ME 7247: Advanced Control Systems **Fall 2021–22**

Lecture 24: Steady state LEQR and the S-procedure

December 10, 2021

Lecturer: Laurent Lessard *Scribe: Michael Shaham*

This lecture covers the steady state LEQR and the S-procedure for determining when one set described by a quadratic inequality is contained within another.

1 Steady state linear exponential quadratic regulator

Consider the linear system given by

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + w_t \\ z_t &= Fx_t + Hu_t. \end{aligned}$$

We saw in Lecture 23 two different formulations of a robust optimization problem that trades off a higher average cost for a lower variance. These two formulations are the linear exponential quadratic regulator, given by

$$\underset{u_0,u_1,\dots,u_{N-1}}{\text{minimize}} \quad \gamma^2 \log \mathbf{E}\left[\exp\left(\frac{1}{\gamma^2}\sum_{t=0}^{N-1}\|z_t\|^2\right)\right] \tag{1}$$

and the dynamic game formulation, given by

$$\begin{aligned} \underset{u_0,\dots,u_{N-1}}{\text{minimize}} \quad \underset{w_0,\dots,w_{N-1}}{\text{maximize}} \quad & \sum_{t=0}^{N-1}\left(x_t^\mathsf{T} Q x_t + u_t^\mathsf{T} R u_t - \gamma^2\|w_t\|^2\right) + x_N^\mathsf{T} Q_f x_N \\ \text{subject to} \quad & x_{t+1} = Ax_t + Bu_t + w_t, \quad t = 0,\dots,N-1. \end{aligned} \tag{2}$$

In both of these formulations, γ is a parameter that trades off between mean cost and cost variance. We also learned that these two problems have identical solutions that can be found via dynamic programming. The optimal controller for each of these finite-horizon robust control problems is given by the recurrence relation

$$\begin{aligned} P_t &= A^\mathsf{T}\tilde{P}_{t+1}A + Q - A^\mathsf{T}\tilde{P}_{t+1}B\left(B^\mathsf{T}\tilde{P}_{t+1}B + R\right)^{-1}B^\mathsf{T}\tilde{P}_{t+1}A \\ \tilde{P}_{t+1} &= \left(P_{t+1}^{-1} - \gamma^{-2}I\right)^{-1} \\ K_t &= -\left(B^\mathsf{T}\tilde{P}_{t+1}B + R\right)^{-1}B^\mathsf{T}\tilde{P}_{t+1}A \\ u_t &= K_t x_t \end{aligned}$$

where $P_t \prec \gamma^2 I$ for all t.

At steady state, this becomes

$$\begin{aligned} P &= A^\top \tilde{P} A + Q - A^\top \tilde{P} B \left(B^\top \tilde{P} B + R\right)^{-1} B^\top \tilde{P} A \\ \tilde{P} &= \left(P^{-1} - \gamma^{-2} I\right)^{-1} \\ K &= -\left(B^\top \tilde{P} B + R\right)^{-1} B^\top \tilde{P} A \\ u_t &= K x_t. \end{aligned} \tag{3}$$

where $P \prec \gamma^2 I$. Why is this controller useful? We know that the LQR controller was also a stabilizing controller (under closed-loop state feedback, the resulting system is stable, i.e., the eigenvalues of $A + BK$ are all of magnitude less than 1). For the infinite horizon LEQR controller, we not only have a stabilizing controller, but we also achieve a root mean square (RMS) gain less than γ. Recall that for a system given by

$$\begin{aligned} x_{t+1} &= A x_t + B w_t \\ y_t &= C x_t, \end{aligned}$$

the bounded-real lemma states that

$$\underset{w_0, w_1, \ldots}{\text{maximize}} \; \frac{\sum_{t=0}^{\infty} \|y_t\|^2}{\sum_{t=0}^{\infty} \|w_t\|^2} < \gamma^2$$

