Convex Risk Averse Control Design

Krishnamurthy Dvijotham, Maryam Fazel and Emanuel Todorov

Abstract— In this paper, we show that for arbitrary stochastic linear dynamical systems, the problem of optimizing parameters of a feedback control policy can be cast as a convex optimization problem when a risk-averse objective (similar to LEQG) is used. The only restriction is a condition relating the control cost, risk factor and noise in the system. The resulting approach allows us to synthesize risk-averse controllers efficiently for finite horizon problems. For the standard quadratic costs in infinite horizon, the resulting problems become degenerate if the uncontrolled system is unstable. As an alternative, we propose using a discount-based approach that ensures that costs do not blow up. We show that the discount factor can effectively be used as a homotopy parameter to gradually synthesize stabilizing controllers for unstable systems. We also propose extensions where non-quadratic costs can be used for controller synthesis, and in this case, as long as the costs are bounded, the optimization problems are well-posed and have non-trivial solutions.

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 [K+60]. 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structure arising from the need to implement control in a decentralized fashion), the dynamic programming approach is no longer applicable.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 [Wit68] and that the general problem of designing linear feedback gains subject to constraints is NP-hard [BT97].

Previous approaches to synthesizing structured controllers can be broadly categorized into three types: Frequency Domain [RL02][QSVK04][RL06][Sha13], Dynamic Programming [FSJ94][SL10][LL13] and Nonconvex optimization [WZ13][LFJ13][ANR08][BHLO06]. 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 subject to constraints on the controller, using nonconvex optimization

K. Dvijotham is with the Department of Computing and Mathematical Sciences, Caltech, Pasadena, California `dvij@cs.washington.edu`

M. Fazel is with the Department of Electrical Engineering, University of Washington, Seattle, Washington `mfazel@uw.edu`

E. Todorov is with the Department of Computer Science and Engineering & the Department of Applied Mathematics, University of Washington, Seattle, Washington `todorov@cs.washington.edu`

techniques. These are generally applicable, but are susceptible to local minima and slow convergence.

Recently, there has been work on trying to solve the structured controller design problem directly as a convex optimization problem. In [FML][Lav13], the authors formulate the decentralized controller problem as a rank-constrained optimization problem and develop SDP relaxations of the optimal decentralized control problem. If the SDP relaxation has a rank-1 solution, the relaxation is exact. This is not true in general but the authors prove that a solution of rank at most 3 can be found. They then propose various heuristics for obtaining an approximately optimal rank-1 solution based on the rank-3 solution, but no guarantees can be made about the quality of the solution in the worst case.

Along similar lines, we have developed alternate convex objectives for structured controller design [DTTF13] [DTF13] that replace the standard LQR/LEQG objectives with different objectives that lead to convex optimization formulations of structured controller design. We also provided theoretical and empirical evidence of the effectiveness of the convex surrogates and their relationship to the standard control objectives.

In this paper, we revisit the LEQG objective and show that without any modifications to the objective, the controller design problem is convex under *arbitrary* convex constraints on the feedback matrices, provided that a condition relating the control cost, risk factor and noise covariance is satisfied.

Apart from [DTTF13], we are not aware of any other work that attempts to solve the structured controller design problem in a risk-averse setting. Even in [DTTF13], risk-aversion was combined with a particular adversarial perturbation model in order to obtain a convex controller synthesis approach. In this work, we prove convexity of the controller design problem for the standard LEQG objectives provided a condition relating the risk factor, control cost and noise is satisfied. We extend the approach to the infinite horizon setting. Given the hardness of designing stabilizing controllers, our approach understandably breaks down in infinite horizon setting leading to provably infeasible problems if the uncontrolled system is unstable. However, as an alternative, we propose using a discounting approach that ensures that costs remain bounded even in infinite horizon. We show that the discount factor can effectively be used as a continuation parameter to synthesize stabilizing controllers by starting with a small discount factor (heavy discounting) and gradually increasing it until a stabilizing controller is obtained. We verify our results numerically on simplified power systems models.

II. FINITE HORIZON PROBLEMS

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

$$x_1 = \epsilon_0, x_{t+1} = A_t x_t + B_t u_t + \epsilon_t, \quad t = 1, 2, \ldots, N-1.$$

Here $t = 0, 1, 2, \ldots, N$ is the discrete time index, $x_t \in \mathbf{R}^n$ is the plant state, $\epsilon_t \in \mathbf{R}^n$ is an exogenous disturbance and $u_t \in \mathbf{R}^{n_u}$ is the control input. For simplicity, we will work with static state feedback $u_t = K_t x_t$ although our approach extends to dynamic feedback as well. Let $\mathbf{K} = \{K_t : t = 1, 2, \ldots, N-1\}$ and denote the closed-loop system dynamics by

$$\tilde{A}_t(K_t) = A_t + B_t K_t.$$

I denotes the identity matrix. Boldface upper case letters denote trajectories:

$$\mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_N \end{pmatrix}, \boldsymbol{\epsilon} = \begin{pmatrix} \epsilon_0 \\ \vdots \\ \epsilon_{N-1} \end{pmatrix}$$

$\mathcal{N}(\mu, \Sigma)$ denotes a Gaussian distribution with mean μ and covariance matrix Σ. The notation $\|y\|_M$ denotes the norm under the metric induced by the positive definite matrix $M \succ 0$, that is, $\|y\|_M = \sqrt{y^T M y}$.

