

Convex Control Design via Covariance Minimization

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 ($\mathcal{H}_2, \mathcal{H}_\infty, \text{LQR}$) and previous works have focussed on characterizing classes of structural constraints that allow efficient solution through convex optimization or dynamic programming techniques. In this paper, we propose a new control objective based on eigenvalues of the covariance matrix of trajectories of the system 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 distributed 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 that it agrees well with standard notions of control. We numerically validate the our approach on problems arising in power systems and simple mechanical systems.

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structure arising from the need to implement control in a decentralized fashion), the dynamic programming approach is no longer applicable.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]. Previous approaches to synthesizing structured controllers can be broadly categorized into three types: Frequency Domain Approaches[4], [5], [6], [7], Dynamic Programming Approaches [8], [9], [10] and Nonconvex optimization methods [11], [12], [13]. The first two classes of approaches find *exact* solutions to structured control problems for special cases. The third class of approaches tries to directly solve the optimal control problem (minimizing the $\mathcal{H}_2, \mathcal{H}_\infty$ norm) subject to constraints on the controller, using nonconvex optimization techniques. These are generally applicable, but are susceptible to local minima and slow convergence

K. Dvijotham is with the Department of Computer Science and Engineering, University of Washington, Seattle, WA, 98195 USA e-mail: dvij@cs.washington.edu

E. Todorov is with the Departments of Computer Science and Engineering and Applied Mathematics at the University of Washington, Seattle. e-mail: todorov@cs.washington.edu

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

(especially for nonsmooth norms such as \mathcal{H}_∞).

Our approach involves the following steps:

- a We formulate the control problem in terms of the minimizing the eigenvalues of the covariance matrix of the system trajectories.
- b We show that the above problem can be solved via an equivalent semidefinite programming (SDP) formulation.
- c We solve the SDP formulation of the problem using standard semidefinite programming techniques. Further, we show that the problem has sparsity structure that enables efficient solution: The computational complexity grows linearly in the time horizon considered.

A. Notation

We use $\mathbf{x} \in \mathbf{R}^n$ to denote states, $\mathbf{u} \in \mathbf{R}^{\text{nu}}$ for controls and $\epsilon \in \mathbf{R}^n$ for process noise. We work with discrete-time systems and denote integer-valued time with t ($1 \leq t \leq T$). T denoting the horizon of the finite-horizon control problem. The time-indices on quantities are indicated as subscripts ($\mathbf{x}_t, \mathbf{u}_t$, etc.). We use $\lambda_{\max}(M)$ to denote the maximum eigenvalue of an $l \times l$ symmetric matrix M , $\lambda_{\min}(M)$ to denote the minimum eigenvalue and $\lambda_i(M)$ the i -th eigenvalue in ascending order:

$$\lambda_1(M) = \lambda_{\min}(M) \leq \lambda_2(M) \leq \dots \\ \dots \leq \lambda_{\max}(M) = \lambda_l(M).$$

We denote by $\text{Cov}_P(Y)$ the covariance matrix of the random variable $Y \in \mathbf{R}^l$ with distribution P .

II. PROBLEM FORMULATION

We deal with linear systems of the form

$$\mathbf{x}_1 \sim \mathcal{N}(0, \Sigma_0)$$

$$\mathbf{x}_{t+1} = A_t \mathbf{x}_t + B_t \mathbf{u}_t + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, \Sigma_t), t = 1, 2, \dots, T-1$$

where $\mathbf{x}_t, \epsilon_t \in \mathbf{R}^n, A_t \in \mathbf{R}^{n \times n}, B_t \in \mathbf{R}^{n \times \text{nu}}, \mathbf{u}_t \in \mathbf{R}^{\text{nu}}$. We will assume that Σ_t is full rank for all $0 \leq t \leq T-1$ (0 refers to the distribution of the initial state \mathbf{x}_1). We denote $S_i = \Sigma_i^{-1}$.

We seek to design feedback matrices $K_t \in \mathbf{R}^{\text{nu} \times n}$ so that the control policy $\mathbf{u}_t = K_t \mathbf{x}_t$ drives the system state \mathbf{x}_t towards the origin. Let $\mathbf{K} = \{K_t\}$ and denote by $P_{\mathbf{K}}(\mathbf{X})$ the joint Gaussian density over the trajectories $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]$ sampled from this stochastic dynamical system. Let P_0 denote the probability density over trajectories in the absence of feedback ($\mathbf{K} = 0$). The feedback \mathbf{K} is chosen to minimize the deviations from the mean trajectory $[0, 0, \dots, 0]$

as quantified by the covariance matrix of the trajectories $E_{P_{\mathbf{K}}}[\mathbf{X}\mathbf{X}^T]$. More concretely, we solve an optimization problem of the form:

$$\begin{aligned} & \text{Minimize } \lambda_{\max} \left(\left(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}) \right)^{-1} \right) \\ & \text{Subject to } \text{Convex Structural Constraints/Costs on } \mathbf{K} \end{aligned}$$

We show (in theorem 3.1) that this problem can be solved via convex programming (specifically semidefinite programming) techniques. In the following section, we discuss the choice of control objective and how it relates to more standard formulations of control.

