



Analysis of the optimization landscape of Linear Quadratic Gaussian (LQG) control

Yujie Tang¹  · Yang Zheng²  · Na Li³ 

Received: 9 September 2021 / Accepted: 31 January 2023 / Published online: 20 March 2023

© Springer-Verlag GmbH Germany, part of Springer Nature and Mathematical Optimization Society 2023

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 dynamic 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 that are not controllable and not observable; (3) All controllable and observable 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.

The first two authors contributed to this work equally and are listed in alphabetical order. This work was supported by National Science Foundation AI Institute 2112085 and CNS 2003111, and Office of Naval Research Young Investigator Program N00014-19-1-2217. The work of Yang Zheng is supported by National Science Foundation ECCS-2154650. Yujie Tang worked on this project while he was a postdoctoral fellow at the School of Engineering and Applied Sciences at Harvard University. This paper is an extended version of the conference paper [35].

✉ Yujie Tang
yujietang@pku.edu.cn

Yang Zheng
zhengy@eng.ucsd.edu

Na Li
nali@seas.harvard.edu

¹ Department of Industrial Engineering & Management, Peking University, Beijing 100871, China

² Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA 92093, USA

³ School of Engineering and Applied Sciences, Harvard University, Allston, MA 02134, USA

Keywords Optimization landscape · Non-convex optimization · Linear quadratic Gaussian · Connectivity · Stationary points

Mathematics Subject Classification 49N10 · 90C26 · 90C30 · 93B52 · 93E20

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 [3, 9, 48]. Despite the non-convexity of the LQG problem, a globally optimal controller can be found by solving two algebraic Riccati equations [48], or a convex semidefinite program based on a change of variables [15, 25, 32].

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 problem, linear quadratic regulator (LQR), can lead to fruitful and profound results, especially for model-free controller synthesis [11, 23, 24, 27, 37, 38, 42]. 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 the gradient dominance property [6, 11]. These properties are fundamental for establishing convergence guarantees of gradient-based algorithms and their model-free RL extensions for solving LQR [24, 27]. Note that the LQR problem considers a linear system with a *fully observable* state, which can impose severe limitations for its applications in many practical scenarios where the system's state is only *partially observable* due to constraints in sensing or communication.

This paper aims to analyze the optimization landscape of the LQG problem, which considers the optimal control of a *partially observable* linear system. Unlike LQR whose optimal solution is a static feedback policy, the optimal controller of the LQG problem is no longer static. We need to search over dynamic controllers for LQG problems. This makes its optimization landscape richer and yet much more complicated than LQR. Furthermore, LQG has a natural symmetry structure induced by similarity transformations that do not change the input-output behavior of dynamic controllers, which is not the case for LQR.

Some recent studies [7, 16, 22, 30, 34] 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 [43] for a recent review. For the LQG problem, we expect that the symmetry associated with similarity transformations can bring some important properties of its non-convex optimization landscape, such as the existence of spurious stationary points, the topology of the set of globally optimal points, etc.

We also note that the notion of *controllable and observable controllers*¹ is a unique feature in controller synthesis of *partially observable* dynamical systems, making the optimization landscape of LQG distinct from many machine learning problems.

1.1 Our contributions

In this paper, we view the classical LQG problem from a modern optimization perspective, and study two aspects of its optimization landscape. First, we characterize the connectivity of the feasible region of the LQG problem, i.e., the set of strictly proper stabilizing dynamic 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 1) that are diffeomorphic under a similarity transformation (Theorem 2). We further present a sufficient condition under which \mathcal{C}_n is always connected, and this condition becomes necessary for LQG problems with a single input or a single output (Theorem 3). As a corollary, we show that \mathcal{C}_n is always connected when the plant is open-loop stable (Corollary 1).

Second, we investigate structural properties of the stationary points of the LQG cost function. By exploiting the symmetry induced by similarity transformations, we show that the LQG cost may have many *strictly suboptimal stationary points* that are not controllable and not observable (Theorem 4). For LQG with an open-loop stable plant, we explicitly construct a family of such strictly suboptimal stationary points, and investigate the eigenvalues of the corresponding Hessian (Theorem 5). In contrast, we prove that all *controllable and observable* stationary points are globally optimal to the LQG problem (Theorem 6); this can be viewed as a special case of existing results on first-order necessary conditions for optimal reduced-order controllers [48, Theorem 20.6], [17, Sect. II]. We also show that these controllable and observable stationary points are identical up to similarity transformations, and form a submanifold of dimension n^2 that has two path-connected components (Proposition 3). This result implies that if local search iterates converge to a stationary point that corresponds to a controllable and observable controller, then the algorithm has found a globally optimal solution (Corollary 3). Finally, we construct an example showing that the second-order shape of the LQG cost function can be ill-behaved around a controllable and observable stationary point in the sense that its Hessian has a very large condition number (see Example 7).

1.2 Related work

Optimization landscape of LQR: The Linear-Quadratic Regulator (LQR) has recently re-attracted increasing interest [8, 11, 24, 31, 37, 38] in the study of RL techniques for control systems. For model-free policy optimization methods, the optimization landscape of LQR is essential for establishing their performance guarantees. In [11, 24, 27], it is shown that both continuous-time and discrete-time LQR problems enjoy the gradient dominance property, and that model-free gradient-based algorithms con-

¹ Such controllers are also called *minimal* in classical control theory. In this paper, we do not use this terminology to avoid confusion with optimal controllers.

verge to the optimal LQR controller under mild conditions. The authors in [42] have examined the optimization landscape of a class of risk-sensitive state-feedback control problems and the convergence of corresponding policy optimization methods. Furthermore, it is shown in [13] that a class of finite-horizon output-feedback linear quadratic control problems also satisfies the gradient dominance property. Some recent studies have examined the connectivity of stabilizing static output feedback policies [6, 10, 12]. It is shown in [12] that the set of stabilizing static output feedback policies can be highly disconnected, which poses a significant challenge for decentralized LQR problems. For general decentralized LQR, policy optimization methods can only be guaranteed to reach some stationary point [23].

We note that many landscape properties of LQR are derived using classical control tools [10, 13, 27, 42]. Our work leverages ideas from classical control tools [15, 25, 32, 48] to analyze the optimization landscape of the LQG problem.

Reinforcement learning for LQG and controller parameterization: Recent studies have also started to investigate LQG with unknown dynamics, including offline robust control [4, 36, 44] and online adaptive control [19, 20, 33]. The line of studies on offline robust control first estimates a system model as well as a bound on the estimation error (see, e.g., [29, 36, 46]), and then design a robust LQG controller that stabilizes the plant against model uncertainty. For online adaptive control, the recent work [33] has introduced an online gradient descent algorithm to update LQG controller parameters with a sub-linear regret; see [19, 20] for further developments. For both lines of works, a convex reformulation of the LQG problem is essential for algorithm design as well as performance analysis. For example, the works [19, 20, 33] employ the classical Youla parameterization [40], while the works [4, 44] adopt the recent system-level parameterization (SLP) [39] and input-output parameterization (IOP) [14], respectively. The Youla parameterization, SLP, and IOP are able to recast the LQG problem into equivalent convex formulations in the frequency domain [45], but they all rely on the underlying system dynamics explicitly. Thus, a system identification procedure is required *a priori* in [4, 33, 36, 44], and these methods are all model-based.

In this work, we consider a natural model-free controller parameterization for LQG in the state-space domain. This parameterization does not depend on the system dynamics explicitly but leads to a non-convex formulation. Our results contribute to the understanding of this non-convex optimization landscape, which shed light on performance analysis of model-free RL methods for solving LQG.

Non-convex optimization with symmetry: Recent works [22, 43] have revealed the significance of symmetry properties in understanding the geometry of many non-convex optimization problems in machine learning. For example, the phase retrieval [34] and low-rank matrix factorization [7, 22] problems have rotational symmetries, while sparse dictionary learning [30] and tensor decomposition [16] exhibit discrete symmetries; see [43] for a recent survey. These symmetries enable identifying the local curvature of stationary points, and contribute to the tractability of the associated non-convex optimization problems.

In this paper, we highlight the symmetry defined by similarity transformations of dynamic output-feedback controllers, which enables us to derive novel results on the optimization landscape of LQG. While the notion of similarity transformation has

been extensively studied in classical control theory, its utilization in analyzing the non-convex optimization landscape of LQG (and other control problems) is limited in existing literature. We note that the symmetry defined by similarity transformations also holds for other dynamic output feedback controller design problems, suggesting that our results and analysis may be generalized or adapted to the optimization landscape analysis of other important control problems (such as \mathcal{H}_2 and \mathcal{H}_∞ optimal control).

1.3 Paper outline

In Sect. 2, we present the problem statement of Linear Quadratic Gaussian (LQG) control. We introduce our main results on the connectivity of stabilizing controllers in Sect. 3, and present our main results on the structure of stationary points of LQG problems in Sect. 4. We conclude the paper in Sect. 5. Some technical proofs are presented in the appendix.

Notations: The set of $k \times k$ real symmetric matrices is denoted by \mathbb{S}^k . The set of $k \times k$ real invertible matrices is denoted by GL_k . $\|M\|_F$ denotes the Frobenius norm for any matrix M . For any $M_1, M_2 \in \mathbb{S}^k$, we use $M_1 \prec M_2$ and $M_2 \succ M_1$ to mean that $M_2 - M_1$ is positive definite, and use $M_1 \preceq M_2$ and $M_2 \succeq M_1$ to mean that $M_2 - M_1$ is positive semidefinite. We use I_k to denote the $k \times k$ identity matrix, and use $0_{k_1 \times k_2}$ to denote the $k_1 \times k_2$ zero matrix; we sometimes omit their subscripts if the dimensions can be inferred from the context.

2 Problem statement

In this section, we first introduce the linear quadratic Gaussian control problem, and then present the problem statement of our work.

2.1 The linear quadratic Gaussian (LQG) problem

Consider a plant described by a continuous-time linear dynamical system²

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

where $x(t) \in \mathbb{R}^n$ represents the state vector, $u(t) \in \mathbb{R}^m$ represents the control input, $y(t) \in \mathbb{R}^p$ represents the output signal, and $w(t) \in \mathbb{R}^n$, $v(t) \in \mathbb{R}^p$ are process and measurement noises at time t . It is assumed that $w(t)$ and $v(t)$ are white Gaussian noises with intensity matrices $W \succeq 0$ and $V \succ 0$. For notational simplicity, we will drop the argument t when it is clear in the context.

² This paper focuses on the continuous-time setup. Discussion and results for discrete-time LQG are provided in our online report [47].

The classical linear quadratic Gaussian (LQG) problem is formulated as

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

where $Q \succeq 0$ and $R \succ 0$. Here, the input $u(t)$ is allowed to depend on all past observation $y(\tau)$ with $\tau < t$. We make the following standard assumption on the problem setup throughout the paper.

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) problem, static feedback policies in general do not achieve the optimal objective value, and we need to consider the class of dynamic controllers in the form of

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

Here $\xi(t) \in \mathbb{R}^q$ is the internal state and A_K, B_K, C_K are matrices that specify the controller's dynamics. We refer to the dimension q of the internal state ξ as the order of the dynamic controller (3). A dynamic controller is said to be *full-order* if $q = n$, and is said to be *reduced-order* if $q < n$. We shall see later that it is unnecessary to consider dynamic controllers with order beyond the system dimension n .

The LQG Problem (2) admits the celebrated separation principle and has a closed-form solution by solving two algebraic Riccati equations [48, Theorem 14.7]. Specifically, the optimal controller is given by

$$\begin{aligned} \dot{\xi} &= (A - BK)\xi + L(y - C\xi), \\ u &= -K\xi. \end{aligned} \quad (4)$$

Here the matrix L is called the Kalman gain which is given by $L = PC^\top V^{-1}$ with P being the unique positive semidefinite solution (see, e.g., [48, Corollary 13.8]) to

$$AP + PA^\top - PC^\top V^{-1} CP + W = 0, \quad (5a)$$

and the matrix K is called the feedback gain, given by $K = R^{-1}B^\top S$ where S is the unique positive semidefinite solution to

$$A^\top S + SA - SBR^{-1}B^\top S + Q = 0. \quad (5b)$$

We see that the optimal LQG controller (4) can be written in the form of (3) with

$$A_K = A - BK - LC, \quad B_K = L, \quad C_K = -K. \quad (6)$$

Thus, the solution from Riccati equations (5) is always full-order, i.e., $q = n$. We note that two dynamic controllers with the same transfer function $\mathbf{K}(s) = C_K(sI - A_K)^{-1}B_K$ lead to the same LQG cost. In general, the optimal LQG controller is only unique in the frequency domain [48, Theorem 14.7] but not unique in the state-space domain; any similarity transformation on (6) leads to another optimal solution that achieves the global minimum cost (see Lemma 7).

2.2 Parameterization of dynamic controllers and the LQG cost function

Recently, model-free reinforcement learning methods have been studied for a range of control problems, such as LQR [11, 27], finite-horizon discrete-time LQG [13], state-feedback risk-sensitive control [42], etc. These works view classical control problems from a modern optimization perspective, and directly optimize over policies based on observed data, without explicit knowledge of the underlying model. In this paper, we adopt a similar angle and view LQG from a model-free optimization perspective.

We consider the natural parameterization of the set of dynamic controllers in (3) by their corresponding matrices (A_K, B_K, C_K) . To formulate the LQG cost as a function of the parameterized dynamic controller (A_K, B_K, C_K) , we first need to specify its domain. By combining (3) with (1), we get the closed-loop system

$$\begin{aligned} \frac{d}{dt} \begin{bmatrix} x \\ \xi \end{bmatrix} &= \begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix} \begin{bmatrix} x \\ \xi \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & B_K \end{bmatrix} \begin{bmatrix} w \\ v \end{bmatrix}, \\ \begin{bmatrix} y \\ u \end{bmatrix} &= \begin{bmatrix} C & 0 \\ 0 & C_K \end{bmatrix} \begin{bmatrix} x \\ \xi \end{bmatrix} + \begin{bmatrix} v \\ 0 \end{bmatrix}. \end{aligned} \quad (7)$$

It is known from classical control theory [48, Chapter 13] that under Assumption 1, the LQG cost is finite if the closed-loop matrix

$$\begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix} = \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix} \quad (8)$$

is stable, i.e., the real parts of all its eigenvalues are negative; dynamic controllers satisfying this condition are said to *internally stabilize* the plant (1). Furthermore, the optimal controller given by (6) is guaranteed to internally stabilize the plant. We therefore define the set of stabilizing controllers with order $q \in \mathbb{N}$ by³

$$\mathcal{C}_q := \left\{ \mathbf{K} = \begin{bmatrix} 0_{m \times p} & C_K \\ B_K & A_K \end{bmatrix} \in \mathbb{R}^{(m+q) \times (p+q)} \middle| (8) \text{ is stable} \right\}, \quad (9)$$

and let $J_q : \mathcal{C}_q \rightarrow \mathbb{R}$ be the function that maps a parameterized dynamic controller in \mathcal{C}_q to its corresponding LQG cost for each $q \in \mathbb{N}$. Since the set of *full-order* stabilizing

³ In (9), for notational simplicity, we lumped the controller parameters into a single matrix; but it should be interpreted as a dynamic controller, represented by (3). Note that this definition allows us to apply block-wise matrix operations; see, e.g., (14).

controllers \mathcal{C}_n contains the optimal controller, we will mainly focus on the properties of \mathcal{C}_n in this paper. We will abbreviate $J_n(\mathbf{K})$ as $J(\mathbf{K})$ when no confusions occur.

The following lemma shows that the set \mathcal{C}_q can be treated as an open set when it is nonempty. This is a direct consequence of the fact that the Routh–Hurwitz stability criterion returns a set of strict polynomial inequalities in terms of the elements of $(A_{\mathbf{K}}, B_{\mathbf{K}}, C_{\mathbf{K}})$.

Lemma 1 *Let $q \geq 1$ such that \mathcal{C}_q is nonempty. Then, \mathcal{C}_q is an open subset of the linear space*

$$\mathcal{V}_q := \left\{ \begin{bmatrix} D_{\mathbf{K}} & C_{\mathbf{K}} \\ B_{\mathbf{K}} & A_{\mathbf{K}} \end{bmatrix} \in \mathbb{R}^{(m+q) \times (p+q)} \mid D_{\mathbf{K}} = 0_{m \times p} \right\}. \quad (10)$$

We also have the following observation on the set of full-order stabilizing controllers \mathcal{C}_n , whose proof is postponed to Appendix A.1.

Lemma 2 *The set \mathcal{C}_n is non-empty, unbounded, and can be non-convex.*

The following two lemmas give useful characterizations of the LQG cost function J_q . Lemma 3 is known in the literature (see, e.g., [2]); Lemma 4 follows directly from Lemma 3, and we provide a short proof in Appendix A.2.

Lemma 3 *Fix $q \in \mathbb{N}$ such that $\mathcal{C}_q \neq \emptyset$. Given $\mathbf{K} \in \mathcal{C}_q$, we have*

$$J_q(\mathbf{K}) = \text{tr} \left(\begin{bmatrix} Q & 0 \\ 0 & C_{\mathbf{K}}^T R C_{\mathbf{K}} \end{bmatrix} X_{\mathbf{K}} \right) = \text{tr} \left(\begin{bmatrix} W & 0 \\ 0 & B_{\mathbf{K}} V B_{\mathbf{K}}^T \end{bmatrix} Y_{\mathbf{K}} \right), \quad (11)$$

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

$$\begin{bmatrix} A & BC_{\mathbf{K}} \\ B_{\mathbf{K}} C & A_{\mathbf{K}} \end{bmatrix} X_{\mathbf{K}} + X_{\mathbf{K}} \begin{bmatrix} A & BC_{\mathbf{K}} \\ B_{\mathbf{K}} C & A_{\mathbf{K}} \end{bmatrix}^T + \begin{bmatrix} W & 0 \\ 0 & B_{\mathbf{K}} V B_{\mathbf{K}}^T \end{bmatrix} = 0, \quad (12a)$$

$$\begin{bmatrix} A & BC_{\mathbf{K}} \\ B_{\mathbf{K}} C & A_{\mathbf{K}} \end{bmatrix}^T Y_{\mathbf{K}} + Y_{\mathbf{K}} \begin{bmatrix} A & BC_{\mathbf{K}} \\ B_{\mathbf{K}} C & A_{\mathbf{K}} \end{bmatrix} + \begin{bmatrix} Q & 0 \\ 0 & C_{\mathbf{K}}^T R C_{\mathbf{K}} \end{bmatrix} = 0. \quad (12b)$$

Lemma 4 *For any $q \in \mathbb{N}$ with $\mathcal{C}_q \neq \emptyset$, the function J_q is real analytic on \mathcal{C}_q .*

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

$$\begin{aligned} & \min_{\mathbf{K}} J_n(\mathbf{K}) \\ & \text{subject to } \mathbf{K} \in \mathcal{C}_n. \end{aligned} \quad (13)$$

Based on (13), one may further derive model-free policy gradient algorithms to find a solution to (13). To characterize the performance of policy gradient algorithms, it is necessary to understand the landscape of (13). However, beyond Lemmas 1,2 and 4, 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 focus on the following two topics about the set \mathcal{C}_n and the LQG cost function J_n :

1. *The connectivity of \mathcal{C}_n and its implications* (Sect. 3). Connectivity of the domain is critical for performance analysis of gradient-based algorithms for model-free controller synthesis. Most recent results focus on state-feedback controllers or static output-feedback controllers [6, 11, 12, 27]. It is known that the set of stabilizing state-feedback controllers is connected, which is crucial for gradient-based algorithms to find a good solution. It is also known that the set of stabilizing static output-feedback controllers can be highly disconnected [12]. The connectivity of the set of stabilizing dynamic controllers \mathcal{C}_n , however, has not been discussed before in the literature.
2. *The structure of the stationary points and the global optimum of J_n* (Sect. 4). Classical control theory shows that the optimal feedback gain for LQR is unique under mild assumptions, and recently it is established that the LQR cost function is gradient dominated and has a unique stationary point which is the globally optimal solution [11, 27]. It has also been shown recently that a class of output-feedback controller design problems in finite-time horizon has a unique stationary point [13]. On the other hand, due to the non-uniqueness of optimal LQG controllers in the state-space domain, we do not expect the LQG cost function $J_n(\mathbf{K})$ to have a unique stationary point. We aim to reveal further structural properties of the stationary points of $J_n(\mathbf{K})$ in this work.

