % Low-freq disturbance rejection
W2 = 1.5*fir_error_bound; % Robust stability
W3 = makeweight(.5, 3 * w_c, 5000, 1); % High-freq noise insensitivity
P = augw(G_fir, W1, W2, W3);
K = hinfsyn(P);
```

In Figure 4.2, we plot the open loop gain $L = G_{\text{fir}}K$, sensitivity function $S = 1/(1+L)$, and complementary sensitivity function $T = 1-S$. Here, we see that the cross-over frequency $\omega_c \approx 0.1$.

Next, in Figure 4.3, we plot the μ values for both the reference tracking objective $T_{r \rightarrow e}$ and the noise insensitivity objective $T_{n \rightarrow e}$, and check that both curves lies below 1 for all frequencies. Recall that this means that G in feedback with K is not only exponentially stable, but also satisfies both performance guarantees.

Finally, in Figure 4.4, we plot the output y as a function of a noisy square wave input u , to show the desired reference tracking behavior, on both the closed loop simulation (with G_{fir}), and the actual closed loop behavior (with G). This shows that, while the model G_{fir} was a coarse grained description of G with up to 20% relative error, it was faithful enough to allow for a robust controller design.

4.5.3 Varying truncation length

We next study the effect of truncation length r on controller design. In Figure 4.5, we assume the same setup and performance goals as the previous section, but vary the truncation length $r \in \{10, 30, 50, 70\}$. We also include the result of a controller design which has full knowledge of the true system G , which we label as **Opt**. We see that for $r = 10$, the resulting controller unsurprisingly has undesirable overshoot behavior. However, as r increases the resulting controller mimics the behavior of **Opt** quite closely. This plot shows that, at least for reference tracking behavior, a fairly low-fidelity model suffices. For instance, across different trials,

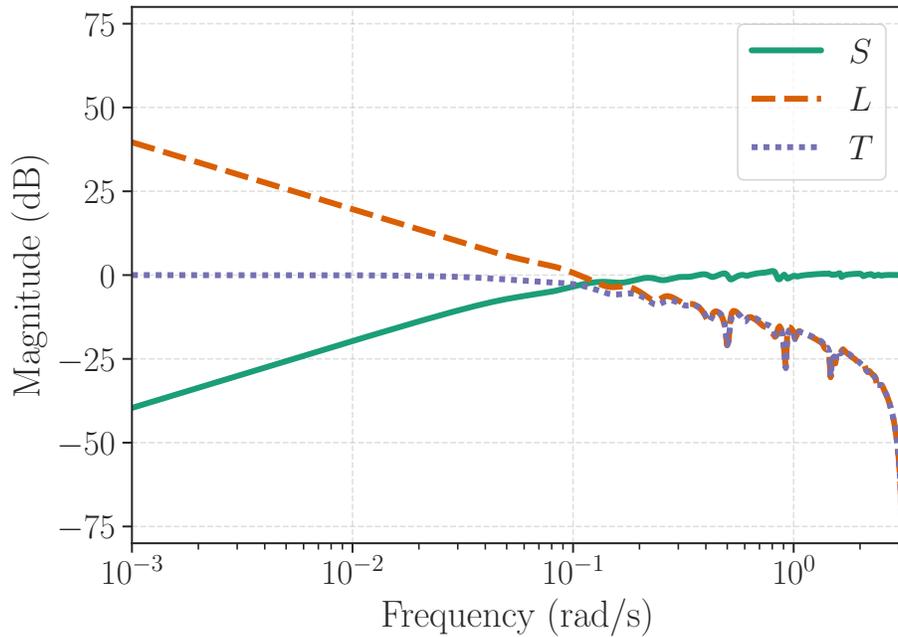


Figure 4.2: Loop-shaping curves from the experimental setup of Figure 4.1. $L = G_{\text{fir}}K$ denotes the open-loop gain, $S = 1/(1 + L)$ is the sensitivity function, and $T = 1 - S$.

the relative error of G_{fir} for $r = 30$ fluctuated between 15% to 30%, but in many cases $r = 30$ was able to provide reasonable reference tracking behavior.

4.6 Conclusion

This chapter explored the use of a coarse-grained FIR model estimated from noisy output data for control design. We showed that sharp bounds on the \mathcal{H}_∞ error between the true unknown plant and the estimated FIR filter can be derived using tools from concentration of measure, and the constant factors on these bounds can be further refined via Monte-Carlo simulation techniques. Finally, we demonstrated empirically that one can perform controller synthesis using only a coarse-grained approximation of the true system while meeting certain performance goals.

There are many possible future extensions of this work. We highlight a few ideas below.

MIMO Systems. While our approach can be generalized to the MIMO case by estimating filters for each input/output pair separately, we believe that when the MIMO transfer matrix has special structure (e.g. low rank), it should be possible to couple the estimation procedure to reduce the n^2 factor increase in sample complexity. This is motivated by the vast literature on compressed sensing, where sparse models embedded in a much larger ambient dimension

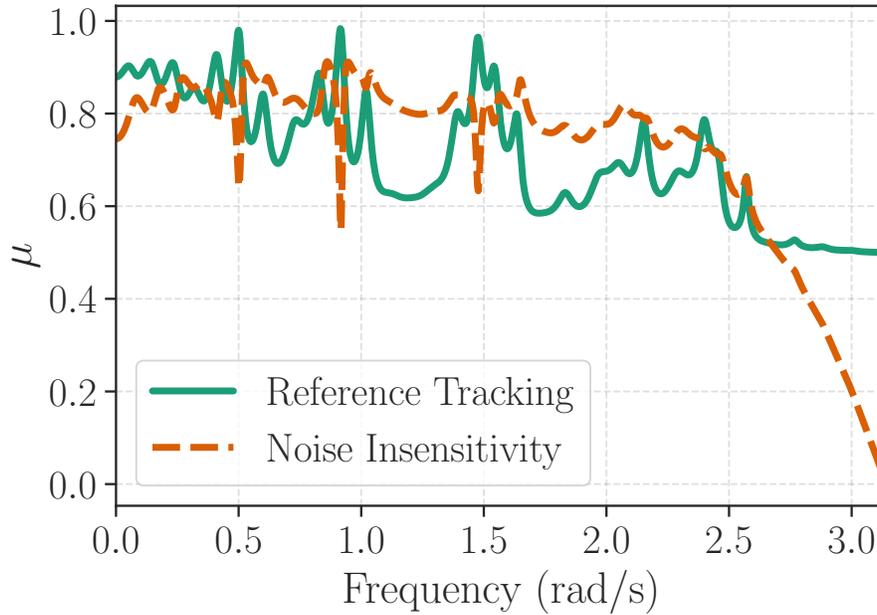


Figure 4.3: The pointwise frequency μ value for both reference tracking $T_{r \rightarrow e}$ and noise insensitivity $T_{n \rightarrow e}$. Robust performance is guaranteed as the curve lies below 1 at all frequencies.

can be uniquely recovered with at most a logarithmic factor more samples than the degree of the intrinsic sparsity.

Nonlinear Systems. An extension of these techniques to nonlinear systems is another exciting direction. One possible idea is to treat a nonlinear system's Jacobian linearization as the target unknown system, and fit a FIR using our techniques by exciting the nonlinear system locally. One would expect that the controller designed on the FIR would be valid in a neighborhood, and upon exiting the neighborhood, the process would repeat itself. The challenge here remains to estimate online the regime for which a controller is valid.

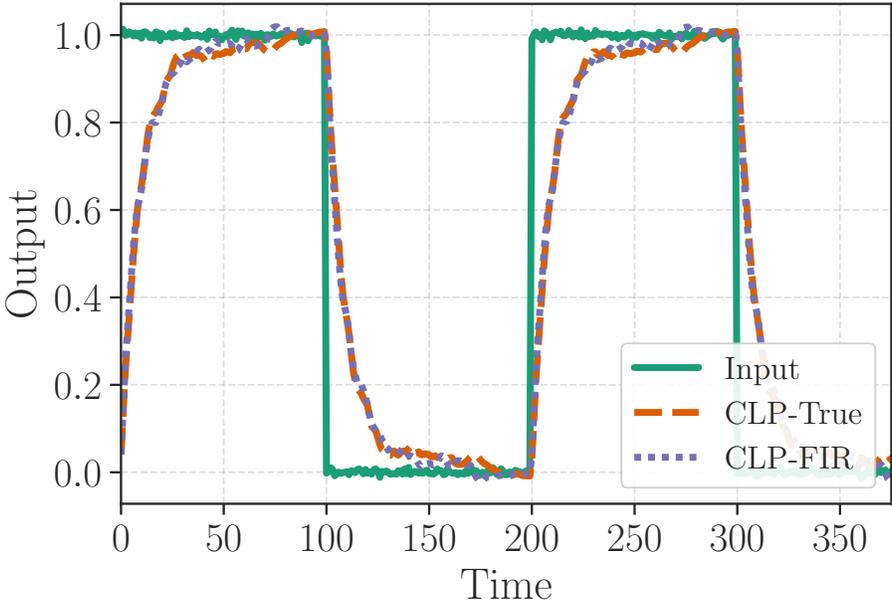


Figure 4.4: Reference tracking behavior of the closed loop with the model G_{fir} and the actual plant G .

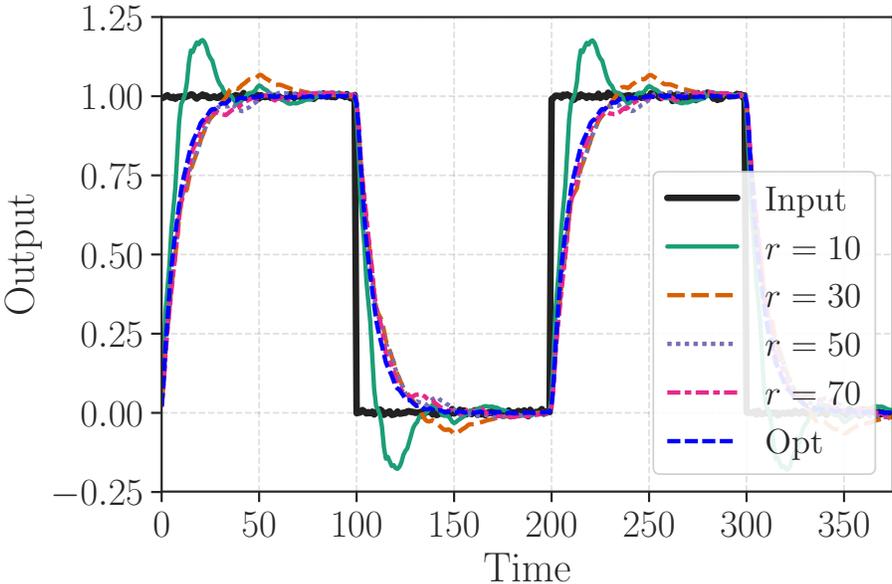


Figure 4.5: Reference tracking behavior as the FIR truncation length is varied, as well as the optimal nominal tracking.

Chapter 5

\mathcal{H}_∞ Bounds for Gain Estimation

5.1 Introduction

Recently, many researchers have proposed algorithms for estimating the \mathcal{H}_∞ -norm of a linear time-invariant (LTI) filter from input/output data [Müller et al., 2017; Oomen et al., 2014; Rallo et al., 2017; Rojas et al., 2012; van Heusden et al., 2007; Wahlberg et al., 2010]. A common property of these algorithms is eschewing model parameter estimation for directly estimating either the worst case ℓ_2 -input signal [Rojas et al., 2012; Wahlberg et al., 2010] or the maximizing frequency [Müller et al., 2017; Rallo et al., 2017]. One of the major motivations behind these algorithms is sample efficiency: since the \mathcal{H}_∞ -norm is a scalar estimate whereas the number of model parameters can be very large and possibly infinite, intuitively one expects that norm estimation can be performed using substantially fewer samples compared with model estimation.

In this chapter, we study the fundamental limits of estimating the \mathcal{H}_∞ -norm of a finite impulse response (FIR) filter, and compare to known bounds for FIR model estimation. We show that for *passive* algorithms which do not adapt their inputs in response to the result of previous queries, it is no more efficient to estimate the \mathcal{H}_∞ -norm than to estimate the model. For *active* algorithms which do adapt their inputs, we show that compared to model estimation, norm estimation is only at most a $\log r$ factor more efficient when the underlying model is a length r FIR filter. Our analysis raises an interesting open question: whether or not there exists an active sampling strategy which attains our stated lower bound.

Based on our theoretical findings, we study the empirical performance of several existing adaptive algorithms compared to a simple (non-adaptive) estimator which first fits a model via least-squares and then returns the norm of the model. Surprisingly, we find that the current adaptive methods do not perform significantly better than the simple estimator.

5.2 Related Work

Data-driven methods for estimating the \mathcal{H}_∞ -norm of an LTI system fall into roughly two major approaches: (a) estimating the worst case ℓ_2 -signal via a power-iteration type algorithm [Oomen et al., 2014; Rojas et al., 2012; Wahlberg et al., 2010] and (b) discretizing the interval $[0, 2\pi)$ and searching for the maximizing frequency [Müller et al., 2017; Rallo et al., 2017].

Algorithms that rely on power iteration take advantage of a clever time-reversal procedure introduced by Wahlberg et al. [2010], which allows one to query the adjoint system G^* with only input/output access to the original system G . One issue with these methods is that the rate of convergence of the top singular value of the truncated Toeplitz matrix to the \mathcal{H}_∞ -norm of the system is typically $O(1/n^2)$ (c.f. Böttcher and Grudsky [2000]), but the constant hidden in the $O(\cdot)$ notation can be quite large as pointed out in Tu et al. [2018b]. A non-asymptotic analysis of the statistical quality of the norm estimate returned by the power iteration remains an open question; asymptotic results can be found in Rojas et al. [2012].

The algorithms developed in Müller et al. [2017] and Rallo et al. [2017] are based on discretizing the frequencies $[0, 2\pi)$ and are rooted in ideas from the multi-armed bandit literature. Here, each frequency can be treated as either an “arm” or an “expert”, and an adaptive algorithm such as Thompson sampling [Müller et al., 2017] or multiplicative weights [Rallo et al., 2017] is applied. While sharp regret analysis for bandits has been developed by the machine learning and statistics communities [Bubeck and Cesa-Bianchi, 2012], one of the barriers to applying this analysis is a lack of a sharp theory for the level of discretization required. In practice, the number of grid points is a parameter that must be appropriately tuned for the problem at hand.

The problem of estimating the model parameters of a LTI system with model error measured in \mathcal{H}_∞ -norm is studied in Tu et al. [2017], and in ℓ_p -norms by Goldenshluger [1998]. For the FIR setting, the matching upper and lower bounds of Chapter 4 will serve as a baseline for us to compare our bounds with in the norm estimation setting. Helmicki et al. [1991] provide lower bounds for estimating both a model in \mathcal{H}_∞ -norm and its frequency response at a particular frequency in a query setting where the noise is worst-case. In this work we consider a less conservative setting with stochastic noise. Müller et al. Müller et al. [2017] prove an asymptotic regret lower bound over algorithms that sample only one frequency at every iteration. Their notion of regret is however defined with respect to the best frequency in a fixed discrete grid and not the \mathcal{H}_∞ -norm. As we discuss in Section 5.3.2, this turns out to be a subtle but important distinction.

5.3 Problem Setup and Main Results

In this section we formulate the problem under consideration and state our main results. We fix a filter length r and consider an unknown length r causal FIR filter $H(g) := \sum_{k=0}^{r-1} g_k z^{-k}$

with $g \in \mathbb{C}^r$. We study the following time-domain input/output query model for $H(g)$: for N rounds, we first choose an input $u^{(t)} \in \mathbb{C}^r$ such that $\|u^{(t)}\|_2 \leq M$, and then we observe a sample $y^{(t)} \sim \mathcal{N}(T(g)u^{(t)}, \sigma^2 I)$, where $T(g)$ denotes the $r \times r$ upper left section of the semi-infinite Toeplitz matrix induced by treating g as an element of ℓ_2 .¹ By $\mathcal{N}(\mu, \sigma^2 I)$ for a complex $\mu \in \mathbb{C}^r$ we mean we observe $(\text{Re}(\mu) + \xi_1) + j(\text{Im}(\mu) + \xi_2)$ where $\xi_i \sim \mathcal{N}(0, \sigma^2 I)$ and are independent.

After these N rounds, we are to return an estimate $\widehat{H} \in \mathbb{R}$ of the norm $\|H(g)\|_{\mathcal{H}_\infty}$ based on the collected data $(u^{(1)}, y^{(1)}, \dots, u^{(N)}, y^{(N)})$. The expected risk of any algorithm for this problem is measured as $\mathbb{E} \left[\left| \widehat{H} - \|H(g)\|_{\mathcal{H}_\infty} \right| \right]$, where the expectation is taken with respect to both the randomness of the algorithm and the noise of the outputs $y^{(t)}$.

Our results distinguish between *passive* and *active* algorithms. A passive algorithm is one in which the distribution of the input $u^{(t)}$ at time t is independent of the history $(u^{(1)}, y^{(1)}, \dots, u^{(t-1)}, y^{(t-1)})$. An active algorithm is one where the distribution of $u^{(t)}$ is allowed to depend on this history.

Given this setup, our first result is a minimax lower bound for the risk attained by any passive algorithm.

Theorem 5.3.1 (Passive lower bound). *Fix a $\gamma > 0$. Let $r \geq c$ for a universal constant $c > 0$ and $N \geq \text{poly}(r, M, 1/\gamma)$. We call a passive algorithm \mathcal{A} admissible if the matrix $\frac{1}{N} \sum_{t=1}^N \mathbb{E}_{u^{(t)}} [T(u^{(t)})^* T(u^{(t)})] \succeq \gamma I$. We have the following minimax lower bound on the risk of any passive admissible algorithm \mathcal{A} :*

$$\inf_{\mathcal{A}} \sup_{g \in \mathbb{C}^r} \mathbb{E} \left[\left| \widehat{H}_{\mathcal{A}} - \|H(g)\|_{\mathcal{H}_\infty} \right| \right] \geq C' \frac{\sigma}{M} \sqrt{\frac{r \log r}{N}}. \quad (5.1)$$

Here, C' is a universal constant. We note the form of the $\text{poly}(r, M, 1/\gamma)$ can be recovered from the proof.

The significance of Theorem 5.3.1 is due to the fact that under our query model, one can easily produce an estimate $\widehat{g} \in \mathbb{C}^r$ such that $\mathbb{E}[\|H(\widehat{g}) - H(g)\|_{\mathcal{H}_\infty}] \leq C'' \frac{\sigma}{M} \sqrt{\frac{r \log r}{N}}$; concretely setting $u^{(t)} = M e_1$ by appealing to the arguments in Chapter 4. That is, for passive algorithms the number of samples to estimate the r model parameters is equal (up to constants) to the number of samples needed to estimate the \mathcal{H}_∞ -norm, at least for the worst-case. The situation changes slightly when we look at active algorithms.

Theorem 5.3.2 (Active lower bound). *The following minimax lower bound on the risk of any active algorithm \mathcal{A} holds:*

$$\inf_{\mathcal{A}} \sup_{g \in \mathbb{C}^r} \mathbb{E} \left[\left| \widehat{H}_{\mathcal{A}} - \|H(g)\|_{\mathcal{H}_\infty} \right| \right] \geq C' \frac{\sigma}{M} \sqrt{\frac{r}{N}}. \quad (5.2)$$

¹We note that our results extend naturally to the setting when $T(g)$ is the $\alpha r \times \alpha r$ upper left section for a positive integer $\alpha \geq 1$. Furthermore, one can restrict both the system coefficients g and the inputs $u^{(t)}$ to be real-valued by considering the discrete cosine transform (DCT) instead of the discrete Fourier transform (DFT) in our proofs.

Here, C' is a universal constant.

We see in the active setting that the lower bound is weakened by a logarithmic factor in r . This bound shows that for low SNR regimes when $M \ll \sqrt{r}$, the gains of being active are minimal in the worst case. On the other hand, for high SNR regimes when $M \gg \sqrt{r}$, the gains are potentially quite substantial. We are unaware of any algorithm that provably achieves the active lower bound; it is currently unclear whether or not the lower bound is loose, or if more careful design/analysis is needed to find a better active algorithm. However, in Section 5.3.2 we discuss special cases of FIR filters for which the lower bound in Theorem 5.3.2 is not improvable.

We note that our proof draws heavily on techniques used to prove minimax lower bounds for functional estimation, specifically in estimating the ℓ_∞ -norm of a unknown mean vector in the Gaussian sequence model. An excellent overview of these techniques is given in Wu [2017, Chapter 22].

5.3.1 Hypothesis Testing and Sector Conditions

An application that is closely related to estimating the \mathcal{H}_∞ -norm is testing whether or not the \mathcal{H}_∞ -norm exceeds a certain fixed threshold τ . Specifically, consider a test statistic $\psi \in \{0, 1\}$ that discriminates between the two alternatives $H_0 : \|H(g)\|_{\mathcal{H}_\infty} \leq \tau$ and $H_1 : \|H(g)\|_{\mathcal{H}_\infty} > \tau$. This viewpoint is useful because it encompasses testing for more fine-grained characteristics of the Nyquist plot $H(\omega)$ via simple transformations. For instance, given finite $a < 0 < b$, one may test if $H(\omega)$ is contained within a circle in the complex plane centered at $(a+b)/2 + 0j$ with radius $(b-a)/2$ by equivalently checking if the \mathcal{H}_∞ -norm of the system $H(g) - (a+b)/2$ is less than $(b-a)/2$; this is known as the $[a, b]$ -sector condition [Gupta and Joshi, 1994].

Due to the connection between estimation and hypothesis testing, our results also give lower bounds on the sum of the Type-I and Type-II errors of any test ψ discriminating between the hypothesis H_0 and H_1 . Specifically, $\Omega(\frac{\sigma^2 r \log r}{\tau^2 M})$ queries in the passive case (when τ is sufficiently small) and $\Omega(\frac{\sigma^2 r}{\tau^2 M})$ queries in the active case are necessary for any test to have Type-I and Type-II error with less than constant probability.

5.3.2 Shortcomings of Discretization

In order to close the gap between the upper and lower bounds, one needs to explicitly deal with the continuum of frequencies on $[0, 2\pi)$; here we argue that if the maximizing frequency is known a-priori to lie in discrete set of points, then the lower bound is sharp.

Suppose we consider a slightly different query model, where at each time t the algorithm chooses a frequency $\omega_t \in [0, 2\pi)$ and receives $y^{(t)} \sim \mathcal{N}(H(\omega_t), 1)$. For simplicity, let us also assume that the \mathcal{H}_∞ -norm is upper bounded by 1. Note that by slightly enlarging $T(g)$ to the $2r \times 2r$ upper left triangle, we can emulate this query model by using a normalized complex sinusoid with frequency ω_t .

If the maximizing frequency for the \mathcal{H}_∞ -norm of the underlying system is located on

the grid $\{2\pi k/r\}_{k=0}^{r-1}$ and its phase is known, this problem immediately reduces to a standard r -arm multi-armed bandit (MAB) problem where each arm is associated with a point on the grid. For this family of instances, the following active algorithm has expected risk upper bounded by $\sqrt{r/N}$ times a constant:

- (a) Run a MAB algorithm that is optimal in the stochastic setting (such as MOSS [Audibert and Bubeck, 2009]) for $N/2$ iterations, with each of the r arms associated to a frequency $2\pi k/r$.
- (b) Sample an index $I \in \{1, \dots, r\}$ where:

$$\mathbb{P}(I = i) = \frac{\# \text{ number of times arm } i \text{ was pulled}}{N/2}.$$

- (c) Query ω_I for $N/2$ times and return the sample mean.

More generally, for any grid of frequencies $\{\omega_k\}$ such that the number of grid points is $O(r)$, this algorithm obtains $O(\sqrt{r/N})$ expected risk. Hence the lower bound of Theorem 5.3.2 is actually sharp with respect to these instances.