A. Justification for the Control Objective

The problem formulation is stated as minimizing the maximum eigenvalue of the inverse covariance - which is equivalent to maximizing the minimum eigenvalue of the covariance (since these eigenvalues are positive and reciprocal to each other). Intuitively, we would expect that doing so would produce controllers that destabilize the system, since it makes the covariance “large”. However, it turns out that given the structure of the inverse covariance matrix $P_{\mathbf{K}}$, this is actually a sensible thing to do.

$P_{\mathbf{K}}$ has a Gaussian density over the space of trajectories \mathbf{X} with the joint inverse covariance $\mathcal{M} = (\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}))^{-1}$. Surprisingly, the structure of the covariance is such that it satisfies $\det(\mathcal{M}) = \prod_{t=0}^{T-1} \det(S_t)$ *irrespective* of the values of A, B, K (lemma 8.1). This is in fact a restatement of a finite horizon version of Bode’s sensitivity integral [14]. Since the determinant is the product of eigenvalues, we have:

$$\lambda_{\max}(\mathcal{M}) = \frac{\prod_{t=0}^T \det(S_t)}{\prod_{i=1}^{nT-1} \lambda_i(\mathcal{M})} = \frac{\prod_{i=2}^{nT} \lambda_i(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}))}{\prod_{t=0}^T \det(\Sigma_t)}$$

where the second equality follows from the fact that eigenvalues of \mathcal{M} are reciprocal to the eigenvalues of $\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X})$. Thus, $\lambda_{\max}(\mathcal{M}) \propto \prod_{i=2}^{nT} \lambda_i(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}))$, the product of the largest $nT - 1$ largest eigenvalues of $\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X})$. Hence, minimizing $\lambda_{\max}(\mathcal{M})$ is equivalent to minimizing $\prod_{i=2}^{nT} \lambda_i(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}))$, which corresponds to our intuition of making the covariance of the trajectories “small”.

An alternate interpretation is that by minimizing $\lambda_{\max}(\mathcal{M})$ or equivalently maximizing $\lambda_{\min}(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}))$, we are effectively minimizing the spread of eigenvalues (since the product of eigenvalues is fixed) and the optimum is achieved when all the eigenvalues are equal (assuming we could change the eigenvalues to any number we desire). Of course, there is no value of \mathbf{K} that would achieve this in general since we can only choose \mathbf{K} and not the eigenvalues themselves directly. However, the optimization criterion we chose will tend to promote uniformity of the eigenvalues. As a concrete example, in figure 1, we compare the eigenvalue range for $\text{Cov}_{P_0}(\mathbf{X})$ and $\text{Cov}_{P_{\mathbf{K}^*}}(\mathbf{X})$ where \mathbf{K} is the optimal solution found by our framework for the simple pendulum system described in section V-A. The

results clearly show that the eigenvalue spread for $P_{\mathbf{K}}$ (the controlled system) is much smaller than that for P_0 (the uncontrolled system).

If we had a state cost of the form $\sum_t \mathbf{x}_t^T \mathbf{x}_t = \mathbf{X}^T \mathbf{X}$, its expected values under the distribution $P_{\mathbf{K}}$ is

$$E_{P_{\mathbf{K}}} \left[\sum_t \mathbf{x}_t^T \mathbf{x}_t \right] = \text{tr} \left(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}) \right) = \sum_{i=1}^{nT} \lambda_i \left(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}) \right).$$

We know that

$$\det \left(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}) \right) = \prod_{i=1}^{nT} \lambda_i \left(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}) \right) = \prod_{i=0}^{nT-1} \det(\Sigma_t)$$

Under this constraint, the above cost is minimized when all the eigenvalues are equal.

Thus, although the objective is not formally equivalent to any of the classical choices (quadratic costs in LQR/ \mathcal{H}_2 , \mathcal{H}_{∞} etc.), the above arguments indicate that it is a sensible choice.

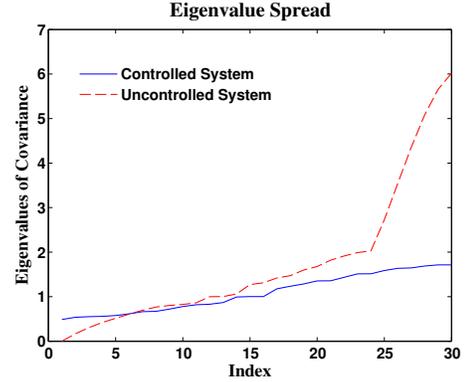


Fig. 1. Eigenvalue Spread

B. Penalizing Control Effort

So far we have only talked about states and not penalized control effort explicitly. This is required for both implementation (control limits) and numerical (conditioning) reasons. In this framework, we propose to do this by penalizing the KL-divergence between the controlled dynamics (with feedback gains \mathbf{K}) and the uncontrolled dynamics (no feedback). This criterion has been used previously in the linearly solvable MDP framework [15] and shown to be related to quadratic control costs for systems with linear dynamics and Gaussian noise. More concretely, it can be shown [16] that

$$\begin{aligned} & \text{KL}(P_0 \parallel P_{\mathbf{K}}) \\ &= E_{P_0} \left[\sum_{t=1}^T \text{KL}(P_0(\mathbf{x}_{t+1} | \mathbf{x}_t) \parallel P_{\mathbf{K}}(\mathbf{x}_{t+1} | \mathbf{x}_t)) \right] \\ &= E_{\mathbf{X} \sim P_0, \mathbf{u}_t = \mathbf{K}_t \mathbf{x}_t} \left[\sum_{t=1}^{T-1} \frac{1}{2} \mathbf{u}_t^T B_t^T \Sigma_t^{-1} B_t \mathbf{u}_t \right] \end{aligned}$$

