Performance Bounds for Constrained Linear Min-Max Control

Tyler H. Summers* and Paul J. Goulart*

Abstract— This paper proposes a method to compute lower performance bounds for discrete-time infinite-horizon min-max control problems with input constraints and bounded disturbances. Such bounds can be used as a performance metric for control policies synthesized via suboptimal design techniques. Our approach is motivated by recent work on performance bounds for stochastic constrained optimal control problems using relaxations of the Bellman equation. The central idea of the paper is to find an unconstrained min-max control problem, with negatively weighted disturbances as in $\mathcal{H}_\infty$ control, that provides the tightest possible lower performance bound on the original problem of interest and whose value function is easily computed. The new method is demonstrated via a numerical example for a system with box constrained input.

I. INTRODUCTION

Consider the discrete-time linear time-invariant system

$$
\begin{aligned}
x_{t+1} &= Ax_t + Bu_t + Gw_t \\
z_t &= Cx_t + Du_t
\end{aligned}
\tag{1}
$$

with state $x_t \in \mathbf{R}^n$, input $u_t \in \mathbf{R}^m$, unknown exogenous disturbance $w \in \mathbf{R}^l$, and costed/controlled output $z_t \in \mathbf{R}^p$ at each time $t = \{0, 1, \dots\}$. We assume throughout that the initial state x_0 is known and that the state is directly observable at each time step. We further assume that the inputs are subject to a compact constraint $u_t \in \mathbb{U}$, that the disturbances are constrained to a compact set $w_t \in \mathbb{W}$, that the pairs (A, B) and (A, C) are controllable and observable, respectively, and that D is full rank. We denote by $\mathcal{W}$ the set of all infinite-horizon disturbance sequences generated by $\mathbb{W}$.

We define a stationary control policy $\pi : \mathbf{R}^n \to \mathbb{U}$ for the system (1) as a decision rule mapping observed states to control inputs, i.e. $u_t = \pi(x_t)$. We denote the set of all such stationary policies satisfying the system constraints for all infinite-horizon disturbance sequences by Π. The linear min-max control problem is to design a control policy $\pi \in \Pi$ that minimizes the worse-case value of some objective function over all admissible disturbance sequences, starting from some known initial state.

Given an initial state $x := x_0$, a control policy π and a disturbance sequence $\mathbf{w} \in \mathcal{W}$, we consider a cost function in the form

$$
J(x_0, \pi, \mathbf{w}) = \sum_{t=0}^{\infty} \alpha^t \ell(z_t, w_t)
\tag{2}
$$

where $\ell : \mathbf{R}^n \times \mathbb{U} \times \mathbb{W} \to \mathbf{R}$ is a cost function and $\alpha \in (0, 1)$ is a discount factor. The optimal value function and policy are given by

$$
\begin{aligned}
V(x_0) &= \min_{\pi \in \Pi} \max_{w \in \mathcal{W}} \sum_{t=0}^{\infty} \alpha^t \ell(z_t, w_t) \\
\pi^*(x) &= \arg\min_{\pi \in \Pi} \max_{w \in \mathcal{W}} \sum_{t=0}^{\infty} \alpha^t \ell(z_t, w_t)
\end{aligned}
\tag{3}
$$

It is well known that, given the preceding assumptions and certain mild conditions on the stage cost ℓ, a stationary optimal control policy π^* exists [1, Ch. 5]. However, it is extremely difficult to compute the optimal cost and policy in general, because the optimization problem to be solved is over an infinite dimensional function space and over the worst case disturbance sequence.

It was shown in [2] that for constrained finite-horizon problems with a cost convex and quadratic in z and concave and quadratic in w, the optimal value function is piecewise quadratic and the optimal control policy piecewise affine over a polytopic partition of the state space, though computing and storing the policy is possible only for very small systems. Recent research has focused on developing tractable but suboptimal policies for systems in which the approach of [2] is intractable. Examples include various formulations of min-max model predictive control [3], [4], [5], [6]. This immediately raises questions about the degree of suboptimality of such policies. To the best of our knowledge there is currently no general method to answer these questions for min-max problems.

We therefore propose a method for computing a lower bound on the optimal cost for input constrained problems, which provides a way to bound the degree of suboptimality of any control policy constructed by any given design method. Our approach is motivated by recent work from Wang and Boyd [7], who propose a similar approach to computing performance bounds for stochastic optimal control problems (i.e. those with an expected value, rather than a min-max, cost). The central idea is to find an unconstrained linear min-max control problem for which the optimal cost can be efficiently computed that best approximates from below the optimal cost of the original constrained problem.

