

Convexity of Optimal Linear Controller Design

Krishnamurthy Dvijotham, Evangelos Theodorou , Emanuel Todorov , Maryam Fazel

Abstract—We develop a general class of stochastic optimal control problems for which the problem of designing optimal linear feedback gains is convex. The class of problems includes arbitrary time varying linear systems and costs that are mixtures of exponentiated quadratics. This allows us to model problems with quadratic state costs and linear constraints on states and state transitions. Further, convexity in the feedback gains lets us impose arbitrary convex constraints or penalties on the feedback matrix: Thus we can model problems like distributed control (by imposing a sparsity structure on the feedback matrix) and variable-stiffness control (by applying time-varying penalties to feedback gain matrices). We show that the convex optimization problem can be solved efficiently by using the structure of the matrices involved. Finally, we present an application of these ideas to a practical problem arising in distributed control of power systems.

I. INTRODUCTION

Linear feedback control synthesis is a classical topic in control theory and has been extensively studied in 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control in a distributed fashion), or constraints on states/state transitions, the dynamic programming approach is no longer applicable. In fact, it has been shown that the optimal feedback may not even be linear [Wit68] and the general problem of designing linear feedback gains subject to constraints [BT97] is NP-hard. In recent years, authors have considered special cases of the optimal feedback synthesis problem that can be solved using convex optimization. In [RL02], the authors introduce the notion of quadratic invariance (QI) which characterizes the set of decentralization constraints under which the feedback synthesis problem can be solved using convex programming techniques. In further work [RL06], the authors prove that QI is a necessary and sufficient condition for convexity of feedback synthesis under certain assumptions, *regardless* of the closed loop system norm (performance metric) minimized. Although interesting, the resulting problems are infinite-dimensional except when the system performance is measured using the \mathcal{H}_2 norm [RL06]. Further, explicit

state-space realizations of the resulting controllers were not available in these works. In [SP10], the authors consider a partial order based structure on the feedback matrices and show that in this case, one can come up with explicit state-space realizations of the resulting controllers. All these approaches are formulated in the frequency domain and solve the infinite horizon problem of constructing a stabilizing feedback controller that minimizes the $\mathcal{H}_2/\mathcal{H}_\infty$ norm of the closed loop system. In very recent work [LL13], the authors show that for decentralization constraints arising from certain nested information structures, the feedback synthesis problem can be solved using dynamic programming techniques directly in state-space form.

Our work here takes a different approach: We ask the question: Is it possible to consider a slightly different formulation of the LQG problem that would allow us to show convexity of the feedback synthesis problem for *arbitrary* convex constraints imposed on the feedback matrix? We develop such a class of problems by taking the standard Linear Exponential Quadratic Gaussian (LEQG) problem [SDJ74] and perturbing it with control-dependent noise. When the perturbation is of a specific form, we show that the problem of synthesizing optimal linear feedback matrices is convex and can be solved to global optimality via convex optimization, with arbitrary convex constraints (or penalties) on the feedback matrices. We consider the finite-horizon formulation, and thus we don't have an explicit notion of stability. Intuition and numerical evidence informs us that our framework tends to produce stabilizing controllers, since the objective function becomes large for unstable systems. However, since our formulation is in finite horizon, we cannot guarantee this. The finite-horizon formulation has the advantage of letting us model time-varying linear systems (which often arise in practice as the linearization of a nonlinear system around a nominal trajectory). Finally, our approach allows us to compose several LEQG-type costs and model arbitrary linear constraints on state space trajectories of the system (via exponential penalties). To the best of our knowledge, these developments are novel and previous approaches do not address the problems we solve here.

The rest of the paper is organized as follows: In section II, we describe the mathematical formulation of this new class of control problems, and how they relate to the classical LEQG formulation. In section III, we present the main technical results proving the convexity of the feedback synthesis problem in this formulation. In section IV, we describe how several problems in control can be modeled within our framework. In section V, we discuss computational issues and show that by exploiting the structure of this problem we scale

K. Dvijotham (dvij@cs.washington.edu) and E. Todorov (todorov@cs.washington.edu) are with the Department of Computer Science and Engineering and Department of Applied Mathematics, University of Washington, Seattle, WA 98195, USA

E. Theodorou (etheodor@usc.edu) is with the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, USA

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

our approach to large systems. We present applications of these ideas to the practical problem of distributed frequency control of a power system in section VI. Finally, in section VII, we discuss directions for future work.

