An Efficient Method to Estimate the Suboptimality of Affine Controllers

Michael J. Hadjiyiannis* Paul J. Goulart Daniel Kuhn***

**Department of Computing*
***Department of Aeronautics*
Imperial College, London SW7 2AZ, UK
{*michael.hadjiyiannis,p.goulart,dkuhn*}*@imperial.ac.uk*

November 1, 2010

Abstract. We consider robust feedback control of time-varying, linear discrete-time systems operating over a finite horizon. For such systems, we consider the problem of designing robust causal controllers that minimize the expected value of a convex quadratic cost function, subject to mixed linear state and input constraints. Determination of an optimal control policy for such problems is generally computationally intractable, but suboptimal policies can be computed by restricting the class of admissible policies to be affine on the observation. By using a suitable re-parameterization and robust optimization techniques, these approximations can be solved efficiently as convex optimization problems. We investigate the loss of optimality due to the use of such affine policies. Using duality arguments and by imposing an affine structure on the dual variables, we provide an efficient method to estimate a lower bound on the value of the optimal cost function for any causal policy, by solving a cone program whose size is a polynomial function of the problem data. This lower bound can then be used to quantify the loss of optimality incurred by the affine policy.

Key words. Stochastic Optimal Control, Affine Control Policies

1 Introduction

We are interested in characterizing the degree of suboptimality of affine feedback policies for linear discrete-time systems with mixed state and input constraints, bounded disturbances and an expected value cost.

The problem of computing an optimal control policy for such systems, either in a minimax or expected value sense, is computationally intractable in general. Methods for solving such problems rely typically on some variation of robust dynamic programming [9] or vertex enumeration methods [23], and typically require the solution of an optimization problem that grows exponentially with the size of the problem data. As a result, significant research effort has focused on methods for finding *suboptimal* control policies that can be computed via solution of a tractable optimization problem.

A common approach is to restrict the class of control policies considered to those based on perturbations to some fixed stabilizing linear controller [19, 21]. More generally, one can compute a controller based on affine disturbance or measurement-error feedback, a technique suggested by a number of authors [20, 25, 6]. In the affine feedback case, characterization of the set of constraint admissible policies is achieved by following the general approach proposed in [6] for robust optimization problems with linear decision rules, leading to a computationally tractable optimization problem.

An attractive feature of such affine parameterizations is that they can be shown to be equivalent (in the state feedback case) to parameterizations of control policies as affine functions of prior states [16, 24], or (in the output feedback case) as affine functions of prior measurements [15, 4]. The idea underpinning these equivalence results is akin to that of the well-known Youla parameterization (or Q-parameterization) in linear systems [28], and relies on a similar nonlinear transformation to produce a convex set of constraint admissible policies over which one can optimize.

Further refinements to the basic idea of affine uncertainty feedback policies have also been proposed, e.g. policies employing "segregated" disturbance feedback based on some partitioning of the uncertainty set [26, 13]. However, a fundamental difficulty with all of the present proposals is that they generally do not provide any estimate of the degree of suboptimality introduced by restricting the optimal control problem to the particular class of control policies proposed. One notable exception to this is in the case of affine disturbance feedback policies for SISO systems, for which such policies can be shown to be optimal for minimax problems in very limited circumstances [10].

In the present paper we provide a general method for estimating the degree of suboptimality of affine feedback controllers when minimizing the expected value of a quadratic cost, for systems with polyhedral state and input constraints and uncertainties whose support is characterized by conic constraints. Our approach follows the general method of [18], and is predicated on a dualization of the original optimal control problem followed by a restriction of the dual variables to those parameterized by a linear decision rule. This results in a tractable optimization problem that provides a *lower* bound on the finite horizon cost achievable, and is a natural counterpart to the *upper* bound on the achievable cost that is found when restricting the class of control policies to those in affine feedback form. The gap between these bounds then serves as an estimate of the worst-case suboptimality of controllers based on affine feedback policies.

The ability to compute the suboptimality gap in an efficient manner provides a valuable insight to anyone wishing to design a control system. If the gap is small, then affine policies are near-optimal and there is little room for improvement. On the contrary, if the gap is large one may consider investing more time and effort to improve on the affine policies, e.g. by using deflected or segregated policies as in [26, 14, 13]. One attractive property of the proposed approach is that due to the symmetry between the upper and lower bound, one may improve both bounds and thus decrease the optimality gap by employing similar techniques on both the primal as well as the dual controllers (e.g. segregation of the uncertainty space [13]).

The paper is organized as follows: Section 2 describes the problem of interest and details a number of standing assumptions. Section 3 outlines the restriction of control policies to those in affine form, and shows how such policies can be calculated via a tractable conic optimization problem. Section 4 describes a method for estimating the degree of suboptimality of affine feedback policies, based on a novel approximation to a dualization of the optimal control problem of interest. We show how this method allows one to compute lower bounds on the optimal cost via solution of a tractable conic optimization problem. Section 5 summarizes a number of general observations about the relationship between the lower- and upper-bounding problems discussed in the previous sections. Section 6 presents numerical examples illustrating the efficacy of the proposed method, with some conclusions drawn in Section 7. An appendix contains the proofs of several technical results presented in the paper.

Notation All random vectors appearing in this paper are defined on an abstract probability space $(\Omega, \mathcal{F}, \mathbb{P})$. $\mathbb{E}(\cdot)$ denotes the expectation operator with respect to $\mathbb{P}$. Random vectors will be represented in boldface, while their realizations will be denoted by the same symbols in normal face. For notational convenience, we denote by $\mathcal{L}_n^2 := \mathcal{L}^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n)$ the space of all $\mathcal{F}$-measurable, square-integrable random vectors valued in $\mathbb{R}^n$. Given a stochastic process $\boldsymbol{z} := (\boldsymbol{z}_0, \ldots, \boldsymbol{z}_T)$, we will denote as $\boldsymbol{z}^t := (\boldsymbol{z}_0, \ldots, \boldsymbol{z}_t)$ the history of the $\boldsymbol{z}$ process up to time t.

The identity matrix in $\mathbb{R}^n$ is denoted $I^{(n)}$ — the superscript will be omitted when the dimension is clear from the context. For matrices $(A, B) \in \mathbb{R}^{m \times n}$, $\ker(A)$ denotes the kernel or null space of A, $A^\dagger$ its pseudo-inverse, and $A \geq B$ denotes component-wise inequality. For every $C \in \mathbb{R}^{n \times n}$ $\mathrm{tr}(C)$ denotes the trace of C. We denote the p-norm in $\mathbb{R}^n$ as $\|\cdot\|_p$ for any $p \in [1, \infty]$. The p-order cone in $\mathbb{R}^{n+1}$ is denoted $\mathcal{K}_p := \{(t, x) \in \mathbb{R}^{n+1} : \|x\|_p \leq t\}$, and its dual cone as $\mathcal{K}_p^* = \mathcal{K}_q$, where $1/p + 1/q = 1$. For any $(y, z) \in \mathbb{R}^{n+1}$ the relation $y \succeq_{\mathcal{K}_p} z$ implies that $(y - z) \in \mathcal{K}_p$. For a matrix $D \in \mathbb{R}^{(n+1)\times(n+1)}$, $D \succeq_{\mathcal{K}_p} 0$ denotes column-wise inclusion in $\mathcal{K}_p$. The vector e_0 denotes the unit vector $e_0 = (1, 0, \ldots, 0)$.

2 Problem Statement and Assumptions

In this paper we are concerned with linear time-varying systems with known state dynamics

$$\boldsymbol{x}_{t+1} = A_t \boldsymbol{x}_t + B_t \boldsymbol{u}_t + C_t \boldsymbol{\xi}_t \quad \text{for} \quad t = 1, \ldots, T-1, \tag{2.1a}$$

where $\boldsymbol{x}_t \in \mathcal{L}_{n_x}^2$ denotes the state of the system with $\boldsymbol{x}_1 = x_1$ known, $\boldsymbol{u}_t \in \mathcal{L}_{n_u}^2$ denotes the control input, and $\boldsymbol{\xi}_t \in \mathcal{L}_{n_\xi}^2$ the process noise. The state of the system is partially observable in the sense that the measured output $\boldsymbol{y}_t \in \mathcal{L}_{n_y}^2$ at time t depends linearly on the state and the noise, i.e.

$$\boldsymbol{y}_t = D_t \boldsymbol{x}_t + E_t \boldsymbol{\xi}_t \quad \text{for} \quad t = 1, \ldots, T-1. \tag{2.1b}$$

The system equations (2.1) can be rewritten compactly in matrix notation as

$$\boldsymbol{x} = B\boldsymbol{u} + C\boldsymbol{\xi} \quad \text{and} \quad \boldsymbol{y} = D\boldsymbol{x} + E\boldsymbol{\xi}\,, \tag{2.2}$$

with

$$\begin{array}{ll} \boldsymbol{x} := (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T) \in \mathcal{L}_{N_x}^2, & N_x := n_x T \\ \boldsymbol{u} := (\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{T-1}) \in \mathcal{L}_{N_u}^2, & N_u := n_u(T-1) \\ \boldsymbol{\xi} := (\boldsymbol{\xi}_0, \ldots, \boldsymbol{\xi}_{T-1}) \in \mathcal{L}_{N_\xi}^2, & N_\xi := 1 + n_\xi(T-1) \\ \boldsymbol{y} := (\boldsymbol{y}_0, \ldots, \boldsymbol{y}_{T-1}) \in \mathcal{L}_{N_y}^2, & N_y := 1 + n_y(T-1), \end{array}$$

where $\boldsymbol{\xi}_0$ and $\boldsymbol{y}_0$ represent degenerate random variables that are almost surely equal to 1. The matrices B, C, D and E in (2.2) are easily constructed from the problem data in (2.1) and are defined in the Appendix. In particular, note that the known initial state x_1 appears as a component of the problem data used in constructing the matrix C. Throughout the paper, we will call $\boldsymbol{x}$ the state process, $\boldsymbol{u}$ the control policy or control process, $\boldsymbol{\xi}$ the noise or disturbance process and $\boldsymbol{y}$ the measurement or observation process. We include the degenerate random variables $(\boldsymbol{\xi}_0, \boldsymbol{y}_0)$ as a conceptual device that will simplify significantly later mathematical developments in the paper.

Remark 1. *We assume without loss of generality that the initial state $\boldsymbol{x}_1$ is known and given by x_1. If that is not the case, we can set $\boldsymbol{x}_1$ to an arbitrary value and identify $\boldsymbol{x}_2$ with the initial state. We can then assign to $\boldsymbol{x}_2$ a prescribed distribution by tying it to $\boldsymbol{\xi}_1$ through a suitable choice of the dynamic system matrices for $t = 1$.*

We consider only physically implementable *causal* control policies, which rely solely on information available by observing the measurement process $\boldsymbol{y}$. To this end, we let $\mathcal{N}$ be the linear space of all *causal* control policies

$$\mathcal{N} := \times_{t=1}^{T-1} \mathcal{L}^2(\Omega, \mathcal{F}_t^y, \mathbb{P}; \mathbb{R}^{n_u}) \subseteq \mathcal{L}^2_{N_u},$$

where $\mathcal{F}_t^y := \sigma(\boldsymbol{y}_0, \ldots, \boldsymbol{y}_t)$ is the σ-algebra generated by the history of the observation process up to time t. We are interested in stochastic optimal control problems of the following type:

$$\begin{array}{ll} \inf & \mathbb{E}\left[\boldsymbol{u}^\top J_u \boldsymbol{u} + \boldsymbol{x}^\top J_x \boldsymbol{x}\right] \\ \text{s.t.} & \boldsymbol{u} \in \mathcal{N},\ \boldsymbol{s} \in \mathcal{L}^2_{N_s} \\ & \left.\begin{array}{l} \boldsymbol{x} = B\boldsymbol{u} + C\boldsymbol{\xi} \\ F_u \boldsymbol{u} + F_x \boldsymbol{x} + F_s \boldsymbol{s} = h \\ \boldsymbol{s} \geq 0 \end{array}\right\} \quad \mathbb{P}\text{-a.s.} \end{array} \tag{$\mathcal{P}$}$$

Our aim is to find a causal control policy $\boldsymbol{u} \in \mathcal{N}$ that minimizes the expectation of the quadratic cost $\boldsymbol{u}^\top J_u \boldsymbol{u} + \boldsymbol{x}^\top J_x \boldsymbol{x}$. The requirement that $\boldsymbol{u} \in \mathcal{N}$ is equivalent to that of each $\boldsymbol{u}_t$ being a function of the observation history $\boldsymbol{y}^t = (\boldsymbol{y}_0, \ldots, \boldsymbol{y}_t)$ available at that time. The control policy is selected subject to linear joint state and control constraints of the form $F_u \boldsymbol{u} + F_x \boldsymbol{x} + F_s \boldsymbol{s} = h$, where $h \in \mathbb{R}^{N_c}$, $F_u \in \mathbb{R}^{N_c \times N_u}$, $F_x \in \mathbb{R}^{N_c \times N_x}$, $F_s \in \mathbb{R}^{N_c \times N_s}$, where N_c is the number of state and control constraints. The random vector $\boldsymbol{s} \in \mathcal{L}^2_{N_s}$ can be interpreted as a vector of N_s slack variables, restricted to be nonnegative. Note that any row of F_s that contains only zeros corresponds to an equality constraint involving only $\boldsymbol{u}$ and $\boldsymbol{x}$. All other constraints include slack variables and thus correspond to inequality constraints for $\boldsymbol{u}$ and $\boldsymbol{x}$. It is assumed without loss of generality that the matrix F_s has full column rank. Finally, we require the dynamic system equation $\boldsymbol{x} = B\boldsymbol{u} + C\boldsymbol{\xi}$ to hold. Recall that this constraint includes the requirement that the initial state $\boldsymbol{x}_1 = x_1$.

