

Policy Optimization over Submanifolds for Linearly Constrained Feedback Synthesis

Shahriar Talebi, *Member, IEEE* and Mehran Mesbahi, *Fellow, IEEE*

Abstract—In this paper, we study linearly constrained policy optimization over the manifold of Schur stabilizing controllers, equipped with a Riemannian metric that emerges naturally in the context of optimal control problems. We provide extrinsic analysis of a generic constrained smooth cost function, that subsequently facilitates subsuming any such constrained problem into this framework. By studying the second order geometry of this manifold, we provide a Newton-type algorithm that does not rely on the exponential mapping nor a retraction, while ensuring local convergence guarantees. The algorithm hinges instead upon the developed stability certificate and the linear structure of the constraints. We then apply our methodology to two well-known constrained optimal control problems. Finally, several numerical examples showcase the performance of the proposed algorithm.

Index Terms—Constrained stabilizing controllers, Optimization over submanifolds, Output-feedback LQR control, Structured LQR control

I. INTRODUCTION

IN recent years, direct Policy Optimization (PO) for different variants of the Linear Quadratic Regulators (LQR) problems have attracted considerable attention in the literature. In the meantime, PO for linearly constrained LQR, e.g., state-feedback Structured LQR (SLQR) and Output-feedback LQR (OLQR), has been less explored due to the intricate geometry of the respective feasible sets and the non-convexity of the cost function. While reparameterization of LQR to a convex setup is possible for unconstrained cases [1], in general, trivial constraints directly on the policy become nontrivial and non-convex after such reparameterizations.¹ Furthermore, the domain of the optimization problems for constrained LQR (and its variants) are generally not only non-convex [3] but also disconnected [4]. As such, there are no guarantees that first-order stationary points are necessarily local minima.

Finding the linear output-feedback policy directly for the OLQR problem was first addressed in [5], a procedure

This work was supported by the AFOSR grant FA9550-20-1-0053 and NSF grant ECCS-2149470. (Corresponding author: Shahriar Talebi.)

The authors are with the William E. Boeing Department of Aeronautics and Astronautics, University of Washington, Seattle, WA 98115 USA. S. Talebi is also with the Department of Mathematics at the University of Washington (e-mail: shahriar@uw.edu; mesbahi@uw.edu).

¹There are exceptions to this statement-like when conditions such as quadratic invariance can be invoked [2].

that involves solving nonlinear matrix equations at each iteration. Since then, there has been on-going research efforts to address this problem adopting distinct perspectives [6]–[13], including its computational complexity [14]–[16]. In this direction, first and second order methods have been adopted for solving SLQR and OLQR problems (see e.g., [9], [10] and references therein). However, these methods often have a number of limitations, including reliance on backtracking line-search techniques at each iteration (which may be computationally expensive or infeasible), absence of convergence guarantees, not utilizing the inherent non-Euclidean geometry of the problem, and finally, not offering a setup for handling general linear constraints on the feedback gain.

Recently, state-feedback LQR problems have been studied through the lens of first order methods, in both discrete-time [17] and continuous-time [18] setups. This point of view was initiated when the LQR cost was shown to satisfy Polyak-Łojasiewicz (PL) (aka *gradient dominance*) property [19], facilitating a global convergence guarantee of first order methods for this problem—despite its non-convexity. Since then, PO using first order methods has been investigated for variants of LQR problem, such as OLQR [20], model-free setup [21], and risk-constrained LQR [22]. The gradient dominance property, however, is only known to be valid with respect to the global optimum of the unconstrained case, and is not necessarily expected for general constrained LQR problems. By merely using the first order information of the cost function, Projected Gradient Descent (PGD) techniques—whenever the projection is possible—can be shown to converge to first order stationary points but with a sublinear convergence rate (e.g., see [17] and [20] for SLQR and OLQR problems, respectively). A sublinear rate is generally unfavorable from a practical point of view, particularly when second order information of the LQR cost can be utilized. Despite computational challenges arising from the non-convexity—and even non-connectedness—of stabilizing feedback gains for general constrained optimal control problems, one may consider developing fast convergent algorithms for efficient identification of feasible local optima.

Note that *some* structure on the policy can be enforced through regularization; however, this approach merely *promotes* structural constraints and does not address problems considered here, as constraints are prescribed as a hard requirement for feasibility of the solution (e.g., see [23]




Fig. 1: The submanifold \tilde{S} of diagonal state-feedback stabilizing controllers ($K = \text{diag}(l_1, l_2)$) for a 2×2 system; superimposed with the level-sets of the constrained LQR cost (h), and the regions on which its *Hessian* is positive definite with respect to the inherent Riemannian geometry ($\text{Hess}h$) and the Euclidean geometry ($\overline{\text{Hess}}h$), respectively. Note that \tilde{S} is covered by the region on which $\text{Hess}h > 0$. See Example 1 for more details.

for an approach pertaining to promoting sparsity for SLQR problem). Here, we aim to utilize the second order information of the cost function to improve the convergence rate, and at the same time, provide a general approach that can address *other* linear constraints from a unifying perspective; in essence, the proposed approach can be adopted for problems such as SLQR and OLQR, well-recognizing their inherent computational complexity [14], [15].

By ignoring the geometry of the problem, one may aim to optimize the *linearly* constrained LQR cost by directly utilizing first or second order methods. Since the domain is non-convex, this might still be possible—depending on the problem scenario—by incorporating an Armijo-type backtracking line-search or requiring that the initial guess be close to the local optima. Here, we preclude from incorporating a line-search in order to systematically exploit the geometry inherent in the LQR cost. This inherent geometry will be precisely captured by the Riemannian metric defined subsequently in §III (see Lemma III.2). Incorporating a backtracking technique is then an immediate extension of our setup. Furthermore, by adopting a geometric perspective in designing direct policy optimization, we aim to also pave the way for future works on including relevant system theoretic criteria that lead to *nonlinear* constraints on the feedback policy; handling such constraints can significantly benefit from the intrinsic geometry of the problem as investigated in this work (see §VII for an example). Generally, the second order behavior of the cost function can be utilized in order to obtain a descent direction—as in the Newton method—as long as the *Hessian* stays positive definite. However, this second order information can be obtained in a variety of ways; for example, with respect to the usual Euclidean geometry (especially for linearly constrained problems) or more interestingly, with respect to the non-Euclidean geometry

inherent to the cost function itself. A simple—yet relevant—example of the above statement is depicted in Figure 1 in relation to the set of *diagonally* constrained stabilizing controllers (denoted by \tilde{S})—which turns out to be non-convex even for this simple example consisting of two inputs and two states. More specifically, this set is the intersection of the 4-dimensional set of stabilizing controllers with a 2-dimensional plane defining the diagonal constraint. In reference to Figure 1, note how the *Euclidean Hessian* of the constrained LQR cost (denoted by $\overline{\text{Hess}}h$) is positive definite on a smaller subset of \tilde{S} —especially in the vicinity of the “minimum.” Therefore, one would expect that the neighborhood of minimum on which the Newton updates using the Euclidean geometry converge, be relatively small. In the meantime, if the second order behavior of the LQR cost is considered through the lens of Riemannian geometry, its *Riemannian Hessian* (denoted by $\text{Hess}h$) captures the behavior of the cost function more effectively, in that it remains positive definite on a larger domain—compared with $\overline{\text{Hess}}h$. Hence, one expects a significant difference in the performance of second order optimization algorithms utilizing these two distinct geometries. To further expand on this example, one needs to define the (Riemannian) $\text{Hess}h$ and devise a second order algorithm that iteratively optimize this constrained problem in the vicinity of a local optimum. This machinery is developed in the rest of this paper, after which we revisit the above example (see Example 1 in §VI).

The literature on optimization over manifolds often relies on having access to either the exponential mapping [24], [25], or a *retraction*² from its tangent bundle onto the manifold itself [26]. However, due to the intricate geometry of the manifold of Schur stabilizing controllers, the exponential mapping is computationally expensive and a suitable retraction is generally not available. Furthermore, many useful constraints for optimal LQR problems (such as OLQR or SLQR) inherit a linear structure. Hence, one may consider the Natural Gradient Descent approach of [27], [28] in order to utilize this inherent geometry; however, this approach is not directly applicable to more involved submanifolds of stabilizing controllers. Nonetheless, it is pertinent to ask whether we can still exploit the intrinsic non-Euclidean geometry—induced by, say, the quadratic cost and linear dynamics, in optimal control problems as illustrated in Figure 1—by circumventing the absence of a computationally feasible retraction while guaranteeing stability.

In this paper, we consider a general optimization problem over the set of linearly constrained stabilizing feedback gains; this setup can easily be tailored to other classes of constrained control synthesis problems. We introduce a Newton-type algorithm that utilizes both the inherent Riemannian geometry as well as the linear structure of the constraints, and provide its convergence analysis to the local minima. Here, in the absence of a computationally feasible (global) retraction from the tangent bundle to the manifold,

²A better terminology would be *graph projection*; however, we adopt *retraction* to be consistent with the manifold optimization literature.

we obtain the so-called *stability certificate* that—together with the linear structure of the constraints—substitute the role that a retraction would generally play, ensuring the feasibility of the next iterate. Finally, as the unit stepsize for the proposed iterates may not be possible in general, we guarantee a linear convergence rate—that eventually becomes quadratic as the iterates converge. Finally, we provide applications of the proposed methodology to the well-known state-feedback SLQR and OLQR problems, followed by numerical examples.

Our contributions can thus be summarized as follows:

(i) We study the second order geometry of the manifold of stabilizing controllers induced by a pertinent Riemannian metric and its associated *connection* (a generalization of directional derivatives [29]), in order to obtain the second order information of a cost through defining the Riemannian Hessian. (ii) We provide extrinsic analysis for first and second-order behavior of a generic smooth cost function constrained to a *Riemannian submanifold*. This, in turn, allows for a general treatment of constrained optimization problems on the manifold of stabilizing controllers. (iii) We introduce Quasi Riemannian-Newton Policy Optimization (QRNPO) algorithm (pronounced as *Kern PO*) with convergence guarantees that exploits the inherent Riemannian geometry in the absence of the exponential mapping or a retraction, effectively providing a stability certificate for linearly constrained feedback gains. (iv) We apply our methodology to SLQR and OLQR problems by first, computing the second order behavior of the LQR cost with respect to the Riemannian connection, and then, explicating the solution to Newton equation for each case using this geometry. (v) While our approach allows for considering any choice of connection, here, we focus on the associated Riemannian connection—we also make a comparison to the ordinary Euclidean connection. (vi) Finally, we provide several numerical examples to showcase the performance and advantages of the proposed methodology that exploits the intrinsic geometry of constrained feedback stabilization. Further applications of the proposed algorithm—e.g. to the problem of controlling diffusion dynamics over a network—has recently appeared in [30].

The rest of the paper is organized as follows. In §II, we introduce the generic (stabilizing) feedback synthesis problem. We then provide the analysis of this problem through the lens of differential geometry in §III. In §IV, we present the algorithm and its convergence analysis for optimization on submanifolds of stabilizing controllers. Applications to SLQR and OLQR problems are then presented in §V. Finally, numerical examples are provided in §VI, followed by concluding remarks in §VII. The appendix contains proofs of the results.

Notation: The space of $m \times n$ matrices over the reals is denoted by $\mathbf{M}(m \times n, \mathbb{R})$ with the trivial smooth structure determined by the atlas consisting of the single chart $(\mathbf{M}(m \times n, \mathbb{R}), \text{vec})$, where $\text{vec}: \mathbf{M}(m \times n, \mathbb{R}) \mapsto \mathbb{R}^{mn}$ denotes the operator that returns a vector obtained by (vertically) stacking the columns of a matrix—from left to right.

We denote the transpose operator and the spectral norm of a matrix by $(\cdot)^\top$ and $\|\cdot\|_2$, respectively. The trace and spectral radius of a square matrix are denoted by $\text{tr}[\cdot]$ and $\rho(\cdot)$. The *Loewner* partial order of symmetric positive (semi-)definite matrices is denoted by \succ (\succeq); we use the same notation to denote positive (semi-)definiteness of 2-tensor fields. The maximum and minimum eigenvalues of symmetric matrices will be designated by $\bar{\lambda}$ and $\underline{\lambda}$, respectively. The set of positive integers less than or equal to m is denoted by $[m]$. By $\mathcal{M} := \{A \in \mathbf{M}(n \times n, \mathbb{R}) \mid \rho(A) < 1\}$, we denote the set of (Schur) stable matrices, and define the *Lyapunov map* $\mathbb{L}: \mathcal{M} \times \mathbf{M}(n \times n, \mathbb{R}) \mapsto \mathbf{M}(n \times n, \mathbb{R})$, that sends the pair (A, Z) to the unique solution X of

$$X = AXA^\top + Z, \quad (1)$$

which has the representation $X = \sum_{i=0}^{\infty} A^i Z (A^\top)^i$; in this case, if $Z \succeq 0$ ($\succ 0$), then $X \succeq 0$ ($\succ 0$). Furthermore, when $Z \succeq 0$, then $X \succ 0$ if and only if $(A, Z^{1/2})$ is controllable (see [31] and references therein). For manifolds we follow the notation and results in [29] and [32] unless stated explicitly.

II. PROBLEM STATEMENT

Given a stabilizable pair (A, B) with $A \in \mathbf{M}(n \times n, \mathbb{R})$ and $B \in \mathbf{M}(n \times m, \mathbb{R})$, we define

$$\mathcal{S} := \{K \in \mathbf{M}(m \times n, \mathbb{R}) \mid \rho(A + BK) < 1\},$$

as the set of stabilizing feedback gains. Subsequently, we will introduce a non-Euclidean geometry over \mathcal{S} using a metric arising naturally in the context of optimal control problems. We are often interested in feedback gains K that lie in a relatively *simple* subset \mathcal{K} of $\mathbf{M}(m \times n, \mathbb{R})$, such that $\tilde{\mathcal{S}} := \mathcal{K} \cap \mathcal{S}$ is an embedded submanifold of \mathcal{S} (see §III-A). A common example of this would be a linear subspace of $\mathbf{M}(m \times n, \mathbb{R})$ characterizing a prescribed sparsity pattern for the admissible controller gains (see §V-A.2). Another example is the optimal output-feedback synthesis considered in §V-A.3.

Herein, we are concerned with the optimization problem,

$$\min_K f(K) \quad \text{s.t.} \quad K \in \tilde{\mathcal{S}}, \quad (2)$$

where $f \in C^\infty(\mathcal{S}, \mathbb{R}) = C^\infty(\tilde{\mathcal{S}})$ and $\tilde{\mathcal{S}}$ is an embedded submanifold of \mathcal{S} , especially when it is endowed with a linear structure. This problem is motivated by parameterized feedback synthesis problems where we optimize the LQR cost f directly for *the policy* $K \in \tilde{\mathcal{S}}$ which takes the form of $f(K) = \frac{1}{2} \text{tr}[P_K \Sigma_K]$ with some $\Sigma_K, P_K \succ 0$ smoothly depending on K (see §V for further details).

Our approach involves using this linear structure with an appropriate Riemannian geometry of \mathcal{S} to circumvent the absence of a (computationally feasible) global retraction from $T\mathcal{S}$ onto \mathcal{S} (or $\tilde{\mathcal{S}}$)—due to the intricate geometry of \mathcal{S} . In this paper, we do not explicitly discuss conditions for the existence of the local (or global) optima for (2); as such, we assume that the minimum exists; see [33], [34] for a few relevant applications.

In order to handle a generic embedded submanifold $\tilde{\mathcal{S}}$, we study the behavior of the restricted function $h := f|_{\tilde{\mathcal{S}}}$ from an extrinsic point of view, an approach that can be generalized to any such submanifold. We note that in general, the function f is not convex and the constraint submanifold $\tilde{\mathcal{S}}$ might be disconnected. Thus, here we focus on local convergence results that aim to exploit the inherent geometry of the problem in order to achieve fast convergence rates—with a relatively reasonable computational complexity.

III. GEOMETRY OF THE SYNTHESIS PROBLEM

In order to examine (2), we analyze the domain manifold using machinery borrowed from differential geometry. Note that embedded submanifolds $\tilde{\mathcal{S}}$ endowed with a linear structure can certainly be investigated without using such a machinery. However, neither the corresponding results can be generalized to submanifolds with *nonlinear* structures, nor the geometry induced by the cost function can be exploited for developing the corresponding optimization algorithms.

Before we proceed, it is worth noting that if we were to directly apply the results developed for optimization over the manifolds (such as [26]), it would have been necessary to access a retraction from the tangent bundle $T\tilde{\mathcal{S}}$ onto $\tilde{\mathcal{S}}$. Unfortunately, due to the intricate geometry of \mathcal{S} , such a mapping is generally not available. Additionally, we will see that the Riemannian exponential map, with respect to the inherent geometry associated with optimal control problems, involves a system of ordinary differential equations whose coefficients are solutions to different Lyapunov equations. Therefore, even though it is possible to compute the exponential mapping, in general, its computational overhead is hard to justify. Nonetheless, we show how we can circumvent this issue when the Riemannian tangential projection onto $T\tilde{\mathcal{S}}$ is available—an operation that is more streamlined.