II. PROBLEM FORMULATION

A. Notation

We use $\mathbf{x} \in \mathbf{R}^n$ to denote states, $\mathbf{u} \in \mathbf{R}^m$ for controls and $\omega \in \mathbf{R}^n$ for disturbances. We work with discrete-time systems and denote integer-valued time with t ($1 \leq t \leq T$) with T denoting the horizon of the finite-horizon control problem. The time-indices on quantities are indicated as subscripts (x_t, u_t , etc.). We use boldface to denote a quantity stacked over time: For example, \mathbf{x} to denote trajectories of a fixed horizon T : $\mathbf{x} = [x_1, \dots, x_T]$.

For a matrix M , let $\mathbf{S}(M) = M + M^T$. We denote by $\text{Cov}_P(Y)$ the covariance matrix of the random variable $Y \in \mathbf{R}^l$ with distribution P . We denote by $\mathbb{E}_{x \sim P}[f(x)]$ the expectation of the function f under the distribution P . We denote by \mathcal{S}^k the space of symmetric matrices of size k . We use $A \geq B$ to denote component-wise inequality and $A \succeq B$ to denote inequality with respect to the cone of positive semidefinite matrices (where $A, B \in \mathcal{S}^k$).

B. Problem

We deal with discrete-time linear systems of the form

$$\mathbf{x}_0 \sim \mathcal{N}(0, \Sigma_{-1}) \quad \mathbf{x}_{t+1} = A_t \mathbf{x}_t + B_t \mathbf{u}_t + \omega_t$$

where $\text{noise}_t = \mathcal{N}(0, \Sigma_t)$, $t = 0, \dots, T-1$, $\mathbf{x}_t \in \mathbf{R}^n$, $A_t \in \mathbf{R}^{n \times n}$, $B_t \in \mathbf{R}^{n \times m}$, $\mathbf{u}_t \in \mathbf{R}^m$ and $\mathcal{N}(\mu, \Sigma)$ denotes a Gaussian distribution of mean μ and covariance Σ . We will assume that Σ_t is full rank for all $-1 \leq t \leq T-1$ (-1 refers to the initial state distribution). We denote $S_i = (\Sigma_i)^{-1}$.

We seek to design feedback matrices $\mathbf{u}_t = K_t \mathbf{x}_t$ that drive the system state \mathbf{x}_t towards minimizing an objective function. Let $\mathbf{K} = \{K_t\}$ and denote by $P_{\mathbf{K}}(\mathbf{x})$ the joint Gaussian density over the trajectories $\mathbf{x} = [x_0, \dots, x_T]$ sampled from this dynamical system. We denote by $P_0(\mathbf{x})$ the distribution of trajectories under the uncontrolled system (with $\mathbf{K} = 0$).

The stochastic control problem we consider is defined as follows:

$$\underset{\mathbf{K}}{\text{Minimize}} \quad \log \left(\mathbb{E}_{P_{\mathbf{K}} \otimes R_{\mathbf{K}}} [J(\mathbf{x})] \right) \quad (1)$$

$$\text{where } J(\mathbf{x}) = \exp \left(\sum_{t=0}^{T-1} \begin{bmatrix} x_t \\ x_{t+1} \\ 1 \end{bmatrix}^T Q_t \begin{bmatrix} x_t \\ x_{t+1} \\ 1 \end{bmatrix} \right)$$

$$= \exp \left(\frac{\mathbf{x}^T \mathbf{Q} \mathbf{x}}{2} + \mathbf{q}^T \mathbf{x} \right)$$

$$R_{\mathbf{K}}(\mathbf{x}) = \exp \left(\sum_{t=0}^{T-1} \frac{1}{2} (B_t K_t x_t)^T (\Sigma_t)^{-1} (B_t K_t x_t) \right)$$

$$P_{\mathbf{K}} \otimes R_{\mathbf{K}}(\mathbf{x}) = \frac{P_{\mathbf{K}}(\mathbf{x}) R_{\mathbf{K}}(\mathbf{x})}{\mathbb{E}_{P_{\mathbf{K}}} [R_{\mathbf{K}}(\mathbf{x})]}$$

$Q_t \in \mathcal{S}^{2n+1} \succeq 0$ and $\mathbf{Q} \in \mathcal{S}^{nT}$, $\mathbf{q} \in \mathbf{R}^{nT}$ are matrices that rearrange the Q_t into a block-diagonal and stacked

vector form, that make it easy to represent $J(\mathbf{x})$ in terms of a quadratic form in \mathbf{x} . Note that the matrix \mathbf{Q} has a block tridiagonal form, since the cost is a sum of terms that involve only successive states $\mathbf{x}_t, \mathbf{x}_{t+1}$.

