

Connectivity of the Feasible and Sublevel Sets of Dynamic Output Feedback Control with Robustness Constraints

Bin Hu¹, and Yang Zheng²

Abstract—This paper considers the optimization landscape of linear dynamic output feedback control with \mathcal{H}_∞ robustness constraints. We consider the feasible set of all the stabilizing full-order dynamical controllers that satisfy an additional \mathcal{H}_∞ robustness constraint. We show that this \mathcal{H}_∞ -constrained set has at most two path-connected components that are diffeomorphic under a mapping defined by a similarity transformation. Our proof technique utilizes a classical change of variables in \mathcal{H}_∞ control to establish a surjective mapping from a set with a convex projection to the \mathcal{H}_∞ -constrained set. This proof idea can also be used to establish the same topological properties of strict sublevel sets of linear quadratic Gaussian (LQG) control and optimal \mathcal{H}_∞ control. Our results bring positive news for gradient-based policy search on robust control problems.

Index Terms—Optimization landscape, sublevel set, direct policy search, \mathcal{H}_∞ control, LQG control

I. INTRODUCTION

Inspired by the impressive successes of reinforcement learning, model-free policy optimization techniques are receiving renewed interests from the controls field. Indeed, we have seen significant recent advances on understanding the theoretical properties of policy optimization methods on benchmark control problems, such as linear quadratic regulator (LQR) [1]–[4], linear robust control [5]–[8], and Markov jump linear quadratic control [9]–[11].

It is well-known that all these control problems are non-convex in the policy space. Classical control theory typically parameterizes the control policies into a convex domain over which efficient optimization algorithms exist [12]. An important recent discovery is that despite non-convexity, many state-feedback control problems (e.g., LQR) admit a useful property of *gradient dominance* [1]. Therefore, model-free policy search methods are guaranteed to enjoy global convergence for these problems [1], [3], [9]. Note that most convergence results require a direct access of the underlying system state, in which a simple change of variables exist to get a convex reformulation of the control problems [13].

For real-world control applications, however, we may only have access to partial output measurements. In the output feedback case, the theoretical results for direct policy search

B. Hu is generously supported by the NSF award CAREER-2048168 and the 2020 Amazon research award.

¹Bin Hu is with the Coordinated Science Laboratory (CSL) and the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, binhu7@illinois.edu

³Yang Zheng is with the Department of Electrical and Computer Engineering, University of California San Diego, zhengy@engr.ucsd.edu

are much fewer and far less complete [14]–[18]. It remains unclear whether model-free policy gradient methods can be modified to yield global convergence guarantees. It has been revealed that the set of stabilizing static output-feedback controllers can be highly disconnected [14]. This is quite different from the state feedback case [19]. Such a negative result indicates that the performance of gradient-based policy search on static output feedback control highly depends on the initialization, and only convergence to stationary points has been established [15]. It is thus natural to investigate dynamical controllers for the output feedback case, and to see whether the corresponding optimization landscape is more favorable for direct policy search methods. The very recent work [16] shows that the set of stabilizing full-order dynamical controllers has at most two path-connected components that are identical in the frequency domain. This brings some positive news and opens the possibility of developing global convergent policy search methods for dynamical output feedback problems, such as linear quadratic Gaussian (LQG) control [16]. Two other recent studies are [17], [18]. In [18], the global convergence of policy search over dynamical filters was proved for a simpler estimation problem.

It is well-known that the optimal LQG controller has no robustness guarantee [20]. It is thus important to explicitly incorporate robustness constraints for the search of dynamical controllers. In this paper, we study the topological properties of the feasible set for linear dynamical output feedback control with \mathcal{H}_∞ robustness constraints. The \mathcal{H}_∞ constraints have been widely used in robust control [12], [21] and risk-sensitive control [22]. Our main result shows that the set of all stabilizing full-order dynamical controllers satisfying an additional input-output \mathcal{H}_∞ constraint has at most two path-connected components, and they are diffeomorphic under a mapping defined by a similarity transformation. Our proof technique is inspired by [16] and relies on a non-trivial but known change of variables for \mathcal{H}_∞ control [23], [24]. If the control cost is invariant under similarity transformation, one can initialize the local policy search anywhere within the feasible set and there is always a continuous path connecting the initial point to a global minimum. Some implications on the connectivity of strict sublevel sets of \mathcal{H}_∞ and \mathcal{H}_2 control are discussed. Our result sheds new light on model-free policy search for robust control tasks.

The rest of this paper is organized as follows. In Section II, we formulate the linear dynamic output feedback control with \mathcal{H}_∞ constraints as a constrained policy optimization problem. Section III presents our main theoretical results. We revisit connectivity of strict sublevel sets for LQG and \mathcal{H}_∞ control

in Section IV. We conclude the paper in Section V.

Notations: The set of $k \times k$ real symmetric matrices is denoted by \mathbb{S}^k , and the determinant of a square matrix M is denoted by $\det M$. 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 dimensions if they are clear from the context. Given a matrix $M \in \mathbb{R}^{k_1 \times k_2}$, M^\top denotes the transpose of M . For any $M_1, M_2 \in \mathbb{S}^k$, we use $M_1 \prec M_2$ ($M_1 \preceq M_2$) and $M_2 \succ M_1$ ($M_2 \succeq M_1$) to mean that $M_2 - M_1$ is positive (semi)definite.

