

On Controller Reduction in Linear Quadratic Gaussian Control with Performance Bounds

Zhaolin Ren

School of Engineering and Applied Science, Harvard University, USA

ZHAOLINREN@G.HARVARD.EDU

Yang Zheng

ECE Department, University of California San Diego, La Jolla, USA

ZHENGY@ENG.UCSD.EDU

Maryam Fazel

ECE Department, University of Washington, Seattle, USA

MFAZEL@UW.EDU

Na Li

School of Engineering and Applied Science, Harvard University, USA

NALI@SEAS.HARVARD.EDU

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

The problem of controller reduction has a rich history in control theory. Yet, many questions remain open. In particular, there exist very few results on the order reduction of general non-observer based controllers and the subsequent quantification of the closed-loop performance. Recent developments in model-free policy optimization for Linear Quadratic Gaussian (LQG) control have highlighted the importance of this question. In this paper, we first propose a new set of sufficient conditions ensuring that a perturbed controller remains internally stabilizing. Based on this result, we illustrate how to perform order reduction of general (non-observer based) output feedback controllers using balanced truncation and modal truncation. We also provide explicit bounds on the LQG performance of the reduced-order controller. Furthermore, for single-input-single-output (SISO) systems, we introduce a new controller reduction technique by truncating *unstable* modes. We illustrate our theoretical results with numerical simulations. Our results will serve as valuable tools to design direct policy search algorithms for control problems with partial observations.

Keywords: Optimal control, model reduction, controller reduction, policy optimization

1. Introduction

In many control applications, low-order controllers are often preferred over high-order controllers, because they are simpler to maintain, more interpretable, and computationally less demanding (Anderson and Liu, 1989). Thus, given a high-order controller, one often would like to find a lower-order approximation that still stabilizes the plant whilst performing similarly on relevant closed-loop performance metrics, such as the Linear Quadratic Gaussian (LQG) cost. This problem is known as *controller reduction*. Traditional approaches to controller reduction in LQG control have focused on reducing the order of observer-based controllers and providing error bounds between the performance of the truncated controller and that of the original controller (Zhou and Chen, 1995).

However, the problem of order-reduction for general non-observer based controllers has been less studied, especially in the context of LQG control. Recent progress in model-free policy optimization for linear control has highlighted the importance of order-reduction for general output feedback controllers (Zheng et al., 2022). In particular, a natural problem in model-free policy optimization is to learn an optimal policy iteratively using policy gradient methods (Hu et al., 2022). It has recently been shown that the optimization landscape of LQG control may contain saddle points in state-space dynamic controllers (Tang et al., 2021). While vanilla policy gradient ensures the convergence to stationary points under mild assumptions, these stationary points may be saddle points that are sub-optimal. As shown recently in Zheng et al. (2022), when a saddle point corresponds to a non-minimal controller, it is possible to escape the saddle point by finding a lower-order controller and adding a suitable random perturbation during policy gradient. It is thus natural to consider

order-reduction for general dynamic controllers, such that we find a lower-order controller with approximately equivalent or lower LQG cost. Moreover, policy gradient for LQG control may meet unstable controllers¹, but the results on order-reduction for unstable controllers are far less complete (Anderson and Liu, 1989; Liu and Anderson, 1986). This motivates the main questions in this paper:

1. *Can we perform controller reduction on general, possibly unstable, LQG controllers, such that the reduced-order controller remains internally stabilizing?*
2. *Can we provide explicit error bounds on the LQG performance of the reduced-order controller compared to the original controller?*

These questions are not only relevant for the reasons relating to policy optimization of LQG control (Tang et al., 2021), but are interesting in their own right for the model and controller reduction literature (Anderson and Liu, 1989; Obinata and Anderson, 2012).

Our contributions. In this paper, we provide positive answers to both questions. We first identify a novel set of sufficient conditions that ensure the stability of a perturbed controller (Theorem 1), and then derive a new bound on the LQG cost of a perturbed controller under the assumption that the truncated component is stable and appropriately small (Theorem 2). For general multiple-input and multiple-output (MIMO) systems, building on Theorem 1 and Theorem 2, we then show (in Section 4 and Section 5.1 respectively) how balanced truncation and modal truncation may be applied to general (non observer-based, possibly unstable) LQG controllers to yield lower-order controllers with bounded LQG performance gap (compared to that of the original controller). Furthermore, for single-input single-output (SISO) systems, we discuss in Section 5.2 how internal stability may be preserved even when the reduced-order controller has fewer unstable poles than the original controller. This opens the path of controller reduction via truncating unstable poles, a novel controller reduction technique that we illustrate both theoretically and empirically.

Related work. A common approach to controller reduction is to truncate the stable part of a controller, whilst keeping its unstable part intact (Anderson and Liu, 1989). Popular methods to perform truncation of the stable part of a controller include modal truncation (Skelton, 1988), balanced truncation (Enns, 1984), and Hankel norm approximation (Glover and Limebeer, 1983). When the difference of the stable portion and its truncated portion satisfies a (frequency-weighted) error bound (see Lemma 3), it guarantees that the truncated controller remains internally stabilizing. However, there appear to be no existing results providing an error bound on the LQG cost of the truncated controller from such a procedure for general (possibly non observer-based) controllers. In contrast, for observer-based controllers, there has been a significant line of work based on coprime factorizations (Vidyasagar, 1975), which not only yields reduced-order controllers that are internally stabilizing but also guarantees LQG performance bounds for the resulting truncated controllers (cf. Liu and Anderson (1986); Anderson and Liu (1989)). However, these methods only work for observer-based controllers.

A closely related but distinct research direction to controller reduction is the topic of open-loop model reduction (Obinata and Anderson, 2012; Benner et al., 2017; Antoulas et al., 2001; Antoulas, 2005). For space reasons, we defer discussion of model reduction to our full report (Ren et al., 2022). Another important related topic is policy optimization for linear control problems. There has been significant recent work studying policy optimization for linear quadratic control problems, for linear-quadratic-regulator (LQR) (Fazel et al., 2018), $\mathcal{H}_2$ linear control with $\mathcal{H}_\infty$ guarantees (Zhang et al., 2020; Hu and Zheng, 2023), as well as LQG problems (Tang et al., 2021; Zheng et al., 2022). In particular, as we explained earlier, the considerations outlined in Zheng et al. (2022) on escaping saddle points of the LQG problem was an important motivation for our work, where controller reduction is required. See Hu et al. (2022) for a recent review.