C. Interpretation of the Model

The traditional LEQG problem [SDJ74] can be phrased as

$$\min_{\mathbf{K}} \mathbb{E}_{P_{\mathbf{K}}} \left[\exp \left(\sum_{t=0}^{T-1} \frac{1}{2} x_t^T Q x_t + \frac{1}{2} u_t^T R u_t + x_T^T Q_f x_T \right) \right].$$

In our formulation, we make three important changes to this:

- We do not have explicit control costs, although they can be modeled through costs on state transitions. The constant term 1 in $\begin{bmatrix} x_t \\ x_{t+1} \\ 1 \end{bmatrix}$ means that J can also include linear terms in x_t .
- We evaluate the expectation with respect to the perturbed distribution $P_{\mathbf{K}} \otimes R_{\mathbf{K}}$ rather than $P_{\mathbf{K}}$ directly.
- We are required to have an initial state distribution that is a Gaussian with mean 0.

We take the expectation of the cost J under a perturbed distribution $P_{\mathbf{K}} \otimes R_{\mathbf{K}}$. Since $R_{\mathbf{K}}$ grows quadratically with the gains \mathbf{K} , this framework resembles control multiplicative noise [Tod05]. In [Tod05], the authors propose a variant of standard LQG models where the covariance of the process noise ω_t is a quadratic function of the controls \mathbf{u} . In the model we use here, we instead directly perturb the joint inverse covariance of the trajectory \mathbf{x} : Specifically, we subtract a positive definite block diagonal matrix (with diagonal blocks given by $(B_t K_t)^T \Sigma_t^{-1} B_t K_t$), thereby reducing the inverse covariance and increasing the covariance. Thus, the perturbation we consider here has a similar effect as control multiplicative noise, although they are not mathematically equivalent.

The form of $R_{\mathbf{K}}$, which can be expressed as $(B_t K_t x_t)^T (\Sigma_t)^{-1} (B_t K_t x_t) = \mathbf{u}_t^T B_t^T (\Sigma_t)^{-1} B_t \mathbf{u}_t$, is reminiscent of that used in path integral control [MN03] [Kap05] [TBS10], with the inverse of the noise showing up in the control cost. In path integral control, this inverse relationship is exploited to transform the Hamilton-Jacobi-Bellman equation from a nonlinear to a linear PDE, allowing efficient solution through Monte-Carlo (sampling based) techniques. Here, we will use it to derive convexity of the objective in the feedback matrix \mathbf{K} .

III. MAIN TECHNICAL RESULTS

Theorem 1: Define

$$c(\mathbf{K}) = \begin{cases} \mathbb{E}_{P_{\mathbf{K}} \otimes R_{\mathbf{K}}} [J(\mathbf{x})] & \text{if } \mathbb{E}_{P_{\mathbf{K}} \otimes R_{\mathbf{K}}} [J(\mathbf{x})] \text{ is finite} \\ \infty & \text{otherwise} \end{cases},$$

i.e., $c(\mathbf{K})$ is an extended-real valued function [BV04]. Then $c(\mathbf{K})$ is convex in \mathbf{K} .

Proof: We need to show two things: the set of \mathbf{K} for which $c(\mathbf{K})$ is finite is convex, and restricted to this domain, $c(\mathbf{K})$ is a convex function.

Let us first assume that $c(\mathbf{K})$ is finite. Now $P_{\mathbf{K}}$ has a Gaussian density over the space of trajectories \mathbf{x} . Denoting $\tilde{A}_t = A_t + B_t K_t$ and $S_t = \Sigma_t^{-1}$, the joint inverse covariance is given by

$$\begin{bmatrix} S_{-1} + \tilde{A}_0^T S_0 \tilde{A}_0 & -\tilde{A}_0^T S_0 & 0 \dots \\ -S_0 \tilde{A}_0 & S_0 + \tilde{A}_1^T S_1 \tilde{A}_1 & -\tilde{A}_1^T S_1 \dots \\ 0 & -S_1 \tilde{A}_1 & \dots \\ \vdots & \vdots & \vdots \end{bmatrix}.$$

