

The feasibility issue in trajectory tracking by means of regions-of-attraction-based gain scheduling^{*}

Leandro R. Lustosa^{*} François Defay^{*} Jean-Marc Moschetta^{*}

^{*} *Institut Supérieur de l'Aéronautique et de l'Espace, Toulouse 31400,
France (e-mail: {leandro-ribeiro.lustosa, francois.defay,
jean-marc.moschetta}@isae.fr).*

Abstract: Linear control theory has been long established and a myriad of techniques are available for designing controllers for linear systems in view of conflicting performance requirements. On the other hand, nonlinear control techniques are often tailored to specific applications and versatile nonlinear control frameworks are still on their infancy. A common approach is to resort to local linearized descriptions at desired set-points over a given desired trajectory and employ linear tools. Furthermore, to enforce stability when switching controllers, the regions-of-attraction approach has gained recent attention. This paper questions whether such method – when applied to a well-posed smooth nonlinear controllable system – always yields a sequence of controllers that successfully tracks a given reference equilibrium trajectory, and an analytic counter-example is provided and thoroughly discussed. Finally, our case study additionally shed light on how gain scheduling fails to track particular trajectories for certain globally controllable systems.

Keywords: Tracking; Switching stability and control; Stability of nonlinear systems.

1. INTRODUCTION

Recent advances in avionics hardware technology allow for novel unmanned aerial vehicles (UAVs) architectures as well as increased maneuverability of traditional ones. For instance, the increasingly popular convertible architectures such as tilt-body and tilt-wing vehicles are capable of switching between helicopter and fixed-wing modes during flight allowing for hovering capabilities in a long-endurance vehicle (Lustosa et al. (2015)). Such vehicles exhibit wide heterogeneous flight envelopes which preclude global linear controllers employment as auto-pilots and/or stability augmentation systems (SAS). Nonlinear control techniques, however, are often tailored to specific applications and general nonlinear control frameworks are still on their infancy.

For instance, feedback linearization methods render a nonlinear system linear by means of appropriate state diffeomorphism transformations. However, such transformation exists only for fully actuated systems (generally absent in aerial robotics) and, additionally, controller design is often challenging due to unstable zero dynamics. Sliding mode control, on the other hand, applies to underactuated systems but model uncertainties may lead to excessive chattering and excitation of unmodeled dynamics. Advances in computing hardware and optimization theory allow for successful model predictive control design but its non-convex nature poses challenges for high dimensional embedded real-time applications.

Gain-scheduling techniques, on the other hand, are based on linearization of nonlinear systems at desired operating points thus profiting of well established and computationally inexpensive linear control techniques. The manual sequential controller design is tedious but has been increasingly replaced by H_∞ -based automated algorithms (Gahinet and Apkarian (2011)). Furthermore, advances in regions-of-attraction computation by means of direct computation of Lyapunov functions using sum-of-squares (SOS) optimization (Johansen (2000); Parrilo (2000)) allow for efficient stability guarantees while switching scheduled controllers. This regions-of-attraction-based gain scheduling (RoA-GS) strategy was recently validated in agile fixed-wing vehicles (Moore et al. (2014)) and promoted to a motion planning technique by means of a sparse randomized tree of scheduled linear quadratic regulators (LQR) (Tedrake et al. (2010)).

The present paper contributes to the RoA-GS framework by investigating whether the RoA-GS planning technique always yield a finite (or even countable infinite) sequence of controllers leading to a desired set-point given a well-posed controllable nonlinear system and a reference equilibrium trajectory. Section 2 revisits the RoA-GS method and sets up notation. Section 3 proposes a well-posed nonlinear system that is globally controllable and fully-actuated almost everywhere, but RoA-GS fails to track a given equilibrium trajectory. Finally, section 4 presents opportunities for future work and conclusion.

^{*} This work is partially supported by the Conselho Nacional de Desenvolvimento Científico e Tecnológico, CNPq (Brazilian National Science Foundation), through the "Ciência sem Fronteiras" program.

