Analysis of the Optimization Landscape of Linear Quadratic Gaussian (LQG) Control

Yujie Tang * YUJIETANG@SEAS.HARVARD.EDU
School of Engineering and Applied Sciences, Harvard University, USA

Yang Zheng * ZHENGY@G.HARVARD.EDU
School of Engineering and Applied Sciences, Harvard University, USA

Na Li † NALI@SEAS.HARVARD.EDU
School of Engineering and Applied Sciences, Harvard University, USA

Editors: A. Jadbabaie, J. Lygeros, G. J. Pappas, P. A. Parrilo, B. Recht, C. J. Tomlin, M. N. Zeilinger

Abstract

This paper revisits the classical Linear Quadratic Gaussian (LQG) control from a modern optimization perspective. We analyze two aspects of the optimization landscape of the LQG problem: 1) connectivity of the set of stabilizing controllers $\mathcal{C}_n$; and 2) structure of stationary points. It is known that similarity transformations do not change the input-output behavior of a dynamical controller or LQG cost. This inherent symmetry by similarity transformations makes the landscape of LQG very rich. We show that 1) the set of stabilizing controllers $\mathcal{C}_n$ has at most two path-connected components and they are diffeomorphic under a mapping defined by a similarity transformation; 2) there might exist many *strictly suboptimal stationary points* of the LQG cost function over $\mathcal{C}_n$ and these stationary points are always *non-minimal*; 3) all *minimal* stationary points are globally optimal and they are identical up to a similarity transformation. These results shed some light on the performance analysis of direct policy gradient methods for solving the LQG problem.

1. Introduction

As one of the most fundamental optimal control problems, Linear Quadratic Gaussian (LQG) control has been studied for decades. Many structural properties of the LQG problem have been established in the literature, such as existence of the optimal controller, separation principle of the controller structure, and no guaranteed stability margin of closed-loop LQG systems (Zhou et al., 1996; Bertsekas, 2017; Doyle, 1978). Despite the non-convexity of the LQG problem, the globally optimal controller can be found by solving two algebraic Riccati equations (Zhou et al., 1996), or a convex semidefinite program based on a change of variables (Gahinet and Apkarian, 1994; Scherer et al., 1997).

While extensive results on LQG have been obtained in classical control, its optimization landscape is less studied, i.e., viewing the LQG cost as a function of the controller parameters and studying its analytical and geometrical properties. On the other hand, recent advances in reinforcement learning (RL) have revealed that the landscape analysis of another benchmark optimal control problem, linear quadratic regulator (LQR), can lead to fruitful and profound results, especially for model-free controller synthesis (Fazel et al., 2018; Malik et al., 2019; Mohammadi et al., 2019; Tu

---

* Y. Tang and Y. Zheng contributed equally to this work and are listed in alphabetical order.

† This work is supported by NSF career, AFOSR YIP, and ONR YIP.

and Recht, 2019; Li et al., 2019b; Umenberger et al., 2019; Zhang et al., 2020a). For instance, it is shown that the set of static stabilizing feedback gains for LQR is connected, and that the LQR cost function is coercive and enjoys an interesting property of gradient dominance (Fazel et al., 2018; Bu et al., 2019). These properties are fundamental for establishing convergence guarantees of gradient-based algorithms for solving LQR and their model-free extensions for RL (Malik et al., 2019; Mohammadi et al., 2019). We note that recent studies have also contributed to establishing performance guarantees of model-based RL techniques for LQR (see, e.g., Dean et al. (2020); Wang and Janson (2020)) as well as LQG control (Tu et al., 2017; Boczar et al., 2018; Zheng et al., 2020; Simchowitz et al., 2020).

This paper aims to analyze the optimization landscape of the LQG problem. Unlike LQR that deals with *fully observed* linear systems whose optimal solution is a static feedback policy, the LQG problem concerns *partially observed* linear systems driven by additive Gaussian noise, and its optimal controller is no longer static. We need to search over dynamical controllers for LQG problems. This makes its optimization landscape richer and yet much more complicated than LQR. Indeed, the set of stabilizing static state feedback policies is connected, but the set of stabilizing static output feedback policies can be highly disconnected (Feng and Lavaei, 2020). The connectivity of stabilizing dynamical output feedback policies, i.e., the feasible region of LQG control, remains unclear. Furthermore, LQG has a natural symmetry structure induced by similarity transformations that do not change the input-output behavior of dynamical controllers, which is not the case for LQR.

Some recent studies (Sun et al., 2018; Chi et al., 2019; Li et al., 2019a; Qu et al., 2019; Ge and Ma, 2017) have demonstrated that symmetry properties play a key role in rendering a large class of non-convex optimization problems in machine learning tractable; see also Zhang et al. (2020b) for a recent review. For the LQG problem, we can expect the inherent symmetry by similarity transformations to bring some important properties of its non-convex optimization landscape. We also note that the notion of *minimal controllers* (i.e., controllable and observable controllers) is a unique feature in controller synthesis of *partially observed* dynamical systems, making the optimization landscape of LQG distinct from many machine learning problems. We provide an extended review on related work in Zheng et al. (2021).

Our contributions We first characterize the connectivity of the feasible region of the LQG problem, i.e., the set of strictly proper stabilizing dynamical controllers, denoted by $\mathcal{C}_n$ (n is the state dimension). We prove that $\mathcal{C}_n$ can be disconnected, but has at most two path-connected components (Theorem 3.1). If $\mathcal{C}_n$ is disconnected, its two path-connected components are diffeomorphic under a mapping defined by a similarity transformation (Theorem 3.2). This brings positive news to gradient-based local search algorithms for the LQG problem, since it makes no difference to search over either path-connected component even if $\mathcal{C}_n$ is disconnected. We further present a sufficient condition under which $\mathcal{C}_n$ is always connected, and this condition becomes necessary for a class of LQG problems with a single input or a single output (Theorem 3.3).