Some conditions need to be imposed for problem $\mathcal{P}$ to be well-defined. We assume that the matrices J_u and J_x are positive semidefinite to ensure convexity of the objective function. Additionally, we assume that the support of the noise process $\boldsymbol{\xi}$ is non-empty, compact and representable as

$$\Xi := \left\{\xi \in \mathbb{R}^{N_\xi} : W_i \xi \succeq_{\mathcal{K}_p} 0,\ i = 1 \ldots l,\ e_0^\top \xi = 1\right\} \tag{2.3}$$

for some matrices $W_i \in \mathbb{R}^{N_\xi \times N_\xi}$. The last constraint in (2.3) ensures that $\xi_0 = 1$ for all $\xi \in \Xi$, which is consistent with our previous assumption that $\boldsymbol{\xi}_0 = 1$ almost surely. We assume that the second-order moment matrix associated with the noise process $\boldsymbol{\xi}$ is known and defined as

$$M := \mathbb{E}\left(\boldsymbol{\xi}\boldsymbol{\xi}^\top\right). \tag{2.4}$$

We further assume that the support Ξ spans the state space $\mathbb{R}^{N_\xi}$ of the noise process. This is equivalent to assuming that there is a $\xi \in \Xi$ that satisfies the strict inequalities $W_i \xi \succ_{\mathcal{K}_p} 0$ for all $i = 1 \ldots l$. Note that this last assumption is non-restrictive and can often be enforced by reducing the dimension of the noise process. This assumption also ensures that M is invertible and consequently that $M \succ 0$.

The cases of greatest general interest are $p \in \{1, 2, \infty\}$. In particular, it is easy to show that any subset of the hyperplane $\{\xi \in \mathbb{R}^{N_\xi} : \xi_0 = 1\}$ that results from a finite intersection of arbitrary ellipsoids and half spaces is representable in the form (2.3) for $p = 2$, so that the class of disturbances we consider includes those with polyhedral support.

2.1 Controller Information Structure

The information available to the causal controllers $\boldsymbol{u} \in \mathcal{N}$ at each time instance t can be interpreted as the σ-algebra $\mathcal{F}_t^y$, i.e. the information available by the history of the observation process $\boldsymbol{y}$ up to time t. For any causal policy $\boldsymbol{u} \in \mathcal{N}$, $\boldsymbol{u}_t$ must be $\mathcal{F}_t^y$-measurable for each $t = 1, \dots, T-1$. Since the observation process $\boldsymbol{y}$ depends on the control policy $\boldsymbol{u}$, the filtration $\mathbb{F}^y := \{\mathcal{F}_t^y\}_{t=0}^{T-1}$ seems to depend on $\boldsymbol{u}$ as well. In order to show that $\mathbb{F}^y$ is in fact independent of $\boldsymbol{u}$, we consider the purified observation process

$$\boldsymbol{\eta} := (\boldsymbol{\eta}_0, \dots, \boldsymbol{\eta}_{T-1}) \in \mathcal{L}^2_{N_y}$$

with $\boldsymbol{\eta} = (DC + E)\boldsymbol{\xi} = G\boldsymbol{\xi}$, where $G := (DC + E) \in \mathbb{R}^{N_y \times N_\xi}$. As described in [3], the purified observation $\boldsymbol{\eta}_t$ can be interpreted as the difference between the actual observation $\boldsymbol{y}_t$ and the observation that would have resulted at time t from a completely noise-free system obeying the same control policy. Note that $\boldsymbol{y}$ is linear in the control policy $\boldsymbol{u}$ and the purified observation process $\boldsymbol{\eta}$, i.e. $\boldsymbol{y} = DB\boldsymbol{u} + \boldsymbol{\eta}$. Now let $\mathcal{F}_t := \sigma(\boldsymbol{\eta}_0, \dots, \boldsymbol{\eta}_t)$ be the σ-algebra generated by the history of the purified observation process up to time t. By construction, $\boldsymbol{\eta}$ and the filtration $\mathbb{F} := \{\mathcal{F}_t\}_{t=0}^{T-1}$ are independent of the control strategy $\boldsymbol{u}$.

Proposition 2.1. *The filtrations $\mathbb{F}^y$ and $\mathbb{F}$ are identical.*

Proposition 2.1 shows that the information structure generated by the observation process is not decision-dependent. Furthermore, it shows that the observation process $\boldsymbol{y}$ and the purified observation process $\boldsymbol{\eta}$ convey the same amount of information. This in turn implies that the set of all control policies that are adapted to $\boldsymbol{y}$ is equivalent to the set of all policies that are adapted to $\boldsymbol{\eta}$. As a result, we can re-define the set of all implementable causal control policies $\mathcal{N}$ in terms of the σ-algebras $\{\mathcal{F}_t\}_{t=0}^{T-1}$.

In the remainder of the paper we will provide tractable methods for calculating upper and lower approximations to the problem $\mathcal{P}$ using affine decision rules. Although no further assumptions are required for the upper bound to be presented in Section 3, some further assumption relating to the purified observation process $\boldsymbol{\eta}$ will be required for the lower bounds presented in Section 4:

A1 (Noise Process) *The expectation of the noise process $\boldsymbol{\xi}$ conditioned on the purified observation process $\boldsymbol{\eta}$ is linear in $\boldsymbol{\eta}$, i.e. there exists a matrix $L \in \mathbb{R}^{N_\xi \times N_y}$ so that*

$$\mathbb{E}\left[\boldsymbol{\xi}\,|\boldsymbol{\eta}\right] = L\,\boldsymbol{\eta}\ \mathbb{P}\text{-a.s.} \tag{2.5}$$

Furthermore, the conditional expectation of $\boldsymbol{\eta}$ given its history $\boldsymbol{\eta}^t := (\boldsymbol{\eta}_0, \dots, \boldsymbol{\eta}_t)$ up to time t, is linear in $\boldsymbol{\eta}^t$, i.e. there exist matrices $H_t \in \mathbb{R}^{N_y \times (1+tn_y)}$ so that

$$\mathbb{E}\left[\boldsymbol{\eta}|\,\boldsymbol{\eta}^t\right] = H_t\boldsymbol{\eta}^t\ \mathbb{P}\text{-a.s. for all } t = 1, \dots, T-1. \tag{2.6}$$

The assumptions of **A1** appear to be restrictive. However, they are satisfiable in two special cases of considerable practical and theoretical interest:

Proposition 2.2 (Elliptically Contoured Distributions)**.** *The conditions of **A1** are satisfied if the distribution of the noise process $\boldsymbol{\xi}$ follows an elliptically contoured distribution.*

Proof. To demonstrate that (2.5) holds, we can use [11, Thm. 1] to show that the joint random vector $(\boldsymbol{\xi}, \boldsymbol{\eta})$, as the image of $\boldsymbol{\xi}$ under the linear transformation $\left(I^{(N_\xi)}, G^\top\right)^\top$, also follows an elliptically contoured distribution. This result allows us to invoke [11, Cor. 5] which guarantees that the conditional expectation of $\boldsymbol{\xi}$ given $\boldsymbol{\eta}$ is almost surely linear in $\boldsymbol{\eta}$. A similar argument can be used to establish (2.6). □

The preceding result confirms that **A1** is very broadly applicable, since the class of elliptically contoured distributions covers a wide range of distributions including multivariate normal, logistic, symmetric stable, generalized-hyperbolic and t-distributions as well as their truncated versions [12].

In the case of perfect state information, a weaker set of conditions are required to ensure satisfaction of **A1**:

Proposition 2.3 (Perfect State Information)**.** *The conditions of **A1** are satisfied if the following conditions all hold:*

i) The system state can be measured perfectly, i.e. $\boldsymbol{y}_t = \boldsymbol{x}_t$ *for all* $t = 1, \ldots, T-1$.

ii) There exist matrices $N_t \in \mathbb{R}^{N_\xi \times (1+tn_\xi)}$ *such that*

$$\mathbb{E}\left[\boldsymbol{\xi} | \, \boldsymbol{\xi}^t\right] = N_t \boldsymbol{\xi}^t \ \mathbb{P}\text{-a.s. for all } \ t = 1, \ldots, T. \tag{2.7}$$

iii) The matrices C_t *have full column rank for all* $t = 1, \ldots, T-1$.

Proof. In the perfect information case under consideration, the system equations (2.1) simplify because $n_y = n_x$ while $D_t = I$ and $E_t = 0$ for all $t = 1, \ldots, T-1$. The matrix G then simplifies to

$$G = \begin{bmatrix} 1 & & & & & \\ x_1 & & & & & \\ \mathcal{A}_1^2 x_1 & C_1 & & & & \\ \mathcal{A}_1^3 x_1 & \mathcal{A}_2^3 C_1 & C_2 & & & \\ \vdots & & & \ddots & & \\ \mathcal{A}_1^T x_1 & \mathcal{A}_2^T C_1 & \mathcal{A}_3^T C_2 & \cdots & C_{T-2} & 0 \end{bmatrix},$$

where the last n_ξ columns contain only zeros. We can now define $\hat{G}$ as the sub-matrix of G which is obtained by removing the last n_ξ columns of G. By construction, we have $\boldsymbol{\eta} = G\boldsymbol{\xi} = \hat{G}\boldsymbol{\xi}^{T-2}$ since $\boldsymbol{\eta}$ does not depend on $\boldsymbol{\xi}_{T-1}$ when $E_{T-1} = 0$. Note that $\hat{G}$ has full column rank, and thus the purified observation process $\boldsymbol{\eta}$ and the truncated noise process $\boldsymbol{\xi}^{T-2}$ generate the same σ-algebra. This implies that

$$\mathbb{E}\left[\boldsymbol{\xi} | \boldsymbol{\eta}\right] = \mathbb{E}\left[\boldsymbol{\xi} | \boldsymbol{\xi}^{T-2}\right] = N_{T-2}\, \boldsymbol{\xi}^{T-2} = N_{T-2}\, \hat{G}^\dagger \, \hat{G}\, \boldsymbol{\xi}^{T-2} = N_{T-2}\, \hat{G}^\dagger \, \boldsymbol{\eta}.$$

A similar argument can be used to establish (2.6). □

Remark 2. *The conditions required in Proposition 2.3 are very mild. Condition (ii) is trivially satisfied if, for example, the noise vectors* $\boldsymbol{\xi}_t$ *are serially independent. Condition (iii) is a common assumption for control problems of the type* $\mathcal{P}$ *with perfect state measurements, and can usually be enforced by reducing the dimension of the noise process.*

3 Upper Approximations of $\mathcal{P}$ and Primal Affine Control Policies

The optimization problem $\mathcal{P}$ generally involves an infinite number of decision variables and constraints. Without a suitable approximation, this problem is not amenable to numerical solution. In this section we impose a restriction on the feasible set to achieve computational tractability. Any reduction of the feasible set corresponds to a conservative approximation of the original problem and thus results in an upper bound on the true optimal value of $\mathcal{P}$. Here, we calculate such an upper bound by restricting the class of causal control policies $\mathcal{N}$ to the subclass of all control policies that are affine with respect to the observations. The affine observation-feedback policies can be written as

$$\boldsymbol{u}_t = u_{t,0} + \sum_{s=1}^{t} U_{t,s}\boldsymbol{y}_s \quad \text{for} \quad t = 1, \ldots, T-1,$$

for some matrices $U_{t,s} \in \mathbb{R}^{n_u \times n_y}$. In compact notation, we have $\boldsymbol{u} = U\boldsymbol{y}$ almost surely (since $\boldsymbol{y}_0 = 1$ almost surely), where

$$U := \begin{bmatrix} u_{1,0} & U_{1,1} & & & \\ u_{2,0} & U_{2,1} & U_{2,2} & & \\ \vdots & & & \ddots & \\ u_{T-1,0} & U_{T-1,1} & U_{T-1,2} & \cdots & U_{T-1,T-1} \end{bmatrix}. \tag{3.8}$$

We denote by $\mathcal{U}$ the linear space of all block lower triangular matrices of the form (3.8).