Therefore, the issue that needs to be understood is whether or not the continuum of frequencies on $[0, 2\pi)$ fundamentally requires additional sample complexity compared to a fixed discrete grid. Note that a naïve discretization argument is insufficient here. For example, it is known (see Bhaskar et al. [2012]) that by choosing P equispaced frequencies one obtains a discretization error bounded by $O(r/P)$, e.g. $\|H(g)\|_{\mathcal{H}_\infty} - |H(\omega_k)| \leq O(r/P)$ for the largest ω_k . This bound is too weak, however, since it requires that the number of arms scale as $O(r/\varepsilon)$ in order to obtain a risk bounded by ε ; in terms of N , the risk would scale $O(1/N^{1/3})$.

To summarize, if one wishes to improve the active lower bound to match the rate given by Theorem 5.3.1, one needs to consider a prior distribution over hard instances where the support of the maximizing frequency is large (possibly infinite) compared to r . On the other hand, if one wishes to construct an algorithm achieving the rate of Theorem 5.3.2, then one will need to understand the function $\omega \mapsto |H(\omega)|$ at a much finer resolution than Lipschitz continuity.

5.4 Proof of Main Results

The proof of Theorem 5.3.1 and Theorem 5.3.2 both rely on a reduction to Bayesian hypothesis testing. While this reduction is standard in the statistics and machine learning communities (see e.g. [Tsybakov, 2009, Chapter 2]), we briefly outline it here, as we believe these techniques are not as widely used in the controls literature.

First, let π_1, π_2 be two prior distributions on \mathbb{C}^r . Suppose that for all $\theta_1 \in \pi_1$ we have $\|H(\theta_1)\|_{\mathcal{H}_\infty} = 0$ and for all $\theta_2 \in \pi_2$ we have $\|H(\theta_2)\|_{\mathcal{H}_\infty} = 2c$ for some $c > 0$. Let \mathbb{P}_{π_i} denote

the joint distribution of $(u^{(1)}, y^{(1)}, \dots, u^{(N)}, y^{(N)})$, which combines the prior distribution π_i with the observation model. Then Markov's inequality implies that,

$$\sup_{\theta \in \mathbb{C}^r} \mathbb{E} \left[\left| \widehat{H} - \|H(\theta)\|_{\mathcal{H}_\infty} \right| \right] \geq \frac{c}{2} (1 - d_{\text{TV}}(\mathbb{P}_{\pi_1}, \mathbb{P}_{\pi_2})),$$

where for two measures \mathbb{P}, \mathbb{Q} we define the total-variation (TV) distance as

$$d_{\text{TV}}(\mathbb{P}, \mathbb{Q}) = \sup_A |\mathbb{P}(A) - \mathbb{Q}(A)|.$$

Hence, if one can construct two prior distributions π_1, π_2 with the aforementioned properties and furthermore show that $d_{\text{TV}}(\mathbb{P}_{\pi_1}, \mathbb{P}_{\pi_2}) \leq 1/2$, then one deduces that the minimax risk is lower bounded by $c/4$. This technique is known as Le Cam's method, and will be our high-level proof strategy.

As working directly with the TV distance is often intractable, one typically computes upper bounds to the TV distance. We choose to work with both the KL-divergence and the χ^2 -divergence. The KL-divergence is defined as $d_{\text{kl}}(\mathbb{P}, \mathbb{Q}) = \int \log \left(\frac{d\mathbb{P}}{d\mathbb{Q}} \right) d\mathbb{P}$, and the χ^2 -divergence is defined as $d_{\chi^2}(\mathbb{P}, \mathbb{Q}) = \int \left(\frac{d\mathbb{P}}{d\mathbb{Q}} - 1 \right)^2 d\mathbb{Q}$ (we assume that $\mathbb{P} \ll \mathbb{Q}$ so these quantities are well-defined). One has the standard inequalities $d_{\text{TV}}(\mathbb{P}, \mathbb{Q}) \leq \sqrt{\frac{1}{2} d_{\text{kl}}(\mathbb{P}, \mathbb{Q})}$ and $d_{\text{TV}}(\mathbb{P}, \mathbb{Q}) \leq \sqrt{d_{\chi^2}(\mathbb{P}, \mathbb{Q})}$ [Tsybakov, 2009].

Omitted proofs in this section can be found in the full paper Tu et al. [2018a].

5.4.1 Proof of Passive Lower Bound (Theorem 5.3.1)

The main reason for working with the χ^2 -divergence is that it operates nicely with mixture distributions, as illustrated by the following lemma.

Lemma 5.4.1 (see e.g. Lemma 22.1 of Wu [2017]). *Let Θ be a parameter space and for each $\theta \in \Theta$ let \mathbb{P}_θ be a measure over \mathcal{X} indexed by θ . Fix a measure \mathbb{Q} on \mathcal{X} and a prior measure π on Θ . Define the mixture measure $\mathbb{P}_\pi = \int \mathbb{P}_\theta \pi(d\theta)$. Suppose for every $\theta \in \Theta$, the measures \mathbb{P}_θ and \mathbb{Q} are both absolutely continuous w.r.t. a fixed base measure μ on \mathcal{X} . Define the function $G(\theta_1, \theta_2)$ as*

$$G(\theta_1, \theta_2) := \int \frac{\frac{d\mathbb{P}_{\theta_1}}{d\mu} \frac{d\mathbb{P}_{\theta_2}}{d\mu}}{\frac{d\mathbb{Q}}{d\mu}} \mu(dx).$$

Then, $d_{\chi^2}(\mathbb{P}_\pi, \mathbb{Q}) = \mathbb{E}_{\theta_1, \theta_2 \sim \pi^{\otimes 2}} [G(\theta_1, \theta_2)] - 1$.

We now specialize this lemma to our setting. Here, our distributions \mathbb{P}_θ are over the data, i.e. $(u^{(1)}, y^{(1)}, \dots, u^{(N)}, y^{(N)})$, for a fixed system parameter $\theta \in \mathbb{C}^r$. The joint distribution \mathbb{P}_θ

has the density (assuming that $u^{(t)}$ has the density $\gamma_t(u^{(t)})$)

$$p_\theta(u^{(1)}, y^{(1)}, \dots, u^{(N)}, y^{(N)}) = \prod_{t=1}^N \gamma_t(u^{(t)}) \phi(y^{(t)}; T(\theta)u^{(t)}),$$

where $\phi(\cdot; \mu)$ denotes the PDF of the multivariate Gaussian $\mathcal{N}(\mu, \sigma^2 I)$. Note that this factorization with $\gamma_t(\cdot)$ independent of θ is only possible under the passive assumption. A straightforward calculation gives the value of G under this density.

Lemma 5.4.2. *Supposing that $u^{(t)} \sim \gamma_t(\cdot)$, we have that*

$$G(\theta_1, \theta_2) = \mathbb{E}_{u^{(t)}} \left[\exp \left(\frac{1}{\sigma^2} \sum_{t=1}^N \operatorname{Re}(\langle T(\theta_1)u^{(t)}, T(\theta_2)u^{(t)} \rangle) \right) \right].$$

We now construct two prior distributions on θ . The first prior will be the system with all coefficients zeros, i.e. $\pi_1 = \{0\}$. The second prior will be more involved. To construct it, we let $\Sigma := \frac{1}{N} \sum_{t=1}^N \mathbb{E}_{u^{(t)} \sim \gamma_t} [T(u^{(t)})^* T(u^{(t)})]$. By the admissibility assumption on the algorithm \mathcal{A} , Σ is invertible. Let $\mathcal{I} \subseteq \{1, \dots, r\}$ denote an index set to be specified. Let $F \in \mathbb{C}^{r \times r}$ denote the unnormalized discrete Fourier transform (DFT) matrix (i.e. $FF^* = rI$ and $F^{-1} = \frac{1}{r}F^*$). We define our prior distribution π_2 as, for some $\tau > 0$ to be chosen, $\pi_2 = \operatorname{Unif}(\{\tau \Sigma^{-1/2} F^{-1} e_i\}_{i \in \mathcal{I}})$. We choose \mathcal{I} as according to the following proposition.

Proposition 5.4.1. *Let $u_1, \dots, u_N \in \mathbb{C}^r$ be independently drawn from N distributions such that $\|u^{(t)}\|_2 \leq M$ almost surely for all $t = 1, \dots, N$. Let $\Sigma = \frac{1}{N} \sum_{t=1}^N \mathbb{E}_{u^{(t)}} [T(u^{(t)})^* T(u^{(t)})]$ and suppose that Σ is invertible. There exists an index set $\mathcal{I} \subseteq \{1, \dots, r\}$ such that $|\mathcal{I}| \geq r/2$ and for all $i \in \mathcal{I}$,*

$$\|F \Sigma^{-1/2} F^{-1} e_i\|_\infty \geq \frac{1}{2M}.$$

Now defining

$$\Delta := \sum_{t=1}^N (T(u^{(t)})^* T(u^{(t)}) - \mathbb{E}_{u^{(t)}} [T(u^{(t)})^* T(u^{(t)})]),$$

we observe that for indices $j_1, j_2 \in \mathcal{I}$,

$$\begin{aligned} \sum_{t=1}^N \langle T(\Sigma^{-1/2} F^{-1} e_{j_1}) u^{(t)}, T(\Sigma^{-1/2} F^{-1} e_{j_2}) u^{(t)} \rangle &\stackrel{(a)}{=} \sum_{t=1}^N \langle T(u^{(t)}) \Sigma^{-1/2} F^{-1} e_{j_1}, T(u^{(t)}) \Sigma^{-1/2} F^{-1} e_{j_2} \rangle \\ &= \sum_{t=1}^N e_{j_1}^\top F^{-*} \Sigma^{-1/2} T(u^{(t)})^* T(u^{(t)}) \Sigma^{-1/2} F^{-1} e_{j_2} \\ &= e_{j_1}^\top F^{-*} \Sigma^{-1/2} (N \Sigma + \Delta) \Sigma^{-1/2} F^{-1} e_{j_2} \\ &= N e_{j_1}^\top F^{-*} F^{-1} e_{j_2} + e_{j_1}^* F^{-*} \tilde{\Delta} F^{-1} e_{j_2} \\ &= \frac{N}{r} \mathbf{1}_{j_1=j_2} + e_{j_1}^\top F^{-*} \tilde{\Delta} F^{-1} e_{j_2}, \end{aligned}$$

where $\tilde{\Delta} := \Sigma^{-1/2} \Delta \Sigma^{-1/2}$. In (a) we used the fact that $T(u)v = T(v)u$, as convolution is commutative. Combining this calculation with Lemma 5.4.2, we see that

$$\begin{aligned} \mathbb{E}_{\theta_i}[G(\theta_1, \theta_2)] &= \mathbb{E}_{j_i, u^{(t)}} \left[\exp \left(\frac{\tau^2 N}{\sigma^2 r} \mathbf{1}_{j_1=j_2} \right) \times \exp \left(\frac{\tau^2}{\sigma^2} \operatorname{Re} \left(e_{j_1}^\top F^{-*} \tilde{\Delta} F^{-1} e_{j_2} \right) \right) \right] \\ &\stackrel{(a)}{\leq} \sqrt{\mathbb{E}_{j_i, u^{(t)}} \left[\exp \left(\frac{2\tau^2 N}{\sigma^2 r} \mathbf{1}_{j_1=j_2} \right) \right]} \times \sqrt{\mathbb{E}_{j_i, u^{(t)}} \left[\exp \left(\frac{2\tau^2}{\sigma^2} \operatorname{Re} \left(e_{j_1}^\top F^{-*} \tilde{\Delta} F^{-1} e_{j_2} \right) \right) \right]} \\ &\stackrel{(b)}{\leq} \sqrt{\exp \left(\frac{2N\tau^2}{\sigma^2 r} \right) \frac{2}{r} + 1 - \frac{2}{r}} \times \sqrt{\mathbb{E}_{j_i, u^{(t)}} \left[\exp \left(\frac{2\tau^2}{\sigma^2} \operatorname{Re} \left(e_{j_1}^\top F^{-*} \tilde{\Delta} F^{-1} e_{j_2} \right) \right) \right]}. \end{aligned}$$

where in (a) we used Cauchy-Schwarz and in (b) we used the fact that $|\mathcal{I}| \geq r/2$. Now condition on j_1, j_2 . For a $1 \leq t \leq N$, define the random variable ψ_t as:

$$\begin{aligned} \psi_t &:= \operatorname{Re} \left(e_{j_1}^\top F^{-*} \Sigma^{-1/2} T(u^{(t)})^* T(u^{(t)}) \Sigma^{-1/2} F^{-1} e_{j_2} \right) \\ &\quad - \operatorname{Re} \left(e_{j_1}^\top F^{-*} \Sigma^{-1/2} \mathbb{E}_{u^{(t)}} [T(u^{(t)})^* T(u^{(t)})] \Sigma^{-1/2} F^{-1} e_{j_2} \right). \end{aligned}$$

We have that $\mathbb{E}_{u^{(t)}}[\psi_t] = 0$ by construction. Furthermore, note that $\|F^{-1}e_j\|_2 = 1/\sqrt{r}$ for $j = 1, \dots, r$ and also that $\|T(u)\| \leq \|H(u)\|_{\mathcal{H}_\infty} \leq \|u\|_1 \leq \sqrt{r}\|u\|_2$ for any vector $u \in \mathbb{C}^r$. These facts, along with the assumption that $\Sigma \succeq \gamma I$, show that $|\psi_t| \leq 2M^2/\gamma$ almost surely. Hence, $\sum_{t=1}^N \psi_t$ is a zero-mean sub-Gaussian random variable with sub-Gaussian parameter $4M^4N/\gamma^2$ (see Vershynin [2018, Chapter 2] for background exposition on sub-Gaussian random variables). Therefore, we know that for any $t > 0$, its moment generating function (MGF) is bounded as $\mathbb{E}_{u^{(t)}|j_i} \left[\exp \left(t \sum_{t=1}^N \psi_t \right) \right] \leq \exp(2t^2M^4N/\gamma^2)$. Hence by iterating expectations and setting $t = 2\tau^2/\sigma^2$, we have:

$$\mathbb{E}_{j_i, u^{(t)}} \left[\exp \left(\frac{2\tau^2}{\sigma^2} \operatorname{Re} \left(e_{j_1}^\top F^{-*} \Sigma^{-1/2} \Delta \Sigma^{-1/2} F^{-1} e_{j_2} \right) \right) \right] \leq \exp(8\tau^4M^4N/(\sigma^4\gamma^2)).$$

Therefore, for any choice of τ such that $\tau^4 \leq (\log(1.1)/8)\sigma^4\gamma^2/(M^4N)$, we have that:

$$\mathbb{E}_{\theta_i, u^{(t)}} [G(\theta_1, \theta_2)] \leq \sqrt{1.1} \sqrt{\exp \left(\frac{2N\tau^2}{\sigma^2 r} \right) \frac{2}{r} + 1 - \frac{2}{r}}.$$

Hence if $r \geq 5$ and if we set $\tau^2 = \frac{\sigma^2 r \log(0.211r)}{2N}$, we have:

$$d_{\chi^2}(\mathbb{P}_\pi, \mathbb{P}_0) = \mathbb{E}_{\theta_i, u^{(t)}} [G(\theta_1, \theta_2)] - 1 \leq 1/4,$$

assuming the previous condition on τ is satisfied. This then implies that $d_{\text{TV}}(\mathbb{P}_\pi, \mathbb{P}_0) \leq 1/2$. We now choose $N \geq (2/\log(1.1))r^2 \log^2(0.211r)M^4/\gamma^2 = \tilde{\Omega}(r^2M^4/\gamma^2)$ so that our condition on τ holds.

To conclude, we need to show a minimum separation between the \mathcal{H}_∞ -norm on π_1 vs. π_2 . Immediately, $\|H(\theta)\|_{\mathcal{H}_\infty} = 0$ on π_1 . On the other hand, for $i \in \mathcal{I}$,

$$\|H(\tau\Sigma^{-1/2}F^{-1}e_i)\|_{\mathcal{H}_\infty} \stackrel{(a)}{\geq} \tau\|F\Sigma^{-1/2}F^{-1}e_i\|_\infty \stackrel{(b)}{\geq} \frac{\tau}{2M},$$

where inequality (a) comes from $\|H(g)\|_{\mathcal{H}_\infty} \geq \|Fg\|_\infty$ for any g and inequality (b) comes from Proposition 5.4.1. Hence we have constructed two prior distributions with a separation of $c = \Omega(\tau/M)$ but a total variation distance less than 1/2. Theorem 5.3.1 now follows.

5.4.2 Proof of Active Lower Bound (Theorem 5.3.2)

For this setting we let $\pi_1 = \{0\}$ and $\pi_2 = \{\tau F^{-1}e_i\}_{i=1}^r$. The proof proceeds by bounding $d_{\text{kl}}(\mathbb{P}_{\pi_1}, \mathbb{P}_{\pi_2})$. To do this, we first bound $d_{\text{kl}}(\mathbb{P}_0, \mathbb{P}_i)$, where \mathbb{P}_0 is the joint distribution induced by the parameter $g = 0$ and \mathbb{P}_i is the joint distribution induced by the parameter $g = \tau F^{-1}e_i$. Proceeding as in the proof of Theorem 1.3 in Tu et al. [2017],

$$\begin{aligned} d_{\text{kl}}(\mathbb{P}_0, \mathbb{P}_i) &= \mathbb{E}_{\mathbb{P}_0} \left[\log \prod_{t=1}^N \frac{\gamma_t(u^{(t)} | \{u^{(k)}, y^{(k)}\}_{k=1}^{t-1}) p_0(y^{(t)} | u^{(t)})}{\gamma_t(u^{(t)} | \{u^{(k)}, y^{(k)}\}_{k=1}^{t-1}) p_i(y^{(t)} | u^{(t)})} \right] \\ &= \sum_{t=1}^N \mathbb{E}_{u^{(t)} \sim \mathbb{P}_0} [d_{\text{kl}}(\mathcal{N}(0, \sigma^2 I), \mathcal{N}(T(\tau F^{-1}e_i)u^{(t)}, \sigma^2 I))] \\ &= \frac{\tau^2}{2\sigma^2} \sum_{t=1}^N \mathbb{E}_{u^{(t)} \sim \mathbb{P}_0} [\|T(F^{-1}e_i)u^{(t)}\|_2^2]. \end{aligned}$$

A straightforward calculation shows that:

$$\sum_{i=1}^r T(F^{-1}e_i)^* T(F^{-1}e_i) = \text{diag} \left(1, \frac{r-1}{r}, \frac{r-2}{r}, \dots, \frac{1}{r} \right),$$

and hence the operator norm of this matrix is bounded by one. Therefore, by convexity of d_{kl} ,

$$\begin{aligned} d_{\text{kl}}(\mathbb{P}_{\pi_1}, \mathbb{P}_{\pi_2}) &\leq \frac{1}{r} \sum_{i=1}^r d_{\text{kl}}(\mathbb{P}_0, \mathbb{P}_i) = \frac{\tau^2}{2\sigma^2 r} \sum_{t=1}^N \mathbb{E}_{u^{(t)} \sim \mathbb{P}_0} \left[(u^{(t)})^* \left(\sum_{i=1}^r T(F^{-1}e_i)^* T(F^{-1}e_i) \right) u^{(t)} \right] \\ &\leq \frac{\tau^2 N M^2}{2\sigma^2 r} \left\| \sum_{i=1}^r T(F^{-1}e_i)^* T(F^{-1}e_i) \right\| \leq \frac{\tau^2 N M^2}{2\sigma^2 r}. \end{aligned}$$

Hence if we set $\tau = \frac{\sigma}{M} \sqrt{\frac{r}{N}}$, we have by Pinsker's inequality that $d_{\text{TV}}(\mathbb{P}_{\pi_1}, \mathbb{P}_{\pi_2}) \leq 1/2$. Finally, we note that $\|H(\tau F^{-1}e_i)\|_{\mathcal{H}_\infty} \geq \tau$ and conclude.

5.5 Experiments

We conduct experiments comparing a simple non-adaptive estimator based on least-squares (which we call the *plugin* estimator) to three active algorithms: two similar algorithms essentially based on the power method [Rojas et al., 2012; Wahlberg et al., 2010] and one based on weighted Thompson Sampling (WTS) [Müller et al., 2017]. Pseudocode for the plugin estimator is shown in Algorithm 1. For completeness, in Appendix C we describe the power method based algorithms in Algorithms 2 and 3, and the WTS algorithm in Algorithm 4. For brevity, we assume input normalization of $\|u^{(t)}\|_2 = 1$; for different SNR the algorithms are modified accordingly.

Algorithm 1: Plugin Estimator

Input: Normalized $\{u^{(t)}\}$.
for $t = 1$ to N **do**
 | Perform the experiment $y^{(t)} = Gu^{(t)} + \eta^{(t)}$.
end
Form \hat{G} from a least-squares fit of $\{y^{(t)}\}$ and $\{u^{(t)}\}$.
return $\hat{H} = \|\hat{G}\|_{\mathcal{H}_\infty}$.

We compare the performance of these four algorithms on a suite of random plants and random draws of noise. We note, however, that it is difficult to place these algorithms on even footing when making a comparison, especially in the presence of output noise. Reasons for this are:

- The parameters deemed “fixed” may be beneficial (or adversarial) to one algorithm or another.
- The amount of “side information” (e.g. noise covariance) an algorithm expects to receive may not be comparable across algorithms.

For an example of the first point, a large experiment budget is beneficial to the plugin and WTS estimators as they generally obtain better estimates with each new experiment while power method estimators hit a “noise floor” and stop improving.

The plants under test are of the form $G(z) = \sum_{k=0}^{r-1} z^{-k} \rho^k \eta_k$ where $\rho \in (0, 1]$ and $\eta_k \stackrel{\text{i.i.d.}}{\sim} \text{Unif}[-1, 1]$. For each suite of tests, we hold all other parameters fixed as shown in Table 5.1. To compare the aggregate performance across suites of random plants, we use *performance profiles*, a tool in the optimization community popularized by Dolan and Moré [2002]. Given a suite of methods $\{m_i\}$ and a metric $d(m_i, m_j)$ for per-instance performance, performance profiles show the percentage of instances where a particular method m is within τ of the best method. In our case, the metric will be the difference in relative error between the method’s estimate and the true \mathcal{H}_∞ -norm of the plant under consideration. As an example, PLUGIN(.05) would be the percentage of instances where the relative error of the

Table 5.1: Experiment Parameters

Type	Value
SNR $\frac{\ u\ _2}{\sigma}$	20 (high), 10 (low)
Experiment budget N	200
Plant length r	10
Input/output data length r'	50
Plant decay ρ	0.75 (decay), 1.0 (no decay)
Number of random plants	100
Noise instances per plant	10

plugin estimator is within 5 percentage points of the smallest relative error on that instance. Performance profiles are meant to show broad differences in algorithm performance and are robust to per-instance variation in the data, when interpreted correctly.

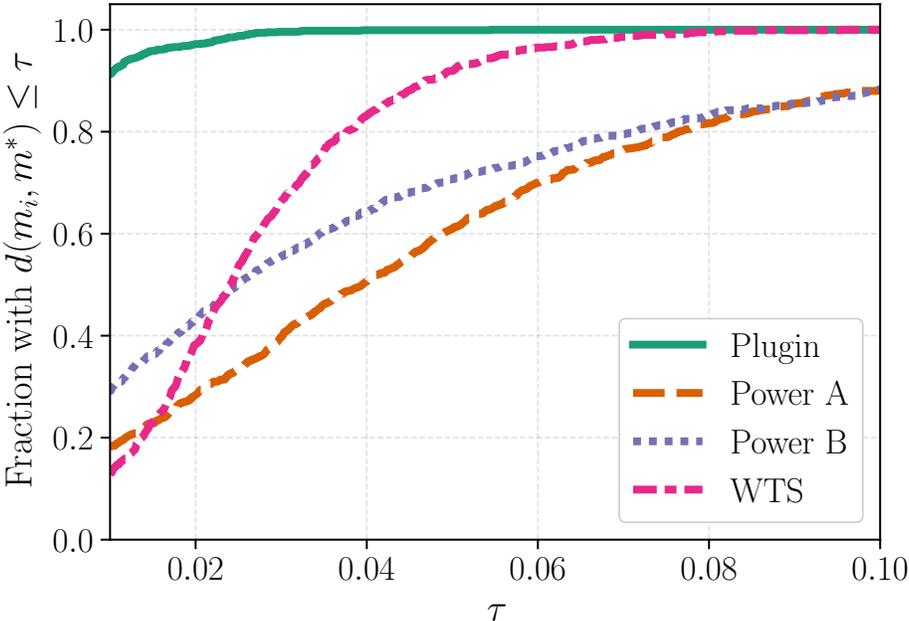
With this in mind, the performance profiles comparing the four algorithms² are shown in Figures 5.1 and 5.2. We see that for the plants with decaying impulse response coefficients, the plugin and WTS estimators are comparable. The latter estimator performs relatively worse in the experiments corresponding to no coefficient decay. As alluded to previously, this can most likely be attributed to our particular experimental setup: the WTS algorithm effectively grids the frequency response curve of the plant, and this set of plants allows for more relative variation in the curve than the ones with decaying coefficients.