II. PRELIMINARIES AND PROBLEM STATEMENT

A. Dynamic output feedback with \mathcal{H}_∞ constraints

We consider a continuous-time linear dynamical system¹

$$\begin{aligned}\dot{x}(t) &= Ax(t) + B_1w(t) + B_2u(t), \\ z(t) &= C_1x(t) + D_{11}w(t) + D_{12}u(t), \\ y(t) &= C_2x(t) + D_{21}w(t),\end{aligned}\tag{1}$$

where $x(t) \in \mathbb{R}^{n_x}$ is the state, $u(t) \in \mathbb{R}^{n_u}$ is the control action, $w(t) \in \mathbb{R}^{n_w}$ is the exogenous disturbance, $y(t) \in \mathbb{R}^{n_y}$ is the measured output, and $z(t) \in \mathbb{R}^{n_w}$ is the regulated performance output. We make the following assumption.

Assumption 1: The state-space model (A, B_2, C_2) in (1) is stabilizable and detectable.

We aim to design a controller that maps the measured output to the control action, in order to minimize some control performance metric, while satisfying stability and/or robustness constraints. Such control design problems can be formulated as a constrained policy optimization of the form

$$\min_{K \in \mathcal{K}} J(K),\tag{2}$$

where the decision variable K is determined by the policy parameterization, the objective function $J(K)$ measures the closed-loop performance, and the feasible set \mathcal{K} is specified by some stability/robustness requirements. We consider the following policy parameterization and robustness constraint:

Decision variable K : Output feedback control problems typically require dynamical controllers, and we consider the full-order dynamical controller in the form of:

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

where $\xi(t)$ is the controller state with the same dimension as $x(t)$, and matrices (A_K, B_K, C_K, D_K) specify the controller dynamics. For convenience, we denote

$$K := \begin{bmatrix} D_K & C_K \\ B_K & A_K \end{bmatrix} \in \mathbb{R}^{(n_u+n_x) \times (n_y+n_x)},\tag{4}$$

but this matrix K should be interpreted as the dynamical controller in (3).

Feasible region: The controller K needs to stabilize the closed-loop system and satisfy a robustness constraint that enforces the \mathcal{H}_∞ norm of the transfer function from $w(t)$ to $z(t)$ smaller than a pre-specified level γ .

We allow a general cost function $J(K)$, which can be an \mathcal{H}_2 performance on some other performance channel, or more general user-specified performance metrics. One advantage for the policy optimization formulation (2) is that it opens the possibility of solving robust control design via model-free policy search methods. This paper aims to characterize connectivity of \mathcal{K} and strict sublevel sets of $J(K)$.

B. Problem statement

We denote the state of the closed-loop system as $\zeta = [x^\top \ \xi^\top]^\top$ after combining (3) with (1). It is not difficult to derive the closed-loop system

$$\begin{aligned}\dot{\zeta}(t) &= A_{cl}\zeta(t) + B_{cl}w(t), \\ z(t) &= C_{cl}\zeta(t) + D_{cl}w(t),\end{aligned}\tag{5}$$

where the matrices $(A_{cl}, B_{cl}, C_{cl}, D_{cl})$ are given by

$$\begin{aligned}A_{cl} &:= \begin{bmatrix} A + B_2D_KC_2 & B_2C_K \\ B_KC_2 & A_K \end{bmatrix}, \\ B_{cl} &:= \begin{bmatrix} B_1 + B_2D_KD_{21} \\ B_KD_{21} \end{bmatrix}, \\ C_{cl} &:= [C_1 + D_{12}D_KC_2 \quad D_{12}C_K], \\ D_{cl} &:= D_{11} + D_{12}D_KD_{21}.\end{aligned}\tag{6}$$

The closed-loop system is internally stable if and only if A_{cl} is Hurwitz [12]. The set of full-order stabilizing dynamical controllers is thus defined as

$$\mathcal{C}_{stab} := \left\{ K \in \mathbb{R}^{(n_u+n_x) \times (n_y+n_x)} \mid A_{cl} \text{ is Hurwitz} \right\}.\tag{7}$$

The transfer function from $w(t)$ to $z(t)$ is

$$\mathbf{T}_{zw}(s) = C_{cl}(sI - A_{cl})^{-1}B_{cl} + D_{cl}.\tag{8}$$

Then, the feasible set is formally specified as

$$\mathcal{K}_\gamma := \{ K \in \mathcal{C}_{stab} \mid \|\mathbf{T}_{zw}\|_\infty < \gamma \},\tag{9}$$

where $\|\mathbf{T}_{zw}\|_\infty$ denotes the \mathcal{H}_∞ norm of \mathbf{T}_{zw} , and can be calculated as $\|\mathbf{T}_{zw}\|_\infty := \sup_\omega \sigma_{\max}(\mathbf{T}_{zw}(j\omega))$, with $\sigma_{\max}(\cdot)$ denoting the maximum singular value. In (9), we explicitly highlight the robustness level γ via the subscript. Under Assumption 1, there exists a finite positive value

$$\gamma^* := \inf_K \|\mathbf{T}_{zw}\|_\infty.$$

Then, \mathcal{K}_γ is non-empty if and only if $\gamma > \gamma^*$. By definition (9), we have $\mathcal{K}_{\gamma_0} \subset \mathcal{C}_{stab}$ for any positive γ_0 .

In (2), it is possible to estimate the gradient of $J(K)$ and $\|\mathbf{T}_{zw}\|_\infty$ from sampled system trajectories, and one may apply model-free gradient-based barrier algorithms to find a solution in an iterative fashion. To understand the performance of such model-free policy search algorithms, we need to characterize the optimization landscape of (2). In particular, we focus on some geometrical properties of the feasible region \mathcal{K}_γ and strict sublevel sets of $J(K)$. It is well-known that \mathcal{K}_γ is in general non-convex, but little is known about their other geometrical properties. Only a very recent work shows that \mathcal{C}_{stab} has at most two path-connected components that are identical up to similarity transformations [16, Theorems 3.1 & 3.2].

¹All topological results can be extended to the discrete-time domain.