We start by stating the problem in finite horizon:

$$\underset{\mathbf{K}=\{K_t\}}{\text{Minimize}} \frac{1}{\alpha} \log\left(\underset{\boldsymbol{\epsilon}}{\mathrm{E}}\left[\exp\left(\alpha\mathcal{L}\right)\right]\right) \qquad (1)$$

$$\text{Subject to } x_1 = \epsilon_0, x_{t+1} = A_t x_t + B_t K_t x_t + \epsilon_t$$

$$\epsilon_t \sim \mathcal{N}(0, \Sigma_t)$$

$$\mathcal{L} = \sum_{t=1}^{N} \ell_t(x_t) + \sum_{t=1}^{N-1} \frac{u_t^T R_t u_t}{2}$$

$$\mathbf{K} \in \mathcal{C}$$

where $\mathcal{C}$ represents the structural constrains on our feedback matrices and $\ell_t : \mathcal{X} \to \mathbf{R}$ is an arbitrary state cost function. This is a generalized version of the LEQG objective [SDJ74] where state costs are allowed to be non-quadratic.

Theorem 2.1 (Convexity of Risk-Averse Control):
Suppose that

$$\Sigma_t \succeq \frac{1}{\alpha} B_t R_t^{-1} B_t^T, t = 1, \ldots, N-1$$

and $\mathcal{C}$ is a convex set. Then, the optimization problem (1) is convex.

Proof: Since $\Sigma_t \succeq \frac{1}{\alpha} B_t R_t^{-1} B_t^T$, let $\Sigma_t' = \Sigma_t - \frac{1}{\alpha} B_t R_t^{-1} B_t^T \succeq 0$. Consider the following system:

$$x_{t+1} = A_t x_t + B_t\left(K_t x_t + \omega_t\right) + \epsilon_t'$$

$$\epsilon_t' \sim \mathcal{N}(0, \Sigma_t'), \omega_t \sim \mathcal{N}\left(0, \frac{1}{\alpha} R_t^{-1}\right) \qquad (2)$$

with ϵ_t, ω_t statistically independent. Define $\epsilon_t = \epsilon_t' + B_t \omega_t$ so that $\epsilon_t \sim \mathcal{N}\left(0, \Sigma_t' + \frac{1}{\alpha} B_t R_t^{-1} B_t^T\right) = \mathcal{N}(0, \Sigma_t)$ (since

the covariance matrices of independent random variables add up). Thus, the above dynamics is equivalent to

$$x_{t+1} = A_t x_t + B_t K_t x_t + \epsilon_t$$

$$\epsilon_t \sim \mathcal{N}(0, \Sigma_t)$$

which is the dynamics stated in (1). Thus, we can work with the modified system (2). Now, define $y_t = K_t x_t + \omega_t$ so that $y_t \sim \mathcal{N}\left(K_t x_t, \frac{1}{\alpha} R_t^{-1}\right)$. Then we can write down the dynamics as

$$x_{t+1} = A_t x_t + B_t y_t + \epsilon_t'$$

$$\epsilon_t' \sim \mathcal{N}(0, \Sigma_t'), y_t \sim \mathcal{N}(K_t x_t, \Sigma_t)$$

Note that with this model, given $\boldsymbol{\epsilon}'$, x_t is a deterministic function of $\mathbf{y}$ and does not depend on K_t. Writing out the expectation in (1) explicitly, we get

$$\underset{\epsilon_t \sim \mathcal{N}(0, \Sigma_t)}{\mathrm{E}} \left[\exp\left(\alpha\mathcal{L}(\mathbf{K})\right)\right] =$$

$$\underset{\omega_t \sim \mathcal{N}\left(0, \frac{1}{\alpha} R_t^{-1}\right), \epsilon_t' \sim \mathcal{N}(0, \Sigma_t')}{\mathrm{E}} \left[\exp\left(\alpha\mathcal{L}(\mathbf{K})\right)\right] =$$

$$\underset{y_t \sim \mathcal{N}\left(K_t x_t, \frac{1}{\alpha} R_t^{-1}\right), \epsilon_t' \sim \mathcal{N}(0, \Sigma_t')}{\mathrm{E}} \left[\exp\left(\alpha\mathcal{L}(\mathbf{K})\right)\right]$$