5.6 Conclusion

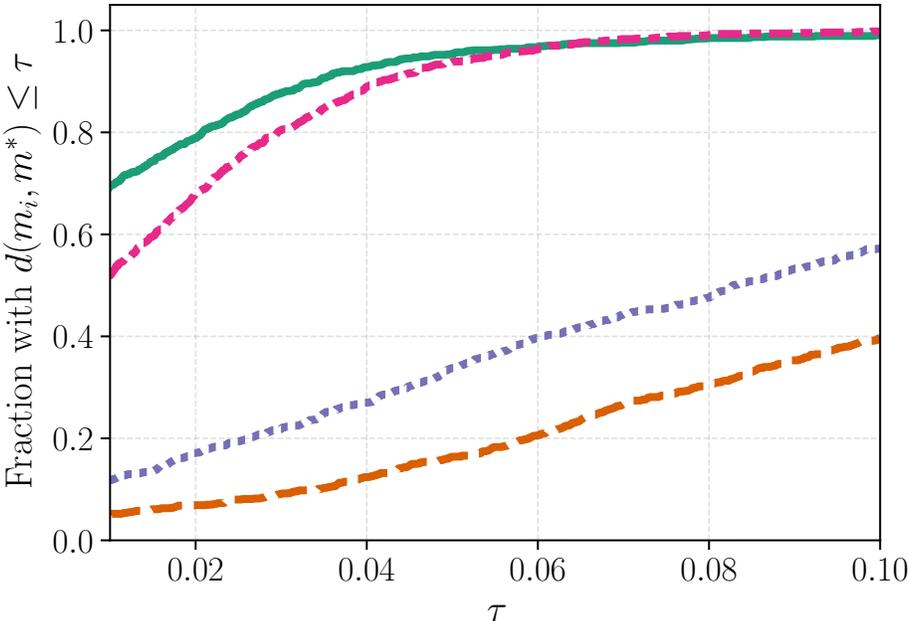
In this chapter, we provide lower bounds for \mathcal{H}_∞ -norm estimation for both passive and active algorithms. Our analysis shows that in the passive case model identification and \mathcal{H}_∞ -norm estimation have the same worst-case sample complexity. In the active setting, the lower bound improves by a factor that is logarithmic in the filter length. Experimentally, we see that the performance of a simple plugin \mathcal{H}_∞ -norm estimator is competitive with the proposed active algorithms for norm estimation in the literature.

Our work raises an interesting question as to whether there exists an active algorithm attaining the lower bound, or if instead the lower bound can be sharpened. In Section 5.3.2, we briefly discussed the technical hurdles that need to be overcome for both cases. Beyond resolving the gap between the lower bounds, an interesting question is how does the sample complexity of both model and norm estimation degrade when the filter length is unknown.

²Code producing these plots can be found at <https://github.com/rjboczar/hinf-lower-bounds-ACC>. The experiments were carried out using the PyWren framework [Jonas et al., 2017].

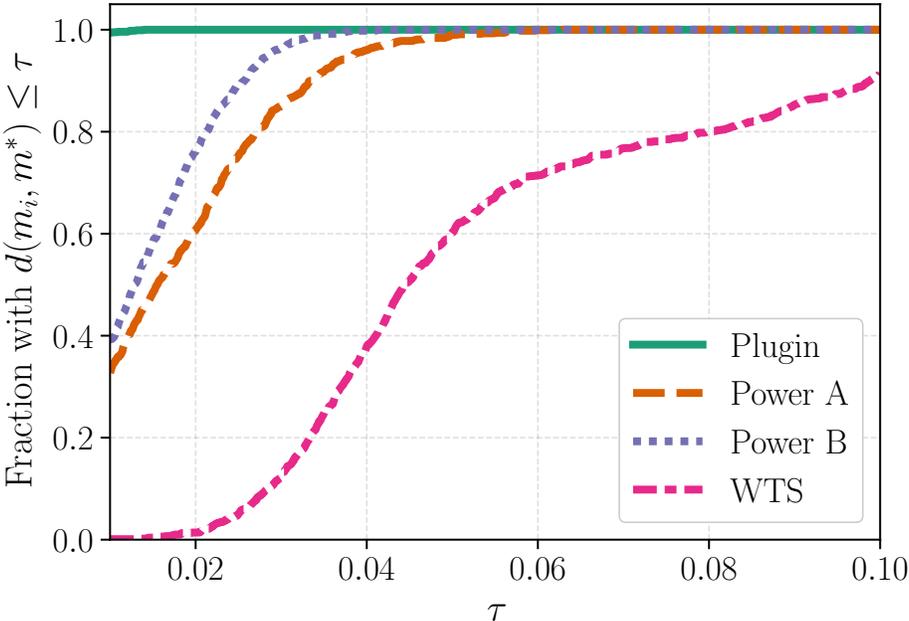


(a) High SNR, decay

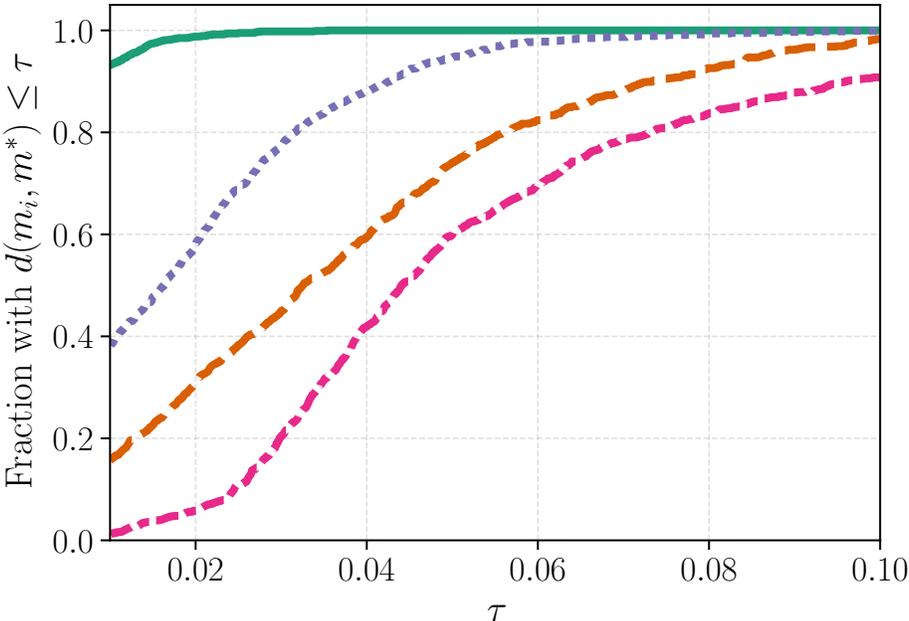


(b) Low SNR, decay

Figure 5.1: Performance profiles for the plugin, power method A, power method B, and weighted Thompson Sampling estimators (with coefficient decay).



(a) High SNR, no decay



(b) Low SNR, no decay

Figure 5.2: Performance profiles for the plugin, power method A, power method B, and weighted Thompson Sampling estimators (without coefficient decay).

Chapter 6

Finite-Data Performance Guarantees for the Output-Feedback Control of an Unknown System

6.1 Introduction

There have been many recent results that apply state-of-the-art machine learning techniques to the control of systems with continuous action spaces [Bansal et al., 2017; Dean et al., 2017; Duan et al., 2016; Fazel et al., 2018; Levine et al., 2016; Marco et al., 2017]. As the systems we control become ever more complex, be it in their dynamics, their scale, or their interaction with the environment, moving to a data-driven approach will be inevitable: in these settings, first-principle modeling becomes either impossible or intractable. However, as promising and exciting as recent empirical demonstrations of these techniques have been, they have, for the most part, lacked the rigorous stability, safety and robustness guarantees that the controls community has always prided itself in providing. Indeed, such guarantees are not only desirable, but *necessary* when such techniques are being proposed for the control of safety critical systems or infrastructures.

This chapter can be seen as a step towards providing such guarantees, albeit in a simplified setting, wherein we establish rigorous baselines of robustness and performance when controlling a single-input-single-output (SISO) system with an unknown transfer function. To do so, we combine contemporary approaches to system identification and robust control into what we term the “coarse-ID control” pipeline. In particular, we leverage the results developed in Chapter 4 to provide finite-sample guarantees on optimally (in a certain sense) estimating a stable single-input single-output linear time-invariant (SISO LTI) system, using input-output data pairs.¹ Such finite-data guarantees are not only in stark contrast to

¹We note that there have been recent results in the system identification literature (for example Campi and Weyer [2002] and Shah et al. [2012]) that also seek to provide non-asymptotic guarantees of model estimation quality.

classical system identification results, which typically only provide asymptotic guarantees of model fidelity (see Ljung [1999] for an overview), but also necessary for the principled integration of these techniques with robust control, as they allow us to quantify the amount of uncertainty that our controller must contend with. We then formulate a robust control problem using the recently developed *system-level synthesis* (SLS) procedure [Wang et al., 2019], which exploits a novel parameterization of stabilizing controllers for LTI systems that allows us to quantify performance degradation in terms of the amount of uncertainty affecting the system [Matni et al., 2017]. Again this is in contrast to classical methods from robust control that are only able to provide robust stability guarantees for a prescribed amount of uncertainty [Zhou et al., 1995].

Main contribution A feature of “coarse-ID control,” as described above, is that we can analyze the end-to-end performance of this pipeline in a non-asymptotic setting. Specifically, we show that the difference in cost between the optimal cost for the true system (an FIR SISO system of length r) and the realized cost induced by instead solving a robust SLS procedure for the approximate system is $O\left(\sqrt{\frac{\sigma^2 r}{m}}\right)$. Here, we assume that the approximate system was estimated using the “optimal” coarse-grained system identification procedure described in Chapter 4, with σ^2 the measurement noise variance and m the number of experiments conducted in order to construct an estimate of the system. Finally, this chapter should be viewed as a step towards generalizing the results in Dean et al. [2017], which provides finite-data end-to-end performance guarantees for the classical LQR optimal control problem, to the output-feedback setting.

In Section 3.2 we fix notation and quickly outline the structure used by common robust control problems. Section 6.3 then gives an overview of the system-level synthesis framework and how it can be used to solve these problems. Finally, in Section 6.4 we combine this framework with recent work on coarse-grained identification to provide quantitative bounds on how the performance of a robust controller synthesized using the SLS framework degrades when the plant to be controlled is only approximately identified. We conclude in Section 6.5 with computational examples.

6.2 Preliminaries

Notation Recall that we use boldface to denote frequency domain signals and transfer functions. Unless otherwise noted, in this chapter $\|\cdot\|$ represents the \mathcal{H}_∞ -norm (the induced $\ell_2 \rightarrow \ell_2$ norm) for elements in \mathcal{RH}_∞ (this reduces to the spectral norm for constant matrices).

6.2.1 The standard robust control problem

We first introduce a standard form for generic robust and optimal control problems, and then show how simple disturbance attenuation and reference tracking problems can be cast into

this standard form. We work with discrete-time LTI systems, but unless stated otherwise, all results extend naturally to the continuous-time setting. A system in standard form can be described by the following equations:

$$\begin{aligned} \mathbf{z} &= \mathbf{P}_{11}\mathbf{w} + \mathbf{P}_{12}\mathbf{u} \\ \mathbf{y} &= \mathbf{P}_{21}\mathbf{w} + \mathbf{P}_{22}\mathbf{u} \\ \mathbf{u} &= \mathbf{K}\mathbf{y}, \end{aligned} \tag{6.1}$$

where \mathbf{z} is the regulated output (e.g., deviations of the system state from a desired set-point), \mathbf{y} is the measured output available to the controller \mathbf{K} to compute the control action $\mathbf{u} = \mathbf{K}\mathbf{y}$, and \mathbf{w} is the exogenous disturbance. We further assume that the full plant \mathbf{P} admits a joint realization,² i.e.

$$\mathbf{P} = \begin{bmatrix} \mathbf{P}_{11} & \mathbf{P}_{12} \\ \mathbf{P}_{21} & \mathbf{P}_{22} \end{bmatrix} = \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & 0 \end{array} \right] \tag{6.2}$$

where $\mathbf{P}_{ij} = C_i(zI - A)^{-1}B_j + D_{ij}$.

The standard optimal control problem of minimizing the gain from exogenous disturbance \mathbf{w} to regulated output \mathbf{z} , subject to internal stability of the closed loop system shown in Figure 6.1, can then be posed as

$$\begin{aligned} &\underset{\mathbf{K}}{\text{minimize}} \quad \|\mathbf{P}_{11} + \mathbf{P}_{12}\mathbf{K}(I - \mathbf{P}_{22}\mathbf{K})^{-1}\mathbf{P}_{21}\| \\ &\text{subject to} \quad \mathbf{K}(I - \mathbf{P}_{22}\mathbf{K})^{-1} \in \mathcal{RH}_\infty. \end{aligned} \tag{6.3}$$

Disturbance rejection Consider the feedback system shown in Figure 6.2, wherein a controller \mathbf{K} is in feedback with a SISO plant \mathbf{G} , with input disturbance \mathbf{d} and measurement noise \mathbf{n} . We can then define the disturbances and outputs as

$$\begin{aligned} \mathbf{w} &= \begin{bmatrix} \mathbf{d} \\ \mathbf{n} \end{bmatrix} \\ \mathbf{z} &= \begin{bmatrix} \mathbf{v} \\ \rho\mathbf{u} \end{bmatrix} \\ \mathbf{y} &= \mathbf{v} + \mathbf{n}, \end{aligned}$$

²We assume throughout that \mathbf{P}_{22} is strictly proper—it follows that $\mathbf{K}(I - \mathbf{P}_{22}\mathbf{K})^{-1} \in \mathcal{RH}_\infty$ is a necessary and sufficient condition for internal stability of the closed loop system shown in Figure 6.1 [Zhou et al., 1995].

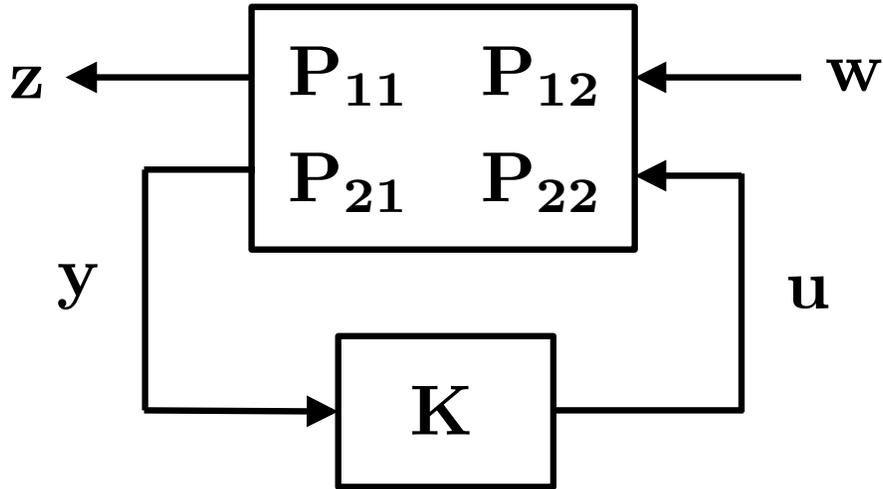


Figure 6.1: The standard optimal control problem (6.3) for the plant \mathbf{P} as defined in equations (6.1) and (6.2).

respectively, where $\rho > 0$. Furthermore, let the plant \mathbf{G} have a state-space realization (A, B, C) . We then have that

$$\begin{aligned} \mathbf{z} &= \begin{bmatrix} \mathbf{G} & 0 \\ 0 & 0 \end{bmatrix} \mathbf{w} + \begin{bmatrix} \mathbf{G} \\ \rho \end{bmatrix} \mathbf{u} \\ &:= \mathbf{P}_{11} \mathbf{w} + \mathbf{P}_{12} \mathbf{u} \\ \mathbf{y} &= \begin{bmatrix} \mathbf{G} & 1 \end{bmatrix} \mathbf{w} + \mathbf{G} \mathbf{u} \\ &:= \mathbf{P}_{21} \mathbf{w} + \mathbf{P}_{22} \mathbf{u}, \end{aligned}$$

from which it follows that the generalized plant \mathbf{P} admits the joint realization

$$\mathbf{P} = \left[\begin{array}{c|cc} A & [B & 0] & B \\ \hline C & \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix} & \begin{bmatrix} 0 \\ \rho \\ 0 \end{bmatrix} \end{array} \right].$$

Reference tracking Now, consider the feedback system shown in Figure 6.3, wherein a controller \mathbf{K} is in feedback with a SISO plant \mathbf{G} , with input disturbance \mathbf{d} and reference

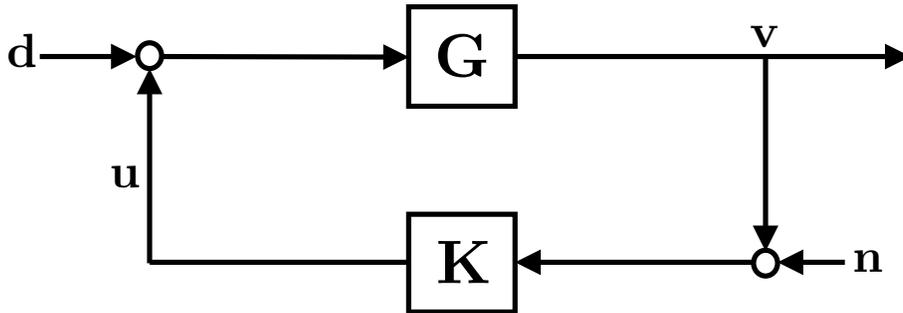


Figure 6.2: The disturbance rejection problem for a SISO plant \mathbf{G} with input disturbance \mathbf{d} and measurement noise \mathbf{n} .

signal \mathbf{r} . We can then define the disturbances and outputs as

$$\begin{aligned}\mathbf{w} &= \begin{bmatrix} \mathbf{d} \\ \mathbf{r} \end{bmatrix} \\ \mathbf{z} &= \begin{bmatrix} \mathbf{e} \\ \rho \mathbf{u} \end{bmatrix} \\ \mathbf{y} &= \mathbf{e},\end{aligned}$$

respectively, where $\rho > 0$. Furthermore, let the plant \mathbf{G} have a state-space realization (A, B, C) . We then have that

$$\begin{aligned}\mathbf{z} &= \begin{bmatrix} \mathbf{G} & -1 \\ 0 & 0 \end{bmatrix} \mathbf{w} + \begin{bmatrix} \mathbf{G} \\ \rho \end{bmatrix} \mathbf{u} \\ &:= \mathbf{P}_{11} \mathbf{w} + \mathbf{P}_{12} \mathbf{u} \\ \mathbf{y} &= [\mathbf{G} \quad -1] \mathbf{w} + \mathbf{G} \mathbf{u} \\ &:= \mathbf{P}_{21} \mathbf{w} + \mathbf{P}_{22} \mathbf{u},\end{aligned}$$

from which it follows that the full plant \mathbf{P} admits the joint realization

$$\mathbf{P} = \left[\begin{array}{c|cc} A & [B \ 0] & B \\ \hline [C] & \begin{bmatrix} 0 & -1 \\ 0 & 0 \\ 0 & -1 \end{bmatrix} & \begin{bmatrix} 0 \\ \rho \\ 0 \end{bmatrix} \end{array} \right].$$

Specialization to FIR plant \mathbf{G} Suppose that \mathbf{G} is strictly proper and has a finite impulse response (FIR) of order r , i.e., that $\mathbf{G} = \sum_{t=1}^{r-1} g_t z^{-t}$ for a collection of real scalars $\{g_t\}_{t=1}^{r-1}$. Defining $g = [g_1, g_2, \dots, g_{r-1}]^\top$, the plant \mathbf{G} admits the state-space realization (Z, e_1, g^\top) where Z is the right-shift operator (i.e., a matrix with ones on the sub-diagonal and zeros

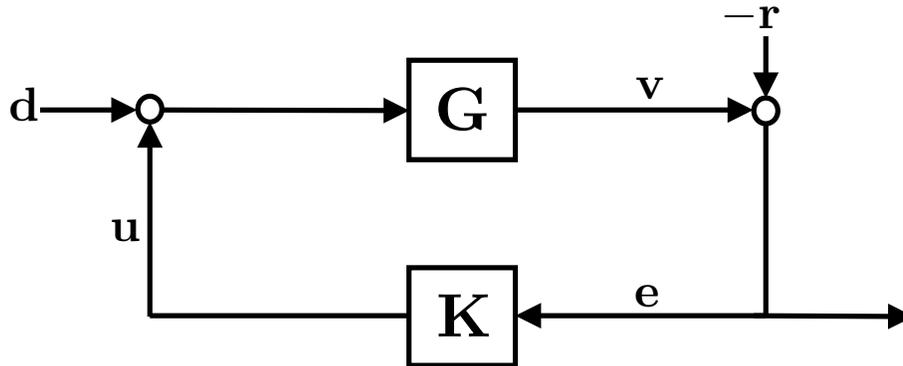


Figure 6.3: The reference tracking problem for a SISO plant \mathbf{G} with disturbance \mathbf{d} and reference \mathbf{r} .

elsewhere). Given the examples presented thus far, going forward we assume that

$$C_1 = [g \ 0]^\top, \quad C_2 = g^\top, \quad (6.4)$$

as well as the standard assumption that $D_{12}^\top C_1 = 0$. Additionally, given that we are considering SISO systems, we can without loss of generality (by suitably rescaling B_2) assume that

$$D_{12} = [0 \ 1]^\top. \quad (6.5)$$

6.2.2 Coarse-grained identification

As our aim is to provide end-to-end guarantees for robust control problems, we must first have a scheme to acquire an approximate plant model $\hat{\mathbf{G}}$. Recall our coarse-grained identification setup³ from Section 4.3:

- (i) carefully choose a series of m inputs $\{\mathbf{u}^{(i)}\}$, where $\mathbf{u}^{(i)} \in \mathcal{U}$, and collect noisy outputs $\{\mathbf{y}^{(i)}\}$ where $\mathbf{y}^{(i)} = \mathbf{G}\mathbf{u}^{(i)} + \xi^{(i)}$ with $\xi^{(i)} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2 I)$
- (ii) form a least-squares estimate $\hat{\mathbf{G}}$ of the impulse response of \mathbf{G} using $\{\mathbf{u}^{(i)}, \mathbf{y}^{(i)}\}$.

We refer to each such pair $(\mathbf{u}^{(i)}, \mathbf{y}^{(i)})$ as an experiment.

However, we need to make slight modifications to the results in Chapter 4. We will need ℓ_2 bounds on the impulse response error, as these are more natural for our problem, and the bounds in Chapter 4 are in term of the \mathcal{H}_∞ -norm. However, while they can conservatively be plugged in verbatim (as $\|g\|_2 \leq \|\mathbf{G}\|_{\mathcal{H}_\infty}$), we will instead modify their proofs slightly to fit our application.

³We use boldface notation for signals to reinforce the frequency-domain aspects in the sequel.