2. REGIONS-OF-ATTRACTION-BASED SWITCHING CONTROL REVISITED

Consider a nonlinear globally controllable dynamical system $f \in C^\infty$ of the form

$$\frac{d\mathbf{x}}{dt} = f(\mathbf{x}, \mathbf{u}) \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{u} \in \mathbb{R}^m$ are, respectively, state-space coordinates and control inputs. Furthermore, consider a given equilibrium point $(\mathbf{x}_0, \mathbf{u}_0)$, i.e.,

$$f(\mathbf{x}_0, \mathbf{u}_0) = \mathbf{0} \quad (2)$$

at which a local linear controller is designed by means of linear control techniques applied at the linearized system

$$\frac{d}{dt}\Delta\mathbf{x}_0 = A_0\Delta\mathbf{x}_0 + B_0\Delta\mathbf{u}_0 \quad (3)$$

where

$$\Delta\mathbf{x}_0 = \mathbf{x} - \mathbf{x}_0 \quad (4)$$

$$\Delta\mathbf{u}_0 = \mathbf{u} - \mathbf{u}_0 \quad (5)$$

$$A_0 = \left. \frac{\partial f}{\partial \mathbf{x}} \right|_{(\mathbf{x}_0, \mathbf{u}_0)} \quad (6)$$

and

$$B_0 = \left. \frac{\partial f}{\partial \mathbf{u}} \right|_{(\mathbf{x}_0, \mathbf{u}_0)} \quad (7)$$

Linear control design techniques (e.g., pole-placement, LQR, H_∞) yield local stabilizing controllers of the form

$$\Delta\mathbf{u}_0 = -K_0\Delta\mathbf{x}_0 \quad (8)$$

which regulate the error $\Delta\mathbf{x}_0$ on a neighborhood of \mathbf{x}_0 . A set $R_0 \subset \mathbb{R}^n$ of initial conditions $\mathbf{x}(0)$ that are regulated by this local controller K_0 is called a *region-of-attraction* (RoA), *basin of attraction* or *attractor* for the given controller around \mathbf{x}_0 (Chiang et al. (1988)). RoA computation is challenging in all but exceedingly simple systems and often conservative estimates are numerically computed instead by means of polynomial Lyapunov functions (Tan and Packard (2008); Topcu and Packard (2009); Topcu et al. (2010)).




Fig. 1. To steer \mathbf{x} from $\mathbf{x}(0)$ towards \mathbf{x}_0 , controller K_1 is employed until \mathbf{x} reaches R_0 . Afterwards, controller K_0 drives \mathbf{x} towards the desired setpoint.

Moreover, consider a second linear regulator K_1 with associated operating point $\mathbf{x}_1 \neq \mathbf{x}_0$ such that $\mathbf{x}_1 \in R_0$ and an initial state $\mathbf{x}(0) \in R_1 \setminus R_0$ (see figure 1). Since $\mathbf{x}(0) \notin R_0$, controller K_0 does not guarantee $\mathbf{x}(t)$ regulation in \mathbf{x}_0 . However, \mathbf{x}_0 -convergence is guaranteed by initially applying controller K_1 until $\mathbf{x}(t) \in R_0$ (which is guaranteed since $\mathbf{x}_1 \in R_0$) and then switching control to K_0 (we denote the time that this occurs as t_1). This methodology yields a RoA-GS controller $K_{0:1}$ with an

enlarged region-of-attraction $R_{0:1} = R_0 \cup R_1$. Finally, N iterations of the aforementioned steps yield a controller $K_{0:N}$ that potentially allows for wide regions-of-attraction

$$R_{0:N} = \bigcup_{i=0}^N R_i \quad (9)$$

as figure 2 illustrates. For instance, Tedrake et al. (2010) applies this concept to motion planning by means of a sparse randomized tree of scheduled LQR controllers.




Fig. 2. Sequence of scheduled controllers $K_{0:N}$ and respective regions-of-attraction $R_{0:N}$, for $N = 3$.