If we fix $\boldsymbol{\epsilon}'$, the expectation over $\mathbf{y}$, upto a constant factor, is equal to

$$\int \exp\left(\sum_{t=1}^{N-1} \frac{\|K_t x_t\|_{\alpha R_t}^2 - \|y_t - K_t x_t\|_{\alpha R_t}^2}{2} + \alpha\ell(x_t)\right) d\mathbf{y}$$

where the first term comes from the Gaussian density of y_t and the remaining terms from $\mathcal{L}(\mathbf{K})$. The term inside the exponent depends on K_t as

$$\frac{K_t^T\left(\alpha R_t - \alpha R_t\right) K_t}{2} + y_t^T K_t x_t$$

which is a linear (and hence convex) of K_t. Thus, the objective is a convex function of $\mathbf{K}$ by composition rules (since $\frac{1}{\alpha} \log\left(\mathrm{E}_X\left[\exp\left(\alpha z f(X)\right)\right]\right)$ is a convex and increasing function of z) for any fixed $\boldsymbol{\epsilon}'$. Taking expectations with respect to $\boldsymbol{\epsilon}'$, we get that the original objective is convex in $\mathbf{K}$. Hence the theorm. ∎

III. INFINITE HORIZON PROBLEMS

In infinite horizon, we cannot obtain finitely parameterized optimal controllers unless the system is time-invariant. Thus, we restrict ourselves to time-invariant systems of the form

$$x_1 = \epsilon_0, x_{t+1} = Ax_t + BKx_t + \epsilon_t, t \geq 1$$

and study the infinite horizon average cost problem

$$\underset{K}{\text{Minimize}} \lim_{N \to \infty} \frac{1}{N} \frac{1}{\alpha} \log\left(\underset{\boldsymbol{\epsilon}}{\mathrm{E}}\left[\exp\left(\alpha\mathcal{L}_N\right)\right]\right) \qquad (3)$$

$$\text{Subject to } x_1 = \epsilon_0, x_{t+1} = Ax_t + BKx_t + \epsilon_t$$

$$\epsilon_t \sim \mathcal{N}(0, \Sigma)$$

$$\mathcal{L}_N = \sum_{t=1}^{N} \ell(x_t) + \sum_{t=1}^{N-1} \frac{u_t^T R u_t}{2}$$

$$\mathbf{K} \in \mathcal{C}$$

It is easy to see that the objective satisfies the following policy-specific Bellman equation:

$$v_K (x) + c = \ell (x) + \frac{\|Kx\|_R^2}{2} + \frac{1}{\alpha} \log \left(\underset{x'}{\mathrm{E}} \left[\exp \left(\alpha v_K \left(x' \right) \right) \right] \right)$$
(4)

where $x' \sim \mathcal{N} \left((A + BK) x, \Sigma \right)$ and c is a constant representing the average cost. We introduce a new discounted version of the Bellman equation:

$$v_K (x) = \ell (x) + \frac{\|Kx\|_R^2}{2} + \frac{1}{\alpha} \log \left(\underset{x'}{\mathrm{E}} \left[\exp \left(\alpha \gamma v_K \left(x' \right) \right) \right] \right)$$
(5)

where $\gamma < 1$ is a discount factor, that decreases the weight on future costs geometrically. In the limit $\alpha \to 0$, this reduces to the Bellman equation for the standard discounted cost setting $\sum_t \gamma^t \ell (x_t, u_t)$. However, when $\alpha > 0$, it leads to a different discounting where the overall objective is given by a cascaded form:

$$\mathrm{E} \left[\ell (x_1, u_1) + \frac{1}{\alpha} \log \left(\mathrm{E} \left[\alpha \gamma \exp \left(\ell (x_2, u_2) + \gamma \alpha \log (\ldots) \right) \right] \right) \right]$$

where $u_t = Kx_t$ and $\ell (x_t, u_t) = \ell (x_t) + \frac{\|u\|_R^2}{2}$. In the limit $\gamma \to 1$, the solution of this problem converges to the solution of (4). This falls into the class of iterated risk measures [Oso12] which have several appealing properties such as time-consistency.

Theorem 3.1 (Convexity in infinite horizon): Let V_K be the solution to (5). The resulting optimization problem $\min_{K \in \mathcal{C}} V_K (x)$ is a convex optimization problem for all x if

$$\Sigma \succeq \frac{1}{\alpha} BR^{-1} B^T$$

Proof: By the same argument as the previous theorem, given the condition of the theorem, we can rewrite the original dynamical system as

$$x_{t+1} = Ax + B \left(Kx + \omega \right) + \epsilon'$$

$$\omega \sim \mathcal{N} \left(0, \frac{1}{\alpha} R^{-1} \right), \epsilon' \sim \mathcal{N} \left(0, \Sigma_t - \frac{1}{\alpha} BR^{-1} B^T \right).$$