A. Analysis of the domain manifold

It is known that \mathcal{S} is contractible [35], and unbounded when $m \geq 2$ with the topological boundary $\partial\mathcal{S} = \{K \in M(m \times n, \mathbb{R}) \mid \rho(A + BK) = 1\}$ as a subset of $M(m \times n, \mathbb{R})$. Furthermore, \mathcal{S} is open in $M(m \times n, \mathbb{R})$ (by continuity of eigenvalues in the entries of the matrix [36, Theorem 5.2] and passing to the quotient [37, Theorem 3.73]); as such \mathcal{S} is a submanifold without boundary.³

In this paper we focus on \mathcal{S} as a manifold on its own. Note that \mathcal{S} can be covered by a single smooth chart and the tangent bundle of \mathcal{S} , denoted by $T\mathcal{S}$, is diffeomorphic to $\mathcal{S} \times \mathbb{R}^{mn}$, which in turn is diffeomorphic to $\mathcal{S} \times M(m \times n, \mathbb{R})$ under the map $\text{Id}_{\mathcal{S}} \times \text{vec}^{-1}$. We refer to this composition of diffeomorphisms as *the usual identification of the tangent bundle* (or $T_K\mathcal{S} \cong M(m \times n, \mathbb{R})$ at any point $K \in \mathcal{S}$) if we need to identify any element of $T\mathcal{S}$ (or $T_K\mathcal{S}$). In particular, let us denote the coordinates of this global chart by $(x^{i,j})$ for \mathcal{S} , its associated global coordinate frame by $(\frac{\partial}{\partial x^{i,j}})$ or simply $(\partial_{i,j})$, and its dual coframe by $(dx^{i,j})$,

³Cf. [38] for an analogous study of Hurwitz stabilizing controllers.

where $i = 1, \dots, m$ and $j = 1, \dots, n$. Moreover, the (k, ℓ) th element of any matrix $A \in M(m \times n, \mathbb{R})$ is denoted by $[A]_{k,\ell}$ or $[A]^{k,\ell}$ depending on viewing A as a point or a tangent vector, respectively. Then, for example, under the usual identification of tangent bundle, for any fixed i and j , we identify $\partial_{i,j}$ as a matrix in $M(m \times n, \mathbb{R})$ whose elements are $[\partial_{i,j}]^{k,\ell} = 1$ if $k = i$ and $\ell = j$, and otherwise $[\partial_{i,j}]^{k,\ell} = 0$. We also use the Einstein summation convention as explained in [29] for double indices; for example, we write $x^{i,j}\partial_{i,j}$ to denote $\sum_{i=1}^m \sum_{j=1}^n x^{i,j}\partial_{i,j}$.⁴ A vector field V on \mathcal{S} is a smooth map $V: \mathcal{S} \mapsto T\mathcal{S}$, usually written as $K \mapsto V_K$, with the property that $V_K \in T_K\mathcal{S}$ for all $K \in \mathcal{S}$. A covariant 2-tensor field is a smooth real-valued multilinear function of 2 vector fields. We denote the set of all vector fields over \mathcal{S} by $\mathfrak{X}(\mathcal{S})$, and the bundle of covariant 2-tensor fields on \mathcal{S} by $T^2(T^*\mathcal{S})$. Finally, for any general mapping $P: \mathcal{S} \mapsto \star$, we use $P_K, P|_K$ or $P(K)$ to denote the element in \star that $K \in \mathcal{S}$ has been mapped to. The following is a frequently used technical lemma.

Lemma III.1. *The subset \mathcal{M} is an open submanifold of $M(n \times n, \mathbb{R})$, the Lyapunov map $\mathbb{L}: \mathcal{M} \times M(n \times n, \mathbb{R}) \mapsto M(n \times n, \mathbb{R})$ is smooth, and its differential acts as*

$$d\mathbb{L}_{(A,Q)}[E, F] = \mathbb{L}(A, E\mathbb{L}(A, Q)A^\top + A\mathbb{L}(A, Q)E^\top + F)$$

on any $(E, F) \in T_{(A,Q)}(\mathcal{M} \times M(n \times n, \mathbb{R}))$ with the identification that follows by $T_{(A,Q)}(\mathcal{M} \times M(n \times n, \mathbb{R})) \cong T_A\mathcal{M} \oplus T_QM(n \times n, \mathbb{R}) \cong M(n \times n, \mathbb{R}) \oplus M(n \times n, \mathbb{R})$. Furthermore, for any $A \in \mathcal{M}$ and $Q, \Sigma \in M(n \times n, \mathbb{R})$ we have, the so-called Lyapunov-trace property,

$$\text{tr}[\mathbb{L}(A^\top, Q)\Sigma] = \text{tr}[\mathbb{L}(A, \Sigma)Q].$$

Next, we note (see (9) in §V) that many optimal control problems such as SLQR, OLQR and even Linear Quadratic Gaussian (LQG) share a similar cost structure, as $f(K) = \frac{1}{2}\text{tr}[P_K\Sigma_K]$, where mappings $P, \Sigma: \mathcal{S} \mapsto M(n \times n, \mathbb{R})$ send K to

$$P_K := \mathbb{L}(A_{\text{cl}}^\top, Q + K^\top RK), \quad \Sigma_K := \Sigma_1 + K^\top \Sigma_2 K, \quad (3)$$

respectively, with the closed-loop system $A_{\text{cl}} := A + BK$, and $\Sigma_1, \Sigma_2 \succeq 0$ as prescribed matrices with appropriate dimensions. By the Lyapunov-trace property, the cost can be recast as $f(K) = \frac{1}{2}\text{tr}[(Q + K^\top RK)\mathbb{L}(A_{\text{cl}}, \Sigma_K)]$. Motivated by this, we define a covariant 2-tensor field on \mathcal{S} which will subsequently be proved to be a Riemannian metric. Riemannian metrics have been used in the literature to efficiently capture the geometry of the problem; e.g., see [39] for optimizing the Rayleigh quotient on the Grassmann manifold, and [40] for natural policy gradient on Markov decision processes. Note that our metric is different from the “Hessian metric” induced by, say, a convex function [41].

⁴Note that $\text{vec}(\cdot)$ does not preserve algebraic operations on its matrix inputs, e.g., $\text{vec}(AB)$ is not a simple function of $\text{vec}(A)$ and $\text{vec}(B)$; as such, we use double indices to maintain the *matrix* structure of points on \mathcal{S} .

Lemma III.2. Let $\langle \cdot, \cdot \rangle : \mathfrak{X}(\mathcal{S}) \times \mathfrak{X}(\mathcal{S}) \mapsto C^\infty(\mathcal{S})$ denote the mapping that, under the usual identification of the tangent bundle, for any $V, W \in \mathfrak{X}(\mathcal{S})$ sets,⁵

$$\langle V, W \rangle|_K := \text{tr}[(V_K)^\top W_K \mathbb{L}(A_{\text{cl}}, \Sigma_K)], \quad \forall K \in \mathcal{S}.$$

Then this map, induced by a smooth symmetric covariant 2-tensor field, is well-defined.

Now, let $g: \mathcal{S} \mapsto T^2(T^*\mathcal{S})$ be a smooth section of the bundle $T^2(T^*\mathcal{S})$ that sends K to $\langle \cdot, \cdot \rangle|_K$. Then, g is in fact a Riemannian metric under mild conditions formalized below.

Proposition III.3. If $(A_{\text{cl}}, \Sigma_K^{1/2})$ is controllable and for all $K \in \mathcal{S}$, $\Sigma_K \succeq 0$, then (\mathcal{S}, g) is a Riemannian manifold. Moreover, if we define the mapping $Y: \mathcal{S} \mapsto M(n \times n, \mathbb{R})$ sending K to

$$Y_K := \mathbb{L}(A_{\text{cl}}, \Sigma_K),$$

then, with respect to the dual coframe $(dx^{i,j})$, $g = g_{(i,j)(k,\ell)} dx^{i,j} \otimes dx^{k,\ell}$, where each $g_{(i,j)(k,\ell)} \in C^\infty(\mathcal{S})$ satisfies $g_{(i,j)(k,\ell)}(K) = [Y_K]_{\ell,j}$ if $i = k$, and 0 otherwise. Furthermore, the inverse matrix $g^{(i,j)(k,\ell)}$ satisfies $g^{(i,j)(k,\ell)}(K) = [Y_K^{-1}]_{\ell,j}$ if $i = k$, and 0 otherwise.

Remark 1. The premise of Proposition III.3 is satisfied if $\Sigma_K \succ 0$ for all $K \in \mathcal{S}$; e.g., when $\Sigma_1 \succ 0$ and $\Sigma_2 \succeq 0$. Also, by some algebraic manipulations, the well-known Hoyer's algorithm [42] can be viewed as a Riemannian quasi-Newton iteration with respect to this Riemannian metric but with the Euclidean connection—also see Remark 5. Finally, we remark that the developed machinery can be generalized by modifying the Riemannian metric to incorporate a “preconditioning” positive definite $\Theta \succ 0$ such as:

$$\langle V, W \rangle|_K := \text{tr}[(V_K)^\top \Theta W_K \mathbb{L}(A_{\text{cl}}, \Sigma_K)], \quad \forall K \in \mathcal{S}.$$

In this case, a judicious choice of Θ may, in general, improve the performance of corresponding numerical schemes. This metric also has a system theoretic interpretation that will be elaborated upon in our subsequent works.

1) *Riemannian connection on $T\mathcal{S}$* : First, consider a Riemannian submanifold $(\tilde{\mathcal{S}}, \tilde{g})$ with $\tilde{g} := \iota_{\tilde{\mathcal{S}}}^* g$, where $\iota_{\tilde{\mathcal{S}}}^*$ denotes the pull-back by inclusion. In order to understand the second order behavior—i.e., the Hessian—of a smooth function on $\tilde{\mathcal{S}}$, we need to study the notion of connection in $T\mathcal{S}$, and how it relates to analogous construct in the tangent bundle $T\tilde{\mathcal{S}}$ —see [29] for further details. Recall that by Fundamental Theorem of Riemannian Geometry, there exists a unique connection $\nabla: \mathfrak{X}(\mathcal{S}) \times \mathfrak{X}(\mathcal{S}) \mapsto \mathfrak{X}(\mathcal{S})$ in $T\mathcal{S}$ that is compatible with g and symmetric, i.e., for all $U, V, W \in \mathfrak{X}(\mathcal{S})$ we have:

- $\nabla_U \langle V, W \rangle = \langle \nabla_U V, W \rangle + \langle V, \nabla_U W \rangle$,
- $\nabla_U V - \nabla_V U \equiv [U, V]$,

where $[V, W] \in \mathfrak{X}(\mathcal{S})$ denotes the Lie bracket of V and W .

Note that, the restriction of ∇ to $\mathfrak{X}(\tilde{\mathcal{S}}) \times \mathfrak{X}(\tilde{\mathcal{S}})$ would not be a connection in $T\tilde{\mathcal{S}}$ as its range does not necessary

⁵The notation $\langle \cdot, \cdot \rangle$ should not be confused with the (ordinary) inner product in inner-product spaces as it is varying over \mathcal{S} .

lie in $\mathfrak{X}(\tilde{\mathcal{S}})$. However, we can denote the (Riemannian) tangential and normal projections by $\pi^\top: T\mathcal{S}|_{\tilde{\mathcal{S}}} \mapsto T\tilde{\mathcal{S}}$ and $\pi^\perp: T\mathcal{S}|_{\tilde{\mathcal{S}}} \mapsto N\tilde{\mathcal{S}}$, respectively, with $N\tilde{\mathcal{S}}$ indicating the normal bundle of $\tilde{\mathcal{S}}$. Then, by Gauss Formula, if $\tilde{\nabla}: \mathfrak{X}(\tilde{\mathcal{S}}) \times \mathfrak{X}(\tilde{\mathcal{S}}) \mapsto \mathfrak{X}(\tilde{\mathcal{S}})$ denotes the Riemannian connection in the tangent bundle $T\tilde{\mathcal{S}}$, computed via,

$$\tilde{\nabla}_U V = \pi^\top \nabla_U V, \quad (4)$$

for any $U, V \in \mathfrak{X}(\tilde{\mathcal{S}})$ arbitrarily extended to vector fields on a neighborhood of $\tilde{\mathcal{S}}$ in \mathcal{S} .

For computational purposes, we also obtain the Christoffel symbols associated with g (denoted by $\Gamma_{(k,\ell)(p,q)}^{(i,j)}$) in the global coordinate frame. This would, in turn, completely characterize the connection ∇ and facilitates its computation in this frame.

Proposition III.4. Consider a point $K \in \mathcal{S}$ and, under the usual identification of $T\mathcal{S}$, define⁶

$$dY_K(p, q) := \mathbb{L}(A_{\text{cl}}, B\partial_{(p,q)} Y_K A_{\text{cl}}^\top + A_{\text{cl}} Y_K \partial_{(p,q)}^\top B^\top + \partial_{(p,q)}^\top \Sigma_2 K + K^\top \Sigma_2 \partial_{(p,q)})$$

for each $(p, q) \in [m] \times [n]$ where $Y_K = \mathbb{L}(A_{\text{cl}}, \Sigma_K)$. Then, the Christoffel symbols associated with the metric g in the global coordinate frame $(\partial_{(i,j)})$ satisfies $\Gamma_{(k,\ell)(p,q)}^{(i,j)}(K) =$

$$\begin{cases} 1/2 [dY_K(p, q) Y_K^{-1}]_{(\ell,j)}, & \text{if } k = i \neq p, \\ 1/2 [dY_K(k, \ell) Y_K^{-1}]_{(q,j)}, & \text{if } p = i \neq k, \\ -1/2 \sum_s [dY_K(i, s)]_{(q,\ell)} [Y_K^{-1}]_{(s,j)}, & \text{if } p = k \neq i, \\ 1/2 \sum_s ([dY_K(i, \ell)]_{(q,s)} + [dY_K(i, q)]_{(\ell,s)} - [dY_K(i, s)]_{(q,\ell)}) [Y_K^{-1}]_{(s,j)}, & \text{if } p = k = i, \\ 0, & \text{otherwise.} \end{cases}$$

Remark 2. Note that with respect to the global coordinates of (\mathcal{S}, g, ∇) , the geodesic equation is a system of (mn) second-order ordinary differential equations whose varying coefficients involve $(mn)^3$ Christoffel symbols $\Gamma_{(i,\ell)(i,q)}^{(i,j)}$ as obtained above. Therefore, computing the Riemannian Exponential mapping is computationally burdensome; as such, in this work, we avoid using it as a retraction.

B. Extrinsic analysis of a smooth function constrained on a Riemannian submanifold

In this subsection, we study the gradient and Hessian operators of a constrained smooth function from an extrinsic point of view, which is yet to be defined. In other words, we consider $(\tilde{\mathcal{S}}, \tilde{g})$ as a Riemannian submanifold of (\mathcal{S}, g) , where $\tilde{g} = \iota_{\tilde{\mathcal{S}}}^* g$, with $\iota_{\tilde{\mathcal{S}}}^*$ denoting the pull-back by inclusion of $\tilde{\mathcal{S}}$ into \mathcal{S} . Then, by considering any smooth function f on \mathcal{S} , we can define its restriction to $\tilde{\mathcal{S}}$ as

$$h := f|_{\tilde{\mathcal{S}}},$$

and examine how its gradient and Hessian operators are related to those of f . In order to answer this question, we

⁶This coincides with the action of $\partial_{(p,q)}|_K$ on the mapping $K \mapsto Y_K$.

utilize the Riemannian connection to analyze the second order behavior of f (or that of h).

First, recall from [29] that the gradient of f with respect to the Riemannian metric g , denoted by $\text{grad } f \in \mathfrak{X}(\mathcal{S})$, is the unique vector field satisfying

$$\langle V, \text{grad } f \rangle = Vf,$$

for any $V \in \mathfrak{X}(\mathcal{S})$. Then, we denote the *Hessian operator* of $f \in C^\infty(\mathcal{S})$ as the map $\text{Hess } f: \mathfrak{X}(\mathcal{S}) \mapsto \mathfrak{X}(\mathcal{S})$ defined by

$$\text{Hess } f[U] := \nabla_U \text{grad } f,$$

for any $U \in \mathfrak{X}(\mathcal{S})$. Note that we use the same notation to denote the gradient and Hessian operators defined on the submanifold $\tilde{\mathcal{S}}$ (see the appendix). Finally, for any normal vector field N (i.e., a smooth section of $N\tilde{\mathcal{S}}$), the *Weingarten map* in the direction of N is a self-adjoint linear map denoted by $\mathbb{W}_N: \mathfrak{X}(\tilde{\mathcal{S}}) \mapsto \mathfrak{X}(\tilde{\mathcal{S}})$, which defines a smooth bundle homomorphism from $T\tilde{\mathcal{S}}$ to itself (linear on each tangent space), characterized by [29],

$$\langle \mathbb{W}_N[V], W \rangle = \langle N, \pi^\perp(\nabla_V W) \rangle, \quad \forall V, W \in \mathfrak{X}(\mathcal{S}).$$

Now, we can formalize this abstract extrinsic analysis as follows.