if and only if there exists $P \succ 0$ such that

$$\begin{bmatrix} A^\top P A - P + C^\top C & A^\top P B \\ B^\top P A & B^\top P B - \gamma^2 I \end{bmatrix} \prec 0,$$

which occurs if and only if the following two conditions are satisfied:

$$A^\top P A - P + C^\top C - A^\top P B \left(B^\top P B - \gamma^2 I\right)^{-1} B^\top P A \prec 0 \tag{4a}$$

$$B^\top P B - \gamma^2 I \prec 0. \tag{4b}$$

This follows from the fact that for block matrices, we have the Schur complement property

$$\begin{bmatrix} A & B \\ B^\top & D \end{bmatrix} \succ 0 \qquad \Longleftrightarrow \qquad D \succ 0 \quad \text{and} \quad A - B D^{-1} B^\top \succ 0$$

We will show that the infinite horizon LEQR satisfies the condition (4). Our system is given by

$$\begin{aligned} x_{t+1} &= A x_t + B u_t + w_t \\ z_t &= \begin{bmatrix} Q^{1/2} \\ 0 \end{bmatrix} x_t + \begin{bmatrix} 0 \\ R^{1/2} \end{bmatrix} u_t. \end{aligned} \tag{5}$$

Note that with this system, we have

$$\sum_{t=0}^{\infty} \|z_t\|^2 = \sum_{t=0}^{\infty} \left(x_t^\top Q x_t + u_t^\top R u_t\right)$$

which is equivalent to the infinite horizon LQR cost. Our claim is that the LEQR controller for the system given by (5) has squared RMS cost no greater than γ^2, where γ is the parameter selected for use in the robust control cost functions given by (1) and (2). The rest of this section is dedicated to proving this.

Since K in (3) is a function of γ, lets define $K_\gamma := K$ to make the notation more clear (we could also use $\tilde{P}_\gamma$ and P_γ to make it even more clear, but we won't do that). Since $u_t = K_\gamma x_t$, (5) becomes

$$\begin{aligned} x_{t+1} &= (A + BK_\gamma)x_t + w_t \\ z_t &= \begin{bmatrix} Q^{1/2} \\ R^{1/2}K_\gamma \end{bmatrix} x_t. \end{aligned} \tag{6}$$

With this system, our "A" matrix is $A + BK$, our "B" matrix is I, and our "C" matrix is $\begin{bmatrix} Q^{1/2} \\ R^{1/2}K_\gamma \end{bmatrix}$. Substituting these in for A, B, and C in the two conditions given by (4), Eq. (4b) becomes

$$P - \gamma^2 I \prec 0,$$

which is already a requirement for our controller given by (3), Eq. (4a) becomes

$$\begin{aligned} &(A + BK_\gamma)^\mathsf{T} P(A + BK_\gamma) - P + Q + K_\gamma^\mathsf{T} RK_\gamma - (A + BK_\gamma)^\mathsf{T} P(P - \gamma^2 I)^{-1} P(A + BK_\gamma) \\ &\quad= (A + BK_\gamma)^\mathsf{T} \left(P - P(P - \gamma^2 I)^{-1} P\right)(A + BK_\gamma) - P + Q + K_\gamma^\mathsf{T} RK_\gamma \\ &\quad= (A + BK_\gamma)^\mathsf{T} \left(P^{-1} - \gamma^{-2} I\right)^{-1} (A + BK_\gamma) - P + Q + K_\gamma^\mathsf{T} RK_\gamma \\ &\quad= (A + BK_\gamma)^\mathsf{T} \tilde{P}(A + BK_\gamma) - P + Q + K_\gamma^\mathsf{T} RK_\gamma \\ &\quad= 0 \end{aligned}$$

where lines 2–3 follow from the matrix inversion lemma, lines 3–4 follow from the definition of $\tilde{P}$ in (3), and lines 4–5 follow from the algebraic Riccati equation in (3) but with K_γ substituted in. The equality at the end actually indicates that γ is exactly the system's RMS bound.