Plugging this into the Bellman equation (5), we get

$$v_K (x) = \ell (x) + \frac{\|Kx\|_R^2}{2}$$
$$+ \frac{1}{\alpha} \log \left(\underset{\epsilon', \omega}{\mathrm{E}} \left[\exp \left(\alpha \gamma v_K \left(Ax + B \left(Kx + \omega \right) + \epsilon' \right) \right) \right] \right)$$

The RHS can be rewritten as (excluding $\ell (x)$)

$$\frac{1}{\alpha} \log \left(\underset{\epsilon', y}{\mathrm{E}} \left[\exp \left(\alpha \gamma v_K \left(Ax + By + \epsilon' \right) + \frac{\alpha}{2} \|Kx\|_R^2 \right) \right] \right)$$

where $y \sim \mathcal{N} \left(Kx, \frac{1}{\alpha} R^{-1} \right)$. Keeping ϵ' fixed, the term inside the exponent depends on K as

$$-\frac{\|y - Kx\|_S^2}{2} + v_K \left(Ax + By + \epsilon' \right) + \frac{\alpha}{2} \|Kx\|_R^2$$

where $S = (\Sigma^u)^{-1}$. Since $\alpha R \succeq (\Sigma^u)^{-1}$, the sum of the first and third term is convex in K. Consider the operator defined by the RHS of (5):

$$\Gamma[f] (x) = \ell (x) + \frac{1}{\alpha} \log \left(\mathrm{E} \left[\exp \left(\alpha \gamma f \left(x' \right) \right) \right] \right).$$

In lemma 1(appendix), we show that this is a contraction mapping, so that starting with any f^0, $\lim_{r \infty} \Gamma^r[f^0]$ is the unique solution to v_K of (5). We can take f^0 to be a constant function which is obviously convex in K for every x. Define $f^r = \Gamma^r[f]$. Suppose that f^r is convex in K for all x for some r. Then, so is

$$-\frac{\|y - Kx\|_S^2}{2} + f^r \left(Ax + By + \epsilon' \right) + \frac{\alpha}{2} \|Kx\|_R^2$$

by the above argument. Taking expectation wrt ϵ and by composition with $\log (\mathrm{E} [\exp (\cdot)])$, the overall mapping is convex and hence so is $f^{r+1} (x)$. The argument is valid for any x, so $f^{r+1} (x)$ is convex in K for all x. Thus, by induction, the solution v_K to the Bellman equation is convex in K for all x. Hence the theorem. ∎

IV. LEQG PROBLEMS

In this section, we apply the above results to LEQG problems where the state cost has a quadratic form $\ell (x) = \frac{x^T Q x}{2}$.

Theorem 4.1 (Bellman Equation for LEQG Problem): If $\ell (x) = \frac{x^T Q x}{2}$, the discounted Bellman equation 5 reduces to $v_K (x) = \frac{x^T V x}{2} + c$ where V satisfies

$$V = Q + K^T RK + \tilde{A}^T \left((\gamma V)^{-1} - \alpha \Sigma \right)^{-1} \tilde{A}$$
(6)

where $\tilde{A} = A + BK$. $c = -\frac{\log \left(\det \left(I - \alpha \gamma \Sigma^{1/2} V \Sigma^{1/2} \right) \right)}{2\alpha(1-\gamma)}$.

Proof: In appendix. ∎

We can now rewrite the control design problem as a determinant maximization problem:

Theorem 4.2: If $\ell (x) = \frac{x^T Q x}{2}$ and $\Sigma = \frac{1}{\alpha} BR^{-1} B^T$, the problem of minimizing $v_K (0)$ reduces to solving the following convex optimization problem:

$$\max_{M, K} \log \left(\det \left(I - \alpha \gamma \Sigma^{1/2} V \Sigma^{1/2} \right) \right)$$
(7)

$$V \succeq Q + \begin{pmatrix} A \\ -RK \end{pmatrix}^T \begin{pmatrix} (\gamma V)^{-1} & -B \\ -B^T & R \end{pmatrix}^{-1} \begin{pmatrix} A \\ -RK \end{pmatrix}$$

$$K \in \mathcal{C}$$

Further, the above problem can be rewritten more explicitly as an LMI:

$$\max_{V, K} \log \left(\det \left(I - \alpha \gamma \Sigma^{1/2} V \Sigma^{1/2} \right) \right)$$
(8)

$$\begin{pmatrix} V - Q - \gamma A^T V A & RK - \gamma B^T V A^T \\ (RK - \gamma B^T V A)^T & R - B^T (\gamma V)^{-1} B \end{pmatrix} \succeq 0$$

$$K \in \mathcal{C}$$

Proof: Let us suppose that the matrix inequality in (7) is satisfied with equality.

$$V = Q + \begin{pmatrix} A \\ -RK \end{pmatrix}^T \begin{pmatrix} (\gamma V)^{-1} - \alpha \Sigma' & -B \\ -B^T & R \end{pmatrix}^{-1} \begin{pmatrix} A \\ -RK \end{pmatrix}$$