Proposition III.5. *Suppose $\tilde{\mathcal{S}}$ is an embedded Riemannian submanifold of \mathcal{S} , both equipped with their respective Riemannian connections. Let $f \in C^\infty(\mathcal{S})$ be any smooth function; then $h := f|_{\tilde{\mathcal{S}}}$ is smooth on $\tilde{\mathcal{S}}$ and we have*

$$\text{grad } h = \pi^\top(\text{grad } f|_{\tilde{\mathcal{S}}}).$$

Furthermore, under the usual identification of $T\tilde{\mathcal{S}} \subset T\mathcal{S}$, for any $V \in \mathfrak{X}(\tilde{\mathcal{S}})$ we have,

$$\text{Hess } h[V] = \pi^\top(\text{Hess } f[V]|_{\tilde{\mathcal{S}}}) + \mathbb{W}_{\pi^\perp(\text{grad } f|_{\tilde{\mathcal{S}}})}[V],$$

where V is arbitrarily extended to vector fields on a neighborhood of $\tilde{\mathcal{S}}$ in \mathcal{S} .

1) *On the choice of connection:* On the manifold (\mathcal{S}, g, ∇) , computing the exponential map requires finding solution to a system of ordinary differential equations of dimension mn . This approach is not only computationally demanding but also does not necessarily provide an exponential map of the submanifold $\tilde{\mathcal{S}}$ (unless it happens to be totally geodesic). In order to avoid the computation of the Riemannian exponential map, it seems reasonable to perform updates by using simpler *retractions* from the tangent bundle to the manifold (cf. [26]); however, in general, we do not have access to such a retraction in our setup. Another computational overhead of utilizing the Riemannian connection associated with the Riemannian metric g pertains to the (mn) -number of Lyapunov equations involved in obtaining the Christoffel symbols at each point.

On the other hand, for applications in which the submanifold appears as $\tilde{\mathcal{S}} = \mathcal{S} \cap \mathcal{K}$, where \mathcal{K} is an affine subspace of $\mathbb{M}(m \times n, \mathbb{R})$, it might seem reasonable to consider the ambient manifold $(\mathcal{S}, g, \bar{\nabla})$, where $\bar{\nabla}$ refers

to the so-called *Euclidean connection*, i.e., the connection whose symbols (with respect to the global coordinates) all vanish (i.e., $\bar{\Gamma}_{(k,\ell)(p,q)}^{i,j} \equiv 0$ on \mathcal{S}). This results in a simpler *Hessian operator* which, however, does not respect the geometry of (\mathcal{S}, g) simply because $\bar{\nabla}$ is not compatible with the metric g —in contrast to its associated Riemannian connection. Nonetheless, for completeness, we also define the *Euclidean Hessian operator* of $f \in C^\infty(\mathcal{S})$ as the map $\bar{\text{Hess}} f: \mathfrak{X}(\mathcal{S}) \mapsto \mathfrak{X}(\mathcal{S})$ defined by

$$\bar{\text{Hess}} f[U] := \bar{\nabla}_U \text{grad } f,$$

for any $U \in \mathfrak{X}(\mathcal{S})$. This operator enjoys similar properties as that of $\text{Hess } f$, but contains different second order information about f (e.g., see Figure 1 for a comparison).

IV. RIEMANNIAN OPTIMIZATION ON SUBMANIFOLDS OF \mathcal{S} WITH LINEAR STRUCTURE

In this section, we propose an optimization algorithm for smooth cost functions, constrained to submanifolds of \mathcal{S} that are endowed with a linear structure; that is, $\tilde{\mathcal{S}} = \mathcal{S} \cap \mathcal{K}$, where \mathcal{K} entails a linear structure in $\mathbb{M}(m \times n, \mathbb{R})$. The proposed algorithm, does not involve the exponential mapping (due to its computational complexity); note that no other retraction from the tangent space onto the manifold \mathcal{S} is known. Instead, we exploit this linear structure together with a geometrically-induced stability certificate that guarantees stability of the iterates by adjusting the respective stepsize.

In what follows, we first introduce this stability certificate and then propose the algorithm. We then show how this certificate can be utilized to choose stepsizes that guarantee a linear convergence rate; furthermore, we will discuss existence of neighborhoods—containing a local minima—on which the algorithm achieves a quadratic rate of convergence.

1) *Stability certificate and (direct) policy optimization:*

Recall that \mathcal{S} is open in $\mathbb{M}(m \times n, \mathbb{R})$. Nonetheless, we provide the following result that quantifies this fact with respect to the problem parameters; an observation that has an immediate utility for analyzing iterative algorithms on \mathcal{S} .

Lemma IV.1. *Consider a smooth mapping $\mathcal{Q}: \mathcal{S} \mapsto \mathbb{M}(n \times n, \mathbb{R})$ that sends K to any $\mathcal{Q}_K \succ 0$. For any direction $G \in T_K \mathcal{S} \cong \mathbb{M}(m \times n, \mathbb{R})$ at any point $K \in \mathcal{S}$, if*

$$0 \leq \eta \leq s_K := \lambda(\mathcal{Q}_K) / (2\bar{\lambda}(\mathbb{L}(A_{\text{cl}}^\top, \mathcal{Q}_K)) \|BG\|_2),$$

then $K^+ := K + \eta G \in \mathcal{S}$; s_K will be referred to as the stability certificate at K .

Proof. As K is stabilizing, for any such \mathcal{Q}_K , there exists a matrix $P = \mathbb{L}(A_{\text{cl}}^\top, \mathcal{Q}_K) \succ 0$ satisfying $P = A_{\text{cl}^+}^\top P A_{\text{cl}^+} + L$ with $A_{\text{cl}^+} := A + BK^+$ and $L := \mathcal{Q}_K + A_{\text{cl}}^\top P A_{\text{cl}} - A_{\text{cl}^+}^\top P A_{\text{cl}^+}$. Therefore, in order to establish that K^+ is stabilizing, as $P \succ 0$ —by the Lyapunov Stability Criterion [43, Theorem8.4]—it suffices to show that $L \succ 0$. Next,

$$\begin{aligned} L &= \mathcal{Q}_K - \eta G^\top B^\top P A_{\text{cl}} - \eta A_{\text{cl}}^\top P B G - \eta^2 G^\top B^\top P B G \\ &\succ \mathcal{Q}_K - a A_{\text{cl}}^\top P A_{\text{cl}} - (1 + 1/a) \eta^2 G^\top B^\top P B G \\ &= (1 + a) \mathcal{Q}_K - a P - (1 + 1/a) \eta^2 G^\top B^\top P B G, \end{aligned} \quad (5)$$

because for any $a > 0$, $P \succ 0$ implies

$$aA_{\text{cl}}^{\top}PA_{\text{cl}} + (\eta^2/a)G^{\top}B^{\top}PBG \succcurlyeq \eta G^{\top}B^{\top}PA_{\text{cl}} + \eta A_{\text{cl}}^{\top}PBG.$$

Now, by recalling the infinite-sum representation of P and the fact that $\mathcal{Q}_K \succeq 0$, we conclude that $\bar{\lambda}(P) \geq \bar{\lambda}(\mathcal{Q}_K) \geq \underline{\lambda}(\mathcal{Q}_K)$. Then, we proceed as follows: if $\bar{\lambda}(P) > \underline{\lambda}(\mathcal{Q}_K)$ then we choose $a = \underline{\lambda}(\mathcal{Q}_K)/(2\bar{\lambda}(P) - 2\underline{\lambda}(\mathcal{Q}_K)) > 0$; otherwise, we choose $a > 2\eta^2\|BG\|_2^2$. Either way, by comparing the minimum eigenvalues of both sides in (5),

$$\lambda(L) \geq \underline{\lambda}(\mathcal{Q}_K)/2 - [2\bar{\lambda}(P)^2/\underline{\lambda}(\mathcal{Q}_K) - \bar{\lambda}(P)]\eta^2\|BG\|_2^2.$$

Therefore, if $|\eta| \leq \underline{\lambda}(\mathcal{Q}_K)/(2\bar{\lambda}(P)\|BG\|_2)$, then $L \succ 0$, hence completing the proof. \square

Remark 3. The proceeding lemma also provides a *conditioning* of the optimization problem in terms of system parameters A, B . In other words, for any choice of $\mathcal{Q}_K \succ 0$ at any $K \in \mathcal{S}$, the ratio $\bar{\lambda}(\mathbb{L}(A_{\text{cl}}^{\top}, \mathcal{Q}_K))/\underline{\lambda}(\mathcal{Q}_K)$ represents a condition number revealing geometric information on the manifold at K . In a sense, this ratio reflects the *Riemannian curvature* of (\mathcal{S}, g, ∇) ; this connection will be further explored in our future work.

Next, we propose an algorithm with convergence guarantees with at least a linear rate (when the iterates are far from the local optima) and eventually a Q-quadratic rate (when the iterates are close enough to the local optima). The complication here is that we do not have access to a retraction with a reasonable computational complexity (see Remark 2). We claim that, starting close enough to a local minimum, a Newton-type method using Riemannian metric and the Euclidean/Riemannian connection must converge quadratically if one could have used the stepsize $\eta = 1$. This is in fact due to the exponential mapping with respect to the Euclidean connection that serves as a retraction with desirable properties. However, the stability certificate suggests that at least away from the local minimum, it might not be possible to use such a large stepsize. Therefore, a stepsize rule has to be deduced—that in turn, hinges upon the stability certificate; the resulting algorithm is summarized in Algorithm 1. Hereafter, we refer to the solution $G \in T_K\tilde{\mathcal{S}}$ of the following equation as the *Newton direction* on $\tilde{\mathcal{S}}$:

$$\text{Hess } h_K[G] = -\text{grad } h_K,$$

where $h = f|_{\tilde{\mathcal{S}}}$; similarly, when $\text{Hess } h$ is replaced by $\overline{\text{Hess}} h$, the corresponding solution is referred to as the *Euclidean Newton direction*.

2) *Linear-quadratic convergence of QRNPO*: In this section, we establish the local linear-quadratic convergence of QRNPO on the submanifold $\tilde{\mathcal{S}}$ using differential geometric techniques [24], [26], [29]. Herein, avoiding the exponential map induced by the Riemannian connection for updating the iterates, and instead relying the stability certificate, adds another layer of complications for the convergence analysis. To proceed, we say that K^* is a critical point of h if $\text{grad } h_{K^*} = 0$; it is *nondegenerate* if $\text{Hess } h_{K^*}$ is

Algorithm 1: Quasi Riemannian-Newton Policy Optimization (QRNPO) for Constrained Problems on \mathcal{S}

- 1: **Initialization:** Problem parameters (A, B) , the linear constraint \mathcal{K} and an initial feasible stabilizing controller $K_0 \in \tilde{\mathcal{S}} = \mathcal{S} \cap \mathcal{K}$
 - 2: Choose a smooth mapping $K \mapsto \mathcal{Q}_K \succ 0$; set $t = 0$
 - 3: **Until stopping criteria are met, do**
 - 4: Find the Newton direction G_t on $\tilde{\mathcal{S}}$ satisfying
$$\text{Hess } h_{K_t}[G_t] = -\text{grad } h_{K_t}$$
 - 5: Use \mathcal{Q}_{K_t} to obtain a stability certificate s_{K_t}
 - 6: Compute step-size $\eta_t = \min\{s_{K_t}, 1\}$
 - 7: Update: $K_{t+1} = K_t + \eta_t G_t$
 - 8: $t \leftarrow t + 1$
-

For examples of the mapping \mathcal{Q} in Line 2 see Remark 6. In Line 4, $\text{Hess } h$ can be replaced by its Euclidean counterpart $\overline{\text{Hess}} h$. The update in Line 7 is possible due to the linear structure of $\tilde{\mathcal{S}}$ induced by \mathcal{K} . For different choices of stopping criteria see [30].

nondegenerate, i.e., $\langle \text{Hess } h_{K^*}[G_1], G_2 \rangle = 0, \forall G_2 \in T_{K^*}\tilde{\mathcal{S}}$ implies that $G_1 = 0 \in T_{K^*}\tilde{\mathcal{S}}$.

Lemma IV.2. *Suppose K^* is a nondegenerate local minimum of $h := f|_{\tilde{\mathcal{S}}}$. Then, it is isolated, $\text{grad } h_{K^*} = 0$, and there exists a neighborhood of K^* on which $\text{Hess } h$ is positive definite. Furthermore, $\overline{\text{Hess}} h_{K^*} = \overline{\text{Hess}} h_{K^*}$.*

Theorem IV.3. *Suppose K^* is a nondegenerate local minimum of $h := f|_{\tilde{\mathcal{S}}}$ over the submanifold $\tilde{\mathcal{S}} = \mathcal{S} \cap \mathcal{K}$ for some linear constraint \mathcal{K} . Then, there exists a neighborhood $\mathcal{U}^* \subset \tilde{\mathcal{S}}$ of K^* with the following property: whenever $K_0 \in \mathcal{U}^*$, the sequence $\{K_t\}$ generated by QRNPO remains in \mathcal{U}^* (therefore, it is stabilizing), and it converges to K^* at least at a linear rate—and eventually—with a quadratic one.*

Remark 4. The above result implies that there exist neighborhoods containing each nondegenerate local minimum of the constrained cost function on which the convergence of QRNPO is guaranteed. The usefulness of this result is that the initial iterate K_0 is not required to be in a (small) neighborhood of the optimum on which the step-size $\eta = 1$ is feasible. Instead, by carefully incorporating the stability certificate (Lemma IV.1), we can obtain a larger basin of attraction for the iterates (see Figure 3 in §VI). Finally, even though the convergence rate is initially linear, as the algorithm proceeds, a quadratic convergence rate is achieved.

V. FEEDBACK SYNTHESIS VIA QRNPO

In this section, we discuss applications of the developed methodology for optimizing the LQR cost over two distinct submanifolds, namely those induced by Structured LQR (SLQR) and Output-feedback LQR (OLQR) problems. Consider a discrete-time linear time-invariant dynamics

$$\mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k, \quad \mathbf{y}_k = C\mathbf{x}_k, \quad (6)$$

where $A \in \mathbb{M}(n \times n, \mathbb{R})$, $B \in \mathbb{M}(n \times m, \mathbb{R})$ and $C \in \mathbb{M}(d \times n, \mathbb{R})$ are the system parameters for some integers n, m and d ; $\mathbf{x}_k, \mathbf{y}_k$ and \mathbf{u}_k denote the states, output and

input vectors at time k , respectively, and \mathbf{x}_0 is given. Conventionally, the Linear Quadratic Regulators (LQR) problem is to find the sequence $\mathbf{u} = (\mathbf{u}_k)_0^\infty \in \ell_2$ that minimizes the following quadratic cost

$$J_{\mathbf{x}_0}(\mathbf{u}) = 1/2 \sum_{k=0}^{\infty} \mathbf{x}_k^\top Q \mathbf{x}_k + \mathbf{u}_k^\top R \mathbf{u}_k, \quad (7)$$

subject to (6), where $Q = Q^\top \succcurlyeq 0$ and $R = R^\top \succ 0$ are prescribed cost parameters. It is well known (see e.g., §22.7 in [44]) that the optimal state-feedback solution \mathbf{u}^* to this problem reduces to solving the Discrete-time Algebraic Riccati Equation (DARE) for the optimal cost matrix P_{LQR} . This results in a linear state-feedback optimal control $\mathbf{u}_k^* = K_{\text{LQR}} \mathbf{x}_k$, where $K_{\text{LQR}} \in \mathbb{M}(m \times n, \mathbb{R})$ is the optimal LQR gain (policy) obtained from P_{LQR} . Furthermore, the associated optimal cost can be obtained as $J_{\mathbf{x}_0}(\mathbf{u}^*) = 1/2 \mathbf{x}_0^\top P_{\text{LQR}} \mathbf{x}_0$.

Naturally, one could think of the LQR cost as a map $K \mapsto J_{\mathbf{x}_0}(\mathbf{u})|_{\mathbf{u}=K\mathbf{x}}$; however, this would depend on \mathbf{x}_0 and generally, its value can still be finite while K is not necessarily stabilizing (i.e., when $K \notin \mathcal{S}$). Instead, in order to avoid the dependency on the initial state while considering the constraints on the policy directly, we pose the following constrained optimization problem,

$$\min_K f(K) := \mathbb{E}_{\mathbf{x}_0 \sim \mathcal{D}} J_{\mathbf{x}_0}(\mathbf{u}) \quad (8)$$

$$\text{s.t. } \mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k, \mathbf{u}_k = K\mathbf{x}_k, \forall k \geq 0, K \in \tilde{\mathcal{S}},$$

where $\tilde{\mathcal{S}}$ is an embedded submanifold of \mathcal{S} , and \mathcal{D} denotes a distribution of zero-mean multivariate random variables of dimension n with covariance matrix Σ_1 so that $0 \prec \Sigma_1 = \Sigma_1^\top \in \mathbb{M}(n \times n, \mathbb{R})$.