3 Connectivity of the set of stabilizing controllers

In this section, we examine the connectivity of the set of stabilizing controllers \mathcal{C}_n . We summarize the main results regarding the connectivity of \mathcal{C}_n in Sect. 3.1, and provide their proofs in the subsequent subsections.

3.1 Main results

We first introduce the notion of similarity transformation that is central in linear control theory. Given $q \geq 1$ such that $\mathcal{C}_q \neq \emptyset$, we define the mapping $\mathcal{T}_q : \text{GL}_q \times \mathcal{C}_q \rightarrow \mathcal{C}_q$ that represents similarity transformations on \mathcal{C}_q by

$$\mathcal{T}_q(T, \mathbf{K}) := \begin{bmatrix} I_m & 0 \\ 0 & T \end{bmatrix} \begin{bmatrix} 0 & C_{\mathbf{K}} \\ B_{\mathbf{K}} & A_{\mathbf{K}} \end{bmatrix} \begin{bmatrix} I_p & 0 \\ 0 & T \end{bmatrix}^{-1} = \begin{bmatrix} 0 & C_{\mathbf{K}} T^{-1} \\ T B_{\mathbf{K}} & T A_{\mathbf{K}} T^{-1} \end{bmatrix} \quad (14)$$

(recall that GL_q denotes the set of $k \times k$ real invertible matrices). It is not hard to verify that given $\mathbf{K} \in \mathcal{C}_q$, $\mathcal{T}_q(T, \mathbf{K})$ is also in \mathcal{C}_q for any $T \in \text{GL}_q$. We can also check that \mathcal{T}_q is indefinitely differentiable on $\text{GL}_q \times \mathcal{C}_q$, and that

$$\mathcal{T}_q(T_2, \mathcal{T}_q(T_1, \mathbf{K})) = \mathcal{T}_q(T_2 T_1, \mathbf{K}) \quad (15)$$

for any $T_1, T_2 \in \text{GL}_q$. This implies that for any fixed $T \in \text{GL}_n$, the mapping $\mathbf{K} \mapsto \mathcal{T}_q(T, \mathbf{K})$ admits an inverse given by $\mathbf{K} \mapsto \mathcal{T}_q(T^{-1}, \mathbf{K})$. Therefore, we have the following result.

Lemma 5 *Given $q \geq 1$ such that $\mathcal{C}_q \neq \emptyset$, for any $T \in \text{GL}_q$, the mapping $\mathbf{K} \mapsto \mathcal{T}_q(T, \mathbf{K})$ is a diffeomorphism from \mathcal{C}_q to itself.*

This section will mainly focus on the case $q = n$. For notational simplicity, for any $T \in \text{GL}_n$, we let $\mathcal{T}_T : \mathcal{C}_n \rightarrow \mathcal{C}_n$ denote the linear mapping given by $\mathcal{T}_T(\mathbf{K}) := \mathcal{T}_n(T, \mathbf{K})$.

We are now ready to present the main technical results.

Theorem 1 *The set \mathcal{C}_n has at most two path-connected components.*

Theorem 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 \mathcal{T}_T for any invertible matrix $T \in \mathbb{R}^{n \times n}$ with $\det T < 0$.*

Theorem 2 shows that even if \mathcal{C}_n has two path-connected components, there exists a linear bijection defined by a similarity transformation \mathcal{T}_T between these two components; this linear bijection will be orthogonal if T is orthogonal with $\det T = -1$. The following theorem then gives a sufficient condition for \mathcal{C}_n to be path-connected; this condition becomes necessary when the plant is single-input or single-output.

Theorem 3 *The following statements hold.*

1. *\mathcal{C}_n is path-connected if there exists a reduced-order stabilizing controller, i.e., $\mathcal{C}_{n-1} \neq \emptyset$.*
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 $\mathcal{C}_{n-1} \neq \emptyset$.*

One main idea in our proofs is based on a classical change of variables for dynamic controllers (see, e.g., [25, 32]). 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 follow from the fact that convex sets are path-connected. The potential disconnectivity of \mathcal{C}_n comes from the fact that the set of real invertible matrices $\text{GL}_n = \{\Pi \in \mathbb{R}^{n \times n} \mid \det \Pi \neq 0\}$ has two path-connected components [21]: $\text{GL}_n^+ = \{\Pi \in \mathbb{R}^{n \times n} \mid \det \Pi > 0\}$, $\text{GL}_n^- = \{\Pi \in \mathbb{R}^{n \times n} \mid \det \Pi < 0\}$. The proof of Theorem 3 is based on the observation that a reduced-order controller can be augmented to a full-order controller that is invariant under a similarity transformation with $\det T < 0$; for single-input or single-output plants, we use the determinant of the observability or the controllability matrix of the controller to characterize whether its order can be reduced. The full proofs are technically involved, and we postpone them to Sects. 3.2 to 3.4.

Example 1 (Disconnectivity of stabilizing controllers) Given any open-loop unstable plant with state dimension $n = 1$, it is straightforward to see that there exist no reduced-order stabilizing controllers, i.e., $\mathcal{C}_{n-1} = \emptyset$. Thus, Theorem 3 indicates that its associated set of stabilizing controllers \mathcal{C}_n is not path-connected.

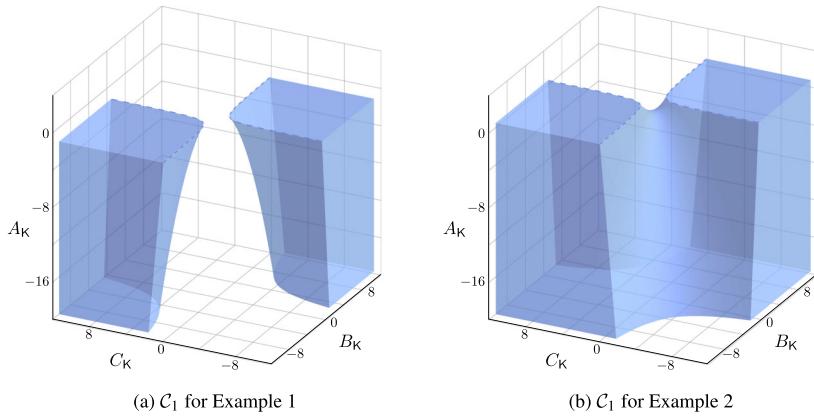


Fig. 1 The set of stabilizing controllers \mathcal{C}_1 for Examples 1 and 2

As an example, consider the plant with $A = 1$, $B = 1$, $C = 1$, and Theorem 3 indicates that the corresponding \mathcal{C}_n is not path-connected. Indeed, using the Routh–Hurwitz stability criterion, it is straightforward to derive that

$$\mathcal{C}_1 = \left\{ \mathbf{K} = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathbb{R}^{2 \times 2} \mid A_K < -1, B_K C_K < A_K \right\}. \quad (16)$$

This set has two path-connected components: $\mathcal{C}_1 = \mathcal{C}_1^+ \cup \mathcal{C}_1^-$ with $\mathcal{C}_1^+ \cap \mathcal{C}_1^- = \emptyset$, where

$$\mathcal{C}_1^\pm := \left\{ \mathbf{K} = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathbb{R}^{2 \times 2} \mid A_K < -1, B_K C_K < A_K, \pm B_K > 0 \right\}.$$

In addition, as expected from Theorem 2, it is easy to verify that \mathcal{C}_1^+ and \mathcal{C}_1^- are diffeomorphic under the mapping \mathcal{T}_T for any $T < 0$. Fig. 1a illustrates the region of the set \mathcal{C}_1 in (16). \square

Theorem 3 also suggests the following corollary.

Corollary 1 *Given any open-loop stable plant (1), the corresponding set of stabilizing controllers \mathcal{C}_n is path-connected.*

Proof This corollary follows from the fact that for any open-loop stable plant, the reduced-order controller $(A_K, B_K, C_K) = (-I_{n-1}, 0_{(n-1) \times p}, 0_{m \times (n-1)})$ is internally stabilizing, showing that $\mathcal{C}_{n-1} \neq \emptyset$. \square

Example 2 (Stabilizing controllers for an open-loop stable system) Consider an open-loop stable plant (1) with $A = -1$, $B = 1$, $C = 1$. Since it is open-loop stable, Corollary 1 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

$$\mathcal{C}_1 = \left\{ \mathbf{K} = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathbb{R}^{2 \times 2} \mid A_K < 1, B_K C_K < -A_K \right\}. \quad (17)$$

This set is path-connected, as illustrated in Fig. 1b. \square

Remark 1 We provide some remarks on the implications of the connectivity of the domain of LQR/LQG for gradient-based algorithms. Recent studies have revisited the classical LQR problem from a modern optimization perspective and designed policy gradient algorithms for model-free controller synthesis [11, 27, 42]. For policy gradient algorithms, the connectivity of the domain (the set of stabilizing controllers) becomes important since gradient-based methods typically cannot jump between different connected components. It is known that the set of stabilizing static state-feedback gains $\{K \in \mathbb{R}^{m \times n} \mid A - BK \text{ is stable}\}$ is connected [6], and this is one critical factor in justifying the performance of the algorithms in [11, 27, 42]. On the other hand, the set of stabilizing static output feedback policies $\{K \in \mathbb{R}^{m \times p} \mid A - BKC \text{ is stable}\}$ can be highly disconnected [12], posing a significant challenge for gradient-based algorithms. In Theorems 1 to 3, we have shown that the set of stabilizing controllers \mathcal{C}_n for LQG has at most two path-connected components that are diffeomorphic to each other under some similarity transformation. Since similarity transformation does not change the input/output behavior of a controller, it makes no difference to search over either path-connected component in \mathcal{C}_n even if \mathcal{C}_n is not path-connected. This brings positive news to gradient-based local search algorithms for LQG.

3.2 Proof of Theorem 1

The basic idea of analyzing the path-connectivity of \mathcal{C}_n for LQG is in some sense similar to the analysis for LQR [6]: We first adopt a classical change of variables for constructing convex reformulation of the controller synthesis problem, and then path-connectivity results generally follow from the path-connectivity of convex sets. But compared to the analysis for LQR, here we need to use a more complicated change of variables for dynamic controllers in the state-space domain.

Specifically, we adopt the change of variables presented in [25, 32]. Given the plant dynamics (A, B, C) in (1), we first introduce the following convex set

$$\begin{aligned} \mathcal{F}_n := & \left\{ (X, Y, M, H, F) \mid X, Y \in \mathbb{S}^n, M \in \mathbb{R}^{n \times n}, H \in \mathbb{R}^{n \times p}, F \in \mathbb{R}^{m \times n}, \right. \\ & \left. \begin{bmatrix} X & I \\ I & Y \end{bmatrix} \succ 0, \begin{bmatrix} AX+BF & A \\ M & YA+HC \end{bmatrix} + \begin{bmatrix} AX+BF & A \\ M & YA+HC \end{bmatrix}^\top \prec 0 \right\}, \end{aligned} \quad (18)$$

and the “extended” set

$$\mathcal{G}_n := \left\{ Z = (X, Y, M, H, F, \Pi, \Xi) \mid \begin{array}{l} (X, Y, M, H, F) \in \mathcal{F}_n, \\ \Pi, \Xi \in \mathbb{R}^{n \times n}, \Xi \Pi = I - YX \end{array} \right\}. \quad (19)$$

We shall later see that there exists a continuous surjective map from \mathcal{G}_n to \mathcal{C}_n , and the path-connectivity of the convex set \mathcal{F}_n plays a key role in analyzing the path-connected components of \mathcal{C}_n . Before proceeding, we note the following observation for each element in \mathcal{G}_n .

Lemma 6 For any $(X, Y, M, H, F, \Pi, \Xi) \in \mathcal{G}_n$, Π and Ξ are always invertible, and consequently, the block triangular matrices $\begin{bmatrix} I & 0 \\ YB & \Xi \end{bmatrix}$ and $\begin{bmatrix} I & CX \\ 0 & \Pi \end{bmatrix}$ are invertible.

Proof By definition, for all $(X, Y, W, H, F, \Pi, \Xi) \in \mathcal{G}_n$, we have $\begin{bmatrix} X & I \\ I & Y \end{bmatrix} \succ 0$, implying that

$$\det(YX - I) = \det X \det(Y - X^{-1}) = \det \begin{bmatrix} X & I \\ I & Y \end{bmatrix} > 0.$$

Thus, $\det(\Pi) \neq 0$ and $\det(\Xi) \neq 0$, indicating they are both invertible. The invertibility of the other two block triangular matrices is straightforward. \square

We now define a mapping from \mathcal{G}_n to a subset of $\mathbb{R}^{(m+n) \times (p+n)}$.

Definition 1 For each $Z = (X, Y, M, H, F, \Pi, \Xi)$ in \mathcal{G}_n , let

$$\Phi(Z) = \begin{bmatrix} 0 & \Phi_C(Z) \\ \Phi_B(Z) & \Phi_A(Z) \end{bmatrix} := \begin{bmatrix} I & 0 \\ YB & \Xi \end{bmatrix}^{-1} \begin{bmatrix} 0 & F \\ H & M - YAX \end{bmatrix} \begin{bmatrix} I & CX \\ 0 & \Pi \end{bmatrix}^{-1}. \quad (20)$$

We point out that this mapping (20) is essentially the change of variables presented in [25, 32], which is critical for deriving convex reformulations of output-feedback controller synthesis problems. The following result builds an explicit connection between \mathcal{G}_n and \mathcal{C}_n via the mapping Φ .

Proposition 1 The mapping Φ defined by (20) is a continuous and surjective mapping from \mathcal{G}_n to \mathcal{C}_n .

Proposition 1 has been effectively proved in [25, 32]. We provide a rigorous proof in Appendix A.3.

After establishing the continuous surjection from \mathcal{G}_n to \mathcal{C}_n , it is now clear that we can study the path-connectivity of \mathcal{C}_n via the path-connectivity of \mathcal{G}_n : Any continuous path in \mathcal{G}_n will be mapped to a continuous path in \mathcal{C}_n , and thus any path-connected component of \mathcal{G}_n has a path-connected image under the mapping Φ . Consequently, the number of path-connected components of \mathcal{C}_n will be no more than the number of path-connected components of \mathcal{G}_n .

We now proceed to provide results on the path-connectivity of the set \mathcal{G}_n .

Proposition 2 The set \mathcal{G}_n has two path-connected components, given by

$$\begin{aligned} \mathcal{G}_n^+ &= \{(X, Y, M, H, F, \Pi, \Xi) \in \mathcal{G}_n \mid \det \Pi > 0\}, \\ \mathcal{G}_n^- &= \{(X, Y, M, H, F, \Pi, \Xi) \in \mathcal{G}_n \mid \det \Pi < 0\}. \end{aligned}$$

Proof First, the convexity of \mathcal{F}_n implies that the set \mathcal{F}_n is path-connected. We then notice that the set of real invertible matrices $\text{GL}_n = \{\Pi \in \mathbb{R}^{n \times n} \mid \det \Pi \neq 0\}$ has two path-connected components [21]

$$\text{GL}_n^+ = \{\Pi \in \mathbb{R}^{n \times n} \mid \det \Pi > 0\}, \quad \text{GL}_n^- = \{\Pi \in \mathbb{R}^{n \times n} \mid \det \Pi < 0\}.$$

Therefore the Cartesian product $\mathcal{F}_n \times \mathrm{GL}_n$ has two path-connected components. Finally, it is not hard to verify that the following mapping

$$(X, Y, M, H, F, \Pi) \mapsto (X, Y, M, H, F, \Pi, (I - YX)\Pi^{-1})$$

is a homeomorphism from $\mathcal{F}_n \times \mathrm{GL}_n$ to \mathcal{G}_n . Therefore \mathcal{G}_n also has two path-connected components, and their expressions are evident. \square

Proposition 2 then implies that \mathcal{C}_n has at most two path-connected components. Precisely, upon defining $\mathcal{C}_n^+ = \Phi(\mathcal{G}_n^+)$, $\mathcal{C}_n^- = \Phi(\mathcal{G}_n^-)$, the two path-connected components of \mathcal{C}_n are just given by \mathcal{C}_n^+ and \mathcal{C}_n^- , if \mathcal{C}_n is not path-connected. This completes the proof of Theorem 1.

3.3 Proof of Theorem 2

We have shown in the previous subsection that \mathcal{C}_n^+ and \mathcal{C}_n^- are the two path-connected components if \mathcal{C}_n is not connected. To prove Theorem 2, it suffices to show that, regardless of the path-connectivity of \mathcal{C}_n , for any $T \in \mathbb{R}^{n \times n}$ with $\det T < 0$, the mapping \mathcal{T}_T restricted on \mathcal{C}_n^+ is a diffeomorphism from \mathcal{C}_n^+ to \mathcal{C}_n^- . And since \mathcal{T}_T is a diffeomorphism from \mathcal{C}_n to itself with inverse $\mathcal{T}_{T^{-1}}$, and \mathcal{C}_n^+ and \mathcal{C}_n^- are two open subsets of \mathcal{C}_n , we only need to show that $\mathcal{T}_T(\mathcal{C}_n^+) \subseteq \mathcal{C}_n^-$ and $\mathcal{T}_{T^{-1}}(\mathcal{C}_n^-) \subseteq \mathcal{C}_n^+$ when $\det T < 0$.