Using the block matrix inversion formula, the inverse of the block matrix above

$$\begin{pmatrix} M' & -BM'R^{-1} \\ -R^{-1}B^T M' & R^{-1}B^T M'BR^{-1} + R^{-1} \end{pmatrix}$$

where $\alpha\Sigma = BR^{-1}B^T, M' = \left((\gamma V)^{-1} - \alpha\Sigma \right)^{-1}$. It can now be verified using simple algebra that the matrix equality stated above is equivalent to (6). Further, the objective is a constant factor time $c = v_K(0)$, so that (7) solves the problem $\min_{K \in \mathcal{C}} v_x(K)$ if the LMI is satisfied with equality. However, if the LMI is satisfied with inequality, then one can pick $\tilde{Q} \succeq Q$ so that the LMI is satisfied with equality, and then the problem has solved $\min_{K \in \mathcal{C}} v_x(K)$ with a larger state cost - meaning that c must be larger than necessary (hence contradicting optimality). Thus, at optimality, the LMI is satisfied with equality and (7) minimizes $\min_{K \in \mathcal{C}} v_x(K)$.

We now show why (7) is a convex optimization problem and can be written as (8). We can also write the block inverse in another form, giving

$$\begin{pmatrix} (\gamma V)^{-1} & -B \\ -B^T & R \end{pmatrix}^{-1}$$
$$= \begin{pmatrix} I & -\gamma V B \\ 0 & I \end{pmatrix} \begin{pmatrix} (\gamma V)^{-1} & 0 \\ 0 & R - B^T(\gamma V)^{-1}B \end{pmatrix}^{-1} \begin{pmatrix} I & -\gamma V B \\ 0 & I \end{pmatrix}^T.$$

Writing the whole LMI, we get

$$V \succeq$$
$$Q + \begin{pmatrix} A \\ -\gamma B^T V A - RK \end{pmatrix}^T \begin{pmatrix} \gamma V & 0 \\ 0 & R - B^T(\gamma V)^{-1}B \end{pmatrix}^{-1} \begin{pmatrix} A \\ -\gamma B^T V A - RK \end{pmatrix}.$$

The RHS evaluates to

$$\gamma A^T V A +$$
$$\left(RK - \gamma B^T V A \right)^T \left(R - B^T(\gamma V)^{-1}B \right)^{-1} \left(RK - \gamma B^T V A \right)$$

Using Schur-complements, it can be shown that the RHS is a $\succeq$-convex function of V, K. Thus, the given optimization problem is convex. ∎

Remark 1: We needed to assume that $\Sigma = \frac{BR^{-1}B^T}{\alpha}$ for the above theorem, while in previous theorems we allowed a semidefinite inequality there. We believe the proof can be extended to the inequality case as well, although we do not have a rigorous result for the inequality case at the time of writing this paper.

Remark 2: Note that theorem 3.1 guarantees that the problem can be solved using convex optimization. However, this result strengthens it saying that one need not solve the Bellman equation (which generally calls for an iterative procedure like value iteration). Instead, one can simply pose the control design problem as a determinant maximization problem subject to convex constraints.

Corollary 1: The controller synthesis formulation (8) is infeasible unless the autonomous system $x_{t+1} = \sqrt{\gamma}Ax_t$ is stable.

Proof: If (8) is feasible, then from the LMI constraint, we know that the top left block $V - Q - \gamma A^T V A$ must be positive semidefinite. Thus,

$$V - Q - \gamma A^T V A = \tilde{P}$$

for some $\tilde{P} \succeq 0$. Rearranging and setting $\tilde{Q} = Q + \tilde{P} \succeq 0$, this gives us

$$(\sqrt{\gamma}A)^T V (\sqrt{\gamma}A) - V + \tilde{Q} = 0$$

which means that V is a Lyapunov matrix for the system $x_{t+1} = \sqrt{\gamma}Ax_t$ implying that it is stable. Thus, the feasibility of (8) implies the stability of the system, so if $\sqrt{\gamma}A$ is unstable, (8) is infeasible. ∎

A. Stabilizing Controllers Via Continuation

The corollary indicates that synthesizing stabilizing controllers in this framework is impossible, since one would need to take the limit $\gamma \to 1$ (to allow for the standard infinite horizon LEQG stabilization problem), which makes the synthesis formulation (8) infeasible.

In this section, we propose a heuristic algorithm that uses the discount factor γ as a continuation parameter (algorithm 1). It starts with a small value of γ (say .8) and solves (8). It then increases the value of γ by a factor β (we choose it to be 1.02 in our experiments), folds the resulting controller into the dynamics ($A_i = A + BK_{i-1}$) and solves a new problem with the increased value of γ.

We do not have a formal justification or guarantees that this approach will find a stabilizing controller although empirical results (section V) are encouraging.