In particular, we construct an auxiliary convex-concave quadratic cost function for the system (1) without constraints, and exploit this auxiliary function to compute a lower bound on the value function V of the constrained system. The particular choice of a convex-concave auxiliary cost function is motivated by the fact that such cost functions appear

naturally in $\mathcal{H}_\infty$ control design problems, for which exact performance value functions can be computed. Our resulting bounds are not generic, but are instead computed individually for each problem instance. Such bounds can be useful in evaluating the performance of suboptimal policies in cases where the optimal policy cannot be efficiently computed: when the gap between the lower bound and the performance of a suboptimal policy is small, one can conclude that the suboptimal policy is nearly optimal.

The paper is organized as follows. Section II briefly reviews solution methods for unconstrained min-max problems, which are closely related to the $\mathcal{H}_\infty$ control problem. Section III describes the method for computing the proposed bound. Section IV provides a numerical example for which the bound is computed. Finally, Section V gives concluding remarks.

II. Unconstrained Problems with Quadratic Cost

This section reviews solution methods for unconstrained min-max problems whose objective functions are convex-concave in (z, w), which will be used to compute lower bounds for input constrained problems in Section III. If the stage cost for some unconstrained problem is less than or equal to the stage cost for the constrained problem, then we will show that the unconstrained optimal cost is a lower bound on the constrained optimal cost for a certain set of initial states. One can then optimize the parameters in the stage cost of the unconstrained problem to find tighter lower bounds for the original constrained problem.

We begin by recalling that a special case, closely related to the $\mathcal{H}_\infty$ optimal control problem, in which the min-max optimal policy can be efficiently computed is when the stage cost is convex and quadratic in (x, u) and concave and quadratic in w, with discount factor $\alpha = 1$, i.e.

$$\bar{J} = \sum_{t=0}^{\infty} \|Cx_t + Du_t\|^2 - \gamma^2\|w_t\|^2, \quad \gamma > 0 \quad (4)$$

and there are no constraints, i.e. when $\mathbb{U} = \mathbf{R}^m$ and $\mathbb{W} = \mathbf{R}^l$. In this case, the optimal value function is quadratic in the initial state, the optimal state feedback policy is linear, and the closed-loop system has ℓ_2 gain bounded by γ, whenever γ is larger than the $\mathcal{H}_\infty$ optimal gain. Furthermore, the optimal cost and feedback gain can be computed efficiently from the problem data using a variety of well-known techniques, see e.g. [8], [9], [10].

One standard solution method for this problem is based on dynamic programming recursion using the Isaacs equation, which is the discrete-time dynamic game counterpart to the Bellman equation from dynamic programming [10]. Suppose that $(P, \bar{P}) \in \mathbf{R}^{n \times n}$ are symmetric matrices that satisfy the generalized algebraic Riccati equation

$$P = Q + A^T\bar{P}A - A^T\bar{P}B(R + B^T\bar{P}B)^{-1}B^T\bar{P}A \quad (5a)$$
$$\bar{P} := P + PG(\gamma^2 I - G^TPG)^{-1}G^TP, \quad (5b)$$

where we have assumed for simplicity that $Q = C^TC > 0$, $R = D^TD > 0$, $C^TD = 0$ and $D^TC = 0$. If $P \succ 0$ is a solution to (5), then $V(x_0) = x_0^TPx_0$ is the optimal value function, and the optimal controller and disturbance policies are $u = Kx$ and $w = K_wx$, respectively, where

$$\begin{aligned} K &= -(R + B^T\bar{P}B)^{-1}B^T\bar{P}A \\ K_w &= (\gamma^2 I - G^TPG)^{-1}G^TP(A + BK). \end{aligned} \quad (6)$$

For $\alpha \in (0, 1)$ the optimal cost and policies have the same structure but with some problem data scaled by the discount factor: $A \to \sqrt{\alpha}A$, $B \to \sqrt{\alpha}B$ and $\gamma \to \gamma/\sqrt{\alpha}$ (see the Appendix). In other words, the optimal cost and controllers for the discounted problem can be computed by solving an associated undiscounted problem.

Alternative solution methods involve relaxing the Riccati equation (5a) to an inequality and solving a semidefinite programming problem. If (5a) is relaxed to the right (i.e by replacing $=$ with $\succeq$), it can be shown that the resulting inequality is equivalent to the linear matrix inequality obtained by applying the Bounded Real Lemma to the system (1) in close-loop with the optimal controllers $u_t = Kx_t$ and $w_t = K_wx_t$ [9], [8] (see Appendix). The optimal cost and state feedback gain matrices are given by the solution to the following optimization problem in variables P and K:

$$\begin{aligned} \text{minimize} \quad & \mathbf{Tr}(P) \\ \text{subject to} \quad & P \succ 0 \\ & \begin{bmatrix} -P^{-1} & A + BK & G & 0 \\ (A + BK)^T & -P & 0 & (C + DK)^T \\ G^T & 0 & -\gamma I & 0 \\ 0 & C + DK & 0 & -\gamma I \end{bmatrix} \prec 0. \end{aligned} \quad (7)$$