Consider now a given desired quasi-static trajectory¹ parametrized by $\mathbf{x}_d(s) : [0, 1] \rightarrow \mathbb{R}^n$, that is, we want to steer the state towards $\mathbf{x}_d(0)$ from $\mathbf{x}_d(1)$ by closely tracking the curve $\mathbf{x}_d(s)$. We chose s as parameter symbol to reinforce the absence of time performance requirements in the trajectory definition.




Fig. 3. Desired quasi-static trajectory $\mathbf{x}_d(s)$ and illustrative iterations of the regions-of-attraction based switching linear control algorithm.

A gain-scheduling control design strategy for $\mathbf{x}_d(s)$ tracking is proposed (Burrige et al. (1999)) as follows. Firstly, we design a controller K_0 for $\mathbf{x}_d(s_0)$, where $s_0 = 0$ (see figure 3), by means of any preferred linear control technique. Let $s_{0,1}$ be the point where R_0 intersects $\mathbf{x}_d(s)$, i.e.,

$$s_{0,1} \triangleq \sup\{s \in [0, 1] : \mathbf{x}_d(s) \in R_0\} \quad (10)$$

and, secondly, we design a controller K_1 at $\mathbf{x}_d(s_1)$ where

$$s_1 = \mu(s_{0,1} - s_0) + s_0 \quad (11)$$

with $0 < \mu < 1$. Appropriate μ design is dependent on the particular $\mathbf{x}_d(s)$ parametrization, K_0 and system

¹ A quasi-static trajectory is system-dependent and means herein that $f(\mathbf{x}_d(s)) = \mathbf{0}$ for all $s \in [0, 1]$, i.e., it is a trajectory composed of equilibrium points.

dynamics. Independently, small μ yield small contribution to the region-of-attraction expansion whereas overly large μ might preclude robustness. The aforementioned RoA-GS procedure is iterated until s_n reaches 1 (see figure 3) and delivers a controller $K_{0:N}$ that follows the desired trajectory.

An important question is whether the aforementioned algorithm terminates for $N < \infty$. The answer is key not only to avoid infinite loops in practical implementations but, more importantly, to better understand, as we shall soon discuss, when gain scheduling is not suitable. Figure 4 illustrates a conceptual case where the regions-of-attraction R_i are successfully advancing forward but never reaching the end goal. Is that a possibility for smooth controllable nonlinear systems? Indeed, the next section proves these phenomena possible and further discusses the matter.




Fig. 4. Illustration of the hypothesis: a sequence of controllers K_i where s_i fail to converge to 1.

3. ROA-GS FEASIBILITY – A CASE STUDY

Consider the following nonlinear system

$$\frac{d}{dt} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = f(\mathbf{x}, \mathbf{u}) = \begin{pmatrix} x_1 u_1 + x_2 \\ u_2 \end{pmatrix} \quad (12)$$

where $\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$ and $\mathbf{u} = (u_1, u_2) \in \mathbb{R}^2$ are, respectively, state variable and control input. Straight-forward linearization with respect to $(\mathbf{x}_i, \mathbf{u}_i)$ yields (for subsequent local linear control design in section 3.2)

$$A_i = \begin{bmatrix} u_{1,i} & 1 \\ 0 & 0 \end{bmatrix} \quad B_i = \begin{bmatrix} x_{1,i} & 0 \\ 0 & 1 \end{bmatrix} \quad (13)$$

where the subscript notation alludes

$$\begin{pmatrix} x_{1,i} \\ x_{2,i} \end{pmatrix} = \mathbf{x}_i \quad \begin{pmatrix} u_{1,i} \\ u_{2,i} \end{pmatrix} = \mathbf{u}_i \quad (14)$$