When this is multiplied with $R_{\mathbf{K}}$ and renormalized, the resulting Gaussian distribution has $(B_t K_t)^T (\Sigma_t)^{-1} (B_t K_t)$ subtracted from the t -th diagonal block of the inverse covariance:

$$\left(\begin{matrix} \text{Cov}(\mathbf{x}) \\ P_{\mathbf{K}} \otimes R_{\mathbf{K}} \end{matrix} \right)^{-1} = \mathcal{L}(\mathbf{K}) = \left(\begin{matrix} \text{Cov}(\mathbf{x}) \\ P_0 \end{matrix} \right)^{-1} + \begin{bmatrix} \mathbf{S} (A_0^T S_0 B_0 K_0) & -K_0^T B_0^T S_0 & 0 \dots \\ -S_0 B_0 K_0 & \mathbf{S} (A_1^T S_1 B_1 K_1) & -K_1^T B_1^T S_1 \dots \\ 0 & -S_1 B_1 K_1 & \dots \\ \vdots & \vdots & \vdots \end{bmatrix}.$$

Thus, the inverse covariance of $P_{\mathbf{K}} \otimes R_{\mathbf{K}}$ is a linear function of \mathbf{K} . We denote this linear map as $\mathcal{L}(\mathbf{K})$. The objective can then be rewritten as

$$\log \left(\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0, \mathcal{L}(\mathbf{K})^{-1})} \left[\exp \left(\frac{1}{2} \mathbf{x}^T \mathbf{Q} \mathbf{x} + \mathbf{q}^T \mathbf{x} \right) \right] \right)$$

where $\mathbf{Q} \in \mathcal{S}^{nT}$, $\mathbf{q} = \mathbf{R}^{nT}$ are obtained by assembling the stage-wise costs into a big block tri-diagonal matrix (or a long vector). It is trivial to see that $\mathbf{Q} \succeq 0$. By theorem 2, the above function is convex in $\mathcal{L}(\mathbf{K})$ and thereby in \mathbf{K} .

From theorem 1, it is easy to see that the objective is finite if and only if $\mathcal{L}(\mathbf{K}) - \mathbf{Q} \succ 0$, which is a convex domain. Hence, $c(\mathbf{K})$ is convex. ■

A. Discussion

The essential ingredient in our convexity proof is the fact that the joint inverse covariance is a linear function of the feedback matrices \mathbf{K} . This allows us to directly prove convexity in \mathbf{K} , as opposed to alternative approaches that perform a nonlinear transformation on \mathbf{K} , which generally precludes one from enforcing constraints on \mathbf{K} , except ones with special structure like Quadratic Invariance [RL06]. Our approach allows us to impose *arbitrary* convex constraints and penalties on \mathbf{K} , something that was not possible under previous approaches. The caveat is that we need to solve a perturbed version of the standard LEQG problem. However, as we have argued, this perturbation changes the problem in a meaningful way (through control-multiplicative noise). Further, in section VI, we will present numerical examples showing that even though we're solving a perturbed problem, the feedback matrices computed by our approach work well

even on the original unperturbed linear system. We note one restriction here: The formulation studied in the papers [RL06][SP10] considers the general case of dynamic output feedback. In this paper, we restrict ourselves to static state feedback $u_t = \mathbf{K}_t x_t$.

IV. APPLICATIONS

The fact that the control objective is convex in \mathbf{K} is very powerful, since this allows us to leverage the full power of convex optimization based modeling. We can impose *arbitrary* convex costs and constraints on the feedback matrix \mathbf{K} . This has several applications in control, which we describe in this section.

A. Distributed Control

We can model distributed control by imposing specific sparsity patterns on the feedback matrix \mathbf{K}_t . For example, in the simple example of fully distributed control $u_{t,i} = k_{t,i} x_{t,i}$, we simply require \mathbf{K}_t to be a diagonal matrix $\text{diag}(\{k_{i,t}\})$. We can also model delays in this framework: If we augment the state to include not just the current state but also the state at the last k time steps: $\tilde{\mathbf{x}} = [x_t; x_{t-1}; \dots; x_{t-k+1}]$, we constrain K_t to have zeros in all state dimensions except those corresponding to the x_{t-k+1} , so that the control is a function of x_{t-k+1} rather than x_t .

B. Variable Impedance Control

The framework also allows us to impose convex costs on the gains K_t (for example the Frobenius norm $\|K_t\|_F$). An example of where this is valuable is variable impedance control, which requires penalizing control gains in a time-varying manner. Concrete applications can be found in robotics [BTSS10]. In particular, it is important that robots perform tasks safely. High gain control creates instabilities for systems with many interacting bodies while it also amplifies sensor noise. On the other hand, low gain control sacrifices tasks performance. Our method could be applied to control scenarios for intelligent gain scheduling that meets the performance and safety requirements.