In many cases, it is desirable to explicitly encode some robustness guarantee for the feasible region [20]–[22]. However, the connectivity of the \mathcal{H}_∞ -constrained set \mathcal{K}_γ remains unknown. In this paper, we focus on topological properties of \mathcal{K}_γ and their implications to gradient-based policy search. We will show that \mathcal{K}_γ shares similar properties with $\mathcal{C}_{\text{stab}}$.

Remark 1: The dynamical controller (3) is proper. Depending on the cost function $J(K)$ (e.g., LQG [12]), we may want to confine the policy space to strictly proper dynamical controllers. Then the feasible set is defined as

$$\tilde{\mathcal{K}}_\gamma := \{K \in \mathcal{K}_\gamma \mid D_K = 0\}. \quad (10)$$

Our analysis technique works for both $\tilde{\mathcal{K}}_\gamma$ and \mathcal{K}_γ , and we show that $\tilde{\mathcal{K}}_\gamma$ and \mathcal{K}_γ have similar topological properties. \square

III. PATH-CONNECTIVITY OF \mathcal{K}_γ

In this section, we present our main results on the topological properties of \mathcal{K}_γ . We first have a simple observation.

Lemma 1: Let $\gamma > \gamma^*$. The set \mathcal{K}_γ is non-empty, open, unbounded and can be non-convex.

This fact is well-known. Then openness of \mathcal{K}_γ follows from the continuity of the \mathcal{H}_∞ norm. It is unbounded since \mathcal{H}_∞ norm is invariant under similarity transformations that are unbounded in the state-space domain. The non-convexity is also known, and we illustrate it using the example below.

Example 1: Consider an open-loop unstable dynamical system (1) with $A = B_1 = B_2 = C_1 = C_2 = D_{21} = D_{12} = 1$, and $D_{11} = 0$. It is easy to verify that the following dynamical controllers $K^{(1)} = \begin{bmatrix} 0 & 2 \\ -2 & -2 \end{bmatrix}$, $K^{(2)} = \begin{bmatrix} 0 & -2 \\ 2 & -2 \end{bmatrix}$ satisfy $\|C_{\text{cl}}(sI - A_{\text{cl}})^{-1}B_{\text{cl}} + D_{\text{cl}}\|_\infty < 3.33$, and thus we have $K^{(1)} \in \mathcal{K}_{3.33}$, $K^{(2)} \in \mathcal{K}_{3.33}$. However, $\frac{1}{2}(K^{(1)} + K^{(2)}) = \begin{bmatrix} 0 & 0 \\ 0 & -2 \end{bmatrix}$, fails to stabilize the system, and is outside $\mathcal{K}_{3.33}$.

Despite the non-convexity, \mathcal{K}_γ has some nice connectivity property which will be established in this section.

A. Main results

Our first main technical result is stated as follows.

Theorem 1: Given any $\gamma > \gamma^*$, the set \mathcal{K}_γ has at most two path-connected components.

Before presenting a formal proof for Theorem 1, we first give some high-level ideas. Based on the bounded real lemma [21], we have $K \in \mathcal{K}_\gamma$ if and only if the matrix inequality,

$$\begin{bmatrix} A_{\text{cl}}^T P + P A_{\text{cl}} & P B_{\text{cl}} & C_{\text{cl}}^T \\ B_{\text{cl}}^T P & -\gamma I & D_{\text{cl}}^T \\ C_{\text{cl}} & D_{\text{cl}} & -\gamma I \end{bmatrix} \prec 0, \quad P \succ 0, \quad (11)$$

is feasible. Clearly, the condition (11) is not convex in K and P . Our result in Theorem 1 relies on the fact that (11) can be convexified into a linear matrix inequality (LMI) (that is convex and hence path-connected), using a non-trivial but known change of variables for \mathcal{H}_∞ control [23], [24]. The only potential of disconnectivity comes from the fact that the set of invertible matrices corresponding to similarity transformations has two path-connected components. Our proof is inspired by the recent work [16] that characterizes $\mathcal{C}_{\text{stab}}$ only, with

the main difference being that we need to analyze a more complicated \mathcal{H}_∞ constraint (11).

We now illustrate this idea for the case of state feedback (i.e. $y(t) = x(t)$ and $u(t) = Kx(t)$ with $K \in \mathbb{R}^{n_u \times n_x}$). In this case, it is known that (11) is feasible² if and only if

$$M_\gamma(Q, L) \prec 0, \quad Q \succ 0 \quad (12)$$

is feasible, where $M_\gamma(Q, L)$ is defined as

$$\begin{bmatrix} QA^T + AQ + L^T B_2^T + B_2 L & B_1 & (C_1 Q + D_{12} L)^T \\ B_1^T & -\gamma I & 0 \\ C_1 Q + D_{12} L & 0 & -\gamma I \end{bmatrix}.$$

Using a simple change of variables $K = LQ^{-1}$, we have

$$\begin{aligned} & \{K \in \mathbb{R}^{n_x \times n_u} \mid (11) \text{ is feasible}\} \\ \iff & \{K = LQ^{-1} \in \mathbb{R}^{n_x \times n_u} \mid (12) \text{ is satisfied}\}. \end{aligned}$$

Since the set of (Q, L) satisfying LMI (12) is convex and the map $K = LQ^{-1}$ is continuous, the set $\{K \in \mathbb{R}^{n_x \times n_u} \mid (11) \text{ is feasible}\}$ is path-connected.