The usual approach to solving (7) is to transform this problem into a semidefinite program in variables P^{-1} and KP^{-1}. However, in the next section we will allow D to vary in order to obtain the tightest possible lower bound for the problem (3), so an alternative to (7) is preferred. As shown in [9], the optimal cost can also be computed separately from the optimal state feedback via the semidefinite program $(\mathcal{SDP})$ in variables P and X. The matrices N and M in $(\mathcal{SDP})$ are bases for the null space of B^T and its orthogonal complement, respectively. Note that in contrast to (7), the problem $(\mathcal{SDP})$ is affine in R^{-1}. This property will be useful when computing lower bounds for the original problem (3).

Remark 1: The problem $(\mathcal{SDP})$ is in contrast with the computation of performance bounds for constrained linear stochastic control problems in [7], in which bounds were obtained by relaxing the standard linear quadratic regulator Riccati equation in the other direction and maximizing over the resulting linear matrix inequality. If (5a) is relaxed to the left (i.e. by replacing $=$ with $\preceq$), one obtains the following

$$\text{minimize} \quad \mathbf{Tr}(P)$$
$$\text{subject to} \quad P \succ 0, \quad X \succ 0$$
$$\begin{bmatrix} N^T(AXA^T - X) & N^TAX & N^T(AXA^T - X)M & N^TG \\ XA^TN & \gamma Q^{-1} & XA^TM & 0 \\ M^T(AXA^T - X)N & M^TAX & M^T(AXA^T - X - \gamma BR^{-1}B^T)M & M^TG \\ G^TN & 0 & G^TM & -\gamma I \end{bmatrix} \prec 0 \qquad (\mathcal{SDP})$$
$$-\gamma I + G^TPG \succ 0, \quad \begin{bmatrix} X & I \\ I & P \end{bmatrix} \succeq 0.$$

optimization problem in variables P and $\bar{P}$

$$\text{maximize} \quad \mathbf{Tr}(P)$$
$$\text{subject to} \quad \begin{bmatrix} R + B^T\bar{P}B & B^T\bar{P}A \\ A^T\bar{P}B & Q + A^T\bar{P}A - P \end{bmatrix} \succeq 0 \qquad (8)$$
$$\bar{P} = P + PG(\gamma^2 I - G^TPG)^{-1}G^TP$$
$$P \succeq 0.$$

Unfortunately, this problem is not convex. For stochastic linear quadratic regulator problems considered in [7] (recovered with $\gamma \to \infty$ and thus $\bar{P} \to P$), the situation is easier since the problem becomes a semidefinite program in P and is also concave in (Q, R).

III. PERFORMANCE BOUND FOR PROBLEMS WITH INPUT CONSTRAINTS

This section describes a method to obtain lower bounds on performance for min-max problems with input constraints. The central idea is to find an unconstrained problem whose optimal cost, which as shown in the previous section can be efficiently computed, is a lower bound on the optimal cost of the constrained problem. To this end, we associate $Q \succ 0$, $R \succ 0$, $\gamma > 0$ with the stage cost of some unconstrained problem and denote by P, K and K_w the corresponding optimal cost and gain matrices given by (5)–(6). The closed-loop system for this unconstrained problem is $x^+ = A_{cl}x$ where $A_{cl} = A + BK + GK_w$. We also define the set

$$X = \{x \in \mathbf{R}^n | K_w(A + BK + GK_w)^k x \in \mathbb{W}, \ \forall k \in \mathbb{N}\}, \quad (9)$$

which corresponds to the set of states for which the optimal unconstrained disturbance sequence always remains in $\mathbb{W}$. We have the following result, which gives a basic lower performance bound for initial states in X:

Lemma 1: Suppose Q, R, γ and $s \in \mathbf{R}$ satisfy

$$x^TQx + u^TRu - \gamma^2 w^Tw + s \leq \ell(z, w),$$
$$\forall x \in \mathbf{R}^n, \ \forall u \in \mathbb{U}, \ \forall w \in \mathbb{W}. \qquad (10)$$

Then we have for $\alpha \in (0, 1)$

$$x_0^TPx_0 + \frac{s}{1-\alpha} \leq V(x_0), \quad \forall x_0 \in X \qquad (11)$$

where P is the optimal cost matrix for the unconstrained min-max problem associated with Q, R, and γ.