Proposition 3.1. *An upper bound to the problem $\mathcal{P}$ can be found by solving the following optimization problem, where the control policy $\boldsymbol{u}$ is parameterized as an affine function of the outputs $\boldsymbol{y}$:*

$$\begin{array}{ll} \inf & \mathbb{E}\left[\boldsymbol{u}^\top J_u \boldsymbol{u} + \boldsymbol{x}^\top J_x \boldsymbol{x}\right] \\ s.t. & U \in \mathcal{U},\ \boldsymbol{s} \in \mathcal{L}^2_{N_s} \\ & \left.\begin{array}{l} \boldsymbol{x} = B\boldsymbol{u} + C\boldsymbol{\xi},\ \boldsymbol{u} = U\boldsymbol{y} \\ F_u\boldsymbol{u} + F_x\boldsymbol{x} + F_s\boldsymbol{s} = h \\ \boldsymbol{s} \geq 0 \end{array}\right\} \quad \mathbb{P}\text{-}a.s. \end{array}$$

The above problem is itself equivalent to the following optimization problem, where the control policy $\boldsymbol{u}$ is parameterized as an affine function of the purified outputs $\boldsymbol{\eta} = G\boldsymbol{\xi}$:

$$\begin{array}{ll} \inf & \mathbb{E}\left[\boldsymbol{u}^\top J_u \boldsymbol{u} + \boldsymbol{x}^\top J_x \boldsymbol{x}\right] \\ s.t. & Q \in \mathcal{U},\ \boldsymbol{s} \in \mathcal{L}^2_{N_s} \\ & \left.\begin{array}{l} \boldsymbol{x} = B\boldsymbol{u} + C\boldsymbol{\xi},\ \boldsymbol{u} = QG\boldsymbol{\xi} \\ F_u\boldsymbol{u} + F_x\boldsymbol{x} + F_s\boldsymbol{s} = h \\ \boldsymbol{s} \geq 0 \end{array}\right\} \quad \mathbb{P}\text{-}a.s. \end{array} \tag{$\mathcal{P}_u$}$$

Note that the two equivalent problem formulations in Proposition 3.1 differ only in their representation of the control policy $\boldsymbol{u}$. In the first case, the control input is parameterized as a causal affine function of prior measurements $\boldsymbol{y}$, while in the second case the input is parameterized as a causal affine function of the purified outputs $\boldsymbol{\eta} = G\boldsymbol{\xi}$. In both cases, the set of feasible control policies is a subset of the set of all possible control policies $\mathcal{N}$.

Before proceeding to the proof of Proposition 3.1, we note that problem $\mathcal{P}_u$ remains seemingly intractable due to the presence of functional decision variables and constraints. We therefore require an additional result that provides a method for solving the problem $\mathcal{P}_u$ as a tractable convex optimization problem.

Proposition 3.2. *A solution to problem $\mathcal{P}_u$ can be found by solving the following equivalent convex optimization problem:*

$$
\begin{array}{ll}
\inf & \operatorname{tr}\left(G^\top Q^\top (J_u + B^\top J_x B) Q G M + 2 C^\top J_x B Q G M + C^\top J_x C M\right) \\
s.t. & Q \in \mathcal{U},\ S \in \mathbb{R}^{N_s \times N_\xi},\ \Lambda_i \in \mathbb{R}^{N_\xi \times N_s},\ \mu \in \mathbb{R}^{N_s} \\
& (F_u + F_x B) Q G + F_x C + F_s S - h e_0^\top = 0 \\
& S = \mu e_0^\top + \sum_{i=1}^{l} \Lambda_i^\top W_i \\
& \Lambda_i \succeq_{\mathcal{K}_p^*} 0,\ i = 1 \ldots l \\
& \mu \geq 0.
\end{array}
\qquad (\tilde{\mathcal{P}}_u)
$$

The problem $\tilde{\mathcal{P}}_u$ is a tractable conic optimization problem and can be solved in polynomial time for any $p \in [1, +\infty]$.

The problem $\tilde{\mathcal{P}}_u$ is of interest because it provides a tractable (though suboptimal) solution to $\mathcal{P}$ by restricting the class of control policies over which the optimization is performed. However, there are at present no sophisticated methods available for estimating the degree of suboptimality introduced by imposing such a restriction. In Section 4, we develop a method for calculating a *lower* bound on $\mathcal{P}$ by formulating a problem dual to $\mathcal{P}$ and restricting the dual variables to affine form to ensure tractability. This will enable us to bound the suboptimality in $\tilde{\mathcal{P}}_u$ in Section 5.

Remark 3. *The second part of Proposition 3.1 provides a method for optimizing over affine output feedback policies $\boldsymbol{u} = \mathcal{U}\boldsymbol{y}$ by re-parameterizing the problem into one of optimizing over a different but equivalent class of parametric policies, for which $\boldsymbol{u}$ and $\boldsymbol{x}$ become affine functions of the parameters [4, 15, 24]. The resulting optimization problem can then be solved using Proposition 3.2. The idea underpinning this re-parameterization follows a similar approach to the classical Youla- or Q-parameterization procedure [28]. The benefit of this approach is that it avoids direct optimization over the parameter $U \in \mathcal{U}$, which is impractical since both $\boldsymbol{u}$ and $\boldsymbol{x}$ become nonlinear (rational) functions of the coefficients of U. In order to keep the paper self-contained and to facilitate comparison with the lower bounding methods of Section 4, we sketch the proofs using our notation.*

Corollary 3.3. *A trivial upper bound to $\mathcal{P}$ can be found by considering only stationary (i.e. open-loop) policies:*

$$
\begin{array}{ll}
\inf & u^\top (J_u + B^\top J_x B) u + 2 u^\top B^\top J_x C \mathbb{E}[\boldsymbol{\xi}] + \operatorname{tr}\left(C^\top J_x C M\right) \\
s.t. & u \in \mathbb{R}^{N_u},\ S \in \mathbb{R}^{N_s \times N_\xi},\ \Lambda_i \in \mathbb{R}^{N_\xi \times N_s},\ \mu \in \mathbb{R}^{N_s} \\
& (F_u + F_x B) u e_0^\top + F_x C + F_s S - h e_0^\top = 0 \\
& S = \mu e_0^\top + \sum_{i=1}^{l} \Lambda_i^\top W_i \\
& \Lambda_i \succeq_{\mathcal{K}_p^*} 0,\ i = 1 \ldots l \\
& \mu \geq 0.
\end{array}
\qquad (\mathcal{P}_{uu})
$$

where $u \in \mathbb{R}^{N_u}$ is the vector of the stationary controls. It is easy to verify that

$$\mathcal{P} \leq \mathcal{P}_u \leq \mathcal{P}_{uu}.$$

Problem $\mathcal{P}_{uu}$ is a control optimization problem over open-loop robust controllers. Note that $\mathcal{P}_{uu}$ is a restriction of $\mathcal{P}_u$. One can derive $\mathcal{P}_{uu}$ from $\mathcal{P}_u$ by restricting the control policies to be stationary instead of affine on $\boldsymbol{y}$. As a result $\mathcal{P}_u$ is a tighter upper bound to $\mathcal{P}$ than $\mathcal{P}_{uu}$.

In the remainder of this section we supply proofs for Propositions 3.1 and 3.2.

Proof of Proposition 3.1

Proof of the first part of Proposition 3.1 is straightforward, since any control law in the affine form $\boldsymbol{u} = U\boldsymbol{y}$ is causal due to the lower triangular structure imposed by the constraint $U \in \mathcal{U}$, hence $\boldsymbol{u} \in \mathcal{N}$.

To prove the second part, we show that any control law in the affine output feedback form $\boldsymbol{u} = U\boldsymbol{y}$ for some $U \in \mathcal{U}$ can be matched exactly by a control law in the purified output feedback form $\boldsymbol{u} = Q\boldsymbol{\eta} = QG\boldsymbol{\xi}$ for some $Q \in \mathcal{U}$, and vice-versa. To prove the first case, assume that $\boldsymbol{u} = U\boldsymbol{y}$ and solve for the control input $\boldsymbol{u}$ in terms of the purified outputs $\boldsymbol{\eta}$, yielding

$$\boldsymbol{u} = (I - UDB)^{-1}U\boldsymbol{\eta}.$$

We then set $Q = (I - UDB)^{-1}U$, noting that the required matrix inverse exists since UDB is strictly lower triangular, which also ensures that $Q \in \mathcal{U}$. Proof of the second case is provided by a similar argument. □

Proof of Proposition 3.2

Both $\boldsymbol{u}$ and $\boldsymbol{x}$ are linear in $\boldsymbol{\xi}$ in the problem $\mathcal{P}_u$. By pre-multiplying the joint state and control constraint of $\mathcal{P}_u$ with the left inverse of F_s (which exists since F_s is assumed to have full column rank), it can be seen that the vector of slack variables $\boldsymbol{s}$ appearing in $\mathcal{P}_u$ is also linear in $\boldsymbol{\xi}$, so that $\boldsymbol{s} = S\boldsymbol{\xi}$ for some matrix $S \in \mathbb{R}^{N_s \times N_\xi}$. Noting also that $h = he_0^\top \boldsymbol{\xi}$ almost surely, we can rewrite the problem $\mathcal{P}_u$ as

$$\begin{array}{ll}
\inf & \mathbb{E}\left[\boldsymbol{u}^\top J_u \boldsymbol{u} + \boldsymbol{x}^\top J_x \boldsymbol{x}\right] \\
\text{s.t.} & Q \in \mathcal{U},\ S \in \mathbb{R}^{N_s \times N_\xi} \\
& \left.\begin{array}{l}
\boldsymbol{x} = B\boldsymbol{u} + C\boldsymbol{\xi},\ \ \boldsymbol{u} = QG\boldsymbol{\xi} \\
(F_u QG + F_x BQG + F_x C + F_s S - he_0^\top)\boldsymbol{\xi} = 0 \\
S\boldsymbol{\xi} \geq 0
\end{array}\right\} \quad \mathbb{P}\text{-a.s.}
\end{array} \tag{3.9}$$

Observe that due to the continuity of the constraint functions in $\boldsymbol{\xi}$, the almost sure constraints in (3.9) must hold for all $\xi \in \Xi$. We can therefore replace the almost sure constraints with semi-infinite constraints to obtain the equivalent semi-infinite problem

$$\begin{array}{ll}
\inf & \operatorname{tr}\left(G^\top Q^\top (J_u + B^\top J_x B) QGM + 2(C^\top J_x BQGM) + C^\top J_x CM\right) \\
\text{s.t.} & Q \in \mathcal{U},\ S \in \mathbb{R}^{N_s \times N_\xi} \\
& \left.\begin{array}{l}
(F_u QG + F_x BQG + F_x C + F_s S - he_0^\top)\xi = 0 \\
S\xi \geq 0
\end{array}\right\} \quad \forall \xi \in \Xi,
\end{array} \tag{3.10}$$

where we have also eliminated $\boldsymbol{u}$ and $\boldsymbol{x}$ and rewritten the objective function in terms of the second-order moment matrix $M = \mathbb{E}\left[\boldsymbol{\xi}\boldsymbol{\xi}^\top\right]$.

