

# Convex Structured Controller Design in Finite Horizon

Krishnamurthy Dvijotham, Emanuel Todorov, and Maryam Fazel

**Abstract**—We consider the problem of synthesizing optimal linear feedback policies subject to arbitrary convex constraints on the feedback matrix. This is known to be a hard problem in the usual formulations ( $\mathcal{H}_2$ ,  $\mathcal{H}_\infty$ , LQR) and previous works have focused on characterizing classes of structural constraints that allow an efficient solution through convex optimization or dynamic programming techniques. In this paper, we propose a new control objective for finite horizon discrete-time problems and show that this formulation makes the problem of computing optimal linear feedback matrices convex under arbitrary convex constraints on the feedback matrix. This allows us to solve problems in decentralized control (sparsity in the feedback matrices), control with delays, and variable impedance control. Although the control objective is nonstandard, we present theoretical and empirical evidence showing that it agrees well with standard notions of control. We show that the theoretical approach carries over to nonlinear systems, although the computational tractability of the extension is not investigated in this paper. We present numerical experiments validating our approach.

**Index Terms**—Convex functions, decentralized control, optimal control, optimization methods.

## I. INTRODUCTION

LINEAR feedback control synthesis is a classical topic in control theory and has been extensively studied in the literature. From the perspective of stochastic optimal control theory, the classical result is the existence of an optimal linear feedback controller for systems with linear dynamics, quadratic costs, and Gaussian noise (LQG systems) that can be computed via dynamic programming [1]. However, if one imposes additional constraints on the feedback matrix (such as a sparse structure arising from the need to implement control in a decentralized fashion), the dynamic programming approach is no longer applicable. In fact, it has been shown that the optimal control policy may not even be linear [2] and that the general problem of designing linear feedback gains subject to constraints is NP-hard [3].

Manuscript received September 20, 2013; revised September 29, 2013 and April 7, 2014; accepted September 19, 2014. Date of publication November 5, 2014; date of current version March 13, 2015. This work was supported by the National Science Foundation. Recommended by Associate Editor A. Nedich.

K. Dvijotham is with the Center for Mathematics of Information, Caltech, Pasadena, CA 91125 USA (e-mail: dvij@cs.washington.edu).

E. Todorov is with the Department of Computer Science and Engineering and Applied Mathematics at the University of Washington, Seattle, WA 98195 USA (e-mail: todorov@cs.washington.edu).

M. Fazel is with the Department of Electrical Engineering at the University of Washington, Seattle, WA 98195 USA (e-mail: mfazel@uw.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCNS.2014.2367359

Previous approaches to synthesizing structured controllers either try to find exact solutions for special cases [4]–[10] or use heuristic methods (based on nonconvex optimization) [11]–[13] which lack convergence or optimality guarantees. In this paper, we take a different approach: We reformulate the structured control problem using a family of new control objectives for finite horizon discrete-time problems (Section IV). We develop these new objectives as follows: Finite horizon discrete-time versions of the  $\mathcal{H}_2$  and  $\mathcal{H}_\infty$  norms can be expressed as functions of singular values of the linear mapping from disturbance trajectories to state trajectories. This mapping is a highly nonlinear function of the feedback gains. However, the inverse of this mapping has a simple linear dependence on the feedback gains. Further, the determinant of the mapping has a fixed value independent of the closed-loop dynamics—this is, in fact, a finite horizon version of Bode’s sensitivity integral and has been studied in [14]. By exploiting both of these facts, we develop upper bounds on the  $\mathcal{H}_2$ ,  $\mathcal{H}_\infty$  norms in terms of the singular values of the inverse mapping. We show that these upper bounds have several properties that make them desirable control objectives. For the new family of objectives, we show that the resulting problem of designing an optimal linear state feedback matrix, under arbitrary convex constraints, is convex (Section V). Further, we prove suboptimality bounds on how the solutions of the convex problems compare to the optima of the original problem. Our approach is directly formulated in state-space terminology and does not make any reference to frequency-domain concepts. Thus, it applies directly to time-varying systems. Since the formulation presented in this paper is in finite horizon, we do not make any stability guarantees. Indeed, controllers synthesized by this formulation may not be stabilizing. However, we show numerically that they achieve good performance on the finite horizon objectives and outperform standard nonlinear optimization approaches applied directly to the original nonconvex objectives (Section VII). We also present a generalization of our approach to nonlinear systems (Section VIII). The resulting formulation, albeit convex, is not computationally tractable for general nonlinear systems. We leave computational aspects of the nonlinear formulation for future work.

## II. DISCUSSION AND RELATED WORK

There have been three major classes of prior work in synthesizing structured controllers: frequency-domain approaches, dynamic programming, and nonconvex optimization approaches. We compare the relative merits of the different approaches in this section.

In frequency-domain approaches, problems are typically formulated as

$$\underset{K}{\text{Minimize}} \quad \|\text{Closed loop system with feedback } K\|$$

Subject to  $K$  Stabilizing,  $K \in \mathcal{C}$

where  $\|\cdot\|$  is typically the  $\mathcal{H}_2$  or  $\mathcal{H}_\infty$  norm. In general, these are solved by reparameterizing the problem in terms of a Youla parameter (via a nonlinear transformation), and imposing special conditions on  $\mathcal{C}$  (like quadratic invariance) that guarantee that the constraints  $\mathcal{C}$  can be translated into convex constraints on the Youla parameter [5], [6]. There are multiple limitations of these approaches:

- 1) Only specific kinds of constraints can be imposed on the controller. Many of the examples have the restriction that the structure of the controller mirrors that of the plant.
- 2) They result in infinite dimensional convex programs in general. One can solve them using a sequence of convex programming problems, but these approaches are susceptible to numerical issues and the degree of the resulting controllers may be ill-behaved, leading to practical problems in terms of implementing them.
- 3) The approaches rely on frequency-domain notions and cannot handle time-varying systems. In the special case of poset-causal systems (where the structure of the plant and controller can be described in terms of a partial order [7]), the problem can be decomposed when the performance metric is the  $\mathcal{H}_2$  norm and explicit state-space solutions are available by solving Ricatti equations for subsystems and combining the results. For the  $\mathcal{H}_\infty$  norm, a state-space solution using an LMI approach was developed in [15].

Another thread of work on decentralized control looks at special cases where dynamic programming techniques can be used despite the decentralization constraints. The advantage of these approaches is that they directly handle finite horizon and time-varying systems. For the LEQG cost-criterion, a dynamic programming approach was developed in [8] for the case of 1-step delay in a 2-agent decentralized control problem. In [9], the authors show that the case of 2 agents can be solved via dynamic programming. In [10], the authors develop a dynamic programming solution that generalizes this and applies it to general ‘‘partially-nested’’ systems allowing for both sparsity and delays.

All of the aforementioned methods work for special structures on the plant and controller (quadratic invariance/partial nestedness) under which decentralized controllers can be synthesized using either convex optimization or dynamic programming methods.

In [16], the authors pose decentralized control (in the discrete-time, finite horizon, linear quadratic setting) as a rank-constrained semidefinite programming problem. By dropping the rank constraint, one can obtain a convex relaxation of the problem. The relaxed problem provides a solution to the original problem only when the relaxed problem has a rank-1 solution. However, it is unknown when this can be guaranteed,

and how a useful controller can be recovered from a higher-rank solution. Further, the SDP posed in this paper grows very quickly with the problem dimension, leading to computational difficulties.

Our work differs from these previous works in one fundamental way: Rather than looking for special decentralization structures that can be solved tractably under standard control objectives, we formulate a new control objective that helps us solve problems with *arbitrary* decentralization constraints. In fact, we can handle *arbitrary convex constraints*-decentralization constraints that impose a sparsity pattern on  $\mathbf{K}$  and are a special case of this. We can also handle time-varying linear systems. Although the objective is nonstandard, we provide theoretical (Section IV-A) and numerical (Section VII) evidence that it is a sensible control objective. The only other approaches that handle all of these problems are nonconvex approaches [11], [12], [17], which lack convergence rate and optimality guarantees. We show (Section VII) that our approach outperforms a standard nonconvex approach, both in terms of performance of the resulting controller and in computational time.