6.3 System-Level Synthesis

The System-Level Synthesis (SLS) framework, proposed by Wang et al. [2019], provides a parameterization of stabilizing controllers that achieve specified responses between disturbances and outputs. We briefly review here the SLS framework, and later show in Section 6.4.1 how it can be modified to solve a robust optimal control problem subject to bounded uncertainty on the FIR coefficients g .

For an LTI system with dynamics described by (6.2), we define a system response $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ to be the maps satisfying

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \begin{bmatrix} \delta_x \\ \delta_y \end{bmatrix}. \quad (6.6)$$

where $\delta_x := B_1 \mathbf{w}$ is the process noise, and $\delta_y := D_{21} \mathbf{w}$ is the measurement noise.

We call a system response $\Theta = \{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ *stable and achievable* with respect to a plant \mathbf{P} if there exists an internally stabilizing controller \mathbf{K} such that the control rule $\mathbf{u} = \mathbf{K}\mathbf{y}$ leads to closed loop behavior consistent with (6.6). It was shown in [Wang et al., 2019] that the parameterization of all stable and achievable system responses $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ is defined by the following affine space:

$$[zI - A \quad -B_2] \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} = [I \quad 0] \quad (6.7a)$$

$$\begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \begin{bmatrix} zI - A \\ -C_2 \end{bmatrix} = \begin{bmatrix} I \\ 0 \end{bmatrix} \quad (6.7b)$$

$$\mathbf{R}, \mathbf{M}, \mathbf{N} \in \frac{1}{z} \mathcal{RH}_\infty, \quad \mathbf{L} \in \mathcal{RH}_\infty. \quad (6.7c)$$

We call equations (6.7a) - (6.7c) the *SLS constraints*. The parameterization of all internally stabilizing controllers is given by the following theorem.

Theorem 6.3.1 ([Wang et al., 2019, Thm. 2]). *Suppose that a system response $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ satisfies the SLS constraints (6.7a) - (6.7c). Then, $\mathbf{K} = \mathbf{L} - \mathbf{M}\mathbf{R}^{-1}\mathbf{N}$ is an internally stabilizing controller for the plant (6.2) that yields the desired system response (6.6). Furthermore, the solutions of (6.7a) - (6.7c) with the implementation $\mathbf{K} = \mathbf{L} - \mathbf{M}\mathbf{R}^{-1}\mathbf{N}$ parameterize all internally stabilizing controllers for the plant (6.2).*

Using this parameterization, we can recast the standard optimal control problem (6.3) as the SLS problem

$$\boxed{\begin{array}{l} \text{minimize } J(G, \Theta) := \left\| [C_1 \quad D_{12}] \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix} + D_{11} \right\| \\ \text{subject to } (6.7a) - (6.7c). \end{array}} \quad (6.8)$$

In the FIR case, we use the abbreviated notation $J(g, \Theta)$ for the case where G is the plant (Z, e_1, g^\top) .

Remark 6. Although we focus on the \mathcal{H}_∞ optimal control problem posed in equation (6.8), the results that follow carry over naturally to \mathcal{H}_2 (LQG) and \mathcal{L}_1 optimal control problems as well.

6.4 Sample Complexity Bounds

We now provide finite-data performance guarantees for a controller synthesized using the system identification and robust synthesis procedures described in the previous sections. Prior to stating our main results, we recall the problem set up and Coarse-ID Control pipeline.

We consider the identification and control of the system (Z, e_1, g^\top) , which is assumed to be FIR of order r . We begin with the simplified setting that the order r of the true system is known, and we use the Coarse-grained identification procedure described in Section 6.2.2 to identify an approximate system (Z, e_1, \tilde{g}^\top) , also of order r , using a series of m experiments. We then use this approximate system (Z, e_1, \tilde{g}^\top) , as well as high-probability bounds on the estimation error $\|g - \tilde{g}\|_2$, in a robust SLS problem (see (6.13) in Section 6.4.1) to compute a controller with provable suboptimality guarantees, as formalized in the following theorem.⁴

Theorem 6.4.1. Let Θ_0 be the optimal solution of the SLS problem (6.8) for the plant (Z, e_1, g^\top) , and let \tilde{g} be an estimate of g obtained using coarse-grained identification (σ^2 -variance output noise only) with m experiments, where $\|\mathbf{u}^{(i)}\|_p \leq 1 \forall i$. Let $(\tilde{\Theta}_*, \alpha_*)$ be the optimal solution to the robust SLS problem (6.13) for (Z, e_1, \tilde{g}^\top) , and let $\hat{\Theta}_*$ be the response achieved on the true system g by the synthesized controller $\tilde{\mathbf{K}}_* = \tilde{\mathbf{L}}_* - \tilde{\mathbf{M}}_* \tilde{\mathbf{R}}_*^{-1} \tilde{\mathbf{N}}_*$. Then, if $m \gtrsim \sigma^2 r \|\mathbf{N}_0\|^2 \log(\eta^{-1})^{\frac{1}{2}}$, with probability at least $1 - \eta$, the controller $\tilde{\mathbf{K}}_*$ stabilizes the true system (Z, e_1, g^\top) and has a suboptimality gap bounded by

$$\begin{aligned} & J(g, \hat{\Theta}_*) - J(g, \Theta_0) \\ & \leq 8 \sqrt{\log 2 \frac{\sigma^2 r^{2/\max(p,2)}}{m}} \left(1 + \sqrt{2 \log \eta^{-1}}\right) \left\| \begin{bmatrix} 1 + g^\top \mathbf{N}_0 \\ \mathbf{L}_0 \end{bmatrix} \right\| \|\mathbf{R}_0 B_1 + \mathbf{N}_0 D_{21}\|. \end{aligned}$$

Corollary 6.4.1. Assume that we are in the setting of Theorem 6.4.1, and further let there be process noise with variance σ_w^2 that enters the system via the same channel as the control input (i.e., $B_1 = B_2$) and measurement noise with variance σ_ξ^2 . Then, Theorem 6.4.1 holds with $\sigma^2 \leftarrow \sigma_w^2 \|\mathbf{G}\|^2 + \sigma_\xi^2$.

We can further generalize these results to the setting where the order r of the underlying system is not known, and that the true system is approximated by a length- \tilde{r} FIR filter with coefficients \tilde{g} where $\tilde{r} < r$. In this case, applying the triangle equality

$$\|\delta\|_2 \leq \|g_{0:\tilde{r}-1} - \tilde{g}\|_2 + \|\tilde{g}_{\tilde{r}:r-1}\|_2$$

⁴Here, \gtrsim hides universal constants; see Lemma 6.4.3 for an explicit characterization.

gives a similar sample complexity bound, albeit one where the cost difference does not tend to zero as the number of experiments m tends to infinity.

Corollary 6.4.2. *Assume that we are in the setting of Theorem 6.4.1, except let \tilde{g} be a length- \tilde{r} (where $\tilde{r} < r$) FIR estimate of g obtained using the prescribed coarse-grained identification. Furthermore, assume that $\|g_{\tilde{r}-1:r-1}\|_2 < (4\mathbf{N}_0)^{-1}$. Then, if*

$$m \gtrsim \sigma^2 r \left(\frac{\|\mathbf{N}_0\|}{1 - 4\|\mathbf{N}_0\|\|g_{\tilde{r}-1:r-1}\|_2} \right)^2 \sqrt{\log(\eta^{-1})},$$

with probability at least $1 - \eta$ the controller $\tilde{\mathbf{K}}_* = \tilde{\mathbf{L}}_* - \tilde{\mathbf{M}}_* \tilde{\mathbf{R}}_*^{-1} \tilde{\mathbf{N}}_*$ stabilizes the true system (Z, e_1, g^\top) and has a suboptimal cost bounded by

$$\begin{aligned} & J(g, \hat{\Theta}_*) - J(g, \Theta_0) \\ & \leq \left(8 \sqrt{\log 2} \frac{\sigma^2 r^{2/\max(p,2)}}{m} \left(1 + \sqrt{2 \log \eta^{-1}} \right) \left\| \begin{bmatrix} 1 + g^\top \mathbf{N}_0 \\ \mathbf{L}_0 \end{bmatrix} \right\| + 4 \|g_{\tilde{r}-1:r-1}\|_2 \right) \\ & \quad \times \|\mathbf{R}_0 B_1 + \mathbf{N}_0 D_{21}\|. \end{aligned}$$

To prove the above results, we first derive a robust variant of the SLS framework presented in Section 6.3, and then show how it can be used to pose a robust synthesis problem that admits suboptimality guarantees. In particular, these guarantees characterize the degradation in performance of the synthesized controller as a function of the size of the uncertainty on the transfer function coefficients g . We then combine this characterization of performance degradation with high-probability bounds on the estimation error produced by the coarse-grained identification procedure to provide an end-to-end analysis of the Coarse-ID control procedure.

6.4.1 Robust SLS

As we only have access to approximately identified plants, we need a robust variant of Theorem 6.3.1. First, we introduce a robust version of (6.7b),

$$\begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \\ \tilde{\mathbf{M}} & \tilde{\mathbf{L}} \end{bmatrix} \begin{bmatrix} zI - A \\ -C_2 \end{bmatrix} = \begin{bmatrix} I + \Delta_1 \\ \Delta_2 \end{bmatrix}. \quad (6.9)$$

We call equations (6.7a), (6.9), and (6.7c) the *robust SLS constraints*. We now have the ingredients needed to connect the main and robust SLS constraints. The proof is mostly algebraic and is thus deferred to the Appendix.

Lemma 6.4.1 (Robust Equivalence). *Consider system responses $\tilde{\Theta} = \{\tilde{\mathbf{R}}, \tilde{\mathbf{M}}, \tilde{\mathbf{N}}, \tilde{\mathbf{L}}\}$ and $\hat{\Theta} = \{\hat{\mathbf{R}}, \hat{\mathbf{M}}, \hat{\mathbf{N}}, \hat{\mathbf{L}}\}$, where the latter is given by*

$$\begin{aligned}\hat{\mathbf{R}} &= (I + \mathbf{\Delta}_1)^{-1} \tilde{\mathbf{R}} \\ \hat{\mathbf{M}} &= \tilde{\mathbf{M}} - \mathbf{\Delta}_2 (I + \mathbf{\Delta}_1)^{-1} \tilde{\mathbf{R}} \\ \hat{\mathbf{N}} &= (I + \mathbf{\Delta}_1)^{-1} \tilde{\mathbf{N}} \\ \hat{\mathbf{L}} &= \tilde{\mathbf{L}} - \mathbf{\Delta}_2 (I + \mathbf{\Delta}_1)^{-1} \tilde{\mathbf{N}},\end{aligned}\tag{6.10}$$

where by assumption $(I + \mathbf{\Delta}_1)^{-1}$ exists and is in \mathcal{RH}_∞ . Let $\mathbf{G} = (A, B, C, D)$ be a given plant, and consider the following statements.

- (i) $\tilde{\Theta}$ satisfies the robust SLS constraints for \mathbf{G} .
- (ii) $\hat{\Theta}$ satisfies the SLS constraints for \mathbf{G} .

Under the assumptions, (i) \implies (ii). Furthermore, let $\mathbf{G}' = (A, B, C', D)$, and let

$$\begin{aligned}\mathbf{\Delta}_1 &= -\tilde{\mathbf{N}}(C - C') \\ \mathbf{\Delta}_2 &= -\tilde{\mathbf{L}}(C - C').\end{aligned}$$

Then, (i) is equivalent to a third statement (iii): $\tilde{\Theta}$ satisfies the SLS constraints for \mathbf{G}' .

A chain of corollaries follow from Lemma 6.4.1 that will be useful in quantifying the performance achieved on the true system of a controller designed using an approximate system model. Unless otherwise noted, let $\tilde{\Theta}$, $\hat{\Theta}$ be defined as in Lemma 6.4.1.

Corollary 6.4.3. *Suppose that $\tilde{\Theta}$ satisfies the robust SLS constraints for the system (6.2). Then, the controller $\tilde{\mathbf{K}} = \tilde{\mathbf{L}} - \tilde{\mathbf{M}}\tilde{\mathbf{R}}^{-1}\tilde{\mathbf{N}}$ stabilizes the system (6.2) and achieves the closed-loop system response $\hat{\Theta}$ if and only if $(I + \mathbf{\Delta}_1)^{-1} \in \mathcal{RH}_\infty$.*

Proof. First note that the robust SLS constraints imply that $\mathbf{\Delta}_2 \in \mathcal{RH}_\infty$. Next, assume $(I + \mathbf{\Delta}_1)^{-1} \in \mathcal{RH}_\infty$. By Lemma 6.4.1, $\hat{\Theta}$ satisfies the SLS constraints for (6.2). Thus, by Theorem 6.3.1, $\hat{\mathbf{K}} = \hat{\mathbf{L}} - \hat{\mathbf{M}}\hat{\mathbf{R}}^{-1}\hat{\mathbf{N}}$ is stabilizing and achieves the closed-loop response $\hat{\Theta}$. Moreover, $\hat{\mathbf{K}}$ is precisely equal to $\tilde{\mathbf{K}}$.

Conversely, assume $(I + \mathbf{\Delta}_1)^{-1}$ exists but is not in \mathcal{RH}_∞ (if it does not exist the system response (6.10) is obviously not well-defined). It then follows that $\hat{\mathbf{R}} = (I + \mathbf{\Delta}_1)^{-1}\tilde{\mathbf{R}}$ is not in \mathcal{RH}_∞ as $\tilde{\mathbf{R}}$ is square and invertible. \blacksquare

This immediately gives us a sufficient condition for robustness of the SLS procedure.

Corollary 6.4.4. *Suppose that $\tilde{\Theta}$ satisfies the robust SLS constraints for the system (6.2). A sufficient condition for the controller $\tilde{\mathbf{K}} = \tilde{\mathbf{L}} - \tilde{\mathbf{M}}\tilde{\mathbf{R}}^{-1}\tilde{\mathbf{N}}$ to stabilize the system (6.2) and achieve closed-loop response $\hat{\Theta}$ is that $\|\mathbf{\Delta}_1\| < 1$, for any induced norm $\|\cdot\|$.*

Proof. It follows from the small-gain theorem [Zhou et al., 1995] that $\|\Delta_1\| < 1$ implies $(I + \Delta_1)^{-1} \in \mathcal{RH}_\infty$, and thus Corollary 6.4.3 applies. ■

We now specialize our results to the case where the plant \mathbf{P} , as defined in (6.2), is FIR. In this case, the modeling error arises only in the coefficient vector g defining the impulse response. To that end, we define the estimated plant $\tilde{\mathbf{G}}$ with the realization (Z, e_1, \tilde{g}^\top) and note that the resulting error arises only in the C_1 and C_2 terms of the corresponding estimated plant $\tilde{\mathbf{P}}$, where these state-space parameters are defined as in (6.4). To that end, we define the estimation error vector $\delta := g - \tilde{g}$, allowing us to further specialize Corollary 6.4.4.

Corollary 6.4.5. *Suppose $\tilde{\Theta}$ satisfies the SLS constraints for the estimated system (Z, e_1, \tilde{g}^\top) . If $\|\tilde{\mathbf{N}}\delta^\top\| < 1$ for any induced norm $\|\cdot\|$, then the controller $\tilde{\mathbf{K}} = \tilde{\mathbf{L}} - \tilde{\mathbf{M}}\tilde{\mathbf{R}}^{-1}\tilde{\mathbf{N}}$ stabilizes the true system (Z, e_1, g^\top) and achieves the closed-loop response $\tilde{\Theta}$ as specified in (6.10). Additionally, if the induced norm $\|\cdot\|$ is either the \mathcal{H}_∞ or \mathcal{L}_1 norm, the response $\tilde{\Theta}$ simplifies to*

$$\begin{bmatrix} \hat{\mathbf{R}} & \hat{\mathbf{N}} \\ \hat{\mathbf{M}} & \hat{\mathbf{L}} \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \\ \tilde{\mathbf{M}} & \tilde{\mathbf{L}} \end{bmatrix} + \frac{1}{1 - \delta^\top \tilde{\mathbf{N}}} \begin{bmatrix} \tilde{\mathbf{N}}\delta^\top \\ \tilde{\mathbf{L}}\delta^\top \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \end{bmatrix}. \quad (6.11)$$

Proof. By Lemma 6.4.1, $\tilde{\Theta}$ satisfies the robust SLS constraints for (Z, e_1, g^\top) . The sufficient condition then follows by applying Corollary 6.4.4. Furthermore, if the induced norm $\|\cdot\|$ used in Corollary 6.4.3 and Corollary 6.4.4 is either the \mathcal{H}_∞ or \mathcal{L}_1 norm, it follows from Hölder's inequality that $\|\tilde{\mathbf{N}}\delta^\top\| < 1$ implies that $|\delta^\top \tilde{\mathbf{N}}| < 1$ on \mathbb{D}^c . Hence, we can use the Sherman-Morrison identity,

$$(I - xy^\top)^{-1} = I + \frac{xy^\top}{1 - y^\top x},$$

to simplify the response (6.10) achieved by the approximate controller $\tilde{\mathbf{K}} = \tilde{\mathbf{L}} - \tilde{\mathbf{M}}\tilde{\mathbf{R}}^{-1}\tilde{\mathbf{N}}$ to the expression (6.11). ■

We now use this robust parameterization to formulate a robust SLS problem that yields a controller with stability and performance guarantees.

We will use these two facts without fanfare in the following sections.

Proposition 6.4.1. *Let $x \in \mathbb{R}^n$, $y \in \mathcal{RH}_\infty^n$. Then $\|x^\top y\| \leq \|x\|_2 \|y\|$.*

Proposition 6.4.2. *$\|(I + \mathbf{A})^{-1}\| \leq (1 - \|\mathbf{A}\|)^{-1}$ for all $\|\mathbf{A}\| < 1$.*

Now, define $J(g, \Theta)$ to be the performance (i.e. the objective in (6.8)) of the controller $\mathbf{K} = \mathbf{L} - \mathbf{M}\mathbf{R}^{-1}\mathbf{N}$ induced by $\Theta = \{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ when placed in closed-loop with the FIR plant G specified by impulse response coefficients g . Now, assume we design a response $\tilde{\Theta}$, with corresponding controller $\tilde{\mathbf{K}}$, that satisfies the SLS constraints specified by the estimate system \tilde{g} . We saw in the previous section that under suitable conditions, the response on the true system g is given by $\hat{\Theta}$, as specified in Corollary 6.4.5. By the triangle inequality,

Corollary 6.4.5, and our parametric assumption (6.4)-(6.5), we can then bound the difference between expectation $J(\tilde{g}, \tilde{\Theta})$ and reality $J(g, \hat{\Theta})$ as follows:

$$\begin{aligned}
J(g, \hat{\Theta}) &= \left\| [C_1 \ D_{12}] \begin{bmatrix} \hat{\mathbf{R}} & \hat{\mathbf{N}} \\ \hat{\mathbf{M}} & \hat{\mathbf{L}} \end{bmatrix} \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix} + D_{11} \right\| \\
&= \left\| \begin{bmatrix} g^\top & 0 \\ 0 & 1 \end{bmatrix} \left(\begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \\ \tilde{\mathbf{M}} & \tilde{\mathbf{L}} \end{bmatrix} + \begin{bmatrix} \frac{\tilde{\mathbf{N}}\delta^\top}{1-\delta^\top\tilde{\mathbf{N}}} \\ \frac{\tilde{\mathbf{L}}\delta^\top}{1-\delta^\top\tilde{\mathbf{N}}} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \end{bmatrix} \right) \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix} + D_{11} \right\| \\
&= \left\| \left(\left(\begin{bmatrix} \tilde{g}^\top & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} \delta^\top & 0 \\ 0 & 0 \end{bmatrix} \right) \begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \\ \tilde{\mathbf{M}} & \tilde{\mathbf{L}} \end{bmatrix} + \begin{bmatrix} g^\top & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \frac{\tilde{\mathbf{N}}\delta^\top}{1-\delta^\top\tilde{\mathbf{N}}} \\ \frac{\tilde{\mathbf{L}}\delta^\top}{1-\delta^\top\tilde{\mathbf{N}}} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \end{bmatrix} \right) \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix} + D_{11} \right\| \\
&\leq J(\tilde{g}, \tilde{\Theta}) + \left\| \left(\begin{bmatrix} \delta^\top \\ 0 \end{bmatrix} + \begin{bmatrix} g^\top & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \frac{\tilde{\mathbf{N}}\delta^\top}{1-\delta^\top\tilde{\mathbf{N}}} \\ \frac{\tilde{\mathbf{L}}\delta^\top}{1-\delta^\top\tilde{\mathbf{N}}} \end{bmatrix} \right) \begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \\ \tilde{\mathbf{M}} & \tilde{\mathbf{L}} \end{bmatrix} \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix} \right\| \\
&\leq J(\tilde{g}, \tilde{\Theta}) + \frac{1}{1 - \|\delta^\top\tilde{\mathbf{N}}\|} \left\| \begin{bmatrix} \delta^\top + \tilde{g}^\top\tilde{\mathbf{N}}\delta^\top \\ \tilde{\mathbf{L}}\delta^\top \end{bmatrix} \right\| \|\tilde{\mathbf{R}}B_1 + \tilde{\mathbf{N}}D_{21}\| \\
&\leq J(\tilde{g}, \tilde{\Theta}) + \frac{\|\delta\|}{1 - \|\delta^\top\tilde{\mathbf{N}}\|} \left\| \begin{bmatrix} 1 + \tilde{g}^\top\tilde{\mathbf{N}} \\ \tilde{\mathbf{L}} \end{bmatrix} \right\| \|\tilde{\mathbf{R}}B_1 + \tilde{\mathbf{N}}D_{21}\|,
\end{aligned}$$

where we assume $\|\tilde{\mathbf{N}}\delta^\top\| < 1$ for the bound to be valid.

For any estimated response \tilde{g} satisfying $\|\delta\| \leq \epsilon$, it then follows that

$$J(g, \hat{\Theta}) \leq J(\tilde{g}, \tilde{\Theta}) + \epsilon\alpha \|\tilde{\mathbf{R}}B_1 + \tilde{\mathbf{N}}D_{21}\| \quad (6.12)$$

for any α satisfying

$$\epsilon\alpha \|\tilde{\mathbf{N}}\| + \left\| \begin{bmatrix} 1 + \tilde{g}^\top\tilde{\mathbf{N}} \\ \tilde{\mathbf{L}} \end{bmatrix} \right\| \leq \alpha,$$

noting that $\epsilon\|\tilde{\mathbf{N}}\| < 1$, which implies $\|\tilde{\mathbf{N}}\delta^\top\| < 1$, is equivalent to $\alpha > 0$. We denote the right-hand side of this bound as

$$Q(\tilde{g}, \tilde{\Theta}, \alpha) := J(\tilde{g}, \tilde{\Theta}) + \epsilon\alpha \|\tilde{\mathbf{R}}B_1 + \tilde{\mathbf{N}}D_{21}\|.$$

The bound (6.12) then suggests the following robust controller synthesis procedure, which balances between solving for the optimal controller for the approximate system \tilde{g} and controlling a perturbative term. We call this problem the *robust SLS problem for \tilde{g}* .

$$\begin{array}{l}
\text{minimize } Q(\tilde{g}, \tilde{\Theta}, \alpha) \\
\{\tilde{\mathbf{R}}, \tilde{\mathbf{M}}, \tilde{\mathbf{N}}, \tilde{\mathbf{L}}\} \\
\alpha > 0 \\
\text{subject to } \tilde{\Theta} \text{ satisfies SLS constraints} \\
(6.7a) - (6.7c) \text{ for } (Z, e_1, \tilde{g}^\top) \\
\epsilon\alpha \|\tilde{\mathbf{N}}\| + \left\| \begin{bmatrix} 1 + \tilde{g}^\top\tilde{\mathbf{N}} \\ \tilde{\mathbf{L}} \end{bmatrix} \right\| \leq \alpha.