Next, we can reformulate (8) as follows. For each stabilizing controller $K \in \mathcal{S}$, from (6) and (7) we have that

$$J_{\mathbf{x}_0}(K\mathbf{x}) = \frac{1}{2} \sum_{k=0}^{\infty} \mathbf{x}_0^\top (A_{\text{cl}}^k)^\top [Q + K^\top R K] A_{\text{cl}}^k \mathbf{x}_0,$$

where $A_{\text{cl}} := A + BK$. Since, A_{cl} is a stability matrix, the sum $\sum_{k=0}^{\infty} (A_{\text{cl}}^k)^\top [Q + K^\top R K] A_{\text{cl}}^k$ converges, which is equal to the unique solution $P_K := \mathbb{L}(A_{\text{cl}}^\top, K^\top R K + Q)$. Therefore, $f(K) = \frac{1}{2} \mathbb{E}_{\mathbf{x}_0 \sim \mathcal{D}} \text{tr}[P_K \mathbf{x}_0 \mathbf{x}_0^\top] = \frac{1}{2} \text{tr}[P_K \Sigma_1]$, and thus the problem in (8) reduces to

$$\min_K f(K) = \frac{1}{2} \text{tr}[P_K \Sigma_1] \quad \text{s.t. } K \in \tilde{\mathcal{S}}. \quad (9)$$

This reformulation of the LQR cost function has been previously adopted in the literature (see e.g., [17], [19], [45]) but the inherent geometry of the submanifold $\tilde{\mathcal{S}}$ has generally been overlooked. In the absence of constraints, i.e., $\tilde{\mathcal{S}} = \mathcal{S}$, Hwer's algorithm is known to converge to the optimal state feedback gain at a quadratic rate [42], given controllability of (A, B) and stability of the initial controller K_0 —see [46] for further study regarding the initial controller. Otherwise, $\tilde{\mathcal{S}}$ may have disconnected components, and in general, the constrained cost function may have stationary points that are not local minima. Nonetheless, in this section, we apply the techniques developed in §III and Algorithm 1 to the constraint arising in the well-known SLQR and OLQR problems. Note that both of these problems can be cast as an optimization in (9) with $\tilde{\mathcal{S}}$ denoting a specific submanifold

of \mathcal{S} that will be further discussed in §V-A.2 and §V-A.3, respectively.

A. Solving for the Newton direction

In order to solve for the Newton direction at any $K \in \tilde{\mathcal{S}}$, suppose that the tuple $(\tilde{\partial}_{(p,q)}|_{(p,q) \in D})$ denotes a smooth local frame for $\tilde{\mathcal{S}}$ on a neighborhood of K , where D is a subset of $[m] \times [n]$ depending on the dimension of $\tilde{\mathcal{S}}$.⁷ In fact, by Proposition III.5, the Newton direction $G = [G]^{k,\ell} \tilde{\partial}_{(k,\ell)}|_K \in T_K \tilde{\mathcal{S}}$ (interpreted as a subspace of $T_K \mathcal{S}$) can be computed by solving the following system of $|D|$ -linear equations (for each index $(p, q) \in D$),

$$\sum_{(k,\ell) \in D} [G]^{k,\ell} h_{;(k,\ell)(p,q)}(K) = - \left\langle \pi^\top (\text{grad } f|_K), \tilde{\partial}_{(p,q)}|_K \right\rangle,$$

where $h_{;(k,\ell)(p,q)}$ denote the coordinates of $\tilde{\nabla}^2 h$ with respect to the local coframe dual to $(\tilde{\partial}_{(p,q)}|_{(p,q) \in D})$ (in tensor notation; e.g., see [29, Example 4.22]). Thus, by (13), $h_{;(k,\ell)(p,q)}(K) = \left\langle \text{Hess } h_K [\tilde{\partial}_{(k,\ell)}|_K], \tilde{\partial}_{(p,q)}|_K \right\rangle$; or with $\text{Hess } h$ replaced by $\overline{\text{Hess } h}$, depending on the connection.

1) *Analysis for the special cost function:* We now turn our attention towards the analysis of the following cost function specific to optimal control problems—see [47] for proofs of the results in this subsection. In order to specialize the results obtained so far to this case, we set $\Sigma_2 = 0$ in (3).

Proposition V.1. *On the Riemannian manifold (\mathcal{S}, g, ∇) , define $f \in C^\infty(\mathcal{S})$ with $f(K) = \frac{1}{2} \text{tr}[P_K \Sigma_1]$, where $P_K = \mathbb{L}(A_{\text{cl}}^\top, K^\top R K + Q)$. Then, f is smooth and under the usual identification of the tangent bundle*

$$\text{grad } f_K = RK + B^\top P_K A_{\text{cl}}.$$

Furthermore, $\text{Hess } f$ and $\overline{\text{Hess } f}$ are both self-adjoint operators such that, for any $E, F \in T_K \mathcal{S}$,

$$\begin{aligned} \langle \text{Hess } f_K [E], F \rangle &= \langle B^\top (S_K [F]) A_{\text{cl}}, E \rangle \\ &+ \langle (R + B^\top P_K B) E + B^\top (S_K [E]) A_{\text{cl}}, F \rangle \\ &- \left\langle \text{grad } f_K, [E]^{k,\ell} [F]^{p,q} \Gamma_{(k,\ell)(p,q)}^{i,j}(K) \partial_{i,j} \right\rangle, \\ \langle \overline{\text{Hess } f}_K [E], F \rangle &= \langle B^\top (S_K [F]) A_{\text{cl}}, E \rangle \\ &+ \langle (R + B^\top P_K B) E + B^\top (S_K [E]) A_{\text{cl}}, F \rangle, \end{aligned}$$

with $\Gamma_{(k,\ell)(p,q)}^{i,j}$ denoting the Christoffel symbols of g and

$$S_K [E] := \mathbb{L}(A_{\text{cl}}^\top, E^\top \text{grad } f_K + (\text{grad } f_K)^\top E).$$

Remark 5. For comparison, in the absence of any constraint (i.e., when $\tilde{\mathcal{S}} = \mathcal{S}$), the Hwer's update $K_{t+1} = -(B^\top P_{K_t} B + R)^{-1} B^\top P_{K_t} A$ in [42] can be written as

$$K_{t+1} = K_t + \hat{G}_t,$$

with the Riemannian quasi-Newton direction \hat{G}_t satisfying,

$$\hat{H}_{K_t} [\hat{G}_t] = - \text{grad } f_{K_t}$$

⁷Each $T_K \tilde{\mathcal{S}}$ can be viewed as a subspace of $T_K \mathcal{S}$ as $\tilde{\mathcal{S}} \subset \mathcal{S}$ is embedded.

where $\widehat{H}_{K_t} := (R + B^\top P_{K_t} B)$ is a positive definite approximation of $\text{Hess } f_{K_t}$ and $\overline{\text{Hess}} f_{K_t}$. The *algebraic coincidence* is that the unit stepsize remains stabilizing throughout these quasi-Newton updates. In general, and particularly on constrained submanifolds $\tilde{\mathcal{S}}$, one needs to instead utilize the stability certificate developed in Lemma IV.1.

As a consequence of Proposition III.5, the next corollary is the analogue of Proposition V.1 on any submanifold of \mathcal{S} .

Corollary V.2. *Under the premise of Proposition V.1, let $h = f|_{\tilde{\mathcal{S}}}$, where $\tilde{\mathcal{S}} \subset \mathcal{S}$ is an embedded Riemannian submanifold with the induced connection. Then, h is smooth and under the usual identification of the tangent bundle*

$$\text{grad } h_K = \pi^\top (RK + B^\top P_K A_{\text{cl}}).$$

Furthermore, $\text{Hess } h$ is a self-adjoint operator and can be characterized as follows: for any $E, F \in T_K \tilde{\mathcal{S}} \subset T_K \mathcal{S}$,

$$\begin{aligned} \langle \text{Hess } h_K[E], F \rangle &= \langle B^\top (S_K[F]) A_{\text{cl}}, E \rangle \\ &+ \langle (R + B^\top P_K B)E + B^\top (S_K[E]) A_{\text{cl}}, F \rangle \\ &- \left\langle \text{grad } h_K, [E]^{k,\ell} [F]^{p,q} \Gamma_{(k,\ell)(p,q)}^{i,j}(K) \partial_{i,j} \right\rangle, \end{aligned}$$

where $\Gamma_{(k,\ell)(p,q)}^{i,j}$ are the Christoffel symbols of g and

$$S_K[E] := \mathbb{L}(A_{\text{cl}}^\top, E^\top \text{grad } f_K + (\text{grad } f_K)^\top E).$$

Remark 6. If $Q \succ 0$, then we can choose the mapping $\mathcal{Q}: K \mapsto \mathcal{Q}_K$ to be $\mathcal{Q}_K = Q + K^\top RK$, and thus the stability certificate s_K as defined in Lemma IV.1 satisfies,

$$s_K \geq \lambda(Q) \lambda(\Sigma_1) / (4f(K) \|BG\|_2),$$

where f denotes the LQR cost. This is due to the fact that $R, Q, \Sigma_1 \succ 0$ and so is $P_K \succ 0$; hence, by the trace inequality, $f(K) \geq (1/2) \lambda(\Sigma_1) \bar{\lambda}(P_K)$. The claimed lower-bound on the stability certificate then follows by combining the last inequality with the definition of the stability certificate. Otherwise if $Q \succeq 0$, one can leverage observability of the pair $(A, Q^{1/2})$ to derive analogous results.

2) State-feedback SLQR: Any desired sparsity pattern on the controller gain K imposes a linear constraint set, denoted by \mathcal{K}_D , which indicates a linear subspace of $\mathbb{M}(m \times n, \mathbb{R})$ with nonzero entries only for a prescribed subset D of entries, i.e., for any $K \in \mathcal{K}_D$ and $(i, j) \notin D$ we must have $[K]_{i,j} = 0$. Then, $\tilde{\mathcal{S}} = \mathcal{S} \cap \mathcal{K}_D$ is a properly embedded submanifold of dimension $|D|$. Furthermore, at any point $K \in \tilde{\mathcal{S}}$ and for any tangent vector $E \in T_K \mathcal{S}$, we can compute the tangential projection $\pi^\top: T_K \mathcal{S} \mapsto T_K \tilde{\mathcal{S}}$ as the unique solution of

$$\text{Proj}_{\mathcal{K}_D} [(E - \tilde{E}^*) Y_K] = 0, \quad (10)$$

where $\text{Proj}_{\mathcal{K}_D}$ denotes the Euclidean projection onto the sparsity pattern \mathcal{K}_D . Note that at each $K \in \tilde{\mathcal{S}}$, the last equality consists of $|D|$ nontrivial linear equations involving $|D|$ unknowns (as the nonzero entries of \tilde{E}^*), which can be solved efficiently. Finally, if $\tilde{\partial}_{(i,j)}$ (as described in §V-A) is taken to be $\tilde{\partial}_{(i,j)} = \partial_{(i,j)}$ for $(i, j) \in D$, then

$(\tilde{\partial}_{(i,j)}|_{(i,j) \in D})$ forms a global smooth frame for $T\tilde{\mathcal{S}}$. Thus, for each $(k, \ell), (p, q) \in D$, the coordinates $h_{;(k,\ell)(p,q)}(K)$ simplifies to

$$\begin{aligned} h_{;(k,\ell)(p,q)}(K) &= \langle B^\top (S_K[\partial_{(p,q)}]) A_{\text{cl}}, \partial_{(k,\ell)} \rangle \\ &+ \langle (R + B^\top P_K B) \partial_{(k,\ell)} + B^\top (S_K[\partial_{(k,\ell)}]) A_{\text{cl}}, \partial_{(p,q)} \rangle \\ &- \left\langle \pi^\top \text{grad } f_K, \Gamma_{(k,\ell)(p,q)}^{i,j}(K) \partial_{i,j} \right\rangle. \end{aligned}$$

3) Output-feedback LQR (OLQR): The OLQR problem can be formulated as the optimization problem in (8) with the submanifold $\tilde{\mathcal{S}} = \mathcal{S} \cap \mathcal{K}_C$, where the constraint set \mathcal{K}_C is defined as

$$\mathcal{K}_C := \{K \in \mathbb{M}(m \times n, \mathbb{R}) \mid K = LC, L \in \mathbb{M}(m \times d, \mathbb{R})\},$$

and $C \in \mathbb{M}(d \times n, \mathbb{R})$ is the prescribed output matrix. For simplicity of presentation, we suppose that C has full rank equal to $d \leq n$. Then, $\tilde{\mathcal{S}} = \mathcal{S}$ is a properly embedded submanifold of \mathcal{S} with dimension md . Also, we can canonically identify each tangent space at $K \in \tilde{\mathcal{S}}$ with $T_K \tilde{\mathcal{S}} \cong \mathcal{K}_C$. Finally, at any $K \in \tilde{\mathcal{S}}$ and for any $E \in T_K \mathcal{S}$, the tangential projection of E is $\pi^\top E = L^* C$ with $L^* \in \mathbb{M}(m \times d, \mathbb{R})$ being the unique solution of the following linear equation

$$L^* C Y_K C^\top = E Y_K C^\top. \quad (11)$$

4) Complexity of QRNPO: The building block of our algorithm is solving DAREs; using the Bartels–Stewart algorithm for solving Sylvester equations, we have a complexity of $\mathcal{O}(n^3)$ for solving each DARE. In the worst case, solving the Newton direction requires the Hessian coefficients, a tangential projection, Christoffel symbols, and the system of $|D|$ -linear equations. The complexity of solving a system of n -linear equations and a $n \times n$ matrix multiplication are both $\mathcal{O}(n^3)$ (without relying on more sophisticated algorithms). By symmetries of $h_{;(k,\ell)(p,q)}$, these require solving $|D|$ number of DAREs. Hence, both the tangential projection and the system of linear equations require $\mathcal{O}(|D|^3)$ operations, and the Hessian coefficients require $\mathcal{O}(n^3 |D|)$ operations.

The computational complexity of QRNPO at each iteration is determined by the complexity of solving for the Newton direction G as the stability certificate s_K involves solving only one DARE within $\mathcal{O}(n^3)$ operations—say using the Bartels–Stewart algorithm. Thus, each iteration of QRNPO with $\overline{\text{Hess}}$ has a rudimentary computational complexity of $\mathcal{O}(n^3 |D| + |D|^3 + n^3) \approx \mathcal{O}(n^4 m + n^3 m^3)$ for the largest possible $|D| = nm$. On the other hand, QRNPO with Hess requires additional computations of approximately $|D|$ number of DAREs for the Christoffel symbols, but resulting in the same complexity. Finally, we note that efficient computation of the Christoffel symbols for these problems can benefit from additional structures such as sparsity; see [47] for further discussion.

VI. NUMERICAL RESULTS

In this section, we provide numerical examples for optimizing the LQR cost over submanifolds induced by SLQR and OLQR problems. Recall that, for each of these problems,

we can compute the coordinate functions of the covariant Hessian $h_{;(k,\ell)(p,q)}(K)$ with respect to the corresponding coordinate frame described in previous subsections. Therefore, finding the Newton direction G at any point $K \in \tilde{\mathcal{S}}$ reduces to solving the system of linear equations for the unknowns $[G]^{k,\ell}$, as described in §V-A, and forming the Newton direction as $G = [G]^{k,\ell} \hat{\partial}_{(k,\ell)}|_K \in T_K \tilde{\mathcal{S}}$.

For each of SLQR and OLQR problems, we have simulated three different algorithms, the first two are the variants of QRNPO where we use Riemannian connection or Euclidean connection to compute $\text{Hess } h$ or $\overline{\text{Hess}} h$, respectively. Note that although $\text{Hess } h_{K^*} = \overline{\text{Hess}} h_{K^*}$ whenever $\text{grad } h_{K^*} = 0$ (as shown in Lemma IV.2), this would not necessarily be the case where $\text{grad } h$ does not vanish; therefore, we expect $\text{Hess } h$ and $\overline{\text{Hess}} h$ to contain distinct information on neighborhoods of isolated local minima, that directly influence the performance of QRNPO as will be discussed below. The third algorithm is the Projected Gradient Descent (PGD) as studied in [17]. PGD is feasible for constraints in our examples as, under relevant assumptions, one is able to perform PGD updates by having access to merely the projection onto linear subspace of matrices—see for example [17, Theorem 7.1]. Note that the Lipschitz estimate provided in [17, Lemma 7.9] is conservative. Here, instead we choose a larger constant step size which generates stabilizing iterates and improves the performance of PGD—but not the convergence rate which remains sublinear. Finally, for comparison, we also implement the Natural Projected Gradient Descent (NPGD) algorithm with the same constant stepsize where, comparing to PGD, the Euclidean projected gradient is replaced by the tangential projection of the Riemannian gradient. Note that, both PGD and NPGD algorithms have similar (initial) sublinear rates [17], [20]; as such, despite their respective progress at the onset of the iterates, these algorithms—without further stepsize adjustment—become rather slow over time and not practically convergent. The code for generating these results can be found at [48].