We also believe that this was the first approach to exploit a fundamental limitation (Bode’s sensitivity integral) to develop efficient control design algorithms. The fact that the spectrum of the input output map satisfies a conservation law (the sum of the logs of singular values is fixed) is a limitation which says that reducing some of the singular values is bound to increase the others. However, this limitation allows us to approximate the difficult problem of minimizing the finite horizon  $\mathcal{H}_2$  or  $\mathcal{H}_\infty$  norm with the easier problem of minimizing a convex surrogate, leading to an efficient solution.

The work presented here is an extension of a recent conference publication [18]. However, this paper contains a significant reformulation of the results presented there and also has new results. The conference paper [18] was formulated in terms of the eigenvalues of the covariance matrix of trajectories. In this paper, we look at the linear map from noise to state trajectories (denoted by  $F$  in this paper), which is a Cholesky factor of the covariance matrix. This leads to simpler proofs of convexity and a more general class of objectives. For example, the nuclear norm (and more generally Ky-Fan norms) of  $F$  is a valid objective in our formulation while it was not in [18], since it is equal to the sum of square roots of eigenvalues of the covariance matrix which is a nonconvex function. Further, we provide an analysis quantifying how well the solutions to our convex objectives perform in terms of the original nonconvex  $\mathcal{H}_2, \mathcal{H}_\infty$  objectives. We present numerical results comparing results of our formulation to other nonconvex approaches for structured controller synthesis.

### III. NOTATION

Let  $\lambda_{\max}(M)$  denote the maximum eigenvalue of an  $l \times l$  symmetric matrix  $M$ ,  $\lambda_{\min}(M)$  the minimum eigenvalue and  $\lambda_i(M)$  the  $i$ th eigenvalue in descending order

$$\lambda_l(M) = \lambda_{\min}(M) \leq \lambda_{l-1}(M) \leq \dots \leq \lambda_{\max}(M) = \lambda_1(M).$$

Similarly, singular values of a general rank  $l$  matrix  $M$  are

$$\sigma_l(M) = \sigma_{\min}(M) \leq \sigma_{l-1}(M) \leq \dots \leq \sigma_{\max}(M) = \sigma_1(M).$$

We use  $I$  to denote the identity matrix. For  $z \in \mathbf{R}^n$

$$\begin{aligned} \text{Var}(z) &= \frac{1}{n} \sum_{i=1}^n \left( z_i - \frac{\sum_{i=1}^n z_i}{n} \right)^2 \\ z_{[i]} &= i\text{-th largest component of } z \\ |z| &= (|z_1| \quad |z_2| \quad \dots \quad |z_n|)^T \end{aligned}$$

$\mathcal{N}(\mu, \Sigma)$  denotes a Gaussian distribution with mean  $\mu$  and covariance matrix  $\Sigma$ .

#### IV. PROBLEM FORMULATION

Consider a finite-horizon discrete-time linear system in state-space form

$$x_1 = D_0 w_0$$

$$x_{t+1} = A_t x_t + B_t u_t + D_t w_t, \quad t = 1, 2, \dots, N-1.$$

Here,  $t = 0, 1, 2, \dots, N$  is the discrete time index,  $x_t \in \mathbf{R}^n$  is the plant state,  $w_t \in \mathbf{R}^n$  is an exogenous disturbance, and  $u_t \in \mathbf{R}^{n_u}$  is the control input. We employ static state feedback  $u_t = K_t x_t$ . Let  $K = \{K_t : t = 1, 2, \dots, N-1\}$  and denote the closed-loop system dynamics by  $A_t(K_t) = A_t + B_t K_t$ . For any linear system of this form, there is a linear mapping between disturbance trajectories  $\mathbf{w} = (w_0 w_1 \dots w_{N-1})$  and state trajectories  $(x_1 x_2 \dots x_N)$ . This will play a key role in our paper, and we denote it by

$$\mathbf{x} = F(\mathbf{K})\mathbf{w},$$

where  $F(\mathbf{K}) =$

$$\begin{bmatrix} D_0 & 0 & \dots & 0 \\ \tilde{A}_1 D_0 & D_1 & \dots & 0 \\ \tilde{A}_2 \tilde{A}_1 D_0 & \tilde{A}_2 D_1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ \prod_{\tau=1}^{N-1} \tilde{A}_{N-\tau} D_0 & \prod_{\tau=2}^{N-1} \tilde{A}_{N-\tau} D_1 & \dots & D_{N-1} \end{bmatrix}.$$

Note that this holds for *any* discrete-time linear system of the form described before and does not assume any special structure in the system dynamics.

Our formulation differs from standard control formulations in the following ways:

- 1) We assume that the controller performance is measured in terms of the norm of the system trajectory  $\mathbf{x}^T \mathbf{x}$  (see the Appendix for an extension that includes control costs).
- 2) As mentioned earlier, we restrict ourselves to have static state feedback  $u_t = K_t x_t$  (The Appendix discusses dynamic output feedback.)
- 3) We assume that  $D_t$  is square and invertible. The assumption of invertibility is required because our convex approximations will be based on  $F(\mathbf{K})^{-1}$ , which does not exist unless  $D_t$  is invertible. Effectively, this means that it is possible to uniquely recover the disturbance trajectory from the state trajectory.

Finite-horizon versions of the  $\mathcal{H}_2$  and  $\mathcal{H}_\infty$  norms of the system are given by

$$\begin{aligned} q_2(\mathbf{K}) &= \mathbb{E}_{w_t \sim \mathcal{N}(0, I)} \left[ \sum_{t=1}^N x_t^T x_t \right] = \mathbb{E} [\text{tr}(\mathbf{x}\mathbf{x}^T)] \\ &= \text{tr}(F(\mathbf{K})^T F(\mathbf{K})) = \sum_{i=1}^{nN} \sigma_i(F(\mathbf{K}))^2 \end{aligned} \quad (1)$$

$$\begin{aligned} q_\infty(\mathbf{K}) &= \sqrt{\max_{w \neq 0} \frac{\sum_{t=1}^N x_t^T x_t}{\sum_{t=0}^{N-1} w_t^T w_t}} \\ &= \max_{w \neq 0} \frac{\|F(\mathbf{K})\mathbf{w}\|}{\|\mathbf{w}\|} = \sigma_{\max}(F(\mathbf{K})). \end{aligned} \quad (2)$$

Note that we can easily include an *invertible* weighting matrix  $P_t$  into the objective as well.

If there are no constraints on  $\mathbf{K}$ , these problems can be solved using standard dynamic programming techniques. However, we are interested in synthesizing structured controllers. We formulate this very generally: We allow *arbitrary* convex constraints on the set of feedback matrices:  $\mathbf{K} \in \mathcal{C}$  for some convex set  $\mathcal{C}$ . Then, the control synthesis problem becomes

$$\underset{\mathbf{K} \in \mathcal{C}}{\text{Minimize}} \quad q_2(\mathbf{K}) \quad (3)$$

$$\underset{\mathbf{K} \in \mathcal{C}}{\text{Minimize}} \quad q_\infty(\mathbf{K}). \quad (4)$$

The general problem of synthesizing stabilizing linear feedback control, subject even to simple bound constraints on the entries of  $K$ , is known to be hard [3]. Several hardness results on linear controller design can be found in [19]. Although these results do not cover the problems (3), (4), they suggest that (3), (4) are hard optimization problems. In this paper, we propose an alternate objective function based on the singular values of the inverse mapping  $F(\mathbf{K})^{-1}$  and prove that this objective can be optimized using convex programming techniques under *arbitrary* convex constraints on the feedback matrices  $\mathbf{K} = \{K_t\}$ . Given the aforementioned hardness results, it is clear that the optimal solution to the convex problem will not match the optimal solution to the original problem. However, we present theoretical and numerical evidence to suggest that the solutions of the convex problem we propose approximate the solution to the original problems well for several problems.