2 The S-procedure

2.1 The lossless S-procedure

The (lossless) S-procedure is used to determine if

$$x^\mathsf{T} P_1 x \le 0 \implies x^\mathsf{T} P_0 x \le 0 \tag{7}$$

where $P_1, P_0 \in \mathbb{R}^{n \times n}$ are symmetric matrices. Note that this is the only requirement—we make no assumptions on the definiteness of these two matrices. We can also express this problem in a different way by defining the two sets

$$\begin{aligned} S_0 &:= \{x \in \mathbb{R}^n \mid x^\mathsf{T} P_0 x \le 0\} \\ S_1 &:= \{x \in \mathbb{R}^n \mid x^\mathsf{T} P_1 x \le 0\} \end{aligned}$$

and then asking the question: "is S_1 contained in S_0?" or, equivalently, "is it true that $S_1 \subseteq S_0$?". Both S_0 and S_1 are *cones*. That is, they have the property that $\alpha x \in S$ for all $x \in S$ and $\alpha \geq 0$.

If we can prove inclusion for cones, then we get inclusion for any set defined by a quadratic constraint (i.e., possibly with linear or constant terms). To see why, consider the more general case

$$x^\mathsf{T} P_1 x + 2q_1^\mathsf{T} x + r_1 \leq 0 \quad \Longrightarrow \quad x^\mathsf{T} P_0 x + 2q_0^\mathsf{T} x + r_0 \leq 0. \tag{8}$$

We can write Eq. (8) as

$$\begin{bmatrix} x \\ z \end{bmatrix}^\mathsf{T} \begin{bmatrix} P_1 & q_1 \\ q_1^\mathsf{T} & r_1 \end{bmatrix} \begin{bmatrix} x \\ z \end{bmatrix} \leq 0 \quad \Longrightarrow \quad \begin{bmatrix} x \\ z \end{bmatrix}^\mathsf{T} \begin{bmatrix} P_0 & q_0 \\ q_0^\mathsf{T} & r_0 \end{bmatrix} \begin{bmatrix} x \\ z \end{bmatrix} \leq 0 \tag{9}$$

The goal is to prove (8) for $z = 1$ and for all x. However, because the inequalities are homogeneous, this is equivalent to having them hold for all x and z. So the problem of containment of shifted quadratic forms is equivalent to the problem of containment of cones. See Fig. 1.

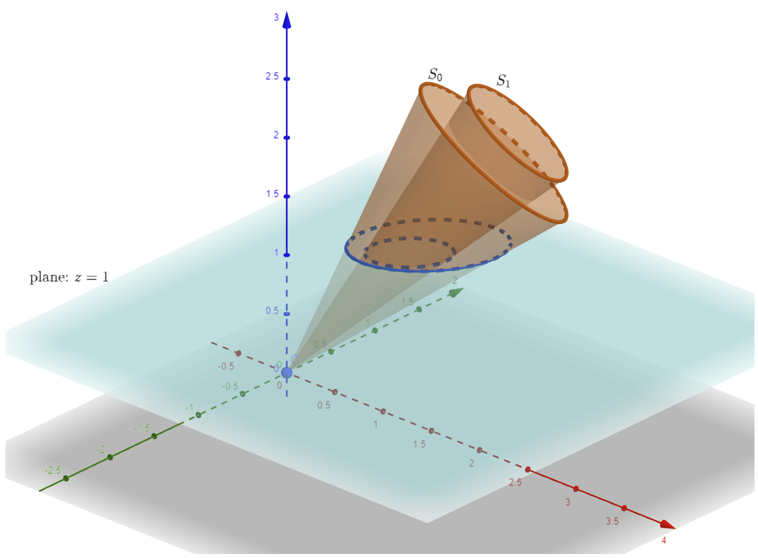

Figure 1: Illustration of a scenario where $S_1 \subseteq S_0$. Both sets are cones. If we slice the cones, for example by intersecting them with the plane $z = 1$ as shown, we get two ellipses that satisfy the same containment property.

