

# LQR-Trees: Feedback Motion Planning via Sums-of-Squares Verification\*

Russ Tedrake, Ian R. Manchester, Mark Tobenkin, John W. Roberts  
Computer Science and Artificial Intelligence Lab  
Massachusetts Institute of Technology  
Cambridge, MA, 02139  
Email: {russt,irm,mmt,jwr}@mit.edu

February 27, 2010

## Abstract

Advances in the direct computation of Lyapunov functions using convex optimization make it possible to efficiently evaluate regions of attraction for smooth nonlinear systems. Here we present a feedback motion planning algorithm which uses rigorously computed stability regions to build a sparse tree of LQR-stabilized trajectories. The region of attraction of this nonlinear feedback policy “probabilistically covers” the entire controllable subset of the state space, verifying that all initial conditions that are capable of reaching the goal will reach the goal. We numerically investigate the properties of this systematic nonlinear feedback design algorithm on simple nonlinear systems, prove the property of probabilistic coverage, and discuss extensions and implementation details of the basic algorithm.

## 1 Introduction

This paper aims to build on advances from systems theory and from randomized motion planning to design efficient and general algorithms for nonlinear feedback control synthesis in constrained nonlinear systems. Specifically, the controls verification community has developed a number of efficient algorithms for direct computation of Lyapunov functions for smooth nonlinear systems using convex optimization[12, 25]. These tools can be used to estimate conservative regions of attraction around a planned and stabilized trajectory for even very complicated dynamical systems. Combinations of this formal control verification with ideas from randomized motion planning open up a number of interesting possibilities

---

\*A preliminary version of this paper was presented at the 2009 Robotics: Science and Systems Conference [35].

for algorithm development. In particular, we present the LQR-Tree algorithm, which uses locally optimal linear feedback control policies to stabilize planned trajectories computed by local trajectory optimizers, and verification based on a sums-of-squares method to estimate the local regions of stability.

Specifically, the LQR-Tree algorithm operates by growing a tree of stabilized and verified trajectories backwards from a goal state. At each step of the algorithm, a random sample is drawn from a bounded distribution over state space. If that point is already inside the region of attraction of one of the trajectories, then it is discarded. If it is outside, then a local trajectory optimizer attempts to find a new trajectory which connects this random sample to the tree (and therefore to the goal). This new trajectory is stabilized and verified, and then the process repeats. We demonstrate that this algorithm has an important notion of completeness - it provably covers the controllable subset of state space with a stabilizing controller. We call this property “probabilistic feedback coverage”.

The aim of this work is to generate a class of algorithms capable of computing verified feedback policies for complicated<sup>1</sup> nonlinear systems which are not amenable to feedback linearization and with dimensionality beyond what might be accessible to grid-based algorithms like dynamic programming. The use of local trajectory optimizers and local feedback stabilization scales well to higher-dimensions, and reasoning about the feedback regions allows the algorithm to cover a bounded subset of state space which is controllable to the a goal with a relatively sparse set of trajectories. In addition, the algorithms operate directly on the continuous state and action spaces and perform verification algebraically, and thus are not subject to the pitfalls of discretization which limit accuracy and scalability. By considering feedback during the planning process, the resulting plans are robust to disturbances and quite suitable for implementation on real robots.

The following sections introduce the details of the LQR-Tree algorithm. After a review of related work in Section 2, Section 3 starts by describing the design, stabilization, and local verification of a single feasible trajectory. Section 4 then introduces the method for growing the tree backwards from the goal, yielding a nonlinear feedback policy over the state space. Section 5 demonstrates the algorithm on a few model robotic control problems. Section 6 contains the proof of probabilistic feedback coverage. In the final sections, we discuss a few of the more subtle implementation details, and discuss possibilities for future extensions.

## 2 Background

### 2.1 Feedback motion planning

The last decade has seen rapid progress in motion planning algorithms, spurred by advances in sample-based, randomized algorithms like the probabilistic

---

<sup>1</sup>E.g., systems subject to underactuation constraints, input saturations, and/or other kinodynamic constraints




Figure 1: Cartoon of motion planning with funnels in the spirit of [5].

roadmaps (PRMs) [13] and the rapidly-exploring random trees (RRTs)[16]. These algorithms have demonstrated the ability to produce feasible open-loop trajectories from an initial condition to a goal state in impressively high-dimensional and convincingly non-convex search problems [15, 8, 32].