#### A. Control Objective

The problems (3), (4) are nonconvex optimization problems, because of the nonlinear dependence of  $F(\mathbf{K})$  on  $\mathbf{K}$ . In this section, we will derive convex upper bounds on the singular values of  $F(\mathbf{K})$  that can be optimized under arbitrary convex constraints  $\mathcal{C}$ . We have the following results (Appendix, theorems A. 1, A.2):

$$\begin{aligned} q_\infty(\mathbf{K}) &\leq \left( \prod_{t=0}^{N-1} \det(D_t) \right) \left( \frac{\sum_{i=1}^{nN-1} \sigma_i(F(\mathbf{K})^{-1})}{nN-1} \right)^{nN-1} \\ q_2(\mathbf{K}) &\leq nN \left( \prod_{t=0}^{N-1} \det(D_t) \right)^2 (\sigma_{\max}(F(\mathbf{K})^{-1}))^{2(nN-1)}. \end{aligned}$$

To illustrate the behavior of these upper bounds, we generated a generated linear system with a 10-D state vector

$$x_{t+1} = (A + B(\alpha K)) x_t + w_t$$




Fig. 1. Convex surrogate versus original objective (rescaled to lie in  $[0,1]$ ):  $q_2$  versus  $UB_2$ .

where the entries of  $A, B$  are chosen to be iid Gaussian variables. We choose a horizon  $N = 100$ .  $\alpha$  is a scaling factor on the feedback matrix  $K$  and we study the behavior of  $q_2$  and  $UB_2$  (the convex upper bound on  $q_2$ ) as  $\alpha$  varies. In Fig. 1(a), we generate  $K$  with a random sparsity pattern (with 20% nonzero entries) and plot  $q_2$  and  $UB_2$  as a function of  $\alpha$ . The figure shows that  $q_2$  is indeed a nonconvex objective (it has a local minimum). We then generate 100 random  $K$  (each with a random sparsity pattern with 20% nonzeros) and iid Gaussian entries. We compute (using brute force search) the global minimum of  $q_2$  with respect to  $\alpha$  ( $q_2^*$ ) and the  $\alpha$  that minimizes our convexified objective (denoted CON). We then compute the suboptimality of the solution found by our approach relative to  $q_2^*$ :  $q_2 : q_2 - q_2^*/q_2^*$ . We plot a histogram of this suboptimality over the 100 trials in Fig. 1(b). This shows that for a large number of (>85%), our convexification is able to find a solution with an objective value within 20% of the optimal solution. Although based only on the search along a single direction, these plots suggest that the original objective is nonconvex and could get stuck in local minima while our convexification can find a solution that is near-optimal. We formalize this result by giving suboptimality bounds in Section IV-A and provide numerical comparisons to nonconvex approaches in Section VII.

### B. General Class of Control Objectives

Inspired by the upper bounds of the previous section, we formulate the controller design problem as follows:

$$\text{Minimize}_{\mathbf{K} \in \mathcal{C}} q_2^c(\mathbf{K}) = \sigma_{\max}(F(\mathbf{K})^{-1}) \text{ (surrogate to } q_2) \quad (5)$$

$$\text{Minimize}_{\mathbf{K} \in \mathcal{C}} q_{\infty}^c(\mathbf{K}) = \sum_{i=1}^{nN-1} \sigma_i(F(\mathbf{K})^{-1}) \text{ (surrogate to } q_{\infty}). \quad (6)$$

The objectives (5), (6) are just two of the control objectives that are allowed in our framework. We can actually allow a general class of objectives that can be minimized for control design. From [20], we know that for any convex function  $f(x)$  on  $\mathbf{R}^n$  that is invariant to sign changes and permutations of the coordinates of  $x$ , the function  $g(X) = f(\sigma X)$  on  $\mathbf{R}^{n \times n}$  is convex. This motivates us to consider a generalized control objective

$$\begin{aligned} & \text{Minimize}_{\mathbf{K}} \underbrace{f(\sigma(F(\mathbf{K})^{-1}))}_{\text{Controller Performance}} + \underbrace{R(\mathbf{K})}_{\text{Minimize Control Effort}} \\ & \text{Subject to } \mathbf{K} \in \mathcal{C} \end{aligned} \quad (7)$$

where  $\mathcal{C}$  is a convex set encoding the structural constraints on  $\mathbf{K}$ , and  $R(\mathbf{K})$  is a convex penalty on the feedback gains  $\mathbf{K}$ . We show (in theorem V.1) that this problem is a convex optimization problem. Common special cases for  $f$  are:

- 1)  $f(x) = \|x\|_{\infty}$  which gives rise to the spectral norm  $\|(F(\mathbf{K}))^{-1}\| = \sigma_{\max}((F(\mathbf{K}))^{-1})$ , the same as (5).
- 2)  $f(x) = \|x\|_1$  which gives rise to the nuclear norm  $\|((F(\mathbf{K}))^{-1})_* = \sum_i \sigma_i((F(\mathbf{K}))^{-1})$ .
- 3)  $f(x) = \sum_{i=1}^k |x_{[i]}|$  which gives rise to the Ky Fan  $k$ -norm  $\sum_{i=1}^k \sigma_i((F(\mathbf{K}))^{-1})$ . In particular,  $f(x) = \sum_{i=1}^{nN-1} |x_{[i]}|$  corresponds to (6).

A common choice for  $R(\mathbf{K})$  is  $\|\mathbf{K}\|^2$ . For decentralized control,  $\mathcal{C}$  would be of the form  $\mathcal{C} = \{\mathbf{K} : \mathbf{K}_t \in S\}$ , where  $S$  is the set of matrices with a certain sparsity pattern corresponding to the decentralization structure required. We now present our main theorem proving the convexity of the generalized problem (7).

## V. MAIN TECHNICAL RESULTS

### A. Proof of Convexity

*Theorem V.1:* Suppose  $f : \mathbf{R}^{nN} \rightarrow \mathbf{R}$  is a convex function and is invariant to sign changes and permutations of its input, that is

$$f(x_1, \dots, x_{nN}) = f(|x_{\pi(1)}|, \dots, |x_{\pi(nN)}|)$$

for every permutation  $\pi$ . Let  $R(\mathbf{K})$  be a convex function and  $\mathcal{C}$  be a convex set, then the problem (7) is a convex optimization problem.

*Proof:* The proof relies on the structure of  $F(\mathbf{K})^{-1}$ . Rewriting the discrete-time dynamics equations, we have

$$w_0 = D_0^{-1}x_1, \quad w_t = D_t^{-1}x_{t+1} - D_t^{-1}\tilde{A}_t x_t \text{ for } t \geq 1.$$

It can be shown that  $F(\mathbf{K})^{-1}$  is given by

$$\begin{bmatrix} D_0^{-1} & 0 & \dots & \dots & 0 \\ -D_1^{-1}\tilde{A}_1 & D_1^{-1} & \dots & \dots & 0 \\ 0 & -D_2^{-1}\tilde{A}_2 & D_2^{-1} & \dots & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots \\ 0 & 0 & 0 & \dots & D_{N-1}^{-1} \end{bmatrix}.$$

This can be verified by simple matrix multiplication. Now, the convexity is obvious since  $A_t = A_t + B_t K_t$  is a linear function

of  $K$ , and so is  $F(\mathbf{K})^{-1}$ . Since  $f$  is invariant to sign changes and permutations and is convex, the function  $X \mapsto f(\sigma(X))$  is convex [20]. Thus,  $f(\sigma((F(\mathbf{K}))^{-1}))$  is the composition of an affine function in  $K$  with a convex function and is hence convex. The function  $\mathbf{R}(\mathbf{K})$  is known to be convex and so are the constraints  $K \in \mathcal{C}$ . Hence, the overall problem is a convex optimization problem.  $\square$

### B. Suboptimality Bounds

We are using convex surrogates for the  $q_2, q_\infty$  norms. Thus, it makes sense to ask the question: How far are the optimal solutions to the convex surrogates from those of the original problem? We answer this question by proving multiplicative suboptimality bounds: We prove that the ratio of the  $q_2$  norm of the convex surrogate solution and the  $q_2$ -optimal solution is bounded above by a quantity that decreases as the variance of the singular vector of  $(F(\mathbf{K}))^{-1}$  at the optimum. Although these bounds may be quite loose, they provide qualitative guidance about when the algorithm would perform well.