Now, let's try to find a sufficient condition: a condition that when true, proves that (7) holds. Suppose there exists $\lambda \geq 0$ such that

$$x^\mathsf{T} P_0 x \leq \lambda x^\mathsf{T} P_1 x \quad \text{for all } x. \tag{10}$$

Then (7) holds. Proof: Clearly, if $x^\mathsf{T} P_1 x \leq 0$, then $\lambda x^\mathsf{T} P_1 x \leq 0$. Since $x^\mathsf{T} P_0 x \leq \lambda x^\mathsf{T} P_1 x$, we must have $x^\mathsf{T} P_0 x \leq 0$ when $x^\mathsf{T} P_1 x \leq 0$, which gives us (7).

This condition generalizes to arbitrary functions, not just quadratics. Specifically: for any two functions $f : \mathbb{R}^n \to \mathbb{R}$ and $g : \mathbb{R}^n \to \mathbb{R}$, if there exists $\lambda \geq 0$ such that $g(x) \leq \lambda f(x)$ for all x, then for any x satisfying $f(x) \leq 0$, we must have $g(x) \leq 0$.

Now, if (10) holds, then we have $x^\mathsf{T}(P_0 - \lambda P_1)x \le 0$ for all x, which implies that

$$P_0 \preceq \lambda P_1.$$

Thus, if we can find $\lambda \ge 0$ such that $P_0 \preceq \lambda P_1$, then (7) holds. Note that $P_0 \preceq \lambda P_1$ is a linear matrix inequality (LMI), and finding $\lambda \ge 0$ that satifies this LMI can be formulated as a convex optimization problem, specifically, a semidefinite program (SDP). However, for this to be even more useful, we need to show the converse, i.e., we would like to show that our sufficient condition is also a necessary condition for (7) to hold.

If we can show the converse, then we have that

$$x^\mathsf{T} P_1 x \le 0 \implies x^\mathsf{T} P_0 x \le 0$$

holds if and only if there exists $\lambda \ge 0$ such that

$$P_0 \preceq \lambda P_1.$$

This is called the lossless S-procedure. It was easy to show the forward direction (sufficiency). The reverse direction (that the condition is also necessary) is a bit more difficult to prove, and the proof can be found in Prof. Lessard's S-procedure supplementary note.

Note on non-convexity. Although Fig. 1 illustrates the case of convex cones in 3 dimensions, the result also holds when the P_k matrices are indefinite, so the cones could be non-convex in general, and far more complicated. The resulting sets S_k could also be non-convex, and possibly even disconnected! (Imagine both sides of a cone, like an hourglass shape, and a plane that intersects both sides of it...)

2.2 The lossy S-procedure

The S-procedure can be generalized to the following form: we would like to determine if

$$x^\mathsf{T} P_k x \le 0 \quad \text{for } k = 1, \dots, m \qquad \implies \qquad x^\mathsf{T} P_0 x \le 0. \tag{11}$$

The condition from the lossless S-procedure can be generalized to: if there exists $\lambda_k \ge 0,\ k = 1, \dots, m$ such that

$$x^\mathsf{T} P_0 x \le \sum_{k=1}^{m} \lambda_k x^\mathsf{T} P_k x \quad \text{for all } x,$$

then (11) holds. This condition can also be formulated as the LMI

$$P_0 \preceq \sum_{k=1}^{m} \lambda_k P_k,$$

and finding $\lambda \ge 0$ to satisfy this LMI is again an SDP. However, the existence of $\lambda \ge 0$ satisfying this LMI is only a sufficient condition for (11) to hold, but it is not necessary (one can construct counter examples for any $m > 1$ to show that the condition is only sufficient). This is known as the "lossy" S-procedure.