Algorithm 1 Stabilizing Controller Synthesis via Continuation

---
Input matrices A, B, K_0, γ_0.
for $i = 1, \dots, T$ **do**
  $\bar{K}_i \leftarrow$ solution of (8) with
    $A_i = A + BK_{i-1}, B_i = B, \gamma = \gamma_i$
  $K_i \leftarrow \bar{K}_i + K_{i-1}, \gamma_{i+1} = \beta\gamma_i.$
end for $K \leftarrow K_T$

V. NUMERICAL RESULTS

In this section, we present preliminary numerical results illustrating the efficacy of our approach on various control synthesis problems. We will work exclusively with sparsity constraints of the form

$$\mathcal{C} = \{K : K_{ij} = 0 \quad \forall (i, j) \notin S\}.$$

We also restrict ourselves to the infinite horizon formulation (8). We first take a simple pendulum, linearize it about its unstable equilibrium and construct an Euler discretization of the dynamics to get an unstable discrete-time dynamical system. In this example, there are no constraints, we only use this to test whether our algorithm is capable of synthesizing stabilizing controllers in the absence of constraints. The maximum eigenvalue of the closed loop dynamics (in terms of absolute value) is a measure of distance to instability - when this value gets smaller than one, we have stable closed loop dynamics. We plot this as a function of γ (see figure 1).
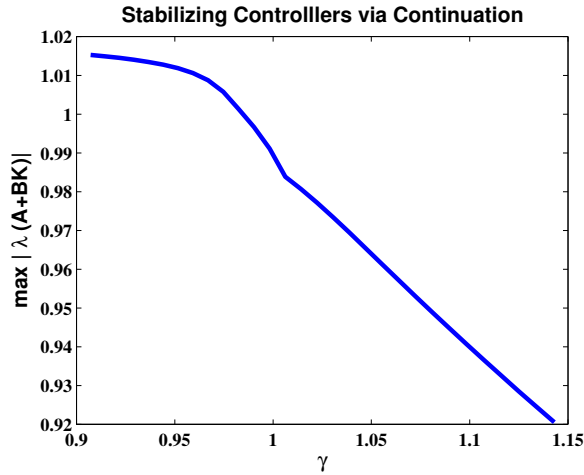
Stabilizing Controlllers via Continuation

Fig. 1: Continuation Controller Synthesis for Pendulum

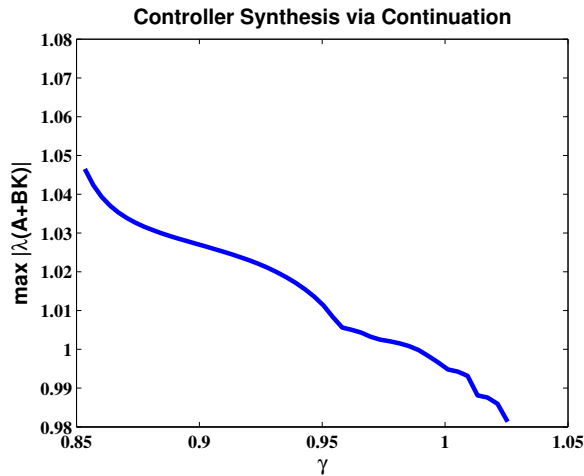**Controller Synthesis via Continuation**

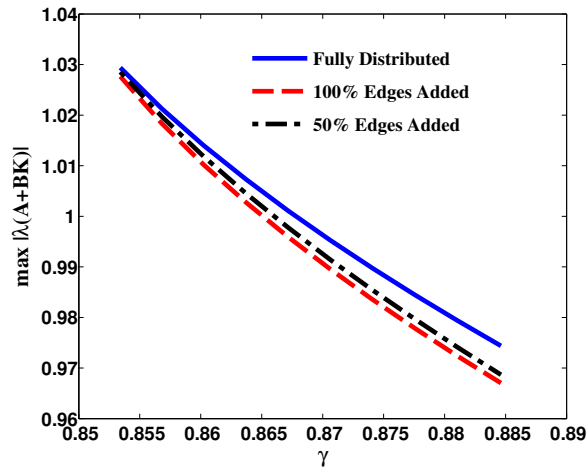Fig. 2: Continuation Controller Synthesis for IEEE 9-bus Power System

Fig. 3: Continuation Controller Synthesis for IEEE 9-bus Power System

We then consider controllers for frequency regulation in power systems. The power system dynamics is described by the swing equations [BH81]. We use the IEEE 9 bus system available from the MATPOWER package [mat] and linearize the system dynamics about an equilibrium point (OPF solution). We scale up the ressistances on the transmission lines to make the system unstable. We do not impose any structural constraints here either. Again, we plot the maximum absolute eigenvalue as a function of γ through iterations of algorithm 1 (see figure 2). In both problems, our approach is able to successfully find stabilizing controllers (reduce the maximum absolute eigenvalue to a number smaller than 1).