Second, we investigate structural properties of the stationary points of the LQG cost function. It is known that the LQG cost is invariant under similarity transformations on the controller (see Lemma 4.1). One natural consequence is that the globally optimal solutions to the LQG problem are not unique, not isolated, and can be disconnected in the state-space domain. For a class of LQG problems, we show that the set of globally optimal solutions is a submanifold of dimension n^2 and it has two path-connected components (Proposition 4.1). When characterizing the stationary points with vanishing gradients, the notion of *minimal controllers* plays an important role. In Theorem 4.1, we show that it is very likely there exist many *strictly suboptimal stationary points* of the LQG cost over $\mathcal{C}_n$, and these stationary points are always non-minimal. In contrast, we prove that all *minimal* stationary points are globally optimal to the LQG problem (Theorem 4.2). These minimal stationary points are identical up to similarity transformations. This is expected from the classical result that the globally optimal LQG controller is unique in the frequency domain (Zhou et al., 1996, Theorem 14.7). Our analysis implies that if local search iterates converge to a critical point that corresponds to a controllable and observable controller, then the algorithm has found a globally optimal solution to the LQG problem (Corollary 4.2). However, it requires further investigation on whether local search algorithms can escape saddle points of LQG (Lee et al., 2019).

2. Problem Statement

2.1. The Linear Quadratic Gaussian (LQG) problem

Consider a continuous-time[1] linear dynamical system

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

where $x \in \mathbb{R}^n, u \in \mathbb{R}^m, y \in \mathbb{R}^p$ represents the state, input, and output, respectively, and $w \in \mathbb{R}^n, v \in \mathbb{R}^p$ are process and measurement noises. It is assumed that w and v are white Gaussian noises with intensity matrices $W \succeq 0$ and $V \succ 0$. The classical LQG problem is defined as

$$\begin{aligned} \min_{u(t)} \quad & J := \lim_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_{t=0}^{T} \left(x^{\mathsf{T}}Qx + u^{\mathsf{T}}Ru\right) dt\right] \\ \text{subject to} \quad & (1), \end{aligned} \tag{2}$$

where $Q \succeq 0$ and $R \succ 0$. In (2), the input $u(t)$ is allowed to depend on all past observations $y(\tau)$ with $\tau < t$. Throughout the paper, we make the following standard assumption of *minimal* systems.

Assumption 1 *(A, B) and $(A, W^{1/2})$ are controllable, and (C, A) and $(Q^{1/2}, A)$ are observable.*

Unlike the linear quadratic regulator (LQR), static policies in general do not achieve the optimal cost, and we need to consider the class of dynamical controllers in the form of

$$\begin{aligned} \dot{\xi}(t) &= A_{\mathsf{K}}\xi(t) + B_{\mathsf{K}}y(t), \\ u(t) &= C_{\mathsf{K}}\xi(t), \end{aligned} \tag{3}$$

where $\xi(t) \in \mathbb{R}^q$ is the internal state of the controller, and $A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}$ are matrices of proper dimensions that specify the dynamics of the controller. We refer to the dimension q of the internal state variable ξ as the order of the dynamical controller (3). A dynamical controller is called *full-order* if its order is the same as the system dimension, i.e., $q = n$; if $q < n$, we call (3) a *reduced-order* or *lower-order* controller. We shall see later that it is unnecessary to consider dynamical controllers with order beyond the system dimension n.

It is well-known that the LQG problem (2) has a closed-form solution by solving two algebraic Riccati equations (Zhou et al., 1996, Theorem 14.7). Precisely, the optimal controller is given by

$$\dot{\xi} = (A - BK)\xi + L(y - C\xi), \qquad u = -K\xi. \tag{4}$$

1. We only consider the continuous-time case; see our extended version (Zheng et al., 2021) for the discrete-time case.

In (4), the matrix L is called the *Kalman gain*, computed as $L = PC^{\mathsf{T}}V^{-1}$ where P is the unique positive semidefinite solution [see, e.g., Zhou et al. (1996, Corollary 13.8)] to

$$AP + PA^{\mathsf{T}} - PC^{\mathsf{T}}V^{-1}CP + W = 0, \tag{5a}$$

and the matrix K is called the *feedback gain*, computed as $K = R^{-1}B^{\mathsf{T}}S$ where S is the unique positive semidefinite solution to

$$A^{\mathsf{T}}S + SA - SBR^{-1}B^{\mathsf{T}}S + Q = 0. \tag{5b}$$

It is clear that the optimal solution from Ricatti equations (5) is always full-order, i.e., $q = n$.

2.2. Parametrization of Dynamical Controllers and the LQG Cost Function

Here, we view the cost in (2) as a function of the parametrized dynamical controller $(A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}})$. By combining (3) with (1), we get the closed-loop system

$$\begin{aligned}\frac{d}{dt}\begin{bmatrix} x \\ \xi \end{bmatrix} &= \begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix}\begin{bmatrix} x \\ \xi \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & B_{\mathsf{K}} \end{bmatrix}\begin{bmatrix} w \\ v \end{bmatrix}, \\ \begin{bmatrix} y \\ u \end{bmatrix} &= \begin{bmatrix} C & 0 \\ 0 & C_{\mathsf{K}} \end{bmatrix}\begin{bmatrix} x \\ \xi \end{bmatrix} + \begin{bmatrix} v \\ 0 \end{bmatrix}.\end{aligned} \tag{6}$$

We denote the set of stabilizing controllers with order $q \in \mathbb{N}$ by

$$\mathcal{C}_q := \left\{ \mathsf{K} = \begin{bmatrix} 0_{m\times p} & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{(m+q)\times(p+q)} \middle| \begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix} \text{ is stable} \right\}, \tag{7}$$

and let $J_q : \mathcal{C}_q \to \mathbb{R}$ denote the function that maps a parameterized dynamical controller in $\mathcal{C}_q$ to its corresponding LQG cost for each $q \in \mathbb{N}$. The following two lemmas give useful characterizations of the cost function J_q; see Zheng et al. (2021) for a short proof.