Comparing this to the standard LQR control cost $\frac{\mathbf{u}_t^T R_t \mathbf{u}_t}{2}$, we see that this is the expected control cost with $R_t =$

$B_t^T \Sigma_t^{-1} B_t$, but for trajectories sampled from the uncontrolled dynamics. Thus, it is not the same a standard quadratic control cost used in the LQR framework, but has a similar effect of choosing feedback gains that do not produce excessively large controls. Using the formula for KL-divergence between Gaussian densities, it can be shown that the KL divergence between P_0 and $P_{\mathbf{K}}$ is given by

$$\text{KL}(P_0 \parallel P_{\mathbf{K}}) = \text{tr} \left(\mathcal{M} \left(\begin{matrix} \text{Cov}_{P_0}(\mathbf{X}) \\ P_0 \end{matrix} \right) \right) - nT.$$

Note that this is a linear (and hence convex) function of \mathcal{M} .

C. Overall optimization problem

The overall control design problem can now be cast as follows:

$$\begin{aligned} & \text{Minimize } \underbrace{\lambda_{\max} \left(\left(\begin{matrix} \text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}) \\ P_{\mathbf{K}} \end{matrix} \right)^{-1} \right)}_{\text{Minimize Trajectory Variance}} + \underbrace{\mu \text{KL}(P_0 \parallel P_{\mathbf{K}})}_{\text{Minimize Control Effort}} \\ & \text{Subject to } \mathbf{K} \in \mathcal{C} \end{aligned} \quad (1)$$

where \mathcal{C} is a convex set encoding the structural constraints on \mathbf{K} and $\mu > 0$ is a regularization parameter controlling the trade-off between performance (small covariance of trajectories) and control effort.

D. More General Objective Functions

The results we develop here are applicable more generally. The general form of the problem we solve is:

$$\begin{aligned} & \text{Minimize } \mathcal{F} \left[\left(\begin{matrix} \text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}) \\ P_{\mathbf{K}} \end{matrix} \right)^{-1} \right] \\ & \text{Subject to } \mathbf{K} \in \mathcal{C} \end{aligned} \quad (2)$$

where $\mathcal{F}[M]$ is an arbitrary *convex* function of M and *increasing* with respect to the \succeq -semidefinite ordering and \mathcal{C} is an arbitrary convex set encoding constraints on the feedback. We show (in theorem 3.1) that this problem can be solved via convex programming (specifically semidefinite programming) techniques. Since $\lambda_{\max}(\mathcal{M})$ and $\text{KL}(P_0 \parallel P_{\mathbf{K}}) = \text{tr}(\mathcal{M}(\text{Cov}_{P_0}(\mathbf{X}))) - nT$ are convex and \succeq -increasing in \mathcal{M} , so is their sum and hence (1) is a special case of this problem.

Again, the fact that $\mathcal{F}[M]$ is monotone-increasing suggest that minimizing it will tend to make $\mathcal{M} = (\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}))^{-1}$ small or $\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X})$ large. A concrete class of functions satisfying this are convex, symmetric and increasing functions of eigenvalues:

$$\mathcal{F}[\mathcal{M}] = f(\lambda_1(\mathcal{M}), \dots, \lambda_{nT}(\mathcal{M}))$$

with f convex and increasing in each of its arguments. We know that $\det \mathcal{M} = \prod_{i=1}^{nT} \lambda_i(\mathcal{M}) = \prod_{t=0}^{nT-1} \det(S_t)$. Relaxing the problem assuming that we can free pick $y_i = \log(\lambda_i(\mathcal{M}))$ (without the restriction that y_i are logarithms of eigenvalues of \mathcal{M}) results in the following problem:

$$\begin{aligned} & \text{Minimize}_y f(\exp(y_1), \exp(y_2), \exp(y_3), \dots, \exp(y_{nT})) \\ & \text{Subject to } \sum_i y_i = \sum_{i=0}^{nT-1} \log(\det(S_t)) \end{aligned}$$

This is a convex optimization problem since f is itself convex and increasing, and composed with a convex function (the \exp function). Writing down the KKT conditions for this problem, we get $\partial_i f(y) \exp(y_i) = \lambda$ where ∂_i denotes the partial derivative with respect to the i -th coordinate and λ is the Lagrange multiplier corresponding to the constraint. Since f is symmetric in its coordinates, $\partial_i f$ is symmetric as well, and hence the y_i are all equal at the optimum. This shows that for this general class of functions, the relaxed problem of optimizing over the log-eigenvalues of \mathcal{M} (assuming they are free variables) leads to an optimal solution where the eigenvalues of \mathcal{M} are all equal. As argued in section II-A, this optimum also minimizes $\mathbb{E}[\sum_t \mathbf{x}_t^T \mathbf{x}_t]$.