The analysis above hinges upon the fact that in the state-feedback case, the non-convex condition (11) can be convexified using the simple change of variables $K = LQ^{-1}$. In the output feedback case, a similar condition can be derived using a more complicated change of variables in [24]. We will leverage this fact to prove Theorem 1. Specifically, it is known that a controller $K \in \mathcal{K}_\gamma$ can be constructed from the solution of the following LMI condition:

$$\begin{bmatrix} X & I \\ I & Y \end{bmatrix} \succ 0, \quad M_\gamma(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}) \prec 0, \quad (13)$$

where $X \in \mathbb{S}^{n_x}$, $Y \in \mathbb{S}^{n_x}$, $\hat{\mathbf{A}} \in \mathbb{R}^{n_x \times n_x}$, $\hat{\mathbf{B}} \in \mathbb{R}^{n_x \times n_y}$, $\hat{\mathbf{C}} \in \mathbb{R}^{n_y \times n_x}$, and $\hat{\mathbf{D}} \in \mathbb{R}^{n_u \times n_y}$, are decision variables. The linear mapping $M_\gamma(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}})$ is defined as

$$M_\gamma(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}) = \begin{bmatrix} M_{11} & M_{12} & M_{13} & M_{14} \\ M_{12}^T & M_{22} & M_{23} & M_{24} \\ M_{13}^T & M_{23}^T & M_{33} & M_{34} \\ M_{14}^T & M_{24}^T & M_{34}^T & M_{44} \end{bmatrix}, \quad (14)$$

where the blocks M_{ij} are given by

$$\begin{aligned} M_{11} &= AX + XA^T + B_2\hat{\mathbf{C}} + (B_2\hat{\mathbf{C}})^T, \\ M_{12} &= \hat{\mathbf{A}}^T + (A + B_2\hat{\mathbf{D}}C_2), \\ M_{13} &= B_1 + B_2\hat{\mathbf{D}}D_{21}, \\ M_{14} &= (C_1X + D_{12}\hat{\mathbf{C}})^T, \\ M_{22} &= A^T Y + YA + \hat{\mathbf{B}}C_2 + (\hat{\mathbf{B}}C_2)^T, \\ M_{23} &= YB_1 + \hat{\mathbf{B}}D_{21}, \\ M_{24} &= (C_1 + D_{12}\hat{\mathbf{D}}C_2)^T, \\ M_{33} &= -\gamma I, \\ M_{34} &= (D_{11} + D_{12}\hat{\mathbf{D}}D_{21})^T, \\ M_{44} &= -\gamma I. \end{aligned} \quad (15)$$

²In the state-feedback case, $(A_{\text{cl}}, B_{\text{cl}}, C_{\text{cl}}, D_{\text{cl}})$ should be calculated from some formulas which are different from (6). We omit the details.

Based on LMI (13), we introduce two useful sets:

$$\mathcal{F}_\gamma := \left\{ (X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}) \mid (13) \text{ is satisfied} \right\}, \quad (16)$$

$$\begin{aligned} \mathcal{G}_\gamma := \left\{ (X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi) \mid \Pi, \Xi \in \mathbb{R}^{n_x \times n_x}, \right. \\ \left. (X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}) \in \mathcal{F}_\gamma, \Xi \Pi = I - YX \right\}. \end{aligned} \quad (17)$$

It is obvious that \mathcal{F}_γ is convex and hence path-connected. Together with the fact that the set of $n_x \times n_x$ invertible matrices has two path-connected components, this guarantees that \mathcal{G}_γ has exactly two path-connected components. We shall see that there exists a continuous surjective map from \mathcal{G}_γ to \mathcal{K}_γ , and thus \mathcal{K}_γ has at most two path-connected components. A detailed proof is provided in the appendix.

Remark 2: Our analysis relies on the LMI conditions (11) and (13) from [24] which are specialized for \mathcal{H}_∞ control and are more complicated than the ones in the proof of [16, Proposition 3.1] for characterizing stability.

B. Implications for \mathcal{H}_∞ -constrained policy optimization

To understand the implications of Theorem 1 for policy optimization, we need to formalize the relationship between \mathcal{K}_γ^+ and \mathcal{K}_γ^- . For this, we introduce the notion of similarity transformation that is widely used in control. For any $T \in \text{GL}_{n_x}$, let $\mathcal{T}_T : \mathcal{C}_{\text{stab}} \rightarrow \mathcal{C}_{\text{stab}}$ denote the mapping given by

$$\mathcal{T}_T(\mathbf{K}) := \begin{bmatrix} D_{\mathbf{K}} & C_{\mathbf{K}} T^{-1} \\ T B_{\mathbf{K}} & T A_{\mathbf{K}} T^{-1} \end{bmatrix},$$

which represents similarity transformations on $\mathcal{C}_{\text{stab}}$.

We have a result that is similar to [16, Theorem 3.2].

Theorem 2: If \mathcal{K}_γ has two path-connected components \mathcal{K}_γ^+ and \mathcal{K}_γ^- , then \mathcal{K}_γ^+ and \mathcal{K}_γ^- are diffeomorphic under the mapping \mathcal{T}_T , for any $T \in \text{GL}_{n_x}$ with $\det T < 0$.

The proof of Theorem 2 is adapted from [16]. We present the details in our arXiv report [25] due to page limit.

Furthermore, similar to [16, Theorem 3.3], we have sufficient conditions to certify the path-connectedness of \mathcal{K}_γ .

Theorem 3: Let $\gamma > \gamma^*$. The following statements hold.

- 1) \mathcal{K}_γ is path-connected if it has one dynamical controller with non-minimal state-space description.
- 2) Suppose the plant (1) is single-input or single-output, i.e., $m = 1$ or $p = 1$. The set \mathcal{K}_γ is path-connected if and only if it has a non-minimal dynamical controller.