1. A dynamic controller that has unstable modes itself but internally stabilizes the plant.

Paper outline. The rest of this paper is structured as follows. We present the problem statement in Section 2. In Section 3, we first introduce Theorem 1, which provides sufficient conditions such that a perturbed controller $\mathbf{K}_r$ of $\mathbf{K}$ is internally stabilizing. We further derive an upper bound on the LQG cost $J(\mathbf{K}_r)$ (Theorem 2). In Section 4, we study balanced truncation on the stable part of a controller. In Section 5, we discuss modal truncation, where Section 5.1 studies modal truncation on the stable part of a controller, and Section 5.2 discusses controller reduction via truncation of unstable poles for SISO systems. Finally, we end with numerical experiments illustrating our theoretical results in Section 6. Many technical proofs are postponed to our full report (Ren et al., 2022).

Notation: We denote the set of real-rational proper stable transfer functions as $\mathcal{RH}_\infty$ (i.e., all the poles are on the open left-half complex plane). For simplicity, we omit the dimension of transfer matrices. We define the $\mathcal{L}_\infty$ norm for a transfer function $\mathbf{G}(s)$ as $\|\mathbf{G}\|_{\mathcal{L}_\infty} := \sup_{w \in \mathbb{R}} \sigma_{\max}(\mathbf{G}(jw))$, where $\sigma_{\max}(\cdot)$ denotes the maximum singular value. When $\mathbf{G} \in \mathcal{RH}_\infty$, its $\mathcal{H}_\infty$ norm is the same as its $\mathcal{L}_\infty$ norm (Zhou et al., 1996, Chapter 4.3). We define the $\mathcal{L}_2$ norm for $\mathbf{G}(s)$ as $\|\mathbf{G}\|_{\mathcal{L}_2} := \sqrt{\frac{1}{2\pi} \int_{-\infty}^{\infty} \text{Tr}(\mathbf{G}(-jw)^\top \mathbf{G}(jw)) dw}$. When $\mathbf{G}$ is stable and strictly proper, its $\mathcal{H}_2$ norm is the same as its $\mathcal{L}_2$ norm (Zhou et al., 1996, Chapter 4.3).

2. Preliminaries and Problem Statement

2.1. General non-observer based controllers

Consider a strictly proper linear time-invariant (LTI) plant²

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Bu(t) + Bw(t), \\ y(t) &= Cx + v(t), \end{aligned} \tag{1}$$

where $x(t) \in \mathbb{R}^n$, $u(t) \in \mathbb{R}^m$, $y(t) \in \mathbb{R}^p$ are the state vector, control action, and measurement vector at time t , respectively; $w(t) \in \mathbb{R}^m$ and $v(t) \in \mathbb{R}^p$ are external disturbances on the state and measurement vectors at time t , respectively. One basic yet fundamental control problem is to design a feedback controller (or policy) to stabilize the plant (1). A standard approach for this problem is to use an observer-based controller of the form

$$\begin{aligned} \dot{\xi}(t) &= A\xi(t) + Bu(t) + L(y(t) - C\xi(t)) \\ u(t) &= -K\xi(t), \end{aligned} \tag{2}$$

where $\xi(t) \in \mathbb{R}^n$ is an estimated state, $L \in \mathbb{R}^{n \times p}$ is an observer gain, and $K \in \mathbb{R}^{m \times n}$ is a feedback gain. The observer and feedback gains are chosen such that $A - LC$ and $A - BK$ are stable, and this guarantees the closed-loop internal stability when applying the controller (2) to the plant (1) (Zhou et al., 1996, Chapter 3.5). Note that the order of this observer-based controller must be the same as the system plant (i.e., the controller state $\xi(t)$ and the system state $x(t)$ have the same dimension). Order reduction for controllers in the form of (2) is discussed in Liu and Anderson (1986); Anderson and Liu (1989); Zhou and Chen (1995). In this paper, we consider a general non-observer based dynamic controller of the form

$$\begin{aligned} \dot{\xi}(t) &= A_K \xi(t) + B_K y(t), \\ u(t) &= C_K \xi(t), \end{aligned} \tag{3}$$

where $\xi(t) \in \mathbb{R}^q$ is the internal state of the controller, and A_K, B_K, C_K are matrices of proper dimensions that specify the dynamics of the controller. The dimension q of the internal control variable ξ is called the order of the dynamic controller (3). The controller in (3) is more suitable for model-free policy optimization as it does not explicitly depend on the system dynamics (Tang et al.,

2. For simplicity, we assume that there is a common B matrix in front of both the input term $u(t)$ and noise term $w(t)$.

2021; Zheng et al., 2022). It is clear that the observer-based controller (2) is a special case of (3) by taking $q = n$, and $A_K = A - BK - LC$, $B_K = L$, $C_K = -K$. By combining (3) with (1), the closed-loop system is internally stable if and only if the following closed-loop matrix is stable (Zhou et al., 1996, Lemma 5.2):

$$A_{\text{cl}} := \begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix}. \quad (4)$$

2.2. Problem statement

Given an internally stabilizing controller (A_K, B_K, C_K) , the *controller reduction problem* is to find a new controller $(\hat{A}_K, \hat{B}_K, \hat{C}_K)$ of lower order $\hat{q} < q$ such that it still internally stabilizes the plant and does not significantly affect the closed-loop performance. In particular, we consider a normalized LQG control performance (Mustafa and Glover, 1991; Jonckheere and Silverman, 1983), defined as

$$J = \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \int_0^T (Cx(t))^T (Cx(t)) + u(t)^T u(t) dt \right], \quad (5)$$

We make the following two assumptions.