Notice that (A_i, B_i) is controllable for all $(\mathbf{x}_i, \mathbf{u}_i) \in \mathbb{R}^2 \times \mathbb{R}^2$ since the associated Kalman controllability matrix is full rank, i.e.,

$$\text{rank}([B_i \ A_i B_i]) = \text{rank} \left(\begin{bmatrix} x_{1,i} & 0 & u_{1,i} x_{1,i} & 1 \\ 0 & 1 & 0 & 0 \end{bmatrix} \right) = 2 \quad (15)$$

3.1 Nonlinear controllability and underactuation analysis

By inspection we conclude that $f(\mathbf{x}, \mathbf{u})$ is fully actuated in $\mathbb{A} \subset \mathbb{R}^2$, where \mathbb{A} is given by

$$\mathbb{A} = \{(x_1, x_2) \in \mathbb{R}^2 : x_1 \neq 0\} \quad (16)$$

and it is underactuated in $\mathbb{U} \subset \mathbb{R}^2$, where \mathbb{U} is given by

$$\mathbb{U} = \mathbb{R}^2 \setminus \mathbb{A} = \{(x_1, x_2) \in \mathbb{R}^2 : x_1 = 0\} \quad (17)$$

Additionally, notice that $x_2(t)$ possesses decoupled and marginally stable controllable linear dynamics. On the other hand, if $u_1(t)$ is held constant and zero, then $\mathbf{x}(t)$ dynamics becomes marginally stable linear controllable with respect to $u_2(t)$. Consequently, the complete nonlinear system $f(\mathbf{x}, \mathbf{u})$ is controllable.

To further illustrate the aforementioned controllability property, consider an initial state $\mathbf{x}_i \in \mathbb{U}$ and a desired final state $\mathbf{x}_f \in \mathbb{U}$ (see figure 5). Notice that the straight trajectory between them is unattainable since u_1 actuation in \dot{x}_1 is canceled by $x_1 = 0$, and $x_2 \neq 0$ drives \mathbf{x} out of \mathbb{U} imperatively. A feasible trajectory is to let \mathbf{x} drive away from \mathbb{U} , and steer \mathbf{x} back to \mathbb{U} at appropriate reentry points by means of the fully actuated \mathbf{u} (which is possible in arbitrary trajectories in \mathbb{A}).




Fig. 5. An example of a feasible trajectory starting and ending at the underactuated region \mathbb{U} .

3.2 LQR design

Consider the problem of tracking the rectilinear equilibrium trajectory $\mathbf{x}_d(s) : [0, 1] \rightarrow \mathbb{R}^2$ given by

$$\mathbf{x}_d(s) = \begin{pmatrix} x_i \\ 0 \end{pmatrix} + \begin{pmatrix} x_f - x_i \\ 0 \end{pmatrix} s \quad (18)$$

with $x_i > 0$ and $x_f < 0$. Accordingly, we apply RoA-GS with LQR as the chosen linear control design technique for all controllers K_i due to its simple methodology (numerous K_i might be required depending on the application) and adequate stability margins (Safonov and Athans (1977)). Each LQR design iteration yields a suboptimal linear control policy $\Delta \mathbf{u}_i^*(t) = -K_i \Delta \mathbf{x}_i$ that locally minimizes the cost

$$J(\Delta \mathbf{x}_i(t_i), \Delta \mathbf{u}_i) = \int_{t_i}^{\infty} (\Delta \mathbf{x}_i^T Q_i \Delta \mathbf{x}_i + \Delta \mathbf{u}_i^T R_i \Delta \mathbf{u}_i) dt \quad (19)$$

where $\Delta \mathbf{x}_i(t_i) \in \mathbb{R}^2$, $Q_i, R_i \in \mathbb{R}^{2 \times 2}$, are, respectively, initial state of the K_i controller at the switching moment instant t_i , positive semi-definite state penalty and positive definite actuator penalty matrices. For instance, let us consider

$$Q_i = \begin{cases} \begin{bmatrix} 1 & -1 \\ -1 & \frac{|x_{1,i}|}{2} \\ \frac{|x_{1,i}|}{2} & x_{1,i}^2 \end{bmatrix}, & \text{if } x_{1,i} \neq 0 \\ \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, & \text{if } x_{1,i} = 0 \end{cases} \quad (20)$$