For implementation on real robots, open-loop trajectories generated by a motion planning system are typically stabilized by a trajectory stabilizing feedback system<sup>2</sup>. While this decoupled approach works for most problems, it is possible that a planned trajectory is not stabilizable, or very costly to stabilize compared to other, more desirable trajectories. *Feedback-motion-planning algorithms*, which explicitly consider the feedback stabilization during the planning process, can avoid this pitfall, and as we will see, can potentially use a local understanding of the capabilities of the feedback system to guide and optimize the search in a continuous state space.

Mason popularized the metaphor of a “funnel” (see Figure 1) for a feedback policy which collapses a large set of initial conditions into a smaller set of final conditions [18]. In [26] this concept was formalized as “nested Lyapunov functions” and used to analyze the time convergence of nonlinear systems in which the equations of dynamics could be represented as Fourier series. Burrige, Rizzi, and Koditschek then painted a beautiful picture of feedback motion planning as a sequential composition of locally valid feedback policies, or funnels, which take a broad set of initial conditions to a goal region [5]. At the time, the weakness of this approach was the difficulty in computing, or estimating by trial-and-error, the region of applicability - the mouth of the funnel, or preimage - for each local controller in a nonlinear system. Consequently, besides the particular solution in [5], these ideas have mostly been limited to reasoning about vector-fields on systems without dynamics [17].

<sup>2</sup>Note that an increasingly plausible alternative is real-time, dynamic re-planning.

## 2.2 Direct computation of Lyapunov functions

Burridge et al. also pointed out the strong connection between Lyapunov functions and these motion planning funnels [5]. For a dynamical system,  $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$  with  $\mathbf{f}(0) = 0$ , a Lyapunov function is a differentiable positive-definite output function,  $V(\mathbf{x})$ , for which  $\frac{\partial V}{\partial \mathbf{x}} \mathbf{f}(\mathbf{x})$  is negative-definite. If these conditions are met over some ball in state space,  $B_r$ , containing the origin, then the origin is asymptotically stable. The ball,  $B_r$ , can then be interpreted as the preimage of the funnel. Lyapunov functions have played a central role in nonlinear control theory, since they allow verification of a system’s stability without ever computing solutions of the system [14]. Indeed, if a dynamical system is asymptotically stable within a region of attraction, then there exists a Lyapunov function certifying that fact. However, computing such a Lyapunov function for nonlinear systems is a challenging problem, and has been the focus of much of the history of nonlinear control theory.

A number of computational approaches to computing Lyapunov functions for nonlinear systems have emerged in the last decade, often based on convex optimization (e.g., [12, 25]). One of these techniques, which forms the basis of the results reported here, is based on the realization that one can check the uniform positive-definiteness of a polynomial expression using a *sums-of-squares* (SOS) optimization program [25].

The main idea of SOS is quite straightforward. Suppose we want to check if the following polynomial inequality is true:

$$x^4 + 2x^3 + 3x^2 - 2x + 2 \geq 0 \quad \forall x \in \mathbb{R}. \quad (1)$$

One way to verify that this is in fact true is to notice that

$$x^4 + 2x^3 + 3x^2 - 2x + 2 = \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}^T \begin{bmatrix} 2 & -1 & 0 \\ -1 & 3 & 1 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}$$

which can be checked by matching coefficients. Now, the matrix on the right is positive definite (its eigenvalues are approximately 3.88, 1.65, and 0.47) which implies that for any value of  $x$  the polynomial (1) is positive.

However, this is not the only way to represent (1) as a quadratic form in the “monomials”  $1, x, x^2$ . In general, we could have a matrix of the following form:

$$x^4 + 2x^3 + 3x^2 - 2x + 2 = \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}^T \begin{bmatrix} a & b & c \\ b & d & e \\ c & e & f \end{bmatrix} \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}$$

where by matching coefficients we have  $a = 2, 2b = -2, d + 2c = 3, 2e = 2, f = 1$ . Notice that one can trade off the affine constraint between  $d$  and  $c$  and still represent the same polynomial. This same approach can be extended to multivariable polynomials, and sums-of-squares programming refers to checking non-negativity of a polynomial by searching for a positive semidefinite matrix

that satisfies a set of affine constraints that match coefficients. Furthermore, one can perform this optimization when the polynomial coefficients are free (decision) parameters. This results in a semidefinite program and can be solved efficiently using interior point methods [22].