Assumption 1 *The plant (1) is minimal, i.e., (A, B) is controllable and (C, A) is observable.*

Assumption 2 *In plant (1), the signals $w(t) \in \mathbb{R}^m$ and $v(t) \in \mathbb{R}^p$ are zero mean Gaussian white noise, each with a spectrum equal to the identity.*

Assumption 1 is standard and guarantees the existence of internally stabilizing controllers³. If the plant is not minimal, we can always perform a lower-order minimal realization before designing a controller. Assumption 2 was used to define the *normalized LQG control problem* in (Mustafa and Glover, 1991; Jonckheere and Silverman, 1983). As we shall see next, this assumption simplifies the expression of LQG cost (5) in the frequency domain. For the controller reduction problem, it may not be easy to work directly with the internal stability condition (4) in the state-space domain due to non-uniqueness of state-space realizations. It is more convenient to consider equivalent conditions in the frequency domain. In particular, the controller (3) can be represented as a transfer function $\mathbf{K} := C_K(sI - A_K)^{-1} B_K$. Let us define a new performance signal $\tilde{y} = Cx$. Some simple manipulations show that the closed-loop transfer function from (w, v) to $(\tilde{y}, u)$ is

$$\begin{bmatrix} \tilde{y} \\ u \end{bmatrix} = \begin{bmatrix} (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G} & (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}\mathbf{K} \\ \mathbf{K}(I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G} & \mathbf{K}(I - \mathbf{G}\mathbf{K})^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{w} \\ \mathbf{v} \end{bmatrix}, \quad (6)$$

where $\mathbf{G}(s) = C(sI - A)^{-1}B$. Then, we have the following condition for internal stability.

Lemma 1 ((Zhou et al., 1996, Lemma 5.3)) *The controller $\mathbf{K}$ in (3) internally stabilizes the plant (1) if and only if the closed-loop transfer function from (w, v) to $(\tilde{y}, u)$ is stable⁴, i.e.,*

$$\begin{bmatrix} (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G} & (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}\mathbf{K} \\ \mathbf{K}(I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G} & \mathbf{K}(I - \mathbf{G}\mathbf{K})^{-1} \end{bmatrix} \in \mathcal{RH}_\infty$$

Furthermore, the normalized LQG control cost (5) can be expressed conveniently in the frequency domain. For completeness, we provide a proof of Lemma 2 in our full report (Ren et al., 2022).

Lemma 2 *Under Assumption 2, given an internally stabilizing controller $\mathbf{K}$ in (3), the normalized LQG cost (5) can be expressed as follows*

$$J(\mathbf{K}) = \left\| \begin{bmatrix} (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G} & (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}\mathbf{K} \\ \mathbf{K}(I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G} & \mathbf{K}(I - \mathbf{G}\mathbf{K})^{-1} \end{bmatrix} \right\|_{\mathcal{H}_2}^2.$$

3. The existence of internally stabilizing controllers only requires stabilizability and detectability.

4. The standard result in (Zhou et al., 1996, Lemma 5.3) uses a slightly different set of closed-loop transfer functions. Simple manipulations via $(I - \mathbf{G}\mathbf{K})^{-1} = I + (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}\mathbf{K}$ can show the equivalence.

3. Robust stability and LQG performance

Our main goal is to properly perturb the controller $\mathbf{K}$ to get a lower-order controller $\mathbf{K}_r$ such that the closed-loop performance remains similar. In this section, we present two technical results that underpin our controller reduction results in Sections 4 and 5: 1) a new robust stability result (Theorem 1), and 2) an upper bound on the LQG performance for the new controller $\mathbf{K}_r$ (Theorem 2).

3.1. A novel sufficient condition for internal stability

Classical results on controller reduction often study the case when the truncated controller $\mathbf{K}_r$ has the same number of unstable poles as the original controller $\mathbf{K}$ (cf. [Liu and Anderson \(1986\)](#)). In particular, a widely-used condition is as follows.

Lemma 3 ([\(Anderson and Liu, 1989, Section II.A\)](#)) *Let $\mathbf{G}$ be the transfer function of an LTI plant (1), the controller $\mathbf{K}$ (3) internally stabilize the plant, and $\mathbf{K}_r$ be another controller. Denote $\Delta := \mathbf{K}_r - \mathbf{K}$. If*

1. $\mathbf{K}$ and $\mathbf{K}_r$ have the same number of poles in $\text{Re}(s) > 0$, and no poles on the imaginary axis,
2. either $\|\Delta \mathbf{G}(I - \mathbf{K}\mathbf{G})^{-1}\|_{\mathcal{L}_\infty} < 1$ or $\|(I - \mathbf{G}\mathbf{K})^{-1} \mathbf{G} \Delta\|_{\mathcal{L}_\infty} < 1$,

then $\mathbf{K}_r$ also internally stabilizes the plant (1).

This classical result underpins many controller reduction techniques in the literature; see [Anderson and Liu \(1989\)](#) for a review. Lemma 3 can be proved via Nyquist stability criterion; see the discussions in [\(Doyle and Stein, 1981, Section IV\)](#). If the controller $\mathbf{K}$ is stable in the first place, then the first condition in Lemma 3 can be naturally satisfied by using any stable controller $\mathbf{K}_r$. When the controller $\mathbf{K}$ is unstable (i.e., A_K in (3) is unstable), we might always need to preserve the unstable part in $\mathbf{K}$ in order to use Lemma 3. However, it is unclear if it is necessary for a lower-order truncated controller $\mathbf{K}_r$ to have the same number of unstable poles as $\mathbf{K}$ in order to maintain closed-loop stability. In this section, we provide a novel set of sufficient conditions in Theorem 1 ensuring that $\mathbf{K}_r$ is still stabilizing, which makes no explicit assumptions on whether $\mathbf{K}_r$ and $\mathbf{K}$ have the same number of unstable poles. This technical result may be of independent interest.

Theorem 1 *Let $\mathbf{G}$ be the transfer function of an LTI plant (1), the controller $\mathbf{K}$ (3) internally stabilize the plant, and let $\mathbf{K}_r$ denote another controller. Denote $\Delta := \mathbf{K}_r - \mathbf{K}$. If $\Delta(I - \mathbf{G}\mathbf{K})^{-1}$ is stable and*

$$\max \left\{ \|(I - \mathbf{G}\mathbf{K})^{-1} \mathbf{G} \Delta\|_{\mathcal{H}_\infty}, \|\Delta(I - \mathbf{G}\mathbf{K})^{-1} \mathbf{G}\|_{\mathcal{H}_\infty} \right\} < 1, \quad (7)$$

then, $\mathbf{K}_r$ also internally stabilizes $\mathbf{G}$.