Consider an arbitrary point $K = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathcal{C}_n^+$. By the definition of \mathcal{C}_n^+ , there exists $Z = (X, Y, M, H, F, \Pi, \mathcal{E}) \in \mathcal{G}_n^+$ such that $\Phi(Z) = K$. Now let

$$\hat{\Pi} = T\Pi, \quad \hat{\mathcal{E}} = \mathcal{E}T^{-1}, \quad \hat{Z} = (X, Y, M, H, F, \hat{\Pi}, \hat{\mathcal{E}}).$$

It is not difficult to verify that $\hat{Z} \in \mathcal{G}_n$. Since $\det \hat{\Pi} = \det T \cdot \det \Pi < 0$, we have $\hat{Z} \in \mathcal{G}_n^-$. Then,

$$\begin{aligned} \Phi(\hat{Z}) &= \begin{bmatrix} 0 & \Phi_C(\hat{Z}) \\ \Phi_B(\hat{Z}) & \Phi_A(\hat{Z}) \end{bmatrix} = \begin{bmatrix} I & 0 \\ YB & \hat{\mathcal{E}} \end{bmatrix}^{-1} \begin{bmatrix} 0 & F \\ H & M - YAX \end{bmatrix} \begin{bmatrix} I & CX \\ 0 & \hat{\Pi} \end{bmatrix}^{-1} \\ &= \begin{bmatrix} I & 0 \\ 0 & T \end{bmatrix} \begin{bmatrix} I & 0 \\ YB & \mathcal{E} \end{bmatrix}^{-1} \begin{bmatrix} 0 & F \\ H & M - YAX \end{bmatrix} \begin{bmatrix} I & CX \\ 0 & \Pi \end{bmatrix}^{-1} \begin{bmatrix} I & 0 \\ 0 & T^{-1} \end{bmatrix} \\ &= \begin{bmatrix} I & 0 \\ 0 & T \end{bmatrix} \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & T^{-1} \end{bmatrix} = \begin{bmatrix} 0 & C_K T^{-1} \\ TB_K & TA_K T^{-1} \end{bmatrix} = \mathcal{T}_T(K), \end{aligned}$$

which implies that $\mathcal{T}_T(K) \in \Phi(\mathcal{G}_n^-) = \mathcal{C}_n^-$ and consequently $\mathcal{T}_T(\mathcal{C}_n^+) \subseteq \mathcal{C}_n^-$.

The proof of $\mathcal{T}_{T^{-1}}(\mathcal{C}_n^-) \subseteq \mathcal{C}_n^+$ is similar by noting that $\det T^{-1} < 0$ if and only if $\det T < 0$.

3.4 Proof of Theorem 3

We first show that the non-emptiness of \mathcal{C}_{n-1} implies the path-connectivity of \mathcal{C}_n . Indeed, suppose there exists $\tilde{K} \in \mathcal{C}_{n-1}$. Then it can be augmented to be a full-order controller in \mathcal{C}_n by

$$K = \begin{bmatrix} 0 & \tilde{C}_K & 0 \\ \tilde{B}_K & \tilde{A}_K & 0 \\ 0 & 0 & -1 \end{bmatrix} \in \mathcal{C}_n.$$

Now let $T = \begin{bmatrix} I_{n-1} & 0 \\ 0 & -1 \end{bmatrix}$. By the proof of Theorem 2, we can see that $K \in \mathcal{C}_n^\pm$ implies $\mathcal{T}_T(K) \in \mathcal{C}_n^\mp$. On the other hand, we can directly check that $\mathcal{T}_T(K) = K$. Therefore we have $K \in \mathcal{C}_n^+ \cap \mathcal{C}_n^-$, indicating that $\mathcal{C}_n^+ \cap \mathcal{C}_n^-$ is nonempty. Consequently, \mathcal{C}_n is path-connected.

We then consider the case when the plant is single-input or single-output. The goal is to find a reduced-order controller in \mathcal{C}_{n-1} when \mathcal{C}_n is connected. Here we only prove the single-output case; the single-input case can be proved similarly.

Let T be any real $n \times n$ matrix with $\det T < 0$. Let $K^{(0)} \in \mathcal{C}_n$ be arbitrary, and let $K^{(1)} = \mathcal{T}_T(K^{(0)})$. If \mathcal{C}_n is path-connected, then there exists a continuous path

$$K(t) = \begin{bmatrix} 0 & C_K(t) \\ B_K(t) & A_K(t) \end{bmatrix}, \quad t \in [0, 1]$$

in \mathcal{C}_n such that $K(0) = K^{(0)}$ and $K(1) = K^{(1)}$. Now for each $t \in [0, 1]$, let $C(t)$ be the controllability matrix for $(A_K(t), B_K(t))$, i.e.,

$$C(t) = [B_K(t) \ A_K(t)B_K(t) \ \cdots \ A_K(t)^{n-1}B_K(t)] \in \mathbb{R}^{n \times n},$$

where the dimension of $C(t)$ is $n \times n$ since the plant is single-output (i.e., the controller is single-input). We then have $C(1) = T C(0)$, and thus $\det C(1) \cdot \det C(0) \leq 0$. On the other hand, it can be seen that $\det C(t)$ is a continuous function over $t \in [0, 1]$. Therefore $\det C(\tau) = 0$ for some $\tau \in [0, 1]$, implying that $(A_K(\tau), B_K(\tau))$ is not controllable. This indicates that the transfer function $C_K(\tau)(sI_n - A_K(\tau))^{-1}B_K(\tau)$ can be realized by a state-space representation with dimension at most $n - 1$, and consequently $\mathcal{C}_{n-1} \neq \emptyset$.

4 Structure of stationary points

In this section, we proceed to characterize the stationary points of the LQG cost function. Section 4.1 discusses the invariance of the LQG cost J_q under similarity transformation and its implications. Section 4.2 shows how to compute the gradient and the Hessian of the LQG cost J_q . In Sect. 4.3, we investigate a class of spurious stationary points that are not controllable and not observable. We characterize the

controllable and observable stationary points for LQG over \mathcal{C}_n in Sect. 4.4. Finally, in Sect. 4.5, we discuss the second-order behavior of $J_n(\mathbf{K})$ around its controllable and observable stationary points.

4.1 Invariance of LQG cost under similarity transformation

As shown in Lemma 5, the similarity transformation $\mathcal{T}_q(T, \cdot)$ is a diffeomorphism from \mathcal{C}_q to itself for any invertible matrix $T \in \text{GL}_q$. Then together with (15), we can see that the set of similarity transformations is a group that is isomorphic to GL_q . We can therefore define the *orbit* of $\mathbf{K} \in \mathcal{C}_q$ by

$$\mathcal{O}_{\mathbf{K}} := \{\mathcal{T}_q(T, \mathbf{K}) \mid T \in \text{GL}_q\}.$$

Since the controllers $\mathbf{K} = \begin{bmatrix} 0 & C_{\mathbf{K}} \\ B_{\mathbf{K}} & A_{\mathbf{K}} \end{bmatrix}$ and $\mathcal{T}_q(T, \mathbf{K}) = \begin{bmatrix} 0 & C_{\mathbf{K}}T^{-1} \\ TB_{\mathbf{K}} & TA_{\mathbf{K}}T^{-1} \end{bmatrix}$ have identical input-output behavior regardless of $T \in \text{GL}_q$, we have the following lemma that the LQG cost is invariant under similarity transformations.

Lemma 7 *Let $q \geq 1$ such that $\mathcal{C}_q \neq \emptyset$. Then $J_q(\mathbf{K}) = J_q(\mathcal{T}_q(T, \mathbf{K}))$ for any $\mathbf{K} \in \mathcal{C}_q$ and $T \in \text{GL}_q$.*

Consequently, the LQG cost is constant over an orbit $\mathcal{O}_{\mathbf{K}}$ for any $\mathbf{K} \in \mathcal{C}_q$.

Now consider $\mathbf{K} = \begin{bmatrix} 0 & C_{\mathbf{K}} \\ B_{\mathbf{K}} & A_{\mathbf{K}} \end{bmatrix} \in \mathcal{C}_q$ such that $(A_{\mathbf{K}}, B_{\mathbf{K}})$ is controllable and $(C_{\mathbf{K}}, A_{\mathbf{K}})$ is observable. The following proposition shows that every orbit $\mathcal{O}_{\mathbf{K}}$ corresponding to a controllable and observable controller has dimension q^2 with two path-connected components. The proof is given in Appendix A.4.

Proposition 3 *Suppose $\mathbf{K} \in \mathcal{C}_q$ represents a controllable and observable controller. Then the orbit $\mathcal{O}_{\mathbf{K}}$ is a submanifold of \mathcal{C}_q of dimension q^2 , and has two path-connected components given by*

$$\mathcal{O}_{\mathbf{K}}^+ = \{\mathcal{T}_q(T, \mathbf{K}) \mid \det T > 0\}, \quad \mathcal{O}_{\mathbf{K}}^- = \{\mathcal{T}_q(T, \mathbf{K}) \mid \det T < 0\}.$$

From Lemma 7 and Proposition 3, one interesting consequence is that given a globally optimal LQG controller $\mathbf{K}^* \in \mathcal{C}_n$, the points in its orbit $\mathcal{O}_{\mathbf{K}^*}$ are all globally optimal, and if \mathbf{K}^* is controllable and observable, the orbit $\mathcal{O}_{\mathbf{K}^*}$ is a submanifold in \mathcal{V}_n of dimension n^2 , and it has two path-connected components. Figure 2 demonstrates the orbits of globally optimal LQG controllers for an open-loop unstable plant and an open-loop stable plant, showing that the set of globally optimal LQG controllers are non-isolated and disconnected in \mathcal{C}_n .

We conclude this subsection by noting that the LQG cost function $J_q(\mathbf{K})$ is not coercive in the sense that there might exist sequences of stabilizing controllers $\mathbf{K}_j \in \mathcal{C}_q$ with $\|\mathbf{K}_j\|_F \rightarrow +\infty$ or $\inf_{\mathbf{K}' \in \partial \mathcal{C}_q} \|\mathbf{K}_j - \mathbf{K}'\|_F \rightarrow 0$ such that $\limsup_{j \rightarrow \infty} J_q(\mathbf{K}_j)$ is finite. Indeed, from Proposition 3, the orbit $\mathcal{O}_{\mathbf{K}}$ can be unbounded while $J_q(\mathbf{K})$ is constant for any controller in the same orbit. Furthermore, the following example

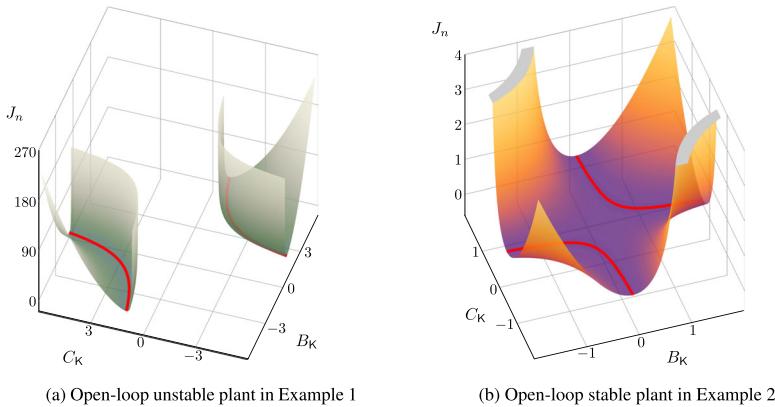


Fig. 2 Non-isolated and disconnected globally optimal LQG controllers. In both cases, we set $Q = 1$, $R = 1$, $V = 1$, $W = 1$. (a) LQG cost for Example 1 when fixing $A_K = -1 - 2\sqrt{2}$. (b) LQG cost for Example 2 when fixing $A_K = 1 - 2\sqrt{2}$. The red curves represent the set of globally optimal LQG controllers

shows that the LQG cost might converge to a finite value even when the controller K approaches the boundary of \mathcal{C}_q .

Example 3 (Non-coercivity of the LQG cost) Consider the plant in Example 2 given by $A = -1$, $B = 1$, $C = 1$, and we let $Q = 1$, $R = 1$, $V = 1$, $W = 1$. The set \mathcal{C}_1 is given by (17). Let $K_\epsilon = \begin{bmatrix} 0 & \epsilon \\ -\epsilon & 0 \end{bmatrix} \in \mathcal{C}_1$ for any $\epsilon \neq 0$, and it is not hard to see that $\lim_{\epsilon \rightarrow 0} K_\epsilon \in \partial \mathcal{C}_1$. On the other hand, by solving the Lyapunov Eq (12a), we get the unique solution $X_{K_\epsilon} = \frac{1}{2} \begin{bmatrix} \epsilon^2 + 1 & \epsilon \\ \epsilon & \epsilon^2 + 2 \end{bmatrix}$ and the corresponding LQG cost $J(K_\epsilon) = \frac{1}{2}(1 + 3\epsilon^2 + \epsilon^4)$, indicating $\lim_{\epsilon \rightarrow 0} J(K_\epsilon) = 1/2 < +\infty$. \square

4.2 The gradient and the Hessian of the LQG cost

The following lemma gives a closed-form expression for the gradient of the LQG cost function J_q , and its proof is given in Appendix A.7.

Lemma 8 (Gradient of LQG cost J_q) Fix $q \geq 1$ such that $\mathcal{C}_q \neq \emptyset$. For every $K = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathcal{C}_q$, the gradient of $J_q(K)$ is given by

$$\nabla J_q(K) = \begin{bmatrix} 0 & \frac{\partial J_q(K)}{\partial C_K} \\ \frac{\partial J_q(K)}{\partial B_K} & \frac{\partial J_q(K)}{\partial A_K} \end{bmatrix},$$

with

$$\frac{\partial J_q(\mathbf{K})}{\partial A_{\mathbf{K}}} = 2 \left(Y_{12}^T X_{12} + Y_{22} X_{22} \right), \quad (21a)$$

$$\frac{\partial J_q(\mathbf{K})}{\partial B_{\mathbf{K}}} = 2 \left(Y_{22} B_{\mathbf{K}} V + Y_{22} X_{12}^T C^T + Y_{12}^T X_{11} C^T \right), \quad (21b)$$

$$\frac{\partial J_q(\mathbf{K})}{\partial C_{\mathbf{K}}} = 2 \left(R C_{\mathbf{K}} X_{22} + B^T Y_{11} X_{12} + B^T Y_{12} X_{22} \right), \quad (21c)$$

where $X_{\mathbf{K}}$ and $Y_{\mathbf{K}}$, partitioned as

$$X_{\mathbf{K}} = \begin{bmatrix} X_{11} & X_{12} \\ X_{12}^T & X_{22} \end{bmatrix}, \quad Y_{\mathbf{K}} = \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{12}^T & Y_{22} \end{bmatrix} \quad (22)$$

are the unique positive semidefinite solutions to (12a) and (12b), respectively.

We next consider the Hessian of $J_q(\mathbf{K})$. Let \mathbf{K} be any controller in \mathcal{C}_q , and we use $\text{Hess}_{\mathbf{K}} : \mathcal{V}_q \times \mathcal{V}_q \rightarrow \mathbb{R}$ to denote the bilinear form of the Hessian of J_q at \mathbf{K} , so that for any $\Delta \in \mathcal{V}_q$, we have

$$J_n(\mathbf{K} + \Delta) = J_n(\mathbf{K}) + \text{tr} \left(\nabla J_q(\mathbf{K})^T \Delta \right) + \frac{1}{2} \text{Hess}_{\mathbf{K}}(\Delta, \Delta) + o(\|\Delta\|_F^2)$$

as $\|\Delta\|_F \rightarrow 0$. Obviously, $\text{Hess}_{\mathbf{K}}$ is symmetric in the sense that for all $x, y \in \mathcal{V}_n$, $\text{Hess}_{\mathbf{K}}(x, y) = \text{Hess}_{\mathbf{K}}(y, x)$. The following lemma shows how to compute the quantity $\text{Hess}_{\mathbf{K}}(\Delta, \Delta)$ for any $\Delta \in \mathcal{V}_q$ by solving three Lyapunov equations, whose proof is given in Appendix A.7.

Lemma 9 Fix $q \geq 1$ such that $\mathcal{C}_q \neq \emptyset$. Let $\mathbf{K} = \begin{bmatrix} 0 & C_{\mathbf{K}} \\ B_{\mathbf{K}} & A_{\mathbf{K}} \end{bmatrix} \in \mathcal{C}_q$. Then for any $\Delta = \begin{bmatrix} 0 & \Delta_{C_{\mathbf{K}}} \\ \Delta_{B_{\mathbf{K}}} & \Delta_{A_{\mathbf{K}}} \end{bmatrix} \in \mathcal{V}_q$, we have

$$\begin{aligned} \text{Hess}_{\mathbf{K}}(\Delta, \Delta) = 2 \text{tr} & \left(2 \begin{bmatrix} 0 & B \Delta_{C_{\mathbf{K}}} \\ \Delta_{B_{\mathbf{K}}} C & \Delta_{A_{\mathbf{K}}} \end{bmatrix} X'_{\mathbf{K}, \Delta} \cdot Y_{\mathbf{K}} + 2 \begin{bmatrix} 0 & 0 \\ 0 & C_{\mathbf{K}}^T R \Delta_{C_{\mathbf{K}}} \end{bmatrix} \cdot X'_{\mathbf{K}, \Delta} \right. \\ & \left. + \begin{bmatrix} 0 & 0 \\ 0 & \Delta_{B_{\mathbf{K}}} V \Delta_{B_{\mathbf{K}}}^T \end{bmatrix} Y_{\mathbf{K}} + \begin{bmatrix} 0 & 0 \\ 0 & \Delta_{C_{\mathbf{K}}}^T R \Delta_{C_{\mathbf{K}}} \end{bmatrix} X_{\mathbf{K}} \right), \end{aligned}$$

where $X_{\mathbf{K}}$ and $Y_{\mathbf{K}}$ are the solutions to the Lyapunov Eqs (12a) and (12b), and $X'_{\mathbf{K}, \Delta} \in \mathbb{R}^{(n+q) \times (n+q)}$ is the solution to the following Lyapunov equation

$$\begin{bmatrix} A & B C_{\mathbf{K}} \\ B_{\mathbf{K}} C & A_{\mathbf{K}} \end{bmatrix} X'_{\mathbf{K}, \Delta} + X'_{\mathbf{K}, \Delta} \begin{bmatrix} A & B C_{\mathbf{K}} \\ B_{\mathbf{K}} C & A_{\mathbf{K}} \end{bmatrix}^T + M_1(X_{\mathbf{K}}, \Delta) = 0, \quad (23)$$

with

$$M_1(X_K, \Delta) := \begin{bmatrix} 0 & B\Delta_{C_K} \\ \Delta_{B_K}C & \Delta_{A_K} \end{bmatrix} X_K + X_K \begin{bmatrix} 0 & B\Delta_{C_K} \\ \Delta_{B_K}C & \Delta_{A_K} \end{bmatrix}^T + \begin{bmatrix} 0 & 0 \\ 0 & B_K V \Delta_{B_K}^T + \Delta_{B_K} V B_K^T \end{bmatrix}.$$

From Lemma 9, one can further get $\text{Hess}_K(\Delta_1, \Delta_2)$ for any $\Delta_1, \Delta_2 \in \mathcal{V}_n$ by

$$\text{Hess}_K(\Delta_1, \Delta_2) = \frac{1}{2} (\text{Hess}_K(\Delta_1 + \Delta_2, \Delta_1 + \Delta_2) - \text{Hess}_K(\Delta_1, \Delta_1) - \text{Hess}_K(\Delta_2, \Delta_2)).$$

4.3 Spurious stationary points

In this part, we show that the LQG cost $J_n(K)$ over the full-order stabilizing controller \mathcal{C}_n may have many spurious stationary points that are not controllable and not observable.

We first investigate the gradient of $J_q(K)$ under similarity transformation. Given any $T \in \text{GL}_q$, recall the definition of the linear mapping of similarity transformation $\mathcal{T}_q(T, K)$ in (14). The following lemma gives an explicit relationship among the gradients of $J_q(\cdot)$ at K and $\mathcal{T}_q(T, K)$, whose proof is given in Appendix A.6.