The equality constraints in (3.10) must hold for all $\xi \in \Xi$, meaning that the linear hull of Ξ belongs to the null space of the linear operator $(F_u + F_x B)QG + F_s S + F_x C - he_0^\top$. Since Ξ is assumed to span the entire space $\mathbb{R}^{N_\xi}$, these semi-infinite equality constraints are equivalent to the matrix equality $(F_u + F_x B)QG + F_s S + F_x C - he_0^\top = 0$. Next, we simplify the semi-infinite inequality constraints by using techniques that are commonly used in robust optimization [5, 7, 8]. The following lemma, which is a special case of [8, Theorem 3.1], describes the key enabling mechanism:

Lemma 3.4. *For any $z \in \mathbb{R}^{N_\xi}$, the following statements are equivalent:*

i) $z^\top \xi \geq 0$ *for all* $\xi \in \Xi$

ii) There exist vectors $\lambda_i \in \mathbb{R}^{N_\xi}$ *for* $i \in \{1 \ldots l\}$ *and a scalar* $\mu \geq 0$ *such that*

$$z^\top = \mu e_0^\top + \sum_{i=1}^{l} \lambda_i^\top W_i$$
$$\lambda_i \succeq_{\mathcal{K}_p^*} 0\ ,\ i = 1 \ldots l.$$

Proof of this lemma is included in the appendix.

Let $z_\kappa^\top$ denote the κ-th row of the matrix S in problem (3.10). We can then use Lemma 3.4 to replace the semi-infinite constraint $z_\kappa^\top \xi \geq 0 \, \forall \xi \in \Xi$ by finitely many linear constraints involving the new decision variables $\mu_\kappa \in \mathbb{R}$ and $\lambda_{\kappa,i} \in \mathbb{R}^{N_\xi}$ for $i = 1, \ldots, l$. The constraint $S\xi \geq 0 \, \forall \xi \in \Xi$ can thus be reformulated in terms of $\mu = (\mu_1, \ldots, \mu_{N_s})^\top \in \mathbb{R}^{N_s}$ and $\Lambda_i = (\lambda_{1,i} \ \cdots \ \lambda_{N_s,i}) \in \mathbb{R}^{N_\xi \times N_s}$, $i = 1, \ldots, l$ to yield the optimization problem $\tilde{\mathcal{P}}_u$. Solvability of $\tilde{\mathcal{P}}_u$ in polynomial time for any $p \in [1, +\infty]$ is ensured by [27]. □

4 Lower Approximations of $\mathcal{P}$ and Dual Affine Control Policies

The use of primal affine control policies leads to a conservative approximation for the original problem $\mathcal{P}$, whose computational complexity scales gracefully with the horizon length and the dimension of the system (2.1). It also yields an implementable control policy which is feasible, but typically suboptimal, in $\mathcal{P}$. The main goal of this paper is to estimate the loss of optimality due to the use of affine policies. To this end, we now establish a lower bound on $\mathcal{P}$ by extending the dual linear decision rule techniques that were recently introduced in the context of stochastic programming; see [18]. The loss of optimality incurred by using affine controllers can then be bounded by the difference between the upper and lower bounds on the optimization problem $\mathcal{P}$. In order to derive the lower bound, we first reformulate $\mathcal{P}$ as a min-max problem by introducing a dual control policy $\boldsymbol{\nu} \in \mathcal{L}^2_{N_c}$ and by moving the joint state and control constraints to the objective function, to obtain

$$\begin{array}{lll} \inf & \sup\limits_{\boldsymbol{\nu} \in \mathcal{L}^2_{N_c}} \mathbb{E}\left[\boldsymbol{u}^\top J_u \boldsymbol{u} + \boldsymbol{x}^\top J_x \boldsymbol{x} + \boldsymbol{\nu}^\top \left[F_u \boldsymbol{u} + F_x \boldsymbol{x} + F_s \boldsymbol{s} - h\right]\right] & \\ \text{s.t.} & \boldsymbol{u} \in \mathcal{N},\ \boldsymbol{s} \in \mathcal{L}^2_{N_s} & \\ & \left.\begin{array}{l} \boldsymbol{x} = B\boldsymbol{u} + C\boldsymbol{\xi} \\ \boldsymbol{s} \geq 0 \end{array}\right\} \ \mathbb{P}\text{-a.s.} & \end{array} \tag{$\mathcal{P}'$}$$

The inner maximization over the dual control policies in $\mathcal{P}'$ ensures that any violation of the state and control constraints on a set of strictly positive probability incurs an infinite penalty. The two problems $\mathcal{P}$ and $\mathcal{P}'$ are equivalent. In order to derive a tractable lower bounding approximation for $\mathcal{P}$, we will adopt the same approach as in the primal problem in Section 3 and restrict the space of all dual control policies to those that depend linearly on the noise process, and are thus representable as $\boldsymbol{\nu} = Y\boldsymbol{\xi}$ for some $Y \in \mathbb{R}^{N_c \times N_\xi}$. We will then show that the resulting approximation problem can be solved by solving an equivalent tractable conic optimization problem.

The first of these results forms a natural lower-bounding counterpart to Proposition 3.1:

Proposition 4.1. *A lower bound to the problem $\mathcal{P}$ (equivalently $\mathcal{P}'$) can be found by solving the following optimization problem, where the dual policy variable $\boldsymbol{\nu}$ is parameterized as an affine function of the disturbance $\boldsymbol{\xi}$:*

$$
\begin{array}{lll}
\inf & \sup\limits_{Y \in \mathbb{R}^{N_c \times N_\xi}} & \mathbb{E}\left[\boldsymbol{u}^\top J_u \boldsymbol{u} + \boldsymbol{x}^\top J_x \boldsymbol{x} + \boldsymbol{\nu}^\top \left[F_u \boldsymbol{u} + F_x \boldsymbol{x} + F_s \boldsymbol{s} - h\right]\right] \\
s.t. & \multicolumn{2}{l}{\boldsymbol{u} \in \mathcal{N},\ \boldsymbol{s} \in \mathcal{L}^2_{N_s}} \\
& \multicolumn{2}{l}{\left.\begin{array}{l} \boldsymbol{x} = B\boldsymbol{u} + C\boldsymbol{\xi} \\ \boldsymbol{\nu} = Y\boldsymbol{\xi},\ \boldsymbol{s} \geq 0 \end{array}\right\} \quad \mathbb{P}\text{-}a.s.}
\end{array}
$$

The above min-max problem is itself equivalent to the following minimization problem, whose solution provides a lower bound to the problem $\mathcal{P}$:

$$
\begin{array}{ll}
\inf & \mathbb{E}\left[\boldsymbol{u}^\top J_u \boldsymbol{u} + \boldsymbol{x}^\top J_x \boldsymbol{x}\right] \\
s.t. & \boldsymbol{u} \in \mathcal{N},\ \boldsymbol{s} \in \mathcal{L}^2_{N_s} \\
& \left.\begin{array}{l} \boldsymbol{x} = B\boldsymbol{u} + C\boldsymbol{\xi} \\ \mathbb{E}\left[(F_u \boldsymbol{u} + F_x \boldsymbol{x} + F_s \boldsymbol{s} - h)\,\boldsymbol{\xi}^\top\right] = 0 \\ \boldsymbol{s} \geq 0 \end{array}\right\} \quad \mathbb{P}\text{-}a.s.
\end{array}
\tag{$\mathcal{P}_\ell$}
$$

Before proceeding to the proof of Proposition 4.1, we note that although $\mathcal{P}_\ell$ has only finitely many state and control constraints, it remains seemingly intractable since it involves functional decision variables and functional inequality constraints. We therefore require a method for solving the problem $\mathcal{P}_\ell$ as a tractable convex optimization problem. The following result provides a solution to this problem and forms a natural lower-bounding counterpart to Proposition 3.2:

Proposition 4.2. *If **A1** holds, then a lower-bounding solution to problem $\mathcal{P}_\ell$ can be found by solving the following convex optimization problem:*

$$
\begin{array}{ll}
\inf & \operatorname{tr}\left(G^\top Q^\top (J_u + B^\top J_x B) Q G M + 2C^\top J_x B Q G M + C^\top J_x C M\right) \\
s.t. & Q \in \mathcal{U},\ S \in \mathbb{R}^{N_s \times N_\xi} \\
& (F_u + F_x B) Q G + F_s S + F_x C - h e_0^\top = 0 \\
& W_i M S^\top \succeq_{\mathcal{K}_p} 0,\ i = 1 \ldots l \\
& e_0^\top M S^\top \geq 0.
\end{array}
\tag{$\tilde{\mathcal{P}}_\ell$}
$$

The problem $\tilde{\mathcal{P}}_\ell$ is a tractable conic optimization problem and can be solved in polynomial time for any $p \in [1, +\infty]$.

Remark 4. *The problem $\mathcal{P}_\ell$ can be shown to be equivalent to the finite-dimensional problem $\tilde{\mathcal{P}}_\ell$ whenever $\mathcal{P}_\ell$ contains a strictly feasible point. Proof of this claim is based on the observation that the constraints in $\tilde{\mathcal{P}}_\ell$ represent the closure of a set of constraints in a problem equivalent to $\mathcal{P}_\ell$ (and which is constructed in the proof). We omit details of this argument for the sake of brevity.*

The problem $\tilde{\mathcal{P}}_\ell$ provides a tractable (though suboptimal) method for finding a lower bound to $\mathcal{P}$, and therefore allows us to bound the degree of approximation in the affine policy optimization problem $\tilde{\mathcal{P}}_u$. Furthermore, there is an attractive structural similarity to the constraints in $\tilde{\mathcal{P}}_\ell$ and $\tilde{\mathcal{P}}_u$ – we elaborate on these points in Section 5.

Corollary 4.3. *A trivial lower bound to $\mathcal{P}$ can be found by considering a certainty equivalence problem:*

$$
\begin{array}{ll}
\inf & \overline{u}^\top J_u \overline{u} + \overline{x}^\top J_x \overline{x} \\
s.t. & \overline{u} \in \mathbb{R}^{N_u},\ \overline{s} \in \mathbb{R}^{N_s} \\
& \overline{x} = B\overline{u} + C\mathbb{E}\left[\boldsymbol{\xi}\right] \\
& F_u \overline{u} + F_x \overline{x} + F_s \overline{s} - h = 0 \\
& \overline{s} \geq 0
\end{array}
\qquad . \qquad (\mathcal{P}_{\ell\ell})
$$

It is easy to verify that $\mathcal{P} \geq \mathcal{P}_\ell \geq \mathcal{P}_{\ell\ell}$.

Problem $\mathcal{P}_{\ell\ell}$ is a certainty equivalence problem where only the mean of the uncertainty is considered. Note that the feasible region of $\mathcal{P}_{\ell\ell}$ is a subset of the feasible region of $\mathcal{P}_\ell$. One can retrieve the feasible region of $\mathcal{P}_{\ell\ell}$ by restricting the dual policy variable $\boldsymbol{\nu} \in \mathcal{L}^2_{N_c}$ in $\mathcal{P}'$ to be stationary instead of affine on $\boldsymbol{\xi}$. Furthermore, Jensen's Inequality ensures that the objective function in $\mathcal{P}_{\ell\ell}$ is a lower bound to the objective function in $\mathcal{P}_\ell$. As a result $\mathcal{P}_\ell$ is a tighter lower bound to $\mathcal{P}$ than $\mathcal{P}_{\ell\ell}$.

In the remainder of this section we supply proofs for Propositions 4.1 and 4.2.

Proof of Proposition 4.1

Proof of the first part of Proposition 4.1 is straightforward, since any dual variable in the affine form $\boldsymbol{\nu} = Y\boldsymbol{\xi}$ also satisfies $\boldsymbol{\nu} \in \mathcal{L}^2_{N_c}$. Hence the inner maximization over Y is a lower bound on the inner maximization that appears in the problem $\mathcal{P}'$.