Proof: Observe that (10) implies

$$\min_{\pi \in \Pi} \max_{w \in \mathcal{W}} \sum_{t=0}^{\infty} \alpha^t[x_t^TQx_t + u_t^TRu_t - \gamma^2 w_t^Tw_t + s]$$
$$\leq \min_{\pi \in \Pi} \max_{w \in \mathcal{W}} \sum_{t=0}^{\infty} \alpha^t l(z, w). \qquad (12)$$

The inequality (12) remains valid if the input constraint in the set of admissible policies on the outer left-hand-side minimization is dropped. If a given x_0 is in the set X then the constraint in the inner left-hand-side maximization over disturbance sequences can be dropped without changing the value because the constraint is never active by the definition of X. The term involving s in the left-hand-side of (12) can be taken out of the optimizations and evaluated to yield the second term on the left-hand-size of (11). The rest of the left-hand side is then the unconstrained optimal min-max cost and the right-hand-side is the optimal cost of the constrained problem, and the result follows immediately. ∎

A. Optimizing the bound

Lemma 1 gives a basic but efficiently computable lower bound on the value function V for our constrained minimax problem. We next consider the problem of maximizing this lower bound by allowing the parameter R^{-1} in $(\mathcal{SDP})$ to vary[1]. Let $J^*(R^{-1})$ denote the optimal value of $(\mathcal{SDP})$. To optimize the lower bound, we can solve the following optimization problem over the parameters R^{-1} and s

$$\text{maximize} \quad J^*(R^{-1}) + \frac{s}{1-\alpha} \qquad (13)$$
$$\text{subject to} \quad (10), \quad R^{-1} \succ 0.$$

This problem presents two difficulties. First, since $(\mathcal{SDP})$ is a minimization problem and (13) is a maximization problem, the optimization cannot be done jointly over P, X, R^{-1} and s. We can however form the dual of $(\mathcal{SDP})$ and maximize jointly over R^{-1}, s and the dual variables. This results in the following bilinear problem with variables R^{-1},

---

[1]One could also contemplate allowing Q^{-1} to vary. Here, since we do not consider state constraints, (10) must hold on all of $\mathbf{R}^n$, and one does not gain anything by allowing Q^{-1} to vary. Therefore, as in [7], we fix Q to the cost for the constrained problem.

s and dual variables Z, Φ and Λ:

$$\max \quad -\gamma \mathbf{Tr}(Z) + \mathbf{Tr}(F(R^{-1})\Lambda) + 2\mathbf{Tr}(\Phi_{12}) + \frac{s}{1-\alpha}$$

subject to $\quad \Phi_{22} + I + G^T Z G = 0, \quad (10), \quad R^{-1} \succ 0$

$$\Phi_{11} + A^T N \Lambda_{11} N^T A - N \Lambda_{11} N^T + 2 A^T N \Lambda_{12}$$
$$+ 2 A^T N \Lambda_{13} M^T A - 2 N \Lambda_{13} M^T + \Lambda_{22}$$
$$+ 2\Lambda_{23} M^T A + A^T M \Lambda_{33} M^T A$$
$$- M \Lambda_{33} M^T = 0$$

$$Z \succeq 0, \quad \Phi = \begin{bmatrix} \Phi_{11} & \Phi_{12} \\ \Phi_{12}^T & \Phi_{22} \end{bmatrix} \succeq 0$$

$$\Lambda = \begin{bmatrix} \Lambda_{11} & \Lambda_{12} & \Lambda_{13} & \Lambda_{14} \\ \Lambda_{12}^T & \Lambda_{22} & \Lambda_{23} & \Lambda_{24} \\ \Lambda_{13}^T & \Lambda_{23}^T & \Lambda_{33} & \Lambda_{34} \\ \Lambda_{14}^T & \Lambda_{24}^T & \Lambda_{34}^T & \Lambda_{44} \end{bmatrix} \succeq 0.$$

(14)

where $F(R^{-1}) :=$

$$\begin{bmatrix} 0 & 0 & 0 & N^T G \\ 0 & -\gamma Q^{-1} & 0 & 0 \\ 0 & 0 & -\gamma M^T B R^{-1} B^T M & M^T G \\ G^T N & 0 & G^T M & -\gamma I \end{bmatrix}.$$

(15)

Observe that there is a product term involving Λ_{33} and R^{-1} in the dual objective, so this problem is bilinear, and consequently non-convex. However, local and global methods for solving such bilinear optimization problems (e.g. PENBMI [11]) are available. In the next section we adopt the simplest approach, of alternately solving semidefinite programs by fixing one variable in the product and optimizing over the other. We show via numerical example that this method can be used to produce improved lower bounds.

The second difficulty is that (10) is a semi-infinite constraint in general since it must hold for all points in $\mathbf{R}^n$, $\mathbb{U}$ and $\mathbb{W}$. As shown in [7], we recall in the following subsections that when $\ell(z, w)$ is quadratic the constraint (10) can be enforced exactly when $\mathbb{U}$ is a finite set, and that it can be replaced with a conservative approximation via the S-procedure in other cases.