\end{array} \quad (6.13)$$

Although this problem is not jointly convex in α and the system responses $\tilde{\Theta}$, one can use a golden section search on α in practice. Moreover, the sum of norms can be split into two norm constraints using an epigraph formulation (see [Boyd and Vandenberghe \[2004, Chapter 3\]](#)).

6.4.2 Sub-optimality guarantees for robust SLS

We now show a bound on the change in the optimal control cost when the controller is synthesized using the robust SLS problem (6.13).

Proposition 6.4.3. *Let Θ_0 and $(\tilde{\Theta}_*, \alpha_*)$, as well as $\hat{\Theta}_*$, be defined as in Theorem (6.4.1), and let $\|\delta\| \leq \epsilon$. If $\epsilon < (2\|\mathbf{N}_0\|)^{-1}$, we have that*

$$J(g, \hat{\Theta}_*) - J(g, \Theta_0) \leq \frac{2\epsilon}{1 - 2\epsilon\|\mathbf{N}_0\|} \left\| \begin{bmatrix} 1 + g^\top \mathbf{N}_0 \\ \mathbf{L}_0 \end{bmatrix} \right\| \|\mathbf{R}_0 B_1 + \mathbf{N}_0 D_{21}\|. \quad (6.14)$$

To prove this proposition, we require a technical lemma that ensures that the true controller \mathbf{K}_0 stabilizes the estimate system specified by the FIR coefficients \tilde{g} , i.e. that the optimal system response Θ_0 can be used to construct a feasible solution to the approximate SLS synthesis problem (6.13).

Lemma 6.4.2. *Let Θ_0 and its induced controller \mathbf{K}_0 be as defined in Theorem 6.4.1, and let $\|\delta\| \leq \epsilon$, with $\epsilon < (2\|\mathbf{N}_0\|)^{-1}$. Then*

$$\alpha_0 := \frac{1}{1 - 2\epsilon\|\mathbf{N}_0\|} \left\| \begin{bmatrix} 1 + g^\top \mathbf{N}_0 \\ \mathbf{L}_0 \end{bmatrix} \right\|$$

is strictly positive, and the controller \mathbf{K}_0 is stabilizing for the estimate system specified by (Z, e_1, \tilde{g}^\top) and achieves the system response $\hat{\Theta}_0$ defined by

$$\begin{aligned} \hat{\mathbf{R}}_0 &= (I + \mathbf{N}_0 \delta^\top)^{-1} \mathbf{R}_0 \\ \hat{\mathbf{M}}_0 &= \mathbf{M}_0 - \mathbf{L}_0 \delta^\top (I + \mathbf{N}_0 \delta^\top)^{-1} \mathbf{R}_0 \\ \hat{\mathbf{N}}_0 &= (I + \mathbf{N}_0 \delta^\top)^{-1} \mathbf{N}_0 \\ \hat{\mathbf{L}}_0 &= \mathbf{L}_0 - \mathbf{L}_0 \delta^\top (I + \mathbf{N}_0 \delta^\top)^{-1} \mathbf{N}_0. \end{aligned}$$

Furthermore, $(\hat{\Theta}_0, \alpha_0)$ are feasible solutions to the approximate SLS synthesis problem (6.13).

Proof. Both of these points are conditional on $(I + \mathbf{N}_0 \delta^\top)^{-1}$ existing in \mathcal{RH}_∞ :

- By assumption, Θ_0 satisfies the SLS constraints for (Z, e_1, g^\top) ; by Lemma 6.4.1, it equivalently satisfies the robust SLS constraints for (Z, e_1, \tilde{g}^\top) , where $\Delta_1 = \mathbf{N}_0 \delta^\top$ and $\Delta_2 = \mathbf{L}_0 \delta^\top$ (note the switched roles of g and \tilde{g}). By Corollary 6.4.3, \mathbf{K}_0 then stabilizes (Z, e_1, \tilde{g}^\top) and achieves the system response $\hat{\Theta}_0$.

- By Lemma 6.4.1, $\widehat{\Theta}_0$ satisfies the SLS constraints for (Z, e_1, \tilde{g}^\top) , and is thus part of a feasible point for the approximate SLS synthesis problem (6.13). Now, we need to check that the corresponding α_0 is also part of a feasible solution. Toward that end, by the Sherman-Morrison identity, we see that

$$\|\epsilon \widehat{\mathbf{N}}_0\| = \left\| \epsilon \left(I - \frac{\mathbf{N}_0 \delta^\top}{1 + \delta^\top \mathbf{N}_0} \right) \mathbf{N}_0 \right\| \leq \epsilon \|\mathbf{N}_0\| + \frac{(\epsilon \|\mathbf{N}_0\|)^2}{1 - \epsilon \|\mathbf{N}_0\|} = \frac{\epsilon \|\mathbf{N}_0\|}{1 - \epsilon \|\mathbf{N}_0\|}.$$

Furthermore,

$$\left\| \begin{bmatrix} 1 + \tilde{g}^\top \widehat{\mathbf{N}}_0 \\ \widehat{\mathbf{L}}_0 \end{bmatrix} \right\| = \left\| \begin{bmatrix} \frac{1 + g^\top \mathbf{N}_0}{1 + \delta^\top \mathbf{N}_0} \\ \frac{\mathbf{L}_0}{1 + \delta^\top \mathbf{N}_0} \end{bmatrix} \right\| \leq \frac{1}{1 - \epsilon \|\mathbf{N}_0\|} \left\| \begin{bmatrix} 1 + g^\top \mathbf{N}_0 \\ \mathbf{L}_0 \end{bmatrix} \right\|.$$

Therefore,

$$\frac{1}{1 - \|\epsilon \widehat{\mathbf{N}}_0\|} \left\| \begin{bmatrix} 1 + \tilde{g}^\top \widehat{\mathbf{N}}_0 \\ \widehat{\mathbf{L}}_0 \end{bmatrix} \right\| \leq \alpha_0,$$

the final feasibility condition of (6.13).

It therefore remains to verify that $(I + \mathbf{N}_0 \delta^\top)^{-1}$ exists and is in \mathcal{RH}_∞ . As we have seen, a sufficient condition is $\|\mathbf{N}_0 \delta^\top\| < 1$, and this condition is implied by the assumption that $\alpha_0 > 0$. ■

Proof. [Proof of Proposition 6.4.3] We immediately invoke Lemma 6.4.2 by noting that our assumption on ϵ ensures $\alpha_0 > 0$, and we are assured that $(\widehat{\Theta}_0, \alpha_0)$ is a feasible point for the approximate SLS synthesis problem (6.13). From inequality (6.12), we then have that

$$\begin{aligned} J(g, \widehat{\Theta}_*) &\leq Q(\tilde{g}, \tilde{\Theta}_*, \alpha_*) \leq Q(\tilde{g}, \widehat{\Theta}_0, \alpha_0) \\ &= J(\tilde{g}, \widehat{\Theta}_0) + \epsilon \alpha_0 \|\widehat{\mathbf{R}}_0 B_1 + \widehat{\mathbf{N}}_0 D_{21}\| \\ &\leq J(\tilde{g}, \widehat{\Theta}_0) + \epsilon \alpha_0 \frac{\|\mathbf{R}_0 B_1 + \mathbf{N}_0 D_{21}\|}{1 - \epsilon \|\mathbf{N}_0\|}, \end{aligned} \quad (6.15)$$

where the second inequality follows from the optimality of $(\tilde{\Theta}_*, \alpha_*)$, and the final inequality from the definitions of $\widehat{\mathbf{R}}_0$ and $\widehat{\mathbf{N}}_0$. Now, we repeat the argument used to derive (6.12) with expectation and reality reversed: this time we assume our design expectation was $J(g, \Theta_0)$ but our reality is $J(\tilde{g}, \widehat{\Theta}_0)$. This is a valid analogy as Θ_0 satisfies the SLS equations for (Z, e_1, g^\top) . With the true and estimated parameters reversed, we can thus bound $J(\tilde{g}, \widehat{\Theta}_0)$ by

$$J(\tilde{g}, \widehat{\Theta}_0) \leq J(g, \Theta_0) + \frac{\epsilon}{1 - \epsilon \|\mathbf{N}_0\|} \left\| \begin{bmatrix} 1 + g^\top \mathbf{N}_0 \\ \mathbf{L}_0 \end{bmatrix} \right\| \|\mathbf{R}_0 B_1 + \mathbf{N}_0 D_{21}\|. \quad (6.16)$$

Finally, combining bounds (6.15) and (6.16) and plugging in α_0 gives

$$\begin{aligned} J(g, \widehat{\Theta}^*) - J(g, \Theta_0) &\leq \left(\alpha_0 + \left\| \begin{bmatrix} 1 + g^\top \mathbf{N}_0 \\ \mathbf{L}_0 \end{bmatrix} \right\| \right) \frac{\epsilon \|\mathbf{R}_0 B_1 + \mathbf{N}_0 D_{21}\|}{1 - \epsilon \|\mathbf{N}_0\|} \\ &= \frac{2\epsilon}{1 - 2\epsilon \|\mathbf{N}_0\|} \left\| \begin{bmatrix} 1 + g^\top \mathbf{N}_0 \\ \mathbf{L}_0 \end{bmatrix} \right\| \|\mathbf{R}_0 B_1 + \mathbf{N}_0 D_{21}\|. \end{aligned}$$

■

6.4.3 Coarse-grained ID and the proof of Theorem 6.4.1

First, to prove the sample complexity of synthesizing a stabilizing controller based on an approximate system, we require an intermediary lemma on how well coarse-grained identification can identify the true system. The proof of the lemma (and the related change necessary for Corollary 6.4.1) is deferred to the Appendix.

Lemma 6.4.3. *Assume we estimate the system g by a length- r FIR system \tilde{g} using coarse-grained identification (output noise only) on m experiments, where the inputs $\mathbf{u}^{(i)}$ are constrained to lie in a unit ℓ_p^T ball. Then, with probability at least $1 - \eta$,*

$$\|\delta\|_2 \leq 2\sqrt{\log 2} \frac{\sigma^2 r^{2/\max(p,2)}}{m} \left(1 + \sqrt{2\log \eta^{-1}}\right).$$

Assuming we take m large enough such that $\|\delta\|_2 = \epsilon < (4\|\mathbf{N}_0\|)^{-1}$ (implied by taking $m \gtrsim \sigma^2 r \|\mathbf{N}_0\|^2 \sqrt{\log(\eta^{-1})}$), we have

$$\frac{2\epsilon}{1 - 2\epsilon \|\mathbf{N}_0\|} \leq 8\sqrt{\log 2} \frac{\sigma^2 r^{2/\max(p,2)}}{m} \left(1 + \sqrt{2\log \eta^{-1}}\right).$$

Finally, we show that $\tilde{\mathbf{K}}_*$ is stabilizing for the true system (Z, e_1, g^\top) . Since $(\tilde{\Theta}_*, \alpha_*)$ is optimal for the approximate SLS synthesis problem for \tilde{g} , it is feasible, and thus $\alpha_* > 0$ allows us to invoke Corollary 6.4.5, as we have that $\|\tilde{\mathbf{N}}_* \delta^\top\| < 1$.

6.5 Experiments

The robust SLS procedure analyzed in the previous section requires solving an infinite-dimensional optimization problem as the responses $\{\tilde{\mathbf{R}}, \tilde{\mathbf{N}}, \tilde{\mathbf{M}}, \tilde{\mathbf{L}}\}$ are not required to be FIR. However, as an approximation, we limit them to be FIR responses of a prescribed length T . By making this restriction, the resulting optimization problem is then finite-dimensional and admits an efficient solution using off-the-shelf convex optimization solvers.⁵

⁵Code for these computations can be found at <https://github.com/rjboczar/OF-end-to-end-CDC>.

Figure 6.4 shows a quantification of this approximation. In this experiment, for each r , we chose random FIR plants with impulse response coefficients uniformly distributed in $[-1, 1]^r$. We then computed the smallest $T(r)$ such that the robust performance returned by the SLS program was within 2% relative error of the performance calculated by MATLAB's `hinfsyn` with relative tolerance 10^{-8} . Figure 6.4 also shows this calculation when each plant was normalized to have unit \mathcal{H}_∞ -norm.

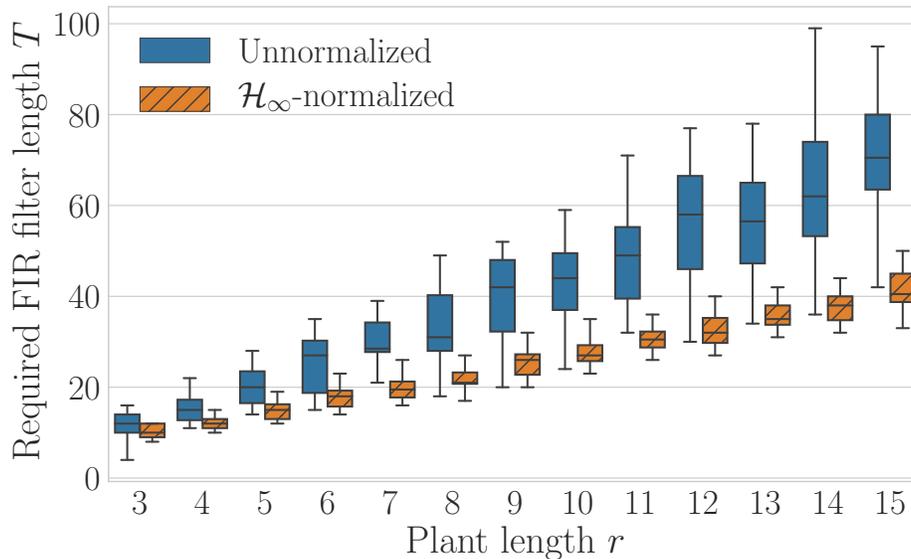


Figure 6.4: FIR approximation length T required to achieve small relative error in the robust performance objective. For each plant length r , the boxes denote the middle quartiles for 10 random plants, and the whiskers show the extents of the data.

6.5.1 Optimization Model

Let $\text{vec}(\mathbf{F}) := [\mathbf{F}_0^\top \dots \mathbf{F}_{n-1}^\top]^\top$, $\overline{\text{vec}}(\mathbf{F}) := [\mathbf{F}_0 \dots \mathbf{F}_{n-1}]$, and \mathbf{I} be the (static) identity transfer function. Furthermore, appealing to the SDP characterization of \mathcal{H}_∞ -bounded FIR systems (Dumitrescu [2017] Thm. 5.8), define

$$H(Q, \mathbf{F}, \gamma) := \begin{bmatrix} Q & \text{vec}(\mathbf{F}) \\ \text{vec}(\mathbf{F})^\top & \gamma I \end{bmatrix}.$$

Then, under the approximate assumption that $\tilde{\mathbf{R}}, \tilde{\mathbf{N}}, \tilde{\mathbf{M}}, \tilde{\mathbf{L}}$ are FIR of length T , and using the notation $Q_{[j,k]}$ for the (j, k) -th block of Q , we can write the full optimization problem of

solving (6.13) for a fixed α :

$$\begin{aligned}
& \underset{\substack{\{\tilde{\mathbf{R}}, \tilde{\mathbf{M}}, \tilde{\mathbf{N}}, \tilde{\mathbf{L}}\}, \\ t_1, t_2, t_3, t_4}}{\text{minimize}} & t_1 + t_2 \\
\text{subject to} & \begin{aligned} & \begin{bmatrix} \overline{\text{vec}}(\tilde{\mathbf{R}}) & 0 \end{bmatrix} - \begin{bmatrix} 0 & A \overline{\text{vec}}(\tilde{\mathbf{R}}) \end{bmatrix} = \begin{bmatrix} 0 & B_2 \overline{\text{vec}}(\tilde{\mathbf{M}}) + \overline{\text{vec}}(\mathbf{I}) \end{bmatrix} \\ & \begin{bmatrix} \overline{\text{vec}}(\tilde{\mathbf{N}}) & 0 \end{bmatrix} - \begin{bmatrix} 0 & A \overline{\text{vec}}(\tilde{\mathbf{N}}) \end{bmatrix} = \begin{bmatrix} 0 & B_2 \overline{\text{vec}}(\tilde{\mathbf{L}}) \end{bmatrix} \\ & \begin{bmatrix} \text{vec}(\tilde{\mathbf{R}}) \\ 0 \end{bmatrix} - \begin{bmatrix} 0 \\ \text{vec}(\tilde{\mathbf{R}})A \end{bmatrix} = \begin{bmatrix} 0 \\ \text{vec}(\tilde{\mathbf{N}})C_2 + \text{vec}(\mathbf{I}) \end{bmatrix} \\ & \begin{bmatrix} \text{vec}(\tilde{\mathbf{M}}) \\ 0 \end{bmatrix} - \begin{bmatrix} 0 \\ \text{vec}(\tilde{\mathbf{M}})A \end{bmatrix} = \begin{bmatrix} 0 \\ \text{vec}(\tilde{\mathbf{L}})C_2 \end{bmatrix} \\ & \begin{cases} 1 \leq i \leq 4 \\ 0 \leq k \leq n \end{cases} \begin{cases} H(Q^{(i)}, \mathbf{F}^{(i)}, t_i) \succeq 0 \\ \sum_{j=1}^{T-k} Q_{[j+k, j]}^{(i)} = \delta_k t_i I \end{cases} \\ & \mathbf{F}^{(1)} = \begin{bmatrix} \tilde{g}^\top & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \\ \tilde{\mathbf{M}} & \tilde{\mathbf{L}} \end{bmatrix} \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix} + D_{11} \\ & \mathbf{F}^{(2)} = \epsilon \alpha \left(\tilde{\mathbf{R}}B_1 + \tilde{\mathbf{N}}D_{21} \right) \\ & \mathbf{F}^{(3)} = \epsilon \alpha \tilde{\mathbf{N}}, \quad \mathbf{F}^{(4)} = \begin{bmatrix} 1 + \tilde{g}^\top \tilde{\mathbf{N}} \\ \tilde{\mathbf{L}} \end{bmatrix} \\ & t_3 + t_4 \leq \alpha. \end{aligned}
\end{aligned}$$

6.5.2 Computational Results

Instead of using Lemma 6.4.3 directly, we use a simulation-based technique⁶ (based on looking at the empirical histogram) to achieve tighter probabilistic tail bounds on $\|\delta\|_2$.

In what follows, we consider the following quantities:

- (i) J_{nominal} : the cost achieved on the true system when the controller was designed using `hifsyn` with the approximate system
- (ii) \hat{J} : the cost achieved on the true system when the controller was designed using the approximate SLS synthesis procedure
- (iii) $\delta J = \frac{J_{\text{nominal}} - \hat{J}}{J_{\text{nominal}}}$: the relative improvement of the approximate SLS synthesis procedure
- (iv) $\Delta J, \widehat{\Delta J}$: the theoretical sub-optimality bound (6.14) and the actual sub-optimality gap, respectively.

⁶The technique involves inverting the Chernoff bound to generate random variable tail bounds that hold with high probability with respect to the simulated instances. See the Appendix of Tu et al. [2017] for details.

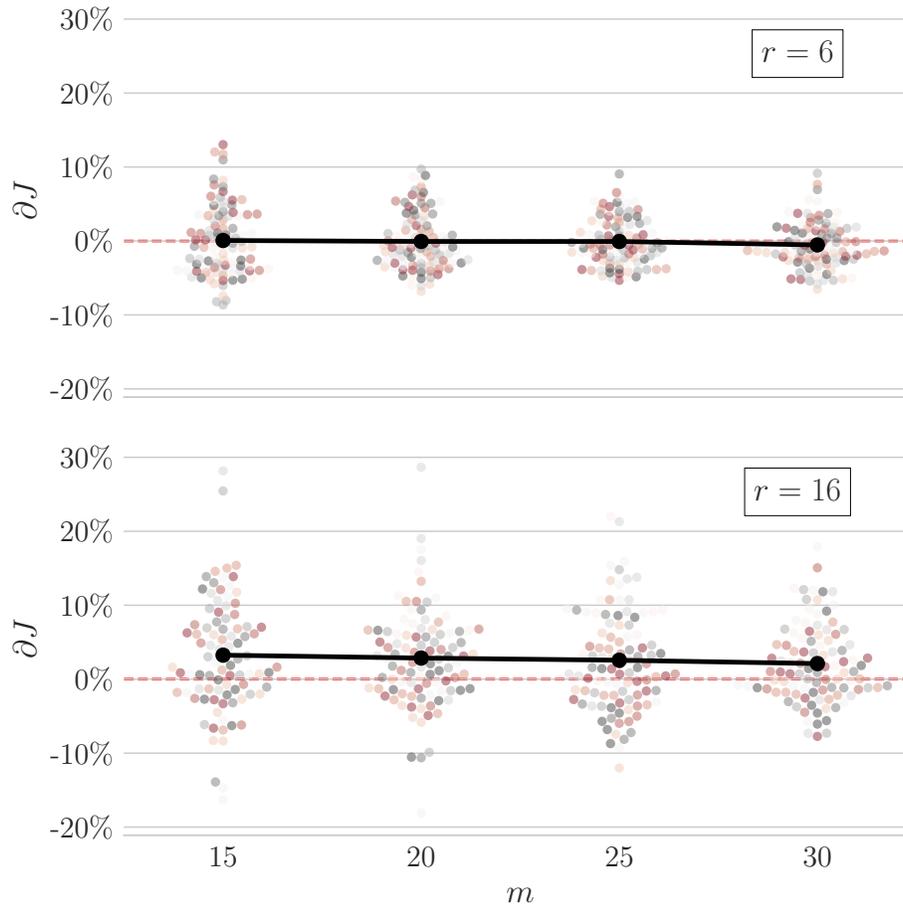


Figure 6.5: “Swarm” plot of relative improvement δJ from using the approximate SLS procedure across multiple random instances of plants and output noise (8×10 , $\sigma = 0.1$), for number of experiments $m \in \{15, 20, 25, 30\}$. Each color dot represents a different plant.

Figure 6.5 shows δJ for random instances of \mathcal{H}_∞ -normalized plants of different lengths r , swept across number of experiments m .

While it is difficult to draw precise quantitative conclusions from a suite of random plants, for the longer plants ($r = 16$) the approximate SLS procedure does perform better on average. We hypothesize that the performance depends on an effective signal-to-noise ratio (SNR): $\frac{m}{\sigma^2 r}$. At low SNR, there may not be enough data for the approximate SLS procedure to be valid (i.e. for $\epsilon < (2\|\mathbf{N}_0\|)^{-1}$). At high SNR, \tilde{g} is very close to g and thus the approximate SLS procedure may be too conservative. Thus, large improvements may be hard to come by, which is seen as m increases or r decreases in Figure 6.5. Therefore, in between these cases may be where the procedure is most effective—the $r = 16$ case may lie in this regime.

Using the data generated for Figure 6.5, Figure 6.6 shows the looseness of the end-to-end

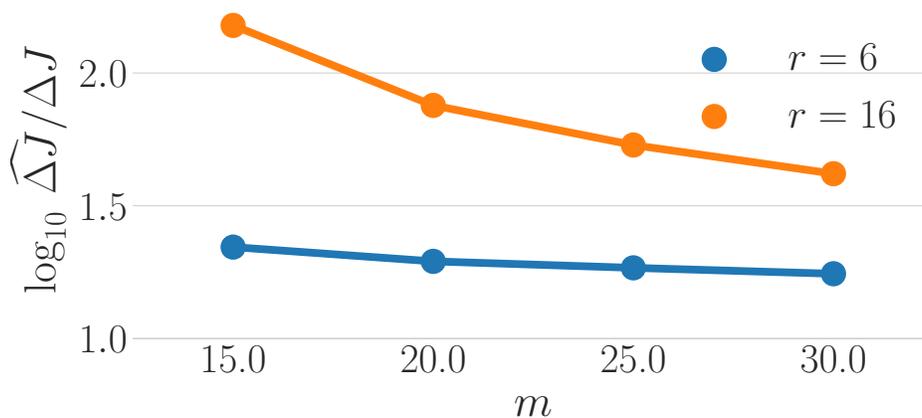


Figure 6.6: Comparison on upper bound ΔJ and actual suboptimality gap $\widehat{\Delta J}$.