Lemma 2.1 *Fix $q \in \mathbb{N}$ such that $\mathcal{C}_q \neq \varnothing$. Given $\mathsf{K} \in \mathcal{C}_q$, we have*

$$J_q(\mathsf{K}) = \operatorname{tr}\left(\begin{bmatrix} Q & 0 \\ 0 & C_{\mathsf{K}}^{\mathsf{T}}RC_{\mathsf{K}} \end{bmatrix} X_{\mathsf{K}}\right) = \operatorname{tr}\left(\begin{bmatrix} W & 0 \\ 0 & B_{\mathsf{K}}VB_{\mathsf{K}}^{\mathsf{T}} \end{bmatrix} Y_{\mathsf{K}}\right), \tag{8}$$

where X_{K} and Y_{K} are the unique positive semidefinite solutions to the following Lyapunov equations

$$\begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix} X_{\mathsf{K}} + X_{\mathsf{K}}\begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix}^{\mathsf{T}} + \begin{bmatrix} W & 0 \\ 0 & B_{\mathsf{K}}VB_{\mathsf{K}}^{\mathsf{T}} \end{bmatrix} = 0, \tag{9a}$$

$$\begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix}^{\mathsf{T}} Y_{\mathsf{K}} + Y_{\mathsf{K}}\begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix} + \begin{bmatrix} Q & 0 \\ 0 & C_{\mathsf{K}}^{\mathsf{T}}RC_{\mathsf{K}} \end{bmatrix} = 0. \tag{9b}$$

Lemma 2.2 *Fix $q \in \mathbb{N}$ such that $\mathcal{C}_q \neq \varnothing$. Then, J_q is a real analytic function on $\mathcal{C}_q$.*

Now, given the dimension n of the system state variable, the LQG problem (2) can be reformulated into a constrained optimization problem:

$$\begin{aligned} \min_{\mathsf{K}} \quad & J_n(\mathsf{K}) \\ \text{subject to} \quad & \mathsf{K} \in \mathcal{C}_n. \end{aligned} \tag{10}$$

After reformulating the LQG problem (2) into (10), a tentative approach for model-free reinforcement learning of LQG is to conduct policy gradient on (10), with the gradient of $J_n(\mathsf{K})$ estimated from system trajectories. To characterize the performance of policy gradient algorithms, it is necessary to understand the optimization landscape of (10). It is well-known that $\mathcal{C}_n$ is in general non-convex and Lemma 2.2 indicates that J_n is a real analytical function. However, little is known about their further geometrical and analytical properties, especially those that are fundamental for establishing convergence of gradient-based algorithms. In this paper, we characterize 1) the connectivity of $\mathcal{C}_n$ (Section 3) and 2) the stationary points of the LQG cost function $J_n(\mathsf{K})$ (Section 4).

3. Connectivity of the Set of Full-Order Stabilizing Controllers

We first have the following observation on some basic properties of the set of full-order stabilizing controllers $\mathcal{C}_n$.

Lemma 3.1 *Under Assumption 1, the set $\mathcal{C}_n$ is non-empty, unbounded, and can be non-convex.*

Example 1 (Non-convexity of stabilizing controllers) *Consider a dynamical system (1) with $A = 1,\ B = 1,\ C = 1$. The set of stabilizing controllers $\mathcal{C}_n = \mathcal{C}_1$ is given by*

$$\mathcal{C}_1 = \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| \begin{bmatrix} 1 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \text{ is stable} \right\}.$$

The following dynamical controllers $\mathsf{K}^{(1)} = \begin{bmatrix} 0 & 2 \\ -2 & -2 \end{bmatrix}, \mathsf{K}^{(2)} = \begin{bmatrix} 0 & -2 \\ 2 & -2 \end{bmatrix}$ stabilize the plant and thus belong to $\mathcal{C}_1$. However, $\hat{\mathsf{K}} = \frac{1}{2}\left(\mathsf{K}^{(1)} + \mathsf{K}^{(2)}\right) = \begin{bmatrix} 0 & 0 \\ 0 & -2 \end{bmatrix}$ fails to stabilize the plant.

We now introduce the notion of similarity transformations that has been widely-used in control theory. Let GL_q denote the set of $q \times q$ real invertible matrices. Given $q \geq 1$ such that $\mathcal{C}_q \neq \varnothing$, we define the mapping $\mathscr{T}_q : \mathrm{GL}_q \times \mathcal{C}_q \to \mathcal{C}_q$ that represents similarity transformations on $\mathcal{C}_q$ by

$$\mathscr{T}_q(T, \mathsf{K}) := \begin{bmatrix} I_m & 0 \\ 0 & T \end{bmatrix} \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \begin{bmatrix} I_p & 0 \\ 0 & T \end{bmatrix}^{-1} = \begin{bmatrix} 0 & C_{\mathsf{K}} T^{-1} \\ T B_{\mathsf{K}} & T A_{\mathsf{K}} T^{-1} \end{bmatrix}. \tag{11}$$