To summarize, let $P_0, P_1, \ldots, P_m \in \mathbb{R}^{n\times n}$. Then

$$\begin{array}{ccc} x^\mathsf{T} P_k x \le 0, \quad \text{for } k = 1,\ldots,m & \Longrightarrow & x^\mathsf{T} P_0 x \le 0 \\ \Updownarrow : \text{"lossless"} & & \Uparrow : \text{"lossy"} \\ \multicolumn{3}{c}{\text{There exists } \lambda_k \ge 0 \text{ such that } \quad P_0 \preceq \sum_{k=1}^{m} \lambda_k P_k.} \end{array}$$

The lossless case occurs only when $m = 1$, and the lossy case occurs for all $m > 1$.

2.3 Solving the lossless S-procedure using (convex) optimization

The lossless S-procedure seeks to determine if (7) holds. Consider the optimization problem

$$\begin{aligned} p^\star \quad = \quad & \underset{x}{\text{maximize}} \quad x^\mathsf{T} P_0 x \\ & \text{subject to} \quad x^\mathsf{T} P_1 x \le 0. \end{aligned} \tag{12}$$

This is a quadratically constrained quadratic program (QCQP). If $p^\star \le 0$, Eq. (7) holds, i.e., if $p^\star \le 0$ we have $x^\mathsf{T} P_1 x \le 0 \implies x^\mathsf{T} P_0 x \le 0$. However, the optimization problem given by (12) is, in general, nonconvex (it is convex if and only if P_0 is negative semidefinite and P_1 is positive semidefinite). Typically, it is extremely difficult to find globally optimal solutions to nonconvex problems. The S-procedure will show us that it is actually very easy to find globally optimal solutions to this problem.

Let's start by finding the dual problem for (12). The dual function for (12) is given by

$$\begin{aligned} F(\lambda) &= \max_x \quad \left(x^\mathsf{T} P_0 x - \lambda x^\mathsf{T} P_1 x\right) \\ &= \max_x \quad x^\mathsf{T}(P_0 - \lambda P_1 x)x \\ &= \begin{cases} 0, & P_0 - \lambda P_1 \preceq 0 \\ \infty, & P_0 - \lambda P_1 \succ 0. \end{cases} \end{aligned}$$

Thus, the dual problem is

$$\begin{aligned} d^\star \quad = \quad & \underset{\lambda}{\text{minimize}} \quad 0 \\ & \text{subject to} \quad P_0 - \lambda P_1 \preceq 0 \\ & \qquad\qquad\quad \lambda \ge 0. \end{aligned} \tag{13}$$

This is a feasibility problem (any feasibility problem can be written as an optimization problem with a constant objective). This is identical to the necessary and sufficient condition we found for the lossless S-procedure: to show that (7) holds, we try to find any $\lambda \ge 0$ such that $P_0 \preceq \lambda P_1$. If such a λ exists, the solution to (13) is zero, i.e., $d^\star = 0$. Otherwise, we say the solution is infinity ($d^\star = \infty$), since the problem is infeasible.

From weak duality, we know that $p^\star \le d^\star$ (note: this is the weak duality result for maximization problems, whereas the weak duality we derived in a previous lecture was for minimization problems).

To be more explicit, we have

$$p^\star = \left(\begin{array}{ll} \underset{x}{\text{maximize}} & x^\mathsf{T} P_0 x \\ \text{subject to} & x^\mathsf{T} P_1 x \le 0 \end{array} \right) \le \left(\begin{array}{ll} \underset{\lambda}{\text{minimize}} & 0 \\ \text{subject to} & P_0 - \lambda P_1 \preceq 0 \\ & \lambda \ge 0. \end{array} \right) = d^\star. \tag{14}$$

We've already discussed how the solution to the dual problem is either zero or infinity, i.e., $d^\star = 0$ or $d^\star = \infty$. This is also true for the primal problem, i.e., $p^\star = 0$ or $p^\star = \infty$. To see this, first note that we can lower bound $p^\star$ by taking $x = 0$, so $p^\star \ge 0$. Now suppose there exists a primal feasible x such that $x^\mathsf{T} P_0 x > 0$. By positively scaling x by some constant $\alpha > 1$, we maintain primal feasibility and increase the objective function. Taking $\alpha \to \infty$, we then have $p^\star = \infty$. Thus, we must have either $p^\star = 0$ or $p^\star = \infty$.