C. Model Errors

If the A_t, B_t, Σ_t matrices are not known with certainty, we can define a set of allowable models $\{A_t, B_t, \Sigma_t\} \in \mathcal{M}$.

We can then define an objective of the kind

$$\max_{\{A_t, B_t, \Sigma_t\} \in \mathcal{M}} \log \left(\mathbb{E}_{\mathbf{x} \sim P_{\mathbf{K}} \otimes R_{\mathbf{K}}(\{A_t, B_t, \Sigma_t\})} [J(\mathbf{x})] \right)$$

which remains convex in \mathbf{K} , since it is the supremum of a set of convex functions. If \mathcal{M} is a finite set, the maximum can be computed simply by enumeration. Otherwise, a finite approximation of \mathcal{M} obtained by sampling a set of models can already provide a sufficient degree of robustness.

D. Linear constraints on the trajectory

Our framework also allows us to model constraints on the trajectory: $\{\mathbf{q}_i^T \mathbf{x} \leq c_i : i = 1, \dots, n_c\}$ by imposing them as exponential penalties. Assuming we have a common quadratic cost \mathbf{Q} , we can define an objective

$$\max_{1 \leq i \leq n_c} \log \left(\mathbb{E}_{\mathbf{x} \sim P_{\mathbf{K}} \otimes R_{\mathbf{K}}} [\exp(\mathbf{x}^T \mathbf{Q} \mathbf{x} + \mathbf{q}_i^T \mathbf{x} - c_i)] \right)$$

which is again convex, being the supremum of a set of convex functions.

Imposing these constraints directly on trajectories rather than individual states allows us to model constraints coupling states: For example, we can have a ramping constraint like $x_{t+1} \leq x_t + \delta$.

V. ALGORITHMS AND COMPUTATION

In this section, we discuss algorithms to compute the objective (1) and its gradient efficiently. We first rewrite the problem (1) in a more explicit form making the objective a clear function of \mathbf{K} :

$$\text{Minimize}_{\mathbf{K}} - \frac{\log \left(\frac{\det(\mathcal{L}(\mathbf{K}) - \mathbf{Q})}{\det(\mathcal{L}(\mathbf{K}))} \right)}{2} + \mathbf{q}^T (\mathcal{L}(\mathbf{K}) - \mathbf{Q})^{-1} \mathbf{q} \quad (2)$$

$$\mathcal{L}(\mathbf{K}) = \left(\text{Cov}_{P_0}(\mathbf{x}) \right)^{-1} + \begin{bmatrix} \mathbf{S} (A_0^T S_0 B_0 K_0) & -K_0^T B_0^T S_0 & 0 \dots \\ -S_0 B_0 K_0 & \mathbf{S} (A_1^T S_1 B_1 K_1) & -K_1^T B_1^T S_1 \dots \\ 0 & -S_1 B_1 K_1 & \dots \\ \vdots & \vdots & \vdots \end{bmatrix}$$

In order to ensure that the objective is well defined, we need to impose the constraint $\mathcal{L}(\mathbf{K}) - \mathbf{Q} \succ 0$. However, the objective function already includes the term $-\frac{\log(\det(\mathcal{L}(\mathbf{K}) - \mathbf{Q}))}{2}$, which plays the role of a log-barrier that prevents the solution from violating the constraint. If $(\text{Cov}_{P_0}(\mathbf{x}))^{-1} \succ \mathbf{Q}$, $\mathbf{K} = 0$ is a guaranteed feasible starting point. If not, we can initialize using an infeasible-start interior point method [BV04].

Further, even though the size of $\mathcal{L}(\mathbf{K})$ is $nT \times nT$, the number of nonzero entries only grows linearly with T . Computing the gradient requires inversion of $\mathcal{L}(\mathbf{K})$ and $\mathcal{L}(\mathbf{K}) - \mathbf{Q}$. By exploiting the block-tridiagonal structure of these matrices, it is possible to compute the block tridiagonal parts of the inverse without computing the rest. The algorithm first performs a Cholesky factorization (which itself is guaranteed to retain the block-tridiagonal structure), and uses the factors to compute the relevant blocks of the inverse. A similar algorithm can also be used to compute the Newton step to a KKT system in an interior point method. This block structure we see here is a special case of a *chordal sparsity pattern* and efficient algorithms have been proposed for these [ADV10]. We can leverage this work to develop

an interior point method for this problem with each iteration of the interior point (solution of the KKT system) can be performed in time and memory that grows linearly with T .