*Theorem V.2:* Let the solution to the convex optimization and original problem be

$$\mathbf{K}_c^* = \underset{\mathbf{K} \in \mathcal{C}}{\operatorname{argmin}} \sigma_{\max} \left( (F(\mathbf{K}))^{-1} \right) \quad (\text{Convex Opt})$$

$$\mathbf{K}^* = \underset{\mathbf{K} \in \mathcal{C}}{\operatorname{argmin}} \sum_i (\sigma_i(F(\mathbf{K})))^2 \quad (\text{Original Opt})$$

respectively. Let  $F^* = (F(\mathbf{K}^*))^{-1}$ ,  $F_c^* = (F(\mathbf{K}_c^*))^{-1}$ . Let

$$\sigma_c^* = \left[ \left( \frac{\sigma_2(F_c^*)}{\sigma_{nN}(F_c^*)} \right)^2, \dots, \left( \frac{\sigma_2(F_c^*)}{\sigma_2(F_c^*)} \right)^2 \right]$$

$$\sigma^* = \left[ \left( \frac{\sigma_{nN}(F^*)}{\sigma_{nN}(F^*)} \right)^2, \dots, \left( \frac{\sigma_{nN}(F^*)}{\sigma_2(F^*)} \right)^2 \right].$$

Then

$$\frac{q_2(\mathbf{K}_c^*)}{q_2(\mathbf{K}^*)} \leq \left( \frac{nN}{nN-1} \right) \exp \left( \frac{\operatorname{Var}(\sigma_c^*) - \operatorname{Var}(\sigma^*)}{2} \right).$$

*Proof:* The proof relies on Holder's defect formula which quantifies the gap in the AM-GM inequality [21]. For any numbers  $0 < a_m \leq \dots \leq a_1$ , we have

$$\left( \frac{\sum_{i=1}^m a_i}{m} \right) \exp \left( -\frac{\mu}{2} \operatorname{Var}(a) \right) = \left( \prod_{i=1}^m a_i \right)^{\frac{1}{m}}$$

where  $\mu \in [(1/a_1)^2, (1/a_m)^2]$ . Plugging in the lower and upper bounds for  $\mu$ , we obtain

$$\left( \frac{\sum_{i=1}^m a_i}{m} \right) \exp \left( -\frac{\operatorname{Var}\left(\frac{a}{a_1}\right)}{2} \right) \geq \left( \prod_{i=1}^m a_i \right)^{\frac{1}{m}}$$

$$\left( \frac{\sum_{i=1}^m a_i}{m} \right) \exp \left( -\frac{\operatorname{Var}\left(\frac{a}{a_m}\right)}{2} \right) \leq \left( \prod_{i=1}^m a_i \right)^{\frac{1}{m}}.$$

Using this inequality with  $a_i = (\sigma_{nN-i+1}(F^*))^{-2}$ ,  $i = 1, 2, 3, \dots, nN-1$ , we obtain

$$\begin{aligned} \frac{q_2(\mathbf{K}_c^*)}{nN-1} &\geq \frac{1}{nN-1} \sum_{i=2}^{nN} \frac{1}{(\sigma_i(F_c^*))^2} \\ &\geq \exp \left( \frac{\operatorname{Var}(\sigma_c^*)}{2} \right) \left( \prod_{i=2}^{nN} \frac{1}{(\sigma_i(F_c^*))^2} \right)^{\frac{1}{nN-1}} \\ &= c \exp \left( \frac{\operatorname{Var}(\sigma_c^*)}{2} \right) (\sigma_{\max}(F_c^*))^{\frac{2}{nN-1}} \end{aligned}$$

where  $c = (\prod_{t=0}^{N-1} \det(D_t))^{2/nN-1}$  and the last equality follows since  $\det(F^*) = \prod_{t=0}^{N-1} \det(D_t)$ . Since  $\mathbf{K}_c^*$  minimizes  $\sigma_{\max}(F(\mathbf{K})^{-1})$ , we have

$$\begin{aligned} \frac{q_2(\mathbf{K}_c^*)}{nN-1} &\geq c \exp \left( \frac{\operatorname{Var}(\sigma_c^*)}{2} \right) (\sigma_{\max}(F_c^*))^{\frac{2}{nN-1}} \\ &\geq \exp \left( \frac{\operatorname{Var}(\sigma_c^*)}{2} \right) \left( \prod_{i=2}^{nN} \left( \frac{1}{\sigma_i(F_c^*)} \right)^2 \right)^{\frac{1}{nN-1}} \\ &\geq \exp \left( \frac{\operatorname{Var}(\sigma_c^*)}{2} - \frac{\operatorname{Var}(\sigma^*)}{2} \right) \left( \frac{\sum_{i=2}^{nN} \frac{1}{(\sigma_i(F_c^*))^2}}{nN-1} \right) \\ &\geq \left( \frac{nN-1}{nN} \right) \exp \left( \frac{\operatorname{Var}(\sigma_c^*)}{2} - \frac{\operatorname{Var}(\sigma^*)}{2} \right) \frac{q_2(\mathbf{K}_c^*)}{nN-1}. \end{aligned}$$

The result follows from simple algebra now.  $\square$

*Theorem V.3:* Let the solution to the convex optimization and original problem be

$$\mathbf{K}_c^* = \underset{\mathbf{K} \in \mathcal{C}}{\operatorname{argmin}} \sum_{i=1}^{nN-1} \sigma_i \left( (F(\mathbf{K}))^{-1} \right) \quad (\text{Convex Opt})$$

$$\mathbf{K}^* = \underset{\mathbf{K} \in \mathcal{C}}{\operatorname{argmin}} \sigma_{\max}(F(\mathbf{K})) \quad (\text{Original Opt})$$

respectively. Let  $F^* = (F(\mathbf{K}^*))^{-1}$ ,  $F_c^* = (F(\mathbf{K}_c^*))^{-1}$ . Let

$$\sigma_c^* = \left[ \frac{\sigma_{nN-1}(F_c^*)}{\sigma_1(F_c^*)}, \dots, \frac{\sigma_1(F_c^*)}{\sigma_1(F_c^*)} \right]$$

$$\sigma^* = \left[ \frac{\sigma_{nN-1}(F^*)}{\sigma_{nN-1}(F^*)}, \dots, \frac{\sigma_1(F^*)}{\sigma_{nN-1}(F^*)} \right].$$

Then

$$\frac{q_\infty(\mathbf{K}_c^*)}{q_\infty(\mathbf{K}^*)} \leq \exp \left( (nN-1) \left( \frac{\operatorname{Var}(\sigma^*) - \operatorname{Var}(\sigma_c^*)}{2} \right) \right).$$

*Proof:* The proof follows a similar structure as the previous theorem and relies on Holder's defect formula. Let  $c = \prod_{t=0}^{N-1} \det(D_t)$ . Using the same inequalities with  $a_i = \sigma_i(F^*)$ ,  $i = 1, 2, \dots, nN-1$ , we obtain

$$\begin{aligned} (q_\infty(\mathbf{K}_c^*))^{\frac{1}{nN-1}} &= c \left( \prod_{i=1}^{nN-1} \sigma_i(F_c^*) \right)^{\frac{1}{nN-1}} \\ &\geq \frac{c \exp \left( -\frac{\operatorname{Var}(\sigma^*)}{2} \right)}{nN-1} \left( \sum_{i=1}^{nN-1} \sigma_i(F^*) \right) \end{aligned}$$

where  $c = (\prod_{t=0}^{N-1} \det(D_t))^{2/nN-1}$ . Since  $\mathbf{K}_c^*$  minimizes  $\sum_{i=1}^{nN-1} \sigma_i(F(\mathbf{K})^{-1})$ , we have

$$\begin{aligned} (q_\infty(\mathbf{K}^*))^{\frac{1}{nN-1}} &\geq \frac{c \exp\left(\frac{-\text{Var}(\sigma^*)}{2}\right)}{nN-1} \left(\sum_{i=1}^{nN-1} \sigma_i(F_c^*)\right) \\ &\geq c \exp\left(\frac{\text{Var}(\sigma_c^*)}{2} - \frac{\text{Var}(\sigma^*)}{2}\right) \\ &\quad \times \left(\prod_{i=1}^{nN-1} \sigma_i(F_c^*)\right)^{\frac{1}{nN-1}} \\ &= \exp\left(\frac{\text{Var}(\sigma_c^*)}{2} - \frac{\text{Var}(\sigma^*)}{2}\right) (q_\infty(\mathbf{K}_c^*))^{\frac{1}{nN-1}}. \end{aligned}$$

The result follows from simple algebra now.  $\square$

### C. Interpretation of Bounds

The bounds have the following interpretation: Since the product of singular values is constrained to be fixed, stable systems (with small  $\mathcal{H}_2, \mathcal{H}_\infty$  norm) would have all of their singular values close to each other. Thus, if the singular values at the solution discovered by our algorithm are close to each other, we can expect that our solution is close to the true optimum. Further, the bounds say that the only thing that matters is the spread of the singular values relative to the spread of singular values at the optimal solution. A side effect of the analysis is that it suggests that the spectral norm of  $(F(\mathbf{K}))^{-1}$  be used as a surrogate for the  $q_2$  norm and the nuclear norm be a surrogate for the  $q_\infty$  norm, since optimizing these surrogates produces solutions with suboptimality bounds on the original objectives.