Unlike the proof of Lemma 3 that is based on Nyquist stability [\(Doyle and Stein, 1981, Section IV\)](#), Theorem 1 can be proved directly from Lemma 1. If we can show

$$\begin{bmatrix} (I - \mathbf{G}\mathbf{K}_r)^{-1} \mathbf{G} & (I - \mathbf{G}\mathbf{K}_r)^{-1} \mathbf{G} \mathbf{K}_r \\ \mathbf{K}_r (I - \mathbf{G}\mathbf{K}_r)^{-1} \mathbf{G} & \mathbf{K}_r (I - \mathbf{G}\mathbf{K}_r)^{-1} \end{bmatrix} \in \mathcal{RH}_\infty, \quad (8)$$

Lemma 1 confirms that $\mathbf{K}_r$ internally stabilizes $\mathbf{G}$. Since $\mathbf{K}$ internally stabilizes $\mathbf{G}$, we know that

$$\begin{bmatrix} (I - \mathbf{G}\mathbf{K})^{-1} \mathbf{G} & (I - \mathbf{G}\mathbf{K})^{-1} \mathbf{G} \mathbf{K} \\ \mathbf{K} (I - \mathbf{G}\mathbf{K})^{-1} \mathbf{G} & \mathbf{K} (I - \mathbf{G}\mathbf{K})^{-1} \end{bmatrix} \in \mathcal{RH}_\infty \quad (9)$$

Motivated by [\(Zheng et al., 2021, Appendix C\)](#), the key idea in our proof is to relate the transfer functions in (8) with those in (9), and then to show that all four subblocks in (8) are stable. For example, it is not difficult to verify

$$(I - \mathbf{G}\mathbf{K}_r)^{-1} \mathbf{G} = (I - \mathbf{X}\Delta)^{-1} (I - \mathbf{G}\mathbf{K})^{-1} \mathbf{G}, \quad (10)$$

where $\mathbf{X} := (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}$. The condition (7) is equivalent to $\|\mathbf{X}\Delta\|_{\mathcal{H}_\infty} < 1$, $\|\Delta\mathbf{X}\|_{\mathcal{H}_\infty} < 1$. According to the small-gain theorem (Zhou et al., 1996, Theorem 9.1), we know that both $(I - \mathbf{X}\Delta)^{-1}$ and $(I - \Delta\mathbf{X})^{-1}$ exist and they are stable. Thus, we know that $(I - \mathbf{G}\mathbf{K}_r)^{-1}\mathbf{G}$ is stable. Similarly, we can show all four subblocks in (8) are stable under the conditions in Theorem 1. For space reasons, we defer the full proof, as well as a remark comparing Lemma 3 and Theorem 1, to our full report (Ren et al., 2022).

3.2. A new bound on the perturbed LQG cost

Theorem 1 presents sufficient conditions to guarantee the closed-loop stability using the reduced-order controller $\mathbf{K}_r$. In many situations (such as policy optimization for LQG in Tang et al. (2021); Zheng et al. (2022)), we also need to understand the closed-loop performance under this new controller $\mathbf{K}_r$. Our next technical result show that if the error $\Delta := \mathbf{K}_r - \mathbf{K}$ is stable with an appropriately bounded $\mathcal{H}_\infty$ norm, the change of the LQG cost (5) can also be bounded. The proof builds on the analysis techniques in Theorem 1, and we defer the proof to our full report (Ren et al., 2022).

Theorem 2 *Let $\mathbf{G}$ be the transfer function of an LTI plant (1), the controller $\mathbf{K}$ (3) internally stabilize the plant, and $\mathbf{K}_r$ denote another controller. Denote $\Delta := \mathbf{K}_r - \mathbf{K}$. If*

$$\|\Delta\|_{\mathcal{H}_\infty} < \frac{1}{\|(I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}\|_{\mathcal{H}_\infty}}, \quad (11)$$

then the controller $\mathbf{K}_r$ internally stabilizes the plant (1), and the resulting LQG cost (5) satisfies

$$J(\mathbf{K}_r) \leq \frac{1}{(1 - \|\mathbf{X}\|_{\mathcal{H}_\infty}\|\Delta\|_{\mathcal{H}_\infty})^2} (J(\mathbf{K}) + S_1 + S_2), \quad (12)$$

where with the notation $\mathbf{Y} := (I - \mathbf{G}\mathbf{K})^{-1}$, and $\mathbf{X} := (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}$, we have

$$\begin{aligned} S_1 &:= 2\|\Delta\|_{\mathcal{H}_\infty}\|\mathbf{X}\|_{\mathcal{H}_2}(\|\mathbf{X}\mathbf{K}\|_{\mathcal{H}_2}) + 2\|\Delta\|_{\mathcal{H}_2}(\|\mathbf{K}\mathbf{Y}\|_{\mathcal{H}_2}\|\mathbf{Y}\|_{\mathcal{H}_\infty} + \|\mathbf{K}\mathbf{X}\|_{\mathcal{H}_2}\|\mathbf{X}\|_{\mathcal{H}_\infty})(1 + \|\mathbf{K}\mathbf{X}\|_{\mathcal{H}_\infty}), \\ S_2 &:= \|\Delta\|_{\mathcal{H}_\infty}^2\|\mathbf{X}\|_{\mathcal{H}_2}^2 + \|\Delta\|_{\mathcal{H}_2}^2(\|\mathbf{Y}\|_{\mathcal{H}_\infty}^2 + \|\mathbf{X}\|_{\mathcal{H}_\infty}^2)(1 + \|\mathbf{K}\mathbf{X}\|_{\mathcal{H}_\infty})^2. \end{aligned} \quad (13)$$

Theorem 2 shows that as long as the truncation error is bounded as in (11), the reduced-order controller $\mathbf{K}_r$ still internally stabilizes the plant. Furthermore, the upper bound (12) implies that

$$\left| \frac{J(\mathbf{K}_r) - J(\mathbf{K})}{J(\mathbf{K})} \right| \leq \mathcal{O}(\|\Delta\|_{\mathcal{H}_\infty}).$$

When the truncation error is small (measured by $\|\Delta\|_{\mathcal{H}_\infty}$), the change of LQG cost is also small. Similar bounds like (12) seem to be less studied for general non-observer based controllers in the literature (Anderson and Liu, 1989). Most existing bounds assume an observed-based controller in (2) (cf. Liu and Anderson (1986); Anderson and Liu (1989); Zhou and Chen (1995)), and most of the techniques therein rely on coprime factorization (Vidyasagar, 1975). Theorem 1 and Theorem 2 work for any perturbed controller $\mathbf{K}_r$ satisfying the assumptions therein. In the next two sections, we show how to use balanced and modal truncation strategies to derive suitable reduced-order controllers $\mathbf{K}_r$.