In the numerical examples we present here, we used an off-the-shelf LBFGS implementation [Sch05], with the gradient computed efficiently as outlined above. We plan to experiment with a full interior point method in future work.

VI. NUMERICAL EXAMPLES

We present an application of this framework to designing feedback controllers for frequency control in power grids. The grid can be viewed as a collection of oscillators (rotating generators) coupled electromechanically through the network. The frequency of oscillation at each node in the network needs to be close to the system frequency (50/60 Hz) for the system to be stable. Deviations from the system frequency are related to imbalances in generation and demand. Sudden changes in generation (due to a fault or outage) need to be compensated for rapidly so as to prevent excursions of frequency from nominal frequency. At these time scales (a few milliseconds to seconds), it is not possible to do a centralized redispatch of generation. We thus need distributed control for this problem.

A. Distributed Control of Power Systems

In this section, we consider the problem of frequency stabilization in a power system. We use the IEEE 14 bus benchmark [Pow] as a test grid. The states consist of the rotor angular positions and frequency deviations (from the system frequency of 50/60 Hz) at each node in the network: $\mathbf{x} = \begin{pmatrix} \theta \\ \dot{\theta} \end{pmatrix}$. For small time-intervals following a fault, the dynamics of the rotating generators can be described using the swing equation [BH81]. The linearized system dynamics is given by

$$\dot{\mathbf{x}} = \begin{pmatrix} 0 & I \\ L & 0 \end{pmatrix} \mathbf{x} + \begin{pmatrix} 0 \\ M \end{pmatrix} \mathbf{u} + \text{noise}$$

where L is a weighted Laplacian system bus matrix, and M is a matrix with ones on the diagonal entries corresponding to controlled generators. We discretize this system with a time-step $dt = 10ms$. We have a quadratic objective penalizing controls and one penalizing frequency deviations $\dot{\theta}$. The controls consist of regulating the output of generators so as to stabilize system frequency in the presence of fluctuations (due to fluctuations in power generation and demand, faults etc.). We formulate this as a distributed control problem by imposing sparsity structure on the feedback matrix. We consider 4 different kinds of sparsity structures:

- Fully centralized control (LEQG): The feedback gains here are unconstrained, and hence can be computed using the standard LEQG formulation.
- Neighborhood feedback control (Nb): The feedback matrix here consists of two blocks: The feedback on the rotor angles θ has the same sparsity structure as the network graph on K , so that the controls at a given

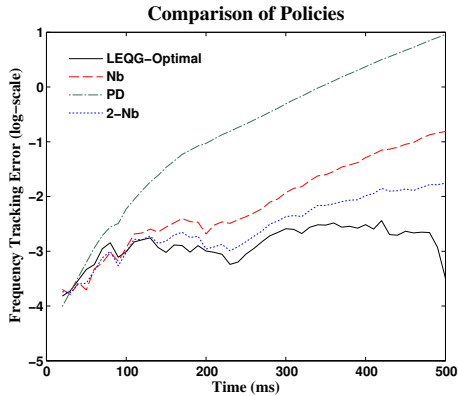


Fig. 1. Performance of various Feedback Policies: The y-axis represents the frequency tracking error on a log-scale. The x-axis is time (in milliseconds). The Nb, 2-Nb policies achieve performance that is very close to the LEQG solution.

generator node only rely on the state dimensions corresponding to the node and its neighbors. Additionally, we allow proportional feedback on the local frequency deviations $\hat{\theta}$.

- 2-Neighborhood feedback control (2-Nb): The same as above, except that we now allow feedback to depend not just on immediate neighbors, but on neighbors two nodes away.
- Fully distributed (PD): Here, the feedback matrix consists of 2 diagonal blocks, so that the control at each generator only depends on the local rotor angle θ_i and frequency deviations $\hat{\theta}_i$.
- No Control: We compare to the autonomous system $K = 0$ as a baseline.

We compare the performance of each policy in the following manner: We treat the LEQG solution as the optimal and the autonomous system as a baseline. For each policy, we then compute the following performance metric:

$$\text{SubOpt}(\text{Policy}) = \frac{\text{Cost}(\text{Policy}) - \text{Cost}(\text{LEQG})}{\text{Cost}(\text{Autonomous}) - \text{Cost}(\text{LEQG})}.$$

This is the percentage loss in performance (degree of suboptimality) relative to the baseline autonomous (uncontrolled) system. The performance is computed by averaging costs over 1000 time domain simulations of the system with each policy plugged in. Note that the simulations are carried out using the original system dynamics (not the perturbed version of the problem we solve here). The table contains the values for the 3 distributed policies (with the gains obtained by our convex programming algorithm) and compares it to the naive alternative of taking the LEQG feedback gains and zeroing-out the entries that don't conform to the policy sparsity pattern (LEQG truncated).