The restriction is that we cannot treat the eigenvalues as free variables and depending on the choice of f and problem constraints (the values of A, B, C), a different optimum would be reached. However, this result hints that our optimization criterion is trying to do something similar to a standard quadratic cost criterion and would hopefully lead to satisfactory results in many problems of practical interest. We also believe that formal results characterizing the quality of the optimal solution with respect to the standard quadratic cost metric are possible and plan to pursue these results in future work. Further, the choice of f may well affect the quality of the final control solution, and since we have freedom in this regard, investigating the effects of various choice of f is another direction for further work.

III. MAIN TECHNICAL RESULTS

In this section, we present our main theorem proving the convexity of the general optimization problem (2), and hence the convexity of the problem (1) follows since it is a special case.

Theorem 3.1: The optimal solution to problem (2) can be obtained by solving the following convex programming problem:

$$\begin{aligned} & \text{Minimize}_{\{\mathcal{M}_i\}, \mathbf{K}} \mathcal{F}[\mathcal{M}] \\ & \text{Subject to } \mathcal{M} = \begin{bmatrix} S_0 + \mathcal{M}_1 & -\tilde{A}_1^T S_1 & 0 \dots \\ -S_1 \tilde{A}_1 & S_2 + \mathcal{M}_2 & -\tilde{A}_2^T S_2 \dots \\ 0 & -S_2 \tilde{A}_2 & \dots \\ \vdots & \vdots & \vdots \end{bmatrix} \\ & \begin{bmatrix} S_i^{-1} & \tilde{A}_i^T \\ \tilde{A}_i & \mathcal{M}_i \end{bmatrix} \succeq 0, i = 1, \dots, T-1 \\ & \tilde{A}_t = A_t + B_t K_t \\ & \mathbf{K} \in \mathcal{C}. \end{aligned}$$

Proof: Let $\mathbf{K}^*, \{\mathcal{M}_t^*\}$ be an optimal solution and let $\mathcal{M}^*, \{\tilde{A}_t^*\}$ be defined accordingly. Let $\tilde{\mathcal{M}}_t = \tilde{A}_t^{*T} S_t \tilde{A}_t^*$.

Construct $\tilde{\mathcal{M}}$ as follows:

$$\tilde{\mathcal{M}} = \begin{bmatrix} S_0 + \tilde{\mathcal{M}}_1 & -\tilde{\mathbf{A}}_1^T S_1 & 0 \dots \\ -S_1 \tilde{\mathbf{A}}_1 & S_2 + \tilde{\mathcal{M}}_2 & -\tilde{\mathbf{A}}_2^T S_2 \dots \\ 0 & -S_2 \tilde{\mathbf{A}}_2 & \dots \\ \vdots & \vdots & \vdots \end{bmatrix}.$$

Then $\mathbf{K}^*, \{\tilde{\mathcal{M}}_t\}, \tilde{\mathcal{M}}, \{\tilde{\mathbf{A}}_t^*\}$ satisfy the problem constraints. The last set of constraints implies (by Schur complements), that $\mathcal{M}_t^* \succeq \tilde{\mathbf{A}}_t^{*T} S_t \tilde{\mathbf{A}}_t^* = \tilde{\mathcal{M}}_t$. Since $\mathcal{M}_t^* \succeq \tilde{\mathcal{M}}_t$ for all t and the off-diagonal blocks of $\tilde{\mathcal{M}}$ and \mathcal{M}^* are identical, $\mathcal{M}^* \succeq \tilde{\mathcal{M}}$. Since $\mathcal{F}[\mathcal{M}]$ is \succeq -increasing, $\mathcal{F}[\mathcal{M}^*] \geq \mathcal{F}[\tilde{\mathcal{M}}]$. Since \mathcal{M}^* is optimal, this inequality can only be satisfied if $\mathcal{F}[\mathcal{M}^*] = \mathcal{F}[\tilde{\mathcal{M}}]$. Further, by construction, $\tilde{\mathcal{M}} = (\text{Cov}_{P_{\mathbf{K}^*}}(\mathbf{X}))^{-1}$ (see lemma 8.1). Finally, for any $\mathbf{K}' \in \mathcal{C}, \mathbf{K}' \neq \mathbf{K}^*$, one can construct $\tilde{\mathbf{A}}_t' = A_t + B_t \mathbf{K}_t$ and $\mathcal{M}'_t = \tilde{\mathbf{A}}_t'^T S_t \tilde{\mathbf{A}}_t'$ and \mathcal{M}' (as above) to satisfy the problem constraints. Again by construction, $\mathcal{M}' = (\text{Cov}_{P_{\mathbf{K}'}}(\mathbf{X}))^{-1}$. Since \mathcal{M}^* is optimal,

$$\begin{aligned} \mathcal{F} \left[\left(\text{Cov}_{P_{\mathbf{K}'}}(\mathbf{X}) \right)^{-1} \right] &= \mathcal{F}[\mathcal{M}'] \geq \mathcal{F}[\mathcal{M}^*] = \mathcal{F}[\tilde{\mathcal{M}}] \\ &= \mathcal{F} \left[\left(\text{Cov}_{P_{\mathbf{K}^*}}(\mathbf{X}) \right)^{-1} \right]. \end{aligned}$$