For any invertible matrix $T \in \mathrm{GL}_q$ and $\mathsf{K} \in \mathcal{C}_q$, $\mathscr{T}_q(T, \mathsf{K})$ is indeed a stabilizing controller of order q and thus is in $\mathcal{C}_q$. We can also check that $\mathscr{T}_q$ is indefinitely differentiable on $\mathrm{GL}_q \times \mathcal{C}_q$, and that

$$\mathscr{T}_q(T_2, \mathscr{T}_q(T_1, \mathsf{K})) = \mathscr{T}_q(T_2 T_1, \mathsf{K}) \tag{12}$$

for any $T_1, T_2 \in \mathrm{GL}_q$. This implies that for any fixed $T \in \mathrm{GL}_q$, the map $\mathsf{K} \mapsto \mathscr{T}_q(T, \mathsf{K})$ admits an indefinitely differentiable inverse given by $\mathsf{K} \mapsto \mathscr{T}_q(T^{-1}, \mathsf{K})$. We therefore have:

Lemma 3.2 *Given $q \geq 1$ such that $\mathcal{C}_q \neq \varnothing$, for any invertible matrix $T \in \mathrm{GL}_q$, the map $\mathsf{K} \mapsto \mathscr{T}_q(T, \mathsf{K})$ is a diffeomorphism from $\mathcal{C}_q$ to itself.*

For notational simplicity, for any fixed $T \in \mathrm{GL}_n$, we let $\mathscr{T}_T : \mathcal{C}_n \to \mathcal{C}_n$ denote the mapping given by $\mathscr{T}_T(\mathsf{K}) := \mathscr{T}_n(T, \mathsf{K})$ between the set of *full-order* stabilizing controllers. We are now ready to present the main results.

Theorem 3.1 *Under Assumption 1, the set of full-order stabilizing controllers $\mathcal{C}_n$ has at most two path-connected components.*

Theorem 3.2 *If $\mathcal{C}_n$ has two path-connected components $\mathcal{C}_n^{(1)}$ and $\mathcal{C}_n^{(2)}$, then $\mathcal{C}_n^{(1)}$ and $\mathcal{C}_n^{(2)}$ are diffeomorphic under the mapping $\mathscr{T}_T$ for any invertible matrix $T \in \mathbb{R}^{n\times n}$ with $\det T < 0$.*

Theorem 3.2 shows that even if $\mathcal{C}_n$ has two path-connected components, there exists a linear bijection defined by a similarity transformation $\mathscr{T}_T$ between these two components. In the following theorem, we present a sufficient condition under which $\mathcal{C}_n$ is path-connected. This condition becomes necessary for a class of dynamical systems with single input or single output.

Theorem 3.3 *Under Assumption 1, the following statements are true.*

1. *The set of full-order stabilizing controllers $\mathcal{C}_n$ is path-connected if there exists a reduced-order stabilizing controller,* i.e., *$\mathcal{C}_{n-1} \neq \varnothing$.*
2. *Suppose the plant (1) is single-input or single-output, i.e., $m = 1$ or $p = 1$. Then the set $\mathcal{C}_n$ is path-connected if and only if $\mathcal{C}_{n-1} \neq \varnothing$.*

One main idea in our proofs[2] is based on a classical change of variables that has been widely used for developing convex reformulation of various controller synthesis problems (Scherer et al., 1997; Gahinet and Apkarian, 1994). We adopt the change of variables to construct a set with a convex projection and a surjective mapping from that set to $\mathcal{C}_n$, and then path-connectivity results generally follow from the fact that a convex set is path-connected. The potential disconnectivity of $\mathcal{C}_n$ comes from the fact that the set of real invertible matrices $\mathrm{GL}_n = \{\Pi \in \mathbb{R}^{n\times n} \mid \det \Pi \neq 0\}$ has two path-connected components (Lee, 2013): $\mathrm{GL}_n^+ = \{\Pi \in \mathbb{R}^{n\times n} \mid \det \Pi > 0\}$, $\mathrm{GL}_n^- = \{\Pi \in \mathbb{R}^{n\times n} \mid \det \Pi < 0\}$. The full proofs are technically involved, which are provided in Zheng et al. (2021, Section 3.2 – Section 3.4).

We note that given any open-loop unstable first-order dynamical system, *i.e.*, $n = 1$, and $A > 0$ in (1), it is easy to see that there exist no reduced-order stabilizing controllers, *i.e.*, $\mathcal{C}_{n-1} = \varnothing$. Thus, Theorem 3.3 indicates that its associated set of stabilizing controllers $\mathcal{C}_n$ is not path-connected. We provide an explicit single-input and single-output (SISO) example below.

Example 2 (Disconectivity of stabilizing controllers) *Consider the system in Example 1: $A = 1,\ B = 1,\ C = 1$. The above reasoning indicates that $\mathcal{C}_n$ is not path-connected. Indeed, using the Routh–Hurwitz stability criterion, it is straightforward to derive that*

$$\mathcal{C}_1 = \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| A_{\mathsf{K}} < -1,\ B_{\mathsf{K}} C_{\mathsf{K}} < A_{\mathsf{K}} \right\}. \tag{13}$$

This set has two path-connected components $\mathcal{C}_1^+$ and $\mathcal{C}_1^-$ given by

$$\mathcal{C}_1^+ := \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| A_{\mathsf{K}} < -1,\ B_{\mathsf{K}} C_{\mathsf{K}} < A_{\mathsf{K}},\ B_{\mathsf{K}} > 0 \right\},$$
$$\mathcal{C}_1^- := \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| A_{\mathsf{K}} < -1,\ B_{\mathsf{K}} C_{\mathsf{K}} < A_{\mathsf{K}},\ B_{\mathsf{K}} < 0 \right\}.$$

2. The controllers in (3) and (7) are *strictly proper*, which is sufficient for the LQG problem (2). For closed-loop stability, we can consider *proper* dynamical controllers. Unlike $\mathcal{C}_n$ that might be disconnected, the set of stabilizing proper dynamical controllers is always path-connected [see the appendix of Zheng et al. (2021)]. Our proof techniques via the change of variables work for both cases.