Now, let's consider the sufficient condition for the lossless S-procedure, which states that if there exists $\lambda \ge 0$ such that $P_0 \preceq \lambda P_1$, then $x^\mathsf{T} P_1 x \le 0 \implies x^\mathsf{T} P_0 x \le 0$. This is clearly portrayed in this weak duality relationship. If there exists $\lambda \ge 0$ such that $P_0 \preceq \lambda P_1$, then we have $d^\star = 0$. Since $p^\star \le d^\star$ and $p^\star \in \{0, \infty\}$, we must have $p^\star = 0$.

On the other hand, if $x^\mathsf{T} P_1 x \le 0 \implies x^\mathsf{T} P_0 x \le 0$, then we have $p^\star = 0$. The lossless S-procedure tells us that $x^\mathsf{T} P_1 x \le 0 \implies x^\mathsf{T} P_0 x \le 0$ if and only if there exists $\lambda \ge 0$ such that $P_0 - \lambda P_1 \preceq 0$. Thus, if $p^\star = 0$, we must have $d^\star = 0$, and strong duality actually holds here. Therefore,

$$p^\star = \left(\begin{array}{ll} \underset{x}{\text{maximize}} & x^\mathsf{T} P_0 x \\ \text{subject to} & x^\mathsf{T} P_1 x \le 0 \end{array} \right) = \left(\begin{array}{ll} \underset{\lambda}{\text{minimize}} & 0 \\ \text{subject to} & P_0 - \lambda P_1 \preceq 0 \\ & \lambda \ge 0. \end{array} \right) = d^\star. \tag{15}$$

What (15) is telling us is that if we want to solve the nonconvex QCQP given by (12), we can instead solve (13), which is a convex SDP. This is one of the few examples in optimization where there exists strong duality between a nonconvex primal problem and its convex dual problem. This result only holds for QCQPs with only a single quadratic constraint. If there are more quadratic constraints (the $m > 1$ case), we lose strong duality. Even further, we lose strong duality if there are any additional linear inequality constraints added to (12).

2.4 Application of the S-procedure: the Lur'e problem

Consider the closed-loop system given by

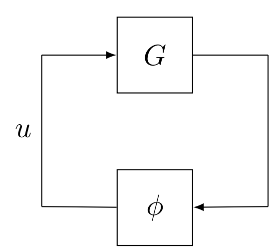

where G is your known system and ϕ is some nonlinearity. Now suppose we know that this nonlinearity belongs to some set S, i.e., we know $\phi \in S$. We want to answer the question: given G and $\phi \in S$, where S known, is our closed-loop system guaranteed to be stable?

A simple example of this is

$$x_{t+1} = Ax_t$$

where

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} + \varepsilon \end{bmatrix}, \quad |\varepsilon| < 1.$$

Typically, we can check stability by looking at the eigenvalues of A. However, there is now an infinite number of possible versions of A, making this a more difficult task. In this case, it may be possible to solve this algebraically, but for larger systems, this becomes infeasible. Further, in this example, our closed-loop system still remains linear. The formulation above using ϕ is more general, allowing for the closed-loop system to be nonlinear.

We will consider a class of systems where the nonlinearities are sector-bounded. Consider the graph shown below.