4. Controller reduction via balanced truncation

In this section, we discuss controller reduction strategies using balanced truncation and apply Theorem 1 and Theorem 2 to derive stability and performance guarantees.

4.1. Balanced truncation

We first recall that for asymptotically stable transfer functions, under appropriate assumptions, a reduced-order transfer function resulting from balanced truncation is also asymptotically stable.

Lemma 4 ((Pernebo and Silverman, 1982, Theorem 3.2), (Antoulas, 2005, Theorem 7.9)) *Let $\mathbf{P}$ be an asymptotically stable transfer function with a balanced minimal state-space realization⁵.*

$$\mathbf{P} = \left[\begin{array}{cc|c} A_{11} & A_{12} & B_1 \\ A_{21} & A_{22} & B_2 \\ \hline C_1 & C_2 & D \end{array} \right], \text{ s.t. } \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix} + \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}^\top = - \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} \begin{bmatrix} B_1 \\ B_2 \end{bmatrix}^\top$$

$$\begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} + \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}^\top \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix} = - \begin{bmatrix} C_1^\top \\ C_2^\top \end{bmatrix} \begin{bmatrix} C_1^\top \\ C_2^\top \end{bmatrix}^\top$$

where $\Sigma_i \succ 0$ is a positive-definite diagonal matrix for each $i \in \{1, 2\}$. If Σ_1 and Σ_2 share no eigenvalues in common, then A_{11} and A_{22} are both asymptotically stable. Furthermore,

$$\|\mathbf{P} - \mathbf{P}_r\|_{\mathcal{H}_\infty} \leq 2\text{trace}(\Sigma_2), \text{ where } \mathbf{P}_r = \left[\begin{array}{c|c} A_{11} & B_1 \\ \hline C_1 & D \end{array} \right]. \quad (14)$$

4.2. Controller reduction

In general, the dynamic controller $\mathbf{K}$ is not stable itself, i.e., A_K in (3) has unstable eigenvalues. The standard balanced truncation procedure cannot be applied to unstable systems directly. Our strategy is to divide the controller $\mathbf{K}$ into a stable part and unstable part

$$\mathbf{K} = \mathbf{K}_< + \mathbf{K}_\geq, \quad (15)$$

where $\mathbf{K}_<$ of order n_1 contains all stable poles (i.e., those on the open left-half plane) and $\mathbf{K}_\geq$ of order n_2 contains the remaining poles (i.e., those on the closed right-half plane), and $n_1 + n_2 = n$. In this section, we assume the controller contains at least one stable pole ($n_1 \geq 1$).

The separation (15) is always possible by computing the Jordan normal form of A_K such that

$$A_K = Q_K \hat{A}_K Q_K^{-1}, \quad \text{with } \hat{A}_K = \begin{bmatrix} \hat{A}_{K,<} & 0 \\ 0 & \hat{A}_{K,\geq} \end{bmatrix}, \quad (16)$$

where $Q_K \in \mathbb{R}^{n \times n}$ is an invertible coordinate transformation⁶, the eigenvalues of $\hat{A}_{K,<} \in \mathbb{R}^{n_1 \times n_1}$ are in the open left-half plane, and the eigenvalues of $\hat{A}_{K,\geq} \in \mathbb{R}^{n_2 \times n_2}$ are in the closed right-half plane, and $n_1 + n_2 = n$. Therefore, the stable and unstable parts in (15) can be expressed as

$$\mathbf{K}_< = \left[\begin{array}{c|c} \hat{A}_{K,<} & \hat{B}_{K,<} \\ \hline \hat{C}_{K,<} & 0 \end{array} \right], \quad \mathbf{K}_\geq = \left[\begin{array}{c|c} \hat{A}_{K,\geq} & \hat{B}_{K,\geq} \\ \hline \hat{C}_{K,\geq} & 0 \end{array} \right], \quad (17)$$

where $\hat{C}_K := C_K Q_K$ and $\hat{B}_K := Q_K^{-1} B_K$ are partitioned into $\hat{C}_K = [\hat{C}_{K,<} \quad \hat{C}_{K,\geq}]$, $\hat{B}_K = \begin{bmatrix} \hat{B}_{K,<} \\ \hat{B}_{K,\geq} \end{bmatrix}$

with $\hat{C}_{K,<} \in \mathbb{R}^{m \times n_1}$, $\hat{C}_{K,\geq} \in \mathbb{R}^{m \times (n-n_1)}$, and $\hat{B}_{K,<} \in \mathbb{R}^{n_1 \times p}$, $\hat{B}_{K,\geq} \in \mathbb{R}^{(n-n_1) \times p}$. We can then perform a balanced truncation on the stable part $\mathbf{K}_<$ and get a reduced-order controller $\mathbf{K}_{<,r} = \left[\begin{array}{c|c} \tilde{A}_{<,11} & \tilde{B}_{<,1} \\ \hline \tilde{C}_{<,1} & 0 \end{array} \right]$, where the order is $n_r < n_1$. The final reduced-order controller becomes

$$\mathbf{K}_r = \mathbf{K}_{<,r} + \mathbf{K}_\geq = \left[\begin{array}{cc|c} \tilde{A}_{<,11} & 0 & \tilde{B}_{<,1} \\ 0 & \hat{A}_{K,\geq} & \hat{B}_{K,\geq} \\ \hline \tilde{C}_{<,1} & \hat{C}_{K,\geq} & 0 \end{array} \right] \quad (18)$$

which has order $r := n_r + n_2 < n$. For space reasons, we omit a summary of this procedure here, and summarize it instead in algorithmic form in our full report (Ren et al., 2022). Based on Theorems 1 and 2, under appropriate conditions, the resulting controller $\mathbf{K}_r$ in (18) remains a stabilizing controller and has a similar LQG cost compared to the original controller $\mathbf{K}$.