Remark 2: The reason why the dual of $(\mathcal{SDP})$ is preferred to the dual of (7) is that the latter has primal variables multiplying D which we allow to vary in the dual problem. As a consequence, the associated dual problem has bilinear *equality* constraints. Since R^{-1} enters as a constant term in $(\mathcal{SDP})$, the dual problem has the bilinearity in the objective, which we find more convenient to work with.

B. Verifying valid initial states

The bound (11) is valid only for initial states $x_0 \in X$, which is characterized by an infinite set of constraints in (9). We now describe a procedure to verify that a given x_0 is indeed contained in X. Suppose we have a symmetric matrix S and scalar $\beta \in \mathbf{R}$ such that $\bar{\mathbb{W}} = \{w \in \mathcal{R}^l | w^T S w \le \beta\} \subseteq \mathbb{W}$. Suppose also that $x^T H x$ is a dissipated quantity for the optimal unconstrained closed-loop system $x^+ = A_{cl}x$, where $A_{cl} = A + BK + GK_w$ (i.e. H satisfies $A_{cl}^T H A_{cl} - H \preceq 0$).

To show $x_0 \in X$, it is sufficient to show

$$x^T H x \le x_0 H x_0 \Rightarrow x^T K_w^T S K_w x \le \beta.$$

By the S-procedure (see e.g. [12]) this is equivalent to the existence of $\lambda \ge 0$ satisfying

$$K_w^T S K_w \preceq \lambda H, \quad -\beta \le -\lambda x_0^T H x_0.$$

Thus, to verify $x_0 \in X$ it is sufficient to show that there exists H and $1/\lambda$ satisfying the LMI conditions

$$(1/\lambda)K_w^T S K_w \preceq H, \quad -\beta/\lambda \le -x_0^T H x_0,$$
$$A_{cl}^T H A_{cl} - H \preceq 0, \quad 1/\lambda > 0, \quad H \succeq 0.$$

(16)

If a feasible point is found, the bound is valid for all initial states in the set $\{x \in \mathbf{R}^n | x^T H x \le x_0^T H x_0\}$. For a particular initial state, this procedure may be somewhat conservative. The conservatism can be reduced by evaluating a finite number of the constraints in X and then solving the feasibility problem from the last state evaluated.

C. Finite input sets

Suppose that $\ell(z, w)$ is a quadratic function in the form

$$\ell(z, w) = x^T Q_0 x + u^T R_0 u - \gamma_0^2 w^T w$$

(with $\gamma_0 = 0$ possible) and the input constraint set is finite ($\mathbb{U} = \{u_1, ..., u_k\}$). If we fix $Q = Q_0$ and $\gamma \ge \gamma_0$, then (10) reduces to

$$u_i^T R u_i + \frac{s}{1-\alpha} \le u_i^T R_0 u_i, \quad i = 1, ..., k, \quad (17)$$

which can be expressed using Schur complements as LMIs in R^{-1} and s:

$$\begin{bmatrix} \left(u_i^T R_0 u_i - \frac{s}{1-\alpha} \right) & u_i^T \\ u_i & R^{-1} \end{bmatrix} \succeq 0, \quad i = 1, ..., k. \quad (18)$$

Replacing (10) in (14) with (18) and solving the bilinear problem can yield an improved lower bound.

D. S-Procedure relaxation

Suppose again the stage costs are quadratic and now that we have $R_1, ..., R_M$ and $s_1, ..., s_M$ for which

$$\mathbb{U} \subseteq \bar{U} = \{u | u^T R_i u + s_i \le 0, \ i = 1, ..., M\}. \quad (19)$$

Setting $Q = Q_0$ and $\gamma \ge \gamma_0$ again, a sufficient condition via the S-procedure (see [12] and [7]) for (10) to hold is the existence of nonnegative $\lambda_1, ..., \lambda_M$ satisfying

$$R - R_0 \preceq \sum_{i=1}^M \lambda_i R_i, \quad \frac{s}{1-\alpha} \le \sum_{i=1}^M \lambda_i s_i. \quad (20)$$

The first inequality can be written using Schur complements as an LMI in R^{-1}

$$\begin{bmatrix} R_0 + \sum_{i=1}^M \lambda_i R_i & 0 \\ 0 & R^{-1} \end{bmatrix} \succeq 0. \quad (21)$$

Again, replacing (10) in (14) with (21) and solving the bilinear problem can yield an improved lower bound.