Suppose that we know a lower bound  $L$  on the quantity  $\text{Var}(\sigma(q_2(\mathbf{K})))$  (such a bound can be found, for example, using a convex relaxation approach) that holds for all  $K \in \mathcal{C}$  (we can always choose  $L = 0$  if no better lower bounds are known). We can then obtain a useful bound for the  $q_2$  case by simply dropping the effect of the negative term so that

$$\frac{q_2(\mathbf{K}_c^*)}{q_2(\mathbf{K}^*)} \leq \left(\frac{nN}{nN-1}\right) \exp\left(\frac{\text{Var}(\sigma_c^*) - L}{2}\right)$$

which can be computed after solving the convex problem to obtain  $\mathbf{K}_c^*$ . From a practical point of view, this would give a bound on how far the solution found by our convexification approach is compared to the true optimum, without having to solve an NP-hard nonconvex optimization problem. This can be used effectively as a diagnostic tool to see when the solutions found by our approach can be trusted to be close to the optimal structured controller (note that this may still not yield satisfactory performance since it is possible that no controller in  $\mathcal{C}$  achieves good performance). A similar approach is possible for the  $q_\infty$  case. Since  $\text{Var}(\sigma^*)$  appears with a positive sign in the bound, one would need *upper* bounds on  $\text{Var}(\sigma^*)$  to obtain a useful suboptimality bound. Once again, this can be computed using a convex relaxation.

Because of the exponential dependence of the bounds on the difference in variance, the suboptimality bounds can be quite

loose. Indeed, we do observe that in some cases our approach does find bad solutions (Section VII). However, the bounds give the intuition that if our approach finds a solution at which the variance in singular values is small, we can rest assured that we have found a near-optimal structured controller. We intend to conduct a finer analysis and get rid of the exponential dependence in the bound in future work, at least for special classes of systems.

## VI. ALGORITHMS AND COMPUTATION

In this paper, our primary focus is to discuss the properties of the new convex formulation of structured controller synthesis we developed here. Algorithms for solving the resulting convex optimization problem (7) are a topic that we will investigate in depth in future work. In most cases, problem (7) can be reformulated as a semidefinite programming problem and solved using off-the-shelf interior point methods. However, although theoretically polynomial time, off-the-shelf solvers tend to be inefficient in practice and do not scale. In this section, we lay out some algorithmic options, including the one we used in our numerical experiments (Section VII).

When the objective used is the nuclear norm  $\sum_{i=1}^{nN} \sigma_i((F(\mathbf{K}))^{-1})$ , we show that it is possible to optimize the objective using standard Quasi-Newton approaches. The nuclear norm is a nonsmooth function in general, but given the special structure of the matrices appearing in our problem, we show that it is differentiable. Thus, one can use standard gradient descent and Quasi Newton methods to minimize it. These methods are orders of magnitude more efficient than other approaches (reformulating as an SDP and using off-the-shelf interior point methods). They still require computing the SVD of an  $nN \times nN$  matrix at every iteration, which will get prohibitively expensive when  $nN$  is on the order of several thousands. However, the structure of  $F(\mathbf{K})^T F(\mathbf{K})$  is block-tridiagonal and efficient algorithms have been proposed for computing the eigenvalues of such matrices (see [22] and the references therein). Since the singular values of  $F(\mathbf{K})$  are simply square roots of eigenvalues of  $F(\mathbf{K})F(\mathbf{K})$ , this approach could give us efficient algorithms for computing the SVD of  $F(\mathbf{K})$ .

When the objective is the spectral norm  $\sigma_{\max}((F(\mathbf{K}))^{-1})$ , we can reformulate the problem as a semidefinite programming problem (SDP):  $\text{Minimize}_{t, \mathbf{K} \in \mathcal{C}} t + R(\mathbf{K})$

$$\text{Subject to } tI \geq \begin{pmatrix} 0 & (F(\mathbf{K}))^{-1T} \\ (F(\mathbf{K}))^{-1} & 0 \end{pmatrix}.$$

The log-barrier for the semidefinite constraint can be rewritten as  $\log(\det(t^2 - F(\mathbf{K})^{-1}F(\mathbf{K})^{-1}))$  using Schur complements. The matrix  $(F(\mathbf{K}))^{-1T}(F(\mathbf{K})^{-1})$  is a symmetric positive definite block-tridiagonal matrix, which is a special case of a chordal sparsity pattern [23]. This means that computing the gradient and Newton step for the log-barrier is efficient, with complexity growing as  $O(N)$ . Thus, at least for the case where the objective is the spectral norm, we can develop efficient interior point methods.

## VII. NUMERICAL EXPERIMENTS

## A. Comparing Algorithms: Decentralized Control

In this section, we compare different approaches to controller synthesis. We work with discrete-time LTI systems over a fixed horizon  $N$  with  $A_t = A, B_t = B = I, D_t = D = I$ . Let  $S \subset \{1, \dots, n\} \times \{1, \dots, n\}$  denote the set of allowed nonzero indices for the feedback matrix  $K$ . Further, we will use  $\mathcal{C} = \{K : K_{ij} = 0 \forall (i, j) \notin S\}$ . The control design methodologies we compare are **NCON**: This refers to nonconvex approaches for the  $q_2$  and  $q_\infty$  norms. The  $q_2$  norm is a differentiable function and we use a standard LBFGS method [24] to minimize it. The  $q_\infty$  norm is nondifferentiable, but only at points where the maximum singular value of  $F(\mathbf{K})$  is not unique. We use a nonsmooth Quasi Newton method [25] to minimize it (using the freely available software implementation HANSO [26]). **CON**: The convex control synthesis described here. In the experiments described here, we use the following objective:

$$\begin{aligned} & \text{Minimize}_K \frac{1}{nN} \left( \sum_{i=1}^m \sigma_i(F(K)^{-1}) \right) \\ & \text{where} \\ & F(K)^{-1} = \begin{bmatrix} I & 0 & \dots & 0 \\ -(A+BK) & I & \dots & 0 \\ 0 & -(A+BK) & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & I \end{bmatrix} \\ & \text{Subject to} \\ & K_{ij} = 0 \quad \forall (i, j) \notin S \end{aligned} \quad (8)$$

with  $m = nN - 1$  as a surrogate for the  $q_\infty$  norm and  $m = 1$  for the  $q_2$  norm. Although these objectives are non differentiable, we find that an off-the-shelf LBFGS optimizer [24] works well and use it in our experiments here. **OPT**: The optimal solution to the problem in the absence of the constraint  $\mathcal{C}$ . This is simply the solution to a standard LQR problem for the  $q_2$  case. For the  $q_\infty$  norm, this is computed by solving a series of LQ games with objective  $\sum_{t=1}^N x_t^T x_t - \sum_{t=0}^{N-1} \gamma^2 w_t^T w_t$  where the controller chooses  $w_t$  to minimize the cost while an adversary chooses  $w_t$  in order to maximize the cost. There is a critical value of  $\gamma$  below which the upper value of this game is unbounded. This critical value of  $\gamma$  is precisely the  $q_\infty$  norm and the resulting policies for the controller at this value of  $\gamma$  are the  $q_\infty$ -optimal control policy. For any value of  $\gamma$ , the solution of the game can be computed by solving a set of Riccati equations backward in time [27]. We work with a dynamical system formed by coupling a set of systems with unstable dynamics  $A^i \in \mathbf{R}^{2 \times 2} : x_{t+1}^i = A^i x_t^i + \sum_j \eta_{ij} x_t^j + u_t^i + w_t^i$ , where  $x^i$  denotes the state of the  $i$ th system and  $\eta_{ij}$  is a coupling coefficient between systems  $i$  and  $j$ . The objective is to design controls  $u = \{u^i\}$  in order to stabilize the overall system. In our examples, we use  $N = 5$  systems giving us a 10-D state space. The  $A^i, \eta_{ij}$  are generated randomly, with each entry having a Gaussian distribution with mean 0 and variance 10. The sparsity pattern  $S$  is also generated randomly by picking 20% of the off-diagonal entries of  $K$  and setting them to 0. For the CON and NCON problems, we initialize the optimizer at the same point  $\mathbf{K} = 0$ . For the  $q_\infty$  norm, we present results comparing




Fig. 2. Comparison of algorithms for  $q_\infty$ -norm controller synthesis. The blue bars represent histograms and the red curves represent kernel density estimates of the distribution of values.

the approaches over 100 trials. The  $q_\infty$  norm of the solution is obtained by the CON approach to that found by NCON, OPT in Fig. 2. We plot histograms of how the  $q_\infty$  is compared between the CON, NCON, and OPT approaches. The red curves show kernel-density estimates of the distribution of values being plotted. The results show that CON consistently outperforms NCON and often achieves performance close to the centralized OPT solution. The  $x$ -axis denotes the ratio between objectives on a log scale. The  $y$ -axis shows the frequency with which a particular ratio is attained (out of 100 trials). We also plot a histogram of computation times with the log of ratio of CPU times for the CON and NCON algorithms on the  $x$ -axis. Again, in terms of CPU times, the CON approach is consistently superior except for a small number of outliers. For the  $q_2$  norm, we plot the results in Fig. 3. Here, the NCON approach does better and beats the CON approach for most trials. However, in more than 70% of the trials, the  $q_2$  norm of the solution found by CON is within 2% of that found by NCON. In terms of computation time, the CON approach retains superiority.