5. Any stable transfer function has a balanced minimal realization. For completeness, we review this in (Ren et al., 2022)

6. Since any real-valued matrix can be expressed in a Jordan canonical form, such a transformation Q_K always exists.

Algorithm 1 Modal truncation

Require: 1) A controller $\mathbf{K}$ with a minimal order- n state-space realization $\mathbf{K} = \left[\begin{array}{c|c} A_K & B_K \\ \hline C_K & 0 \end{array} \right]$; 2)

the order reduction parameter r_{red} (a positive integer less than k)

- 1: Convert A_K into the standard Jordan normal form $A_K = \text{diag}(A_1, \dots, A_k)$, where each A_i is a Jordan block of order n_i , and $\sum_{i=1}^k n_i = n$. Let λ_i denote the eigenvalue associated with A_i .
- 2: Decompose $\mathbf{K}$ as $\mathbf{K}(s) = \sum_{i=1}^k C_i(sI - A_i)^{-1} B_i$.
- 3: For each i , compute

$$d_i = \begin{cases} \|C_i(sI - A_i)^{-1} B_i\|_{\mathcal{H}_\infty} & \text{if } \lambda_i < 0 \\ \|C_i A_i^{-1} B_i\|_2 & \text{if } \lambda_i > 0. \end{cases} \quad (20)$$

- 4: Let o_i be the ranking of i based on $\{d_i\}_{i=1}^k$, i.e. $o_i := j$ if d_i is the j -th smallest value in $\{d_i\}_{i=1}^k$.
- 5: Set $\Delta = \sum_{i \in [k], o_i \leq r_{\text{red}}} C_i(sI - A_i)^{-1} B_i$.
- 6: **return** the reduced order controller $\mathbf{K}_r := \mathbf{K} - \Delta$.

Corollary 1 Consider a minimal n -th order controller $\mathbf{K}$ which stabilizes the plant $\mathbf{G}$. Suppose we obtain an r -th order controller $\mathbf{K}_r$ where $r < n$, such that $\mathbf{K}_r = \mathbf{K}_{<,r} + \mathbf{K}_{\geq}$, where $\mathbf{K}_{<,r}$ is a lower-order balanced truncation of $\mathbf{K}_{<}$. Suppose that $\Sigma_{<,1}$ and $\Sigma_{<,2}$ in the balanced truncation of $\mathbf{K}_{<}$ share no eigenvalues. If

$$\sigma_{n_r+1} + \dots + \sigma_{n_1} < \frac{1}{2\|(I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}\|_{\mathcal{H}_\infty}}, \quad (19)$$

where $\sigma_{n_r+1}, \dots, \sigma_{n_1}$ are the diagonal elements of $\Sigma_{<,2}$ in the balanced truncation of $\mathbf{K}_{<}$, then the reduced-order controller $\mathbf{K}_r$ internally stabilizes (1), and the resulting LQG cost (5) satisfies

$$J(\mathbf{K}_r) \leq \frac{1}{(1 - \|\mathbf{X}\|_{\mathcal{H}_\infty} \|\Delta\|_{\mathcal{H}_\infty})^2} (J(\mathbf{K}) + S_1 + S_2),$$

where S_1 and S_2 are defined in (13).

The proof is based on the analysis in Theorem 2. The only difference is that Corollary 1 truncates the singular values in the stable part $\mathbf{K}_r$ and imposes the condition (19), which is the same as condition (11) in Theorem 2 when applying the bound from (14) in Lemma 4.

5. Controller reduction via modal truncation

In this section, we proceed to discuss controller reduction by modal truncation, which may apply to the truncation of either stable or unstable component(s) in a controller. In particular, we first apply modal truncation on the stable part of a controller in Section 5.1, and then discuss the performance of modal truncation on possibly unstable component(s) for SISO systems in Section 5.2.

5.1. Modal truncation on stable component(s)

The basic idea of *modal truncation* begins with writing the controller $\mathbf{K}$ (3) into $\mathbf{K}(s) = \sum_{i=1}^k C_i(sI - A_i)^{-1} B_i$, where A_i contains a mode corresponding to an eigenvalue of λ_i in $\mathbf{K}$. We then directly remove some modes that are less significant according the criterion defined in (20). The detailed steps are listed in Algorithm 1. As a counterpart to balanced truncation, we can derive upper bounds on the LQG cost when performing modal truncation on the stable part of a controller $\mathbf{K}$.

Corollary 2 Consider a minimal n -th order controller $\mathbf{K}$ which stabilizes the plant $\mathbf{G}$. Consider the decomposition $\mathbf{K} = \mathbf{K}_{<} + \mathbf{K}_{\geq}$ in (15), and suppose we obtain a lower-order approximation $\mathbf{K}_{r,<}$ of $\mathbf{K}_{<}$ using the modal truncation algorithm in Algorithm 1. Let $\mathbf{K}_r := \mathbf{K}_{r,<} + \mathbf{K}_{\geq}$. Denote $\Delta = \mathbf{K}_{<} - \mathbf{K}_{r,<}$. Suppose that

$$\|\Delta\|_{\mathcal{H}_\infty} < \frac{1}{\|(I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}\|_{\mathcal{H}_\infty}}. \quad (21)$$

Then, we have that $\mathbf{K}_r$ internally stabilizes $\mathbf{G}$, and the resulting LQG cost (5) satisfies

$$J(\mathbf{K}_r) \leq \frac{1}{(1 - \|\mathbf{X}\|_{\mathcal{H}_\infty} \|\Delta\|_{\mathcal{H}_\infty})^2} (J(\mathbf{K}) + S_1 + S_2),$$

where S_1 and S_2 are defined in (13).