Hence, for any $\mathbf{K}' \in \mathcal{C}, \mathbf{K}' \neq \mathbf{K}^*$, $\mathcal{F} \left[(\text{Cov}_{P_{\mathbf{K}'}}(\mathbf{X}))^{-1} \right] \geq \mathcal{F} \left[(\text{Cov}_{P_{\mathbf{K}^*}}(\mathbf{X}))^{-1} \right]$. Hence, \mathbf{K}^* is an optimal solution for problem (2) as well. ■

In fact, we can prove a stronger result using concepts from matrix convexity: The function $\mathcal{F}[\mathcal{M}]$ can be directly shown to be a convex function of \mathbf{K} (theorem 8.2, section VIII). However, the above formulation allows us to reduce the general convex programming problem using to semidefinite programming, which allows us to formulate and solve them in modeling languages like cvx [17].

IV. EXPLOITING STRUCTURE: EFFICIENT ALGORITHMS

The SDPs we describe have a lot of structure: All the matrices involved have block-tridiagonal structure, which is a special case of a chordal sparsity structure [18], allowing for efficient computation of Cholesky factors and gradients/Hessians of log-barrier functions used in interior point methods for semidefinite programming. Thus, given the structure of the problem, we can compute a Newton step in the interior point method in time that grows linearly with T and cubically with n in the worst case - $O(n^3 T)$. This complexity is of the same order as solving time-varying Ricatti equations for a standard LQR problem. However, we will have to do this at every iteration of the interior point method - which typically converge in a small number of iterations (< 40 , largely independent of problem dimension).

Further, in many cases of practical interest, the matrices A, B, K would themselves have additional structure (sparsity etc.) which would allow further improvements in computation time. In this work, our focus is on the modeling and problem formulation. Hence, we do not investigate issues of computational complexity in depth and use off-the-shelf solvers [17]. However, this is an interesting area for future work and is critical to making this promising framework applicable to large-scale real-world distributed control problems.

V. NUMERICAL EXPERIMENTS

In this section we present numerical results that demonstrate that our framework can produce stabilizing controllers for dynamical systems under non-trivial structural constraints on the feedback matrix.

A. Mechanical Systems: Simple Pendulum

We first tested the algorithm on the simple pendulum: The system is characterized by the angular position and velocity $\mathbf{x} = [\theta; \dot{\theta}]$ and the control input \mathbf{u} is the external torque applied. Assuming a unit mass, the system dynamics is given by

$$\ddot{\theta} = u + g \sin(\theta).$$

where θ is measured clockwise from the upright position of the pendulum. The control task is to stabilize the pendulum at the topmost position, which is an unstable equilibrium. Linearizing about this equilibrium $\theta = 0, \dot{\theta} = 0$, the system dynamics is given by

$$\dot{\mathbf{x}} = \begin{pmatrix} 0 & 1 \\ g & 0 \end{pmatrix} \mathbf{x} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} \mathbf{u}.$$

Using an Euler-discretization with time-step h , the discrete-time dynamics is given by

$$A = I + h \begin{pmatrix} 0 & 1 \\ g & 0 \end{pmatrix}, B = h \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

We solve the optimization problem (1) with these matrices with the constraint that $K_1 = K_2 = \dots = K_T$ (the feedback is time-invariant), $T = 40$ and $h = .01$ (implying a 400 ms horizon). We compare it to the solution obtained by solving the infinite-horizon discrete-time algebraic ricatti equation (the optimal solution for the infinite horizon LQR problem).

The results for a particular initial displacement (1 radian) are plotted in figure 2. We get qualitatively similar behavior from both controllers. The controller produced by our framework takes longer to stabilize but it utilizes smaller (in magnitude) controls initially. The trade-off between controls and performance is different in our framework and in LQR and they cannot be directly compared - the illustration here is simply meant to show that our framework discovers a stabilizing controller with qualitatively similar behavior as an LQR controller.

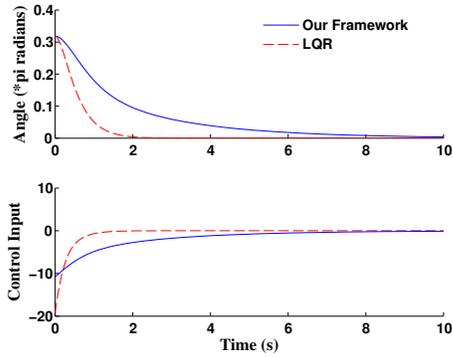


Fig. 2. Pendulum: LQR vs Our Framework

B. Power Systems

As a second application, we study distributed control of generators in a power system. The power grid is described by an undirected graph with nodes i representing generators or loads (or both) and edges representing transmission links between neighboring nodes. We denote by $\text{Neb}(i)$ the set of nodes that are neighbors of i in the graph. The model we consider is the structure preserving model [19]. Under this model, the dynamics of the voltage angles in the generators is given by