Lemma 10 *Let $K = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathcal{C}_q$ be arbitrary. For any $T \in \text{GL}_q$, we have*

$$\nabla J_q \Big|_{\mathcal{T}_q(T, K)} = \begin{bmatrix} I_m & 0 \\ 0 & T^{-T} \end{bmatrix} \cdot \nabla J_q \Big|_K \cdot \begin{bmatrix} I_p & 0 \\ 0 & T^T \end{bmatrix}. \quad (24)$$

As expected, a direct consequence of Lemma 10 is that, if $K \in \mathcal{C}_q$ is a stationary point of J_q , then any controller in the orbit \mathcal{O}_K is also a stationary point of J_q . In addition, Lemma 10 allows us to establish an interesting result that any stationary point of J_q can be augmented to stationary points of $J_{q+q'}$ for any $q' > 0$ with the same objective value.

Theorem 4 *Let $q \geq 1$ be arbitrary. Suppose there exists $K^* = \begin{bmatrix} 0 & C_K^* \\ B_K^* & A_K^* \end{bmatrix} \in \mathcal{C}_q$ such that $\nabla J_q(K^*) = 0$. Then for any $q' \geq 1$ and any stable $\Lambda \in \mathbb{R}^{q' \times q'}$, the following controller*

$$\tilde{K}^* = \begin{bmatrix} 0 & C_K^* & 0 \\ B_K^* & A_K^* & 0 \\ 0 & 0 & \Lambda \end{bmatrix} \in \mathcal{C}_{q+q'} \quad (25)$$

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

Proof Since $K^* \in \mathcal{C}_q$, we have $\tilde{K}^* \in \mathcal{C}_{q+q'}$ by construction. It is straightforward to verify that

$$\mathcal{T}_{q+q'}(T, \tilde{K}^*) = \tilde{K}^* \quad \text{with} \quad T = \begin{bmatrix} I_q & 0 \\ 0 & -I_{q'} \end{bmatrix}.$$

Therefore, by Lemma 10, we have

$$\nabla J_{q+q'} \Big|_{\tilde{K}^*} = \nabla J_{q+q'} \Big|_{\mathcal{J}_{q+q'}(T, \tilde{K}^*)} = \begin{bmatrix} I_{m+q} & 0 \\ 0 & -I_{q'} \end{bmatrix} \cdot \nabla J_{q+q'} \Big|_{\tilde{K}^*} \cdot \begin{bmatrix} I_{p+q} & 0 \\ 0 & -I_{q'} \end{bmatrix},$$

which implies that, excluding the the bottom right $q' \times q'$ block, the last q' rows and the last q' columns of $\nabla J_{q+q'} \Big|_{\tilde{K}^*}$ are zero. On the other hand, it can be checked that

$$J_{q+q'} \left(\begin{bmatrix} K & 0 \\ 0 & \Lambda \end{bmatrix} \right) = J_q(K), \quad \forall K \in \mathcal{C}_q,$$

and since $\nabla J_q(K^*) = 0$, we can see that the upper left $(m+q) \times (p+q)$ block of $\nabla J_{q+q'} \Big|_{\tilde{K}^*}$ is equal to zero. Then, from Lemma 3, it is not difficult to verify that the value $J_q(\tilde{K}^*)$ is independent of the $q' \times q'$ stable matrix Λ , and thus the bottom right $q' \times q'$ block of $\nabla J_{q+q'} \Big|_{\tilde{K}^*}$ is zero.

We can now see that $\nabla J_{q+q'} \Big|_{\tilde{K}^*} = 0$. This completes the proof. \square

Theorem 4 indicates that from any stationary point of J_q over reduced-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 (25) are not controllable and not observable. Therefore, assuming that the optimal LQG controller is controllable and observable, as long as there exists some $q < n$ such that the problem of finding an optimal reduced-order controller $\min_{K \in \mathcal{C}_q} J_q(K)$ has a solution, we can then augment this solution to obtain a family of stationary points in \mathcal{C}_n that are not controllable and not observable, and consequently are spurious stationary points.

The following theorem explicitly constructs a family of stationary points for J_n with an open-loop stable plant, and also provides a criterion for checking whether the corresponding Hessian is indefinite or vanishing.

Theorem 5 Suppose the plant (1) is open-loop stable. Let $\Lambda \in \mathbb{R}^{n \times n}$ be stable, and let

$$K^* = \begin{bmatrix} 0 & 0 \\ 0 & \Lambda \end{bmatrix}.$$

Then K^* is a stationary point of $J_n(K)$ over $K \in \mathcal{C}_n$, and the corresponding Hessian Hess_{K^*} is either indefinite or zero.

Furthermore, suppose Λ is diagonalizable, and let $\text{eig}(-\Lambda)$ denote the set of (distinct) eigenvalues of $-\Lambda$. Let X_{op} and Y_{op} be the solutions to the following Lyapunov equations

$$AX_{\text{op}} + X_{\text{op}}A^T + W = 0, \quad A^T Y_{\text{op}} + Y_{\text{op}}A + Q = 0, \quad (26)$$

and let

$$\mathcal{Z} = \left\{ s \in \mathbb{C} \mid CX_{\text{op}}(sI - A^T)^{-1}Y_{\text{op}}B = 0 \right\}. \quad (27)$$

Then, the Hessian of J_n at K^* is indefinite if and only if $\text{eig}(-\Lambda) \not\subseteq \mathcal{Z}$, and is zero if and only if $\text{eig}(-\Lambda) \subseteq \mathcal{Z}$.

The fact that $K^* = \begin{bmatrix} 0 & 0 \\ 0 & \Lambda \end{bmatrix}$ is a stationary point can be proved similarly as Theorem 4. Appendix A.9 gives a detailed proof regarding the properties of the Hessian. A recent study [26] also shows that the LQG cost in terms of observer-based controllers has a zero gradient when $K = 0, L = 0$ for open-loop stable systems.

Theorem 5 constructs a family of strict saddle points or stationary points with vanishing Hessians for LQG with open-loop stable plants. We present two examples illustrating the Hessians at stationary points that are not controllable and observable.

Example 4 (Strict saddle point) Consider the plant in Example 2 with $A = -1, B = 1, C = 1$, and we choose $Q = R = 1, W = V = 1$. By Theorem 5, given any $a < 0$, the controller $K^* = \begin{bmatrix} 0 & 0 \\ 0 & a \end{bmatrix} \in \mathbb{R}^{2 \times 2}$ is a stationary point of $J_1(K)$ over the set of full-order stabilizing controller \mathcal{C}_1 . Furthermore, we can be check that

$$CX_{\text{op}}(sI - A^T)^{-1}Y_{\text{op}}B = \frac{1}{4(s+1)}.$$

Therefore the Hessian of J_1 at K^* is indefinite by Theorem 5, indicating that K^* is a strict saddle point. Indeed, by using (11), we can directly compute the LQG cost and obtain

$$J_1\left(\begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix}\right) = \frac{A_K^2 - A_K(1 + B_K^2 C_K^2) - B_K C_K(1 - 3B_K C_K + B_K^2 C_K^2)}{2(-1 + A_K)(A_K + B_K C_K)}.$$

The Hessian at K^* can then be represented as

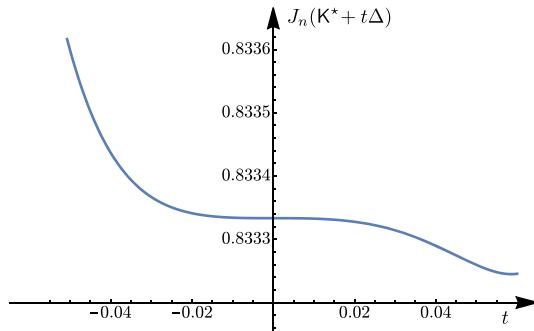
$$\begin{bmatrix} \frac{\partial^2 J(K)}{\partial A_K^2} & \frac{\partial^2 J(K)}{\partial A_K \partial B_K} & \frac{\partial^2 J(K)}{\partial A_K \partial C_K} \\ \frac{\partial^2 J(K)}{\partial B_K \partial A_K} & \frac{\partial^2 J(K)}{\partial B_K^2} & \frac{\partial^2 J(K)}{\partial B_K \partial C_K} \\ \frac{\partial^2 J(K)}{\partial C_K \partial A_K} & \frac{\partial^2 J(K)}{\partial C_K \partial B_K} & \frac{\partial^2 J(K)}{\partial C_K^2} \end{bmatrix} \Big|_{K=K^*} = \frac{1}{2(1-a)} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix},$$

which has eigenvalues 0 and $\pm \frac{1}{2(1-a)}$. □

Example 5 (Stationary point with a vanishing Hessian) Consider the following SISO system:

$$A = \begin{bmatrix} -1 & 0 \\ 1 & -2 \end{bmatrix}, \quad B = \begin{bmatrix} -1 \\ 1 \end{bmatrix}, \quad C = \begin{bmatrix} -2 & 11 \end{bmatrix}, \quad W = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad V = 1,$$

Fig. 3 The function $t \mapsto J_n(K^* + t\Delta)$ for Example 5



and let $Q = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, $R = 1$. It can be checked that

$$CX_{\text{op}}(sI - A^T)^{-1}Y_{\text{op}}B = \frac{5(s-1)}{36(s+1)(s+2)}.$$

By Theorem 5, the point $K^* = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}$ is a stationary point of J_n with a vanishing Hessian. In Fig. 3, we plot the graph of the function $t \mapsto J_n(K^* + t\Delta)$ for $\Delta = \begin{bmatrix} 0 & 2 & 1/2 \\ -1 & 1 & 3 \\ 3 & 0 & 0 \end{bmatrix}$. Figure 3 suggests that K^* is a saddle point of J_n with a vanishing Hessian but non-vanishing third-order partial derivatives. \square

Remark 2 Some recent studies have shown that many gradient-based algorithms can automatically escape strict saddle points under mild conditions [18]. However, Example 5 shows that the LQG cost function may have non-strict saddle points, and further analysis is required to examine whether gradient-based methods can also escape such stationary points. In addition, the existence of local minima is also important and relevant for the convergence of first-order algorithms, which we leave as future work.

4.4 Controllable and observable stationary points are globally optimal

In this section, we will show that all *controllable and observable* stationary points are globally optimal to the LQG problem (2).

We first give a useful lemma for controllable and observable stabilizing controllers (see Appendix A.8 for a proof).

Lemma 11 Fix $q \in \mathbb{N}$ such that $\mathcal{C}_q \neq \emptyset$, and let $K \in \mathcal{C}_q$ be controllable and observable. Under Assumption 1, the solutions X_K and Y_K to (12) are positive definite.

By letting the gradient (21) be equal to zero, we can derive closed-form expressions for full-order controllable and observable stationary points $K \in \mathcal{C}_n$ and show that they are globally optimal. This result is formally summarized below.

Theorem 6 Under Assumption 1, all controllable and observable stationary points of $J_n(K)$ over $K \in \mathcal{C}_n$ are globally optimal, and they are in the form of

$$A_K = T(A - BK - LC)T^{-1}, \quad B_K = -TL, \quad C_K = KT^{-1}, \quad (28)$$

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

$$K = R^{-1}B^T S, \quad L = PC^T V^{-1}, \quad (29)$$

with P and S being the unique positive definite solutions to the Riccati equations (5).

Proof The theorem can be viewed as a special case of [48, Theorem 20.6], [17, Section II] that analyze first-order necessary conditions for optimal reduced-order controllers. Following the analysis in [48, Chapter 20], we give an adapted proof.

Consider a stationary point $K = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathcal{C}_n$ such that the gradient (21) vanishes. If the controller K is controllable and observable, we know by Lemma 11 that the solutions X_K and Y_K to (12a) and (12b) are unique and positive definite. Upon partitioning X_K and Y_K in (22), by the Schur complement, the following matrices are well-defined and positive definite

$$P := X_{11} - X_{12}X_{22}^{-1}X_{12}^T \succ 0, \quad S := Y_{11} - Y_{12}Y_{22}^{-1}Y_{12}^T \succ 0. \quad (30)$$

We further define $T := Y_{22}^{-1}Y_{12}^T$. By (21a), we know that matrix T is invertible, and $T^{-1} = -X_{12}X_{22}^{-1}$. Now, by letting $\frac{\partial J_n(K)}{\partial B_K} = 0$ and noting (21b), we have

$$\begin{aligned} B_K &= -(X_{12}^T + Y_{22}^{-1}Y_{12}^T X_{11})C^T V^{-1} = -(X_{12}^T + TX_{11})C^T V^{-1} \\ &= -T(X_{11} - X_{12}X_{22}^{-1}X_{12}^T)C^T V^{-1} = -TPC^T V^{-1}. \end{aligned} \quad (31)$$

Similarly, from (21c), we have

$$C_K = -R^{-1}B^T(Y_{11}X_{12}X_{22}^{-1} + Y_{12}) = R^{-1}B^T S T^{-1}. \quad (32)$$

Furthermore, since X_K is the solution to the Lyapunov equation (12a), by plugging in the blocks of X_K we get

$$0 = AX_{11} + X_{11}A + BC_K X_{12}^T + X_{12}C_K^T B_K^T + W, \quad (33a)$$

$$0 = AX_{12} + BC_K X_{22} + X_{11}C^T B_K^T + X_{12}A_K^T, \quad (33b)$$

$$0 = A_K X_{22} + X_{22}A_K^T + B_K C X_{12} + X_{12}^T C^T B_K^T + B_K V B_K^T. \quad (33c)$$

Now, (33c) + $T \times$ (33b) leads to

$$\begin{aligned} &A_K X_{22} + X_{22}A_K^T + B_K C X_{12} + X_{12}^T C^T B_K^T + B_K V B_K^T \\ &+ T(AX_{12} + BC_K X_{22} + X_{11}C^T B_K^T + X_{12}A_K^T) = 0, \end{aligned}$$

which is the same as

$$A_K X_{22} + X_{22} A_K^\top - T P C^\top V^{-1} C X_{12} - X_{12}^\top C^\top V^{-1} C P T + T P C^\top V^{-1} C P T \\ + T(AX_{12} + BR^{-1}B^\top ST^{-1}X_{22} - X_{11}C^\top V^{-1}C P T + X_{12}A_K^\top) = 0.$$

By the definition of T , we have $TX_{12} = -X_{22}$. Then, the equation above becomes

$$A_K X_{22} - T P C^\top V^{-1} C X_{12} - X_{12}^\top C^\top V^{-1} C P T + T P C^\top V^{-1} C P T \\ + T(AX_{12} + BR^{-1}B^\top ST^{-1}X_{22} - X_{11}C^\top V^{-1}C P T) = 0,$$

leading to

$$A_K = T P C^\top V^{-1} C X_{12} X_{22}^{-1} + X_{12}^\top C^\top V^{-1} C P T X_{22}^{-1} - T P C^\top V^{-1} C P T X_{22}^{-1} \\ - T(AX_{12} + BR^{-1}B^\top ST^{-1}X_{22} - X_{11}C^\top V^{-1}C P T) X_{22}^{-1} \\ = T(A - P C^\top V^{-1} C - B R^{-1} B^\top S) T^{-1}. \quad (34)$$

From (31), (32) and (34), upon defining K and L in (29), it is easy to see that the stationary points are in the form of (28). It remains to prove that P and S defined in (30) are the unique positive definite solutions to the Riccati Eqs (5a) and (5b).

We multiply (33c) by T^{-1} on the left and by $T^{-\top}$ on the right, and by noting that $B_K = -T P C^\top V^{-1}$ and $T^{-1} = -X_{12} X_{22}^{-1}$, we get

$$0 = X_{12} X_{22}^{-1} A_K X_{12}^\top + X_{12} A_K^\top X_{22}^{-1} X_{12}^\top \\ + P C^\top V^{-1} C X_{12} X_{22}^{-1} X_{12}^\top + X_{12} X_{22}^{-1} X_{12}^\top C^\top V^{-1} C P + P C^\top V^{-1} C P.$$

Since $P = X_{11} - X_{12} X_{22}^{-1} X_{12}^\top$, we further get

$$0 = X_{12} X_{22}^{-1} A_K X_{12}^\top + X_{12} A_K^\top X_{22}^{-1} X_{12}^\top \\ + P C^\top V^{-1} C X_{11} + X_{11} C^\top V^{-1} C P - P C^\top V^{-1} C P. \quad (35)$$

Next, we multiply (33b) by $-T^{-\top} = X_{22}^{-1} X_{12}^\top$ on the right and get

$$0 = A X_{12} X_{22}^{-1} X_{12}^\top + B C_K X_{12}^\top + X_{11} C^\top V^{-1} C^\top P + X_{12} A_K^\top X_{22}^{-1} X_{12}^\top.$$

By plugging this equality into (35), we get

$$0 = -A X_{12} X_{22}^{-1} X_{12}^\top - B C_K X_{12}^\top - X_{12} X_{22}^{-1} X_{12}^\top A - X_{12} C_K^\top B^\top - P C^\top V^{-1} C P.$$

Then, we plug the above equality into (33a) and get

$$0 = A(X_{11} - X_{12} X_{22}^{-1} X_{12}^\top) + (X_{11} - X_{12} X_{22}^{-1} X_{12}^\top) A - P C^\top V^{-1} C P + W,$$

and since $P = X_{11} - X_{12}X_{22}^{-1}X_{12}^T$, we see that P satisfies the Riccati Eq (5a). By similar steps, we can derive from (12b) that S satisfies the Riccati equation (5b).

Finally, from classical control theory [48, Theorem 14.7], a globally optimal controller to the LQG problem (13) is given by (6), and any similarity transformation leads to another equivalent controller with the same LQG cost. Therefore, any controllable and observable stationary point, given by (28), is globally optimal. \square

We note that *controllability and observability* are required in the proof of Theorem 6, as they guarantee that the matrices (30) are well-defined and the solutions (31) and (32) are unique.

Theorem 6 implies that, if the LQG problem (13) has a globally optimal solution in \mathcal{C}_n that is also controllable and observable, then the globally optimal controller is unique modulo similarity transformations. This is expected from the classical result that the globally optimal LQG controller is unique in the frequency domain [48, Theorem 14.7]. Theorem 6 also allows us to establish the following corollaries.

Corollary 2 *The following statements hold:*

1. *If $J_n(\mathbf{K})$ has a controllable and observable stationary point in \mathcal{C}_n , then any stationary point that is not controllable or observable is strictly suboptimal.*
2. *If $J_n(\mathbf{K})$ has a globally optimal point in \mathcal{C}_n that is not controllable or observable, then all stationary points of $J_n(\mathbf{K})$ are not controllable or observable.*

We have already seen LQG cases with strictly suboptimal stationary points that are not controllable and not observable in Example 4 and Example 5. It should be noted that, even with Assumption 1, the LQG problem (13) might have no controllable and observable stationary points; this happens if the controller from the Riccati equations (5) is not controllable or observable.

Example 6 (Globally optimal controllers that are not controllable or observable) Here we give an example from [41], whose optimal LQG controller does not have a full-order realization in \mathcal{C}_n that is controllable and observable. Consider the linear system (1) with

$$A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & -1 \end{bmatrix}, \quad W = \begin{bmatrix} 1 & -1 \\ -1 & 16 \end{bmatrix}, \quad V = 1,$$

and let the LQG cost be defined by $Q = \begin{bmatrix} 4 & 0 \\ 0 & 0 \end{bmatrix}$, $R = 1$. This LQG problem satisfies Assumption 1, and the globally optimal controller is given by

$$A_K = \begin{bmatrix} -3 & 0 \\ 5 & -4 \end{bmatrix}, \quad B_K = L = \begin{bmatrix} 1 \\ -4 \end{bmatrix}, \quad C_K = -K = \begin{bmatrix} -2 & 0 \end{bmatrix}. \quad (36)$$

It is not hard to see that (C_K, A_K) is not observable. Consequently, by Corollary 2, all stationary points of J_n are not controllable or observable for this example.