bound in Proposition 6.4.3. This looseness is somewhat unsurprising, as both the formulation of the approximate SLS problem and the proof of Proposition 6.4.3 feature multiple uses of the triangle inequality.

6.6 Conclusion

In this chapter, we provide a computational tool for optimal output feedback control for the coarse-ID setting, as well as a non-asymptotic analysis of its performance. Future work involves relaxing assumptions to allow IIR or unstable plants, the latter of which may require significant modifications to the coarse-ID analysis.

Bibliography

- Pierre Apkarian and Dominikus Noll. IQC analysis and synthesis via nonsmooth optimization. *Systems & Control Letters*, 55(12):971–981, 2006.
- Jean-Yves Audibert and Sébastien Bubeck. Minimax policies for adversarial and stochastic bandits. In *COLT*, pages 217–226, Montreal, Canada, 2009.
- Somil Bansal, Roberto Calandra, Ted Xiao, Sergey Levine, and Claire J. Tomiim. Goal-Driven Dynamics Learning Via Bayesian Optimization. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 5168–5173, December 2017. doi:[10.1109/CDC.2017.8264425](https://doi.org/10.1109/CDC.2017.8264425).
- Badri Narayan Bhaskar, Gongguo Tang, and Benjamin Recht. Atomic Norm Denoising with Applications to Line Spectral Estimation. 2012, [arXiv:1204.0562](https://arxiv.org/abs/1204.0562).
- Ross Boczar, Laurent Lessard, and Benjamin Recht. Exponential Convergence Bounds Using Integral Quadratic Constraints. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 7516–7521, 2015.
- Ross Boczar, Laurent Lessard, Andrew Packard, and Benjamin Recht. Exponential Stability Analysis via Integral Quadratic Constraints. 2017, [arXiv:1706.01337](https://arxiv.org/abs/1706.01337).
- Ross Boczar, Nikolai Matni, and Benjamin Recht. Finite-Data Performance Guarantees for the Output-Feedback Control of an Unknown System. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 2994–2999, 2018.
- Albrecht Böttcher and Sergie M. Grudsky. *Toeplitz Matrices, Asymptotic Linear Algebra, and Functional Analysis*. 2000.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- Stephen Boyd and Laurent El Ghaoui. Method of Centers for Minimizing Generalized Eigenvalues. *Linear Algebra and its Applications*, 188:63–111, 1993.
- Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

- Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*, volume 15. SIAM, 1994.
- S. Bubeck and N. Cesa-Bianchi. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- Marco C Campi and Erik Weyer. Finite Sample Properties of System Identification Methods. *IEEE Transactions on Automatic Control*, 47(8), 2002.
- Jie Chen and Guoxiang Gu. *Control-Oriented System Identification: An \mathcal{H}_∞ Approach*. Wiley-Interscience, 2000.
- Jie Chen and Carl N. Nett. The Carathéodory-Fejér Problem and \mathcal{H}_∞ Identification: A Time Domain Approach. In *Proceedings of 32nd IEEE Conference on Decision and Control*, 1993.
- M. Corless and G. Leitmann. Bounded Controllers for Robust Exponential Convergence. *Journal of Optimization Theory and Applications*, 76(1):1–12, 1993. ISSN 0022-3239, 1573-2878. doi:[10.1007/BF00952819](https://doi.org/10.1007/BF00952819).
- F.J. D’Amato, M.A. Rotea, A.V. Megretski, and U.T. Jönsson. New results for analysis of systems with repeated nonlinearities. *Automatica*, 37(5):739–747, 2001.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the Sample Complexity of the Linear Quadratic Regulator. 2017, [arXiv:1710.01688](https://arxiv.org/abs/1710.01688).
- Elizabeth Dolan and Jorge Moré. Benchmarking Optimization Software with Performance Profiles. *Mathematical Programming*, 91(2):201–213, 2002.
- J. Doyle. Analysis of Feedback Systems with Structured Uncertainties. *IEE Proc., Part D*, 129(6):242–250, November 1982.
- John Doyle, Bruce Francis, and Allen Tannenbaum. *Feedback Control Theory*. 1990.
- Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking Deep Reinforcement Learning for Continuous Control. In *International Conference on Machine Learning*, pages 1329–1338, 2016.
- Bogdan Dumitrescu. *Positive Trigonometric Polynomials and Signal Processing Applications*. Springer, 2017.
- Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Finite Time Identification in Unstable Linear Systems. October 2017, [arXiv:1710.01852](https://arxiv.org/abs/1710.01852).
- Maryam Fazel, Rong Ge, Sham M Kakade, and Mehran Mesbahi. Global Convergence of Policy Gradient Methods for Linearized Control Problems. 2018, [arXiv:1801.05039](https://arxiv.org/abs/1801.05039).

- Randy A Freeman. Noncausal Zames-Falb Multipliers for Tighter Estimates of Exponential Convergence Rates. In *2018 Annual American Control Conference (ACC)*, pages 2984–2989, 2018.
- László Gerencsér. $AR(\infty)$ Estimation and Nonparametric Stochastic Complexity. *IEEE Transactions on Information Theory*, 38(6), 1992.
- Alexander Goldenshluger. Nonparametric Estimation of Transfer Functions: Rates of Convergence and Adaptation. *IEEE Transactions on Information Theory*, 44(2):644–658, March 1998. ISSN 0018-9448.
- Alexander Goldenshluger and Assaf Zeevi. Nonasymptotic Bounds for Autoregressive Time Series Modeling. *The Annals of Statistics*, 29(2), 2001.
- Michael Grant and Stephen Boyd. Graph Implementations for Nonsmooth Convex Programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. http://stanford.edu/~boyd/graph_dcp.html.
- Michael Grant and Stephen Boyd. CVX: Matlab Software for Disciplined Convex Programming, version 2.1. 2014. <http://cvxr.com/cvx>.
- Sandeep Gupta and Suresh M. Joshi. Some Properties and Stability Results for Sector-Bounded LTI Systems. In *Proceedings of 1994 33rd IEEE Conference on Decision and Control*, 1994.
- Moritz Hardt, Tengyu Ma, and Benjamin Recht. Gradient Descent Learns Linear Dynamical Systems. 2016, [arXiv:1609.05191](https://arxiv.org/abs/1609.05191).
- Elad Hazan, Karan Singh, and Cyril Zhang. Learning Linear Dynamical Systems via Spectral Filtering. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 6702–6712, 2017. <http://papers.nips.cc/paper/7247-learning-linear-dynamical-systems-via-spectral-filtering.pdf>.
- Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. Spectral Filtering for General Linear Dynamical Systems. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 4634–4643, 2018. <http://papers.nips.cc/paper/7714-spectral-filtering-for-general-linear-dynamical-systems.pdf>.
- W. P. Heath and A. G. Wills. Zames-Falb multipliers for quadratic programming. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 963–968, 2005.

- W. P. Heath, J. Carrasco, and D. A. Altshuler. Stability Analysis of Asymmetric Saturation via Generalised Zames–Falb Multipliers. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 3748–3753, December 2015. doi:[10.1109/CDC.2015.7402801](https://doi.org/10.1109/CDC.2015.7402801).
- Arthur J. Helmicki, Clas A. Jacobson, and Carl N. Nett. Control Oriented System Identification: A Worst-Case/Deterministic Approach in \mathcal{H}_∞ . *IEEE Transactions on Automatic Control*, 36(10):1163–1176, 1991.
- Haitham Hindi, Chang-Yun Seong, and Stephen Boyd. Computing Optimal Uncertainty Models from Frequency Domain Data. In *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, 2002.
- B. Hu and P. Seiler. Exponential Decay Rate Conditions for Uncertain Linear Systems Using Integral Quadratic Constraints. *IEEE Transactions on Automatic Control*, 61(11):3631–3637, 2016. ISSN 0018-9286. doi:[10.1109/TAC.2016.2521781](https://doi.org/10.1109/TAC.2016.2521781).
- Bin Hu. *A Robust Control Perspective on Optimization of Strongly-Convex Functions*. PhD thesis, 2016.
- J.L. Jerez, P.J. Goulart, S. Richter, G.A. Constantinides, E.C. Kerrigan, and M. Morari. Embedded Online Optimization for Model Predictive Control at Megahertz Rates. *IEEE Transactions on Automatic Control*, 59(12):3238–3251, 2014. ISSN 0018-9286. doi:[10.1109/TAC.2014.2351991](https://doi.org/10.1109/TAC.2014.2351991).
- Eric Jonas, Qifan Pu, Shivaram Venkataraman, Ion Stoica, and Benjamin Recht. Occupy the Cloud: Distributed Computing for the 99%. In *Proceedings of the 2017 Symposium on Cloud Computing*, 2017.
- Ulf Jönsson. A Nonlinear Popov Criterion. In *Proceedings of the 36th IEEE Conference on Decision and Control*, volume 4, pages 3523–3527, 1997.
- Jean-Pierre Kahane. *Some Random Series of Functions*. 1994.
- Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, third edition, 2002.
- Keiji Konishi and Hideki Kokame. Robust stability of Lure systems with time-varying uncertainties: A linear matrix inequality approach. *International Journal of Systems Science*, 30(1):3–9, 1999. ISSN 0020-7721. doi:[10.1080/002077299292605](https://doi.org/10.1080/002077299292605). <http://www.tandfonline.com/doi/abs/10.1080/002077299292605>.
- Nicholas Kottenstette and Panos J Antsaklis. Relationships Between Positive Real, Passive Dissipative, & Positive Systems. In *American Control Conference*, pages 409–416, 2010.
- Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and Design of Optimization Algorithms Via Integral Quadratic Constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.

- Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-End Training of Deep Visuomotor Policies. *Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- Lennart Ljung. *System Identification: Theory for the User*. Prentice Hall, 1999.
- Lennart Ljung and Bo Wahlberg. Asymptotic properties of the least-squares method for estimating transfer functions and disturbance spectra. *Advances in Applied Probability*, 24(2):412–440, June 1992.
- Lennart Ljung and Zhen-Dong Yuan. Asymptotic Properties of Black-Box Identification of Transfer Functions. *IEEE Transactions on Automatic Control*, 30(6), 1985.
- A. Marco, P. Hennig, S. Schaal, and S. Trimpe. On the Design of Lqr Kernels for Efficient Controller Learning. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 5193–5200, December 2017. doi:[10.1109/CDC.2017.8264429](https://doi.org/10.1109/CDC.2017.8264429).
- N. Matni, Y. S. Wang, and J. Anderson. Scalable System Level Synthesis for Virtually Localizable Systems. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 3473–3480, December 2017. doi:[10.1109/CDC.2017.8264168](https://doi.org/10.1109/CDC.2017.8264168).
- Daniel J McDonald, Cosma Rohilla Shalizi, and Mark Schervish. Nonparametric Risk Bounds for Time-Series Forecasting. *Journal of Machine Learning Research*, 18, 2017.
- D. McFarlane and K. Glover. A Loop Shaping Design Procedure Using H_∞ Synthesis. *IEEE Transactions on Automatic Control*, 37(6), 1992.
- Mark W. Meckes. On the Spectral Norm of a Random Toeplitz Matrix. *Electronic Communications in Probability*, 12, 2007.
- Alexandre Megretski and Anders Rantzer. System Analysis Via Integral Quadratic Constraints. *IEEE Transactions on Automatic Control*, 42(6):819–830, 1997.
- MOSEK ApS. *The MOSEK Optimization Toolbox for MATLAB Manual. Version 7.1 (Revision 28)*. 2015. <http://docs.mosek.com/7.1/toolbox/index.html>.
- Matías I. Müller, Patricio E. Valenzuela, Alexandre Proutiere, and Cristian R. Rojas. A stochastic multi-armed bandit approach to nonparametric \mathcal{H}_∞ -norm estimation. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017.
- Tom Oomen, Rick van der Maas, Cristian R. Rojas, and Håkan Hjalmarsson. Iterative Data-Driven Norm Estimation of Multivariable Systems With Application to Robust Active Vibration Isolation. *IEEE Transactions on Control Systems Technology*, 22(6):2247–2260, 2014.
- Samet Oymak and Necmiye Ozay. Non-asymptotic Identification of LTI Systems from a Single Trajectory. June 2018, [arXiv:1806.05722](https://arxiv.org/abs/1806.05722).

- Andrew Packard and John Doyle. The Complex Structured Singular Value. *Automatica*, 29(1), 1993.
- Friedrick Pukelsheim. *Optimal Design of Experiments*. 1993.
- Gianmarco Rallo, Simone Formentin, Cristian R. Rojas, Tom Oomen, and Sergio M. Savaresi. Data-driven \mathcal{H}_∞ -norm estimation via expert advice. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017.
- A. Rantzer. Friction analysis based on integral quadratic constraints. *International Journal of Robust and Nonlinear Control*, 11(7):645–652, 2001. ISSN 1099-1239. doi:10.1002/rnc.621. <http://dx.doi.org/10.1002/rnc.621>.
- Anders Rantzer. On the Kalman–Yakubovich–Popov lemma. *Systems & Control Letters*, 28(1):7–10, 1996.
- Anders Rantzer and Alexandre Megretski. System Analysis via Integral Quadratic Constraints, Part II. Technical Report ISRN LUTFD2/TFRT-.1em-7559-.1em-SE, Department of Automatic Control, Lund University, Sweden, 1997.
- R. T. Rockafellar. Characterization of the Subdifferentials of Convex Functions. *Pacific Journal of Mathematics*, 17(3):497–510, 1966. <https://projecteuclid.org:443/euclid.pjm/1102994514>.
- Cristian R. Rojas, Tom Oomen, Håkan Hjalmarsson, and Bo Wahlberg. Analyzing iterations in identification with application to nonparametric \mathcal{H}_∞ -norm estimation. *Automatica*, 48(11):2776–2790, 2012.
- Michael G. Safonov and Vishwesh V. Kulkarni. Zames–Falb Multipliers for MIMO Nonlinearities. In *Proceedings of the 2000 American Control Conference*, volume 6, pages 4144–4148, 2000.
- R. Salem and A. Zygmund. Some Properties of Trigonometric Series Whose Terms Have Random Signs. *Acta Mathematica*, 91, 1954.
- Tuhin Sarkar and Alexander Rakhlin. Near optimal finite time identification of arbitrary linear dynamical systems. December 2018, [arXiv:1812.01251](https://arxiv.org/abs/1812.01251).
- Tuhin Sarkar, Alexander Rakhlin, and Munther A. Dahleh. Finite-Time System Identification for Partially Observed LTI Systems of Unknown Order. February 2019, [arXiv:1902.01848](https://arxiv.org/abs/1902.01848).
- Peter Seiler. Stability Analysis with Dissipation Inequalities and Integral Quadratic Constraints. *IEEE Transactions on Automatic Control*, 60(6):1704–1709, 2015. ISSN 0018-9286. doi:10.1109/TAC.2014.2361004.

- Parikshit Shah, Badri Narayan Bhaskar, Gongguo Tang, and Benjamin Recht. Linear System Identification via Atomic Norm Regularization. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 2012.
- Ritei Shibata. Asymptotically Efficient Selection of the Order of the Model for Estimating Parameters of a Linear Process. *The Annals of Statistics*, 8(1), 1980.
- Max Simchowitz, Ross Boczar, and Benjamin Recht. Learning Linear Dynamical Systems with Semi-Parametric Least Squares. February 2019, [arXiv:1902.00768](https://arxiv.org/abs/1902.00768).
- Erin Summers. *Performance Analysis of Nonlinear Systems Combining Integral Quadratic Constraints and Sum-of-Squares Techniques*. PhD thesis, University of California, Berkeley, 2012. <http://oskicat.berkeley.edu/record=b20753491>.
- Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. 2009.
- Stephen Tu, Ross Boczar, Andrew Packard, and Benjamin Recht. Non-Asymptotic Analysis of Robust Control from Coarse-Grained Identification. 2017, [arXiv:1707.04791](https://arxiv.org/abs/1707.04791).
- Stephen Tu, Ross Boczar, and Benjamin Recht. Minimax Lower Bounds for \mathcal{H}_∞ -Norm Estimation. 2018a, [arXiv:1809.10855](https://arxiv.org/abs/1809.10855).
- Stephen Tu, Ross Boczar, and Benjamin Recht. On the Approximation of Toeplitz Operators for Nonparametric \mathcal{H}_∞ -norm Estimation. In *2018 Annual American Control Conference (ACC)*, 2018b.
- Klaske van Heusden, Alireza Karimi, and Dominique Bonvin. Data-Driven Estimation of the Infinity Norm of a Dynamical System. In *2007 IEEE 46th Conference on Decision and Control*, 2007.
- Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. 2018.
- M. Vidyasagar and Rajeeva L Karandikar. A learning theory approach to system identification and stochastic adaptive control. *Journal of Process Control*, 18(3), 2008.
- Bo Wahlberg, Märta Barenthin Syberg, and Håkan Hjalmarsson. Non-parametric methods for L_2 -gain estimation using iterative experiments. *Automatica*, 46(8):1376–1381, 2010.
- Yuh-Shyang Wang, Nikolai Matni, and John C. Doyle. A System Level Approach to Controller Synthesis. *IEEE Transactions on Automatic Control*, 2019.
- Siep Weiland and Carsten Scherer. *Linear Matrix Inequalities in Control*. 2015. <https://www.imng.uni-stuttgart.de/mst/files/LectureNotes.pdf>.
- Jan C. Willems. Dissipative Dynamical Systems Part I: General Theory. *Archive for Rational Mechanics and Analysis*, 45(5):321–351, 1972.

- Yihong Wu. Lecture notes for ECE598YW: Information-theoretic methods for high-dimensional statistics. 2017. <http://www.stat.yale.edu/~yw562/teaching/598/ref.html>.
- G. Zames. On the Input-Output Stability of Time-Varying Nonlinear Feedback Systems—Part II: Conditions Involving Circles in the Frequency Plane and Sector Nonlinearities. *IEEE Transactions on Automatic Control*, 11(2,3):465–476, 1966.
- G. Zames and P. L. Falb. Stability Conditions for Systems with Monotone and Slope-Restricted Nonlinearities. *SIAM Journal on Control*, 6(1):89–108, 1968. ISSN 0036-1402. doi:10.1137/0306007.
- Huanshui Zhang, Lihua Xie, and Yeng Chai Soh. Discrete J-spectral factorization. *Systems & Control Letters*, 43(4):275–286, 2001.
- K. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control*. 1995.

Appendix A

Supplementary Material for Chapter 3

A.1 Proof of Proposition 3.3.1

In these appendices, we will restate theorems, propositions, and lemmas for convenience.

Proposition 3.3.1. *Suppose $G(z)$ has a minimal realization (A, B, C, D) . If the interconnection in Figure 3.3 is BIBO stable, then the interconnection in Figure 3.1 with initial state x_0 is exponentially stable.*

Suppose the interconnection of Figure 3.3 is stable. Then, there exists some $K > 0$ such that for any choice of the signals e and f and for all T ,

$$\sum_{k=0}^T (\|w_k\|^2 + \|v_k\|^2) \leq K \sum_{k=0}^T (\|e_k\|^2 + \|f_k\|^2). \quad (\text{A.1})$$

The proof will follow by carefully choosing e and f to transform Figure 3.3 into Figure 3.1. To this end, note that (A, B) is controllable by assumption. Thus, there exists a finite sequence of inputs u_0, \dots, u_{n-1} and corresponding outputs y_0, \dots, y_{n-1} that drives the state of G from $\xi_0 = 0$ to $\xi_n = x_0$. Therefore, if we set

$$e_k = \begin{cases} \rho^{-k} u_k & 0 \leq k < n \\ 0 & k \geq n \end{cases}, \quad f_k = \begin{cases} -\rho^{-k} y_k & 0 \leq k < n \\ 0 & k \geq n \end{cases},$$

we obtain $\xi_n = x_0$ in the interconnection of Figure 3.3. Moreover, $\rho_- \circ \rho_+$ is the identity operator. It follows that for $k \geq n$, the two interconnections become identical and therefore $\xi_k = x_{k-n}$.

Substituting into (A.1), we conclude that

$$\sum_{k=0}^T (\|w_k\|^2 + \|v_k\|^2) \leq K \sum_{k=0}^{n-1} (\|e_k\|^2 + \|f_k\|^2). \quad (\text{A.2})$$

The right-hand side of (A.2) is independent of T , but (A.2) holds for all T , so we must have

$$\lim_{k \rightarrow \infty} \|w_k\| = 0 \quad \text{and} \quad \lim_{k \rightarrow \infty} \|v_k\| = 0.$$

Now, for $k \geq n$, we have $w_k = \rho^{-k}u_k$ and $v_k = \rho^{-k}y_k$. Therefore, there exists some constant $c > 0$ such that

$$\|u_k\| \leq c\rho^k \quad \text{and} \quad \|y_k\| \leq c\rho^k.$$

Since (A, C) is observable by assumption, we may choose L such that the eigenvalues of $A + LC$ are all zero. Then, rewrite the dynamics of G as

$$x_{k+1} = \bar{A}x_k + \bar{B}h_k, \tag{A.3}$$

where $\bar{A} := A + LC$, $\bar{B} := [LD + B \quad -L]$, and $h_k := [u_k^\top \quad y_k^\top]^\top$. Iterating (A.3), we obtain

$$x_k = \bar{A}^k x_0 + \sum_{i=0}^{k-1} \bar{A}^{k-1-i} \bar{B} h_i. \tag{A.4}$$

Since all eigenvalues of \bar{A} are zero, \bar{A} is nilpotent and so $\bar{A}^n = 0$. For $k \geq n$, (A.4) then becomes

$$x_k = \sum_{i=0}^{n-1} \bar{A}^{n-1-i} \bar{B} h_{k-n+i}.$$

We can now finally bound the size of the state using the triangle inequality:

$$\begin{aligned} \|x_k\| &\leq \underbrace{\|[\bar{A}^{n-1}\bar{B} \quad \dots \quad \bar{A}\bar{B} \quad \bar{B}]\|}_{\gamma} \sum_{i=k-n}^{k-1} \|h_i\| \\ &\leq 2\gamma c \left(\frac{\rho^{-n} - 1}{1 - \rho} \right) \rho^k. \end{aligned}$$

■

A.2 Proof of Theorem 3.6.1 and Related Extensions

Theorem 3.6.1 (Zames–Falb ρ -IQC). *Suppose Δ is static and slope-restricted on $[\alpha, \beta]$. Then $\Delta \in \text{IQC}(\Pi(z), \rho)$ where Π is the Zames–Falb IQC (3.12) and \hat{h} also satisfies the additional constraint*

$$\sum_{k=0}^{\infty} \rho^{-2k} |h_k| \leq 1.$$

A.2.1 $[0, \infty]$ -slope-restricted Case

We will prove this general result by first considering the simpler case where the slope restriction is on $[\alpha, \beta] = [0, \infty]$ and $H(z) = \pm\gamma_l\rho^{2j}z^{-l}$ for some constants $0 < \gamma_l \leq 1$. Note that this choice trivially satisfies (3.6.1), and the extensions of (3.6.1) follow for specific restrictions on γ_k and mixed-sign h_k . In this case, the Π from (3.12) (first taking the positive sign in $H(z)$) becomes

$$\Pi = \begin{bmatrix} 0 & 1 - \gamma_j\rho^{2j}\bar{z}^{-l} \\ 1 - \gamma_l\rho^{2l}z^{-l} & 0 \end{bmatrix} \quad (\text{A.5})$$

where \bar{z} denotes the complex conjugate of z . We call (A.5) the “off-by- l ” Zames–Falb IQC. We would like to show that $\Delta \in \text{IQC}(\Pi(z), \rho)$. Appealing to Definition 3.3.2, Remark 3, and Proposition 3.6.1, this amounts to proving that

$$\sum_{k=0}^{\infty} \rho^{-2k} u_k^\top (y_k - \gamma_l \rho^{2l} y_{k-l}) \geq 0. \quad (\text{A.6})$$

We will prove (A.6) by borrowing the approach from Lessard et al. [2016]. If Δ is multidimensional, we require that Δ be the gradient of a potential function [Heath and Wills, 2005]. By the assumption that Δ is slope-restricted on $[0, \infty]$, we have that Δ is monotone, i.e.