and

$$R_i = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (21)$$

The perhaps peculiar choice of Q_i for $x_{1,i} \neq 0$ is justified by the following identity

$$\begin{aligned} \Delta \mathbf{x}_i^T Q_i \Delta \mathbf{x}_i &= \begin{pmatrix} \Delta x_{1,i} \\ \Delta x_{2,i} \end{pmatrix}^T \begin{bmatrix} 1 & -1 \\ -1 & \frac{|x_{1,i}|}{2} \\ \frac{|x_{1,i}|}{2} & x_{1,i}^2 \end{bmatrix} \begin{pmatrix} \Delta x_{1,i} \\ \Delta x_{2,i} \end{pmatrix} = \\ &= \begin{pmatrix} \Delta x_{1,i} \\ \Delta x_{2,i} \\ \frac{|x_{1,i}|}{2} \end{pmatrix}^T \begin{bmatrix} 1 & -1 \\ -1 & 2 \end{bmatrix} \begin{pmatrix} \Delta x_{1,i} \\ \Delta x_{2,i} \\ \frac{|x_{1,i}|}{2} \end{pmatrix} \quad (22) \end{aligned}$$

which illustrates larger allowance of $\Delta x_{2,i}$ usage in view of distant (from the origin) operating points $x_{1,i}$.

In view of the chosen weights, by means of the Hamilton-Jacobi-Bellman equation, one can show (Anderson and Moore (1990)) that the control policy

$$\Delta \mathbf{u}_i^*(t) = \arg \min_{\Delta \mathbf{u}_i} J(\Delta \mathbf{x}_i(t_i), \Delta \mathbf{u}_i) \quad (23)$$

is independent of $\Delta \mathbf{x}_i(t_i)$ and is given by

$$\Delta \mathbf{u}_i^*(t) = -B_i^T P_i \Delta \mathbf{x}_i(t) \quad (24)$$

where P_i is the unique semi-positive definite solution of the algebraic Riccati equation (ARE)

$$P_i A_i + A_i^T P_i - P_i B_i B_i^T P_i + Q_i = 0 \quad (25)$$

if (A_i, B_i) is stabilizable and $(A_i, Q_i^{1/2})$ is detectable. The solution P_i is normally computed numerically (Laub (1978); Petkov et al. (1991)), but our reachability assessment study calls for an analytical expression. Accordingly, we denote P_i by

$$P_i = \begin{bmatrix} \alpha & \beta \\ \beta & \gamma \end{bmatrix} \quad (26)$$

and substitute it back in (25) to obtain the following nonlinear algebraic system of equations

$$\begin{cases} 1 - \beta^2 - \alpha^2 x_{1,i}^2 = 0 \\ \alpha - \frac{1}{|x_{1,i}|} - \beta\gamma - \alpha\beta x_{1,i}^2 = 0 \\ -\beta^2 x_{1,i}^2 + 2\beta - \gamma^2 + \frac{2}{x_{1,i}^2} = 0 \end{cases} \quad (27)$$

by additionally recalling that A_i and B_i (equation 13) over the trajectory described by (18) reduce to

$$A_i = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \quad B_i = \begin{bmatrix} x_{1,i} & 0 \\ 0 & 1 \end{bmatrix} \quad (28)$$

We invite the reader to check that the following P_i is a solution to (27)

$$P_i = \begin{bmatrix} \frac{1}{|x_{1,i}|} & 0 \\ 0 & \frac{\sqrt{2}}{|x_{1,i}|} \end{bmatrix} \quad (29)$$

which is positive-definite and thus the unique solution of (27). Substitution of (28) and (29) into (24) yields

$$\Delta \mathbf{u}_i(t) = - \begin{bmatrix} \frac{x_{1,i}}{|x_{1,i}|} & 0 \\ 0 & \frac{\sqrt{2}}{|x_{1,i}|} \end{bmatrix} \Delta \mathbf{x}_i(t) \quad (30)$$

We omitted the (rather simple to solve) case $x_{1,i} = 0$ since the RoA-GS algorithm never reaches $\mathbf{x}_i = \mathbf{0}$, as we shall prove next.