In this case, the globally optimal controllers in \mathcal{C}_n are not all connected by similarity transformations. For example, it can be verified that the following two controllers are both globally optimal:

$$K_1 = \begin{bmatrix} 0 & -2 & 0 \\ 1 & -3 & 0 \\ -4 & 5 & -4 \end{bmatrix}, \quad K_2 = \begin{bmatrix} 0 & -2 & 0 \\ 1 & -3 & 0 \\ 0 & 0 & -1 \end{bmatrix},$$

but there exists no similarity transformation between K_1 and K_2 since $\begin{bmatrix} -3 & 0 \\ 5 & -4 \end{bmatrix}$ and $\begin{bmatrix} -3 & 0 \\ 0 & -1 \end{bmatrix}$ have different sets of eigenvalues. \square

Now let us consider a sequence of gradient descent iterates that converges to a point. Then Theorem 6 also allows us to check whether the limit point is a globally optimal solution to the LQG problem.

Corollary 3 *Consider the gradient descent iteration $K_{t+1} = K_t - \alpha_t \nabla J_n(K_t)$ for the LQG problem (13), where $\alpha_t > 0$ is the step size. Suppose $\inf_t \alpha_t > 0$ and the iterates K_t converge to a point K^* . Then K^* is globally optimal if it is a controllable and observable controller.*

Remark 3 Corollary 3 proposes checking the controllability and observability of K^* for verifying global optimality when the gradient descent iterates converge to K^* . In practice, the limit K^* cannot be directly computed, and one tentative approach to check its controllability (observability) is to check whether the smallest singular value of the controllability (observability) matrix of the last iterate K_T is sufficiently bounded away from zero. A rigorous justification of this approach will be of interest for future work.

Remark 4 Note that Corollary 3 does not discuss under what conditions will the gradient descent iterates converge. The results in [1] guarantee that if the cost function is analytic over the whole Euclidean space, then the gradient descent with step sizes satisfying the Wolfe conditions will either converge to a stationary point or diverge to infinity. In our case, however, the cost function $J_n(K)$ is only analytic over a subset $\mathcal{C}_n \subset \mathcal{V}_n$. Furthermore, $J_n(K)$ is not coercive as shown in Example 3. Whether the gradient descent with properly chosen step sizes can converge to a stationary point of $J_n(K)$ requires further investigation.

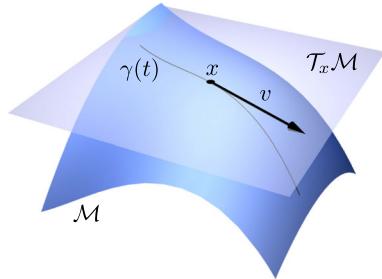
4.5 Hessian of $J_n(K)$ at controllable and observable stationary points

Finally, we turn to characterizing the second-order behavior of J_n around a globally optimal controller K^* by investigating the eigenvalues and eigenspaces of the Hessian Hess_{K^*} . We assume K^* is controllable and observable throughout this subsection.

Proposition 3 guarantees that for any controllable and observable $K \in \mathcal{C}_n$, the orbit \mathcal{O}_K is a submanifold of dimension n^2 in \mathcal{C}_n , which allows us to define the tangent space of \mathcal{O}_K .⁴ For each controllable and observable $K \in \mathcal{C}_n$, we use $\mathcal{T}\mathcal{O}_K$ to denote

⁴ A visualization of a manifold \mathcal{M} and its tangent space $\mathcal{T}_x\mathcal{M}$ at one point $x \in \mathcal{M}$ is provided in Fig. 4.

Fig. 4 Visualization of a manifold \mathcal{M} and its tangent space $T_x\mathcal{M}$ at some $x \in \mathcal{M}$. Here $\gamma(t)$ is an arbitrary C^∞ curve in \mathcal{M} that passes through x , and v is the tangent vector of $\gamma(t)$ at x . The tangent space $T_x\mathcal{M}$ consists of all such vectors v



the tangent space of \mathcal{O}_K at K , and treat it as a subspace of \mathcal{V}_n ; recall that \mathcal{V}_n is defined by (10). The dimension of $T\mathcal{O}_K$ is then $\dim T\mathcal{O}_K = \dim \mathcal{O}_K = n^2$. We denote the orthogonal complement of $T\mathcal{O}_K$ in \mathcal{V}_n by $T\mathcal{O}_K^\perp$. The following proposition characterizes the tangent space $T\mathcal{O}_K$ and its orthogonal complement $T\mathcal{O}_K^\perp$ at a controllable and observable controller $K \in \mathcal{C}_n$, whose proof is given in Appendix A.5.

Proposition 4 *Let $K = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathcal{C}_n$ represent a controllable and observable controller. Then*

$$T\mathcal{O}_K = \left\{ \begin{bmatrix} 0 & -C_K H \\ H B_K & H A_K - A_K H \end{bmatrix} \mid H \in \mathbb{R}^{n \times n} \right\},$$

$$T\mathcal{O}_K^\perp = \left\{ \Delta = \begin{bmatrix} 0 & \Delta_{B_K} \\ \Delta_{C_K} & \Delta_{A_K} \end{bmatrix} \in \mathcal{V}_n \mid \Delta_{A_K} A_K^\top - A_K^\top \Delta_{A_K} + \Delta_{B_K} B_K^\top - C_K^\top \Delta_{C_K} = 0 \right\}.$$

We now present the following lemma, which shows that the tangent space $T\mathcal{O}_{K^*}$ is a subspace of the null space of Hess_{K^*} , defined as

$$\text{null Hess}_{K^*} = \{x \in \mathcal{V}_n \mid \text{Hess}_{K^*}(x, y) = 0, \forall y \in \mathcal{V}_n\}.$$

Lemma 12 *Suppose K^* is controllable and observable. Then $T\mathcal{O}_{K^*} \subseteq \text{null Hess}_{K^*}$.*

This lemma is a direct corollary of [22, Theorem 2], and can be viewed as a local version of Lemma 7 indicating the invariance of J_n along the orbit \mathcal{O}_K . Consequently, the dimension of the null space of Hess_{K^*} is at least n^2 . On the other hand, we also have the following result.

Lemma 13 *Suppose K^* is controllable and observable, and let $\Delta \in T\mathcal{O}_{K^*}^\perp$. Then for all sufficiently small $t > 0$,*

$$J_n(K^* + t\Delta) - J_n(K^*) > 0.$$

Proof We prove by contradiction. Suppose for any $\delta > 0$, there always exists $t \in (0, \delta)$ such that $J_n(K^* + t\Delta) = J_n(K^*)$. Then we can find a positive sequence $(t_j)_{j \geq 1}$ such that $t_j \rightarrow 0$ and $J_n(K^* + t_j\Delta) = J_n(K^*)$. Denote $K_j = K^* + t_j\Delta$. Since Δ is orthogonal to $T\mathcal{O}_{K^*}$, there must exist some $j \geq 1$ such that $K_j \notin \mathcal{O}_{K^*}$. By [48, Theorem 3.17],

we can see that the transfer function of K_j will be different from the transfer function of K^* . Then by the uniqueness of the transfer function solution to the LQG problem, K_j cannot be a global minimum of J_n , contradicting $J_n(K_j) = J_n(K^*)$. \square

Combining the observations from Lemmas 12 and 13, we can see that, while the Hessian Hess_{K^*} is degenerate and its null space has a nontrivial subspace \mathcal{TO}_{K^*} , the degeneracy associated with \mathcal{TO}_{K^*} does not cause much trouble for optimizing J_n , as the directions in \mathcal{TO}_{K^*} correspond to similarity transformations that lead to other globally optimal controllers, while along the directions orthogonal to \mathcal{TO}_{K^*} , the optimal controller of J_n is locally unique.

We are therefore interested in the behavior of Hess_{K^*} restricted to the subspace $\mathcal{TO}_{K^*}^\perp$. Specifically, we let rcond_{K^*} denote the reciprocal condition number of Hess_{K^*} restricted to the subspace $\mathcal{TO}_{K^*}^\perp$, i.e.,

$$\text{rcond}_{K^*} := \frac{\min_{\Delta \perp \mathcal{TO}_{K^*}} \text{Hess}_{K^*}(\Delta, \Delta) / \|\Delta\|_F^2}{\max_{\Delta \perp \mathcal{TO}_{K^*}} \text{Hess}_{K^*}(\Delta, \Delta) / \|\Delta\|_F^2}. \quad (37)$$

Intuitively, if rcond_{K^*} is bounded away from zero, then we can expect gradient-based methods to achieve good local convergence behavior for optimizing J_n . However, we give an explicit example below showing that rcond_{K^*} can be very bad even if the original plant seems entirely normal.

Example 7 Let $\epsilon > 0$ be arbitrary, and let

$$A = \frac{3}{2} \begin{bmatrix} -1 & 0 \\ 0 & -1 - \epsilon \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 1 + \epsilon \end{bmatrix}, \quad C = [1 \ 1],$$

and

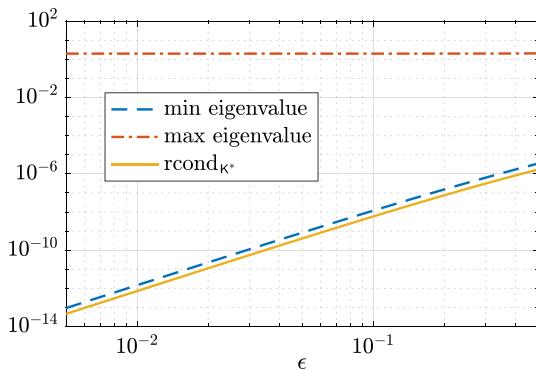
$$Q = \begin{bmatrix} 4 & 1 \\ 1 & 4 \end{bmatrix}, \quad W = \begin{bmatrix} 4 & 1 + \epsilon \\ 1 + \epsilon & 4(1 + \epsilon)^2 \end{bmatrix}, \quad V = R = 1.$$

For this plant, the positive definite solutions to the Riccati equations (5) are given by $P = \begin{bmatrix} 1 & 0 \\ 0 & 1 + \epsilon \end{bmatrix}$, $S = \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{1+\epsilon} \end{bmatrix}$, and the optimal controller K^* is then given by

$$K^* = \begin{bmatrix} 0 & -1 & -1 \\ 1 & -\frac{7}{2} & -2 \\ 1 + \epsilon & -2(1 + \epsilon) & -\frac{7}{2}(1 + \epsilon) \end{bmatrix}.$$

It can be checked that the optimal controller provided by the Riccati equations is controllable and observable when $\epsilon \neq 0$. In Fig. 5, we plot the minimum and maximum eigenvalues of Hess_{K^*} restricted to $\mathcal{TO}_{K^*}^\perp$, as ϵ varies in $[0.005, 0.5]$. It can be seen that rcond_{K^*} degrades rapidly as ϵ approaches zero. Moreover, even if we set $\epsilon = 0.5$, the reciprocal condition number rcond_{K^*} is still below 1.7×10^{-6} . On the other hand,

Fig. 5 The minimum and maximum eigenvalues of Hess_{K^*} restricted on $\mathcal{TO}_{K^*}^\perp$, and the reciprocal condition number rcond_{K^*}



if we plug in $\epsilon = 0.5$, the resulting plant's parameters as well as the controllability and observability matrices

$$[B \ AB] = \begin{bmatrix} 1 & -1.5 \\ 1.5 & -3.375 \end{bmatrix}, \quad \begin{bmatrix} C \\ CA \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ -1.5 & -2.25 \end{bmatrix}$$

seem entirely normal. \square

The observations in Example 7 suggest that, if we apply the vanilla gradient descent algorithm to the optimization problem (13), it may take a large number of iterations for the iterate to converge to a globally optimal controller for certain LQG problems that appear entirely normal.

To conclude this section, we provide some final remarks on the symmetry structures in LQG compared to existing literature on the landscapes of some non-convex machine learning problems.

Remark 5 (Symmetry structures in LQG) Due to the symmetry induced by similarity transformations, the landscape of LQG shares some similarities with the landscapes of non-convex machine learning problems with rotational symmetries such as phase retrieval, matrix factorization [22, 34, 43]. For example, the stationary points of these non-convex problems are non-isolated, and the tangent space of the orbit associated with the symmetry group is a subspace of the null space of the Hessian (see Lemma 12). On the other hand, for phase retrieval [34] and matrix factorization [22], the classification of all stationary points as well as their local curvatures (Hessian) seem to be relatively well understood, while there remain many open questions regarding the stationary points of LQG, such as the existence of local optimizers that are not globally optimal, whether all non-globally-optimal stationary points have the form of (25) up to similarity transformations. Finally, in addition to the apparent algebraic intricacy of LQG and control-theoretic notions such as controllable and observable controllers, the non-compactness of the group of similarity transformations may also render the landscape of LQG distinct from the non-convex machine learning problems with rotational symmetries.

5 Conclusion

In this paper, we investigated the optimization landscape structures of the LQG problem, including the connectivity of the set of stabilizing dynamic controllers \mathcal{C}_n , and some structural properties of the stationary points of the LQG cost function. These results reveal rich yet complicated optimization landscape properties of LQG.

Ongoing work includes establishing convergence conditions for gradient descent algorithms and investigating whether local search algorithms can escape saddle points of the LQG problem. Also, we note that the optimization landscape of LQG depends on the parameterization of dynamic controllers, and it will be interesting to investigate different parameterizations (see our technical report [47] for relevant numerical results). It would also be interesting to investigate criteria for the existence of controllable and observable stationary points of J_n in \mathcal{C}_n , as well as how to certify the global optimality of a stationary point without knowing the system order n . Finally, we hope our results will facilitate future research on the design of a full model-free policy gradient algorithm for LQG with performance guarantees.

A Technical Proofs

A.1 Proof of Lemma 2

It is a well-known fact in control theory that $\mathcal{C}_n \neq \emptyset$ under Assumption 1. In particular, any pole assignment algorithm or solving the Riccati equations (5a) and (5b) can find a feasible point in \mathcal{C}_n . To show the unboundedness of \mathcal{C}_n , we introduce the following set

$$\mathcal{S}_n = \left\{ K = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathbb{R}^{(m+n) \times (p+n)} \middle| \begin{array}{l} A_K = A - BK - LC, B_K = L, C_K = -K, \\ A - BK \text{ and } A - LC \text{ are stable} \end{array} \right\}.$$

It has been established in classical control theory that $\mathcal{S}_n \subset \mathcal{C}_n$ [48, Chapter 3.5] and the set $\{K \mid A - BK \text{ is stable}\}$ is unbounded (see, e.g., [6, Observation 3.6]). Thus, the set \mathcal{S}_n is unbounded, and so is \mathcal{C}_n . Non-convexity of \mathcal{C}_n is also known and can be illustrated by the explicit counterexample in Example 8.

Example 8 (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

$$\mathcal{C}_n = \left\{ K = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathbb{R}^{2 \times 2} \middle| \begin{bmatrix} 1 & C_K \\ B_K & A_K \end{bmatrix} \text{ is stable} \right\}.$$

It is easy to verify that the two controllers $K^{(1)} = \begin{bmatrix} 0 & 2 \\ -2 & -2 \end{bmatrix}$ and $K^{(2)} = \begin{bmatrix} 0 & -2 \\ 2 & -2 \end{bmatrix}$ internally stabilize the plant and thus belong to \mathcal{C}_1 . However, $\hat{K} = \frac{1}{2}(K^{(1)} + K^{(2)}) = \begin{bmatrix} 0 & 0 \\ 0 & -2 \end{bmatrix}$ fails to stabilize the plant. \square

A.2 Proof of Lemma 4

Upon vectorizing the Lyapunov equation (12a), we have

$$(I_{n+q} \otimes A_{\text{cl},K} + A_{\text{cl},K} \otimes I_{n+q}) \text{vec}(X_K) = -\text{vec} \left(\begin{bmatrix} W & 0 \\ 0 & B_K V B_K^\top \end{bmatrix} \right).$$

Since $A_{\text{cl},K}$ is stable, we know that $I_{n+q} \otimes A_{\text{cl},K} + A_{\text{cl},K} \otimes I_{n+q}$ is invertible, and thus we have

$$\text{vec}(X_K) = - (I_{n+q} \otimes A_{\text{cl},K} + A_{\text{cl},K} \otimes I_{n+q})^{-1} \text{vec} \left(\begin{bmatrix} W & 0 \\ 0 & B_K V B_K^\top \end{bmatrix} \right).$$

It is not difficult to see that each element of $(I_{n+q} \otimes A_{\text{cl},K} + A_{\text{cl},K} \otimes I_{n+q})^{-1}$ is a rational function of the elements of K . Therefore, the LQG cost function

$$J_q(K) = \text{tr} \left(\begin{bmatrix} Q & 0 \\ 0 & C_K^\top R C_K \end{bmatrix} X_K \right)$$

is a rational function of the elements of K , which is real analytical.

A.3 Proof of Proposition 1

The following Lyapunov stability criterion [5] will be used in our proof: A square real matrix M is stable if and only if the Lyapunov inequality $MP + PM^\top \prec 0$ has a positive definite solution $P \succ 0$.

It is straightforward to see that $\Phi(\cdot)$ is continuous since each element of $\Phi(Z)$ is a rational function in terms of the elements of Z (a ratio of two polynomials). To show that Φ is a mapping onto \mathcal{C}_n , we need to prove the following statements:

1. For all $K \in \mathcal{C}_n$, there exists $Z = (X, Y, M, H, F, \Pi, \Xi) \in \mathcal{G}_n$ such that $\Phi(Z) = K$.
2. For all $Z = (X, Y, M, H, F, \Pi, \Xi) \in \mathcal{G}_n$, we have $\Phi(Z) \in \mathcal{C}_n$.