Proof of the second part follows from the first since

$$
\begin{aligned}
\sup_Y \mathbb{E}\left[\boldsymbol{\xi}^\top Y^\top \left[F_u \boldsymbol{u} + F_x \boldsymbol{x} + F_s \boldsymbol{s} - h\right]\right] &= \sup_Y \mathbb{E}\left[\operatorname{tr}\left\{\left[F_u \boldsymbol{u} + F_x \boldsymbol{x} + F_s \boldsymbol{s} - h\right] \boldsymbol{\xi}^\top Y^\top\right\}\right] \\
&= \begin{cases} 0 & \text{if } \left[F_u \boldsymbol{u} + F_x \boldsymbol{x} + F_s \boldsymbol{s} - h\right] \boldsymbol{\xi}^\top = 0 \ \mathbb{P}\text{-a.s.} \\ \infty & \text{otherwise.} \end{cases}
\end{aligned}
$$

Problem $\mathcal{P}_\ell$ is clearly a relaxation of $\mathcal{P}$ since any $(\boldsymbol{u}, \boldsymbol{s})$ satisfying the almost sure state and control constraints in $\mathcal{P}$ will therefore also satisfy the expectation constraint in $\mathcal{P}_\ell$. □

Proof of Proposition 4.2

We first eliminate the state process $\boldsymbol{x}$ in $\mathcal{P}_\ell$ by substituting the dynamic system constraint into both the objective function and the expectation constraint in $\mathcal{P}_\ell$. Recalling the definition of the second-order moment matrix M, the approximate problem $\mathcal{P}_\ell$ can then be rewritten as

$$
\begin{array}{ll}
\inf & \mathbb{E}\left[\boldsymbol{u}^\top (J_u + B^\top J_x B)\boldsymbol{u}\right] + 2\operatorname{tr}(C^\top J_x B \mathbb{E}\left[\boldsymbol{u}\boldsymbol{\xi}^\top\right] + C^\top J_x C M) \\
\text{s.t} & \boldsymbol{u} \in \mathcal{N},\ \boldsymbol{s} \in \mathcal{L}^2_{N_s} \\
& (F_u + F_x B)\mathbb{E}\left[\boldsymbol{u}\boldsymbol{\xi}^\top\right] + F_s \mathbb{E}\left[\boldsymbol{s}\boldsymbol{\xi}^\top\right] + (F_x C - h e_0^\top)M = 0 \\
& \boldsymbol{s} \geq 0 \quad \mathbb{P}\text{-a.s.}
\end{array}
\qquad (4.11)
$$

The optimization problem (4.11) contains a variety of terms composed of the functional decision variables $(\boldsymbol{u}, \boldsymbol{s})$. In the remainder of the proof, we introduce a number of technical lemmas that will enable us to eliminate each of these terms in turn, replacing them with finite-dimensional variables. Proof of each of these lemmas is included in the appendix.

The first result will allow us to eliminate terms in the form $\mathbb{E}[\boldsymbol{u}\boldsymbol{\xi}^\top]$ and $\mathbb{E}[\boldsymbol{s}\,\boldsymbol{\xi}^\top]$:

Lemma 4.4. *For every $\boldsymbol{s} \in \mathcal{L}^2_{N_s}$ there exists a matrix $S \in \mathbb{R}^{N_s \times N_\xi}$ that satisfies*

$$SM = \mathbb{E}[\boldsymbol{s}\,\boldsymbol{\xi}^\top]. \tag{4.12}$$

*Likewise, if **A1** holds, then for every $\boldsymbol{u} \in \mathcal{N}$ there exists a block lower triangular matrix $Q \in \mathcal{U}$ that satisfies*

$$QGM = \mathbb{E}[\boldsymbol{u}\boldsymbol{\xi}^\top]. \tag{4.13}$$

Remark 5. *For any $Q \in \mathcal{U}$, the affine controller $\boldsymbol{u} = Q\boldsymbol{\eta} \in \mathcal{N}$ satisfies (4.13) since*

$$\mathbb{E}[Q\boldsymbol{\eta}\,\boldsymbol{\xi}^\top] = QG\mathbb{E}[\boldsymbol{\xi}\boldsymbol{\xi}^\top] = QGM.$$

The second part of Lemma 4.4 investigates the converse situation in which some $\boldsymbol{u} \in \mathcal{N}$ (not necessarily affine in $\boldsymbol{\eta}$) is given, and we seek some $Q \in \mathcal{U}$ satisfying (4.13).

The results of Lemma 4.4 allow us to add new decision variables $Q \in \mathcal{U}$ and $S \in \mathbb{R}^{N_s \times N_\xi}$ to problem (4.11) and append (4.12) and (4.13) as additional constraints without constraining the choice $\boldsymbol{u} \in \mathcal{N}$ and $\boldsymbol{s} \in \mathcal{L}^2_{N_s}$ in (4.11). Therefore (4.11) is equivalent to the optimization problem

$$\begin{array}{ll} \inf & \mathbb{E}\left[\boldsymbol{u}(J_u + B^\top J_x B)\boldsymbol{u}^\top\right] + 2\,\mathrm{tr}(C^\top J_x BQGM + C^\top J_x CM) \\ \text{s.t.} & \boldsymbol{u} \in \mathcal{N},\ Q \in \mathcal{U},\ S \in \mathbb{R}^{N_s \times N_\xi} \\ & (F_u + F_x B)QGM + F_s SM + (F_x C - he_0^\top)M = 0 \\ & QGM = \mathbb{E}\left[\boldsymbol{u}\,\boldsymbol{\xi}^\top\right] \\ & \exists \boldsymbol{s} \in \mathcal{L}^2_{N_s} : SM = \mathbb{E}\left[\boldsymbol{s}\,\boldsymbol{\xi}^\top\right],\ \boldsymbol{s} \geq 0 \quad \mathbb{P}\text{-a.s.} \end{array} \tag{4.14}$$

We next introduce a technical Lemma that will enable us to eliminate the functional decision variable $\boldsymbol{u}$ in (4.14) by explicitly minimizing the part of the objective function containing this term:

Lemma 4.5. *Consider the convex optimization problem*

$$\begin{array}{ll} \min & \mathbb{E}\left[\boldsymbol{u}^\top J \boldsymbol{u}\right] \\ s.t. & \boldsymbol{u} \in \mathcal{N} \\ & QGM = \mathbb{E}\left[\boldsymbol{u}\,\boldsymbol{\xi}^\top\right] \end{array} \tag{4.15}$$

where J is positive semi-definite and $Q \in \mathcal{U}$ is fixed. A minimizer for this problem is $\boldsymbol{u}^* = QG\boldsymbol{\xi}$.

We can apply Lemma 4.5 with $J := (J_u + B^\top J_x B)$ to (4.14) and replace $\boldsymbol{u}$ by $QG\boldsymbol{\xi}$. This yields the following optimization problem:

$$\begin{array}{ll} \inf & \mathrm{tr}\left(G^\top Q^\top (J_u + B^\top J_x B)QGM + 2C^\top J_x BQGM + C^\top J_x CM\right) \\ \text{s.t.} & Q \in \mathcal{U},\ S \in \mathbb{R}^{N_s \times N_\xi} \\ & (F_u + F_x B)QG + F_s S + F_x C - he_0^\top = 0 \\ & \exists \boldsymbol{s} \in \mathcal{L}^2_{N_s} : SM = \mathbb{E}\left[\boldsymbol{s}\,\boldsymbol{\xi}^\top\right],\ \boldsymbol{s} \geq 0 \quad \mathbb{P}\text{-a.s.} \end{array} \tag{4.16}$$

Note also that we removed the second-order moment matrix M from the state and control constraint by post-multiplying the corresponding constraint in (4.14) with M^{-1}.

The last constraint in (4.16) requires the solution of N_s moment problems, i.e. for a given $S \in \mathbb{R}^{N_s \times N_\xi}$ we must assert the existence of N_s non-negative Borel measures whose vectors of zero- and first-order moments coincide with the rows of SM and which have square integrable densities with respect to $\mathbb{P}$. The following lemma provides a means for dealing with this constraint by replacing it with a set of conic constraints that enclose its feasible region:

Lemma 4.6. *Define the cones* $\mathcal{C} \subseteq \mathbb{R}^{N_\xi}$ *and* $\mathcal{C}_{\boldsymbol{\xi}} \subseteq \mathbb{R}^{N_\xi}$ *as*

$$\mathcal{C}_{\boldsymbol{\xi}} := \left\{ z \in \mathbb{R}^{N_\xi} \ : \ \exists \boldsymbol{s} \in \mathcal{L}_1^2 \text{ with } z = \mathbb{E}\left[\boldsymbol{s}\boldsymbol{\xi}\right], \boldsymbol{s} \geq 0 \ \mathbb{P}\text{-a.s} \right\}$$
$$\mathcal{C} := \operatorname{cone}(\Xi) = \left\{ z \in \mathbb{R}^{N_\xi} \ : \ W_i\, z \succeq_{\mathcal{K}_p} 0, i = 1 \ldots l, e_0^\top z \geq 0 \right\}.$$

Then $\emptyset \neq \operatorname{int} \mathcal{C} \subseteq \mathcal{C}_{\boldsymbol{\xi}} \subseteq \mathcal{C}$.

Using the definition of $\mathcal{C}_{\boldsymbol{\xi}}$ in Lemma 4.6, we can rewrite (4.16) in the form

$$\begin{array}{ll} \inf & \operatorname{tr}\left(G^\top Q^\top (J_u + B^\top J_x B) Q G M + 2 C^\top J_x B Q G M + C^\top J_x C M\right) \\ \text{s.t.} & Q \in \mathcal{U}, \ S \in \mathbb{R}^{N_s \times N_\xi} \\ & (F_u + F_x B) Q G + F_s S + F_x C - h e_0^\top = 0 \\ & (M S^\top)_i \in C_{\boldsymbol{\xi}}, \ \forall i \in \{1, \ldots, N_s\}, \end{array} \tag{4.17}$$

where $(MS^\top)_i$ denotes the i^{th} column of $MS^\top$. The problem (4.17) is therefore equivalent to $\mathcal{P}_\ell$. A lower bound to (4.17) can then be found by taking the closure of its constraints via application of Lemma 4.6, replacing $\mathcal{C}_{\boldsymbol{\xi}}$ with $\mathcal{C}$, which results immediately in the problem $\tilde{\mathcal{P}}_\ell$. It can be shown that the application of this closure operation does not affect the optimal value except for pathological cases of little practical interest (cf. Remark 4), so that in most cases $\tilde{\mathcal{P}}_\ell$ is actually equivalent to (4.17). Note that $\mathcal{C}$ is the cone generated by Ξ.

Solvability of $\tilde{\mathcal{P}}_\ell$ in polynomial time for any $p \in [1, +\infty]$ is ensured by [27]. □

5 Geometric Properties

In this section we make a number of general observations about the relationship between the lower- and upper-bounding problems $\mathcal{P}_\ell$ and $\mathcal{P}_u$. Where necessary for clarity, we will make the dependence of the problem $\mathcal{P}$ on the initial state x_1 explicit by denoting it $\mathcal{P}(x_1)$, and its optimal value $\mathcal{P}^*(x_1)$. We adopt a similar notation in relation to the problems $\mathcal{P}_\ell$ and $\mathcal{P}_u$ and their tractable reformulations $\tilde{\mathcal{P}}_\ell$ and $\tilde{\mathcal{P}}_u$. The first result is a natural consequence of Propositions 3.2 and 4.2 and is the central result in the paper:

Theorem 5.1. *The optimal values of the problem* $\mathcal{P}$ *and the tractable problems* $\tilde{\mathcal{P}}_u$ *and* $\tilde{\mathcal{P}}_\ell$ *satisfy*

$$\tilde{\mathcal{P}}_\ell^*(x_1) \leq \mathcal{P}^*(x_1) \leq \tilde{\mathcal{P}}_u^*(x_1) \quad \forall x_1 \in \mathbb{R}^{n_x}.$$

The above result is significant because it allows one to bound the degree of suboptimality incurred when employing affine decision rules to solve approximately the intractable optimization problem $\mathcal{P}(x_1)$.

Further insight is possible by comparing the sets of feasible decision variables in the finite dimensional upper- and lower-bounding problems $\tilde{\mathcal{P}}_u(x_1)$ and $\tilde{\mathcal{P}}_\ell(x_1)$. Define the sets of feasible decision variables for these problems as

$$\Pi_u(x_1) := \left\{ (Q, S) \ : \ \text{problem } \tilde{\mathcal{P}}_u(x_1) \text{ is feasible for some } \mu \geq 0, \ \Lambda_i \succeq_{\mathcal{K}_p^*} 0, \ i \in \{1 \ldots l\} \right\}$$
$$\Pi_\ell(x_1) := \left\{ (Q, S) \ : \ \text{problem } \tilde{\mathcal{P}}_\ell(x_1) \text{ is feasible} \right\},$$

and define the sets of initial states x_1 for which a feasible policy can be found for these problems as

$$\begin{aligned} X_u &:= \{x_1 \in \mathbb{R}^{n_x} \;:\; \Pi_u(x_1) \neq \emptyset\} \\ X_\ell &:= \{x_1 \in \mathbb{R}^{n_x} \;:\; \Pi_\ell(x_1) \neq \emptyset\}. \end{aligned}$$

The next result guarantees that these sets are nested:

Proposition 5.2. *The policy sets $\Pi_u(x_1)$ and $\Pi_\ell(x_1)$ satisfy*

$$\Pi_u(x_1) \subseteq \Pi_\ell(x_1) \quad \forall x_1 \in \mathbb{R}^{n_x},$$

and the initial state sets X_u and X_ℓ satisfy $X_u \subseteq X_\ell$.