The numerical results indicate that the convex surrogates work well in many cases. However, they do fail in particular cases. In general, the surrogates seem to perform better on the  $q_\infty$  norm than the  $q_2$  norm. The initial results are promising but we believe that further analytical and numerical work is required to exactly understand when the convex objectives proposed in this paper are good surrogates for the original nonconvex  $q_2$  and  $q_\infty$  objectives.




Fig. 3. Comparison of algorithms for  $q_2$ -norm controller synthesis. The blue bars represent histograms and the red curves represent kernel density estimates of the distribution of values.

### VIII. GENERALIZATION TO NONLINEAR SYSTEMS

We now present a generalization of our approach to nonlinear systems. The essential idea is to study a nonlinear system in terms of sensitivities of system trajectories with respect to disturbances. Consider a control-affine nonlinear discrete-time system

$$\begin{aligned} x_1 &= D_0 w_0 \\ x_{t+1} &= a_t(x_t) + B_t(x_t)u_t + D_t w_t \quad (1 \leq t \leq N-1) \end{aligned}$$

where  $a_t : \mathbf{R}^n \mapsto \mathbf{R}^n$  and  $B : \mathbf{R}^n \mapsto \mathbf{R}^{n \times n_u}$ ,  $D_t \in \mathbf{R}^{n \times n_w}$ ,  $x_t \in \mathbf{R}^n$ ,  $w_t \in \mathbf{R}^{n_w}$ ,  $u_t \in \mathbf{R}^{n_u}$ . Suppose that 0 is an equilibrium point (if not, we simply translate the coordinates to make this the case). Now we seek to design a controller  $u_t = K_t \phi(x_t)$ , where  $\phi$  is any set of fixed “features” of the state on which we want the control to depend that minimizes deviations from the constant trajectory  $[0, 0, \dots, 0]$ . We can look at the closed-loop system

$$x_{t+1} = a_t(x_t) + B_t(x_t)K_t \phi_t(x_t) + D_t w_t$$

where  $\phi_t(x_t) \in \mathbf{R}^m$ ,  $K_t \in \mathbf{R}^{n_u \times m}$ . As before, let  $\mathbf{K} = \{K_t : 1 \leq t \leq N-1\}$ . Let  $F(\mathbf{K})(\mathbf{w})$  denote the (nonlinear) mapping from a sequence of disturbances  $\mathbf{w} = [w_0, \dots, w_{N-1}]$  to the state-space trajectory  $\mathbf{x} = [x_1, \dots, x_N]$ . The finite-horizon  $q_\infty$  norm for a nonlinear system can be defined analogously as for a linear system

$$\max_{\mathbf{w} \neq 0} \frac{\|F(\mathbf{K})(\mathbf{w})\|}{\|\mathbf{w}\|}. \quad (9)$$

Given a state trajectory  $x = [x_1, \dots, x_N]$ , we can recover the noise sequence as

$$\begin{aligned} w_0 &= D_0^{-1} x_1 \\ w_t &= D_t^{-1} (x_{t+1} - a_t(x_t) - B_t(x_t)K_t \phi_t(x_t)), \quad t > 0. \end{aligned} \quad (10)$$

Thus, the map  $F(\mathbf{K})$  is invertible. Let  $F(\mathbf{K})^{-1}$  denote the inverse. It can be shown (theorem A.4) that the objective (9) (assuming it is finite) can be bounded above by

$$\sup_{\mathbf{x}} \left( \frac{\sum_{i=1}^{nN-1} \sigma_i \left( \left( \frac{\partial(F(\mathbf{K}))^{-1}(\mathbf{x})}{\partial \mathbf{x}} \right) \right)}{nN-1} \right)^{nN-1}.$$

In the linear case, the maximization over  $\mathbf{x}$  is unnecessary since the term being maximized is independent of  $\mathbf{x}$ . However, for a nonlinear system, the Jacobian of  $(F(\mathbf{K}))^{-1}(\mathbf{x})$  is a function of  $\mathbf{x}$  and explicit maximization needs to be performed to compute the objective. Thus, we can formulate the control design problem as

$$\min_{\mathbf{K} \in \mathcal{C}} \sup_{\mathbf{x}} \left( \frac{\sum_{i=1}^{nN-1} \sigma_i \left( \left( \frac{\partial(F(\mathbf{K}))^{-1}(\mathbf{x})}{\partial \mathbf{x}} \right) \right)}{nN-1} \right)^{nN-1}. \quad (11)$$

The convexity of the aforementioned objective follows using a very similar proof as the linear case (see theorem A.3). Computing the objective (maximizing over  $\mathbf{x}$ ) in general would be a hard problem, so this result is only of theoretical interest in its current form. However, in future work, we hope to explore the computational aspects of this formulation more carefully.

### IX. CONCLUSION

We have argued that the framework developed seems promising and overcomes limitations of previous works on computationally tractable approaches to structured controller synthesis. Although the control objective used is nonstandard, we have argued why it is a sensible objective, and we also presented numerical examples showing that it produces controllers outperforming other nonconvex approaches. Further, we proved suboptimality bounds that give guidance on when our solution is good even with respect to the original  $(\mathcal{H}_2/\mathcal{H}_\infty)$  metrics. There are four major directions for future work: 1) investigating the effect of various objectives in our family of control objectives; 2) developing efficient solvers for the resulting convex optimization problems; 3) finding conditions under which our approach is guaranteed to produce stabilizing controllers; and 4) deriving computationally efficient algorithms for nonlinear systems.

### APPENDIX A

#### A. Penalizing Control Effort

To model control costs, define an augmented problem with  $\bar{x}_t \in \mathbf{R}^{n+n_u}$ ,  $\bar{w}_t \in \mathbf{R}^{n+n_w}$

$$\begin{aligned} \bar{A}_t &= \begin{pmatrix} A_t & 0 \\ 0 & 0 \end{pmatrix}, \quad \bar{B}_t = \begin{pmatrix} B_t \\ R_t \end{pmatrix}, \quad \bar{D}_t = \begin{pmatrix} D_t & 0 \\ 0 & \gamma I \end{pmatrix} \\ \bar{x}_{t+1} &= \bar{A}_t \bar{x}_t + \bar{B}_t u_t + \bar{D}_t \bar{w}_t. \end{aligned}$$



Partitioning the new state  $\bar{x}_t = \begin{pmatrix} x_t \\ \tilde{x}_t \end{pmatrix}$ ,  $\bar{w}_t = \begin{pmatrix} w_t \\ \tilde{w}_t \end{pmatrix}$ , we have  $x_{t+1} = A_t x_t + B_t u_t + D_t w_t$ ,  $\tilde{x}_{t+1} = R_t u_t + \gamma \tilde{w}_t$ . Given this  $\sum_{t=1}^N \bar{x}_t^T \bar{x}_t = \sum_{t=1}^N x_t^T x_t + \sum_{t=1}^{N-1} (R_t u_t + \gamma \tilde{w}_t)^T (R_t u_t + \gamma \tilde{w}_t) + \gamma^2 w_0^T w_0$ . In the limit  $\gamma \rightarrow 0$ , we recover the standard LQR cost. However, setting  $\gamma = 0$  violates the condition of invertibility. Thus, solving the problem with an augmented state  $\bar{x} \in \mathbf{R}^{n_u + \nu}$ ,  $\bar{w} \in \mathbf{R}^{n_u + \nu}$

$$\bar{A}_t = \begin{pmatrix} A_t & 0 \\ 0 & 0 \end{pmatrix}, \bar{B}_t = \begin{pmatrix} B_t \\ R_t \end{pmatrix}, \bar{D}_t = \begin{pmatrix} D_t & 0 \\ 0 & \gamma I \end{pmatrix}$$

solves the problem with a quadratic control cost in the limit  $\gamma \rightarrow 0$ . The caveat is that the problems (5), (6) become increasingly ill-conditioned as  $\gamma \rightarrow 0$ . However, we should be able to solve the problem for a small value of  $\gamma$ , which models the quadratic controls cost closely but still leads to a sufficiently well-conditioned problem that we can solve numerically.