Proof: If $\mathbf{K} \in \mathcal{K}_\gamma$ is non-minimal, then its minimal realization has dimension less than n_x . In particular, we can find a reduced-order controller $(\tilde{A}_\mathbf{K}, \tilde{B}_\mathbf{K}, \tilde{C}_\mathbf{K}, D_\mathbf{K})$ with dimension $(n_x - 1)$ such that

$$\tilde{C}_\mathbf{K}(sI - \tilde{A}_\mathbf{K})^{-1}\tilde{B}_\mathbf{K} + D_\mathbf{K} = C_\mathbf{K}(sI - A_\mathbf{K})^{-1}B_\mathbf{K} + D_\mathbf{K}.$$

Then, this reduced-order controller can be augmented to be a full-order controller in \mathcal{K}_γ as

$$\tilde{\mathbf{K}} = \begin{bmatrix} \tilde{D}_\mathbf{K} & \tilde{C}_\mathbf{K} & 0 \\ \tilde{B}_\mathbf{K} & \tilde{A}_\mathbf{K} & 0 \\ 0 & 0 & -1 \end{bmatrix} \in \mathcal{K}_\gamma.$$

Define a matrix $T = \text{diag}(I_{n_x-1}, -1)$. We can directly verify $\det T < 0$ and $\mathcal{T}_T(\mathbf{K}) = \mathbf{K}$. By Theorem 2, we can see that $\mathbf{K} \in \mathcal{K}_\gamma^\pm$ implies $\mathcal{T}_T(\mathbf{K}) \in \mathcal{K}_\gamma^\mp$, indicating $\mathbf{K} \in \mathcal{K}_\gamma^+ \cap \mathcal{K}_\gamma^-$. Thus, $\mathcal{K}_\gamma^+ \cap \mathcal{K}_\gamma^-$ is nonempty, and \mathcal{K}_γ is path-connected.

The proof for the second statement is identical to the proof of [16, Theorem 3.3], and hence is omitted here. \blacksquare

Theorems 2 and 3 bring positive news on local policy search methods for \mathcal{H}_∞ -constrained optimization (2). If \mathcal{K}_γ is path-connected, it makes sense to initialize the policy search from any point in the feasible set. If \mathcal{K}_γ has two path-connected components, then the initial point may fall into either of the components. If $J(\mathbf{K})$ is invariant with respect to similarity transformations (e.g. the LQG cost), then both components include global minima. It becomes reasonable to initialize the policy search within either path-connected component. The following corollary is immediate.

Corollary 1: Suppose the cost function $J(\mathbf{K})$ is invariant with respect to similarity transformations, then there exists a continuous path connecting any feasible point $\mathbf{K} \in \mathcal{K}_\gamma$ to a global minimum of (2) if it exists.

C. The case of strictly proper controllers

We briefly discuss the case of strictly proper dynamical controllers with $D_\mathbf{K} = 0$, which is required in some classical control problems, including the continuous-time LQG problem [12]. The topological properties of $\tilde{\mathcal{K}}_\gamma$ in (9) and \mathcal{K}_γ in (10) are identical. To see this, we let

$$\begin{aligned} \tilde{\mathcal{F}}_\gamma &= \{(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}) \in \mathcal{F}_\gamma \mid \hat{\mathbf{D}} = 0\}, \\ \tilde{\mathcal{G}}_\gamma &= \{(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi) \in \mathcal{G}_\gamma \mid \hat{\mathbf{D}} = 0\}. \end{aligned}$$

Minor modification of the proofs in the appendix can show that $\tilde{\mathcal{F}}_\gamma$ is path-connected, and that $\tilde{\mathcal{G}}_\gamma$ has two path-connected components. The same mapping Φ in (21) is a continuous and surjective mapping from $\tilde{\mathcal{G}}_\gamma$ to $\tilde{\mathcal{K}}_\gamma$. Therefore, we conclude that $\tilde{\mathcal{K}}_\gamma$ has at most two path-connected components and they are diffeomorphic under the similarity transformation with $\det T < 0$.

IV. REVISIT SUBLVEL SETS IN LQG AND \mathcal{H}_∞ CONTROL

The results in Section III can be also interpreted as the connectivity of strict sublevel sets in optimal \mathcal{H}_∞ control. Based on (8), \mathbf{T}_{zw} can be viewed as a function of \mathbf{K} , and the optimal \mathcal{H}_∞ synthesis [12] can be formulated as

$$\begin{aligned} \min_{\mathbf{K}} \quad & \|\mathbf{T}_{zw}\|_\infty \\ \text{subject to} \quad & \mathbf{K} \in \mathcal{C}_{\text{stab}}. \end{aligned} \quad (18)$$

Now, \mathcal{K}_γ in (9) is exactly the γ -level strict sublevel set of the optimal \mathcal{H}_∞ control (18). Thus, Theorems 1 to 3 characterize the strict sub-level sets of optimal \mathcal{H}_∞ control.

In addition to (18), the proof idea of using the change of variables (21) can be applied to other output feedback control problems to establish connectivity of their strict sublevel sets. For example, we can consider an \mathcal{H}_2 formulation of the LQG control [16] as follows

$$\begin{aligned} \min_{\mathbf{K}} \quad & \|\mathbf{T}_{zw}\|_2^2 \\ \text{subject to} \quad & \mathbf{K} \in \mathcal{C}_{\text{stab}} \cap \{\mathbf{K} \mid D_\mathbf{K} = 0\}, \end{aligned} \quad (19)$$

where $\|\mathbf{T}_{zw}\|_2$ denotes the \mathcal{H}_2 norm of \mathbf{T}_{zw} . This problem (19) covers the LQG control as a special case when the dynamics in (1) are chosen appropriately (this fact is well-known; see [26] for early discussions). The same proof techniques in Section III can establish the connectivity of the strict sublevel sets of (19):

$$\mathcal{L}_\gamma = \{\mathbf{K} \in \mathcal{C}_{\text{stab}} \mid D_{\mathbf{K}} = 0, \|\mathbf{T}_{zw}\|_2^2 < \gamma\}. \quad (20)$$

We have the following result (see our report [25] for details).