Proof. To prove the first part, assume that $(Q, S) \in \Pi_u(x_1)$ is specified, so that $S = \mu e_0^\top + \sum_{i=1}^l \Lambda_i^\top W_i$ for some $\Lambda_i \succeq_{\mathcal{K}_p^*} 0$, $i = 1, \ldots, l$, and $\mu \geq 0$. Then

$$\begin{aligned} W_i M S^\top &= W_i M(e_0 \mu^\top + \textstyle\sum_{i=1}^l W_i^\top \Lambda_i) \\ &= \mathbb{E}\left[W_i \boldsymbol{\xi}\right] \mu^\top + \textstyle\sum_{i=1}^l W_i M W_i^\top \Lambda_i \\ &= \mathbb{E}\left[W_i \boldsymbol{\xi}\right] \mu^\top + \textstyle\sum_{i=1}^l \mathbb{E}\left[(W_i \boldsymbol{\xi})(\Lambda_i^\top W_i \boldsymbol{\xi})^\top\right] \succeq_{\mathcal{K}_p} 0, \end{aligned}$$

since $W_i \xi \succeq_{\mathcal{K}_p} 0$ and $\Lambda_i^\top W_i \xi \geq 0$ $\forall \xi \in \Xi$, $i = 1, \ldots, l$. This ensures that the conic inequality in $\tilde{\mathcal{P}}_\ell(x_1)$ is satisfied. Furthermore,

$$\begin{aligned} e_0^\top M S^\top &= e_0^\top M(e_0 \mu^\top + \textstyle\sum_{i=1}^l W_i^\top \Lambda_i) \\ &= \mu^\top + \textstyle\sum_{i=1}^l e_0^\top M W_i^\top \Lambda_i \\ &= \mu^\top + \textstyle\sum_{i=1}^l \mathbb{E}\left[(\Lambda_i^\top W_i \boldsymbol{\xi})^\top\right] \geq 0. \end{aligned}$$

This ensures that the final inequality constraint in $\tilde{\mathcal{P}}_\ell(x_1)$ is satisfied. Since all other constraints in $\tilde{\mathcal{P}}_\ell(x_1)$ and $\tilde{\mathcal{P}}_\ell(x_1)$ are identical, $(Q, S) \in \Pi_\ell(x_1)$.

The second part of Proposition 5.2 follows immediately from the first part. □

Remark 6. *An important application for solutions to robust finite horizon problems such as $\mathcal{P}$ is in receding horizon control (RHC). If one equips problem $\mathcal{P}$ (alternatively, $\tilde{\mathcal{P}}_u$) with appropriate terminal conditions on the constraints and objective function, then a RHC law synthesized from repeated solutions to $\mathcal{P}$ can be shown to endow the resulting closed-loop system with desirable stability and invariance properties; cf. [22, §4],[17, 26].*

It is difficult in general to assess the degradation of performance (if any) of such a controller if one substitutes receding horizon implementation of solutions to problem $\mathcal{P}$ with solutions to its sub-optimal approximation $\tilde{\mathcal{P}}_u$. In particular, there is no obvious method for directly inferring stability properties (e.g. input-to-state gain) from the value function $\mathcal{P}^*$, which typically plays the role of a Lyapunov function in RHC. On the other hand, if one defines $X := \{x \;:\; \mathcal{P}(x) \text{ is feasible}\}$, then clearly $X_u \subseteq X \subseteq X_\ell$. Given appropriate terminal conditions, the set X (alternatively, X_u) is the region of attraction of such a RHC controller, and the set difference $X_u \backslash X_\ell$ provides an estimate of the conservatism of a RHC synthesized from $\tilde{\mathcal{P}}_u$ with respect to region of attraction.

Finally, we show that the upper and lower bounds calculable via $\tilde{\mathcal{P}}_u$ and $\tilde{\mathcal{P}}_\ell$ coincide for problems with equality constraints only, which demonstrates that our approximation methods are not unnecessarily conservative in this case:

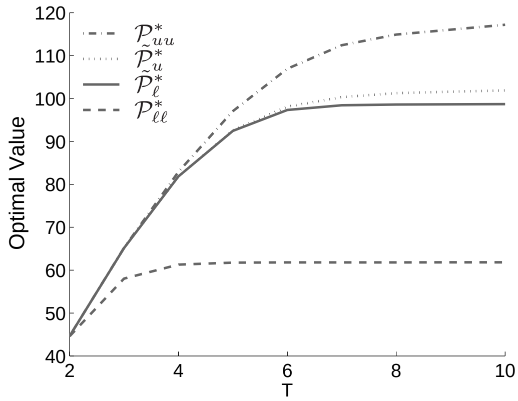

Figure 1: Upper and lower bounds for the perfect state measurements example.

Corollary 5.3. *If $F_s = 0$ then the optimal values achieved in the upper- and lower-bounding problems $\tilde{\mathcal{P}}_u$ and $\tilde{\mathcal{P}}_\ell$ coincide, i.e.*

$$\tilde{\mathcal{P}}_\ell^*(x_1) = \mathcal{P}^*(x_1) = \tilde{\mathcal{P}}_u^*(x_1)\ .$$

Proof. If $F_s = 0$, then both of the problems $\tilde{\mathcal{P}}_\ell$ and $\tilde{\mathcal{P}}_u$ are constrained only by the causality condition $Q \in \mathcal{U}$ and the equality $(F_u + F_x B)QG + F_x C - he_0^\top = 0$. This ensures that $\Pi_u(x_1) = \Pi_l(x_1)$, and consequently $\tilde{\mathcal{P}}_\ell^*(x_1) = \tilde{\mathcal{P}}_u^*(x_1)$. The result then follows from Proposition 3.2, 4.2 and Theorem 5.1. □

Remark 7. *The above result ensures that affine policies computed using both the primal and dual methods of Section 3 and 4 are optimal for control problems of the type $\mathcal{P}$ if $\boldsymbol{u}$ and $\boldsymbol{x}$ are restricted only to a subspace. In the particular case that (J_x, J_u) are block diagonal and $(F_x, F_u, F_s) = 0$, both problems reduce to the standard Linear Quadratic Regulator problem, for which linear feedback policies are well-known to be optimal [1].*

6 Numerical Examples

6.1 Perfect State Measurements

We consider the following time-invariant, discrete-time linear system, which allows for perfect state measurements:

$$\boldsymbol{x}_{t+1} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \boldsymbol{x}_t + \begin{bmatrix} 0.2 \\ 1 \end{bmatrix} \boldsymbol{u}_t + \begin{bmatrix} 0.1 \\ 0.1 \end{bmatrix} \boldsymbol{\xi}_t,\ \boldsymbol{y}_t = \boldsymbol{x}_t,$$

where $t = 1, \ldots, T-1$. The initial state of the system is set to $\boldsymbol{x}_1 = (4.7, 0)^\top$. We assume that the $\boldsymbol{\xi}_t$ are independent and uniformly distributed on $[0, 2]$ for $t = 1, \ldots, T$, while $\boldsymbol{\xi}_0 = 1\,\mathbb{P}$-a.s. Our objective is to minimize $\mathbb{E}\left[\boldsymbol{u}^\top \boldsymbol{u} + \boldsymbol{x}^\top \boldsymbol{x}\right]$ subject to

$$\left.\begin{array}{l} (-5,-5)^\top \le \boldsymbol{x}_t \le (5,5)^\top,\ t = 1,\ldots,T \\ (1,1)\,\boldsymbol{x}_t \le 5,\ t = 1,\ldots,T \\ (1,-1)\,\boldsymbol{x}_t \le 5,\ t = 1,\ldots,T \\ |\boldsymbol{u}_t| \le 1,\ t = 1,\ldots,T-1 \end{array}\right\} \quad \mathbb{P}\text{-a.s.}$$

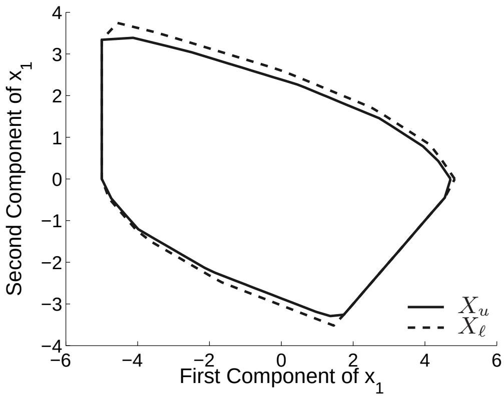

Figure 2: Sets of initial states X_ℓ and X_u for horizon $T = 8$, for the perfect state measurements example.

We solve the approximate problems $\tilde{\mathcal{P}}_\ell$ and $\tilde{\mathcal{P}}_u$ as well as the trivial bound problems $\mathcal{P}_{\ell\ell}$ and $\mathcal{P}_{uu}$ for $T = 2, \ldots, 10$. Figure 1 shows results for the optimal value achieved by solving these problems for the given horizon lengths. For this particular example, the upper and lower bounds $\tilde{\mathcal{P}}_u$ and $\tilde{\mathcal{P}}_\ell$ are close to each other but diverge slowly. The bounds do not coincide suggesting that affine policies may not be optimal, especially as the time horizon increases beyond 4. The trivial bounds $\mathcal{P}_{\ell\ell}$ and $\mathcal{P}_{uu}$ fail to detect the near-optimality of affine policies for time horizons less than 4. Figure 2 shows the two regions X_ℓ and X_u of initial states x_1, for which a feasible policy can be found for $\tilde{\mathcal{P}}_\ell$ and $\tilde{\mathcal{P}}_u$ respectively, for a time horizon of $T = 8$. Observe that $X_u \subseteq X_\ell$.

6.2 Imperfect State Measurements

We consider a similar system but with imperfect state information and measurement errors. The system equations are:

$$\boldsymbol{x}_{t+1} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \boldsymbol{x}_t + \begin{bmatrix} 0.5 \\ 1 \end{bmatrix} \boldsymbol{u}_t + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \boldsymbol{\xi}_t$$

$$\boldsymbol{y}_t = \begin{bmatrix} 1 & 1 \end{bmatrix} \boldsymbol{x}_t + \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} \boldsymbol{\xi}_t$$

for all $t = 1 \ldots T-1$. The initial state is set to $\boldsymbol{x}_1 = (0,1)^\top$ and the support Ξ is defined as

$$\Xi := \left\{ \xi \in \mathbb{R}^{N_\xi} : \ \xi \succeq_{\mathcal{K}_2} 0, \, e_0^\top \xi = 1 \right\}.$$

The process noise $\boldsymbol{\xi}$ is uniformly distributed on Ξ. Our objective is to minimize $\mathbb{E}\left[\boldsymbol{u}^\top \boldsymbol{u} + \boldsymbol{x}^\top \boldsymbol{x}\right]$ subject to

$$\left. \begin{array}{ll} (-5,-5)^\top \leq \boldsymbol{x}_t \leq (5,5)^\top, & t = 1, \ldots, T \\ |\boldsymbol{u}_t| \leq 2, & t = 1, \ldots, T-1 \end{array} \right\} \quad \mathbb{P}\text{-a.s.}$$

We solve the approximate problems $\tilde{\mathcal{P}}_\ell$ and $\tilde{\mathcal{P}}_u$ and the trivial bound problems $\mathcal{P}_{\ell\ell}$ and $\mathcal{P}_{uu}$ for $T = 2, \ldots, 20$. The results are shown in Figure 3. For this example, the bounds $\tilde{\mathcal{P}}_\ell$ and $\tilde{\mathcal{P}}_u$ coincide, meaning there is no optimality gap and thus affine policies are optimal in this instance. This is of course not true in general, but as the proposed bounds are problem-specific, they can attest distinct instances of problems where affine controllers are optimal. Problem $\mathcal{P}_{uu}$ becomes infeasible for $T \geq 5$, rendering the trivial upper bound meaningless for those horizons.

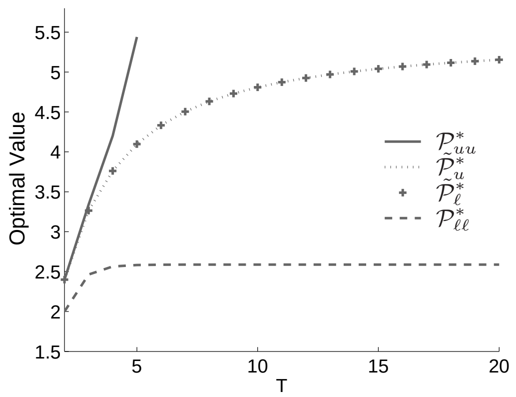

Figure 3: The upper and lower bounds for the imperfect state measurements example. The two bounds $\tilde{\mathcal{P}}_\ell$ and $\tilde{\mathcal{P}}_u$ coincide meaning that for this problem instance affine policies are optimal.