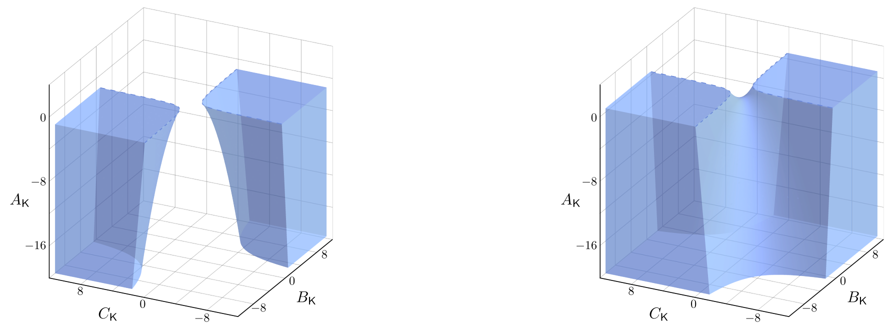

(a) $\mathcal{C}_1$ for Example 2 **(b)** $\mathcal{C}_1$ for Example 3

Figure 1: The set of stabilizing controllers $\mathcal{C}_1$ for Examples 2 and 3: (a) For Example 2, the set $\mathcal{C}_1$ given by (13) has two path-connected components; (b) For Example 3, the set $\mathcal{C}_1$ given by (14) is path-connected.

In addition, as expected from Theorem 3.2, it is easy to verify that $\mathcal{C}_1^+$ and $\mathcal{C}_1^-$ are homeomorphic under the mapping $\mathscr{T}_T$ for any $T < 0$. Figure 1a illustrates the region of the set $\mathcal{C}_1$ in (13).

Theorem 3.3 also suggests the following result: *Given any open-loop stable dynamical system* (1), i.e.*, A is stable, we have that $\mathcal{C}_n$ is path-connected.* An explicit example is shown below.

Example 3 (Stabilizing controllers for an open-loop stable system) *Consider an open-loop stable dynamical system (1) with $A = -1,\ B = 1,\ C = 1$. Since it is open-loop stable, Theorem 3.3 indicates that its associated set of stabilizing controllers $\mathcal{C}_n$ is path-connected. Using the Routh–Hurwitz stability criterion, it is straightforward to derive that*

$$\mathcal{C}_1 = \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| A_{\mathsf{K}} < 1, B_{\mathsf{K}} C_{\mathsf{K}} < -A_{\mathsf{K}} \right\}. \tag{14}$$

This set is path-connected, as illustrated in Figure 1b.

Remark 1 (Connectivity of the feasible region of LQR/LQG) *Some recent studies revisited the classical LQR problem from a modern optimization perspective and designed policy gradient algorithms (Fazel et al., 2018; Mohammadi et al., 2019; Zhang et al., 2020a). The connectivity of the feasible region (i.e., the set of stabilizing controllers) becomes important to local search algorithms (e.g., policy gradient) since they typically cannot jump between different connected components. It is known that the set of stabilizing static state-feedback policies $\{K \in \mathbb{R}^{m\times n} \mid A - BK \text{ is stable}\}$ is connected (Bu et al., 2019), and this is one important factor in justifying the performance of the algorithms in Fazel et al. (2018); Mohammadi et al. (2019); Zhang et al. (2020a). On the other hand, the set of stabilizing static output feedback policies $\{D_{\mathsf{K}} \in \mathbb{R}^{m\times p} \mid A - BD_{\mathsf{K}}C \text{ is stable}\}$ can be highly disconnected (Feng and Lavaei, 2020), posing a significant challenge for local search algorithms. In Theorems 3.1, 3.2 and 3.3, we have shown that the set of stabilizing* dynamical *controllers $\mathcal{C}_n$ in the LQG problem has at most two path-connected components that are diffeomorphic to each other under a particular similarity transformation. Since similarity transformations do not change the input/output behavior of a controller, it makes no difference to search over either path-connected component of $\mathcal{C}_n$ even if $\mathcal{C}_n$ is not path-connected. This brings positive news to gradient-based local search algorithms for the LQG problem.*

4. Structure of Stationary Points

We have shown that $\mathcal{C}_n$ might be disconnected, and that the potential disconnectivity has no harm to gradient-based local search algorithms. In this section, we proceed to characterize the stationary points of $J_n(\mathsf{K})$, which is also important for performance analysis of gradient-based algorithms.

4.1. Invariance of LQG Cost under Similarity Transformations

As shown in Lemma 3.2, the similarity transformation $\mathscr{T}_q(T,\cdot)$ is a diffeomorphism from $\mathcal{C}_q$ to itself for any invertible matrix $T \in \mathrm{GL}_q$. Then together with (12), we can see that the set of similarity transformations is a group isomorphic to GL_q. We can therefore define the *orbit* of $\mathsf{K} \in \mathcal{C}_q$ by

$$\mathcal{O}_{\mathsf{K}} := \{\mathscr{T}_q(T,\mathsf{K}) \mid T \in \mathrm{GL}_q\}.$$

It is known that the LQG cost is invariant under similarity transformations on the controller, and thus is a constant over an orbit $\mathcal{O}_{\mathsf{K}}$ for any $\mathsf{K} \in \mathcal{C}_q$.