### B. Dynamic Output Feedback

To allow for dynamic output feedback based on the last  $k$  measurements  $y_t = C_t x_t$ , create an augmented state  $\bar{x}_t = (x_t \dots x_{t-k})^T$ . Then, define  $\tilde{K}_t \in \mathbf{R}^{n_u \times k m}$  and

$$\tilde{K}_t = K_t \begin{pmatrix} C_t & 0 & \dots & 0 \\ 0 & C_{t-1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & C_{t-k} \end{pmatrix}$$

and augmented dynamics

$$\bar{A}_t = \begin{pmatrix} A_t & 0 & \dots & 0 & 0 \\ I & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & I & 0 \end{pmatrix} \bar{D}_t = \begin{pmatrix} D_t & 0 & \dots & 0 \\ 0 & \gamma I & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \gamma I \end{pmatrix}$$

and  $\bar{B}_t = (B_t \ 0 \ \dots \ 0)^T$ . Again, we need to set  $\gamma = 0$  to exactly match the standard output feedback problem but that violates the assumption of invertibility. We can consider taking  $\gamma \rightarrow 0$  and recovering the solution as a limiting case, as in the previous section.

### C. Proofs

*Theorem A.1:*

$$q_\infty(\mathbf{K}) = \sigma_{\max}(F(\mathbf{K})) \leq \prod_{t=0}^{N-1} \det(D_t) \left( \frac{\sum_{i=1}^{nN-1} \sigma_i(F(\mathbf{K})^{-1})}{nN-1} \right)^{nN-1}.$$

*Proof:* Since  $F(\mathbf{K})$  is a block lower triangular matrix (a reflection of the fact that we have a causal linear system), its determinant is simply the product of determinants of diagonal blocks:  $\det(F(\mathbf{K})) = \prod_t \det(D_t) = c$  independent of the values of  $A_t$ . In fact, this result is a generalization of Bode's classical sensitivity integral result and has been studied in [14]. Since the product of singular values is equal to the determinant, we have

$$\sigma_{\max}(F(\mathbf{K})) = \frac{c}{\prod_{i=2}^{nN} \sigma_i(F(\mathbf{K}))} = c \prod_{i=1}^{nN-1} \sigma_i(F(\mathbf{K})^{-1})$$

where the last equality follows because the singular values of  $F(\mathbf{K})^{-1}$  are simply reciprocals of the singular values of  $F(\mathbf{K})$ . The result now follows using the AM-GM inequality.  $\square$

*Theorem A.2:*

$$q_2(\mathbf{K}) \leq nN \left( \prod_{t=0}^{N-1} \det(D_t) \right)^2 (\sigma_{\max}(F(\mathbf{K})^{-1}))^{2(nN-1)}.$$

*Proof:* Let  $\prod_{t=0}^{N-1} \det(D_t) = c$ . From the above argument, we can express  $\sigma_i(F(\mathbf{K}))$  as

$$c \prod_{j \neq nN-i+1} \sigma_j((F(\mathbf{K}))^{-1}) \leq c (\sigma_{\max}(F(\mathbf{K})^{-1}))^{(nN-1)}.$$

The expression for  $q_2(\mathbf{K})$  is

$$\sum_{i=1}^{nN} (\sigma_i(F(\mathbf{K})))^2 \leq nN c^2 (\sigma_{\max}(F(\mathbf{K})^{-1}))^{2(nN-1)}. \quad \square$$

*Theorem A.3:* For the nonlinear system described in (10), the function  $\sup_{\mathbf{x}} (\sum_{i=1}^{nN-1} \sigma_i((\partial(F(\mathbf{K}))^{-1}(\mathbf{x}))) / (nN-1))^{nN-1}$  is convex in  $K$ .

*Proof:* First fix  $\mathbf{w}$  to an arbitrary value. From (10), we know that  $(\partial(F(\mathbf{K}))^{-1}(\mathbf{x}) / \partial \mathbf{x})$  is of the form

$$\begin{bmatrix} D_0^{-1} & 0 & \dots & \dots & 0 \\ -D_1^{-1} \frac{\partial w_1}{\partial x_1} & D_1^{-1} & \dots & \dots & 0 \\ 0 & -D_2^{-1} \frac{\partial w_2}{\partial x_2} & D_2^{-1} & \dots & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots \\ 0 & 0 & 0 & \dots & D_{N-1}^{-1} \end{bmatrix}.$$

Since  $w_t = D_t(x_{t+1} - a_t(x_t) - B_t(x_t)K_t\phi_t(x_t))$ ,  $w_t$  is an affine function of  $\mathbf{K}$ . Hence, so is  $\partial w_t / \partial x_t$ , for any  $t$ . Thus, the overall matrix  $(\partial(F(\mathbf{K}))^{-1}(\mathbf{x}) / \partial \mathbf{x}) = M(\mathbf{K})$  is an affine function of  $\mathbf{K}$ . Thus, by composition properties,  $(\sum_{i=1}^{nN-1} \sigma_i(M(\mathbf{K}))) / (nN-1)$  is a convex function of  $\mathbf{K}$  for any fixed  $\mathbf{x}$ . Taking a supremum over all  $\mathbf{x}$  preserves convexity, since the pointwise supremum of a set of convex functions is convex.  $\square$

*Theorem A.4:* Consider the nonlinear system described in (10). Suppose that  $\sup_{\mathbf{w} \neq 0} (\|F(\mathbf{K})(\mathbf{w})\| / \|\mathbf{w}\|)$  is finite and the supremum is achieved at  $\mathbf{w}^* \neq 0$  for all values of  $\mathbf{K}$ . Then,  $\sup_{\mathbf{w} \neq 0} (\|F(\mathbf{K})(\mathbf{w})\| / \|\mathbf{w}\|)$  is bounded above by

$$\sup_{\mathbf{x}} \left( \frac{\sum_{i=1}^{nN-1} \sigma_i \left( \left( \frac{\partial(F(\mathbf{K}))^{-1}(\mathbf{x})}{\partial \mathbf{x}} \right) \right)}{nN-1} \right)^{nN-1}.$$

*Proof:* By theorem A.5,  $\sup_{\mathbf{w} \neq 0} (\|F(\mathbf{K})(\mathbf{w})\| / \|\mathbf{w}\|)$  is bounded above by  $\sup_{\mathbf{w} \neq 0} \sigma_{\max}(\partial F(\mathbf{K})(\mathbf{w}) / \partial \mathbf{w})$ . Now,  $M(\mathbf{K}) = \partial F(\mathbf{K})(\mathbf{w}) / \partial \mathbf{w}$  is a lower-triangular matrix (since we have a causal system) and the diagonal blocks are given by  $D_t$ . Thus,  $\det(M(\mathbf{K})) = \prod_{t=0}^{N-1} \det(D_t) = c$ , and we can rewrite  $\sigma_{\max}(M(\mathbf{K}))$  as  $c \prod_{i=1}^{nN-1} \sigma_i((M(\mathbf{K}))^{-1})$ . By the rules of calculus, we know that  $(M(\mathbf{K}))^{-1} = (\partial(F(\mathbf{K}))^{-1}(\mathbf{x}) / \partial \mathbf{x})_{\mathbf{x}=F(\mathbf{K})(\mathbf{w})}$ . Thus, the above objective reduces to  $\sup_{\mathbf{w} \neq 0} \prod_{i=1}^{nN-1} \sigma_i((\partial(F(\mathbf{K}))^{-1}(\mathbf{x}) / \partial \mathbf{x})_{\mathbf{x}=F(\mathbf{K})(\mathbf{w})})$ . Given any  $\mathbf{x}$ , we can find  $\mathbf{w}$  such that  $\mathbf{x} = F(\mathbf{K})(\mathbf{w})$  (simply choose  $\mathbf{w} = (F(\mathbf{K}))^{-1}(\mathbf{x})$ ). Thus, the above quantity is equal to  $\sup_{\mathbf{x}} \prod_{i=1}^{nN-1} \sigma_i((\partial(F(\mathbf{K}))^{-1}(\mathbf{x}) / \partial \mathbf{x}))$ . The result now follows using the AM-GM inequality.