Table I shows the performance of various policies for this problem (in terms of the SubOpt metric). We also plot the frequency tracking error for the various policies on a log-scale in figure 1. The results show that our convex programming solution does significantly better than simply

TABLE I
PERFORMANCE OF FEEDBACK POLICIES

Algorithm	SubOpt	Truncated SubOpt	% nonzeros(K)
PD	1.15 %	100 %	7 %
Nb	.08 %	52 %	16.43 %
2-Nb	.08%	46 %	28.57 %

truncating the LEQG feedback matrices. The gap between the truncated solution and the optimum seems to decrease as the constraints on the feedback matrix become less strict (fewer non-zeros), as expected. Further, for this frequency control problem, the Neighborhood control scheme that uses information from neighbors seems to achieve performance very close to the optimal LEQG performance.

VII. CONCLUSIONS AND FUTURE WORK

A. Conclusions

We have derived a novel framework under which the problem of designing linear feedback gains is a convex optimization problem. We have shown that this framework can capture many control problems of practical interest and go beyond the traditional LQG/LEQG framework: It can deal with constraints on feedback matrices, model errors and linear constraints on the state. We have shown that the resulting convex optimization problem has structure that can be taken advantage of in order to yield efficient convex programming algorithms. Using the practical application of distributed frequency control of power systems, we showed that this framework can successfully solve distributed control problems efficiently, and the solution achieves performance close to unconstrained LEQG, with a relatively small number of nonzeros.

B. Future Work

The new formulation we use requires a specific perturbation of the classical LEQG problem. From our experience so far, it seems that the perturbation encourages low gain solutions. However, more theoretical and empirical work is required to understand the effects of the perturbation.

Besides this understanding, we plan to extend this work in several specific directions:

- *Technical extensions:*
 - Drop the assumption of invertability of the covariance matrix and extend the proof to the general case of degenerate Gaussians.
 - Understand the constraint imposed by $\mathcal{L}(\mathbf{K}) \succ \mathbf{Q}$ and interpret constraints on \mathbf{K} in terms of $\mathbf{Q}, \mathbf{A}, \mathbf{B}$.
 - Extend the results to continuous-time.
- *Modeling extensions:*
 - Our framework lets us deal with exponential cost criteria. One can use this cost in conjunction with Chernoff-bounds to get bounds on the probability of deviation from a ellipsoidal/polytopic constraint sets [ST12]. Thus, chance-constrained control can also be formulated in this framework.

- We described a naive approach of dealing with model errors by looking at the worst case over a finite set of models. It would be useful to extend this to general parameterized uncertainty models.
- Sparsity promoting design: We could also combine this convex formulation with recent advances in sparsity-promoting feedback optimization [FLJ11]. The convex formulation can enable us to develop more efficient algorithms for these problems, as well as perform theoretical analyses of when convex programming can guarantee obtaining the sparsest solutions to these problems [CRPW10].

VIII. ACKNOWLEDGMENTS

This work was supported by the NSF.

REFERENCES

- [ADV10] Martin S Andersen, Joachim Dahl, and Lieven Vandenbergh. Implementation of nonsymmetric interior-point methods for linear optimization over sparse matrix cones. *Mathematical Programming Computation*, 2(3):167–201, 2010.
- [BH81] AR Bergen and David J Hill. A structure preserving model for power system stability analysis. *Power Apparatus and Systems, IEEE Transactions on*, (1):25–35, 1981.
- [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.
- [BTSS10] Jonas Buchli, Evangelos Theodorou, Freek Stulp, and Stefan Schaal. Variable impedance control—a reinforcement learning approach. In *Robotics: Science and Systems Conference (RSS)*, 2010.
- [BV04] Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- [CRPW10] Venkat Chandrasekaran, Benjamin Recht, Pablo A Parrilo, and Alan S Willsky. The convex algebraic geometry of linear inverse problems. In *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pages 699–703. IEEE, 2010.
- [FLJ11] Makan Fardad, Fu Lin, and Mihailo R Jovanovic. Sparsity-promoting optimal control for a class of distributed systems. In *American Control Conference (ACC), 2011*, pages 2050–2055. IEEE, 2011.
- [HJ90] R.A. Horn and C.R. Johnson. *Matrix analysis*. Cambridge university press, 1990.
- [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.
- [Kap05] HJ Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of statistical mechanics: theory and experiment*, 2005(11):P11011, 2005.
- [LL13] Andrew Lamperski and Laurent Lessard. Optimal decentralized state-feedback control with sparsity and delays. *arXiv preprint arXiv:1306.0036*, 2013.
- [MN03] Sanjoy K Mitter and Nigel J Newton. A variational approach to nonlinear estimation. *SIAM journal on control and optimization*, 42(5):1813–1833, 2003.
- [Pow] Ieee 14 bus test case. http://www.ee.washington.edu/research/pstca/pf14/pg_tcal4bus.htm.
- [RL02] Michael Rotkowitz and Sanjay Lall. Decentralized control information structures preserved under feedback. In *Decision and Control, 2002. Proceedings of the 41st IEEE Conference on*, volume 1, pages 569–575. IEEE, 2002.
- [RL06] Michael Rotkowitz and Sanjay Lall. A characterization of convex problems in decentralized control; img src=. *Automatic Control, IEEE Transactions on*, 51(2):274–286, 2006.
- [Sch05] Mark Schmidt. Minfunc, 2005.
- [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.