IV. Numerical Example

In this section we compute a performance bound for an example system in which the input constraint set $\mathbb{U}$ is a box. The disturbance constraint set is taken to be a unit ball, and the cost is taken to be quadratic: $\ell(z,w) = x^T Q_0 x + u^T R_0 u - \gamma_0 w^T w$. System and cost matrices with dimensions $n = l = 4$, $m = 2$, and $p = 6$ were randomly generated, with entries from A, B, G, and Cholesky factors of Q_0 and R_0 drawn from a standard normal distribution. The disturbance weight γ_0 was set to ten percent larger than the unconstrained $\mathcal{H}_\infty$ optimal value, and the discount factor was set to $\alpha = 0.95$.

For the associated unconstrained problem, we set $Q = Q_0$ and $\gamma = \gamma_0$. The problem $(\mathcal{SDP})$ is first solved with $R = R_0$ to obtain a basic lower bound. Then the bound is optimized by solving (14) by alternately fixing R^{-1} and Λ_{33}. The corresponding semidefinite programs were solved using the modeling language YALMIP [13] and the cone solver SeDuMi [14].

A box constraint is defined by

$$\mathbb{U} = \{u \in \mathcal{R}^m | \; \|u\|_\infty \leq U_{max}\}.$$

This constraint can be represented by a set of quadratic inequalities as follows

$$u^T e_i e_i^T u \leq U_{max}^2, \quad i = 1, ..., m,$$

where e_i is the ith unit vector. In relation to (19), we have $R_i = e_i e_i^T$ and $s_i = -U_{max}^2$.

We now describe a specific example. The system matrices are:

$$A = \begin{bmatrix} 0.434 & 0.050 & 0.212 & 0.007 \\ 0.264 & 0.001 & 0.092 & 0.419 \\ 0.307 & 0.255 & 0.371 & 0.359 \\ 0.364 & 0.003 & 0.291 & 0.427 \end{bmatrix}, \; B = \begin{bmatrix} 0.739 & 0.550 \\ 0.371 & 0.748 \\ 0.323 & 0.760 \\ 0.491 & 0.472 \end{bmatrix}$$

$$G = \begin{bmatrix} 0.802 & 0.666 & 0.737 & 0.629 \\ 0.471 & 0.677 & 0.866 & 0.793 \\ 0.203 & 0.9425 & 0.991 & 0.449 \\ 0.576 & 0.7701 & 0.504 & 0.524 \end{bmatrix}, \; R_0 = \begin{bmatrix} 0.262 & 0.560 \\ * & 1.33 \end{bmatrix}$$

$$Q_0 = \begin{bmatrix} 0.105 & 0.286 & 0.221 & 0.271 \\ * & 0.929 & 0.618 & 0.687 \\ * & * & 1.22 & 0.854 \\ * & * & * & 0.873 \end{bmatrix}.$$

Solving $(\mathcal{SDP})$ with $Q = Q_0$, $R = R_0$ and $\gamma = \gamma_0$ we obtain an optimal value of 3.526. We then run an alternating semidefinite programming procedure on (14) and obtain the value 6.42, which is 82% improvement on the basic bound obtained from $(\mathcal{SDP})$. Sets of initial states for which these bounds are valid can be computed via (16).

V. Conclusions

This paper developed a method to compute performance bounds for infinite-horizon linear min-max control problems with input constraints. The method requires a bilinear matrix inequality to be solved. We showed that improved bounds can be computed by solving semidefinite programs in an alternating fashion. Further research directions include computing bounds for problems with state constraints and comparison of the bound with suboptimal policies.

Appendix

A. Infinite-horizon discounted $\mathcal{H}_\infty$ problem

This section derives the solution of the infinite-horizon unconstrained linear min-max optimal control problem with a discounting factor. The quadratic form $z^T Q z$ is denoted by $\|z\|_Q^2$. Given the system

$$x^+ = Ax + Bu + Gw,$$

consider the problem of minimizing the worst-case discounted cost function with discount factor $\alpha \in (0,1)$

$$J(x_0, \pi, \mathbf{w}) := \sum_{k=0}^\infty \alpha^k (\|x\|_Q^2 + \|u\|_R^2 - \gamma^2 \|w\|^2).$$

The Isaacs equation is

$$V = \min_u \max_w [\ell(x, u, w) + \alpha V(Ax + Bu + Gw)],$$
$$= \min_u [\|x\|_Q^2 + \|u\|_R^2 + J(x, u)]$$

where

$$J(x, u) := \max_w [-(\gamma^2)\|w\|^2 + \alpha V(Ax + Bu + Gw)].$$

Assume that $V(x) = x^T P x$. The inner maximization problem can be solved explicitly by rewriting it first as

$$J(x, u)$$
$$= \max_w [-\gamma^2 \|w\|^2 + \alpha \|Ax + Bu + Gw\|_P^2]$$
$$= \max_w [-\|w\|_{(\gamma^2 I - G^T(\alpha P)G)}^2 + w^T G^T(\alpha P)(Ax + Bu)]$$
$$\quad + \|Ax + Bu\|_{\alpha P}^2$$