Example 1 (Trajectories of QRNPO using Hess versus $\overline{\text{Hess}}$). In order to illustrate how the performance of QRNPO is different in terms of using the Riemannian connection ($\text{Hess } h$) versus Euclidean connection ($\overline{\text{Hess}} h$), we consider an example with system parameters $(A|B|Q|R|\Sigma) = \begin{pmatrix} 0.8 & 1.0 & 0.0 & 1.0 & 10.0 & 0.0 & 0.1 & 0.0 & 1.0 & 0.0 \\ 0.0 & 0.9 & 1.0 & 0.0 & 0.0 & 0.5 & 0.0 & 0.1 & 0.0 & 5.0 \end{pmatrix}$. We run QRNPO and PGD algorithms for both SLQR and OLQR problems involving two decision variables, so that we can plot the trajectories of the iterates over the level curves of the associated cost functions from different initial conditions (as illustrated in Figure 2 and Figure 3, respectively). In this example, the stopping criterion for QRNPO is when the error of iterates from optimality is below a small tolerance (10^{-12}); unless the Hessian fails to be positive definite for QRNPO using $\overline{\text{Hess}}$, we run the algorithm for a fixed number of iterations. Since PGD does not practically converge even with large number of iterates due to its sublinear convergence rate, it is terminated when

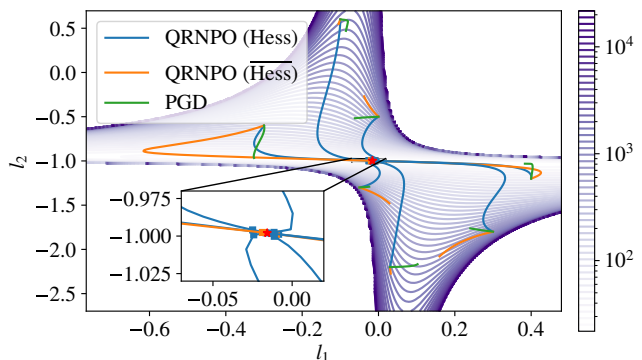


Fig. 2: The trajectories of iterates $K = \text{diag}(l_1, l_2)$ generated by QRNPO (with $\text{Hess } h$ and $\overline{\text{Hess}} h$) and PGD—from different initial points—for the SLQR problem with constraint $D = \{(1, 1), (2, 2)\}$, over the level curves of f in (8).

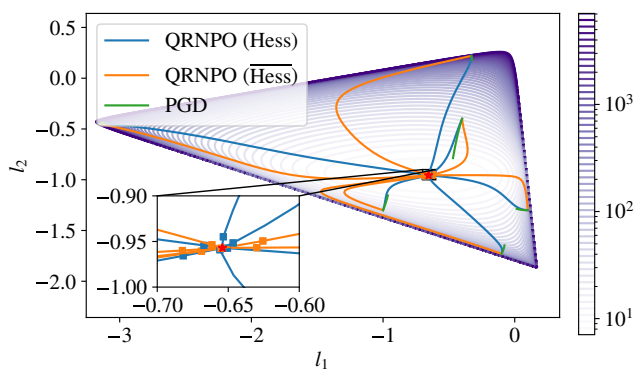


Fig. 3: The trajectories of iterates $L = (l_1 \ l_2)^T$ generated by QRNPO (with $\text{Hess } h$ and $\overline{\text{Hess}} h$) and PGD—from different initial points—for the OLQR problem with output matrix $C = \begin{pmatrix} 1.0 & 1.0 \end{pmatrix}$, over the level curves of f in (8).

QRNPO, using Hess , has converged.

In reference to Figure 2, we first note that QRNPO with $\overline{\text{Hess}}$ does not converge if initialized away from the local minimum (and away from the line $l_2 = -1$) as the Euclidean Hessian fails to be positive definite therein (see Figure 1). On the other hand, QRNPO with Hess successfully captures the inherent geometry of the problem and converges from all initializations. These observations exemplify how QRNPO can exploit the connection compatible with the metric (inherent to the cost function) in order to provide more effective iterate updates. Second, the square marker on each trajectory of QRNPO indicates the first time stepsize $\eta_t = 1$ is guaranteed to be stabilizing (i.e., $s_{K_t} \geq 1$). It can be seen that the neighborhood of the local minimum (zoomed in)—on which the identity stepsize is possible—is relatively small. Whereas, by using the stability certificate, the specific choice of stepsize adopted here enables QRNPO to handle initialization further away from the local minimum.

Next, notice that in both Figures 2 and 3, the trajectories of QRNPO with Hess are much more favorable in comparison to QRNPO with $\overline{\text{Hess}}$, particularly, when initialized from points further away from the local minimum and closer to the boundary. Additionally, similar to Figure 2, the region

on which the unit stepsize is guaranteed to be stabilizing is relatively small in Figure 3.

Example 2 (Randomly selected system parameters). Next, we consider an example with $n = 6$ number of states and $m = 3$ number of inputs, and simulate the behavior of QRNPO and PGD for 100 randomly sampled system parameters. Particularly, the parameters (A, B) are sampled from a zero-mean unit-variance normal distribution, where A is scaled so that the open-loop system is stable, i.e., $K_0 = 0$ is stabilizing, and the pair is controllable. Furthermore, we choose $Q = \Sigma = I_n$ and $R = I_m$ in order to consistently compare the convergence behaviors across different samples. For the SLQR problem, we randomly sample for the sparsity pattern D so that at least half of the entries are zero and all of them have converged from $K_0 = 0$ in less than 30 iterations. For the OLQR problem, we also randomly sample the output matrix C with $d = 2$, where 98% and 92% of the corresponding iterates have converged from $K_0 = 0$ in less than 50 iterations using Hess and $\overline{\text{Hess}}$, respectively.

The minimum, maximum and median progress of the three algorithms for both SLQR and OLQR problems are illustrated in Figure 4a and Figure 4b, respectively. As guaranteed by Theorem IV.3, the linear-quadratic convergence behavior of QRNPO is observed in these problems. Especially, the quadratic behavior starts as soon as the stability certificate exceeds 1. In both cases, QRNPO with Hess h (blue curves) built upon the Riemannian connection has a superior convergence rate compared with the case of using the Euclidean connection (orange curves); this was expected as the Riemannian connection is compatible with the metric induced by the geometry inherent to the cost function itself. This superior performance of QRNPO with Hess h , in the meantime, requires computation of the Christoffel symbols.

VII. CONCLUSIONS AND FUTURE DIRECTIONS

In this work, we considered the problem of optimizing a smooth function over submanifolds of Schur stabilizing controllers \mathcal{S} . In order to treat this problem in a more general setting, we studied the first and second order behavior of a smooth function when constrained to an embedded submanifold from an extrinsic point of view. Subsequently, using the second order information of the restricted function, we developed an algorithm that guarantees convergence to local minima—at least with a linear rate—and eventually with a quadratic rate. Combining this approach with backtracking line-search techniques or positive definite modifications of the Hessian operator [49], [50] can be considered as immediate future directions for a global convergence analysis.

Even though the proposed algorithm depends on the linear structure of $\tilde{\mathcal{S}}$, the machinery developed here can be utilized for other submanifolds, a topic that will be considered in our future works. For example, in contrast to the SLQR and OLQR problems considered, we can explore how a constraint on the *average input energy* translates to a *nonlinear* constraint that pertains to the inherent geometry

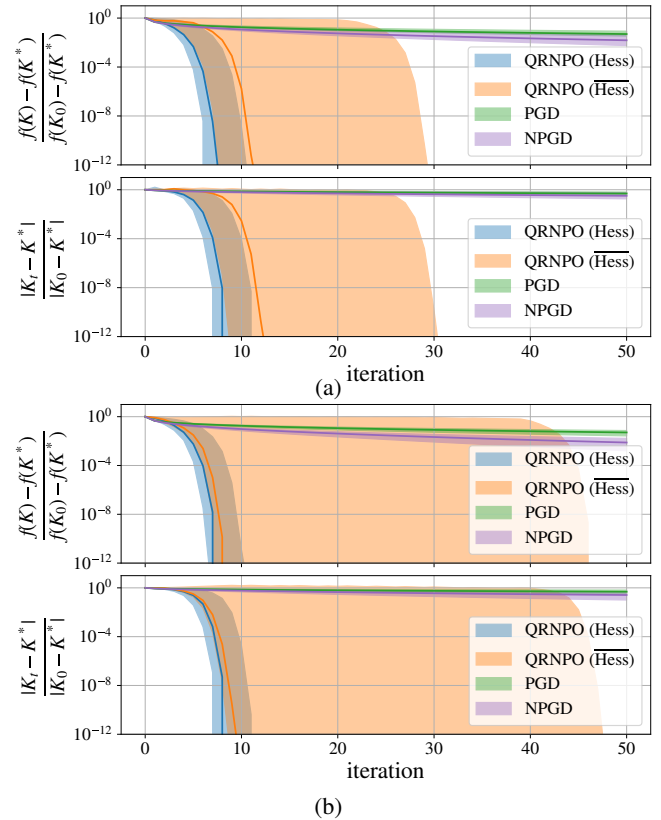


Fig. 4: The min, max and median progress of normalized error of iterates and cost values at each iteration of QRNPO (with Hess h and $\overline{\text{Hess}}$), PGD and NPGD for the (a) SLQR and (b) OLQR problems with 100 different randomly sampled system parameters, sparsity patterns and output matrices.

of the LQR problem. In this direction, we define the *average input energy*, $E_u := \mathbb{E}_{x_0 \sim \mathcal{D}} \|\mathbf{u}\|_{\ell_2}^2$, where $\|\cdot\|_{\ell_2}$ refers to ℓ_2 -norm. If a static linear policy, i.e., $\mathbf{u} = K\mathbf{x}$ for $K \in \mathcal{S}$, is desired for this problem setup, then the closed-loop system assumes the form $\mathbf{x}_k = (A_{cl})^k \mathbf{x}_0$; one can now show that under the usual identification of $T_K \mathcal{S}$ with $M(m \times n, \mathbb{R})$, we have $E_u(K) = |K|_{g_K}^2$. But then, $E_u: \mathcal{S} \mapsto \mathbb{R}$ is smooth by composition, and as such, every regular level set of E_u translates to an upperbound on the average input energy. Regular Level Set Theorem now implies that the average input energy optimal control synthesis can be pursued via an embedded submanifold of \mathcal{S} which has a nonlinear but simple structure, whenever considered in the associated Riemannian geometry. Solving this problem still requires an efficient retraction that would substitute the linear updates possible in SLQR and OLQR problems. The framework discussed in this work allows the integration of such retractions in the synthesis procedure. As such, the proposed work opens up a new approach for solving a wide range of constrained optimization problems over the manifold of Schur stabilizing controllers.

ACKNOWLEDGMENTS

The first author thank Professor John M. Lee for his inspiring lectures on differential geometry and insightful com-

ments on this manuscript. The authors also thank Jingjing Bu for helpful discussions on first order methods for control, as well as the Associate Editor and anonymous reviewers for constructive feedback and suggestions that have been reflected in this manuscript.

APPENDIX

On the Hessian Operator. The Hessian operator (denoted by $\text{Hess } f$) as introduced in §III-B is well-defined and the value of $\text{Hess } f[U]$ at any $K \in \mathcal{S}$ depends only on U_K ; this is due to the property for the connection. Note that

$$\begin{aligned} \langle \text{Hess } f[U], W \rangle &= U \langle \text{grad } f, W \rangle - \langle \text{grad } f, \nabla_U W \rangle \\ &= U(Wf) - (\nabla_U W)f = W(Uf) - (\nabla_W U)f, \end{aligned} \quad (12)$$

$\forall U, W \in \mathfrak{X}(\mathcal{S})$, where the first equality is the consequence of having the Riemannian connection compatible with the metric, the second one is by the definition of $\text{grad } f$, and the last one is due to symmetry of the Riemannian connection. Thus, by (12), the Hessian operator is self-adjoint, i.e.,

$$\langle \text{Hess } f[U], W \rangle = \langle U, \text{Hess } f[W] \rangle.$$

Similarly, we can consider $\text{Hess } h$ for any smooth function $h \in C^\infty(\tilde{\mathcal{S}})$, where we consider the submanifold $\tilde{\mathcal{S}} \subset \mathcal{S}$ with the induced Riemannian metric and the associated Riemannian connection of \mathcal{S} .

Next, for the computational purposes, we would like to introduce the *covariant Hessian* of f with respect to g , denoted by $\nabla^2 f$ [29, Proposition 4.17]. It is a 2-tensor field obtained by taking total covariant derivative of f twice. The Riemannian connection is symmetric, and so is the covariant Hessian. Furthermore, the covariant Hessian and Hessian operator are related as,

$$\nabla^2 f[W, U] = \nabla_U \nabla_W f - (\nabla_U W)f = \langle \text{Hess } f[U], W \rangle, \quad (13)$$

where the last equality follows by (12).

Recall now that $\text{grad } f = (df)^\sharp$, where \sharp denotes the index raising operator, referred to as the *sharp* operator [29, Chapter 2]. Also, $\text{Hess } f: \mathfrak{X}(\mathcal{S}) \mapsto \mathfrak{X}(\mathcal{S})$ can be viewed as the total covariant derivative of $\text{grad } f$, i.e., $\text{Hess } f = \nabla \text{grad } f$. Then

$$\text{Hess } f = \nabla (df)^\sharp = (\nabla(df))^\sharp = (\nabla^2 f)^\sharp, \quad (14)$$

where the equality in the middle follows by the fact the index raising operator commute with the covariant derivative operator, and the last equality is due to the definition of connection for a smooth function $f \in C^\infty(\mathcal{S})$. Note that in (14), the index raising refer to the second argument of $\nabla^2 f$. However, as the covariant Hessian of any smooth function is a symmetric 2-tensor field, the index raising could be with respect to any of the entries. Finally, similar definitions and relations as discussed above are available for h as a smooth function on the embedded Riemannian submanifold $\tilde{\mathcal{S}}$ with the induced metric and corresponding connection; these are omitted for brevity. \square

Proof of Lemma III.1. Since $\rho: \mathbf{M}(n \times n, \mathbb{R}) \mapsto \mathbb{R}$ is a continuous map, \mathcal{M} is an open subset of $\mathbf{M}(n \times n, \mathbb{R})$ and thus an open submanifold. For each $A \in \mathcal{M}$, by Lyapunov Stability Criterion, there exists a unique solution X to (1) which has the infinite-sum representation. But, as for each $A \in \mathcal{M}$, the series converges, each matrix entry of X can be written as a convergent power series of elements of A and Z . Therefore, each matrix entry of X is a real analytic function of several variables (as defined in [51]) on the open subset $\mathcal{M} \times \mathbf{M}(n \times n, \mathbb{R}) \subset \mathbf{M}(n \times n, \mathbb{R}) \times \mathbf{M}(n \times n, \mathbb{R})$. Hence, we conclude that \mathbb{L} is a well-defined smooth map.⁸ Next, under the identification in the premise, it follows that,

$$d \mathbb{L}_{(A,Q)}[E, F] = d \mathbb{L}_{(A,Q)}[E, 0] + d \mathbb{L}_{(A,Q)}[0, F].$$

However, \mathbb{L} is linear in the second entry, so $d \mathbb{L}_{(A,Q)}[0, F] = \mathbb{L}(A, F)$. Also since \mathcal{M} is open, for small enough ε , $\gamma: [0, \varepsilon] \mapsto \mathcal{M} \times \mathbf{M}(n \times n, \mathbb{R})$ with $\gamma(t) = (A + tE, Q)$ is a well-defined smooth curve starting at (A, Q) whose initial velocity is $(E, 0)$. Then,

$$d \mathbb{L}_{(A,Q)}[E, 0] = d/dt|_{t=0} \mathbb{L} \circ \gamma(t).$$

Let $X_t := \mathbb{L} \circ \gamma(t)$ and $X := \mathbb{L} \circ \gamma(0)$; then we obtain,

$$\begin{aligned} X_t - X &= \mathbb{L}(A, t(EXA^\top + AXE^\top) + \mathcal{O}(t^2)) \\ &= t \mathbb{L}(A, EXA^\top + AXE^\top) + \mathcal{O}(t^2), \end{aligned}$$

where the first equality is by direct algebraic manipulation and the second one follows by linearity of \mathbb{L} in the second entry. Therefore, $d \mathbb{L}_{(A,Q)}[E, 0] = \mathbb{L}(A, EXA^\top + AXE^\top)$, and the first claim follows by adding the two computed differentials and using linearity of \mathbb{L} in the second entry again. Finally, note that any square matrix has a spectrum identical to its transpose; therefore if $A \in \mathcal{M}$ then $A^\top \in \mathcal{M}$, and thus the last property follows by the convergent series representations of $\mathbb{L}(A^\top, Q)$ and $\mathbb{L}(A, \Sigma)$, as well as the cyclic permutation property of trace. \square