To show the first statement, let $K = \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \in \mathcal{C}_n$ be arbitrary. The stability of the matrix $\begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix}$ implies that the Lyapunov inequality

$$\begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix} \begin{bmatrix} X & \Pi^\top \\ \Pi & \hat{X} \end{bmatrix} + \begin{bmatrix} X & \Pi^\top \\ \Pi & \hat{X} \end{bmatrix} \begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix}^\top \prec 0 \quad (38)$$

has a solution $\begin{bmatrix} X & \Pi^\top \\ \Pi & \hat{X} \end{bmatrix} \succ 0$. Without loss of generality we may assume that $\det \Pi \neq 0$ (otherwise we can add a small perturbation on Π to make it invertible while still

preserving the inequality (38)). Upon defining

$$\begin{bmatrix} Y & \Xi \\ \Xi^\top & \hat{X} \end{bmatrix} := \begin{bmatrix} X & \Pi^\top \\ \Pi & \hat{X} \end{bmatrix}^{-1}, \quad T := \begin{bmatrix} X & \Pi^\top \\ \Pi & \hat{X} \end{bmatrix}^{-1} \begin{bmatrix} X & I \\ \Pi & 0 \end{bmatrix} = \begin{bmatrix} I & Y \\ 0 & \Xi^\top \end{bmatrix},$$

we can verify that

$$YX + \Xi\Pi = I, \quad T^\top \begin{bmatrix} X & \Pi^\top \\ \Pi & \hat{X} \end{bmatrix} T = \begin{bmatrix} X & I \\ I & Y \end{bmatrix} \succ 0. \quad (39)$$

Upon letting

$$\begin{aligned} M &= YAX + \Xi B_K CX + YBC_K \Pi + \Xi A_K \Pi, \\ H &= \Xi B_K, \quad F = C_K \Pi, \end{aligned} \quad (40)$$

we can also verify that

$$T^\top \begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix} \begin{bmatrix} X & \Pi^\top \\ \Pi & \hat{X} \end{bmatrix} T = \begin{bmatrix} AX + BF & A \\ M & YA + HC \end{bmatrix}. \quad (41)$$

Combining (41) with (38) and (39), we see that $Z = (X, Y, M, H, F, \Pi, \Xi) \in \mathcal{G}_n$ by the definition of \mathcal{G}_n . Note that the change of variables (40) can be compactly represented as

$$\begin{bmatrix} 0 & F \\ H & M \end{bmatrix} = \begin{bmatrix} I & 0 \\ YB & \Xi \end{bmatrix} \begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} \begin{bmatrix} I & CX \\ 0 & \Pi \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & YAX \end{bmatrix},$$

and with the guarantee in Lemma 6, we see that

$$\begin{bmatrix} 0 & C_K \\ B_K & A_K \end{bmatrix} = \begin{bmatrix} I & 0 \\ YB & \Xi \end{bmatrix}^{-1} \begin{bmatrix} 0 & F \\ H & M - YAX \end{bmatrix} \begin{bmatrix} I & CX \\ 0 & \Pi \end{bmatrix}^{-1} = \begin{bmatrix} 0 & \Phi_C(Z) \\ \Phi_B(Z) & \Phi_A(Z) \end{bmatrix} = \Phi(Z).$$

We then prove the second statement. Let $Z = (X, Y, M, H, F, \Pi, \Xi) \in \mathcal{G}_n$ be arbitrary. Let $\hat{X} = \Pi(X - Y^{-1})^{-1}\Pi^\top$, and it is straightforward to see that $\hat{X} \succ 0$ and

$$\begin{bmatrix} X & \Pi^\top \\ \Pi & \hat{X} \end{bmatrix} \begin{bmatrix} I & Y \\ 0 & \Xi^\top \end{bmatrix} = \begin{bmatrix} X & XY + \Pi^\top \Xi^\top \\ \Pi & \Pi Y + \hat{X} \Xi^\top \end{bmatrix} = \begin{bmatrix} X & I \\ \Pi & 0 \end{bmatrix},$$

where we used the fact that

$$\begin{aligned} \Pi Y + \hat{X} \Xi^\top &= \Pi Y + \Pi(X - Y^{-1})^{-1}\Pi^\top \Xi^\top = \Pi Y - \Pi(X - Y^{-1})^{-1}(XY - I) \\ &= \Pi Y - \Pi(X - Y^{-1})^{-1}(X - Y^{-1})Y = 0. \end{aligned}$$

We also have

$$\begin{bmatrix} 0 & F \\ H & M \end{bmatrix} = \begin{bmatrix} I & 0 \\ YB & \Xi \end{bmatrix} \begin{bmatrix} 0 & \Phi_C(Z) \\ \Phi_B(Z) & \Phi_A(Z) \end{bmatrix} \begin{bmatrix} I & CX \\ 0 & \Pi \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & YAX \end{bmatrix}$$

from the definition of Φ . Similarly as showing the equality (41), we can derive that

$$\begin{bmatrix} AX+BF & A \\ M & YA+HC \end{bmatrix} = \begin{bmatrix} I & Y \\ 0 & \Xi^T \end{bmatrix}^T \begin{bmatrix} A & B\Phi_C(Z) \\ \Phi_B(Z)C & \Phi_A(Z) \end{bmatrix} \begin{bmatrix} X & \Pi^T \\ \Pi & \hat{X} \end{bmatrix} \begin{bmatrix} I & Y \\ 0 & \Xi^T \end{bmatrix}.$$

Then from the definition of \mathcal{G}_n , we can further get

$$\begin{bmatrix} A & B\Phi_C(Z) \\ \Phi_B(Z)C & \Phi_A(Z) \end{bmatrix} \begin{bmatrix} X & \Pi^T \\ \Pi & \hat{X} \end{bmatrix} + \begin{bmatrix} X & \Pi^T \\ \Pi & \hat{X} \end{bmatrix} \begin{bmatrix} A & B\Phi_C(Z) \\ \Phi_B(Z)C & \Phi_A(Z) \end{bmatrix}^T \prec 0,$$

and since $X - \Pi^T \hat{X}^{-1} \Pi = Y^{-1} \succ 0$, the matrix $\begin{bmatrix} X & \Pi^T \\ \Pi & \hat{X} \end{bmatrix}$ is positive definite. We can now see that $\begin{bmatrix} A & B\Phi_C(Z) \\ \Phi_B(Z)C & \Phi_A(Z) \end{bmatrix}$ satisfies the Lyapunov inequality and thus is stable, meaning that $\Phi(Z) \in \mathcal{C}_n$.

A.4 Proof of Proposition 3

We have already seen that \mathcal{T}_q gives a smooth Lie group action of GL_q on \mathcal{C}_q . We first show that the isotropy group of K under the group actions in GL_q , defined by $\{T \in \mathrm{GL}_q \mid \mathcal{T}_q(T, \mathsf{K}) = \mathsf{K}\}$, is a trivial group containing only the identity matrix.

Let $T \in \mathrm{GL}_q$ satisfy $\mathcal{T}_q(T, \mathsf{K}) = \mathsf{K}$, i.e.,

$$\begin{bmatrix} 0 & C_{\mathsf{K}} T^{-1} \\ TB_{\mathsf{K}} & TA_{\mathsf{K}} T^{-1} \end{bmatrix} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix}.$$

Then we have $TA_{\mathsf{K}} = A_{\mathsf{K}}T$, and consequently $TA_{\mathsf{K}}^{j+1}B_{\mathsf{K}} = A_{\mathsf{K}}TA_{\mathsf{K}}^jB_{\mathsf{K}}$. By mathematical induction, we can see that $TA_{\mathsf{K}}^jB_{\mathsf{K}} = A_{\mathsf{K}}^jB_{\mathsf{K}}$ for all $j = 0, \dots, q-1$, indicating that any column vector of $A_{\mathsf{K}}^jB_{\mathsf{K}}$ is an eigenvector of T with eigenvalue 1. On the other hand, the controllability of K implies the column vectors of the matrix $\begin{bmatrix} B_{\mathsf{K}} & A_{\mathsf{K}}B_{\mathsf{K}} & \dots & A_{\mathsf{K}}^{q-1}B_{\mathsf{K}} \end{bmatrix}$ span the whole space \mathbb{R}^q . Therefore \mathbb{R}^q is a subspace of the eigenspace of T with eigenvalue 1, meaning that T is just the identity matrix.

Since the isotropy group $\{T \in \mathrm{GL}_q \mid \mathcal{T}_q(T, \mathsf{K}) = \mathsf{K}\}$ only contains the identity, by [21, Proposition 7.26], the mapping $T \mapsto \mathcal{T}_q(T, \mathsf{K})$ is an immersion and the orbit \mathcal{O}_{K} is an immersed submanifold.

We then prove that \mathcal{O}_{K} is closed under the original topology of \mathcal{C}_q . Suppose $(T_j)_{j=1}^{\infty}$ is a sequence in GL_q such that

$$\mathcal{T}_q(T_j, \mathsf{K}) = \begin{bmatrix} 0 & C_{\mathsf{K}} T_j^{-1} \\ T_j B_{\mathsf{K}} & T_j A_{\mathsf{K}} T_j^{-1} \end{bmatrix} \rightarrow \begin{bmatrix} 0 & \tilde{C}_{\mathsf{K}} \\ \tilde{B}_{\mathsf{K}} & \tilde{A}_{\mathsf{K}} \end{bmatrix} = \tilde{\mathsf{K}}, \quad j \rightarrow \infty.$$

Let $\mathbf{G}(s)$ be the transfer function of K , i.e., $\mathbf{G}(s) = C_{\mathsf{K}}(sI - A_{\mathsf{K}})^{-1}B_{\mathsf{K}}$. We notice that for any $j \geq 1$, the matrix $sI - T_j A_{\mathsf{K}} T_j^{-1}$ is invertible if and only if $sI - A_{\mathsf{K}}$ is

invertible. Thus for any fixed $s \in \mathbb{C}$ such that $sI - A_K$ is invertible, we have

$$\lim_{j \rightarrow \infty} C_K T_j^{-1} (sI - T_j A_K T_j^{-1})^{-1} T_j B_K = \tilde{C}_K (sI - \tilde{A}_K)^{-1} \tilde{B}_K.$$

On the other hand, we simply have

$$C_K T_j^{-1} (sI - T_j A_K T_j^{-1})^{-1} T_j B_K = C_K (sI - A_K)^{-1} B_K = \mathbf{G}(s).$$

This shows that the transfer function of \tilde{K} agrees with $\mathbf{G}(s)$ for any $s \in \mathbb{C}$ such that $sI - A_K$ is invertible, and thus is just equal to $\mathbf{G}(s)$. On the other hand, the controllability and observability of $K \in \mathcal{C}_q$ indicates that the transfer function $\mathbf{G}(s)$ has order q , and so any two state-space representations of $\mathbf{G}(s)$ with order q will always be similarity transformations of each other [48, Theorem 3.17]. In other words, there exists $\tilde{T} \in \mathrm{GL}_q$ such that

$$\tilde{K} = \begin{bmatrix} 0 & \tilde{C}_K \\ \tilde{B}_K & \tilde{A}_K \end{bmatrix} = \begin{bmatrix} 0 & C_K \tilde{T}^{-1} \\ \tilde{T} B_K & \tilde{T} A_K \tilde{T}^{-1} \end{bmatrix} = \mathcal{T}_q(\tilde{T}, K),$$

which implies that $\tilde{K} \in \mathcal{O}_K$. We can now conclude that \mathcal{O}_K is a closed subset of \mathcal{C}_q . As a consequence of the closedness of \mathcal{O}_K , the set \mathcal{O}_K equipped with the subspace topology induced from \mathcal{C}_q is a locally compact Hausdorff space.

Now, by combining the above results and applying [28, Theorem 2.13], we can conclude that the mapping $T \mapsto \mathcal{T}_q(T, K)$ is a homeomorphism from GL_q to \mathcal{O}_K . Therefore, the mapping $T \mapsto \mathcal{T}_q(T, K)$ is a diffeomorphism from GL_q to \mathcal{O}_K , and \mathcal{O}_K is an embedded submanifold of \mathcal{C}_q with dimension given by $\dim \mathcal{O}_K = \dim \mathrm{GL}_q = q^2$. Finally, the two path-connected components of \mathcal{O}_K are immediate.

A.5 Proof of Proposition 4

Let $H \in \mathbb{R}^{q \times q}$ be arbitrary. Then for sufficiently small ϵ , we have

$$\mathcal{T}_q(I + \epsilon H, K) = \begin{bmatrix} 0 & C_K(I + \epsilon H)^{-1} \\ (I + \epsilon H)B_K & (I + \epsilon H)A_K(I + \epsilon H)^{-1} \end{bmatrix} = K + \epsilon \begin{bmatrix} 0 & -C_K H \\ H B_K & H A_K - A_K H \end{bmatrix} + o(\epsilon),$$

implying that the tangent map of $\mathcal{T}_q(\cdot, K)$ at the identity is given by

$$H \mapsto \begin{bmatrix} 0 & -C_K H \\ H B_K & H A_K - A_K H \end{bmatrix}.$$

Then since $\mathcal{T}_q(\cdot, K)$ is a diffeomorphism from GL_q to \mathcal{O}_K , the tangent map of $\mathcal{T}_q(\cdot, K)$ at the identity is an isomorphism from $\mathbb{R}^{q \times q}$ (the tangent space of GL_q at the identity) to the tangent space $T\mathcal{O}_K$. Thus

$$T\mathcal{O}_K = \left\{ \begin{bmatrix} 0 & -C_K H \\ H B_K & H A_K - A_K H \end{bmatrix} \mid H \in \mathbb{R}^{q \times q} \right\}.$$

Then the orthogonal complement $\mathcal{T}\mathcal{O}_K^\perp$ is given by

$$\begin{aligned}
 \mathcal{T}\mathcal{O}_K^\perp &= \left\{ \Delta \in \mathcal{V}_q \mid \text{tr}(U^\top \Delta) = 0 \text{ for all } U \in \mathcal{T}\mathcal{O}_K \right\} \\
 &= \left\{ \Delta \in \mathcal{V}_q \mid \text{tr} \left(\begin{bmatrix} 0 & -C_K H \\ H B_K & H A_K - A_K H \end{bmatrix}^\top \Delta \right) = 0, \forall H \in \mathbb{R}^{q \times q} \right\} \\
 &= \left\{ \begin{bmatrix} 0 & \Delta_{B_K} \\ \Delta_{C_K} & \Delta_{A_K} \end{bmatrix} \in \mathcal{V}_q \mid \text{tr} H^\top \left(\Delta_{A_K} A_K^\top - A_K^\top \Delta_{A_K} + \Delta_{B_K} B_K^\top - C_K^\top \Delta_{C_K} \right) = 0 \right\} \\
 &= \left\{ \begin{bmatrix} 0 & \Delta_{B_K} \\ \Delta_{C_K} & \Delta_{A_K} \end{bmatrix} \in \mathcal{V}_q \mid \Delta_{A_K} A_K^\top - A_K^\top \Delta_{A_K} + \Delta_{B_K} B_K^\top - C_K^\top \Delta_{C_K} = 0 \right\}.
 \end{aligned}$$

This completes the proof.

A.6 Proof of Lemma 10

Let $\Delta \in \mathcal{V}_q$ be arbitrary. We have

$$\begin{aligned}
 J_q(\mathcal{T}_q(T, K + \Delta)) - J_q(\mathcal{T}_q(T, K)) \\
 &= J_q(\mathcal{T}_q(T, K) + \mathcal{T}_q(T, \Delta)) - J_q(\mathcal{T}_q(T, K)) \\
 &= \text{tr} \left[\left(\nabla J_q \Big|_{\mathcal{T}_q(T, K)} \right)^\top \cdot \mathcal{T}_q(T, \Delta) \right] + o(\|\Delta\|) \\
 &= \text{tr} \left[\left(\nabla J_q \Big|_{\mathcal{T}_q(T, K)} \right)^\top \cdot \begin{bmatrix} I_m & 0 \\ 0 & T \end{bmatrix} \Delta \begin{bmatrix} I_p & 0 \\ 0 & T^{-1} \end{bmatrix} \right] + o(\|\Delta\|) \\
 &= \text{tr} \left[\left(\begin{bmatrix} I_m & 0 \\ 0 & T \end{bmatrix}^\top \cdot \nabla J_q \Big|_{\mathcal{T}_q(T, K)} \cdot \begin{bmatrix} I_p & 0 \\ 0 & T^{-1} \end{bmatrix}^\top \right) \Delta \right] + o(\|\Delta\|).
 \end{aligned}$$

On the other hand, Lemma 7 shows that the LQG cost stays the same when applying similarity transformation. Thus, we have

$$J_q(\mathcal{T}_T(K + \Delta)) - J_q(\mathcal{T}_T(K)) = J_q(K + \Delta) - J_q(K) = \text{tr} \left[\left(\nabla J_q \Big|_K \right)^\top \cdot \Delta \right] + o(\|\Delta\|).$$

Comparing the two equations leads to the relationship (24).

A.7 The gradient and the Hessian of $J_q(K)$

We first introduce the following lemma.

Lemma 14 Suppose $M : (-\delta, \delta) \rightarrow \mathbb{R}^{k \times k}$ and $G : (-\delta, \delta) \rightarrow \mathbb{S}^k$ are two indefinitely differentiable matrix-valued functions for some $\delta > 0$ and $k \in \mathbb{N} \setminus \{0\}$, and suppose $M(t)$ is stable for all $t \in (-\delta, \delta)$. Let $X(t)$ denote the solution to the following Lyapunov equation

$$M(t)X(t) + X(t)M(t)^\top + G(t) = 0.$$

Then $X(t)$ is indefinitely differentiable over $t \in (-\delta, \delta)$, and its j 'th order derivative at $t = 0$, denoted by $X^{(j)}(0)$, is the solution to the following Lyapunov equation

$$\begin{aligned} & M(0)X^{(j)}(0) + X^{(j)}(0)M(0)^\top \\ & + \left(\sum_{i=1}^j \frac{j!}{i!(j-i)!} \left(M^{(i)}(0)X^{(j-i)}(0) + X^{(j-i)}(0)M^{(i)}(0)^\top \right) + G^{(j)}(0) \right) = 0. \end{aligned} \quad (42)$$

Proof The differentiability of $X(t)$ follows from the observation that it can be written as $\text{vec}(X(t)) = -(I_k \otimes M(t) + M(t) \otimes I_k)^{-1} \text{vec}(G(t))$ by the vectorized form of the Lyapunov equation. Now, since $M(t)$, $G(t)$ and $X(t)$ are indefinitely differentiable, they admit Taylor expansions around $t = 0$ given by

$$\begin{aligned} M(t) &= \sum_{j=0}^a \frac{t^j}{j!} M^{(j)}(0) + o(t^a), \\ G(t) &= \sum_{j=0}^a \frac{t^j}{j!} G^{(j)}(0) + o(t^a), \\ X(t) &= \sum_{j=0}^a \frac{t^j}{j!} X^{(j)}(0) + o(t^a) \end{aligned}$$

for any $a \in \mathbb{N}$. By plugging these Taylor expansions into the original Lyapunov equation, after some algebraic manipulations, we can show that

$$\sum_{j=0}^a t^j \left[\sum_{i=0}^j \frac{1}{i!(j-i)!} \left(M^{(i)}(0)X^{(j-i)}(0) + X^{(j-i)}(0)M^{(i)}(0)^\top \right) + \frac{G^{(j)}(0)}{j!} \right] + o(t^a) = 0.$$

Since the above equality holds for all sufficiently small t , we get

$$\sum_{i=0}^j \frac{1}{i!(j-i)!} \left(M^{(i)}(0)X^{(j-i)}(0) + X^{(j-i)}(0)M^{(i)}(0)^\top \right) + \frac{1}{j!} G^{(j)}(0) = 0,$$

which is the same as (42). Thus, $X^{(j)}(0)$ is a solution to the Lyapunov Eq (42). \square

Now, given any stabilizing controller $K \in \mathcal{C}_q$, we denote the closed-loop matrix as

$$A_{\text{cl},K} = \begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix} = \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix} K \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix}$$

and recall that the LQG cost is given by

$$J_q(K) = \text{tr} \left(\begin{bmatrix} Q & 0 \\ 0 & C_K^\top R C_K \end{bmatrix} X_K \right),$$

where X_K is the unique positive semidefinite solution to the Lyapunov Eq (12a).