The optimizer $w^*(x, u)$ for the above is

$$w^*(x, u) = [\gamma^2 I - G^T(\alpha P)G]^{-1} G^T(\alpha P)(Ax + Bu),$$

so that

$$J(x, u) = (Ax + Bu)\bar{P}(Ax + Bu)$$
$$\bar{P} := (\alpha P) + (\alpha P)G (\gamma^2 I - G^T(\alpha P)G)^{-1} G^T(\alpha P).$$

The value function can then be written as

$$V(x) = \min_u [\|x\|_Q^2 + \|u\|_Q^2 + J(x, u)]$$
$$= \min_u [\|x\|_Q^2 + \|u\|_Q^2 + (Ax + Bu)^T \bar{P}(Ax + Bu)].$$

The rest of the argument then follows along the lines of those in [10], yielding

$$P = Q + A^T \bar{P} A - A^T \bar{P} B(R + B^T \bar{P} B)^{-1} B^T \bar{P} A \quad (22)$$
$$\bar{P} := (\alpha P) + (\alpha P)G (\gamma^2 I - G^T(\alpha P)G)^{-1} G^T(\alpha P). \quad (23)$$

The optimal input and disturbance policies are

$$K_u = -(R + B^T \bar{P} B)^{-1} B^T \bar{P} A,$$
$$K_w = (\gamma^2 I - G^T(\alpha P)G)^{-1} G^T(\alpha P)(A + B K_u).$$

We can now try to equate a solution of the above to some other undiscounted problem. Rewrite (23) as

$$\bar{P} = (\alpha P) + \frac{1}{\alpha}(\alpha P)G\left(\frac{\gamma^2}{\alpha}I - G^T PG\right)^{-1} G^T(\alpha P)$$
$$= \alpha \cdot \left[P + PG\left(\tilde{\gamma}^2 I - G^T PG\right)^{-1} G^T P\right] = \alpha \cdot \bar{\mathcal{P}}$$

where $\tilde{\gamma} := \gamma/\sqrt{\alpha}$ and

$$\bar{\mathcal{P}} := P + PG\left(\tilde{\gamma}^2 I - G^T PG\right)^{-1} G^T P.$$

Then substitute $\bar{P} = \alpha\bar{\mathcal{P}}$ into (22) to obtain

$$P = Q + A^T(\alpha\bar{\mathcal{P}})A - A^T(\alpha\bar{\mathcal{P}})B(R + B^T(\alpha\bar{\mathcal{P}})B)^{-1}B^T(\alpha\bar{\mathcal{P}})A$$
$$= Q + \tilde{A}^T\bar{\mathcal{P}}\tilde{A} - \tilde{A}^T\bar{\mathcal{P}}\tilde{B}(R + \tilde{B}^T\bar{\mathcal{P}}\tilde{B})^{-1}\tilde{B}^T\bar{\mathcal{P}}\tilde{A}$$

where $\tilde{A} := \sqrt{\alpha}A$ and $\tilde{B} := \sqrt{\alpha}B$. The conclusion is that the solution to the discounted $\mathcal{H}_\infty$ problem can be obtained by solving the equations

$$\bar{\mathcal{P}} = P + PG\left(\tilde{\gamma}^2 I - G^T PG\right)^{-1} G^T P$$
$$P = Q + \tilde{A}^T\bar{\mathcal{P}}\tilde{A} - \tilde{A}^T\bar{\mathcal{P}}\tilde{B}(R + \tilde{B}^T\bar{\mathcal{P}}\tilde{B})^{-1}\tilde{B}^T\bar{\mathcal{P}}\tilde{A}$$

i.e. by solving an undiscounted problem using the problem data $\tilde{A} \leftarrow \sqrt{\alpha}A$, $\tilde{B} \leftarrow \sqrt{\alpha}B$ and $\tilde{\gamma} \leftarrow \gamma/\sqrt{\alpha}$.

The value function for the discounted problem is then $V(x) = x^T Px$, where P comes from the solution to the undiscounted $\mathcal{H}_\infty$ problem using the modified problem data.

Remark 3: We note that the control that minimizes the infinite-horizon discounted cost is not necessarily stabilizing for the system (1), even if the optimal cost is finite. However, only the cost is used when determining a lower bound.

B. Equivalence of Riccati Inequality and Bounded-Real LMI

This section demonstrates the equivalence between the Riccati inequality obtained by replacing $=$ with $\succeq$ in (5) and the bounded-real LMI in (7). We first recall a result from [8].