7 Conclusion

We investigate discrete time, time-varying linear systems with input and measurement uncertainty. Finding causal controllers that minimize an expected cost function is computationally intractable in the presence of state constraints. Therefore, we restrict attention to affine controllers over which one can optimize in polynomial time by solving a tractable conic optimization problem. The optimal value of this problem constitutes an upper bound on the minimal cost achievable. Estimating the degree of suboptimality of the best *affine* controller has been a longstanding open problem. By linearizing the dual variables in the original control problem, we devise a relaxed problem whose optimal value provides a lower bound on the minimal cost achievable by any *causal* controller. We argue that this problem is again equivalent to a tractable conic optimization problem. The difference between the two bounds constitutes an a posteriori measure for the degree of suboptimality of the best affine controller. This error estimate depends on the structure of the underlying control problem and can be calculated in polynomial time.

A Appendix

Matrix Definitions

Define $B \in \mathbb{R}^{N_x \times N_u}$, $C \in \mathbb{R}^{N_x \times N_\xi}$, $D \in \mathbb{R}^{N_y \times N_x}$ and $E \in \mathbb{R}^{N_y \times N_\xi}$ as

$$B := \begin{bmatrix} 0 & & & & \\ \mathcal{A}_2^2 B_1 & 0 & & & \\ \mathcal{A}_2^3 B_1 & \mathcal{A}_3^3 B_2 & 0 & & \\ \vdots & & & \ddots & \\ \vdots & & & & 0 \\ \mathcal{A}_2^T B_1 & \mathcal{A}_3^T B_2 & \cdots & \cdots & \mathcal{A}_T^T B_{T-1} \end{bmatrix}, \qquad C := \begin{bmatrix} \mathcal{A}_1^1 x_1 & & & & \\ \mathcal{A}_1^2 x_1 & \mathcal{A}_2^2 C_1 & & & \\ \mathcal{A}_1^3 x_1 & \mathcal{A}_2^3 C_1 & \mathcal{A}_3^3 C_2 & & \\ \vdots & & & \ddots & \\ \mathcal{A}_1^T x_1 & \mathcal{A}_2^T C_1 & \mathcal{A}_3^T C_2 & \cdots & \mathcal{A}_T^T C_{T-1} \end{bmatrix}$$

$$D := \begin{bmatrix} 0 & & & \\ D_1 & 0 & & \\ & \ddots & \ddots & \\ & & D_{T-1} & 0 \end{bmatrix}, \qquad E := \begin{bmatrix} 1 & & & \\ & E_1 & & \\ & & \ddots & \\ & & & E_{T-1} \end{bmatrix}$$

where $\mathcal{A}_s^t := A_{t-1}A_{t-2}\cdots A_s$ for $s < t$, $\mathcal{A}_t^t := I^{(n_x)}$, and x_1 is the given value of the initial state $\boldsymbol{x}_1$.

Proof of Proposition 2.1

The statement is proved by induction. For any causal policy $\boldsymbol{u} \in \mathcal{N}$, $\boldsymbol{u}_t$ must be $\mathcal{F}_t^y$-measurable for each $t = 1, \ldots, T-1$. This is equivalent to asserting the existence of Borel-measurable functions $\varphi_t : \mathbb{R}^{1+n_y t} \to \mathbb{R}^{n_u}$ such that $\boldsymbol{u}_t = \varphi_t(\boldsymbol{y}_0, \ldots, \boldsymbol{y}_t)$, see e.g. [2, Theorem 6.4.2 (c)]. It is clear that $\mathcal{F}_0^y = \mathcal{F}_0$ since $\boldsymbol{y}_0 = \boldsymbol{\eta}_0$, which is 1 almost surely. Assume now that $\mathcal{F}_s^y = \mathcal{F}_s$ for all $0 \leq s < t$ where $0 < t \leq T-1$. Thus, there are Borel-measurable functions ψ_s and χ_s such that $\boldsymbol{y}_s = \psi_s(\boldsymbol{\eta}_0, \ldots, \boldsymbol{\eta}_s)$ and $\boldsymbol{\eta}_s = \chi_s(\boldsymbol{y}_0, \ldots, \boldsymbol{y}_s)$ for all $0 \leq s < t$. Moreover, by the causality of the control policy $\boldsymbol{u}$ we have

$$
\begin{aligned}
\boldsymbol{y}_t &= \sum\nolimits_{s=1}^{t-1} D_t A_{s+1}^t B_s \boldsymbol{u}_s + \sum\nolimits_{s=0}^{t-1} D_t A_{s+1}^t C_s \boldsymbol{\xi}_s + E_t \boldsymbol{\xi}_t \\
&= \sum\nolimits_{s=1}^{t-1} D_t A_{s+1}^t B_s \, \varphi_s(\psi_0(\boldsymbol{\eta}_0), \ldots, \psi_s(\boldsymbol{\eta}_0, \ldots, \boldsymbol{\eta}_s)) + \boldsymbol{\eta}_t \\
&=: \psi_t(\boldsymbol{\eta}_0, \ldots, \boldsymbol{\eta}_t),
\end{aligned} \tag{A.18}
$$

and from (A.18) we infer that

$$
\begin{aligned}
\boldsymbol{\eta}_t &= \boldsymbol{y}_t - \sum\nolimits_{s=1}^{t-1} D_t A_{s+1}^t B_s \, \varphi_s(\boldsymbol{y}_0 \ldots, \boldsymbol{y}_s) \\
&=: \chi_t(\boldsymbol{y}_0, \ldots, \boldsymbol{y}_t).
\end{aligned} \tag{A.19}
$$

The relation (A.18) implies that

$$
\mathcal{F}_t^y = \sigma(\psi_0(\boldsymbol{\eta}_0), \ldots, \psi_t(\boldsymbol{\eta}_0, \ldots, \boldsymbol{\eta}_t)) \subseteq \sigma(\boldsymbol{\eta}_0, \ldots, \boldsymbol{\eta}_t) = \mathcal{F}_t, \tag{A.20a}
$$

and the relation (A.19) implies that

$$
\mathcal{F}_t = \sigma(\chi_0(\boldsymbol{y}_0), \ldots, \chi_t(\boldsymbol{y}_0, \ldots, \boldsymbol{y}_t)) \subseteq \sigma(\boldsymbol{y}_0, \ldots, \boldsymbol{y}_t) = \mathcal{F}_t^y. \tag{A.20b}
$$

The inclusions (A.20a) and (A.20b) yield $\mathcal{F}_t^y = \mathcal{F}_t$, and thus the claim follows. □

Proof of Lemma 3.4

The proposition follows from a strong duality argument. Statement (i) is equivalent to

$$
\begin{aligned}
0 \leq \quad \min \quad & z^\top \xi \\
\text{s.t.} \quad & \xi \in \mathbb{R}^{N_\xi} \\
& W_i \, \xi \succeq_{\mathcal{K}_p} 0, \; i = 1 \ldots l \\
& e_0^\top \xi = 1.
\end{aligned} \tag{A.21}
$$

Due to our assumptions on the support Ξ, the feasible region of the minimization problem (A.21) has nonempty relative interior. By strong conic duality, (A.21) therefore is equivalent to

$$
\begin{aligned}
0 \leq \quad \max \quad & \mu \\
\text{s.t.} \quad & \mu \in \mathbb{R}, \;\; \lambda_i \in \mathbb{R}_\xi^N \\
& z^\top = \textstyle\sum_{i=1}^l \lambda_i^\top W_i + \mu e_0^\top \\
& \lambda_i \succeq_{\mathcal{K}_p^*} 0, \;\; i = 1 \ldots l,
\end{aligned}
$$

which is manifestly equivalent to statement (ii). Thus the claim follows. □

Proof of Lemma 4.4

Proof of the first statement is immediate since M is invertible and we impose no special structure (such as block lower triangularity) on the matrix S.

Proof of the second statement is subdivided into three parts. Before beginning the proof, we introduce a sequence of truncation operators

$$P_t : \mathbb{R}^{N_y} \to \mathbb{R}^{1+tn_y},\ \eta \mapsto \eta^t = (\eta_0, \dots, \eta_t),\ \ t = 1, \dots, T-1. \tag{A.22}$$

It is easy to verify that $P_t\boldsymbol{\eta} = P_t G\boldsymbol{\xi} = G_t\boldsymbol{\xi}$ where $G_t \in \mathbb{R}^{(1+tn_y)\times N_\xi}$ is the sub-matrix of G which is obtained by removing the last $((T-1)-t)n_y$ rows of G.

Part 1. We first show that for every $\boldsymbol{u} \in \mathcal{N}$ there exists a matrix $Q_t \in \mathbb{R}^{n_u\times(1+n_y t)}$ that satisfies

$$Q_t G_t M G_t^\top = \mathbb{E}\left[\boldsymbol{u}_t\ (G_t\boldsymbol{\xi})^\top\right]. \tag{A.23}$$

Since the matrix M is positive definite by assumption, the null space relationship $\ker(G_t M G_t^\top) = \ker([G_t M G_t^\top]^\dagger) = \ker(G_t^\top)$ holds. We set

$$Q_t = E\left[\boldsymbol{u}_t\ (G_t\boldsymbol{\xi})^\top\right](G_t M G_t^\top)^\dagger.$$

Then $Q_t v = 0$ for all $v \in \ker(G_t^\top)$, and $Q_t G_t M G_t^\top v = \mathbb{E}\left[\boldsymbol{u}_t\ (G_t\boldsymbol{\xi})^\top\right] v$ for all v in the orthogonal complement of $\ker(G_t^\top)$, which guarantees (A.23).

Part 2. Next we show that for every $\boldsymbol{u} \in \mathcal{N}$ there exists a $Q \in \mathcal{U}$ so that

$$QGMG = \mathbb{E}\left[\boldsymbol{u}\boldsymbol{\eta}^\top\right]. \tag{A.24}$$

By assumption, there exists a matrix $H_t \in \mathbb{R}^{N_y\times(1+tn_y)}$ so that $\mathbb{E}\left[\boldsymbol{\eta}|\ \boldsymbol{\eta}^t\right] = H_t P_t\,\boldsymbol{\eta} = H_t G_t\,\boldsymbol{\xi}$ $\mathbb{P}$-a.s. Therefore, we find

$$\begin{aligned}
\mathbb{E}\left[\boldsymbol{u}_t\,\boldsymbol{\eta}^\top\right] &= \mathbb{E}\left[\boldsymbol{u}_t\,\mathbb{E}\left[\boldsymbol{\eta}|\ \boldsymbol{\eta}^t\right]^\top\right] && = \mathbb{E}\left[\boldsymbol{u}_t\,(G_t\,\boldsymbol{\xi})^\top\right]\ H_t^\top \\
&= Q_t\,G_t\,M\,G_t^\top\ H_t^\top && = Q_t\,\mathbb{E}\left[G_t\,\boldsymbol{\xi}\,(H_t\,G_t\,\boldsymbol{\xi})^\top\right] \\
&= Q_t\,\mathbb{E}\left[\boldsymbol{\eta}^t\,\mathbb{E}\left[\boldsymbol{\eta}|\ \boldsymbol{\eta}^t\right]^\top\right] && = Q_t\,\mathbb{E}\left[\boldsymbol{\eta}^t\,\boldsymbol{\eta}^\top\right] \\
&= Q_t\,G_t\,\mathbb{E}\left[\boldsymbol{\xi}\,\boldsymbol{\xi}^\top\right]G^\top && = Q_t\,P_t\,G\,M\,G^\top,
\end{aligned} \tag{A.25}$$

where the first line follows from the law of iterated conditional expectations and the $\mathcal{F}_t$ -measurability of $\boldsymbol{u}_t$. The second line holds due to (A.23) and the definition of M, and the third line follows from the definition of H_t and the law of iterated conditional expectations. The last line follows from the definition of G_t and M. Finally, we define

$$Q := \begin{pmatrix} Q_0 P_0 \\ Q_1 P_1 \\ \vdots \\ Q_{T-1} P_{T-1} \end{pmatrix} \in \mathcal{U}.$$

By construction, Q satisfies (A.24).