Lemma 4.1 *Let $q \geq 1$ such that $\mathcal{C}_q \neq \varnothing$. Then we have $J_q(\mathsf{K}) = J_q(\mathscr{T}_q(T,\mathsf{K}))$ for any $\mathsf{K} \in \mathcal{C}_q$ and any invertible matrix $T \in \mathrm{GL}_q$.*

We further have the following proposition characterizing the dimension of every orbit $\mathcal{O}_{\mathsf{K}}$ corresponding to minimal controllers; see Zheng et al. (2021) for the proof.

Proposition 4.1 *Suppose $\mathsf{K} \in \mathcal{C}_q$ represents a controllable and observable controller. Then the orbit $\mathcal{O}_{\mathsf{K}}$ is a submanifold of $\mathcal{C}_q$ of dimension q^2, and has two path-connected components, given by*

$$\begin{aligned}\mathcal{O}_{\mathsf{K}}^{+} &= \{\mathscr{T}_q(T,\mathsf{K}) \mid T \in \mathrm{GL}_q, \det T > 0\},\\ \mathcal{O}_{\mathsf{K}}^{-} &= \{\mathscr{T}_q(T,\mathsf{K}) \mid T \in \mathrm{GL}_q, \det T < 0\}.\end{aligned}$$

From Lemma 4.1 and Proposition 4.1, one interesting consequence is that given a globally optimal LQG controller $\mathsf{K}^* \in \mathcal{C}_n$, any controller in the orbit $\mathcal{O}_{\mathsf{K}^*} := \{\mathscr{T}_n(T,\mathsf{K}^*) \mid T \in \mathrm{GL}_n\}$ is globally optimal. If K^* is minimal (i.e., controllable and observable), the orbit $\mathcal{O}_{\mathsf{K}^*}$ is a submanifold of dimension n^2, and it has two path-connected components. Figure 2 demonstrates the orbit of globally optimal LQG controllers for an open-loop unstable system and another open-loop stable system. Additionally, the LQG cost function $J_q(\mathsf{K})$ is not coercive in the sense that there might exist 1) sequences of stabilizing controllers $\mathsf{K}_j \in \mathcal{C}_q$ where $\lim_{j\to\infty}\mathsf{K}_j = \hat{\mathsf{K}} \in \partial\mathcal{C}_q$ such that $\lim_{j\to\infty} J_q(\mathsf{K}_j) < \infty$, and 2) sequences of stabilizing controllers $\mathsf{K}_j \in \mathcal{C}_q$ where $\lim_{j\to\infty}\|\mathsf{K}_j\|_F = \infty$ such that $\lim_{j\to\infty} J_q(\mathsf{K}_j) < \infty$. The latter fact is easy to see from Proposition 4.1 since the orbit $\mathcal{O}_{\mathsf{K}}$ can be unbounded and $J_q(\mathsf{K})$ is constant over the same orbit. The following example shows that the LQG cost can converge to a finite value even when the controller K goes to the boundary of $\mathcal{C}_q$.

Example 4 (Non-coercivity of the LQG cost) *Consider the open-loop stable SISO system in Example 3, and we fix $Q = 1, R = 1, V = 1, W = 1$ in the LQG formulation. The set of full-order stabilizing controllers $\mathcal{C}_1$ is shown in (14). We consider the following stabilizing controller $\mathsf{K}_\epsilon = \begin{bmatrix} 0 & \epsilon \\ -\epsilon & 0\end{bmatrix} \in \mathcal{C}_1, \forall \epsilon \neq 0$. It is not hard to see that $\lim_{\epsilon\to 0}\mathsf{K}_\epsilon \in \partial\mathcal{C}_1$. By solving the Lyapunov equation (9a), we get the unique solution as $X_{\mathsf{K}_\epsilon} = \frac{1}{2}\begin{bmatrix}\epsilon^2+1 & \epsilon \\ \epsilon & \epsilon^2+2\end{bmatrix}$, and the corresponding LQG cost as $J(\mathsf{K}_\epsilon) = \frac{1+3\epsilon^2+\epsilon^4}{2}$. Therefore, we have $\lim_{\epsilon\to 0} J(\mathsf{K}_\epsilon) = 1/2$, while $\lim_{\epsilon\to 0}\mathsf{K}_\epsilon \in \partial\mathcal{C}_1$.*

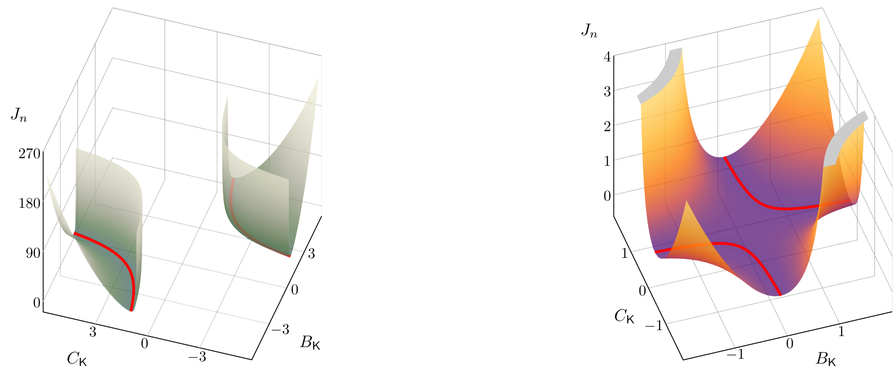

(a) Open-loop unstable system in Example 2 **(b)** Open-loop stable system in Example 3