Finally, we investigate the effect of structural constraints in the problem. We have the same dynamics as the previous example, but use a larger test system, the IEEE-14 bus network (from [mat]). In the first case, we make the control fully distributed so no two nodes can communicate with each other. We then add 50% of the edges (randomly chosen) in the transmission network (since communication between nodes that are directly connected by transmission lines is likely to be most useful for control). Finally we add all the edges in the network. The effect of these various sparsity patterns is plotted in figure 3. Again, we successfully synthesize stabilizing controllers in all cases, although the decay in the maximum eigenvalue increases as the constraints on the system are reduced.

VI. CONCLUSIONS AND FUTURE WORK

We showed that controller design can be made convex under arbitrary convex structural constraints on the set of feedback matrices, provided a condition relating the control cost, noise covariance and risk factor is satisfied. We have both finite and infinite horizon versions of the results. In infinite horizon, if we do not use a discount factor, our formulation leads to provably infeasible problems. This is not surprising, consider the known hardness of synthesizing stabilizing feedback controllers under arbitrary convex constraints on the feedback gains. However, we proposed a heuristic algorithm that uses the discount factor as a continuation parameter to synthesize stabilizing controllers and showed its applicability on some test systems numerically. In future work, we plan to rigorously analyze the heuristic and give conditions under which it is guaranteed to produce a stabilizing controller. We also plan to explore non-LEQG formulations. Theorem 3.1 applies to arbitrary cost functions (not just quadratic). It is possible that the degenrate behavior we see (corollary IV-A) is an artifact of the unbounded quadratic costs. Indeed, if we use globally bounded costs, degeneracy cannot arise in our problems. These results also suggest a broader question: Is the hardness of structured controller synthesis an artefact of the LQR formulation where systems can grow unbounded along with the corresponding state/control costs? Does this hardness go away if we formulate structured controller synthesis in terms of bounded costs? This is an interesting question that will be a topic of our future investigations.

REFERENCES

[ANR08] Pierre Apkarian, Dominikus Noll, and Aude Rondepierre. Mixed h-2/h-inf control via nonsmooth optimization. *SIAM Journal on Control and Optimization*, 47(3):1516–1546, 2008.

[BH81] AR Bergen and David J Hill. A structure preserving model for power system stability analysis. *IEEE Transactions on Power Apparatus and Systems*, (1):25–35, 1981.

[BHLO06] JV Burke, D Henrion, AS Lewis, and ML Overton. Hifoo-a matlab package for fixed-order controller design and h optimization. In *Fifth IFAC Symposium on Robust Control Design, Toulouse*, 2006.

[BT97] Vincent Blondel and John N Tsitsiklis. Np-hardness of some linear control design problems. *SIAM Journal on Control and Optimization*, 35(6):2118–2127, 1997.

[DTF13] Krishnamurthy Dvijotham, Emanuel Todorov, and Maryam Fazel. Convex structured controller design. *arXiv preprint arXiv:1309.7731*, 2013.

[DTTF13] Krishnamurthy Dvijotham, Evangelos Theodorou, Emanuel Todorov, and Maryam Fazel. Convexity of optimal linear controller design. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 2477–2482, Dec 2013.

[FML] Ghazal Fazelnia, Ramtin Madani, and Javad Lavaei. Convex relaxation for optimal distributed control problem.

[FSJ94] Chih-Hai Fan, Jason L Speyer, and Christian R Jaensch. Centralized and decentralized solutions of the linear-exponential-gaussian problem. *IEEE Transactions on Automatic Control*, 39(10):1986–2003, 1994.

[K+60] Rudolph Emil Kalman et al. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960.

[Lav13] Javad Lavaei. Optimal decentralized control problem as a rank-constrained optimization. In *Proceedings of the Allerton Conference on Control and Computing*, 2013.

[LFJ13] 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*, 58(9):2426–2431, September 2013.

[LL13] Andrew Lamperski and Laurent Lessard. Optimal decentralized state-feedback control with sparsity and delays. *arXiv preprint arXiv:1306.0036*, 2013.

[mat] IEEE 1547 Standard for Interconnecting Distributed Resources with Electric Power Systems.

[Oso12] Takayuki Osogami. Iterated risk measures for risk-sensitive markov decision processes with discounted cost. *arXiv preprint arXiv:1202.3755*, 2012.

[QSVK04] Xin Qi, Murti V Salapaka, Petros G Voulgaris, and Mustafa Khammash. Structured optimal and robust control with multiple criteria: A convex solution. *IEEE Transactions on Automatic Control*, 49(10):1623–1640, 2004.

[RL02] Michael Rotkowitz and Sanjay Lall. Decentralized control information structures preserved under feedback. In *Proceedings of the 41st IEEE Conference on Decision and Control*, volume 1, pages 569–575, 2002.