Part 3. To conclude the proof, we use the relation (A.24) in conjunction with assumption **A1** that $\mathbb{E}[\boldsymbol{\xi}|\boldsymbol{\eta}] = L\boldsymbol{\eta}$ $\mathbb{P}$-a.s. We have

$$\begin{aligned}\mathbb{E}\left[\boldsymbol{u}\,\boldsymbol{\xi}^\top\right] &= \mathbb{E}\left[\boldsymbol{u}\,\mathbb{E}\left(\boldsymbol{\xi}|\boldsymbol{\eta}\right)^\top\right]\\ &= \mathbb{E}\left[\boldsymbol{u}\,\boldsymbol{\eta}^\top\right]L^\top \\ &= QGMG^\top L^\top = QGM,\end{aligned} \tag{A.26}$$

where that last equality follows from

$$\begin{aligned}GMG^\top L^\top &= G\mathbb{E}\left[\boldsymbol{\xi}\,\boldsymbol{\xi}^\top\right]G^\top L^\top = \mathbb{E}\left[\boldsymbol{\eta}\,\boldsymbol{\eta}^\top\right]L^\top\\ &= \mathbb{E}\left[\boldsymbol{\eta}\,\mathbb{E}\left(\boldsymbol{\xi}|\boldsymbol{\eta}\right)^\top\right] = \mathbb{E}\left[\boldsymbol{\eta}\,\boldsymbol{\xi}^\top\right]\\ &= GM.\end{aligned} \tag{A.27}$$

This observation completes the proof of the second statement. □

Proof of Lemma 4.5

We first consider the relaxed problem

$$\min_{\boldsymbol{u}\in\mathcal{L}^2_{N_u}} \left\{\mathbb{E}\left[\boldsymbol{u}^\top J\boldsymbol{u}\right] \,:\, QGM = \mathbb{E}\left[\boldsymbol{u}\,\boldsymbol{\xi}^\top\right]\right\}, \tag{A.28}$$

where $\boldsymbol{u}$ is not restricted to be in $\mathcal{N}$. We will show that this problem has a solution in $\mathcal{N}$, which is therefore also a solution for (4.15). Consider the Lagrangian of the relaxed problem (A.28),

$$L(\boldsymbol{u},\Lambda) = \mathbb{E}\left[\boldsymbol{u}^\top J\boldsymbol{u}\right] + \mathrm{tr}\left(\Lambda^\top\left(QGM - \mathbb{E}\left[\boldsymbol{u}\,\boldsymbol{\xi}^\top\right]\right)\right), \tag{A.29}$$

where $\Lambda \in \mathbb{R}^{N_u\times N_\xi}$ is the matrix of Lagrange multipliers associated with the equality constraints. A set of first-order optimality conditions is found by setting the Gâteaux differential of the Lagrangian with respect to $\boldsymbol{u}$ to zero for all descent directions $\boldsymbol{h} \in L^2_{N_u}$.

$$\begin{aligned}\delta L(\boldsymbol{u},\Lambda;\boldsymbol{h}) = \quad & \mathbb{E}\left[\boldsymbol{h}^\top\left(2\,J\,\boldsymbol{u} - \Lambda\,\boldsymbol{\xi}\right)\right] = 0 \quad \forall \boldsymbol{h}\in L^2_{N_u}\\ \iff \quad & 2J\boldsymbol{u} - \Lambda\boldsymbol{\xi} = 0 \quad \mathbb{P}\text{-a.s.}\end{aligned} \tag{A.30}$$

It can be verified that $\Lambda^* = 2JQG$ and $\boldsymbol{u}^* = QG\boldsymbol{\xi} = Q\boldsymbol{\eta}$ satisfy both the optimality conditions (A.30) and the constraints of problem (A.28). Thus $\boldsymbol{u}^*$ constitutes a valid but not necessarily unique solution of the convex problem (A.28). Furthermore, since Q is an element of $\mathcal{U}$, the affine policy $\boldsymbol{u}^*$ is non-anticipative, that is $\boldsymbol{u}^* \in \mathcal{N}$. As a result, $\boldsymbol{u}^*$ is also an optimal solution for problem (4.15), which is more restrictive than (A.28). Thus the claim follows. □

Proof of Lemma 4.6

First, observe that $\emptyset \neq \mathrm{int}\,\mathcal{C}$. This follows from the assumption that Ξ has a non-empty relative interior and spans $\mathbb{R}^{N_\xi}$. In the remainder of the proof we show that $\mathrm{int}\,\mathcal{C} \subseteq \mathcal{C}_{\boldsymbol{\xi}} \subseteq \mathcal{C}$. We will show this using the methodologies used in [18, Proposition 3].

Let $\mathcal{M}^+$ be the set of all nonnegative finite measures on $(\Xi, \mathcal{B}(\Xi))$ with finite second moments, and let $\mathcal{M}^+_{\boldsymbol{\xi}}$ be the subset of all measures in $\mathcal{M}^+$ that have a square-integrable density with respect to the distribution $\mathbb{P}_{\boldsymbol{\xi}}$ of $\boldsymbol{\xi}$. Define two convex cones in $\mathbb{R}^{N_\xi}$ as

$$\tilde{\mathcal{C}} := \left\{\int_\Xi \xi\,\mu(\mathrm{d}\xi) : \mu \in \mathcal{M}^+\right\} \quad \text{and} \quad \tilde{\mathcal{C}}_{\boldsymbol{\xi}} := \left\{\int_\Xi \xi\,\mu(\mathrm{d}\xi) : \mu \in \mathcal{M}^+_{\boldsymbol{\xi}}\right\}.$$

By construction $\tilde{\mathcal{C}}_{\boldsymbol{\xi}} \subseteq \tilde{\mathcal{C}}$. Since the density of any $\mu \in \mathcal{M}^+_{\boldsymbol{\xi}}$ with respect to $\mathbb{P}_{\boldsymbol{\xi}}$ can be identified with a function in $\mathcal{L}^2_1$, it is clear that $\tilde{\mathcal{C}}_{\boldsymbol{\xi}} = \mathcal{C}_{\boldsymbol{\xi}}$. Furthermore, as $\mathcal{M}^+_{\boldsymbol{\xi}}$ is weakly dense in $\mathcal{M}^+$ (since Ξ constitutes the support of $\mathbb{P}_{\boldsymbol{\xi}}$), and the identity mapping $\xi \mapsto \xi$ is continuous, it follows that $\tilde{\mathcal{C}}_{\boldsymbol{\xi}}$ is dense in $\tilde{\mathcal{C}}$. Keeping in mind that $\tilde{\mathcal{C}}_{\boldsymbol{\xi}}$ is also convex, the above findings imply

$$\operatorname{int} \tilde{\mathcal{C}} \subseteq \tilde{\mathcal{C}}_{\boldsymbol{\xi}} = \mathcal{C}_{\boldsymbol{\xi}} \subseteq \tilde{\mathcal{C}}\,. \tag{A.31}$$

We now proceed to show that $\tilde{\mathcal{C}} = \mathcal{C}$, which will complete the proof. For any $z \in \mathcal{C}$ there exists a $z' \in \Xi$ and a scaling factor $\lambda \geq 0$ such that $z = \lambda z'$. Next, we define $\mu_z := \lambda \delta_{z'}$ where $\delta_{z'}$ is the Dirac measure which concentrates unit mass at the point z'. It is easy to verify that $\mu_z \in \mathcal{M}^+$. We then have

$$z = \int_\Xi \xi \, \mu_z(\mathrm{d}\xi) \in \tilde{\mathcal{C}},$$

yielding the relation $\mathcal{C} \subseteq \tilde{\mathcal{C}}$. In order to derive the opposite relation, we select any $z \in \tilde{\mathcal{C}}$. By the definition of $\tilde{\mathcal{C}}$ there exists a $\mu \in \mathcal{M}^+$ so that $z = \int_\Xi \xi \, \mu(\mathrm{d}\xi)$. We set $\lambda := \mu(\Xi)$. If $\lambda = 0$, then $z = 0 \in \mathcal{C}$, since μ is a non-negative measure. From now on we thus assume that $\lambda > 0$ and set $z' := \frac{1}{\lambda} \int_\Xi \xi \, \mu(\mathrm{d}\xi)$. Observe that μ/λ constitutes a probability measure. Then, it is easy to verify that $z' \in \Xi$, since z' is the mean value of a probability distribution supported on the convex set Ξ. Since $\lambda > 0$, we have $z = \lambda z' \in \mathcal{C}$. As the choice of $z \in \tilde{\mathcal{C}}$ was arbitrary, we have $\tilde{\mathcal{C}} \subseteq \mathcal{C}$, and thus we conclude that $\tilde{\mathcal{C}} = \mathcal{C}$. Substituting this result into (A.31) completes the proof. □

References

[1] B. D. O. Anderson and J. B. Moore. *Optimal control: linear quadratic methods.* Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1990.

[2] R.B Ash. *Real analysis and probability.* Academic Press New York, 1972.

[3] A. Ben-Tal, S. Boyd, and A. Nemirovski. Control of uncertainty affected discrete time linear systems via convex programming. Technical report, Minerva Optimization Center,Technion, Haifa, Israel, 2005.

[4] A. Ben-Tal, S. Boyd, and A. Nemirovski. Extending scope of robust optimization: Comprehensive robust counterparts of uncertain problems. *Mathematical Programming*, 107(1–2):63–89, June 2006.

[5] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization.* Princeton Series in Applied Mathematics. Princeton University Press, October 2009.

[6] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

[7] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23:769–805, 1998.

[8] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operation Research Letters*, 25(1):1–13, 1999.

[9] D. Bertsekas and I. Rhodes. On the minimax reachability of target sets and target tubes. *Automatica*, 7(2):233–247, 1971.

[10] D. Bertsimas, D. Iancu, and P. Parrilo. Optimality of affine policies in multi-stage robust optimization. *Mathematics Of Operations Research*, 35(2):363–394, 2010.

[11] S. Cambanis, S. Huang, and G. Simons. On the theory of elliptically contoured distributions. *Journal of Multivariate Analysis*, 11(3):368 – 385, 1981.

[12] M. Chmielewski. Elliptically symmetric distributions: A review and bibliography. *International Statistical Review*, 49(1):67–74, 1981.

[13] A. Georghiou, W. Wiesemann, and D. Kuhn. Generalized decision rule approximations for stochastic programming via liftings. 2010. Available on Optimization Online, submitted for publication (2010).

[14] J. Goh and M. Sim. Distributionally robust optimization and its tractable approximations. *Operations Research*, 58(4-Part-1):902–917, 2010.

[15] P. Goulart and E. Kerrigan. Output feedback receding horizon control of constrained systems. *International Journal of Control*, 80(1):8–20, January 2007.

[16] P. Goulart, E. Kerrigan, and J. Maciejowski. Optimization over state feedback policies for robust control with constraints. *Automatica*, 42(4):523 – 533, 2006.

[17] P. J. Goulart and E.C. Kerrigan. Input-to-state stability of robust receding horizon control with an expected value cost. *Automatica*, 44(4):1171–1174, April 2008.

[18] D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, pages 1–33, 2009. 10.1007/s10107-009-0331-4.

[19] Y. Lee and B. Kouvaritakis. Robust receding horizon predictive control for systems with uncertain dynamics and input saturation. *Automatica*, 36(10):1497–1504, 2000.

[20] J. Löfberg. *Minimax Approaches to Robust Model Predictive Control*. PhD thesis, Linköping University, April 2003.

[21] D. Mayne, M. Seron, and S. Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–24, February 2005.

[22] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, June 2000. Survey paper.

[23] P. Scokaert and D. Mayne. Min-max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic Control*, 43(8):1136–1142, August 1998.

[24] J. Skaf and S. Boyd. Design of affine controllers via convex optimization. *IEEE Transactions on Automatic Control*, 2008. in press.

[25] D. van Hessem and O. Bosgra. A conic reformulation of model predictive control including bounded and stochastic disturbances under state and input constraints. In *Proc. 41st IEEE Conference on Decision and Control*, pages 4643–4648, December 2002.

[26] C. Wang, C.-J. Ong, and M. Sim. Model predictive control using segregated disturbance feedback. *IEEE Transactions on Automatic Control*, 55(4):831 –840, apr. 2010.

[27] G. Xue and Y. Ye. An efficient algorithm for minimizing a sum of p-norms. *SIAM Journal on Optimization*, 10(2):551–579, 1999.

[28] D. Youla, H. Jabr, and J. Bongiorno. Modern Wiener-Hopf design of optimal controllers: Part II. *IEEE Transactions on Automatic Control*, AC-21:319–338, 1976.