**Theorem A.5:** Let  $g(\mathbf{y}) : \mathbf{R}^l \mapsto \mathbf{R}^p$  be any differentiable function. If the function  $\|g(\mathbf{y})\|_2 / \|\mathbf{y}\|_2$  attains its maximum at  $\mathbf{y}^*$

$$\sup_{\mathbf{y}} \frac{\|g(\mathbf{y})\|_2}{\|\mathbf{y}\|_2} \leq \sup_{\mathbf{y} \neq 0} \sigma_{\max} \left( \frac{\partial g(\mathbf{y})}{\partial \mathbf{y}} \right) \quad \forall \mathbf{y} \neq 0.$$

*Proof:* The result follows by writing down the optimality conditions for the optimization problem  $\max \log(\|g(\mathbf{y})\|_2^2 / \|\mathbf{y}\|_2^2)$ .  $\square$

#### ACKNOWLEDGMENT

The authors would like to thank M. Mesbahi, A. Lamperski, and P. Shah for helpful discussions and feedback on earlier versions of this manuscript. They would also like to thank M. Jovanovic for suggestions that significantly improved the numerical results section.

#### REFERENCES

- [1] R. E. Kalman, "A new approach to linear filtering and prediction problems," *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 1960.
- [2] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM J. Control*, vol. 6, no. 1, pp. 131–147, 1968.
- [3] V. Blondel and J. N. Tsitsiklis, "Np-hardness of some linear control design problems," *SIAM J. Control Optimiz.*, vol. 35, no. 6, pp. 2118–2127, 1997.
- [4] M. Rotkowitz and S. Lall, "Decentralized control information structures preserved under feedback," in *Proc. 41st IEEE Conf. Dec. Control*, 2002, vol. 1, pp. 569–575.
- [5] X. Qi, M. V. Salapaka, P. G. Voulgaris, and M. Khammash, "Structured optimal and robust control with multiple criteria: A convex solution," *IEEE Trans. Autom. Control*, vol. 49, no. 10, pp. 1623–1640, Oct. 2004.
- [6] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *IEEE Trans. Autom. Control*, vol. 51, no. 2, pp. 274–286, Feb. 2006.
- [7] P. Shah, "H2-optimal decentralized control over posets: A state-space solution for state-feedback," *IEEE Trans. Autom. Control*, vol. 58, no. 12, pp. 3084–3096, Dec. 2013.
- [8] C.-H. Fan, J. L. Speyer, and C. R. Jaensch, "Centralized and decentralized solutions of the linear-exponential-gaussian problem," *IEEE Trans. Autom. Control*, vol. 39, no. 10, pp. 1986–2003, Oct. 1994.
- [9] J. Swigart and S. Lall, "An explicit dynamic programming solution for a decentralized two-player optimal linear-quadratic regulator," in *Proc. Math. Theory Netw. Syst.*, 2010, pp. 1443–1447.
- [10] A. Lamperski and L. Lessard, "Optimal decentralized state-feedback control with sparsity and delays, 2013. [Online]. Available: arxiv.org/abs/1306.0036"
- [11] F. Lin, M. Fardad, and M. R. Jovanović, "Design of optimal sparse feedback gains via the alternating direction method of multipliers," *IEEE Trans. Autom. Control*, vol. 58, no. 9, pp. 2426–2431, Sep. 2013.
- [12] P. Apkarian, D. Noll, and A. Rondepierre, "Mixed h-2/h-inf control via nonsmooth optimization," *SIAM J. Control Optimiz.*, vol. 47, no. 3, pp. 1516–1546, 2008.
- [13] J. Burke, D. Henrion, A. Lewis, and M. Overton, "Hifoo—a matlab package for fixed-order controller design and h optimization," presented at the 5th IFAC Symp. Robust Control Design, Toulouse, France, 2006.
- [14] P. A. Iglesias, "Tradeoffs in linear time-varying systems: An analogue of bode's sensitivity integral," *Automatica*, vol. 37, no. 10, pp. 1541–1550, 2001.
- [15] C. W. Scherer, "Structured Hinf-Optimal Control for nested interconnections: A state-space solution, May 2013. [Online]. Available: arxiv.org/abs/1305.1746v2"
- [16] J. Lavaei, "Optimal decentralized control problem as a rank-constrained optimization," in *Proc. 51st Annu. Allerton Conf. Control Comput.*, 4Oct. 2013, pp. 39–45.
- [17] G. Zhai, M. Ikeda, and Y. Fujisaki, "Decentralized h-2/h-inf controller design: A matrix inequality approach using a homotopy method," *Automatica*, vol. 37, no. 4, pp. 565–572, 2001.
- [18] K. Dvijotham, E. Todorov, and M. Fazel, "Convex controller design via covariance minimization," in *Proc. Allerton Conf. Control Comput.*, Monticello, IL, USA, 2013, pp. 93–99.
- [19] V. D. Blondel and J. N. Tsitsiklis, "A survey of computational complexity results in systems and control," *Automatica*, vol. 36, no. 9, pp. 1249–1274, 2000.
- [20] A. S. Lewis, "The convex analysis of unitarily invariant matrix functions," *J. Convex Anal.*, vol. 2, no. 1, pp. 173–183, 1995.
- [21] R. Becker, "The variance drain and jensen's inequality," in *CAEPR Working Paper*, 2012.
- [22] A. Sandryhaila and J. M. F. Moura, "Eigendecomposition of block tridiagonal matrices, Jun. 2013. [Online]. Available: arxiv.org/pdf/1306.0217"
- [23] M. S. Andersen, J. Dahl, and L. Vandenbergh, "Implementation of nonsymmetric interior-point methods for linear optimization over sparse matrix cones," *Math. Program. Comput.*, vol. 2, no. 3, pp. 167–201, 2010.
- [24] M. Schmidt, "Minfunc: Unconstrained differentiable multivariate optimization in Matlab 2012. [Online]. Available: http://www.cs.ubc.ca/~schmidtm/Software/minFunc.html"
- [25] A. S. Lewis and M. L. Overton, "Nonsmooth optimization via quasi-newton methods," *Math. Program.*, pp. 1–29, 2012.
- [26] M. Overton, "Hanso 2013. [Online]. Available: http://www.cs.nyu.edu/overton/software/hanso/"
- [27] T. Başar and P. Bernhard, *H-Infinity Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*. New York, USA: Springer, 2008.



in different domains.



**Krishnamurthy Dvijotham** received the B.Tech degree of computer science and engineering from IIT Bombay, Mumbai, India, in 2008, and the M.S. and Ph.D. degrees in computer science and engineering from the University of Washington, Seattle, WA, USA, in 2010 and 2014, respectively.

Currently, he is a Postdoctoral Fellow at the Center for Mathematics of Information at the California Institute of Technology, Pasadena, CA. His current research interests include convex optimization and its applications to optimal control problems arising

**Emanuel Todorov** received the B.S., M.S., and Ph.D. degrees in cognitive neuroscience from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1993, 1995, and 1998, respectively.

Currently, he is an Associate Professor at the Departments of Computer Science and Engineering and Applied Mathematics at the University of Washington, Seattle, WA, USA, since 2009. Prior to joining UW, he also held positions at the Gatsby Computational Neuroscience Unit, USC and UCSD.



**Maryam Fazel** received the B.S. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1995, and the M.S. degree in 1997, and Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 2002.

Currently, she is an Associate Professor of Electrical Engineering at the University of Washington (UW), Seattle, WA, USA, with adjunct appointments in the Departments of Computer Science and Engineering, Mathematics, and Statistics. She was a Research Scientist at Caltech, Pasadena, CA, USA prior to joining UW, Seattle, WA, USA in 2008. Her current research interests include convex optimization and parsimonious models with applications in control, machine learning, and signal processing.

Prof. Fazel is a recipient of the National Science Foundation Career Award and the UW EE Outstanding Teaching Award.