$$M_i \ddot{\theta} + D_i \dot{\theta} = p_i - |V_i| |V_j| \sum_{j \in \text{Neb}(i)} B_{ij} \sin(\theta_i - \theta_j) + G_{ij} \cos(\theta_i - \theta_j)$$

where θ_i refers to the voltage phase at node i , $|V_i|$ is the voltage magnitude, G_{ij}, B_{ij} are the electrical conductance and susceptance of the line connecting i and j , M_i is the inertia of the generator at node i , D_i is a damping term and p_i is the power due to mechanical torque (generated by the turbine in a generator or consumed by a motor load) at node i . We assume that the voltage magnitudes are fixed to unity and ignore the inertias to get an approximate first-order model of the system (akin to [20]) that still captures essential features of the dynamics. It has been argued that this model suffices to predict voltage angle-based instabilities in power systems. Ideally, we want the voltage angles to all be close to each other (for the generators to be “in-phase”). We linearize the dynamics about an equilibrium (computed by solving a quasi-static Optimal Power Flow (OPF) problem using [21]) and proceed with applying our control framework to this problem. The control variable is the power injection p_i , and we consider modifying that in response to changes in θ . We consider 3 different control schemes: a) Full Distributed - p_i changes based only on θ_i , b) Sparse p_i changes based on $\{\theta_i\} \cup \{\theta_j : j \in \text{Neb}(i)\}$ and c) Centralized control (p_i changes based on changes in all θ).

We impose the constraint that $|K_t| \leq 5$, entrywise, which corresponds to enforcing control limits within the expected operating conditions of the problem (reasonable ranges for

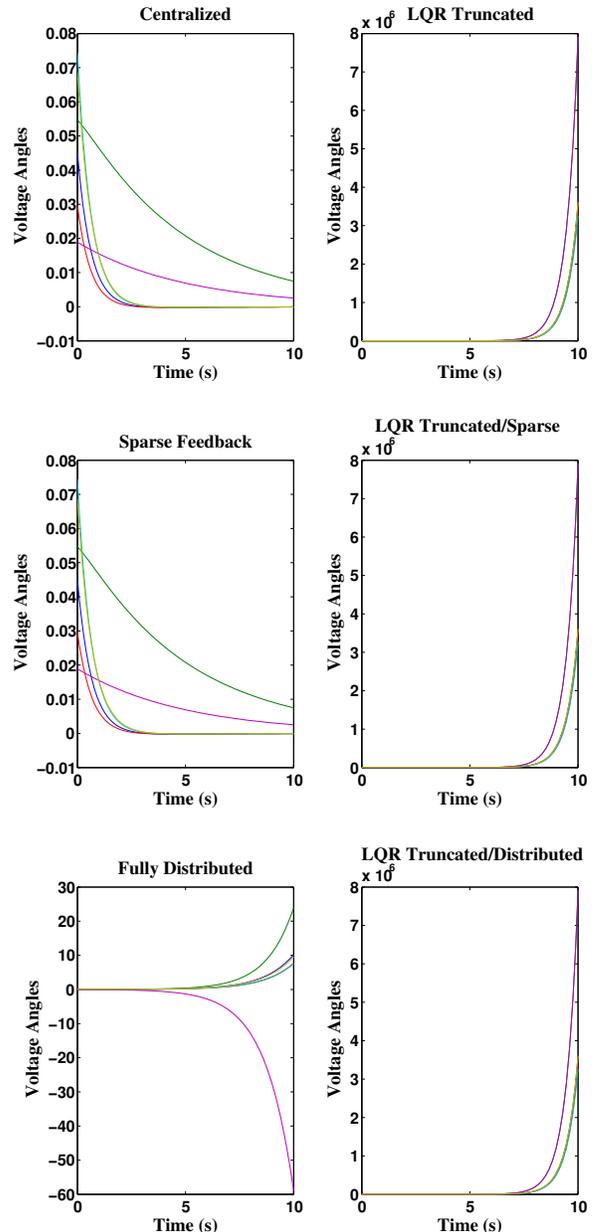


Fig. 3. Distributed Control in Power Systems: The different lines correspond to voltage phase angles of different generators in the system.

θ). Note that this limitation is not possible in the LQR framework, and hence we are forced to make the approximation of truncating the values of the feedback matrix generated by LQR to conform to these limits. The results (figure 3) show that our algorithm can do much better than the simple truncated LQR, which causes the voltage angles to go completely out of phase, while our control design is able to quickly bring the system back into synchronized operation (all angles equal) in the centralized and sparse cases. In the fully distributed case, our control design is unable to keep angle deviations to acceptable levels, but still performs much better than the LQR based design.

Thus, our framework is able to handle nontrivial con-

straints on the feedback matrices, which LQR and other dynamic programming based approaches cannot. Even by playing with the relative values of the state and control costs for LQR, we were unable to obtain a feedback matrix satisfying the constraints $|K_t| \leq 5$ and simple truncation based schemes produce bad results as shown in this example.