[RL06] : Rotkowitz, Michael and Sanjay Lall. A characterization of convex problems in decentralized control. *IEEE Transactions on Automatic Control*, 51(2):274–286, 2006.

[SDJ74] J Speyer, John Deyst, and D Jacobson. Optimization of stochastic linear systems with additive measurement and process noise using exponential performance criteria. *Automatic Control, IEEE Transactions on*, 19(4):358–366, 1974.

[Sha13] P. Shah. H2-optimal decentralized control over posets: A state-space solution for state-feedback. *IEEE Transactions on Automatic Control*, PP(99):1–1, 2013.

[SL10] John Swigart and Sanjay 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.

[Wit68] Hans S Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6(1):131–147, 1968.

[WZ13] M. Wytock and J. Zico Kolter. A Fast Algorithm for Sparse Controller Design. *ArXiv e-prints*, December 2013.

APPENDIX

Lemma 1 (Contraction Mapping Property): The operator Γ acting on functions $f : \mathbf{R}^n \to \mathbf{R}$ is defined by

$$\Gamma[f](x) = q(x) + \frac{1}{\alpha} \log \left(\underset{\mathbf{y} \sim P(x)}{\mathrm{E}} \left[\exp \left(\alpha \gamma f(\mathbf{y}) \right) \right] \right)$$

where $q : \mathbf{R}^n \to \mathbf{R}$, $0 < \gamma < 1$, $P(x)$ is a probability measure for every x. This is a contraction mapping with respect to the metric $\|f - g\|_\infty = \sup_x |f(x) - g(x)|$.

Proof:

$$\Gamma[f](x) - \Gamma[g](x) = \frac{1}{\alpha} \log \left(\frac{\mathrm{E} \left[\exp \left(\alpha \gamma f(x + \omega) \right) \right]}{\mathrm{E} \left[\exp \left(\alpha \gamma g(x + \omega) \right) \right]} \right)$$

For any $\mathbf{y}$, we can write $f(\mathbf{y}) \leq g(\mathbf{y}) + \|f - g\|_\infty$. The RHS of the equation above is bounded above by

$$\frac{1}{\alpha} \log \left(\frac{\mathrm{E} \left[\exp \left(\alpha \gamma g(x + \omega) \right) \right]}{\mathrm{E} \left[\exp \left(\alpha \gamma g(x + \omega) \right) \right]} \right) + \gamma \|f - g\|_\infty$$
$$= \gamma \|f - g\|_\infty.$$

Similarly, we can show that $\Gamma[g](x) - \Gamma[f](x) \leq \gamma \|f - g\|_\infty$. Thus,

$$\|\Gamma[f] - \Gamma[g]\|_\infty = \sup_x |\Gamma[f](x) - \Gamma[g](x)| \leq \gamma \|f - g\|$$

Γ is a contraction mapping. Note that we assumed f, g were bounded in order to make this argument. We can do the same argument with $\tilde{f}(x) = \min \left(\max \left(f(x), m \right), M \right)$ (for $-\infty < m < M < \infty$) and take limits under appropriate technical assumptions. ∎

Proof of theorem 4.1

Proof: Follows by applying Gaussian integrals and the matrix inversion lemma. Plugging in $v_\mathrm{K}(x) = \frac{x^T V x}{2} + c$ into (5), the term inside the logarithm in the RHS evaluates to

$$\underset{\omega \sim \mathcal{N}(0, \Sigma)}{\mathrm{E}} \left[\exp \left(\alpha \gamma \frac{\left(\tilde{A}x + \omega \right)^T V \left(\tilde{A}x + \omega \right) + c}{2} \right) \right]$$

Taking logarithms and dividing by α, this evaluates to

$$\frac{\left(\tilde{A}x \right)^T (\gamma V) \tilde{A}x}{2}$$
$$+ \frac{\left(\tilde{A}x \right)^T \left((\gamma V) \Sigma^{1/2} \left(\frac{I}{\alpha} - \Sigma^{1/2} (\gamma V) \Sigma^{1/2} \right)^{-1} (\gamma V) \right) \tilde{A}x}{2}$$
$$+ \gamma c - \frac{\log \left(\det \left(I - \alpha \gamma \Sigma^{1/2} V \Sigma^{1/2} \right) \right)}{2\alpha}$$

Using the matrix inversion lemma, the matrix inside brackets evaluates to

$$\left((\gamma V)^{-1} - \alpha \Sigma^{1/2} \Sigma^{1/2} \right)^{-1} = \left((\gamma V)^{-1} - \alpha \Sigma \right)^{-1}.$$

Matching the terms in (5), we get

$$V = Q + K^T R K + \tilde{A}^T \left((\gamma V)^{-1} - \alpha \Sigma \right)^{-1} \tilde{A}$$

with $c = \frac{1}{\alpha(1 - \gamma)}$. ∎