Theorem 1 (Thm. 2.2, [8]): For the system

$$\begin{aligned} x^+ &= Ax + Gw \\ z &= Cx, \end{aligned} \tag{24}$$

the following statements are equivalent:

1) A is a stable matrix and $\|C(zI - A)^{-1}G\|_\infty \leq \gamma$.
2) There exists a stabilizing solution $P = P^T \succeq 0$ to the Riccati equation

$$P = A^T PA + \frac{1}{\gamma^2}A^T PG(I - \frac{1}{\gamma^2}G^T PG)^{-1}G^T PA + C^T C. \tag{25}$$

Claim 1: If one substitutes $A \rightarrow (A + BK)$ and $C \rightarrow (C + DK)$ in (25) with K as in (6), then (25) $\Leftrightarrow$ (5).

Proof: We have from (25)

$$P = A^T PA + A^T PG\left[\gamma^2 - G^T PG\right]^{-1} G^T PA + C^T C$$
$$= A^T PA + A^T(\bar{P} - P)A + C^T C = A^T\bar{P}A + C^T C,$$

where the first step uses the definition of $\bar{P}$ from (5).

Next replace $A \rightarrow (A + BK)$ and $C \rightarrow (C + DK)$ to get

$$P = (A + BK)^T\bar{P}(A + BK) + (C + DK)^T(C + DK)$$
$$= (A + BK)^T\bar{P}(A + BK) + Q + K^T RK$$
$$= A^T\bar{P}A + Q + K^T(R + B^T\bar{P}B)K$$
$$\quad + A^T\bar{P}BK + K^T B^T\bar{P}A, \tag{26}$$

where the first step comes from the assumptions $D^T D = R$, $C^T C = Q$, and $C^T D = 0$. Recalling the definition of K in (6), the final three terms in (26) can be rewritten as

$$K^T(R + B^T\bar{P}B)K + A^T\bar{P}BK + K^T B^T\bar{P}A \tag{27}$$
$$= -A^T\bar{P}B(R + B^T\bar{P}B)^{-1}B^T\bar{P}A. \tag{28}$$

Substituting this into (26) results in

$$P = A^T\bar{P}A + Q - A^T\bar{P}B(R + B^T\bar{P}B)^{-1}B^T\bar{P}A,$$

which is identical to (5). ∎

Finally, if we now relax (25) by replacing $=$ with $\succeq$, apply the Schur complement lemma three times and substitute $A \rightarrow (A + BK)$ and $C \rightarrow (C + DK)$, we arrive at the bounded-real LMI in (7).

REFERENCES

[1] D. Bertsekas and S. Shreve. *Stochastic optimal control: The discrete time case*, volume 139. Athena Scientific, 1978.

[2] D.Q. Mayne, SV Raković, RB Vinter, and EC Kerrigan. Characterization of the solution to a constrained $\mathcal{H}_\infty$ optimal control problem. *Automatica*, 42(3):371–382, 2006.

[3] S. Lall and K. Glover. A game theoretic approach to moving horizon control. *Advances in model-based predictive control*, pages 131–144, 1994.

[4] M.V. Kothare, V. Balakrishnan, and M. Morari. Robust constrained model predictive control using linear matrix inequalities. *Automatica*, 32(10):1361–1379, 1996.

[5] J. Lofberg. Approximations of closed-loop minimax MPC. In *Proceedings of the 42nd IEEE Conference on Decision and Control*, pages 1438–1442, 2003.

[6] P.J. Goulart, E.C. Kerrigan, and T. Alamo. Control of constrained discrete-time systems with bounded ℓ_2 gain. *IEEE Transactions on Automatic Control*, 54(5):1105–1111, 2009.

[7] Y. Wang and S. Boyd. Performance bounds for linear stochastic control. *Systems & Control Letters*, 58(3):178–182, 2009.

[8] C.E. de Souza and L. Xie. On the discrete-time bounded real lemma with application in the characterization of static state feedback $\mathcal{H}_\infty$ controllers. *Systems & Control Letters*, 18(1):61–71, 1992.

[9] P. Gahinet and P. Apkarian. A linear matrix inequality approach to h control. *International Journal of Robust and Nonlinear Control*, 4(4):421–448, 1994.

[10] T. Basar and P. Bernhard. $\mathcal{H}_\infty$ *optimal control and related minimax design problems*. Birkhauser Boston, 2004.

[11] M. Kocvara and M. Stingl. PENBMI users guide. *Available from http://www.penopt.com*, 2005.

[12] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear matrix inequalities in system and control theory*, volume 15. Society for Industrial and Applied Mathematics, 1994.

[13] J. Lofberg. Yalmip: A toolbox for modeling and optimization in matlab. In *IEEE International Symposium on Computer Aided Control Systems Design*, pages 284–289. IEEE, 2004.

[14] J.F. Sturm. Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones. *Optimization methods and software*, 11(1-4):625–653, 1999.