VI. 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 follows:

$$\begin{aligned} & \underset{K}{\text{Minimize}} \quad \| \text{Closed loop system with feedback } K \| \\ & \text{Subject to } K \text{ Stabilizing, } K \in \mathcal{C} \end{aligned}$$

where $\|\cdot\|$ is typically the \mathcal{H}_2 or \mathcal{H}_∞ 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 [6][5]. 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 Riccati equations for subsystems and combining the results. For the \mathcal{H}_∞ norm, a state-space solution using an LMI approach was developed in [22].

Another thread of work on decentralized control looks at special cases where dynamic programming techniques can be used in spite of the decentralization constraints. The advantage of these approaches is that they directly handle finite horizon and time-varying approaches. 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 for the case of 2 agents (a block-lower triangular structure in A, B with 2 blocks) can be solved via dynamic programming. In [10], the authors develop a dynamic programming solution that

generalizes this and applies to general “partially-nested” systems allowing for both sparsity and delays.

All the above 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.

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} are a special case of this. We can also handle time-varying linear systems. Although the objective is nonstandard, we have provided theoretical and numerical evidence that it is a sensible control objective.

In this paper, we have presented a formulation of the feedback synthesis problem that is convex under arbitrary convex constraints on the feedback matrix. Designing stabilizing feedback controllers under even simple constraints like bounds on coefficients of \mathbf{K} is known to be NP-hard [3], which was precisely the constraint we imposed in section V-B. Thus, it is clear that we cannot guarantee stability of controllers produced by our framework even for simple constraints on \mathbf{K} . Further, our framework is formulated in finite-horizon, where stabilizing constraints do not show up explicitly (although intuitively we would expect that making the covariance of the trajectories small would tend to produce stabilizing controllers). Further work needs to be done to understand the relationship between the covariance-based metric we have considered here and more standard metrics used in control ($\mathcal{H}_2, \mathcal{H}_\infty, \text{LQR}$ etc.).

VII. CONCLUSION

As argued above, the framework we have developed seems promising and overcomes limitations of previous works on convex optimization based controller synthesis. Although the control objective used is non-standard, we have argued why it is a sensible objective, and we also presented numerical examples showing that it produces controllers that generate “desirable” behavior. Additional work needs to be done to rigorously understand the relationship with more standard control objectives and whether solutions generated by our framework can have guarantees in terms of traditional metrics (stability etc.) We also have some flexibility in choosing the control objective in this framework: Investigating the effects of this choice is another topic worth further investigation. Finally, we hope to develop solvers that exploit the structure of the problem (referred to in section IV) and scale to large-scale real-world control problems.

VIII. ACKNOWLEDGEMENTS

This work was supported by the NSF. We would like to thank Mehran Mesbahi for helpful comments and feedback on this manuscript.

- [1] R. E. Kalman *et al.*, “A new approach to linear filtering and prediction problems,” *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [2] H. S. Witsenhausen, “A counterexample in stochastic optimum control,” *SIAM Journal on 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 Journal on Control and Optimization*, vol. 35, no. 6, pp. 2118–2127, 1997.
- [4] M. Rotkowitz and S. Lall, “Decentralized control information structures preserved under feedback,” in *Proceedings of the 41st IEEE Conference on Decision and Control*, vol. 1, 2002, 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 Transactions on Automatic Control*, vol. 49, no. 10, pp. 1623–1640, 2004.
- [6] 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.
- [7] P. Shah, “H2-optimal decentralized control over posets: A state-space solution for state-feedback,” *IEEE Transactions on Automatic Control*, vol. PP, no. 99, pp. 1–1, 2013.
- [8] C.-H. Fan, J. L. Speyer, and C. R. Jaensch, “Centralized and decentralized solutions of the linear-exponential-gaussian problem,” *IEEE Transactions on Automatic Control*, vol. 39, no. 10, pp. 1986–2003, 1994.
- [9] J. Swigart and S. Lall, “An explicit dynamic programming solution for a decentralized two-player optimal linear-quadratic regulator,” in *Proceedings of mathematical theory of networks and systems*, 2010.
- [10] A. Lamperski and L. Lessard, “Optimal decentralized state-feedback control with sparsity and delays,” *arXiv preprint arXiv:1306.0036*, 2013.
- [11] F. Lin, M. Fardad, and M. R. Jovanović, “Design of optimal sparse feedback gains via the alternating direction method of multipliers,” *IEEE Trans. Automat. Control*, vol. 58, no. 9, pp. 2426–2431, September 2013.
- [12] P. Apkarian, D. Noll, and A. Rondepierre, “Mixed h_2/h_∞ control via nonsmooth optimization,” *SIAM Journal on Control and Optimization*, 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,” in *Fifth IFAC Symposium on Robust Control Design, Toulouse*, 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] E. Todorov, “Efficient computation of optimal actions,” *Proceedings of the national academy of sciences*, vol. 106, no. 28, pp. 11 478–11 483, 2009.
- [16] H. J. Kappen, V. Gómez, and M. Opper, “Optimal control as a graphical model inference problem,” *Machine learning*, vol. 87, no. 2, pp. 159–182, 2012.
- [17] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming, version 2.0 beta,” <http://cvxr.com/cvx>, Sep. 2012.
- [18] M. S. Andersen, J. Dahl, and L. Vandenberghe, “Implementation of nonsymmetric interior-point methods for linear optimization over sparse matrix cones,” *Mathematical Programming Computation*, vol. 2, no. 3, pp. 167–201, 2010.
- [19] A. Bergen and D. J. Hill, “A structure preserving model for power system stability analysis,” *Power Apparatus and Systems, IEEE Transactions on*, no. 1, pp. 25–35, 1981.
- [20] F. Dorfler and F. Bullo, “Synchronization and transient stability in power networks and nonuniform kuramoto oscillators,” *SIAM Journal on Control and Optimization*, vol. 50, no. 3, pp. 1616–1642, 2012.
- [21] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, “Matpower: Steady-state operations, planning, and analysis tools for power systems research and education,” *Power Systems, IEEE Transactions on*, vol. 26, no. 1, pp. 12–19, 2011.
- [22] C. W. Scherer, “Structured \mathcal{H}_∞ -Optimal Control for Nested Interconnections: A State-Space Solution,” *ArXiv e-prints: arXiv:1305.1746*, May 2013.
- [23] A. Ben-Tal and A. Nemirovskii, “Lectures on modern convex optimization: analysis, algorithms, and engineering applications,” <http://www2.isye.gatech.edu/~nemirovs/Lect.ModConvOpt.pdf>, Jun. 2012.