Theorem 4: Under [16, Assumption 1], the strict sublevel set \mathcal{L}_γ (20) has at most two path-connected components $\mathcal{L}_\gamma^{(1)}$ and $\mathcal{L}_\gamma^{(2)}$. If so, $\mathcal{L}_\gamma^{(1)}$ and $\mathcal{L}_\gamma^{(2)}$ are diffeomorphic under the mapping \mathcal{T}_T , for any $T \in \text{GL}_{n_x}$ with $\det T < 0$.

A straightforward implication from Theorem 4 is that there exists a continuous path connecting any feasible point $\mathbf{K} \in \mathcal{L}_\gamma$ to a global minimum of LQG control. Moreover, path connectivity of sublevel sets may imply further landscape properties [27], [28]. For example, using a special definition of minimizing sets in [28, Definition 5.1], Theorem 5.4 in [28] guarantees that \mathcal{H}_∞ control (18) and LQG control (19) have a unique global minimizing set in some weak sense.

Definition 1: A nonempty S is an LT-critical set³ (or LTCS) for the function $\|\mathbf{T}_{zw}\|_2^2$ if 1) $\|\mathbf{T}_{zw}\|_2^2$ is constant, $\forall \mathbf{K} \in S$, and 2) for any $\gamma' > \gamma$ with γ being the value of $\|\mathbf{T}_{zw}\|_2^2$ over S , the strict sublevel set $\mathcal{L}_{\gamma'}$ has a single connected component containing S , and the intersection of all such single connected components with $\gamma' > \gamma$ equals to S .

Definition 2: A LTCS is called a global LT-minimizing set (or global LTMS) if the value of $\|\mathbf{T}_{zw}\|_2^2$ over this set is no greater than the values of $\|\mathbf{T}_{zw}\|_2^2$ for all $\mathbf{K} \in \mathcal{C}_{\text{stab}}$.

Corollary 2: If the global \mathcal{H}_2 optimal controller exists, then the cost function $\|\mathbf{T}_{zw}\|_2^2$ (as a function of \mathbf{K}) has a unique global LTMS, and no other LT-critical sets.

This result is a direct consequence of [28, Theorem 5.4]. A similar result holds for \mathcal{H}_∞ control. Corollary 2 ensures a unique minimizing set only in a weak sense, and does not rule out the normal notion of local minima and saddle points. For (19), saddle points actually exist [16]. A rigorous definition of strict local minima for (18) or (19) requires extra work due to unboundedness of similarity transformations. Discussions on local optimality conditions appear in [29].

V. CONCLUSIONS

We have proved that the set of \mathcal{H}_∞ -constrained full-order dynamical controllers has at most two path-connected components (cf. Theorem 1) and they are diffeomorphic under similarity transformations (cf. Theorem 2). We have also discussed various implications on the strict sublevel sets of LQG and \mathcal{H}_∞ control (cf. Theorem 4). This brings positive news for direct policy search of robust controllers.

REFERENCES

- [1] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi, “Global convergence of policy gradient methods for the linear quadratic regulator,” in *International Conference on Machine Learning*, vol. 80, 2018, pp. 1467–1476.
- [2] D. Malik, A. Pananjady, K. Bhatia, K. Khamaru, P. Bartlett, and M. Wainwright, “Derivative-free methods for policy optimization: Guarantees for linear quadratic systems,” in *International Conference on Artificial Intelligence and Statistics*, 2019, pp. 2916–2925.
- [3] L. Furieri, Y. Zheng, and M. Kamgarpour, “Learning the globally optimal distributed LQ regulator,” in *Learning for Dynamics and Control*, 2020, pp. 287–297.
- [4] H. Mohammadi, A. Zare, M. Soltanolkotabi, and M. R. Jovanovic, “Convergence and sample complexity of gradient methods for the model-free linear quadratic regulator problem,” *IEEE Transactions on Automatic Control*, 2021.
- [5] K. Zhang, B. Hu, and T. Başar, “Policy optimization for \mathcal{H}_2 linear control with \mathcal{H}_∞ robustness guarantee: Implicit regularization and global convergence,” *SIAM Journal on Control and Optimization*, vol. 59, no. 6, pp. 4081–4109, 2021.
- [6] K. Zhang, B. Hu, and T. Başar, “On the stability and convergence of robust adversarial reinforcement learning: A case study on linear quadratic systems,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [7] B. Gravell, P. M. Esfahani, and T. Summers, “Learning optimal controllers for linear systems with multiplicative noise via policy gradient,” *IEEE Transactions on Automatic Control*, vol. 66, no. 11, pp. 5283–5298, 2020.
- [8] K. Zhang, X. Zhang, B. Hu, and T. Başar, “Derivative-free policy optimization for linear risk-sensitive and robust control design: Implicit regularization and sample complexity,” in *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [9] J. P. Jansch-Porto, B. Hu, and G. E. Dullerud, “Convergence guarantees of policy optimization methods for Markovian jump linear systems,” in *American Control Conference*, 2020, pp. 2882–2887.
- [10] S. Rathod, M. Bhadu, and A. De, “Global convergence using policy gradient methods for model-free markovian jump linear quadratic control,” *arXiv preprint arXiv:2111.15228*, 2021.
- [11] J. P. Jansch-Porto, B. Hu, and G. Dullerud, “Policy optimization for markovian jump linear quadratic control: Gradient-based methods and global convergence,” *arXiv preprint arXiv:2011.11852*, 2020.
- [12] K. Zhou, J. C. Doyle, and K. Glover, *Robust and optimal control*. Prentice Hall, 1996.
- [13] Y. Sun and M. Fazel, “Learning optimal controllers by policy gradient: Global optimality via convex parameterization,” in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 4576–4581.
- [14] H. Feng and J. Lavaei, “On the exponential number of connected components for the feasible set of optimal decentralized control problems,” in *2019 American Control Conference (ACC)*, 2019, pp. 1430–1437.
- [15] I. Fatkhullin and B. Polyak, “Optimizing static linear feedback: Gradient method,” *SIAM Journal on Control and Optimization*, vol. 59, no. 5, pp. 3887–3911, 2021.
- [16] Y. Zheng, Y. Tang, and N. Li, “Analysis of the optimization landscape of linear quadratic gaussian (LQG) control,” *arXiv preprint arXiv:2102.04393*, 2021.
- [17] J. Duan, W. Cao, Y. Zheng, and L. Zhao, “On the optimization landscape of dynamical output feedback linear quadratic control,” *arXiv preprint arXiv:2201.09598*, 2022.
- [18] J. Umenberger, M. Simchowitz, J. C. Perdomo, K. Zhang, and R. Tedrake, “Globally convergent policy search over dynamic filters for output estimation,” *arXiv preprint arXiv:2202.11659*, 2022.
- [19] J. Bu, A. Mesbahi, and M. Mesbahi, “On topological and metrical properties of stabilizing feedback gains: the MIMO case,” *arXiv preprint arXiv:1904.02737*, 2019.
- [20] J. C. Doyle, “Guaranteed margins for LQG regulators,” *IEEE Transactions on Automatic Control*, vol. 23, no. 4, pp. 756–757, 1978.
- [21] G. Dullerud and F. Paganini, *A Course in Robust Control Theory: A Convex Approach*. Springer, 1999.
- [22] P. Whittle, *Risk-sensitive Optimal Control*. Wiley Chichester, 1990.
- [23] P. Gahinet and P. Apkarian, “A linear matrix inequality approach to \mathcal{H}_∞ control,” *International journal of robust and nonlinear control*, vol. 4, no. 4, pp. 421–448, 1994.
- [24] C. Scherer, P. Gahinet, and M. Chilali, “Multiobjective output-feedback control via LMI optimization,” *IEEE Transactions on Automatic Control*, vol. 42, no. 7, pp. 896–911, 1997.
- [25] B. Hu and Y. Zheng, “Connectivity of the feasible and sublevel sets of dynamic output feedback control with robustness constraints,” *arXiv preprint arXiv:2203.11177*, 2022.
- [26] V. Kucera, “The LQG and \mathcal{H}_2 design: Two different problems?” in *European Control Conference*, 1993, pp. 334–337.
- [27] J. M. Ortega and W. C. Rheinboldt, *Iterative solution of nonlinear equations in several variables*. SIAM, 2000.