Figure 2: Non-isolated and disconnected globally optimal LQG controllers. In both cases, we set $Q = 1, R = 1, V = 1, W = 1$. (a) LQG cost for Example 2 when fixing $A_{\mathsf{K}} = -1 - 2\sqrt{2}$, for which the set of globally optimal points $\left\{(B_{\mathsf{K}}, C_{\mathsf{K}}) \mid B_{\mathsf{K}} = (1+\sqrt{2})\frac{1}{T}, C_{\mathsf{K}} = -(1+\sqrt{2})T,\ T \neq 0\right\}$ has two connected components. (b) LQG cost for Example 3 when fixing $A_{\mathsf{K}} = 1 - 2\sqrt{2}$, for which the set of globally optimal points $\left\{(B_{\mathsf{K}}, C_{\mathsf{K}}) \mid B_{\mathsf{K}} = (-1+\sqrt{2})\frac{1}{T}, C_{\mathsf{K}} = (1-\sqrt{2})T,\ T \neq 0\right\}$ has two connected components.

4.2. Non-minimal Stationary Points

We provide an interesting result on non-minimal stationary points.

Theorem 4.1 *Let $q \geq 1$ be arbitrary. Suppose there exists $\mathsf{K}^\star = \begin{bmatrix} 0 & C_{\mathsf{K}}^\star \\ B_{\mathsf{K}}^\star & A_{\mathsf{K}}^\star \end{bmatrix} \in \mathcal{C}_q$ such that $\nabla J_q(\mathsf{K}^\star) = 0$. Then for any $q' \geq 1$ and any stable $\Lambda \in \mathbb{R}^{q' \times q'}$, the following controller*

$$\tilde{\mathsf{K}}^\star = \left[\begin{array}{c|cc} 0 & C_{\mathsf{K}}^\star & 0 \\ \hline B_{\mathsf{K}}^\star & A_{\mathsf{K}}^\star & 0 \\ 0 & 0 & \Lambda \end{array}\right] \in \mathcal{C}_{q+q'} \tag{15}$$

is a stationary point of $J_{q+q'}$ over $\mathcal{C}_{q+q'}$ satisfying $J_{q+q'}\big(\tilde{\mathsf{K}}^\star\big) = J_q(\tilde{\mathsf{K}})$.

Theorem 4.1 indicates that from any stationary point of J_q over lower-order stabilizing controllers in $\mathcal{C}_q$, we can construct a family of stationary points of $J_{q+q'}$ over higher-order stabilizing controllers in $\mathcal{C}_{q+q'}$. Moreover, the stationary points constructed by (15) are neither controllable nor observable. This indicates that, if the globally optimal controller of J_n is controllable and observable, and if $\min_{\mathsf{K} \in \mathcal{C}_q} J_q(\mathsf{K})$ has a solution for some $q < n$, then there will exist many *strictly suboptimal stationary points* of J_n over $\mathcal{C}_n$. The proof of Theorem 4.1 relies on similarity transformation and is provided in Zheng et al. (2021). In Zheng et al. (2021, Theorem 4.2), we further construct an explicit family of stationary points for $J_n(\mathsf{K})$ with an open-loop stable plant, and provide a criterion for checking whether the corresponding Hessian is indefinite or vanishing.

4.3. Minimal Stationary Points Are Globally Optimal

Theorem 4.1 shows that there may exist many *non-minimal* stationary points for J_n that are not globally optimal. Here, we show that all *minimal* stationary points are globally optimal to the LQG problem (10). In particular, we have closed-form expressions for full-order minimal stationary points $\mathsf{K} \in \mathcal{C}_n$, which turn out to be globally optimal.

Theorem 4.2 *Under Assumption 1, all minimal stationary points $\mathsf{K} \in \mathcal{C}_n$ to the LQG problem* (10) *are globally optimal, and they are in the form of*

$$A_{\mathsf{K}} = T(A - BK - LC)T^{-1}, \qquad B_{\mathsf{K}} = -TL, \qquad C_{\mathsf{K}} = KT^{-1}, \tag{16}$$

where $T \in \mathbb{R}^{n \times n}$ is an invertible matrix, and

$$K = R^{-1}B^{\mathsf{T}}S, \qquad L = PC^{\mathsf{T}}V^{-1}, \tag{17}$$

with P and S being the unique positive semidefinite solutions to the Riccati equations (5a) *and* (5b).

Theorem 4.2 can be viewed as a special case in Zhou et al. (1996, Theorem 20.6) that presents first-order necessary conditions for optimal reduced-order controllers $\mathsf{K} \in \mathcal{C}_q$; see Zheng et al. (2021) for an adapted proof. Theorem 4.2 indicates that if the LQG problem (10) has a globally optimal solution in $\mathcal{C}_n$ that is also minimal, then the globally optimal controllers are unique in $\mathcal{C}_n$ up to a similarity transformation. This is expected from the classical result that the globally optimal LQG controller is unique in the frequency domain (Zhou et al., 1996, Theorem 14.7). Theorem 4.2 allows us to establish the following corollary.

Corollary 4.1 *The following statements are true:*

1. *If $J_n(\mathsf{K})$ has a minimal stationary point in $\mathcal{C}_n$, then all its non-minimal stationary points $\mathsf{K} \in \mathcal{C}_n$ are strictly suboptimal.*
2. *If $J_n(\mathsf{K})$ has a non-minimal stationary point in $\mathcal{C}_n$ that is globally optimal, then all stationary points $\mathsf{K} \in \mathcal{C}_n$ of $J_n(\mathsf{K})$ are non-minimal.*

We have constructed explicit LQG examples with non-minimal stationary points that are strictly suboptimal in Zheng et al. (2021). It should be noted that, even with Assumption 1, the LQG problem (10) might have no minimal stationary points. This happens if the controller from the Ricatti equations (5) is not minimal; see the example in Zheng et al. (2021) taken from Yousuff and Skelton (1984). Theorem 4.2 also allows us to check whether a sequence of gradient iterates converges to a globally optimal solution.