Proof of Lemma III.2. By Lemma III.1 for each $K \in \mathcal{S}$, $\mathbb{L}(A_{\text{cl}}, \Sigma_K)$ is uniquely determined, symmetric and smooth in K , since A_{cl} is stabilizing. Also, $\mathbb{L}(A_{\text{cl}}, \Sigma_K) \in \mathbf{M}(n \times n, \mathbb{R})$ is positive semidefinite by observing the infinite-sum representation of solution to Lyapunov equation and the fact that $\Sigma_K \succeq 0$. Next, $\text{tr}[(V_K)^\top W_K \mathbb{L}(A_{\text{cl}}, \Sigma_K)]$ is a smooth function of elements of V_K, W_K and $\mathbb{L}(A_{\text{cl}}, \Sigma_K)$. Therefore, for any $V, W \in \mathfrak{X}(\mathcal{S})$, the function $\langle V, W \rangle$, as defined in the premise, is well-defined and smooth on \mathcal{S} . Additionally, by linearity of trace, we observe that $\langle \cdot, \cdot \rangle$ is multilinear over $C^\infty(\mathcal{S})$, i.e.,

$$\langle fU + hV, W \rangle = f \langle U, W \rangle + h \langle V, W \rangle,$$

for any $f, h \in C^\infty(\mathcal{S})$ and $U, V \in \mathfrak{X}(\mathcal{S})$, and similarly for the second entry. Therefore, by Tensor Characterization Lemma [32, Lemma 12.24], it is induced by a smooth

⁸An alternative argument can be provided by the closed form solution of (1) and its *vectorization* involving rational functions of several variables with non-vanishing denominators—cf. Lemma 3.6 in [17].

covariant 2-tensor field. Finally, symmetry follows from the fact that for any $K \in \mathcal{S}$, the cyclic property of trace implies,

$$\begin{aligned} \langle V, W \rangle \Big|_K &= \text{tr}[(\mathbb{L}(A_{\text{cl}}, \Sigma_K))^\top (W_K)^\top V_K] \\ &= \text{tr}[(W_K)^\top V_K (\mathbb{L}(A_{\text{cl}}, \Sigma_K))^\top] = \langle W, V \rangle \Big|_K, \end{aligned}$$

as $\mathbb{L}(A_{\text{cl}}, \Sigma_K)$ is symmetric. \square

Proof of Proposition III.3. We know that \mathcal{S} is a smooth manifold, and by Lemma III.2 and Smoothness Criteria for Tensor Fields [32, Proposition 12.19], g is a smooth symmetric covariant 2-tensor field. Hence, it suffices to show that it is positive definite at each point $K \in \mathcal{S}$. But $\Sigma_K \succeq 0$ and $(A_{\text{cl}}, \Sigma_K)$ is controllable; therefore $Y_K = \mathbb{L}(A_{\text{cl}}, \Sigma_K)$ is a positive definite matrix, implying that $g_K(E, E) = \text{tr}[(EY_K^{1/2})^\top EY_K^{1/2}] \geq 0$, for any $E \in T_K\mathcal{S}$ with equality if and only if E is the zero element. Next, to compute the coordinate representation of g , for each coordinate pairs (i, j) and (k, ℓ) we have,

$$g_{(i,j)(k,\ell)}(K) = g_K(\partial_{i,j}|_K, \partial_{k,\ell}|_K) = \text{tr}[(\partial_{i,j}|_K)^\top \partial_{k,\ell}|_K Y_K],$$

under the usual identification of $T_K\mathcal{S}$. Since under this identification $\partial_{i,j}$ corresponds to the element of $M(m \times n, \mathbb{R})$ with entry 1 in (i, j) -th coordinate and zero elsewhere, the expression for $g_{(i,j)(k,\ell)}$ follows by direct computation of the last equality. Finally, by definition of *inverse matrix*, for each (i, j) and (k, ℓ) , we must have $\sum_{r,s} g^{(i,j)(r,s)} g_{(r,s)(k,\ell)} = 1$ if $(i, j) = (k, \ell)$ and 0 otherwise. Next, for each $i = k$, let $[g^{(k,\cdot)(k,\cdot)}]$ denote the matrix with $g^{(k,j)(k,s)}$ as its (j, s) th entry. Then, by the expression for $g_{(i,j)(k,\ell)}$, it must satisfy $[g^{(k,\cdot)(k,\cdot)}] Y_K = I_n$, and therefore $[g^{(k,\cdot)(k,\cdot)}] = Y_K^{-1}$ as $Y_K \succ 0$. The rest of the reasoning follows by performing a similar computation for each $i \neq k$ and noting the zero pattern in the expression for $g_{(i,j)(k,\ell)}$. \square

Proof of Proposition III.4. We know that

$$\begin{aligned} \Gamma_{(k,\ell)(p,q)}^{(i,j)} &= \sum_{r,s} (g^{(i,j)(r,s)} / 2) (\partial_{(k,\ell)} g_{(r,s)(p,q)} \\ &\quad + \partial_{(p,q)} g_{(r,s)(k,\ell)} - \partial_{(r,s)} g_{(k,\ell)(p,q)}), \end{aligned} \quad (15)$$

for any $(i, j), (k, \ell), (p, q) \in [m] \times [n]$, where $g^{(i,j)(r,s)}$ denotes the inverse matrix of $g_{(i,j)(r,s)}$. If $k \neq i \neq p \neq k$, then by sparsity pattern in the expression for $g_{(i,j)(k,\ell)}$ in Proposition III.3, we obtain $\Gamma_{(k,\ell)(p,q)}^{(i,j)} = 0$. Next, if $k = i \neq p$, then (15) simplifies to

$$\sum_s \frac{g^{(i,j)(i,s)}}{2} (\partial_{(p,q)} g_{(i,s)(i,\ell)}) = \frac{1}{2} (\partial_{(p,q)} [Y_K]_{(\ell,\cdot)}) [Y_K^{-1}]_{(\cdot,j)}.$$

Next, for any fix i, p and q , let $\Gamma_{(i,\cdot)(p,q)}^{(i,\cdot)}$ denote the $n \times n$ matrix with $\Gamma_{(i,\ell)(p,q)}^{(i,j)}$ as its (ℓ, j) entry. Then it must satisfy

$\Gamma_{(i,\cdot)(p,q)}^{(i,\cdot)} = 1/2 (\partial_{(p,q)} Y_K) Y_K^{-1}$, where $\partial_{(p,q)} Y_K$ indicates the action of tangent vector $\partial_{(p,q)}$ on the composite map $K \mapsto (A_{\text{cl}}, \Sigma_K) \xrightarrow{\mathbb{L}} Y_K$. By Lemma III.1, we can compute,

$$\partial_{(p,q)} Y_K = d \mathbb{L}_{(A_{\text{cl}}, \Sigma_K)} \left[B \partial_{(p,q)}, \partial_{(p,q)}^\top \Sigma_2 K + K^\top \Sigma_2 \partial_{(p,q)} \right]$$

under the usual identification of $T_K\mathcal{S}$. Thus, $\partial_{(p,q)} Y_K = dY_K(p, q)$ which proves the second case. The third case follows by the symmetry of the Riemannian connection, i.e., $\Gamma_{(k,\ell)(p,q)}^{(i,j)} = \Gamma_{(p,q)(k,\ell)}^{(i,j)}$. Next, if $p = k \neq i$, then by Proposition III.3, (15) simplifies to

$$\sum_s \frac{g^{(i,j)(i,s)}}{2} (\partial_{(i,s)} g_{(k,\ell)(k,q)}) = \frac{1}{2} \sum_s [\partial_{(i,s)} Y_K]_{(q,\ell)} [Y_K^{-1}]_{(s,j)},$$

with $\partial_{(i,s)} Y_K = dY_K(i, s)$ computed similarly. Finally, if $k = i = p$ then similarly (15) simplifies to

$$\sum_s \frac{g^{(i,j)(i,s)}}{2} (\partial_{(i,\ell)} g_{(i,s)(i,q)} + \partial_{(i,q)} g_{(i,s)(i,\ell)} - \partial_{(i,s)} g_{(i,\ell)(i,q)}),$$

and substituting each term similarly completes the proof. \square

Proof of Proposition III.5. Note that f is smooth and $\tilde{\mathcal{S}}$ is an embedded submanifold of \mathcal{S} . Therefore, $h: \tilde{\mathcal{S}} \mapsto \mathbb{R}$ is smooth by restriction and we can define $\text{grad } h$ and $\text{Hess } h$ on $\tilde{\mathcal{S}}$. But, $\text{grad } h \in \mathfrak{X}(\tilde{\mathcal{S}})$ is the unique vector field on $\tilde{\mathcal{S}}$ such that $\tilde{g}(W, \text{grad } h) = Wh$ for any $W \in \mathfrak{X}(\tilde{\mathcal{S}})$. Unraveling the definition implies that for any $K \in \tilde{\mathcal{S}} \subset \mathcal{S}$,

$$\begin{aligned} d h_K(W_K) &= d f_K(d \iota_{\tilde{\mathcal{S}}}(W_K)) = g_K(d \iota_{\tilde{\mathcal{S}}}(W_K), \text{grad } f_K) \\ &= g_K(d \iota_{\tilde{\mathcal{S}}}(W_K), d \iota_{\tilde{\mathcal{S}}}(\pi^\top \text{grad } f_K)), \end{aligned}$$

as $h = f \circ \iota_{\tilde{\mathcal{S}}}$ and thus $dh = df \circ d\iota_{\tilde{\mathcal{S}}}$, where the last equality follows by the fact that $\iota_{\tilde{\mathcal{S}}}(W_K) \in T_K\mathcal{S}$ is tangent to $\tilde{\mathcal{S}}$. By definition of tangential projection, $\pi^\top(\text{grad } f|_{\tilde{\mathcal{S}}})$ is then a vector field on $\tilde{\mathcal{S}}$ that satisfies

$$Wh = \tilde{g}(W, \pi^\top \text{grad } f|_{\tilde{\mathcal{S}}}),$$

for any $W \in \mathfrak{X}(\tilde{\mathcal{S}})$. Therefore, the first claim follows by uniqueness of the gradient. Next, note that the Hessian operator of $h \in C^\infty(\tilde{\mathcal{S}})$ is defined as $\text{Hess } h[V] := \tilde{\nabla}_V \text{grad } h$, for any $V \in \mathfrak{X}(\tilde{\mathcal{S}})$. But then, the first claim together with (4) and the linearity of connection imply that

$$\begin{aligned} \text{Hess } h[V] &= \pi^\top \nabla_V (\pi^\top (\text{grad } f|_{\tilde{\mathcal{S}}})) \\ &= \pi^\top (\text{Hess } f[V]|_{\tilde{\mathcal{S}}}) - \pi^\top \nabla_V (\pi^\perp (\text{grad } f|_{\tilde{\mathcal{S}}})), \end{aligned}$$

where all $V, \pi^\top(\text{grad } f|_{\tilde{\mathcal{S}}})$ and $\pi^\perp(\text{grad } f|_{\tilde{\mathcal{S}}})$ are extended arbitrarily to vector fields on a neighborhood of $\tilde{\mathcal{S}}$ in \mathcal{S} . Finally, the extrinsic expression of $\text{Hess } h$ follows by The Weingarten Equation [29, Proposition 8.4], indicating that $\pi^\top \nabla_V (\pi^\perp(\text{grad } f|_{\tilde{\mathcal{S}}})) = -\mathbb{W}_{\pi^\perp(\text{grad } f|_{\tilde{\mathcal{S}}})}[V]$. \square

Proof of Lemma IV.2. Let $\tilde{\gamma}: (-\varepsilon, \varepsilon) \mapsto \tilde{\mathcal{S}}$ denote the smooth geodesic curve on the submanifold $\tilde{\mathcal{S}}$ with $\tilde{\gamma}(0) = K^*$ and $\tilde{\gamma}'(0) = F$ for an arbitrary $F \in T_{K^*}\tilde{\mathcal{S}}$. Define $\ell(t) := h \circ \tilde{\gamma}(t): (-\varepsilon, \varepsilon) \mapsto \mathbb{R}$ which is smooth by composition. Then, K^* is a local minimum for h , so is $t = 0$ for $\ell(t)$ following by smoothness of $\tilde{\gamma}$. Therefore,

$$0 = \ell'(0) = d h_{\tilde{\gamma}(0)} \circ \tilde{\gamma}'(0) = \langle \text{grad } h_{K^*}, F \rangle,$$

and as F was an arbitrary tangent vector, we conclude that $\text{grad } h_{K^*} = 0$. Now, recall that non-degenerate critical points are isolated [52, Corollary 2.3]. Next, Taylor's formula for ℓ at $t = 0$ yields

$$\ell(t) = \ell(0) + t \langle \text{grad } h_{K^*}, F \rangle + 1/2 \ell''(s) t^2,$$

for some $s \in (0, t)$. As $\text{grad } h_{K^*} = 0$ and $t = 0$ is a local minimum of $\ell(t)$, we must have $\ell''(s) \geq 0$; by tending $t \mapsto 0$, smoothness of ℓ implies that $\ell''(0) \geq 0$. But,

$$\ell''(t) = \tilde{D}_t \langle \text{grad } h_{\tilde{\gamma}(t)}, \tilde{\gamma}'(t) \rangle = \langle \tilde{D}_t \text{grad } h_{\tilde{\gamma}(t)}, \tilde{\gamma}'(t) \rangle$$

where \tilde{D}_t denotes the covariant derivative along $\tilde{\gamma}$ on $\tilde{\mathcal{S}}$, and the last equality follows by its compatibility with the metric and the fact that $\tilde{\gamma}$ is a geodesic (so that $\tilde{D}_t \tilde{\gamma}'(t) \equiv 0$). As $\text{grad } h|_{\tilde{\gamma}(t)} \in \mathfrak{X}(\tilde{\gamma})$ is clearly extendable, we conclude that

$$\ell''(t) = \langle \nabla_{\tilde{\gamma}'(t)} \text{grad } h|_{\tilde{\gamma}(t)}, \tilde{\gamma}'(t) \rangle = \langle \text{Hess } h_{\tilde{\gamma}(t)}[\tilde{\gamma}'(t)], \tilde{\gamma}'(t) \rangle$$

and thus particularly $\ell''(0) = \langle \text{Hess } h_{K^*}[F], F \rangle$. Since F was arbitrary and K^* is nondegenerate, $\ell''(0) \geq 0$ implies that $\text{Hess } h_{K^*}$ is positive definite. Next, existence of a neighborhood at K^* on which $\text{Hess } h$ is positive definite follows by smoothness—in particular continuity—of the operator $\text{Hess } h_K$ in K . Finally, let $\tilde{\nabla}$ and $\bar{\nabla}$ denote the connections on $T\tilde{\mathcal{S}}$ induced, respectively, by the connections ∇ and $\bar{\nabla}$ on TS . Then by The Difference Tensor Lemma, the difference tensor between $\tilde{\nabla}$ and $\bar{\nabla}$ —defined as $D(U, V) := \tilde{\nabla}_U V - \bar{\nabla}_U V$ for any $U, V \in \mathfrak{X}(\tilde{\mathcal{S}})$ —is indeed a (1,2)-tensor field. That means, as $\text{grad } h_{K^*} = 0$,

$$\text{Hess } h_{K^*}[U_{K^*}] - \bar{\text{Hess}} h_{K^*}[U_{K^*}] = D(U, \text{grad } h)|_{K^*} = 0.$$

The last claim then follows as $U \in \mathfrak{X}(\tilde{\mathcal{S}})$ was arbitrary. \square