Since (21) holds, the condition in (7) holds, and it follows from Theorem 1 that $\mathbf{K}_r$ internally stabilizes $\mathbf{G}$. Then, the same calculations in Theorem 2 can show the upper bound on $J(\mathbf{K}_r)$.

5.2. Modal truncation on unstable component(s)

Next, we introduce the following result, which studies the LQG cost change when truncating unstable mode(s) of a controller $\mathbf{K}$, for single-input single-output (SISO) systems.

Theorem 3 (Order reduction of unstable SISO controllers) Consider a minimal n -th order controller $\mathbf{K}$ which stabilizes the plant $\mathbf{G}$. Suppose both $\mathbf{G}(s)$ and $\mathbf{K}(s)$ are univariate rational polynomial functions (SISO systems). Suppose we obtain an r -order controller $\mathbf{K}_r$ via Algorithm 1, where $r < n$, where we denote $\Delta = \mathbf{K} - \mathbf{K}_r$ as in Algorithm 1; note that Δ may be unstable. Suppose Δ has no unstable mode at 0, and

$$\left\| \frac{1}{1 - (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}\Delta} \right\|_{\mathcal{H}_\infty} < \infty, \quad (22)$$

Then, we have that $\mathbf{K}_r$ internally stabilizes $\mathbf{G}$, and

$$J(\mathbf{K}_r) \leq \|(1 - (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}\Delta)^{-1}\|_{\mathcal{H}_\infty}^2 (J(\mathbf{K}) + S_1 + S_2),$$

where with the notation $\mathbf{Y} := (I - \mathbf{G}\mathbf{K})^{-1}$, $\mathbf{X} := (I - \mathbf{G}\mathbf{K})^{-1}\mathbf{G}$, we have

$$\begin{aligned} S_1 &:= 2\|\Delta\|_{L_\infty} \|\mathbf{X}\|_{\mathcal{H}_2} \|\mathbf{X}\mathbf{K}\|_{\mathcal{H}_2} + 2\|\mathbf{K}\mathbf{Y}\|_{\mathcal{H}_2} (\|\Delta\|_{L_2} \|\mathbf{Y}\|_{\mathcal{H}_\infty}) \\ &\quad + 2\|\mathbf{K}\mathbf{Y}\|_{\mathcal{H}_2} (\|\Delta\|_{L_\infty} (\|\mathbf{Y}\|_{\mathcal{H}_2} + \|\mathbf{K}\mathbf{X}\|_{\mathcal{H}_2} \|\mathbf{Y}\|_{\mathcal{H}_\infty} + \|\mathbf{K}\mathbf{X}\|_{\mathcal{H}_\infty} \|\mathbf{X}\|_{\mathcal{H}_2})) \\ S_2 &:= \|\Delta\|_{L_\infty}^2 \|\mathbf{X}\|_{\mathcal{H}_2}^2 + (\|\mathbf{Y}\|_{\mathcal{H}_\infty} (\|\Delta\|_{L_\infty} \|\mathbf{K}\mathbf{X}\|_{\mathcal{H}_2} + \|\Delta\|_{L_2}))^2 + \|\mathbf{X}\|_{\mathcal{H}_2}^2 (\|\Delta\|_{L_\infty} \|\mathbf{K}\mathbf{X}\|_{\mathcal{H}_\infty} + \|\Delta\|_{L_\infty})^2. \end{aligned}$$

The proof is similar to those of Theorems 1 and 2 but the details are technically involved. We defer the details as well as a discussion of benefits and limitations to our full report (Ren et al., 2022).

6. Numerical Examples

For space reasons, we defer more extensive simulations, including a comparison of balanced truncation and modal truncation, to the full report (Ren et al., 2022).

6.1. Scaling effect of $\|\Delta\|_{\mathcal{H}_\infty}$.

We study how the performance gap behaves as a function of the size of the truncated component. For this analysis, we (randomly) generate five stable and minimal SISO systems of order 4 (which we denote as $\mathbf{K}_r^{(1)}, \mathbf{K}_r^{(2)}, \mathbf{K}_r^{(3)}, \mathbf{K}_r^{(4)}, \mathbf{K}_r^{(5)}$) and augment the system by adding a stable mode Δ , where

$$\Delta = \begin{bmatrix} -1 & \sqrt{\epsilon} \\ \sqrt{\epsilon} & 0 \end{bmatrix}.$$

For each $i \in [5]$, we denote the augmented system as $\mathbf{K}^{(i)} := \mathbf{K}_r^{(i)} + \Delta$. We then generate a stabilizing controller, $\mathbf{G}^{(i)}$, which stabilizes $\mathbf{K}^{(i)}$. Viewing $\mathbf{G}^{(i)}$ as the system plant, we then compare the LQG cost of $\mathbf{K}^{(i)}$ and $\mathbf{K}_r^{(i)}$ on the system $\mathbf{G}^{(i)}$, as the $\mathcal{H}_\infty$ norm of the truncated component, i.e. $\|\Delta\|_{\mathcal{H}_\infty}$, varies (to be precise, we plot 30 Δ 's, each corresponding to a different ϵ , where we let ϵ range (equally spaced) between 0.0001 and 0.05). As seen in Figure 1, for the five set of controllers, there is a linear relationship (with different slopes) between the LQG cost gap ratio $\frac{J(\mathbf{K}_r^{(i)}) - J(\mathbf{K}^{(i)})}{J(\mathbf{K}^{(i)})}$ and the $\mathcal{H}_\infty$ norm of the truncated component, i.e. $\|\Delta\|_{\mathcal{H}_\infty}$; this is consistent with the upper bound on $J(\mathbf{K}_r)$ in Theorem 2.