Consider an arbitrary direction $\Delta = \begin{bmatrix} 0 & \Delta_{C_K} \\ \Delta_{B_K} & \Delta_{A_K} \end{bmatrix} \in \mathcal{V}_q$. For sufficiently small $t > 0$ such that $K + t\Delta \in \mathcal{C}_q$, the corresponding closed-loop matrix is

$$A_{\text{cl},K+t\Delta} = A_{\text{cl},K} + t \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix} \Delta \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix},$$

and we let $X_{K,\Delta}(t)$ denote the solution to the Lyapunov equation (12a) with closed-loop matrix $A_{\text{cl},K+t\Delta}$, i.e.,

$$\begin{aligned} & \left(A_{\text{cl},K} + t \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix} \Delta \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix} \right) X_{K,\Delta}(t) + X_{K,\Delta}(t) \left(A_{\text{cl},K} + t \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix} \Delta \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix} \right)^\top \\ & \quad + \begin{bmatrix} W & 0 \\ 0 & (B_K + t\Delta_{B_K})V(B_K + t\Delta_{B_K})^\top \end{bmatrix} = 0. \end{aligned} \quad (43)$$

By Lemma 14, we see that $X_{K,\Delta}(t)$ admits a Taylor expansion of the form

$$X_{K,\Delta}(t) = X_K + t \cdot X'_{K,\Delta}(0) + \frac{t^2}{2} \cdot X''_{K,\Delta}(0) + o(t^2), \quad (44)$$

and the derivatives $X'_{K,\Delta}(0)$ and $X''_{K,\Delta}(0)$ are the solutions to the following Lyapunov equations

$$A_{\text{cl},K} X'_{K,\Delta}(0) + X'_{K,\Delta}(0) A_{\text{cl},K}^\top + M_1(X_K, \Delta) = 0, \quad (45)$$

$$A_{\text{cl},K} X''_{K,\Delta}(0) + X''_{K,\Delta}(0) A_{\text{cl},K}^\top + 2M_2(X'_{K,\Delta}(0), \Delta) = 0, \quad (46)$$

where

$$M_1(X_K, \Delta) := \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix} \Delta \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix} X_K + X_K \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix}^\top \Delta^\top \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix}^\top + \begin{bmatrix} 0 & 0 \\ 0 & B_K V \Delta_{B_K}^\top + \Delta_{B_K} V B_K^\top \end{bmatrix},$$

$$M_2(X'_{K,\Delta}(0), \Delta) := \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix} \Delta \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix} X'_{K,\Delta}(0) + X'_{K,\Delta}(0) \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix}^\top \Delta^\top \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix}^\top + \begin{bmatrix} 0 & 0 \\ 0 & \Delta_{B_K} V \Delta_{B_K}^\top \end{bmatrix}.$$

Now, by plugging the Taylor expansion (44) into the expression (11) for $J_q(\mathbf{K})$, we get

$$\begin{aligned} J_q(\mathbf{K} + t\Delta) &= \text{tr} \left(\begin{bmatrix} Q & 0 \\ 0 & (C_{\mathbf{K}} + t\Delta_{C_{\mathbf{K}}})^T R (C_{\mathbf{K}} + t\Delta_{C_{\mathbf{K}}}) \end{bmatrix} X_{\mathbf{K}, \Delta}(t) \right) \\ &= J_q(\mathbf{K}) + t \cdot \text{tr} \left(\begin{bmatrix} Q & 0 \\ 0 & C_{\mathbf{K}}^T R C_{\mathbf{K}} \end{bmatrix} X'_{\mathbf{K}, \Delta}(0) + \begin{bmatrix} 0 & 0 \\ 0 & C_{\mathbf{K}}^T R \Delta_{C_{\mathbf{K}}} + \Delta_{C_{\mathbf{K}}}^T R C_{\mathbf{K}} \end{bmatrix} X_{\mathbf{K}} \right) \\ &\quad + \frac{t^2}{2} \cdot \text{tr} \left(\begin{bmatrix} Q & 0 \\ 0 & C_{\mathbf{K}}^T R C_{\mathbf{K}} \end{bmatrix} X''_{\mathbf{K}, \Delta}(0) + 2 \begin{bmatrix} 0 & 0 \\ 0 & C_{\mathbf{K}}^T R \Delta_{C_{\mathbf{K}}} + \Delta_{C_{\mathbf{K}}}^T R C_{\mathbf{K}} \end{bmatrix} X'_{\mathbf{K}, \Delta}(0) \right. \\ &\quad \left. + 2 \begin{bmatrix} 0 & 0 \\ 0 & \Delta_{C_{\mathbf{K}}}^T R \Delta_{C_{\mathbf{K}}} \end{bmatrix} X_{\mathbf{K}} \right) + o(t^2), \end{aligned}$$

from which we can directly recognize $\frac{dJ_q(\mathbf{K}+t\Delta)}{dt} \Big|_{t=0}$ and $\frac{d^2J_q(\mathbf{K}+t\Delta)}{dt^2} \Big|_{t=0}$.

Now suppose X is the solution to the following Lyapunov equation $A_{\text{cl}, \mathbf{K}} X + X A_{\text{cl}, \mathbf{K}}^T + M = 0$ for some $M \in \mathbb{S}^{n+q}$. Then, by [48, Lemma 3.18], the solution to the above Lyapunov equation can be written as $X = \int_0^{+\infty} e^{A_{\text{cl}, \mathbf{K}} s} M e^{A_{\text{cl}, \mathbf{K}}^T s} ds$, and consequently

$$\begin{aligned} \text{tr} \left(\begin{bmatrix} Q & 0 \\ 0 & C_{\mathbf{K}}^T R C_{\mathbf{K}} \end{bmatrix} X \right) &= \int_0^{+\infty} \text{tr} \left(\begin{bmatrix} Q & 0 \\ 0 & C_{\mathbf{K}}^T R C_{\mathbf{K}} \end{bmatrix} e^{A_{\text{cl}, \mathbf{K}} s} M e^{A_{\text{cl}, \mathbf{K}}^T s} \right) ds \\ &= \int_0^{+\infty} \text{tr} \left(e^{A_{\text{cl}, \mathbf{K}}^T s} \begin{bmatrix} Q & 0 \\ 0 & C_{\mathbf{K}}^T R C_{\mathbf{K}} \end{bmatrix} e^{A_{\text{cl}, \mathbf{K}} s} M \right) ds = \text{tr}(Y_{\mathbf{K}} M), \end{aligned}$$

in which we recall that $Y_{\mathbf{K}}$ is the unique positive semidefinite solution to Lyapunov Eq (12b). Therefore the first-order derivative $\frac{dJ_q(\mathbf{K}+t\Delta)}{dt} \Big|_{t=0}$ can be alternatively given by

$$\begin{aligned} \frac{dJ_q(\mathbf{K}+t\Delta)}{dt} \Big|_{t=0} &= \text{tr} \left(Y_{\mathbf{K}} M_1(X_{\mathbf{K}}, \Delta) + \begin{bmatrix} 0 & 0 \\ 0 & C_{\mathbf{K}}^T R \Delta_{C_{\mathbf{K}}} + \Delta_{C_{\mathbf{K}}}^T R C_{\mathbf{K}} \end{bmatrix} X_{\mathbf{K}} \right) \\ &= 2 \text{tr} \left[\left(\begin{bmatrix} 0 & R C_{\mathbf{K}} \\ 0 & 0 \end{bmatrix} X_{\mathbf{K}} \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} + \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix}^T Y_{\mathbf{K}} X_{\mathbf{K}} \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix}^T + \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} Y_{\mathbf{K}} \begin{bmatrix} 0 & 0 \\ B_{\mathbf{K}} V & 0 \end{bmatrix} \right)^T \Delta \right]. \end{aligned}$$

One can readily recognize the gradient $\nabla J_q(\mathbf{K})$ as $\frac{dJ_q(\mathbf{K}+t\Delta)}{dt} \Big|_{t=0} = \text{tr}(\nabla J_q(\mathbf{K})^T \Delta)$. Upon partitioning $X_{\mathbf{K}}$ and $Y_{\mathbf{K}}$ as (22), a few simple calculations lead to the gradient formula of $J_q(\mathbf{K})$ in (21).

Similarly, we can show that the second-order derivative $\frac{d^2 J_q(K+t\Delta)}{dt^2} \Big|_{t=0}$ can be alternatively given by

$$\begin{aligned} \frac{d^2 J_q(K+t\Delta)}{dt^2} \Big|_{t=0} &= 2 \operatorname{tr} \left(Y_K M_2(X'_{K,\Delta}(0), \Delta) + \begin{bmatrix} 0 & 0 \\ 0 & C_K^\top R \Delta C_K + \Delta_{C_K}^\top R C_K \end{bmatrix} X'_{K,\Delta}(0) \right. \\ &\quad \left. + \begin{bmatrix} 0 & 0 \\ 0 & \Delta_{C_K}^\top R \Delta C_K \end{bmatrix} X_K \right) \\ &= 2 \operatorname{tr} \left(2 \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix} \Delta \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix} X'_{K,\Delta}(0) Y_K + 2 \begin{bmatrix} 0 & 0 \\ 0 & C_K^\top R \Delta C_K \end{bmatrix} X'_{K,\Delta}(0) \right. \\ &\quad \left. + \begin{bmatrix} 0 & 0 \\ 0 & \Delta_{B_K} V \Delta_{B_K}^\top \end{bmatrix} Y_K + \begin{bmatrix} 0 & 0 \\ 0 & \Delta_{C_K}^\top R \Delta C_K \end{bmatrix} X_K \right). \end{aligned}$$

Then noticing that $\operatorname{Hess}_K(\Delta, \Delta) = \frac{d^2 J_q(K+t\Delta)}{dt^2} \Big|_{t=0}$ for any $\Delta \in \mathcal{V}_q$, we get the desired expression for the Hessian of J_q .

A.8 Proof of Lemma 11

By [48, Lemma 3.18], given a stable matrix A , if (C, A) is observable, then the solution L to the Lyapunov equation $A^\top L + L A + C^\top C = 0$ is positive definite. Therefore, we only need to prove that

$$\left(\begin{bmatrix} Q^{\frac{1}{2}} & 0 \\ 0 & R^{\frac{1}{2}} C_K \end{bmatrix}, \begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix} \right)$$

is observable. By [48, Theorem 3.3], this is equivalent to showing that the eigenvalues of the following matrix

$$\begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix} + \begin{bmatrix} L_{11} & L_{12} \\ L_{21} & L_{22} \end{bmatrix} \begin{bmatrix} Q^{\frac{1}{2}} & 0 \\ 0 & R^{\frac{1}{2}} C_K \end{bmatrix} = \begin{bmatrix} A + L_{11} Q^{\frac{1}{2}} & BC_K + L_{12} R^{\frac{1}{2}} C_K \\ B_K C + L_{21} Q^{\frac{1}{2}} & A_K + L_{22} R^{\frac{1}{2}} C_K \end{bmatrix}$$

can be arbitrarily assigned by choosing $L_{11}, L_{12}, L_{21}, L_{22}$. This is indeed true by choosing $L_{12} = -B R^{-\frac{1}{2}}$ and observing that $A + L_{11} Q^{\frac{1}{2}}$ and $A_K + L_{22} R^{\frac{1}{2}} C_K$ can be arbitrarily assigned since $(Q^{\frac{1}{2}}, A), (C_K, A_K)$ are both observable. Thus, by [48, Lemma 3.18], the solution Y_K to (12b) is positive definite. Similarly, we can prove X_K is positive definite.

A.9 Proof of Theorem 5

We first show that K^* is a stationary point of $J_n(K)$ over $K \in \mathcal{C}_n$. Since $\mathcal{T}_n(-I_n, K^*) = K^*$, by Lemma 10, we have

$$\nabla J_n \Big|_{K^*} = \nabla J_n \Big|_{\mathcal{T}_n(-I_n, K^*)} = \begin{bmatrix} I_m & 0 \\ 0 & -I_n \end{bmatrix} \cdot \nabla J_n \Big|_{K^*} \cdot \begin{bmatrix} I_p & 0 \\ 0 & -I_n \end{bmatrix}.$$

This equality implies that, excluding the bottom right $n \times n$ block, the last n rows and the last n columns of $\nabla J_n \Big|_{K^*}$ are zero. On the other hand, it is not hard to see that $J_n(K^*)$ does not depend on the choice of Λ as long as Λ is stable. Therefore the bottom right $n \times n$ block of $\nabla J_n \Big|_{K^*}$ is zero. We can now see that $\nabla J_n \Big|_{K^*} = 0$, showing that K^* is a stationary point of J_n .

Let $\Delta = \begin{bmatrix} 0 & \Delta_{C_K} \\ \Delta_{B_K} & \Delta_{A_K} \end{bmatrix} \in \mathcal{V}_n$ be arbitrary, and let

$$\Delta^{(1)} = \begin{bmatrix} 0 & \Delta_{C_K} \\ 0 & 0 \end{bmatrix}, \quad \Delta^{(2)} = \begin{bmatrix} 0 & 0 \\ \Delta_{B_K} & 0 \end{bmatrix}, \quad \Delta^{(3)} = \begin{bmatrix} 0 & 0 \\ 0 & \Delta_{A_K} \end{bmatrix}.$$

By the bilinearity of the Hessian, we have

$$\text{Hess}_{K^*}(\Delta, \Delta) = \sum_{1 \leq i < j \leq 3} \text{Hess}_{K^*}(\Delta^{(i)} + \Delta^{(j)}, \Delta^{(i)} + \Delta^{(j)}) - \sum_{i=1}^3 \text{Hess}_{K^*}(\Delta^{(i)}, \Delta^{(i)}).$$

Since the controllers $K^* + t\Delta^{(i)}$ for $i = 1, 2, 3$ and $K^* + t(\Delta^{(i)} + \Delta^{(3)})$ for $i = 1, 2$ have the same transfer function representation as K^* , we can see that for all sufficiently small t ,

$$\begin{aligned} J_n(K^*) &= J_n(K^* + t\Delta^{(1)}) = J_n(K^* + t\Delta^{(2)}) = J_n(K^* + t\Delta^{(3)}) \\ &= J_n(K^* + t(\Delta^{(1)} + \Delta^{(3)})) = J_n(K^* + t(\Delta^{(2)} + \Delta^{(3)})), \end{aligned}$$

which implies that $\text{Hess}_{K^*}(\Delta^{(i)}, \Delta^{(i)}) = 0$ for all $i = 1, 2, 3$, and that $\text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(3)}, \Delta^{(1)} + \Delta^{(3)}) = \text{Hess}_{K^*}(\Delta^{(2)} + \Delta^{(3)}, \Delta^{(2)} + \Delta^{(3)}) = 0$. Therefore

$$\text{Hess}_{K^*}(\Delta, \Delta) = \text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)}).$$

Now, if $\text{Hess}_{K^*}(\Delta, \Delta) = 0$ for all $\Delta \in \mathcal{V}_n$, then the Hessian Hess_{K^*} is obviously zero. Otherwise, $\text{Hess}_{K^*}(\Delta, \Delta) \neq 0$ for some $\Delta \in \mathcal{V}_n$, which implies that

$$\begin{aligned}\text{Hess}_{K^*}(\Delta^{(1)}, \Delta^{(2)}) &= \frac{1}{2} \left(\text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)}) \right. \\ &\quad \left. - \text{Hess}_{K^*}(\Delta^{(1)}, \Delta^{(1)}) - \text{Hess}_{K^*}(\Delta^{(2)}, \Delta^{(2)}) \right) \\ &= \frac{1}{2} \text{Hess}_{K^*}(\Delta, \Delta) \neq 0.\end{aligned}$$

Note that $\Delta^{(1)}$ and $\Delta^{(2)}$ are linearly independent (otherwise $\text{Hess}_{K^*}(\Delta^{(1)}, \Delta^{(2)})$ will be zero). Together with $\text{Hess}_{K^*}(\Delta^{(i)}, \Delta^{(i)}) = 0$ for $i = 1, 2$, we see that Hess_{K^*} must be indefinite (a symmetric matrix having a 2×2 principal submatrix with zero diagonal entries and non-zero off-diagonal entries must be indefinite).

Now we proceed to the situation where Λ is diagonalizable. We let $e_i^{(k)}$ denote the k -dimensional vector where only the i th entry is 1 and other entries are zero.

Part I: $\text{eig}(-\Lambda) \not\subseteq \mathcal{Z} \implies \text{the Hessian is indefinite.}$ Let $\lambda \in \text{eig}(-\Lambda) \setminus \mathcal{Z}$. Since $\lambda \notin \mathcal{Z}$, there exists some i, j such that

$$G(\lambda) := e_i^{(p)\top} C X_{\text{op}} (\lambda I - A^\top)^{-1} Y_{\text{op}} B e_j^{(m)} \neq 0.$$

We shall only provide the proof for the situation when λ is real. When λ is complex, the proof employs similar techniques but is more complicated, and we refer interested readers to our online report [47].