³This terminology is adopted from [28], and “LT” stands for “less than.”

- [28] D. H. Martin, "Connected level sets, minimizing sets, and uniqueness in optimization," *Journal of Optimization Theory and Applications*, vol. 36, no. 1, pp. 71–91, 1982.
- [29] C. Hoi, C. W. Scherer, E. Van der Meché, and O. Bosgra, "A nonlinear SDP approach to fixed-order controller synthesis and comparison with two other methods applied to an active suspension system," *European Journal of Control*, vol. 9, no. 1, pp. 13–28, 2003.
- [30] J. M. Lee, *Introduction to Smooth Manifolds*, 2nd ed. Springer Science & Business Media, 2013.

APPENDIX

This appendix presents the proof of Theorem 1.

Lemma 2: For any $(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi) \in \mathcal{G}_\gamma$, Π and Ξ are always invertible, and consequently, the block triangular matrices $\begin{bmatrix} I & 0 \\ YB_2 & \Xi \end{bmatrix}$ and $\begin{bmatrix} I & C_2X \\ 0 & \Pi \end{bmatrix}$ are invertible.

The proof is straightforward by observing that $\det(\Pi\Xi) = \det(YX - I) \neq 0$ for any $(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi) \in \mathcal{G}_\gamma$. Based on the change of variables in [24], we can map each element of \mathcal{G}_γ back to a controller $K \in \mathbb{R}^{(n_u+n_x) \times (n_y+n_x)}$. For each $Z = (X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi)$ in \mathcal{G}_γ , we define

$$\begin{aligned} \Phi(Z) &= \begin{bmatrix} \Phi_D(Z) & \Phi_C(Z) \\ \Phi_B(Z) & \Phi_A(Z) \end{bmatrix} \\ &= \begin{bmatrix} I & 0 \\ YB_2 & \Xi \end{bmatrix}^{-1} \begin{bmatrix} \hat{\mathbf{D}} & \hat{\mathbf{C}} \\ \hat{\mathbf{B}} \hat{\mathbf{A}} - YAX & \end{bmatrix} \begin{bmatrix} I & C_2X \\ 0 & \Pi \end{bmatrix}^{-1}. \end{aligned} \quad (21)$$

The nonlinear change of variables in (21) is from [24], which allows us to establish the following essential result. This result can be derived from [24], and we provide a proof for completeness.

Proposition 1: The mapping Φ in (21) is a continuous and surjective mapping from \mathcal{G}_γ to \mathcal{K}_γ .

Proof: It is clear that $\Phi(\cdot)$ is a continuous mapping. To show that Φ is a mapping onto \mathcal{K}_γ , we need to prove the following statements:

- 1) For any arbitrary controller $K \in \mathcal{K}_\gamma$, there exists $Z = (X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi) \in \mathcal{G}_\gamma$ such that $\Phi(Z) = K$.
- 2) For all $Z = (X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi) \in \mathcal{G}_\gamma$, we have $\Phi(Z) \in \mathcal{K}_\gamma$.