6.2. Truncating unstable mode(s)

We here consider an example to illustrate Theorem 3. Consider the following plant $\mathbf{G}$

$$\mathbf{G} = \left[\begin{array}{ccc|c} \begin{bmatrix} -5.86 & -9.50 & 0.56 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} & \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \\ \hline \begin{bmatrix} -7.18 & -25.61 & -8.41 \end{bmatrix} & 0 \end{array} \right],$$

It can be verified that the controller $\mathbf{K} = \left[\begin{array}{c|c} A_K & B_K \\ \hline C_K & 0 \end{array} \right]$, with

$$A_K = \begin{bmatrix} 1.37 & 0 & 0 \\ 0 & -0.37 & 0 \\ 0 & 0 & 0.34 \end{bmatrix}, B_K = \begin{bmatrix} 0.19 \\ 0.04 \\ 0.04 \end{bmatrix}, C_K = \begin{bmatrix} 3.79 \\ 4.14 \\ -1.57 \end{bmatrix}^\top.$$

internally stabilizes $\mathbf{G}$. Applying Algorithm 1, we obtain an order 2 controller $\mathbf{K}_r$, which removes the last (unstable) mode of A_K , with an unstable truncated component Δ taking the form $\Delta = \left[\begin{array}{c|c} 0.34 & 0.04 \\ \hline -1.57 & 0 \end{array} \right]$, which satisfies the bound (22). Thus, Theorem 3 guarantees that the reduced-order controller $\mathbf{K}_r$ still internally stabilizes the plant. Indeed, numerical computation shows that the LQG cost of the original controller, $J(\mathbf{K})$, is 343.2, while the LQG cost of the truncated controller, $J(\mathbf{K}_r)$, is 58.2. In this case, modal truncation not only yields a stabilizing lower-order controller, but also a cost of lower LQG cost.⁷

7. Conclusion

We have presented controller reduction for general non observer-based controllers using balanced and modal truncation. For SISO systems, we demonstrate how LQG control may be performed even when there are no stable components in the controller. We hope our work will be useful not only for policy optimization in LQG control but also for the controller reduction community. Two interesting future directions are 1) extending truncation of unstable modes to MIMO systems and 2) applying the results to escape saddle points in the LQG policy optimization (Zheng et al., 2022).

7. For this instance, the theoretical upper bound in Theorem 3 is significantly larger than the original cost.




Figure 1: Relationship between LQG cost gap $J(\mathbf{K}_r) - J(\mathbf{K})$ and the $\mathcal{H}_\infty$ norm of the truncated component.

Acknowledgments

The work of Zhaolin Ren and Na Li is supported by NSF CAREER: ECCS-1553407, NSF CNS: 2003111, NSF AI institute: 2112085, and ONR YIP: N00014-19-1-2217. The work of Yang Zheng is supported by NSF ECCS 2154650.

References

- B.D.O. Anderson and Y. Liu. Controller reduction: concepts and approaches. *IEEE Trans. Automat. Contr.*, 34(8):802–812, 1989. doi: 10.1109/9.29422.
- Athanasios C Antoulas. *Approximation of large-scale dynamical systems*. SIAM, 2005.
- Athanasios C Antoulas, Danny C Sorensen, and Serkan Gugercin. A survey of model reduction methods for large-scale systems. *Contemporary Mathematics*, 280:193–219, 2001.
- Peter Benner, Mario Ohlberger, Albert Cohen, and Karen Willcox. *Model reduction and approximation: theory and algorithms*. SIAM, 2017.
- John Doyle and Gunter Stein. Multivariable feedback design: Concepts for a classical/modern synthesis. *IEEE Trans. Automat. Contr.*, 26(1):4–16, 1981.
- Dale F Enns. Model reduction with balanced realizations: An error bound and a frequency weighted generalization. In *The 23rd IEEE conference on decision and control*, pages 127–132. IEEE, 1984.
- Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pages 1467–1476. PMLR, 2018.
- Keith Glover and DJN Limebeer. Robust multivariable control system design using optimal reduced order plant models. In *1983 American Control Conference*, pages 644–649. IEEE, 1983.
- Bin Hu and Yang Zheng. Connectivity of the feasible and sublevel sets of dynamic output feedback control with robustness constraints. *IEEE Control Systems Letters*, 7:442–447, 2023. doi: 10.1109/LCSYS.2022.3188008.
- Bin Hu, Kaiqing Zhang, Na Li, Mehran Mesbahi, Maryam Fazel, and Tamer Başar. Towards a theoretical foundation of policy optimization for learning control policies. *arXiv preprint arXiv:2210.04810*, 2022.
- Edmond Jonckheere and Leonard Silverman. A new set of invariants for linear systems—application to reduced order compensator design. *IEEE Trans. Automat. Contr.*, 28(10):953–964, 1983.
- Yi Liu and Brian DO Anderson. Controller reduction via stable factorization and balancing. *International Journal of Control*, 44(2):507–531, 1986.
- Denis Mustafa and Keith Glover. Controller reduction by $\mathcal{H}_\infty$ balanced truncation. *IEEE Trans. Automat. Contr.*, 36(6):668–682, 1991.
- Goro Obinata and Brian DO Anderson. *Model reduction for control system design*. Springer Science & Business Media, 2012.
- Lars Pernebo and Leonard Silverman. Model reduction via balanced state space representations. *IEEE Trans. Automat. Contr.*, 27(2):382–387, 1982.

Zhaolin Ren, Yang Zheng, Maryam Fazel, and Na Li. On controller reduction in Linear Quadratic Gaussian Control with performance bounds. *arXiv preprint arXiv:2211.16291*, 2022.

Robert E Skelton. *Dynamic systems control: linear systems analysis and synthesis*. John Wiley & Sons, Inc., 1988.

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

Mathukumalli Vidyasagar. Coprime factorizations and stability of multivariable distributed feedback systems. *SIAM Journal on Control*, 13(6):1144–1155, 1975.

Kaiqing Zhang, Bin Hu, and Tamer Basar. Policy optimization for $\mathcal{H}_2$ linear control with $\mathcal{H}_\infty$ robustness guarantee: Implicit regularization and global convergence. In *Learning for Dynamics and Control*, pages 179–190, 2020.

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

Yang Zheng, Yue Sun, Maryam Fazel, and Na Li. Escaping high-order saddles in policy optimization for linear quadratic gaussian (LQG) control. *arXiv preprint arXiv:2204.00912*, 2022.

Kemin Zhou and Jie Chen. Performance bounds for coprime factor controller reductions. *Systems & control letters*, 26(2):119–127, 1995.

Kemin Zhou, John Comstock Doyle, Keith Glover, et al. *Robust and optimal control*, volume 40. Prentice hall New Jersey, 1996.