Let T be a real invertible matrix such that $T \Lambda T^{-1} = \begin{bmatrix} -\lambda & 0 \\ 0 & * \end{bmatrix}$. Let $\Delta^{(1)}, \Delta^{(2)} \in \mathcal{V}_n$ be given by

$$\Delta^{(1)} = \begin{bmatrix} 0 & \Delta_{C_K}^{(1)} \\ 0 & 0 \end{bmatrix}, \quad \Delta^{(2)} = \begin{bmatrix} 0 & 0 \\ \Delta_{B_K}^{(2)} & 0 \end{bmatrix},$$

where $\Delta_{C_K}^{(1)} = e_j^{(m)} e_1^{(n)\top} T^{-1}$ and $\Delta_{B_K}^{(2)} = T e_1^{(n)} e_i^{(p)\top}$. Then it's not hard to see that $J_n(K^* + t \Delta^{(1)}) = J_n(K^* + t \Delta^{(2)}) = J_n(K^*)$ for any sufficiently small t , indicating that both $\text{Hess}_{K^*}(\Delta^{(1)}, \Delta^{(1)})$ and $\text{Hess}_{K^*}(\Delta^{(2)}, \Delta^{(2)})$ are equal to zero. On the other hand, we have that the unique solutions to Lyapunov Eqs (12a) and (12b) are $X_{K^*} = \begin{bmatrix} X_{\text{op}} & 0 \\ 0 & 0 \end{bmatrix}$, $Y_{K^*} = \begin{bmatrix} Y_{\text{op}} & 0 \\ 0 & 0 \end{bmatrix}$. By Lemma 9, we can see that

$$\text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)}) = 4 \text{tr} \left(\begin{bmatrix} 0 & B \Delta_{C_K}^{(1)} \\ \Delta_{B_K}^{(2)} C & 0 \end{bmatrix} X'_{K^*, \Delta^{(1)} + \Delta^{(2)}} \begin{bmatrix} Y_{\text{op}} & 0 \\ 0 & 0 \end{bmatrix} \right),$$

where $X'_{K^*, \Delta^{(1)} + \Delta^{(2)}}$ is the solution to the following Lyapunov equation

$$\begin{aligned} & \begin{bmatrix} A & 0 \\ 0 & \Lambda \end{bmatrix} X'_{K^*, \Delta^{(1)} + \Delta^{(2)}} + X'_{K^*, \Delta^{(1)} + \Delta^{(2)}} \begin{bmatrix} A & 0 \\ 0 & \Lambda \end{bmatrix}^T \\ & + \begin{bmatrix} 0 & B \Delta_{C_K}^{(1)} \\ \Delta_{B_K}^{(2)} C & 0 \end{bmatrix} \begin{bmatrix} X_{\text{op}} & 0 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} X_{\text{op}} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & B \Delta_{C_K}^{(1)} \\ \Delta_{B_K}^{(2)} C & 0 \end{bmatrix}^T = 0. \end{aligned}$$

Since

$$\begin{bmatrix} 0 & B \Delta_{C_K}^{(1)} \\ \Delta_{B_K}^{(2)} C & 0 \end{bmatrix} \begin{bmatrix} X_{\text{op}} & 0 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} X_{\text{op}} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & B \Delta_{C_K}^{(1)} \\ \Delta_{B_K}^{(2)} C & 0 \end{bmatrix}^T = \begin{bmatrix} 0 & X_{\text{op}} C^T \Delta_{B_K}^{(2)T} \\ \Delta_{B_K}^{(2)} C X_{\text{op}} & 0 \end{bmatrix},$$

the matrix $X'_{K^*, \Delta^{(1)} + \Delta^{(2)}}$ can be represented by

$$\begin{aligned} X'_{K^*, \Delta^{(1)} + \Delta^{(2)}} &= \int_0^{+\infty} \exp\left(\begin{bmatrix} A & 0 \\ 0 & \Lambda \end{bmatrix} s\right) \begin{bmatrix} 0 & X_{\text{op}} C^T \Delta_{B_K}^{(2)T} \\ \Delta_{B_K}^{(2)} C X_{\text{op}} & 0 \end{bmatrix} \exp\left(\begin{bmatrix} A & 0 \\ 0 & \Lambda \end{bmatrix}^T s\right) ds \\ &= \int_0^{+\infty} \begin{bmatrix} e^{As} & 0 \\ 0 & e^{\Lambda s} \end{bmatrix} \begin{bmatrix} 0 & X_{\text{op}} C^T \Delta_{B_K}^{(2)T} \\ \Delta_{B_K}^{(2)} C X_{\text{op}} & 0 \end{bmatrix} \begin{bmatrix} e^{A^T s} & 0 \\ 0 & e^{\Lambda^T s} \end{bmatrix} ds \\ &= \int_0^{+\infty} \begin{bmatrix} 0 & e^{As} X_{\text{op}} C^T \Delta_{B_K}^{(2)T} e^{A^T s} \\ e^{\Lambda s} \Delta_{B_K}^{(2)} C X_{\text{op}} e^{A^T s} & 0 \end{bmatrix} ds, \end{aligned}$$

which can be shown to lead to

$$\text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)}) = \int_0^{+\infty} 4 \text{tr} \left(B \Delta_{C_K}^{(1)} e^{\Lambda s} \Delta_{B_K}^{(2)} C X_{\text{op}} e^{A^T s} Y_{\text{op}} \right) ds.$$

By the construction of $\Delta_{C_K}^{(1)}$ and $\Delta_{B_K}^{(2)}$, we have $\Delta_{C_K}^{(1)} e^{\Lambda s} \Delta_{B_K}^{(2)} = e^{-\lambda s} e_j^{(m)} e_i^{(p)T}$, and thus

$$\begin{aligned} \text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)}) &= \int_0^{+\infty} 4 e_i^{(p)T} C X_{\text{op}} e^{(A^T - \lambda I)s} Y_{\text{op}} B e_j^{(m)} ds \\ &= 4 e_i^{(p)T} C X_{\text{op}} (\lambda I - A^T)^{-1} Y_{\text{op}} B e_j^{(m)} = 4 G(\lambda), \end{aligned}$$

which is nonzero by assumption. Consequently,

$$\begin{aligned} \text{Hess}_{K^*}(\Delta^{(1)}, \Delta^{(2)}) &= \frac{1}{2} (\text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)}) \\ &\quad - \text{Hess}_{K^*}(\Delta^{(1)}, \Delta^{(1)}) - \text{Hess}_{K^*}(\Delta^{(2)}, \Delta^{(2)})) \\ &= 2G(\lambda) \neq 0. \end{aligned}$$

Together with the fact that $\text{Hess}_{K^*}(\Delta^{(1)}, \Delta^{(1)}) = \text{Hess}_{K^*}(\Delta^{(2)}, \Delta^{(2)}) = 0$, we can see that neither Hess_{K^*} nor $-\text{Hess}_{K^*}$ can be positive semidefinite. Thus Hess_{K^*} has at least one positive eigenvalue and one negative eigenvalue.

Part II: $\text{eig}(-\Lambda) \subseteq \mathcal{Z} \implies \text{the Hessian is zero.}$ In this part, we will show that if $\text{eig}(-\Lambda) \subseteq \mathcal{Z}$, then $\text{Hess}_{K^*}(\Delta, \Delta) = 0$ for any $\Delta \in \mathcal{V}_n$.

Let $\Delta = \begin{bmatrix} 0 & \Delta_{C_K} \\ \Delta_{B_K} & \Delta_{A_K} \end{bmatrix} \in \mathcal{V}_n$ be arbitrary. Let

$$\Delta^{(1)} = \begin{bmatrix} 0 & \Delta_{C_K} \\ 0 & 0 \end{bmatrix}, \quad \Delta^{(2)} = \begin{bmatrix} 0 & 0 \\ \Delta_{B_K} & 0 \end{bmatrix}, \quad \Delta^{(3)} = \begin{bmatrix} 0 & 0 \\ 0 & \Delta_{A_K} \end{bmatrix}.$$

We have already shown that $\text{Hess}_{K^*}(\Delta, \Delta) = \text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)})$. Let T be an invertible $n \times n$ (complex) matrix that diagonalizes Λ as $T \Lambda T^{-1} = \text{diag}(-\lambda_1, \dots, -\lambda_n)$. Define $U_{ik} = e_i^{(m)} e_k^{(n)\top} T^{-1}$, $V_{jk} = T e_k^{(n)} e_j^{(p)\top}$ for each $1 \leq i \leq m$, $1 \leq j \leq p$ and $1 \leq k \leq n$. It is not hard to see that $\{U_{ik} \mid 1 \leq i \leq m, 1 \leq k \leq n\}$ forms a basis of $\mathbb{C}^{m \times n}$, and $\{V_{jk} \mid 1 \leq j \leq n, 1 \leq k \leq n\}$ forms a basis of $\mathbb{C}^{n \times q}$. Therefore Δ_{C_K} and Δ_{B_K} can be expanded as

$$\Delta_{C_K} = \sum_{1 \leq i \leq m} \sum_{1 \leq k \leq n} \alpha_{ik} U_{ik}, \quad \Delta_{B_K} = \sum_{1 \leq j \leq q} \sum_{1 \leq k \leq n} \beta_{jk} V_{jk}.$$

By similar derivations as in Case 1, we can get

$$\text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)}) = \int_0^{+\infty} 4 \text{tr} \left(B \Delta_{C_K} e^{\Lambda s} \Delta_{B_K} C X_{\text{op}} e^{\Lambda^T s} Y_{\text{op}} \right) ds.$$

Then, since

$$\begin{aligned} \Delta_{C_K} e^{\Lambda s} \Delta_{B_K} &= \sum_{1 \leq i \leq m} \sum_{1 \leq j \leq q} \sum_{1 \leq k \leq n} \sum_{1 \leq k' \leq n} \alpha_{ik} \beta_{jk'} U_{ik} e^{\Lambda s} V_{jk'} \\ &= \sum_{1 \leq i \leq m} \sum_{1 \leq j \leq q} \sum_{1 \leq k \leq n} \sum_{1 \leq k' \leq n} \alpha_{ik} \beta_{jk'} e_i^{(m)} e_k^{(n)\top} \begin{bmatrix} e^{-\lambda_1 s} & & & \\ & \ddots & & \\ & & e^{-\lambda_n s} & \end{bmatrix} e_{k'}^{(n)} e_j^{(p)\top} \\ &= \sum_{1 \leq i \leq m} \sum_{1 \leq j \leq q} \sum_{1 \leq k \leq n} \alpha_{ik} \beta_{jk} e^{-\lambda_k s} e_i^{(m)} e_j^{(p)\top}, \end{aligned}$$

we have

$$\begin{aligned} &\text{Hess}_{K^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)}) \\ &= \sum_{1 \leq i \leq m} \sum_{1 \leq j \leq q} \sum_{1 \leq k \leq n} \int_0^{+\infty} 4 \alpha_{ik} \beta_{jk} \cdot e_j^{(p)\top} C X_{\text{op}} e^{(A - \lambda_k I)^T s} Y_{\text{op}} B e_i^{(m)} ds \\ &= \sum_{1 \leq i \leq m} \sum_{1 \leq j \leq q} \sum_{1 \leq k \leq n} 4 \alpha_{ik} \beta_{jk} \cdot e_j^{(p)\top} C X_{\text{op}} (\lambda_k I - A^T)^{-1} Y_{\text{op}} B e_i^{(m)}. \end{aligned}$$

Since $\text{eig}(-\Lambda) \setminus \mathcal{Z} = \emptyset$, we can see that $CX_{\text{op}}(\lambda_k I - A^T)^{-1}Y_{\text{op}}B = 0$ for any $1 \leq k \leq n$. Therefore, we have $\text{Hess}_{\mathcal{K}^*}(\Delta, \Delta) = \text{Hess}_{\mathcal{K}^*}(\Delta^{(1)} + \Delta^{(2)}, \Delta^{(1)} + \Delta^{(2)}) = 0$, which completes the proof.

References

1. Absil, P.A., Mahony, R., Andrews, B.: Convergence of the iterates of descent methods for analytic cost functions. *SIAM J. Optim.* **16**(2), 531–547 (2005)
2. Basuthakur, S., Knapp, C.H.: Optimal constant controllers for stochastic linear systems. *IEEE Trans. Autom. Control* **20**(5), 664–666 (1975)
3. Bertsekas, D.P.: *Dynamic Programming and Optimal Control*, vol. 1. Athena Scientific, Belmont, MA, USA (1995)
4. Boczar, R., Matni, N., Recht, B.: Finite-data performance guarantees for the output-feedback control of an unknown system. In: *Proceedings of the 57th IEEE Conference on Decision and Control (CDC)*, pp. 2994–2999 (2018)
5. Boyd, S., El Ghaoui, L., Feron, E., Balakrishnan, V.: *Linear matrix inequalities in system and control theory*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (1994)
6. Bu, J., Mesbahi, A., Mesbahi, M.: On topological and metrical properties of stabilizing feedback gains: the MIMO case. *arXiv preprint arXiv:1904.02737* (2019)
7. Chi, Y., Lu, Y.M., Chen, Y.: Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Trans. Signal Process.* **67**(20), 5239–5269 (2019)
8. Dean, S., Mania, H., Matni, N., Recht, B., Tu, S.: On the sample complexity of the linear quadratic regulator. *Found. Comput. Math.* **20**, 633–679 (2020)
9. Doyle, J.C.: Guaranteed margins for LQG regulators. *IEEE Trans. Autom. Control* **23**(4), 756–757 (1978)
10. Fatkhullin, I., Polyak, B.: Optimizing static linear feedback: gradient method. *SIAM J. Control. Optim.* **59**(5), 3887–3911 (2021)
11. Fazel, M., Ge, R., Kakade, S., Mesbahi, M.: Global convergence of policy gradient methods for the linear quadratic regulator. In: *Proceedings of the 35th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 80, pp. 1467–1476 (2018)
12. Feng, H., Lavaei, J.: Connectivity properties of the set of stabilizing static decentralized controllers. *SIAM J. Control. Optim.* **58**(5), 2790–2820 (2020)
13. Furieri, L., Zheng, Y., Kamgarpour, M.: Learning the globally optimal distributed LQ regulator. In: *Proceedings of the 2nd Conference on Learning for Dynamics and Control, Proceedings of Machine Learning Research*, vol. 120, pp. 287–297 (2020)
14. Furieri, L., Zheng, Y., Papachristodoulou, A., Kamgarpour, M.: An input-output parametrization of stabilizing controllers: Amidst Youla and system level synthesis. *IEEE Control Syst. Lett.* **3**(4), 1014–1019 (2019)
15. Gahinet, P., Apkarian, P.: A linear matrix inequality approach to \mathcal{H}_∞ control. *Int. J. Robust Nonlinear Control* **4**(4), 421–448 (1994)
16. Ge, R., Ma, T.: On the optimization landscape of tensor decompositions. *Adv. Neural Inf. Process. Syst.*, vol. 30 (2017)
17. Hyland, D., Bernstein, D.: The optimal projection equations for fixed-order dynamic compensation. *IEEE Trans. Autom. Control* **29**(11), 1034–1037 (1984)
18. Jin, C., Ge, R., Netrapalli, P., Kakade, S.M., Jordan, M.I.: How to escape saddle points efficiently. In: *Proceedings of the 34th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 70, pp. 1724–1732 (2017)
19. Lale, S., Azizzadenesheli, K., Hassibi, B., Anandkumar, A.: Logarithmic regret bound in partially observable linear dynamical systems. *Adv. Neural Inf. Process. Syst.* **33**, 20876–88 (2020)
20. Lale, S., Azizzadenesheli, K., Hassibi, B., Anandkumar, A.: Adaptive control and regret minimization in linear quadratic Gaussian (LQG) setting. In: *2021 American Control Conference (ACC)*, pp. 2517–2522 (2021)
21. Lee, J.M.: *Introduction to Smooth Manifolds*, 2nd edn. Springer, New York, NY, USA (2013)

22. Li, X., Lu, J., Arora, R., Haupt, J., Liu, H., Wang, Z., Zhao, T.: Symmetry, saddle points, and global optimization landscape of nonconvex matrix factorization. *IEEE Trans. Inf. Theory* **65**(6), 3489–3514 (2019)

23. Li, Y., Tang, Y., Zhang, R., Li, N.: Distributed reinforcement learning for decentralized linear quadratic control: A derivative-free policy optimization approach. *IEEE Trans. Autom. Control* **67**(12), 6429–6444 (2022)

24. Malik, D., Pananjady, A., Bhatia, K., Khamaru, K., Bartlett, P., Wainwright, M.: Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. In: Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research, vol. 89, pp. 2916–2925 (2019)

25. Masubuchi, I., Ohara, A., Suda, N.: LMI-based controller synthesis: a unified formulation and solution. *Int. J. Robust Nonlinear Control* **8**(8), 669–686 (1998)

26. Mohammadi, H., Soltanolkotabi, M., Jovanović, M.R.: On the lack of gradient domination for linear quadratic gaussian problems with incomplete state information. In: Proceedings of the 60th IEEE Conference on Decision and Control (CDC), pp. 1120–1124 (2021)

27. Mohammadi, H., Zare, A., Soltanolkotabi, M., Jovanović, M.R.: Convergence and sample complexity of gradient methods for the model-free linear-quadratic regulator problem. *IEEE Trans. Autom. Control* **67**(5), 2435–2450 (2021)

28. Montgomery, D., Zippin, L.: *Topological Transformation Groups*. Dover Publications, Mineola, NY, USA (2018)

29. Oymak, S., Ozay, N.: Non-asymptotic identification of LTI systems from a single trajectory. In: 2019 American Control Conference (ACC), pp. 5655–5661 (2019)

30. Qu, Q., Zhai, Y., Li, X., Zhang, Y., Zhu, Z.: Analysis of the optimization landscapes for overcomplete representation learning. *arXiv preprint arXiv:1912.02427* (2019)

31. Recht, B.: A tour of reinforcement learning: the view from continuous control. *Ann. Rev. Control, Robot, Autonom Syst.* **2**, 253–279 (2019)

32. Scherer, C., Gahinet, P., Chilali, M.: Multiobjective output-feedback control via LMI optimization. *IEEE Trans. Autom. Control* **42**(7), 896–911 (1997)

33. Simchowitz, M., Singh, K., Hazan, E.: Improper learning for non-stochastic control. In: Proceedings of Thirty Third Conference on Learning Theory, Proceedings of Machine Learning Research, vol. 125, pp. 3320–3436 (2020)

34. Sun, J., Qu, Q., Wright, J.: A geometric analysis of phase retrieval. *Found. Comput. Math.* **18**, 1131–1198 (2018)

35. Tang, Y., Zheng, Y., Li, N.: Analysis of the optimization landscape of linear quadratic Gaussian (LQG) control. In: Proceedings of the 3rd Conference on Learning for Dynamics and Control, Proceedings of Machine Learning Research, vol. 144, pp. 599–610 (2021)

36. Tu, S., Boczar, R., Packard, A., Recht, B.: Non-asymptotic analysis of robust control from coarse-grained identification. *arXiv preprint arXiv:1707.04791* (2017)

37. Tu, S., Recht, B.: The gap between model-based and model-free methods on the linear quadratic regulator: An asymptotic viewpoint. In: Proceedings of the Thirty-Second Conference on Learning Theory, Proceedings of Machine Learning Research, vol. 99, pp. 3036–3083 (2019)

38. Umenberger, J., Ferizbegovic, M., Schön, T.B., Hjalmarsson, H.: Robust exploration in linear quadratic reinforcement learning. *Adv. Neural Inf. Process. Syst.*, vol. 32 (2019)

39. Wang, Y.S., Matni, N., Doyle, J.C.: A system-level approach to controller synthesis. *IEEE Trans. Autom. Control* **64**(10), 4079–4093 (2019)

40. Youla, D.C., Jabr, H.A., Bongiorno, J.J., Jr.: Modern Wiener-Hopf design of optimal controllers – Part II: the multivariable case. *IEEE Trans. Autom. Control* **21**(3), 319–338 (1976)

41. Youssuff, A., Skelton, R.E.: A note on balanced controller reduction. *IEEE Trans. Autom. Control* **29**(3), 254–257 (1984)

42. Zhang, K., Hu, B., Basar, T.: Policy optimization for \mathcal{H}_2 linear control with \mathcal{H}_∞ robustness guarantee: Implicit regularization and global convergence. In: Proceedings of the 2nd Conference on Learning for Dynamics and Control, Proceedings of Machine Learning Research, vol. 120, pp. 179–190 (2020)

43. Zhang, Y., Qu, Q., Wright, J.: From symmetry to geometry: tractable nonconvex problems. *arXiv preprint arXiv:2007.06753* (2020)

44. Zheng, Y., Furieri, L., Kamgarpour, M., Li, N.: Sample complexity of linear quadratic gaussian (LQG) control for output feedback systems. In: Proceedings of the 3rd Conference on Learning for Dynamics and Control, Proceedings of Machine Learning Research, vol. 144, pp. 559–570 (2021)

45. Zheng, Y., Furieri, L., Papachristodoulou, A., Li, N., Kamgarpour, M.: On the equivalence of Youla, system-level and input-output parameterizations. *IEEE Trans. Autom. Control* **66**(1), 413–420 (2021)
46. Zheng, Y., Li, N.: Non-asymptotic identification of linear dynamical systems using multiple trajectories. *IEEE Control Syst. Lett.* **5**(5), 1693–1698 (2021)
47. Zheng, Y., Tang, Y., Li, N.: Analysis of the optimization landscape of linear quadratic gaussian (LQG) control. Technical report. [arXiv:2102.04393](https://arxiv.org/abs/2102.04393) (2021)
48. Zhou, K., Doyle, J.C., Glover, K.: Robust and Optimal Control. Prentice Hall, Upper Saddle River, NJ, USA (1996)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.