To show the first statement, let $K = \begin{bmatrix} D_K & C_K \\ B_K & A_K \end{bmatrix} \in \mathcal{K}_\gamma$ be arbitrary. By the bounded real lemma [21], there exists $P \succ 0$ such that (11) is feasible. We partition the matrix P as

$$P = \begin{bmatrix} Y & \Xi \\ \Xi^\top & \hat{Y} \end{bmatrix}. \quad (22)$$

Without loss of generality, we assume that $\det \Xi \neq 0$ (otherwise we can add a small perturbation on Ξ thanks to the strict inequality in (11)). We further define

$$\begin{bmatrix} X & \Pi^\top \\ \Pi & \hat{X} \end{bmatrix} := \begin{bmatrix} Y & \Xi \\ \Xi^\top & \hat{Y} \end{bmatrix}^{-1}, \quad T := \begin{bmatrix} X & I \\ \Pi & 0 \end{bmatrix}. \quad (23)$$

we can verify that

$$YX + \Xi\Pi = I, \quad T^\top PT = \begin{bmatrix} X & I \\ I & Y \end{bmatrix} \succ 0. \quad (24)$$

Now we choose $(\hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}})$ as

$$\begin{aligned} \hat{\mathbf{A}} &= Y(A + B_2 D_K C_2)X + \Xi B_K C_2 X \\ &\quad + Y B_2 C_K \Pi + \Xi A_K \Pi, \\ \hat{\mathbf{B}} &= Y B_2 D_K + \Xi B_K, \\ \hat{\mathbf{C}} &= D_K C_2 X + C_K \Pi, \quad \hat{\mathbf{D}} = D_K. \end{aligned} \quad (25)$$

We can verify that $M_\gamma(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}})$ is the same as

$$\begin{bmatrix} T^\top & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{bmatrix} \begin{bmatrix} A_{cl}^\top P + PA_{cl} & PB_{cl} & C_{cl}^\top \\ B_{cl}^\top P & -\gamma I & D_{cl}^\top \\ C_{cl} & D_{cl} & -\gamma I \end{bmatrix} \begin{bmatrix} T & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{bmatrix},$$

which is clearly negative definite due to (11). Thus, we have $Z = (X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi) \in \mathcal{G}_\gamma$ by the definition of \mathcal{G}_γ . Note that (25) can be compactly rewritten as

$$\begin{bmatrix} \hat{\mathbf{D}} & \hat{\mathbf{C}} \\ \hat{\mathbf{B}} & \hat{\mathbf{A}} - YAX \end{bmatrix} = \begin{bmatrix} I & 0 \\ YB_2 & \Xi \end{bmatrix} \begin{bmatrix} D_K & C_K \\ B_K & A_K \end{bmatrix} \begin{bmatrix} I & C_2X \\ 0 & \Pi \end{bmatrix}.$$

Based on Lemma 2, we have

$$\begin{bmatrix} D_K & C_K \\ B_K & A_K \end{bmatrix} = \begin{bmatrix} \Phi_D(Z) & \Phi_C(Z) \\ \Phi_B(Z) & \Phi_A(Z) \end{bmatrix} = \Phi(Z).$$

Therefore, the first statement is true. The second statement reduces to the standard controller construction for LMI-based \mathcal{H}_∞ -synthesis [24]. We complete the proof. ■

Remark 3: The analysis technique via the change of variables (21) in Proposition 1 is from [24]. This analysis can also be used for \mathcal{H}_2 and other costs; see [24] for details.

Based on Proposition 1, any path-connected component of \mathcal{G}_γ has a path-connected image under the surjective mapping Φ . Consequently, the number of path-connected components of \mathcal{K}_γ will be no more than the number of path-connected components of \mathcal{G}_γ . The number of path-connected components of the set \mathcal{G}_γ is given below.

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

$$\mathcal{G}_\gamma^+ = \left\{ (X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi) \in \mathcal{G}_\gamma \mid \det \Pi > 0 \right\},$$

$$\mathcal{G}_\gamma^- = \left\{ (X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, \Xi) \in \mathcal{G}_\gamma \mid \det \Pi < 0 \right\}.$$

Proof: First, \mathcal{F}_γ is path-connected since it is convex. The set of real invertible matrices $GL_{n_x} = \{\Pi \in \mathbb{R}^{n_x \times n_x} \mid \det \Pi \neq 0\}$ has two path-connected components [30]

$$GL_{n_x}^+ = \{\Pi \in \mathbb{R}^{n_x \times n_x} \mid \det \Pi > 0\},$$

$$GL_{n_x}^- = \{\Pi \in \mathbb{R}^{n_x \times n_x} \mid \det \Pi < 0\}.$$

Thus, the Cartesian product $\mathcal{F}_\gamma \times GL_{n_x}$ has two path-connected components. We further observe that the mapping from $(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi)$ to $(X, Y, \hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}}, \Pi, (I - YX)\Pi^{-1})$ is a continuous bijection from $\mathcal{F}_\gamma \times GL_{n_x}$ to \mathcal{G}_γ . This immediately leads to the desired conclusion. ■

The proofs for Proposition 2 and [16, Proposition 3.2] are similar. Proposition 3.2 in [16] can be viewed as a special case of Proposition 2 with $\gamma \rightarrow +\infty$. Now Theorem 1 can be proved by combining Proposition 1 and Proposition 2.

Proof of Theorem 1: We define $\mathcal{K}_\gamma^+ := \Phi(\mathcal{G}_\gamma^+)$ and $\mathcal{K}_\gamma^- := \Phi(\mathcal{G}_\gamma^-)$. We have $\mathcal{K}_\gamma = \mathcal{K}_\gamma^+ \cup \mathcal{K}_\gamma^-$. If \mathcal{K}_γ is not path-connected, the two path-connected components of \mathcal{K}_γ are exactly \mathcal{K}_γ^+ and \mathcal{K}_γ^- . Based on Proposition 1, Theorem 1 holds.