3.3 Region-of-attraction computation

Due to its complexity regions-of-attraction are often computed numerically. However, in the present study, analytical intervals-of-attraction in the direction of the desired trajectory are provided to prove RoA-GS infeasibility. Initially, since we are interested in $\Delta \mathbf{x}_i$ convergence, we rewrite (12) by substituting (4), (5) and (30) to obtain

$$\Delta \dot{x}_{1,i} = -|x_{1,i}| \Delta x_i - \frac{x_{1,i}}{|x_{1,i}|} \Delta x_{1,i}^2 + \Delta x_{2,i} \quad (31)$$

and

$$\Delta \dot{x}_{2,i} = -\frac{\sqrt{2}}{|x_{1,i}|} \Delta x_{2,i} \quad (32)$$

The $\Delta x_{2,i}$ component has decoupled linear dynamics that yields the well-known exponential decay given by

$$\Delta x_{2,i}(t) = \Delta x_{2,i}(t_i) e^{\frac{-\sqrt{2}}{|x_{1,i}|}(t-t_i)} \quad (33)$$

that clearly converges to zero when $t \rightarrow \infty$ for all $x_{i,1} \in \mathbb{R}^*$. Substitution of (33) in (31) yields

$$\Delta \dot{x}_{1,i} = -|x_{1,i}| \Delta x_i - \frac{x_{1,i}}{|x_{1,i}|} \Delta x_{1,i}^2 + \Delta x_{2,i}(t_i) e^{\frac{-\sqrt{2}}{|x_{1,i}|}(t-t_i)} \quad (34)$$

Consider the following subset $\Pi \subset \mathbb{R}^2$ of initial conditions

$$\Pi = \{\Delta \mathbf{x}_i(t_i) \in \mathbb{R}^2 : \Delta x_{2,i}(t_i) = 0\} \quad (35)$$

For $\Delta \mathbf{x}_i(t_i) \in \Pi$, equation 34 reduces to

$$\Delta \dot{x}_{1,i} = -|x_{1,i}| \Delta x_i - \frac{x_{1,i}}{|x_{1,i}|} \Delta x_{1,i}^2 \quad (36)$$

which is nonlinear time-invariant. Its equilibrium points p_1 and p_2 are solutions of the quadratic equation

$$-|x_{1,i}| p_i - \frac{x_{1,i}}{|x_{1,i}|} p_i^2 = 0 \quad (37)$$

i.e.,

$$p_1 = 0, \quad p_2 = -x_{i,1} \quad (38)$$

Its associated phase portrait is illustrated in figures 6 and 7 (be aware of the change of variables from $\Delta x_{1,i}$ to $x_{1,i}$). Interestingly, the phase portrait reveals that $\mathbf{x} = \mathbf{0}$ is not stable thus not included in any R_i for any controller K_i designed at any operation point $x_{1,i} \neq 0$. Therefore the sequence of regions-of-attraction R_i never cross the origin as figure 8 illustrates.

From another point of view, given a K_i controller centered in $x_{i,1} > 0$, we conclude that the next controller, namely K_{i+1} , is scheduled at

$$x_{i+1,1} = (1 - \mu)x_{i,1} \quad (39)$$

such that

$$x_{i+N,1} = (1 - \mu)^N x_{i,1} > 0 \quad (40)$$




Fig. 6. Phase portrait of closed-loop $f(\mathbf{x}, -K_i \Delta \mathbf{x}_i)$ when $\Delta \mathbf{x}_i(t_i) \in \Pi_+$.