Proof of Theorem IV.3. By Lemma IV.2, $\text{grad } h_{K^*} = 0$ and there exists a neighborhood \mathcal{U} of K^* on which $\text{Hess } h_K$ is positive. Furthermore, by continuity of $\text{Hess } h$ (and, if necessary, shrinking \mathcal{U}) we can obtain constant positive scalars m and M such that for all $K \in \mathcal{U}$ and $G \in T_K \tilde{\mathcal{S}}$,

$$m |G|_{g_K}^2 \leq \langle \text{Hess } h_K[G], G \rangle \leq M |G|_{g_K}^2, \quad (16)$$

where $|\cdot|_{g_K}$ denotes the norm induced by g at K . In particular, if $G_t \in T_{K_t} \tilde{\mathcal{S}}$ is the Newton direction at some point $K_t \in \mathcal{U}$, then (by Cauchy-Schwartz inequality at K_t)

$$|G_t|_{g_{K_t}} \leq (1/m) |\text{grad } h_{K_t}|_{g_{K_t}}. \quad (17)$$

Next, define the curve $\gamma: [0, s_{K_t}] \mapsto \tilde{\mathcal{S}}$ with $\gamma(\eta) = K_t + \eta G_t$, and consider a smooth parallel vector field (with respect to the Riemannian connection) $E(\eta)$ along γ —refer to [29] for *parallel vector fields along curves and parallel transport*. Also, define $\phi: [0, s_{K_t}] \mapsto \mathbb{R}$ with $\phi(\eta) := \langle \text{grad } h_{\gamma(\eta)}, E(\eta) \rangle$. Notice that $\text{grad } h$ is smooth, so is ϕ and by compatibility with the metric and that $\text{grad } h_{\gamma(\eta)}$ is clearly extendable, we have

$$\phi'(\eta) = \langle D_\eta \text{grad } h_{\gamma(\eta)}, E(\eta) \rangle = \langle \text{Hess } h_{\gamma(\eta)}[G_t], E(\eta) \rangle,$$

where D_η is the covariant derivative along γ and G_t is extended to the vector field along γ with constant coordinates in the global coordinate frame. Thus, as

$$\phi(\eta) = \phi(0) + \eta \phi'(0) + \int_0^\eta [\phi'(\tau) - \phi'(0)] d\tau,$$

by direct substitution and the fact that G_t is the Newton direction at iteration t , we obtain that

$$\begin{aligned} \phi(\eta) &= (\eta - 1) \langle \text{Hess } h_{K_t}[G_t], E(0) \rangle \\ &\quad + \int_0^\eta \langle [\text{Hess } h_{\gamma(\tau)} - \mathcal{P}_{0,\tau}^\gamma \text{Hess } h_{\gamma(0)}] G_t, E(\tau) \rangle d\tau, \end{aligned}$$

where $\mathcal{P}_{0,\tau}^\gamma$ denotes the parallel transport from 0 to τ along γ . Again, as every parallel transport map along γ is a linear isometry we claim that

$$\begin{aligned} \langle \text{grad } h_{K_{t+1}}, E(\eta_t) \rangle &= (\eta_t - 1) \langle \mathcal{P}_{0,\eta_t}^\gamma \text{Hess } h_{K_t}[G_t], E(\eta_t) \rangle \\ &\quad + \int_0^{\eta_t} \langle \mathcal{P}_{\tau,\eta_t}^\gamma [\text{Hess } h_{\gamma(\tau)} - \mathcal{P}_{0,\tau}^\gamma \text{Hess } h_{\gamma(0)}] G_t, E(\eta_t) \rangle d\tau. \end{aligned}$$

Note that, for each $\tau \in [0, s_{K_t}]$, $\text{Hess } h_{\gamma(\tau)}$ is a self-adjoint operator that is smooth in τ as γ is. So, by (16), we obtain

$$|\mathcal{P}_{0,\eta_t}^\gamma \text{Hess } h_{K_t}[G_t]|_{g_{K_{t+1}}} \leq M |G_t|_{g_{K_t}}$$

and by smoothness there exist a constant $L > 0$ such that

$$|\mathcal{P}_{\tau,\eta_t}^\gamma [\text{Hess } h_{\gamma(\tau)} - \mathcal{P}_{0,\tau}^\gamma \text{Hess } h_{\gamma(0)}] G_t|_{g_{K_{t+1}}} \leq \tau L |G_t|_{g_{K_t}}^2$$

where we used the isometry of parallel transport again in obtaining the bounds. Therefore, by choosing the parallel vector field $E(\eta)$ along γ such that $E(\eta_t) = \text{grad } h_{K_{t+1}}$ we obtain that

$$\begin{aligned} |\text{grad } h_{K_{t+1}}|_{g_{K_{t+1}}} &\leq M |1 - \eta_t| |G_t|_{g_{K_t}} + (\eta_t^2 L/2) |G_t|_{g_{K_t}}^2 \\ &\leq \frac{M |1 - \eta_t|}{m} |\text{grad } h_{K_t}|_{g_{K_t}} + \frac{L \eta_t}{2m^2} |\text{grad } h_{K_t}|_{g_{K_t}}^2 \end{aligned} \quad (18)$$

where the last inequality follows by (17) and since $\eta_t \leq 1$. Next, let $F_{t+1} \in T_{K_{t+1}} \tilde{\mathcal{S}}$ be tangent vector that $\xi(\eta) = \text{exp}_{K_{t+1}}[\eta F_{t+1}]$ is the minimum-length geodesic in $\tilde{\mathcal{S}}$ joining $\xi(0) = K_{t+1}$ to $\xi(1) = K^*$, where exp denotes the exponential map on $\tilde{\mathcal{S}}$. This is certainly possible (by shrinking \mathcal{U} if necessary) because geodesics are locally-minimizing [29]. Similar to the function ϕ , define $\psi: [0, 1] \mapsto \mathbb{R}$ with

$$\psi(\eta) := \langle \text{grad } h_{\xi(\eta)}, E(\eta) \rangle,$$

for some parallel vector $E(\eta)$ along ξ . Then, similarly

$$\psi'(\eta) = \langle D_\eta \text{grad } h_{\xi(\eta)}, E(\eta) \rangle = \langle \text{Hess } h_{\xi(\eta)}[\xi'(\eta)], E(\eta) \rangle.$$

The velocity of any geodesic is a parallel vector field along itself, so by choosing $E(\eta) = \xi'(\eta)$ and using the fundamental lemma of calculus for ψ we obtain that

$$\psi(1) = \langle \text{grad } h_{K_{t+1}}, F_{t+1} \rangle + \int_0^1 \langle \text{Hess } h_{\xi(\tau)}[\xi'(\tau)], \xi'(\tau) \rangle d\tau$$

Note that $\psi(1) = 0$ and $|\xi'(\tau)|_{g_{\xi(\tau)}} = |F_{t+1}|_{g_{K_{t+1}}}$ for all τ as ξ is a geodesic. Thus, by using (16), we conclude that

$$m |F_{t+1}|_{g_{K_{t+1}}} \leq |\text{grad } h_{K_{t+1}}|_{g_{K_{t+1}}} \leq M |F_{t+1}|_{g_{K_{t+1}}} \quad (19)$$

where the last inequality follows by enlarging M when necessary. Finally, combining (18) and (19) at two iterations $t+1$ and t , and noticing $\text{dist}(K_{t+1}, K^*) = |F_{t+1}|_{g_{K_{t+1}}}$ imply that

$$\begin{aligned} \text{dist}(K_{t+1}, K^*) &\leq (|1 - \eta_t| M^2 / m^2) \text{dist}(K_t, K^*) \\ &\quad + (\eta_t L M^2 / 2m^3) \text{dist}(K_t, K^*)^2, \end{aligned} \quad (20)$$

where $\text{dist}(\cdot, \cdot)$ denotes the Riemannian distance function between two points. Next, note that the mapping $K \mapsto \mathcal{Q}_K$ is chosen to be smooth such that $\mathcal{Q}_K \succ 0$, therefore as a result of Lemma III.1, the mapping $K \mapsto \mathbb{L}(A_{\text{cl}}^\top, \mathcal{Q}_K)$ is smooth by composition. By smoothness (in particular continuity) of this mapping and the continuity of the maximum eigenvalue (utilized in the definition of stability certificate s_K in Lemma IV.1), we can shrink \mathcal{U} —if necessary—to obtain a positive constant $c > 0$ such that

$$s_{K_t} \geq c / |G_t|_{g_{K_t}} \geq cm / (M \text{dist}(K_t, K^*)), \quad (21)$$

where the last inequality follows by combining (17), (19) and the fact that $\text{dist}(K_t, K^*) = |F_t|_{g_{K_t}}$. Now, pick $r \in (0, 1)$; if we set $\mathcal{U}^* \subset \mathcal{U} \subset \tilde{\mathcal{S}}$ such that for any $K_0 \in \mathcal{U}^*$ we have

$$\text{dist}(K_0, K^*) < \min\left\{\frac{c(M/m)}{(M/m)^2 - r/2}, \frac{r}{(L/m)(M/m)^2}\right\},$$

then by the choice of stepsize $\eta_t = \min\{s_{K_t}, 1\}$ and the lower-bound in (21), we can claim that

$$|1 - \eta_0| M^2 / m^2 + (\eta_0 L M^2 / 2m^3) \text{dist}(K_0, K^*) < r.$$

But then, (20) implies that $\text{dist}(K_1, K^*) \leq r \text{dist}(K_0, K^*)$. Therefore, $K_1 \in \mathcal{U}^*$ as $r < 1$, and thus by induction we conclude a linear convergence rate to K^* . Consequently, (21) implies that $s_{K_t} \geq 1$ for large enough t , and thus by the choice of step-size, (20) simplifies to

$$\text{dist}(K_{t+1}, K^*) \leq (LM^2 / 2m^3) \text{dist}(K_t, K^*)^2,$$

guaranteeing a quadratic convergence rate. Finally, Lemma IV.2 implies that a critical point is nondegenerate with respect to the induced Riemannian connection on $T\tilde{\mathcal{S}}$ if and only if it is so with respect to the Euclidean one. The proof for QRNPO with $\overline{\text{Hess}}$ then follows similarly by redefining ϕ and ψ using the Euclidean metric under the usual identification of the tangent bundle. \square

Proof of Proposition V.1. By definition, $f: \mathcal{S} \mapsto \mathbb{R}$ can be viewed as the composition:

$$f: K \xrightarrow{\Phi} (A_{\text{cl}}^\top, K^\top RK + Q) \xrightarrow{\mathbb{L}} P_K \xrightarrow{\Psi} \frac{1}{2} \text{tr}[P_K \Sigma_1]. \quad (22)$$

Since the first and last maps are smooth (i.e., linear or quadratic in K), we conclude that $f \in C^\infty(\mathcal{S})$ by composition and Lemma III.1. For any $K \in \mathcal{S}$, we can compute its differential at K , denoted by df_K , using the chain rule:

$$df_K(E) = d\Psi_{P_K} \circ d\mathbb{L}_{(A_{\text{cl}}^\top, K^\top RK + Q)} \circ d\Phi_K(E),$$

for any $E \in T_K \mathcal{S}$. But Ψ is a linear map, and under the usual identification of the tangent bundle we obtain

$$d\Phi_K(E) = (E^\top B^\top, E^\top RK + K^\top RE).$$

Therefore, by Lemma III.1 we claim the followings

$$\begin{aligned} d(\mathbb{L} \circ \Phi)_K(E) &= \mathbb{L}(A_{\text{cl}}^\top, E^\top (B^\top P_K A_{\text{cl}} + RK) \\ &\quad + (K^\top R + A_{\text{cl}}^\top P_K B)E), \quad (23) \\ \implies df_K(E) &= \Psi \circ \mathbb{L}(A_{\text{cl}}^\top, E^\top (B^\top P_K A_{\text{cl}} + RK) \\ &\quad + (K^\top R + A_{\text{cl}}^\top P_K B)E). \end{aligned}$$

Thus, $df_K(E) = \langle E, RK + B^\top P_K A_{\text{cl}} \rangle$ with $Y_K = \mathbb{L}(A_{\text{cl}}, \Sigma_1)$ —by Lyapunov-trace property—which is well-defined and unique as A_{cl} is a stability matrix. As $df_K(E) = Ef$, the expression for $\text{grad} f \in \mathfrak{X}(\mathcal{S})$ then follows by its definition. Next, as the Hessian operator is self-adjoint (see Appendix), in order to obtain $\text{Hess} f$ we can compute (12) for any $U, W \in \mathfrak{X}(\mathcal{S})$. As $\text{Hess} f[U]_K$ only depends on the value of U at K , it suffices to obtain $\text{Hess} f_K[U_K]$ at each $K \in \mathcal{S}$ with $U_K = E$ for arbitrary $E \in T_K \mathcal{S}$. To do so, we compute $\langle \text{Hess} f_K[E], F \rangle$ for an arbitrary vector $F \in T_K \mathcal{S}$ by extending F to the vector field W along the curve $\gamma: t \mapsto K + tE$ with constant coordinates with respect to the global coordinate frame $(\partial_{(i,j)})$. As $\langle \text{Hess} f[U], W \rangle|_K$ only depends on the value of $W_K = F$ and $U_K = E$, how these vector fields have been extended is arbitrary. By properties of the Riemannian connection and the fact that W can be extended with constant coordinates, here we can compute $\nabla_U W|_K$ in the global coordinate frame $(\partial_{i,j})$ and obtain,

$$\nabla_U W|_K = [E]^{k,\ell} [F]^{p,q} \Gamma_{(k,\ell)(p,q)}^{i,j}(K) \partial_{i,j}, \quad (24)$$

where $\Gamma_{(k,\ell)(p,q)}^{i,j}(K)$ denotes the Christoffel symbols associated with the Riemannian metric g at the point $K \in \mathcal{S}$. Therefore, from (12) we have that

$$\langle \text{Hess} f[U], W \rangle|_K = Er - \langle \text{grad} f_K, \nabla_U W|_K \rangle, \quad (25)$$

where $r := \langle \text{grad} f, W \rangle \in C^\infty(\mathcal{S})$. By the expression obtained for $\text{grad} f$ and that W has constant coordinates, the mapping $K \mapsto r(K)$ can be decomposed as:

$$\begin{aligned} K &\xrightarrow{\text{Id} \times \text{Id}} (K, K) \xrightarrow{\text{Id} \times (\mathbb{L} \circ \Phi)} (K, P_K) \xrightarrow{\Xi} \\ &\quad (A_{\text{cl}}^\top, (\text{grad} f_K)^\top F + F^\top \text{grad} f_K) \xrightarrow{\Psi \circ \mathbb{L}} r(K), \end{aligned}$$

where we used the Lyapunov-trace property and invariance of trace under transpose to justify the last mapping. Also note that Φ and Ψ are defined in (22) and $\Xi: \mathbb{M}(m \times n, \mathbb{R}) \times \mathbb{M}(n \times n, \mathbb{R}) \mapsto \mathbb{M}(n \times n, \mathbb{R}) \times \mathbb{M}(n \times n, \mathbb{R})$ is defined as above. Therefore, under the usual identification of tangent bundle, for any $(E, G) \in T_{(K, P_K)}(\mathbb{M}(m \times n, \mathbb{R}) \times \mathbb{M}(n \times n, \mathbb{R}))$ we can compute,

$$\begin{aligned} d\Xi_{(K, P_K)}[E, G] &= \left(E^\top B^\top, E^\top (R + B^\top P_K B)F + A_{\text{cl}}^\top G B F \right. \\ &\quad \left. + F^\top (R + B^\top P_K B)E + F^\top B^\top G^\top A_{\text{cl}} \right). \quad (26) \end{aligned}$$

Therefore, by the chain rule, for any $E \in T_K \mathcal{S}$ we have

$$dr_K(E) = \Psi \circ d\mathbb{L} \circ d\Xi \circ (E, d(\mathbb{L} \circ \Phi)_K(E)),$$

where the base points of differentials are understood and dropped for brevity. But then, by (23) and (26), we obtain,

$$\begin{aligned} d\Xi[E, d(\mathbb{L} \circ \Phi)_K(E)] &= \left(E^\top B^\top, E^\top (R + B^\top P_K B)F \right. \\ &\quad \left. + A_{\text{cl}}^\top (S_K[E])BF + F^\top (R + B^\top P_K B)E + F^\top B^\top (S_K[E])A_{\text{cl}} \right), \end{aligned}$$

where $S_K[E]$ is defined in the premise. Therefore,

$$\begin{aligned} d r_K(E) &= \Psi \circ \mathbb{L} \left(A_{\text{cl}}^\top, E^\top (R + B^\top P_K B) F \right. \\ &\left. + A_{\text{cl}}^\top (S_K[E]) B F + A_{\text{cl}}^\top (S_K[F]) B E + E^\top B^\top (S_K[F]) A_{\text{cl}} \right), \end{aligned}$$

that using the Lyapunov-trace property can be simplified as,

$$\begin{aligned} 2 d r_K(E) &= \text{tr} [E^\top B^\top (S_K[F]) A_{\text{cl}} Y_K + A_{\text{cl}}^\top (S_K[F]) B E Y_K] \\ &\quad + \text{tr} [(E^\top (R + B^\top P_K B) + A_{\text{cl}}^\top (S_K[E]) B) F Y_K] \\ &\quad + \text{tr} [F^\top ((R + B^\top P_K B) E + B^\top (S_K[E]) A_{\text{cl}}) Y_K], \end{aligned}$$

where $Y_K = \mathbb{L}(A_{\text{cl}}, \Sigma_1)$. Noting that Y_K , P_K , $S_K[E]$ and $S_K[F]$ are all symmetric, using the cyclic permutation property of trace, we now obtain,

$$\begin{aligned} d r_K(E) &= \langle (R + B^\top P_K B) E + B^\top (S_K[E]) A_{\text{cl}}, F \rangle \\ &\quad + \langle B^\top (S_K[F]) A_{\text{cl}}, E \rangle. \quad (27) \end{aligned}$$

Then, the expression for Hess f follows by substituting (27) and (24) in (25). Finally, the expression of $\overline{\text{Hess}} f$ can be obtained similarly by threading through the definitions. \square

Proof of Corollary V.2. Smoothness of h and the expression of its gradient follows immediately by Proposition III.5 and Proposition V.1. In order to compute Hess h_K , we can combine its extrinsic representation as obtained in Proposition III.5 with (25), and use the definition of Weingarten map to obtain,

$$\begin{aligned} \langle \text{Hess } h_K[E], F \rangle &= \langle \pi^\top (\text{Hess } f|_{\mathcal{S}})|_{\mathcal{S}}, W \rangle|_K + \langle \mathbb{W}_{\pi^\perp(\text{grad } f|_{\mathcal{S}})}(U), W \rangle|_K \\ &= E r - \langle \text{grad } f_K, \nabla_U W|_K \rangle + \langle \pi^\perp \text{grad } f_K, \pi^\perp \nabla_U W|_K \rangle \\ &= E r - \langle \pi^\top \text{grad } f_K, \nabla_U W|_K \rangle, \end{aligned}$$

for any $E, F \in T_K \tilde{\mathcal{S}} \subset \mathcal{S}$, which are extended to vector fields on a neighborhood in \mathcal{S} with constant coordinates with respect to the global coordinate frame. The claimed expression of Hess h_K then follows by substituting (24) and (27) into the last expression. \square

Proof of SLQR Manifold in §V-A.2. Let the tuple $(x^{(i,j)})_{(i,j) \in [m] \times [n]}$ denote the component functions of the global smooth chart $(\mathbb{M}(m \times n, \mathbb{R}), \text{vec})$, and define $\Phi: \mathbb{M}(m \times n, \mathbb{R}) \mapsto \mathbb{R}^{mn-|D|}$ with $\Phi(K) = \sum_{(i,j) \notin D} [K]_{i,j} x^{(i,j)}$. Then, it is easy to see that Φ is a smooth submersion, and so is $\Phi|_{\mathcal{S}}$ as \mathcal{S} is an open submanifold of $\mathbb{M}(m \times n, \mathbb{R})$. Therefore, as $\tilde{\mathcal{S}} = \mathcal{S} \cap \mathcal{K}_D = (\Phi|_{\mathcal{S}})^{-1}(0)$, by Submersion Level Set Theorem, we conclude that $\tilde{\mathcal{S}}$ is a properly embedded submanifold of dimension $|D|$.