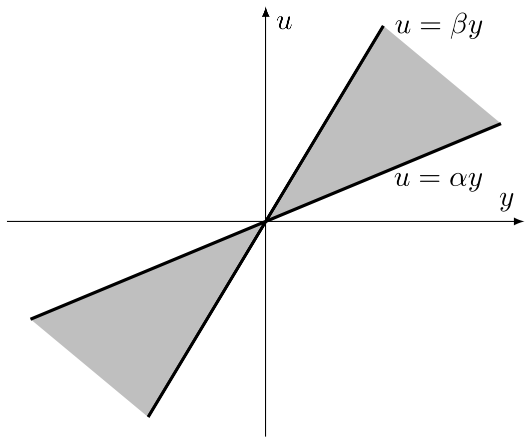

A sector nonlinearity ϕ is any function that falls within the shaded region of this graph. Note that this function can be arbitrarily complex, as long as it falls within this region. Thus, we want to know whether our system is stable when the nonlinearity ϕ falls within this region. We can characterize the region with a single quadratic constraint. Notice that

$$\text{if } y \geq 0 \text{ then } \alpha y \leq u \leq \beta y,$$
$$\text{if } y \leq 0 \text{ then } \beta y \leq u \leq \alpha y.$$

This can be combined into a single quadratic constraint given by

$$(\alpha y - u)(\beta y - u) \leq 0. \tag{16}$$

To see this, note that if we subtract u from both inequalities, when $y \geq 0$, we have $\beta y - u \geq 0$ and $\alpha y - u \leq 0$, and when $y \leq 0$, we have $\beta y - u \leq 0$ and $\alpha y - u \geq 0$. In either case, if we multiply the two constraints together, the product must be zero or negative, which yields (16).

Now, to show our closed loop system is stable, we use Lyapunov theory. We want to find a function V such that

$$V(x_{t+1}) - V(x_t) \leq 0. \tag{17}$$

If we assume our system has linear dynamics given by

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t \\ y_t &= Cx_t \end{aligned}$$

then we can try to find a quadratic function given by $V(x_t) = x^\top P x$, $P \succ 0$ that satisfies the Lyapunov decrease condition given by (17). Substituting all of this into (17), our Lyapunov decrease condition becomes

$$(Ax + Bu)^\top P(Ax + Bu) - x^\top Px \leq 0.$$

Our task is to prove that whenever ϕ belongs to the sector, the Lyapunov decrease condition holds. In other words,

$$(\alpha y - u)^\top(\beta y - u) \leq 0 \quad \Longrightarrow \quad (Ax + Bu)^\top P(Ax + Bu) - x^\top Px \leq 0.$$

From the S-procedure, we know we can determine if this is true by finding $\lambda \geq 0$ such that

$$(Ax + Bu)^\top P(Ax + Bu) - x^\top Px \leq \lambda(\alpha y - u)^\top(\beta y - u). \tag{18}$$

Substituting $y = Cx$, Eq. (18) becomes quadratic in (x, u). Rearranging to put everything on the same side of the inequality, we ultimately seek $\lambda \geq 0$ and $P \succ 0$ such that

$$\begin{bmatrix} x \\ u \end{bmatrix}^\top \begin{bmatrix} A^\top PA - P - \lambda\alpha\beta C^\top C & A^\top PB + \lambda\alpha C^\top \\ B^\top PA + \lambda\beta C & B^\top PB - \lambda I \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix} \leq 0 \qquad \text{for all } x \text{ and } u.$$

Symmetrizing the quadratic form, this is equivalent to finding $\lambda \geq 0$ and $P \succ 0$ such that

$$\begin{bmatrix} A^\top PA - P - \lambda\alpha\beta C^\top C & A^\top PB + \lambda\frac{\alpha+\beta}{2}C^\top \\ B^\top PA + \lambda\frac{\alpha+\beta}{2}C & B^\top PB - \lambda I \end{bmatrix} \preceq 0.$$

Note that all terms in this are linear in P and λ, so this is an LMI, and we can solve this using any SDP solver.