- [SP10] Parikshit Shah and Pablo A Parrilo. H-infinity-optimal decentralized control over posets: A state space solution for state-feedback. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 6722–6727. IEEE, 2010.
- [ST12] Jacob Steinhardt and Russ Tedrake. Finite-time regional verification of stochastic non-linear systems. *The International Journal of Robotics Research*, 31(7):901–923, 2012.
- [TBS10] Evangelos Theodorou, Jonas Buchli, and Stefan Schaal. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 9999:3137–3181, 2010.
- [Tod05] Emanuel Todorov. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural computation*, 17(5):1084–1108, 2005.
- [Wit68] Hans S Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6(1):131–147, 1968.

IX. APPENDIX

Theorem 2: Let $M, N \in \mathcal{S}^p$, $N \succeq 0$, $M \succ 0$. Let $n \in \mathbf{R}^p$. Then

$$\log \left(\mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, M^{-1})} [\exp(\mathbf{y}^T N \mathbf{y} + n^T \mathbf{y})] \right)$$

is convex in M .

Proof: If $M - N \not\succeq 0$, the expectation is infinite, and we simply define the objective to be ∞ . Restricting ourselves to the domain $M \succ N$, we prove that the function is convex. Then the overall function is an extended convex function [BV04] on the space of all symmetric matrices M .

The expression can be evaluated as $-\log \left(\sqrt{\frac{\det(M-N)}{\det M}} \right) + \frac{n^T(M-N)^{-1}n}{2}$. The second term is convex, since its epigraph is convex, as shown by the following semidefinite representation of the epigraph:

$$\{(M, t) : n^T(M-N)^{-1}n \leq t\} = \left\{ \begin{bmatrix} t & n^T \\ n & M-N \end{bmatrix} \succeq 0 \right\}.$$

Dropping the scaling factor of 1/2, the first term can be rewritten as

$$f(M) = -\log(\det(M-N)) - (-\log(\det M)).$$

From [ADV10], we know that the Hessian-vector product for this function can be written as

$$\nabla^2 f(M)[V] = (M-N)^{-1}V(M-N)^{-1} - M^{-1}VM^{-1}.$$

To prove that the Hessian is positive semidefinite, it suffices to show that $\langle \nabla^2 f(M)[V], V \rangle$ is non-negative. Since $N \succeq 0$, $M \succeq (M-N) \implies (M-N)^{-1} \succeq M^{-1}$ [HJ90]. Letting $X = (M-N)^{-1}$, $Y = M^{-1}$, we have

$$\langle \nabla^2 f(M)[V], V \rangle = \text{tr}(XVXV) - \text{tr}(YVYV).$$

$$VXV \succeq 0, X - Y \succeq 0 \implies \text{tr}((X-Y)VXV) \geq 0$$

$$\implies \text{tr}(XVXV) \geq \text{tr}(YVXV) = \text{tr}(XVYV).$$

$$YVY \succeq 0, X - Y \succeq 0 \implies \text{tr}((X-Y)VYV) \geq 0$$

$$\implies \text{tr}(XVYV) \geq \text{tr}(YVYV).$$

Thus, we’ve proved that $\text{tr}(XVXV) \geq \text{tr}(YVYV)$, so $\langle \nabla^2 f(M)[V], V \rangle \geq 0$ for all $V \in \mathcal{S}^{p \times p}$. Hence, the first term of the objective is convex as well. ■