Corollary 4.2 *Consider a gradient descent algorithm $\mathsf{K}_{t+1} = \mathsf{K}_t - \alpha \nabla J(\mathsf{K})$ for the LQG problem* (10)*. Suppose the iterates K_t converge to a point K^*,* i.e., $\lim_{t\to\infty} \mathsf{K}_t = \mathsf{K}^*$. *If K^* is a controllable and observable controller, then it is globally optimal.*

In our extended version (Zheng et al., 2021), numerical experiments are provided to demonstrate empirical performance of gradient descent methods for solving the LQG problem (10).

5. Conclusion and future work

In this paper, we have characterized the connectivity of the set of stabilizing controllers $\mathcal{C}_n$ and provided some structural properties of the LQG cost function. These results reveal rich yet complicated optimization landscape properties of the LQG problem. Ongoing work includes establishing convergence conditions for gradient descent algorithms and investigating whether local search algorithms can escape saddle points of the LQG problem. We note that the optimization landscape of LQG also depends on the parameterization of dynamical controllers. It will be interesting to look into the LQG problem when parameterizing controllers in a canonical form. Finally, our analysis reveals that *minimal* stationary points in $\mathcal{C}_n$ are always globally optimal, and it would be interesting to investigate the existence of minimal stationary points for the LQG problem.

References

Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, 4th edition, 2017.

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.

Jingjing Bu, Afshin Mesbahi, and Mehran Mesbahi. On topological and metrical properties of stabilizing feedback gains: the MIMO case. *arXiv preprint arXiv:1904.02737*, 2019.

Yuejie Chi, Yue M. Lu, and Yuxin Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239–5269, 2019.

Sarah Dean, Horia Mania, Nikolai Matni, et al. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20:633–679, 2020.

John C. Doyle. Guaranteed margins for LQG regulators. *IEEE Transactions on Automatic Control*, 23(4):756–757, 1978.

Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *Proceedings of the 35th International Conference on Machine Learning*, pages 1467–1476, 2018.

Han Feng and Javad Lavaei. Connectivity properties of the set of stabilizing static decentralized controllers. *SIAM Journal on Control and Optimization*, 58(5):2790–2820, 2020.

Pascal Gahinet and Pierre Apkarian. A linear matrix inequality approach to $\mathcal{H}_\infty$ control. *International Journal of Robust and Nonlinear Control*, 4(4):421–448, 1994.

Rong Ge and Tengyu Ma. On the optimization landscape of tensor decompositions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, page 3656–3666, 2017.

Jason D. Lee, Ioannis Panageas, Georgios Piliouras, et al. First-order methods almost always avoid strict saddle points. *Mathematical programming*, 176:311–337, 2019.

John M. Lee. *Introduction to Smooth Manifolds*. Springer Science & Business Media, 2013.

Xingguo Li, Junwei Lu, Raman Arora, et al. Symmetry, saddle points, and global optimization landscape of nonconvex matrix factorization. *IEEE Transactions on Information Theory*, 65(6): 3489–3514, 2019a.

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

Dhruv Malik, Ashwin Pananjady, Kush Bhatia, et al. Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, pages 2916–2925, 2019.

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

Qing Qu, Yuexiang Zhai, Xiao Li, et al. Analysis of the optimization landscapes for overcomplete representation learning. *arXiv preprint arXiv:1912.02427*, 2019.

Carsten Scherer, Pascal Gahinet, and Mahmoud Chilali. Multiobjective output-feedback control via LMI optimization. *IEEE Transactions on Automatic Control*, 42(7):896–911, 1997.

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

Ju Sun, Qing Qu, and John Wright. A geometric analysis of phase retrieval. *Foundations of Computational Mathematics*, 18:1131–1198, 2018.

Stephen Tu and Benjamin Recht. The gap between model-based and model-free methods on the linear quadratic regulator: An asymptotic viewpoint. In *Proceedings of the 32nd Conference on Learning Theory*, pages 3036–3083, 2019.

Stephen Tu, Ross Boczar, Andrew Packard, and Benjamin Recht. Non-asymptotic analysis of robust control from coarse-grained identification. *arXiv preprint arXiv:1707.04791*, 2017.

Jack Umenberger, Mina Ferizbegovic, Thomas B. Schön, and Håkan Hjalmarsson. Robust exploration in linear quadratic reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Feicheng Wang and Lucas Janson. Exact asymptotics for linear quadratic adaptive control. *arXiv preprint arXiv:2011.01364*, 2020.

A. Yousuff and R. E. Skelton. A note on balanced controller reduction. *IEEE Transactions on Automatic Control*, 29(3):254–257, 1984.

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 *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, pages 179–190, 2020a.

Yuqian Zhang, Qing Qu, and John Wright. From symmetry to geometry: Tractable nonconvex problems. *arXiv preprint arXiv:2007.06753*, 2020b.

Yang Zheng, Luca Furieri, Maryam Kamgarpour, and Na Li. Sample complexity of linear quadratic Gaussian (LQG) control for output feedback systems. *arXiv preprint arXiv:2011.09929*, 2020.

Yang Zheng, Yujie Tang, and Na Li. Analysis of the optimization landscape of linear quadratic Gaussian (LQG) control. *arXiv preprint arXiv:2102.04393*, 2021.

Kemin Zhou, John C. Doyle, and Keith Glover. *Robust and optimal control*. Prentice Hall, 1996.