Fig. 7. Phase portrait of closed-loop $f(\mathbf{x}, -K_i \Delta \mathbf{x}_i)$ when $\Delta \mathbf{x}_i(t_i) \in \Pi_-$.




Fig. 8. Sequence of RoA-GS controllers K_i with associated R_i fails to cross the origin and reach x_f .

and never reaches x_f since $x_f < 0$. RoA-GS generates an infinite sequence of controllers K_i with operating points \mathbf{x}_i asymptotically converging to $\mathbf{0}$ but never crossing it. \square

Finally, consider the scenario of gain scheduling design without RoA for tracking the same trajectory with a uniform distribution per unit length ρ of LQR controllers used instead. By the foregoing development we conclude that no ρ renders tracking possible. Therefore, in general, discrete gain scheduling is inappropriate for this example.

4. CONCLUSION

We prove by means of an analytical counterexample that global controllability does not imply successful RoA-GS

nor gain scheduling quasi-static trajectory tracking. Future work might include sufficient conditions for RoA-GS trajectory tracking.

ACKNOWLEDGEMENTS

The authors support free education for everyone, everywhere, and highly praise Prof. Russ Tedrake for his beautiful course on underactuated robotics at edX.org.

REFERENCES

- Anderson, B.D.O. and Moore, J.B. (1990). *Optimal control, linear quadratic methods*. Dover.
- Burrige, R.R., Rizzi, A.A., and Koditschek, D.E. (1999). Sequential composition of dynamically dexterous robot behaviors. *International Journal of Robotics Research*, 18(6), 534–555.
- Chiang, H.D., Hirsh, M.W., and Wu, F.F. (1988). Stability regions of nonlinear autonomous dynamical systems. *IEEE Transactions on Automatic Control*, 33(1), 16–27.
- Gahinet, P. and Apkarian, P. (2011). Decentralized and fixed-structure H_∞ control in MATLAB. In *Proc. IEEE Conference on Decision and Control*, 8205–8210.
- Johansen, T.A. (2000). Computation of Lyapunov functions for smooth nonlinear systems using convex optimization. *Automatica*, 36(11), 1617–1626.
- Laub, A.J. (1978). A Schur method for solving algebraic Riccati equations. Technical Report 859, Massachusetts Institute of Technology, Laboratory for Information and Decision systems, LIDS.
- Lustosa, L.R., Defay, F., and Moschetta, J.M. (2015). Longitudinal study of a tilt-body vehicle: modeling, control and stability analysis. In *Proc. Int. Conf. on Unmanned Aircraft Systems (ICUAS)*, 816–824.
- Moore, J., Cory, R., and Tedrake, R. (2014). Robust post-stall perching with a simple fixed-wing glider using LQR-trees. *Bioinspiration and Biomimetics*, 9(2).
- Parrilo, P.A. (2000). *Structured semidefinite programs and semialgebraic geometry methods in roburobust and optimization*. Ph.D. thesis, California Institute of Technology.
- Petkov, P., Christov, N., and Konstantinov, M. (1991). *Computational methods for linear control systems*. New York: Prentice.
- Safonov, M.G. and Athans, M. (1977). Gain and phase margins for multiloop LQG regulators. *IEEE Transactions on Automatic Control*.
- Tan, W. and Packard, A. (2008). Stability region analysis using polynomial and composite polynomial Lyapunov functions and sum-of-squares programming. *IEEE Transactions on Automatic Control*, 53(2), 565–571.
- Tedrake, R., Manchester, I.R., Tobenkin, M.M., and Roberts, J.W. (2010). LQR-trees: feedback motion planning via sums of squares verification. *International Journal of Robotics Research*, 29, 1038–1052.
- Topcu, U. and Packard, A. (2009). Local stability analysis for uncertain nonlinear systems. *IEEE Transactions on Automatic Control*, 54(5), 1042–1047.
- Topcu, U., Packard, A., Seiler, P., and Balas, G.J. (2010). Robust region-of-attraction estimation. *IEEE Transactions on Automatic Control*, 55(1), 137–142.