$$(\Delta(x) - \Delta(y))^\top (x - y) \geq 0 \quad \text{holds for all } x, y.$$

In other words, Δ is monotone. Furthermore (see Rockafellar [1966]), there exists¹ a convex scalar function g such that $\Delta(x)$ is in the subdifferential of g at x for all x , meaning

$$g(y) \geq g(x) + \Delta(x)^\top (y - x) \quad \text{for all } x, y.$$

Moreover, setting $(x, y) \mapsto (y_k, 0)$ or $(x, y) \mapsto (y_k, y_{k-l})$ leads to the two inequalities

$$u_k^\top y_k \geq g(y_k) \quad (\text{A.7})$$

$$u_k^\top (y_k - y_{k-l}) \geq g(y_k) - g(y_{k-l}). \quad (\text{A.8})$$

¹In the loop-transformed case for $[\alpha, \beta]$ -slope-restricted nonlinearities, the argument is more delicate. See Freeman [2018] for a rigorous discussion.

We will assume for simplicity that $g(x) \geq 0$ for all x , and we will first prove the case where we take the positive sign in $H(z)$. Substituting (A.7) and (A.8) into the left-hand side of (A.6), the partial sum from 0 to T is

$$\begin{aligned}
\sum_{k=0}^T \rho^{-2k} u_k^\top (y_k - \gamma_l \rho^{2l} y_{k-l}) &= \sum_{k=0}^T \rho^{-2k} ((1 - \gamma_l \rho^{2l}) u_k^\top y_k + \gamma_l \rho^{2l} u_k^\top (y_k - y_{k-l})) \\
&\geq \sum_{k=0}^T \rho^{-2k} ((1 - \gamma_l \rho^{2l}) g(y_k) + \gamma_l \rho^{2l} (g(y_k) - g(y_{k-l}))) \\
&= \sum_{k=0}^T \rho^{-2k} (g(y_k) - \gamma_l \rho^{2l} g(y_{k-l})) \\
&= \sum_{k=0}^{T-l} (1 - \gamma_l) \rho^{-2k} g(y_k) + \sum_{k=T-l+1}^T \rho^{-2k} g(y_k) \geq 0.
\end{aligned}$$

Since each partial sum is nonnegative, the infinite sum (which must converge) is also nonnegative, and therefore we have proven (A.6). Now, for the case where we take the negative sign in $H(z)$, further assume that Δ is an odd function, which implies g is an even function. Using this fact and convexity inequality for g with $(x, y) \mapsto (y_k, -y_{k-l})$ leads to the additional inequality

$$u_k^\top (y_k + y_{k-l}) \geq g(y_k) - g(y_{k-l}).$$

The proof of nonnegativity of the partial sums then follows as before. Thus, a $[0, \infty]$ -slope restricted Δ satisfies the off-by- l ρ -IQC (and also the negative version if Δ is assumed to be odd).

Now we consider the case of a more general $\hat{h}(z)$. Suppose $\hat{h}(z) = \sum_{k=0}^{\infty} h_k z^{-k}$ where h_k satisfies $\sum_k \gamma_k^{-1} \rho^{-2k} |h_k| \leq 1$. Then,

$$\begin{aligned}
1 - \hat{h}(z) &= 1 - \underbrace{\sum_{k=0}^{\infty} \gamma_k^{-1} \rho^{-2k} |h_k|}_{\equiv c} + \sum_{h_k \geq 0} \gamma_k^{-1} \rho^{-2k} h_k (1 - \gamma_k \rho^{2k} z^{-k}) \\
&\quad + \sum_{h_k < 0} \gamma_k^{-1} \rho^{-2k} (-h_k) (1 + \gamma_k \rho^{2k} z^{-k}) \\
&= c(1 - \hat{h}_s) + \sum_{h_k \geq 0} \gamma_k^{-1} \rho^{-2k} h_k (1 - \hat{h}_k^+(z)) \\
&\quad + \sum_{h_k < 0} \gamma_k^{-1} \rho^{-2k} (-h_k) (1 - \hat{h}_k^-(z)),
\end{aligned}$$

where $\hat{h}_k^\pm(z) = \pm\gamma_k\rho^{2k}z^{-k}$ and $\hat{h}_s = 0$ (for illustration). Note that $\hat{h}_k(z)$ corresponds the off-by- k Zames–Falb IQC, which we proved above is a ρ -IQC, where the negative version is only used (with corresponding negative h_k) if Δ is assumed to be odd. Also, $1 - \hat{h}_s$ corresponds to the sector IQC, which is also a ρ -IQC. Now note that the general Zames–Falb IQC (3.12) is linear in $1 - \hat{h}$ and $1 - \hat{h}^*$. Therefore, since by assumption $c \geq 0$, $\Pi(z)$ is a positive linear combination of ρ -IQCs and must therefore be a ρ -IQC itself. ■

A.2.2 Specific Zames–Falb Classes

We would now like to generalize this proof (or equivalently, specify further the class of nonlinearities). Now, let us assume that the nonlinearity Δ can be written as $\Delta_2 \circ \Delta_1$, where Δ_1 satisfies a ρ -IQC of the Zames–Falb type in the preceding section where $\sum_k \gamma_k^{-1} \rho^{-2k} |h_k| \leq 1$. Furthermore, assume that Δ_2 (which is possibly time-varying) satisfies

$$(1 - \delta)u_k \leq \Delta_2(u_k) \leq (1 + \delta)u_k \quad \forall u_k, \forall k,$$

or equivalently,

$$\Delta_2(u_k) = (1 + \delta_k)u_k, \quad |\delta_k| \leq \delta$$

for some $\delta < 1$. We would like to show under what conditions $\Delta_2 \circ \Delta_1$ satisfies a Zames–Falb IQC with rate ρ .

To do this, we will show that $\Delta_2 \circ \Delta_1$ satisfies the relevant off-by- l Zames–Falb ρ -IQC, which then extends to general Zames–Falb by the preceding section. As in Section A.2.1, we would like to show that

$$\sum_{k=0}^{\infty} \rho^{-2k} \Delta(y_k)^\top (y_k \mp \gamma_l \rho^{2l} y_{k-l}) \geq 0.$$

Using our prescribed Δ (taking the negative sign for simplicity) and denoting $\Delta_1(y_k) := u_k$, we see that each partial sum satisfies

$$\sum_{k=0}^T \rho^{-2k} (1 + \delta_k) u_k^\top (y_k - \gamma_l \rho^{2l} y_{k-l}) \geq \sum_{k=0}^T \rho^{-2k} (1 + \delta_k) (g(y_k) - \gamma_l \rho^{2l} g(y_{k-l}))$$

by the same argument from the preceding section. This is then equal to

$$\sum_{k=0}^{T-l} (1 + \delta_k - (1 + \delta_{k+l})\gamma_l) \rho^{-2k} g(y_k) + \sum_{k=T-l+1}^T (1 + \delta_k) \rho^{-2k} g(y_k).$$

A sufficient condition for this sum to be positive is

$$\gamma_l \leq \frac{1 + \delta_k}{1 + \delta_{l+k}} \quad \forall k,$$

which is satisfied if, for example,

$$\gamma_l \equiv \gamma \leq \frac{1 - \delta}{1 + \delta}.$$

If so, the partial sums converge and so does the infinite sum. Again, if we further assume that Δ_1 is odd, the negative off-by- l IQC is also satisfied. The argument for general $H(z)$ follows as before. \blacksquare

Proofs for specific classes of nonlinearities in the literature correspond to specific choices of γ_k and δ_k . These are summarized in Table A.1.

Table A.1: Variable Choices for Specific Zames–Falb Proofs

Type	δ_k	γ_k	Notes
$[0, \infty]$ -slope restricted	0	1	
$[\alpha, \beta]$ -slope restricted	0	1	Loop transformation $(y, u) \mapsto (\beta y - u, u - \alpha y)$
Noisy composition	δ	1	
Stiction (slope $1/\epsilon$)	δ	$\frac{1 - \delta}{1 + \delta}$	
Quasi-monotone/odd	$\frac{R_m - 1}{R_m + 1}, \frac{R_o - 1}{R_o + 1}$	$R_m^{-1}, h_k \geq 0$ $R_o^{-1}, h_k < 0$	Notational change in def. from Heath et al. [2015] : $\underline{n}_* \mapsto \frac{2}{R_* + 1} \underline{n}_*$

A.3 Proof of Theorem 3.6.2

Theorem 3.6.2. *Assume Γ is (ρ, H) -diagonally dominant. Then, if a static nonlinearity ϕ is $[\alpha, \beta]$ slope-restricted, the repeated nonlinearity $\Delta(y) = \text{diag}\{\phi(y_i)\}$ satisfies the ρ -IQC*

$$\Pi = \begin{bmatrix} -\alpha\beta(2\Gamma - \hat{H} - \hat{H}^*) & \alpha(\Gamma - \hat{H}) + \beta(\Gamma - \hat{H}^*) \\ \alpha(\Gamma - \hat{H}^*) + \beta(\Gamma - \hat{H}) & -2\Gamma + \hat{H} + \hat{H}^* \end{bmatrix}. \quad (3.13)$$

Moreover, it admits the factorization

$$\Psi = \begin{bmatrix} \beta(\Gamma - \hat{H}) & -(\Gamma - \hat{H}) \\ -\alpha I & I \end{bmatrix}, \quad M = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}.$$

We begin with an elementary lemma; a similar one is used in [D’Amato et al. \[2001\]](#) for the proof of (3.6.2).

Lemma A.3.1. *For a repeated monotone nondecreasing nonlinearity Δ and with $u = \Delta(y)$, we have that*

$$u^{[i]}y^{[j]} + u^{[j]}y^{[i]} \leq u^{[i]}y^{[i]} + u^{[j]}y^{[j]},$$

and if ϕ is odd,

$$|u^{[i]}y^{[j]} + u^{[j]}y^{[i]}| \leq u^{[i]}y^{[i]} + u^{[j]}y^{[j]},$$

for all indices i, j .

Proof. Assume without loss of generality that $y^{[i]} \geq y^{[j]}$. By monotonicity and the fact that the nonlinearity is repeated, we must have that $u^{[i]} \geq u^{[j]}$. Thus:

$$(u^{[i]} - u^{[j]})(y^{[i]} - y^{[j]}) \geq 0.$$

If ϕ is odd, then the second equation is proven by also observing that

$$[u^{[i]}, -u^{[j]}]^\top = \Delta([y^{[i]}, -y^{[j]}]^\top) \implies (u^{[i]} + u^{[j]})(y^{[i]} + y^{[j]}) \geq 0.$$

■

Let E_{ij} denote the standard basis matrix. Now, given a diagonally dominant matrix Γ , define the following symmetric matrices:

$$\begin{aligned} \tilde{\Gamma}_{ij} &= \Gamma_{ij}(E_{ij} + E_{ji}) + (|\Gamma_{ij}| + 1)(E_{ii} + E_{jj}), \quad i \neq j \\ \tilde{\Gamma}_{ii} &= E_{ii} \\ \tilde{H}_{ij}^l(z) &= \rho^{2l} z^{-l} (E_{ij} + E_{ji}), \quad i \neq j \\ \tilde{H}_{ii}^l(z) &= \rho^{2l} z^{-l} E_{ii}. \end{aligned}$$

Proposition A.3.1. Δ satisfies the “ (i, j) off-by- l ρ -IQC” defined by

$$\Pi = \begin{bmatrix} 0 & \tilde{\Gamma}_{ij} - \tilde{H}_{ij}^{l*} \\ \tilde{\Gamma}_{ij} - \tilde{H}_{ij}^l & 0 \end{bmatrix}$$

for all ρ in $(0, 1]$.

Proof. As before, assume the ϕ is the gradient of a potential function f . For notational convenience, define the following symbols:

$$\begin{aligned} d_k &= u_k^{[i]}y_k^{[i]} + u_k^{[j]}y_k^{[j]} \\ c_k &= u_k^{[j]}y_k^{[i]} + u_k^{[i]}y_k^{[j]} \\ p_k^l &= u_k^{[j]}y_{k-l}^{[i]} + u_k^{[i]}y_{k-l}^{[j]} \\ f_k &= f(y_k^{[i]}) + f(y_k^{[j]}). \end{aligned}$$

Using convexity and the fact that the nonlinearity is repeated, we can obtain the inequalities

$$\begin{aligned} d_k - p_k^l &\geq f_k - f_{k-l} \\ d_k &\geq f_k. \end{aligned}$$

Then,

$$\begin{aligned} \sum_{k=0}^T \rho^{-2k} u_k^\top (\tilde{\Gamma}_{ij} - \tilde{H}_{ij}^l) y_k &= \sum_{k=0}^T \rho^{-2k} |\Gamma_{ij}| (d_k - c_k) + \sum_{k=0}^T \rho^{-2k} ((1 - \rho^{2l})d_k + \rho^{2l}(d_k - p_k^{ij,l})) \\ &\geq \sum_{k=0}^T \rho^{-2k} |\Gamma_{ij}| (d_k - c_k) + \sum_{k=0}^T \rho^{-2k} ((1 - \rho^{2l})f_k + \rho^{2l}(f_k - f_{k-l})). \end{aligned}$$

The first sum is nonnegative by Lemma A.3.1 and the second sum is nonnegative by the same arguments as in Appendix A.2.1. \blacksquare

We will now show that for a (ρ, H) -diagonally dominant matrix Γ , the ρ -IQC defined by

$$\Pi = \begin{bmatrix} 0 & \Gamma - \hat{H}^* \\ \Gamma - \hat{H} & 0 \end{bmatrix}$$

is a positive combination of satisfied ρ -IQCs and is thus a satisfied ρ -IQC. Toward this end,

$$\Gamma - \hat{H}(z) = \Gamma - \sum_{i \leq j, k} \rho^{-2k} H_{ij,k} \tilde{\Gamma}_{ij} + \left[\sum_{i \leq j, k} \rho^{-2k} H_{ij,k} (\tilde{\Gamma}_{ij} - \tilde{H}_{ij}^k(z)) \right].$$

The term in brackets is a nonnegative linear combination of satisfied IQCs, so let us focus on the first term:

$$Q := \Gamma - \sum_{i \leq j, k} \rho^{-2k} H_{ij,k} \tilde{\Gamma}_{ij} = \begin{cases} \Gamma_{ii} - \sum_k \left[\sum_{j=1, j \neq i}^n \rho^{-2k} H_{ij,k} (|\Gamma_{ij}| + 1) + \rho^{-2k} H_{ii,k} \right] \\ \quad (ii \text{ indices}) \\ (1 - \sum_k \rho^{-2k} H_{ij,k}) \Gamma_{ij} \\ \quad (ij \text{ indices}). \end{cases}$$

We can see that this constant matrix Q is diagonally dominant, since

$$\begin{aligned} Q_{ii} - \sum_{j=1, j \neq i}^n |Q_{ij}| &= \Gamma_{ii} - \sum_k \left[\sum_{j=1, j \neq i}^n \rho^{-2k} H_{ij,k} (|\Gamma_{ij}| + 1) + \rho^{-2k} H_{ii,k} \right] \\ &\quad - \sum_{j=1, j \neq i}^n (1 - \sum_k \rho^{-2k} H_{ij,k}) |\Gamma_{ij}| \\ &= \Gamma_{ii} - \sum_{j=1, j \neq i}^n |\Gamma_{ij}| - \sum_{j=1}^n \sum_k \rho^{-2k} H_{ij,k} \geq 0, \end{aligned}$$

where the last line follows from the assumption that Γ is (ρ, H) -diagonally dominant. Similar modifications to the assumptions and proof hold for odd and $[\alpha, \beta]$ -sector nonlinearities. Also, note that in the case of diagonal (but not necessarily repeating) Δ , Lemma A.3.1 will not hold in general. Thus, considering where it is used in the proof, we would need to constrain Γ and H to be diagonal.

A.4 Computational Considerations

A.4.1 Homogeneity

We may leverage the structure of the repeated Zames–Falb ρ -IQC's to reduce the size and complexity of the LMI (3.6).

Proposition A.4.1. (*Homogeneity simplification*) *If we are searching over a combination of repeated Zames–Falb IQCs with fixed α, β, ρ and varying over Γ and H matrices, and fixed non-Zames–Falb ρ -IQC's, i.e.*

$$\Pi_A = \left\{ \sum_{\theta=1}^n \lambda_{\theta} \Pi_{ZF}(\Gamma_{\theta}, H_{\theta}) + \sum_{\delta=1}^r \lambda_{\delta} \Pi_{\delta} \mid \lambda_{\theta} \geq 0, \lambda_{\delta} \geq 0, \Gamma_{\theta} \text{ is } (\rho, H_{\theta})\text{-diag. dom.} \right\}.$$

Then, searching over the associated LMI is equivalent to searching over

$$\Pi_B = \left\{ \lambda \Pi_{ZF}(\Gamma, H) + \sum_{\delta=1}^r \lambda_{\delta} \Pi_{\delta} \mid \lambda \geq 0, \lambda_{\delta} \geq 0, \Gamma \text{ is } (\rho, H)\text{-diag. dom.} \right\}.$$

That is, $\Pi_A = \Pi_B$.

Proof. $\Pi_B \subseteq \Pi_A$ is immediate. To check the other direction, first assume $\sum_{\theta} \lambda_{\theta} = \Lambda > 0$ (the $\Lambda = 0$ case is trivial). Now, note that, due to the linearity of the repeated Zames–Falb IQC in terms of Γ and H , see that

$$\sum_{\theta=1}^n \lambda_{\theta} \Pi_{ZF}(\Gamma_{\theta}, H_{\theta}) = \Lambda \Pi_{ZF} \left(\frac{\sum_{\theta} \lambda_{\theta} \Gamma_{\theta}}{\Lambda}, \frac{\sum_{\theta} \lambda_{\theta} H_{\theta}}{\Lambda} \right).$$

All that remains is to check that the diagonal dominance definitions are satisfied; two require care. First, the filter condition on H :

$$\begin{aligned} \sum_k \rho^{-2k} \left| \frac{\sum_{\theta} \lambda_{\theta} H_{ij,k}^{\theta}}{\Lambda} \right| &\leq \sum_k \rho^{-2k} \frac{\sum_{\theta} \lambda_{\theta} |H_{ij,k}^{\theta}|}{\Lambda} \\ &= \frac{\sum_{\theta} \lambda_{\theta}}{\Lambda} \sum_k \rho^{-2k} |H_{ij,k}^{\theta}| \\ &\leq \frac{\sum_{\theta} \lambda_{\theta}}{\Lambda} \quad (\text{by assumption}) \\ &= 1. \end{aligned}$$

Second, the main diagonal dominance condition:

$$\begin{aligned}
& \frac{\sum_{\theta} \lambda_{\theta} \Gamma_{ii}^{\theta}}{\Lambda} - \sum_{j=1, j \neq i}^n \left| \frac{\sum_{\theta} \lambda_{\theta} \Gamma_{ij}^{\theta}}{\Lambda} \right| - \sum_{j=1}^n \sum_k \rho^{-2k} \left| \frac{\sum_{\theta} \lambda_{\theta} H_{ij,k}^{\theta}}{\Lambda} \right| \\
& \geq \frac{\sum_{\theta} \lambda_{\theta} \Gamma_{ii}^{\theta}}{\Lambda} - \frac{\sum_{\theta} \lambda_{\theta}}{\Lambda} \sum_{j=1, j \neq i}^n |\Gamma_{ij}^{\theta}| - \frac{\sum_{\theta} \lambda_{\theta}}{\Lambda} \sum_{j=1}^n \sum_k \rho^{-2k} |H_{ij,k}^{\theta}| \\
& = \Lambda^{-1} \sum_{\theta} \lambda_{\theta} \left(\Gamma_{ii}^{\theta} - \sum_{j=1, j \neq i}^n |\Gamma_{ij}^{\theta}| - \sum_{j=1}^n \sum_k \rho^{-2k} |H_{ij,k}^{\theta}| \right) \\
& \geq 0, \text{ by assumption.}
\end{aligned}$$

The non-Zames–Falb IQCs carry straight through in both directions. Therefore, we see that $\sum_{\theta=1}^n \lambda_{\theta} \Pi_{ZF}(\Gamma_{\theta}, H_{\theta}) + \sum_{\delta} \lambda_{\delta} \Pi_{\delta} \in \Pi_A$, and we have $\Pi_A = \Pi_B$. ■

A.4.2 Convexification for Repeated Nonlinearities

For repeated nonlinearities, the main LMI (3.6) with constraints can be written as

$$\begin{aligned}
& \min_{P, \lambda, \Gamma, H} 0 \\
& \text{s.t. } \mathcal{A}(P) + \lambda \mathcal{M}(\Gamma, H) \prec 0 \\
& P \succeq 0 \\
& \lambda \geq 0 \\
& \Gamma_{ii} \geq 0, \forall i \\
& \Gamma_{ij} \leq 0, \forall i \neq j \\
& H_{ij,k} \geq 0, \forall i, j, k \\
& \sum_{k=0}^{\infty} \rho^{-2k} |H_{ij,k}| \leq 1, \forall i, j \\
& \Gamma_{ii} \geq \sum_{j=1, j \neq i}^n |\Gamma_{ij}| + \sum_{j=1}^n \sum_{k=0}^{\infty} \rho^{-2k} |H_{ij,k}|, \quad \forall i
\end{aligned}$$

for some known linear functions \mathcal{A}, \mathcal{M} (note that the latter is linear in the pair $[\Gamma, H]$). This problem is nonconvex, due to the product of λ and $[\Gamma, H]$.

However, one can show that this program can be reduced to the convex problem

$$\begin{aligned}
 & \min_{\tilde{P}, \zeta, \tilde{\Gamma}, \tilde{H}} 0 \\
 \text{s.t. } & \mathcal{A}(\tilde{P}) + \mathcal{M}(\tilde{\Gamma}, \tilde{H}) \preceq -I \\
 & \tilde{P} \succeq 0 \\
 & \zeta > 0 \\
 & \tilde{\Gamma}_{ii} \geq 0, \forall i \\
 & \tilde{\Gamma}_{ij} \leq 0, \forall i \neq j \\
 & \tilde{H}_{ij,k} \geq 0, \forall i, j, k \\
 & \sum_{k=0}^{\infty} \rho^{-2k} |\tilde{H}_{ij,k}| \leq \zeta, \forall i, j \\
 & \tilde{\Gamma}_{ii} \geq \sum_{j=1, j \neq i}^n |\tilde{\Gamma}_{ij}| + \sum_{j=1}^n \sum_{k=0}^{\infty} \rho^{-2k} |\tilde{H}_{ij,k}|, \forall i.
 \end{aligned}$$

In practice, for numerical considerations it is helpful to replace $\preceq -I$ with $\preceq -tI$ and maximize over t , while placing upper bounds on ζ and t .

A.5 Proof of Theorem 3.8.1

Theorem 3.8.1. *Consider the block interconnection of Figure 3.15. Assume the error signal w_k comes i.i.d. from a distribution with zero mean and covariance matrix Λ_w , and is independent of u_k and x_k . Suppose G is given by (3.17) and ϕ satisfies a ρ -hard IQC defined by (Ψ, M, z_*) where Ψ is given by (3.18) and $0 \leq \rho \leq 1$. Assume x_* is a fixed point of the dynamical system given by (3.19) when $w_k = 0$. Consider the LMI*

$$\begin{bmatrix} \hat{A}^\top P \hat{A} - \rho^2 P & \hat{A}^\top P \hat{B} \\ \hat{B}^\top P \hat{A} & \hat{B}^\top P \hat{B} \end{bmatrix} + \lambda [\hat{C} \quad \hat{D}]^\top M [\hat{C} \quad \hat{D}] \preceq 0. \quad (3.20)$$

If (3.20) is feasible for some $P \succ 0$ and $\lambda \geq 0$, then for any x_0 , we have

$$\mathbb{E} \|x_k - x_*\|^2 \leq \kappa(P) \rho^{2k} \|x_0 - x_*\|^2 + \frac{\text{tr}(B^\top P_{11} B \Lambda_w)}{\lambda_{\min}(P)} \frac{1 - \rho^{2k}}{1 - \rho^2} \quad \forall k.$$

The complete proof follows in the same vein as the main result in Lessard et al. [2016]; we show the modifications needed in order to handle the noisy input case.