Lemma 8.1: Consider the discrete time dynamical system

$$\mathbf{x}_1 \sim \Sigma_0, \mathbf{x}_{t+1} = A_t \mathbf{x}_t + B_t K_t \mathbf{x}_t + \epsilon_t, \epsilon_t \sim \mathcal{N}(0, \Sigma_t)$$

$t = 1, \dots, T - 1$. Let $\tilde{A}_t = A_t + B_t K_t$. Then, the distribution of trajectories $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]$ is jointly Gaussian: $\mathcal{N}(0, \mathcal{M}^{-1})$ where

$$\mathcal{M} = \begin{bmatrix} S_0 + \tilde{A}_1^T S_1 \tilde{A}_1 & -\tilde{A}_1^T S_1 & 0 & \dots \\ -S_1 \tilde{A}_1 & S_1 + \tilde{A}_2^T S_2 \tilde{A}_2 & -\tilde{A}_2^T S_2 & \dots \\ 0 & -S_2 \tilde{A}_2 & \dots & \dots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix}$$

and $S_t = \Sigma_t^{-1}$. Further, $\det(\mathcal{M}) = \prod_{t=1}^{T-1} \det(S_t)$.

Proof: By the Markov Property,

$$\begin{aligned} P(\mathbf{X}) &= P(\mathbf{x}_1) \prod_{t=1}^{T-1} P(\mathbf{x}_{t+1} | \mathbf{x}_t) \\ &= \frac{\exp\left(-\frac{\mathbf{x}_1^T S_0 \mathbf{x}_1}{2}\right)}{\sqrt{(2\pi)^n \det(\Sigma_0)}} \prod_{t=1}^{T-1} \frac{\exp\left(-\frac{(\mathbf{x}_{t+1} - \tilde{A}_t \mathbf{x}_t)^T S_t (\mathbf{x}_{t+1} - \tilde{A}_t \mathbf{x}_t)}{2}\right)}{\sqrt{(2\pi)^n \det(\Sigma_t)}} \\ &= \frac{\exp\left(-\frac{\mathbf{X}^T \mathcal{M} \mathbf{X}}{2}\right)}{\sqrt{\prod_{t=0}^{T-1} (2\pi)^n \det(\Sigma_t)}}. \end{aligned}$$

From the numerator, we see that the resulting density has a Gaussian form with inverse covariance \mathcal{M} . It follows that the normalizing constant is $\sqrt{(2\pi)^{nT} \det(\mathcal{M}^{-1})}$. Matching this with the denominator below, we get

$$\det(\mathcal{M}) = \frac{1}{\prod_{t=0}^{T-1} \det(\Sigma_t)} = \prod_{t=0}^{T-1} \det(S_t).$$

Theorem 8.2: Supposed that $\mathcal{F}[M]$ is a convex function of $M \in \mathbf{R}^{nT \times nT}$ for $M \succeq 0$ and is increasing with respect to the \succeq -ordering, that is, $\mathcal{F}[M] \geq \mathcal{F}[M']$ whenever $M \succeq M'$. Then the function $\mathcal{F}\left[(\text{Cov}_{P_{\mathbf{K}}}(\mathbf{X}))^{-1}\right]$ is a convex function of \mathbf{K} .

Proof: The proof uses the concept of convexity with respect to the \succeq -cone [23] (chapter 3, section 3.8). The block super-diagonal and sub-diagonal elements of \mathcal{M} are \succeq -convex since they are linear functions of \mathbf{K} . The block-diagonal elements are of the form $M^T Q M$, $Q \succeq 0$, which are known to be \succeq -convex themselves. Thus, $\mathcal{M}(\mathbf{K})$ is a sum of \succeq -convex functions, and hence itself a \succeq -convex function of \mathbf{K} . Since $\mathcal{F}[M]$ is a convex and \succeq -increasing function of M , the composition $\mathcal{F}[\mathcal{M}(\mathbf{K})]$ is a convex function of \mathbf{K} . ■