Furthermore, at any point $K \in \tilde{\mathcal{S}}$ and for any tangent vector $E \in T_K \tilde{\mathcal{S}}$, we can compute the tangential projection $\pi^\top: T_K \mathcal{S} \mapsto T_K \tilde{\mathcal{S}}$ as,

$$\pi^\top E = \arg \min_{\tilde{E} \in T_K \tilde{\mathcal{S}}} \langle E - \tilde{E}, E - \tilde{E} \rangle. \quad (28)$$

Since \mathcal{K}_D is a linear subspace of $\mathbb{M}(m \times n, \mathbb{R})$, we can identify $T_K \tilde{\mathcal{S}}$ with \mathcal{K}_D itself (a dimension argument). Then, it follows that the unique solution \tilde{E}^* to the minimization

problem above (with linear constraint and strongly convex cost function, as $Y_K \succ 0$ in Proposition III.3), satisfies $E - \tilde{E}^* \perp \mathcal{K}_D$ with respect to the Riemannian metric at K ; or equivalently (10). \square

Proof of OLQR Manifold in §V-A.3. Note that \mathcal{K}_C is a linear subspace of $\mathbb{M}(m \times n, \mathbb{R})$ whose dimension depends on the rank of C . Now, define $\Psi: \mathbb{M}(m \times n, \mathbb{R}) \mapsto \mathbb{M}(m \times n, \mathbb{R})$ as, $\Psi(K) = K(I_n - C^\dagger C)$, where \dagger denotes the Moore-Penrose inverse. Note that Ψ is a linear map that is surjective onto its range, denoted by \mathcal{R} , which is an $m(n-d)$ dimensional linear subspace of $\mathbb{M}(m \times n, \mathbb{R})$. Therefore, $\Phi: \mathcal{S} \mapsto \mathcal{R}$ defined as the restriction of Ψ both in domain and codomain, is a smooth submersion. Finally, as $CC^\dagger C = C$, we can observe that $\text{Ker}(\Psi) = \mathcal{K}_C$. Therefore, $\tilde{\mathcal{S}} = \mathcal{K}_C \cap \mathcal{S} = \Phi^{-1}(0)$ and thus, by Submersion Level Set Theorem, $\tilde{\mathcal{S}}$ is a properly embedded submanifold of \mathcal{S} with dimension md . We also conclude that at each $K \in \tilde{\mathcal{S}}$, we can canonically identify the tangent space at K as, $T_K \tilde{\mathcal{S}} = \text{Ker}(d\Phi_K) \cong \mathcal{K}_C$.

Next, at any $K \in \tilde{\mathcal{S}}$ and for any $E \in T_K \mathcal{S}$, the tangential projection of E , denoted by $\tilde{E}^* = \pi^\top E$, is the unique solution of a minimization problem similar to (28). Moreover, under the above identification, $E - \tilde{E}^* \perp \mathcal{K}_C$ (with respect to the Riemannian metric) must be satisfied, or equivalently, $\text{tr} [C^\top L^\top (E - \tilde{E}^*) Y_K] = 0, \forall L \in \mathbb{M}(m \times d, \mathbb{R})$. Here, $Y_K = \mathbb{L}(A_{\text{cl}}, \Sigma_1)$ is positive definite and since C is assumed to be full-rank, $CY_K C^\top$ is positive definite. Hence, we conclude that $\pi^\top E = L^* C$ with $L^* \in \mathbb{M}(m \times d, \mathbb{R})$ being the unique solution of (11).

Finally, at each point $K \in \tilde{\mathcal{S}}$, we denote the global coordinate functions of $\mathbb{M}(m \times d, \mathbb{R})$ (with slight abuse of notation) by the tuple $(x^{i,j})$ for $(i,j) \in D := [m] \times [d]$ and its corresponding global coordinate frame by $(\partial_{(i,j)})$. Recall that C has full-rank and consider the identification of $T_K \tilde{\mathcal{S}} \cong \mathcal{K}_C$ described above. Then, the (constant) global vector fields $(\tilde{\partial}_{(i,j)} = \partial_{(i,j)} C)$ with $(i,j) \in D$, form a global smooth frame for $T\tilde{\mathcal{S}}$ as they are linearly independent on $\tilde{\mathcal{S}}$. Therefore, the coordinates of the covariant Hessian $h_{(k,\ell)(p,q)}(K)$ with respect to this frame can be computed by substituting $E = \partial_{(k,\ell)} C$ and $F = \partial_{(p,q)} C$ in Corollary V.2 for each $(k,\ell), (p,q) \in D$ —similar to the SLQR case. It is worth noting that the sparsity pattern in E, F and Christoffel symbols can simplify the computation; we will not delve further into this issue due to space limitations. \square

REFERENCES

- [1] H. Mohammadi, A. Zare, M. Soltanolkotabi, and M. R. Jovanović, "Convergence and sample complexity of gradient methods for the model-free linear-quadratic regulator problem," *IEEE Transactions on Automatic Control*, vol. 67, no. 5, pp. 2435–2450, 2021.
- [2] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control*," *IEEE Transactions on Automatic Control*, vol. 51, no. 2, pp. 274–286, 2006.
- [3] J. Ackermann, "Parameter space design of robust control systems," *IEEE Transactions on Automatic Control*, vol. 25, no. 6, pp. 1058–1072, 1980.

- [4] 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)*, pp. 1430–1437, 2019.
- [5] W. Levine and M. Athans, "On the determination of the optimal constant output feedback gains for linear multivariable systems," *IEEE Transactions on Automatic Control*, vol. 15, no. 1, pp. 44–48, 1970.
- [6] B. Anderson and J. Moore, *Linear Optimal Control*. Englewood Cliffs, New Jersey, Prentice–Hall, 1971.
- [7] D. Moerder and A. Calise, "Convergence of a numerical algorithm for calculating optimal output feedback gains," *IEEE Transactions on Automatic Control*, vol. 30, no. 9, pp. 900–903, 1985.
- [8] H. T. Toivonen, "A globally convergent algorithm for the optimal constant output feedback problem," *International Journal of Control*, vol. 41, no. 6, pp. 1589–1599, 1985.
- [9] P. Makila and H. Toivonen, "Computational methods for parametric LQ problems—a survey," *IEEE Transactions on Automatic Control*, vol. 32, no. 8, pp. 658–671, 1987.
- [10] H. T. Toivonen and P. M. Makila, "Newton's method for solving parametric linear quadratic control problems," *International Journal of Control*, vol. 46, no. 3, pp. 897–911, 1987.
- [11] T. Iwasaki, R. Skelton, and J. Geromel, "Linear quadratic suboptimal control with static output feedback," *Systems & Control Letters*, vol. 23, no. 6, pp. 421–430, 1994.
- [12] T. Rautert and E. W. Sachs, "Computational design of optimal output feedback controllers," *SIAM Journal on Optimization*, vol. 7, no. 3, pp. 837–852, 1997.
- [13] K. Mårtensson and A. Rantzer, "Gradient methods for iterative distributed control synthesis," in *Proceedings of the IEEE Conference on Decision and Control*, pp. 549–554, IEEE, 2009.
- [14] V. L. Syrmos, C. T. Abdallah, P. Dorato, and K. Grigoriadis, "Static output feedback—a survey," *Automatica*, vol. 33, no. 2, pp. 125–137, 1997.
- [15] V. Blondel and J. N. Tsitsiklis, "NP-hardness of some linear control design problems," *SIAM journal on control and optimization*, vol. 35, no. 6, pp. 2118–2127, 1997.
- [16] C. H. Papadimitriou and J. Tsitsiklis, "Intractable problems in control theory," *SIAM Journal on Control and Optimization*, vol. 24, no. 4, pp. 639–654, 1986.
- [17] J. Bu, A. Mesbahi, M. Fazel, and M. Mesbahi, "LQR through the lens of first order methods: Discrete-time case," *arXiv preprint arXiv:1907.08921*, 2019.
- [18] J. Bu, A. Mesbahi, and M. Mesbahi, "Policy gradient-based algorithms for continuous-time linear quadratic control," *arXiv preprint arXiv:2006.09178*, 2020.
- [19] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi, "Global convergence of policy gradient methods for the linear quadratic regulator," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80, pp. 1467–1476, PMLR, 10–15 Jul 2018.
- [20] 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.
- [21] H. Mohammadi, M. Soltanolkotabi, and M. R. Jovanovic, "On the linear convergence of random search for discrete-time LQR," *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 989–994, 2021.
- [22] F. Zhao, K. You, and T. Başar, "Global convergence of policy gradient primal–dual methods for risk-constrained LQRs," *IEEE Transactions on Automatic Control*, vol. 68, no. 5, pp. 2934–2949, 2023.
- [23] Y. Park, R. Rossi, Z. Wen, G. Wu, and H. Zhao, "Structured policy iteration for linear quadratic regulator," in *International Conference on Machine Learning*, pp. 7521–7531, PMLR, 2020.
- [24] D. Gabay, "Minimizing a differentiable function over a differential manifold," *Journal of Optimization Theory and Applications*, vol. 37, no. 2, pp. 177–219, 1982.
- [25] S. T. Smith, "Optimization techniques on Riemannian manifolds," *Fields Institute Communications*, vol. 3, no. 3, pp. 113–135, 1994.
- [26] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2009.
- [27] S.-I. Amari, "Natural gradient works efficiently in learning," *Neural Computation*, vol. 10, no. 2, pp. 251–276, 1998.
- [28] S.-I. Amari and S. Douglas, "Why natural gradient?," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 2, pp. 1213–1216 vol.2, 1998.
- [29] J. M. Lee, *Introduction to Riemannian Manifolds*. Cham, Switzerland: Springer Nature, 2nd ed., 2018.
- [30] S. Talebi and M. Mesbahi, "Riemannian constrained policy optimization via geometric stability certificates," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 1472–1478, 2022.
- [31] Z. Gajic and M. T. J. Qureshi, *Lyapunov Matrix Equation in System Stability and Control*. Courier Corporation, 2008.
- [32] J. M. Lee, *Introduction to Smooth Manifolds*. Springer, 2nd ed., 2013.
- [33] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and design*. Hoboken, NJ, USA: Wiley-Blackwell, 2nd ed., 2005.
- [34] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multi-agent Networks*. Princeton University Press, 2010.
- [35] 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.
- [36] D. Serre, *Matrices: Theory and Applications*. Springer Science and Media, 2nd ed., 2010.
- [37] J. Lee, *Introduction to Topological Manifolds*. Springer Science & Business Media, 2nd ed., 2010.
- [38] A. Ohara and S.-I. Amari, "Differential geometric structures of stable state feedback systems with dual connections," *IFAC Proceedings*, vol. 25, no. 21, pp. 176–179, 1992.
- [39] P. A. Absil, R. Sepulchre, P. Van Dooren, and R. Mahony, "Cubically convergent iterations for invariant subspace computation," *SIAM Journal on Matrix Analysis and App.*, vol. 26, no. 1, pp. 70–96, 2004.
- [40] J. Müller and G. Montúfar, "Geometry and convergence of natural policy gradient methods," *Information Geometry*, pp. 1–39, 2023.
- [41] F. Alvarez, J. Bolte, and O. Brahic, "Hessian Riemannian gradient flows in convex programming," *SIAM Journal on Control and Optimization*, vol. 43, no. 2, pp. 477–501, 2004.
- [42] G. Hewer, "An iterative technique for the computation of the steady state gains for the discrete optimal regulator," *IEEE Transactions on Automatic Control*, vol. 16, no. 4, pp. 382–384, 1971.
- [43] J. P. Hespanha, *Linear Systems Theory*. Princeton Univ. Press, 2018.
- [44] G. C. Goodwin, S. F. Graebe, and M. E. Salgado, *Control System Design*. Upper Saddle River, NJ: Prentice Hall, 2001.
- [45] K. Mårtensson, *Gradient Methods for Large-Scale and Distributed Linear Quadratic Control*. PhD thesis, Lund University, 2012.
- [46] S. Talebi, S. Alemzadeh, N. Rahimi, and M. Mesbahi, "On regularizability and its application to online control of unstable LTI systems," *IEEE Transactions on Automatic Control*, vol. 67, no. 12, pp. 6413–6428, 2022.
- [47] S. Talebi and M. Mesbahi, "Policy optimization over submanifolds for constrained feedback synthesis," *arXiv preprint arXiv:2201.11157*, 2022.
- [48] S. Talebi and M. Mesbahi, "Quasi Riemannian Newton Policy Optimization (QRNPO)," 2022. Available on GitHub at <https://github.com/shahriarta/QRNPO>.
- [49] J. E. Dennis Jr and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Society for Industrial and Applied Mathematics, 1996.
- [50] J. Nocedal and S. Wright, *Numerical Optimization*. Springer Science & Business Media, 2006.
- [51] S. G. Krantz and H. R. Parks, *A Primer of Real Analytic Functions*. Springer Science & Business Media, 2002.
- [52] J. Milnor, *Morse Theory*, vol. 51. Princeton University Press, 1963.



Shahriar Talebi (Student Member, IEEE) received the Ph.D. degree in aeronautics and astronautics, specializing in control theory, and the M.Sc. degree in Mathematics, focusing on differential geometry, both from the University of Washington, Seattle, WA, USA, in 2023. He also received the B.Sc. degree from Sharif University of Technology, Tehran, Iran, in 2014, and the M.Sc. degree from the University of Central Florida, Orlando, FL, in 2017, both in electrical engineering.

His research interests include control theory, differential geometry, learning for control, networked dynamical systems, and game theory.

Dr. Talebi was the recipient of the 2022 Excellence in Teaching Award at UW. He is also a recipient of a number of scholarships, including the William E. Boeing Endowed Fellowship, Paul A. Carlstedt Endowment, and Latvian Arctic Pilot-A. Vagners Memorial Scholarship (UW, 2018–2019), as well as the Frank Hubbard Engineering Scholarship (UCF, 2017).



Mehran Mesbahi (Fellow, IEEE) received his Ph.D. degree in electrical engineering-systems from University of Southern California, Los Angeles, CA, USA, in 1996.

M. Mesbahi was a member of the Guidance, Navigation, and Analysis group at JPL, Pasadena, CA, USA, from 1996–2000 and an Assistant Professor of Aerospace Engineering and Mechanics at the University of Minnesota from 2000–2002. He is currently a Professor of Aeronautics and Astronautics, Adjunct Pro-

fessor of Electrical and Computer Engineering and Mathematics at the University of Washington, and the Executive Director of the Joint Center for Aerospace Technology Innovation. He was the recipient of NSF CAREER Award, NASA Space Act Award, UW Distinguished Teaching Award, UW College of Engineering Innovator Award for Teaching, and is a member of Washington State Academy of Sciences.

His research interests include distributed and networked aerospace systems, autonomy, and system and control theory.