Let $x, u, z \in \ell_2^2$, along with a sequence w , be a set of sequences that satisfies the perturbed system (3.19). Suppose (P, λ) is a solution of (3.20). Suppose z_* and u_* are the values of z and u associated with the fixed point x_* of the *unperturbed* system (i.e. (3.19) with $w_k = 0$). By the definition of fixed point, denoting $\bar{q}_k := q_k - q_*$ for all symbols q , we see that

$$\begin{aligned}\bar{x}_{k+1} &= \hat{A}\bar{x}_k + \hat{B}\bar{u}_k + \begin{bmatrix} B \\ 0 \end{bmatrix} w_k \\ \bar{z}_k &= \hat{C}\bar{x}_k + \hat{D}\bar{u}_k.\end{aligned}\tag{A.9}$$

For convenience, we denote an elementary probability fact.

Lemma A.5.1. *If a_k and b_k are independent and $\mathbb{E}b_k = 0$, then for every constant matrix P ,*

$$\mathbb{E}a_k^\top P a_k = \mathbb{E}(a_k + b_k)^\top P (a_k + b_k) - \mathbb{E}b_k^\top P b_k.$$

If we multiply (3.20) on the left and right by $[\bar{x}_k^\top \ \bar{u}_k^\top]$ and its transpose, respectively, by (A.9) we have that

$$(\hat{A}\bar{x}_k + \hat{B}\bar{u}_k)^\top P (\hat{A}\bar{x}_k + \hat{B}\bar{u}_k) - \rho^2 \bar{x}_k^\top P \bar{x}_k + \lambda \bar{z}_k^\top M \bar{z}_k \leq 0.\tag{A.10}$$

Taking expectations, we can apply Lemma A.5.1 to the first term of (A.10), with

$$a_k := \hat{A}\bar{x}_k + \hat{B}\bar{u}_k, \quad b_k := \begin{bmatrix} B \\ 0 \end{bmatrix} w_k.$$

We are then left with

$$\mathbb{E}[\bar{x}_{k+1}^\top P \bar{x}_{k+1} - \rho^2 \bar{x}_k^\top P \bar{x}_k + \lambda \bar{z}_k^\top M \bar{z}_k] \leq \mathbb{E}w_k^\top B^\top P_{11} B w_k.$$

The rest of the proof follows as in Lessard et al. [2016]. Momentarily switching the index to t , multiplying each side by ρ^{-2t} and summing from $t = 0$ to $k - 1$ gives

$$\mathbb{E}\left[\sum_{t=0}^{k-1} (\rho^{-2t} \bar{x}_{t+1}^\top P \bar{x}_{t+1} - \rho^{-2(t-1)} \bar{x}_t^\top P \bar{x}_t) + \lambda \sum_{t=0}^{k-1} \rho^{-2t} \bar{z}_t^\top M \bar{z}_t\right] \leq \mathbb{E}\sum_{t=0}^{k-1} \rho^{-2t} w_t^\top B^\top P_{11} B w_t.$$

The sum telescopes, and only the $t = 0$ and $t = k - 1$ terms are left. Thus,

$$\mathbb{E}\left[\rho^{-2(k-1)} \bar{x}_k^\top P \bar{x}_k - \rho^2 \bar{x}_0^\top P \bar{x}_0 + \lambda \sum_{t=0}^{k-1} \rho^{-2t} \bar{z}_t^\top M \bar{z}_t\right] \leq \mathbb{E}\sum_{t=0}^{k-1} \rho^{-2t} w_t^\top B^\top P_{11} B w_t.$$

²The precise version of this statement (much like many classical IQC-based arguments) would instead invoke the “extended” space ℓ_{2e} but we leave those technical details aside.

Using the fact that (Ψ, M) is a ρ -hard IQC, we can drop the term dependent on M , and we are left with

$$\begin{aligned}
\mathbb{E}[\rho^{-2(k-1)}\bar{x}_k^\top P\bar{x}_k - \rho^2\bar{x}_0^\top P\bar{x}_0] &\leq \mathbb{E}\sum_{t=0}^{k-1}\rho^{-2t}w_t^\top B^\top P_{11}Bw_t \\
\implies \mathbb{E}[\bar{x}_k^\top P\bar{x}_k] &\leq \rho^{2k}\bar{x}_0^\top P\bar{x}_0 + \sum_{t=0}^{k-1}\rho^{2(k-t-1)}\mathbb{E}w_t^\top B^\top P_{11}Bw_t \\
&= \mathbb{E}[\bar{x}_k^\top P\bar{x}_k] \leq \rho^{2k}\bar{x}_0^\top P\bar{x}_0 + \sum_{t=0}^{k-1}\rho^{2(k-t-1)}\text{tr}(B^\top P_{11}B\Lambda_w) \\
&= \mathbb{E}[\bar{x}_k^\top P\bar{x}_k] \leq \rho^{2k}\bar{x}_0^\top P\bar{x}_0 + \text{tr}(B^\top P_{11}B\Lambda_w)\sum_{t=0}^{k-1}\rho^{2t},
\end{aligned}$$

where the final step reverses the order of summation. Finally, bounding the quadratic terms by the Rayleigh-Ritz inequalities gives

$$\begin{aligned}
\mathbb{E}\|\bar{x}_k\|^2 &\leq \kappa(P)\rho^{2k}\|\bar{x}_0\|^2 + \frac{\text{tr}(B^\top P_{11}B\Lambda_w)}{\lambda_{\min}(P)}\sum_{t=0}^{k-1}\rho^{2t} \\
&= \kappa(P)\rho^{2k}\|\bar{x}_0\|^2 + \frac{\text{tr}(B^\top P_{11}B\Lambda_w)}{\lambda_{\min}(P)}\frac{1-\rho^{2k}}{1-\rho^2} \\
&\leq \kappa(P)\rho^{2k}\|\bar{x}_0\|^2 + \frac{\text{tr}(B^\top P_{11}B\Lambda_w)}{(1-\rho^2)\lambda_{\min}(P)}.
\end{aligned}$$

Appendix B

Supplementary Material for Chapter 4

B.1 Details for Monte–Carlo Simulations

In all of our simulations we are faced with the following problem which we describe in some generality. Let X be a random variable distributed according to the law \mathbb{P} . We assume we have access to iid samples from \mathbb{P} . Our goal is to estimate an upper bound on $\mathbb{P}(X \geq t)$ for a fixed $t \in \mathbb{R}$.

If the law \mathbb{P} admits a density $f(\cdot)$ with respect to the Lebesgue measure, a possible solution could be to solve this problem exactly by numerically integrating

$$\int_{X(\xi) \geq t} X(\xi) f(\xi) d\xi .$$

However, numerical integration does not scale favorably with dimension. For our experiments, ξ is 75-dimensional, which is prohibitive for numerical integration.

An alternative approach to numerical integration is to rely on concentration of measure. Let X_1, \dots, X_N be i.i.d. copies of X , and let \mathbb{P}^N denote the product measure $\mathbb{P}^N = \otimes_{k=1}^N \mathbb{P}$. Using a Chernoff bound and defining $F_t := \mathbb{P}(X \geq t)$, we have

$$\mathbb{P}^N \left(\frac{1}{N} \sum_{k=1}^N \mathbf{1}_{\{X_k \geq t\}} \leq F_t - \varepsilon \right) \leq e^{-N \cdot D(F_t - \varepsilon, F_t)} , \quad (\text{B.1})$$

where $D(p, q) = p \log(p/q) + (1-p) \log((1-p)/(1-q))$ is the KL-divergence between two Bernoulli distributions. Given a $\delta \in (0, 1)$, define the random variable Q as the solution to the implicit equation

$$N \cdot D \left(\frac{1}{N} \sum_{k=1}^N \mathbf{1}_{\{X_k \geq t\}}, Q \right) = \log(1/\delta) . \quad (\text{B.2})$$

Note that, from a realization of X_1, \dots, X_N , the realization of Q from (B.2) can be solved for by numerical root finding. Plugging the definition of Q back into the Chernoff inequality (B.1), we conclude that there exists an event \mathcal{E} (in the product σ -algebra) such that on \mathcal{E} the inequality $F_t \leq Q$ holds, and furthermore $\mathbb{P}^N(\mathcal{E}) \geq 1 - \delta$. This is the methodology which we use to generate all our bounds, with $\delta = 10^{-4}$. Hence, the statements of the form “ $F_t \leq \gamma$ ” in Section 4.5.1 should be understood as operating under the assumption that our implementation of the simulation chose a particular realization which is contained in the simulator event \mathcal{E} described previously. This procedure—“inverting the Chernoff bound” is proven in the sequel.¹

B.2 Inverting the Chernoff Bound

For convenience, we will assume that all infimums in this section are finite, as the general case is even more tedious. Our starting point is as follows: for any $0 < \epsilon \leq F_t$, we have

$$\mathbb{P}^N(F_t \geq \hat{p} + \epsilon) \leq e^{-N \cdot D(F_t - \epsilon, F_t)} := H(F_t - \epsilon, F_t),$$

where $\hat{p} = \frac{1}{N} \sum_i \mathbf{1}_{\{X_i \geq t\}}$ and $F_t = \mathbb{E}\hat{p} = \mathbb{P}(X \geq t)$. Define the quantity

$$\alpha(F_t, \delta) := \inf\{\alpha : H(F_t - \alpha, F_t) > \delta, 0 < \alpha \leq F_t\}.$$

Since $H(u, v)$ is increasing in its first argument u for $u \leq v$, we have that

$$\mathbb{P}^N(F_t \geq \hat{p} + \epsilon) \leq H(F_t - \epsilon, F_t) \leq \delta,$$

for all $\epsilon \in (0, \alpha(F_t, \delta))$. Furthermore, we claim that the function $f(w) := w - \alpha(w, \delta)$ is increasing in w , for each δ . To see this, consider some small $\eta > 0$. Then:

$$\begin{aligned} f(w + \eta) &= w + \eta - \alpha(w + \eta, \delta) \\ &= w + \eta - \inf\{\alpha : H(w + \eta - \alpha, w + \eta) > \delta, 0 < \alpha \leq w + \eta\} \\ &= w - \inf\{\alpha' : H(w - \alpha', w + \eta) > \delta, -\eta < \alpha' \leq w\} \\ &\geq w - \inf\{\alpha' : H(w - \alpha', w + \eta) > \delta, 0 < \alpha' \leq w\} \\ &\geq w - \inf\{\alpha' : H(w - \alpha', w) > \delta, 0 < \alpha' \leq w\} \\ &= f(w), \end{aligned}$$

where the penultimate inequality comes from the fact that $H(u, v + \eta) \leq H(u, v)$ when $u \leq v$.

Next, define the random variable

$$Z(\hat{p}, \delta) := \inf\{z : z - \alpha(z, \delta) \geq \hat{p}\}$$

¹We thank Stephen Tu, Max Simchowitz, and Kevin Jamieson for their help in rigorously analyzing this idea.

and note that the increasing nature of $f(w)$ implies that

$$F_t \geq Z(\hat{p}, \delta) \implies F_t - \alpha(F_t, \delta) \geq Z(\hat{p}, \delta) - \alpha(Z(\hat{p}, \delta), \delta).$$

Thus,

$$\mathbb{P}(F_t \geq Z(\hat{p}, \delta)) \leq \mathbb{P}(F_t - \alpha(F_t, \delta) \geq Z(\hat{p}, \delta) - \alpha(Z(\hat{p}, \delta), \delta)).$$

However, by definition, we have that $Z(\hat{p}, \delta) - \alpha(Z(\hat{p}, \delta), \delta) \geq \hat{p}$, so

$$\begin{aligned} \mathbb{P}(F_t - \alpha(F_t, \delta) \geq Z(\hat{p}, \delta) - \alpha(Z(\hat{p}, \delta), \delta)) &\leq \mathbb{P}(F_t - \alpha(F_t, \delta) \geq \hat{p}) \\ &\leq \mathbb{P}(F_t - \epsilon \geq \hat{p}) \quad (\text{for any } 0 < \epsilon < \alpha(F_t, \delta)) \\ &\leq \delta. \end{aligned}$$

Therefore, $F_t \geq Z(\hat{p}, \delta)$ with probability at most δ , and $Z(\hat{p}, \delta)$ is a computable quantity—it is the solution z^* to the system of equations

$$\begin{aligned} z - \alpha &= \hat{p} \\ H(z - \alpha, z) &= \delta \quad (\alpha < z). \end{aligned}$$

If we let \tilde{z} be the (upper) solution to the equation $H(\hat{p}, z) = \delta$, we see that $(\tilde{z}, \tilde{z} - \hat{p})$ is a solution to these equations. Finally, we see that $F_t \leq \{z : H(\hat{p}, z) = \delta\}$ with probability at least $1 - \delta$.

Appendix C

Supplementary Material for Chapter 5

Algorithm 2: Power Method A [Rojas et al., 2012]

Input: Normalized $u^{(1)}$.

for $t = 1$ to N **do**

 Perform the experiment $y^{(t)} = Gu^{(t)} + \eta^{(t)}$.

 Create the time-reversed $\tilde{y}^{(t)}$.

$\mu^{(t)} = \|\tilde{y}^{(t)}\|_2$.

$u^{(t+1)} = \tilde{y}^{(t)} / \mu^{(t)}$.

$\hat{H}_t = \sqrt{\mu^{(t-1)}(u^{(t-1)})^\top \tilde{y}^{(t)}}$.

end

return \hat{H}_N .

Algorithm 3: Power Method B [Wahlberg et al., 2010]

Input: Normalized $u^{(1)}$.

for $t = 1$ to $N/2$ **do**

 Perform the experiment $y^{(t)} = Gu^{(t)} + \eta^{(t)}$.

 Create the time-reversed $\tilde{y}^{(t)}$.

 Perform the experiment $z^{(t)} = G\tilde{y}^{(t)} + \eta^{(t)}$.

 Create the time-reversed $\tilde{z}^{(t)}$.

$\hat{H}_t = \sqrt{|(u^{(t)})^\top \tilde{z}^{(t)}|}$.

$u^{(t+1)} = \tilde{z}^{(t)} / \|\tilde{z}^{(t)}\|_2$.

end

return $\hat{H}_{(N/2)}$.

Algorithm 4: Weighted Thompson Sampling (WTS) [Müller et al., 2017]

Input: $M, \lambda, \sigma, \rho_k^1 = 1/K \forall k, m^0 = 0, v^0 = \lambda^2 I$.

for $t = 1$ to N **do**

Input design: Create the normalized input signal $u^{(t)}$ proportional to the DFT power profile $p^t = \rho^t$.

Perform the experiment $y^{(t)} = Gu^{(t)} + \eta^{(t)}$ and obtain DFT coefficients

$$X_k^t = Y_k^t / U_k^t.$$

Update the posterior for all k :

$$m_k^{t+1} = \frac{\lambda^2 \sum_{\ell=1}^t p_k^\ell X_k^\ell}{\sigma^2 + \lambda^2 \sum_{\ell=1}^t p_k^\ell}.$$

$$v_k^{t+1} = \lambda^2 / (1 + \lambda^2 / \sigma^2 \sum_{\ell=1}^t p_k^\ell).$$

Update the posterior ρ^{t+1} :

Draw $s^l \sim \mathcal{N}_C(m_k^{t+1}, v_k^{t+1}), \quad l = 1, \dots, M$.

$$\rho_k^{t+1} = \frac{1}{M} \sum_{l=1}^M \#(\arg \max_i \{|s_i^l|\} = k).$$

$$\hat{H}_t = \max_k \sum_{\ell=1}^t p_k^\ell X_k^\ell / \sum_{\ell=1}^t p_k^\ell.$$

end

return \hat{H}_N .

Appendix D

Supplementary Material for Chapter 6

D.1 Proof of Lemma 6.4.1

Lemma 6.4.1 (Robust Equivalence). *Consider system responses $\tilde{\Theta} = \{\tilde{\mathbf{R}}, \tilde{\mathbf{M}}, \tilde{\mathbf{N}}, \tilde{\mathbf{L}}\}$ and $\hat{\Theta} = \{\hat{\mathbf{R}}, \hat{\mathbf{M}}, \hat{\mathbf{N}}, \hat{\mathbf{L}}\}$, where the latter is given by*

$$\begin{aligned}\hat{\mathbf{R}} &= (I + \Delta_1)^{-1} \tilde{\mathbf{R}} \\ \hat{\mathbf{M}} &= \tilde{\mathbf{M}} - \Delta_2 (I + \Delta_1)^{-1} \tilde{\mathbf{R}} \\ \hat{\mathbf{N}} &= (I + \Delta_1)^{-1} \tilde{\mathbf{N}} \\ \hat{\mathbf{L}} &= \tilde{\mathbf{L}} - \Delta_2 (I + \Delta_1)^{-1} \tilde{\mathbf{N}},\end{aligned}\tag{6.10}$$

where by assumption $(I + \Delta_1)^{-1}$ exists and is in \mathcal{RH}_∞ . Let $\mathbf{G} = (A, B, C, D)$ be a given plant, and consider the following statements.

- (i) $\tilde{\Theta}$ satisfies the robust SLS constraints for \mathbf{G} .
- (ii) $\hat{\Theta}$ satisfies the SLS constraints for \mathbf{G} .

Under the assumptions, (i) \implies (ii). Furthermore, let $\mathbf{G}' = (A, B, C', D)$, and let

$$\begin{aligned}\Delta_1 &= -\tilde{\mathbf{N}}(C - C') \\ \Delta_2 &= -\tilde{\mathbf{L}}(C - C').\end{aligned}$$

Then, (i) is equivalent to a third statement (iii): $\tilde{\Theta}$ satisfies the SLS constraints for \mathbf{G}' .

- (i) \iff (iii): The SLS constraints for \mathbf{G}' and the robust SLS constraints for \mathbf{G} are identical, by the definitions of Δ_1 and Δ_2 .
- (i) \implies (ii): Satisfaction of (6.9) under $(\tilde{\Theta}, \mathbf{G})$ implies satisfaction of (6.7b) under $(\hat{\Theta}, \mathbf{G})$ as they are related by a linear transformation defined by

$$V = \begin{bmatrix} (I + \Delta_1)^{-1} & 0 \\ -\Delta_2 (I + \Delta_1)^{-1} & I \end{bmatrix}.$$

We readily see that

$$V \begin{bmatrix} \tilde{\mathbf{R}} & \tilde{\mathbf{N}} \\ \tilde{\mathbf{M}} & \tilde{\mathbf{L}} \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{R}} & \hat{\mathbf{N}} \\ \hat{\mathbf{M}} & \hat{\mathbf{L}} \end{bmatrix} \quad \text{and} \quad V \begin{bmatrix} I + \mathbf{\Delta}_1 \\ \mathbf{\Delta}_2 \end{bmatrix} = \begin{bmatrix} I \\ 0 \end{bmatrix}.$$

Next, satisfaction of (6.7a) under $(\hat{\Theta}, \mathbf{G})$ is implied by satisfaction of (6.7a) and (6.9) under $(\tilde{\Theta}, \mathbf{G})$. To see this, note that multiplying (6.9) on the right by $[zI - A; -C_2]$ gives $B_2 \mathbf{\Delta}_2 = (zI - A) \mathbf{\Delta}_1$. One can then verify that

$$[zI - A \quad -B_2] \begin{bmatrix} \hat{\mathbf{R}} & \hat{\mathbf{N}} \\ \hat{\mathbf{M}} & \hat{\mathbf{L}} \end{bmatrix} = [I \quad 0].$$

Finally, consider (6.7c). The second equation of (6.9), i.e.

$$\tilde{\mathbf{M}}(zI - A) - \tilde{\mathbf{L}}C_2 = \mathbf{\Delta}_2,$$

shows that $\mathbf{\Delta}_2 \in \mathcal{RH}_\infty$. Combined with our assumption that $(I + \mathbf{\Delta}_1)^{-1} \in \mathcal{RH}_\infty$, by the definition of $\hat{\Theta}$ we may assert that $z\hat{\mathbf{R}}, z\hat{\mathbf{M}}, z\hat{\mathbf{N}}, \hat{\mathbf{L}} \in \mathcal{RH}_\infty$.

D.1.1 Proof of Lemma 6.4.3 and Corollary 6.4.1

Lemma 6.4.3. *Assume we estimate the system g by a length- r FIR system \tilde{g} using coarse-grained identification (output noise only) on m experiments, where the inputs $\mathbf{u}^{(i)}$ are constrained to lie in a unit ℓ_p^T ball. Then, with probability at least $1 - \eta$,*

$$\|\delta\|_2 \leq 2\sqrt{\log 2} \frac{\sigma^2 r^{2/\max(p,2)}}{m} \left(1 + \sqrt{2 \log \eta^{-1}}\right).$$

Recall the setup of Chapter 4 and the notation

$$\Sigma(u) := \sum_{k=1}^m \text{Toep}(\mathbf{u}^{(k)})^\top \text{Toep}(\mathbf{u}^{(k)}).$$

We readily see that $\delta \sim \mathcal{N}(0, \sigma^2 S)$ where $S = (\Sigma(u)^{-1})_{[r]}$. Then, noting that $\|\delta\|_2$ is a $\sigma \|S^{1/2}\|$ -Lipschitz function of i.i.d. standard Gaussian random variables, from concentration of measure (see Boucheron et al. [2013, Thm. 5.6]) we have that

$$P(\|\delta\|_2 \geq \mathbb{E}\|\delta\|_2 + t) \leq e^{-t^2/(2\sigma^2\|S\|)},$$

for all $t \geq 0$. Furthermore, by Jensen's inequality,

$$\mathbb{E}\|\delta\|_2 \leq \sqrt{\sum_{i=1}^r \mathbb{E}\delta_i^2} = \sigma \sqrt{\text{tr}(S)}.$$

This then gives

$$\|\delta\|_2 \leq \sigma \left(\|S\|^{1/2} \sqrt{2 \log \eta^{-1}} + \sqrt{\text{tr}(S)} \right)$$

with probability at least $1 - \eta$. To probabilistically guarantee small error, we would then like to minimize the right hand side over input signals. When \mathcal{U} is a unit ℓ_p -ball in \mathbb{R}^T , the arguments around Lemma 4.3.6 provides relevant bounds on S . Since $\|S\| \leq \text{tr}(S)$, in the most general case we have that

$$\|\delta\|_2 \leq 2\sigma \sqrt{\log 2 \frac{r^{2/\max(p,2)}}{m}} \left(1 + \sqrt{2 \log \eta^{-1}} \right)$$

with probability at least $1 - \eta$. If $p \in [1, 2]$, $\|S\|$ can be computed to be $\frac{1}{m}$, and we have

$$\|\delta\|_2 \leq 2\sigma \sqrt{\frac{r}{m}} \left(1 + \sqrt{\frac{2 \log \eta^{-1}}{r}} \right).$$

Finally, with respect to Corollary 6.4.1, as noted in Section 4.3 we have that $\delta \sim \mathcal{N}(0, \Lambda)$, where

$$\Lambda = (Z^\top Z)^{-1} Z^\top (\sigma_w^2 \text{Toep}(g) \text{Toep}(g)^\top + \sigma_\xi^2 I) Z (Z^\top Z)^{-1} \preceq (\sigma_w^2 \|\mathbf{G}\|_{\mathcal{H}_\infty}^2 + \sigma_\xi^2) (Z^\top Z)^{-1},$$

where $Z := [\text{Toep}(\mathbf{u}^{(1)})^\top \ \cdots \ \text{Toep}(\mathbf{u}^{(m)})^\top]^\top \in \mathbb{R}^{Tm \times T}$ and the inequality comes from the stability of \mathbf{G} .