Freely available libraries including SOSTOOLS [28] and the Systems Polynomial Optimization Toolbox [19] make it quite accessible to perform these computations in MATLAB. As we will see, the ability to check uniform positive (or negative) definiteness will offer the ability to verify candidate Lyapunov functions over a region of state space for smooth nonlinear systems.

### 2.3 Feedback Synthesis by Sums-of-Squares Optimization

The compatibility of sums-of-squares methods with semidefinite programming, and the associated interior-point method solvers, has had a notable impact on stability verification (see, e.g., [25, 30, 24, 29, 34, 37, 23] and many others). However, an open and active area of research is focused on extending these approaches to controller synthesis. Given a system,  $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}$ , we wish to simultaneously generate a feedback,  $\mathbf{u} = \pi(\mathbf{x})$  and a Lyapunov function,  $V(\mathbf{x})$ , such that  $\frac{\partial V}{\partial \mathbf{x}} [\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\pi(\mathbf{x})]$  is negative definite. This has proven to be a difficult problem because the set of  $V(\mathbf{x})$  and  $\pi(\mathbf{x})$  satisfying these conditions may not be convex or even connected [30]. To combat this, a number of solutions have been proposed based on density functions [27], and iterative algorithms for model-predictive control design [7].

This paper takes a different approach to control synthesis. Rather than designing a nonlinear feedback controller directly in the optimization, we rely on classical linear quadratic regulator (LQR) synthesis [1] to design a series of locally-valid controllers and compose these controllers using feedback motion planning. This has the advantage that it may work for hopelessly non-convex control problems, such as navigation through a complicated obstacle field, robotic manipulation, or legged locomotion over rough terrain, where the randomized motion planning algorithms have demonstrated success. We prove in Section 6 that, under some mild assumptions, the feedback synthesis will eventually stabilize effectively every point that is controllable to the goal for a class of smooth nonlinear system.

### 2.4 Other related work

In other related work, Atkenson and Stephens [2] used local trajectory optimizers and LQR stabilizers with randomized starting points to try to cover the space, with the hope of verifying global optimality (in the infinite resolution case) by having consistent locally quadratic estimates of the value function on neighboring trajectories. The conditions for adding nodes in that work were based on the magnitude of the value function (not the region of guaranteed stability). In the work described here, we sacrifice direct attempts at obtaining optimal feedback policies in favor of computing good-enough policies which probabilistically cover the controllable subset of state space with a region of

attraction. As a result, we have stronger guarantees of getting to the goal and considerably sparser collections of sample paths.

Region-of-attraction verification for polynomial systems has recently been explored as a search for Lyapunov functions using sums-of-squares techniques [34]. The optimization required is bilinear in the decision variables, and so optimization methods will only find a local solution. Simulation of the system dynamics can help find good initial guesses for the Lyapunov function [37]. Tomlin et al. have a series of work exploring verification and control of nonlinear systems using reachable set computations and hybrid automata[36]. In general this requires computing level sets of a Hamilton-Jacobi-Bellman-Isaacs partial differential equation, which must be approached numerically for all but the simplest systems[21]. A recent approach to control design which has some philosophical connections to the method we propose is that of “patchy Lyapunov functions”, in which a system is stabilized by covering the state space with an ordered family of locally valid control Lyapunov functions and stabilization ensured via a switching feedback control [10].

### 3 Linear Feedback Design and Verification

The LQR-Tree algorithm is based on the ability to efficiently design trajectories of the robot through state space, to stabilize those trajectories using linear quadratic regulator (LQR) feedback design, and to estimate the region of attraction of the feedback controller. In this section, we develop that procedure first for a stabilizable point and then for a single stabilizable trajectory.

#### 3.1 Stabilizing a Goal State

Consider a smooth nonlinear system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}). \quad (2)$$

where  $\mathbf{x} \in \mathbb{R}^n$  is the state of the system and  $\mathbf{u} \in \mathbb{R}^m$  is the control input. We first examine stabilizing a goal-state with an infinite horizon LQR controller, and approximating the closed loop region of attraction. Consider a goal state,  $\mathbf{x}_G$ , with  $\mathbf{u}_G$  defined so that  $\mathbf{f}(\mathbf{x}_G, \mathbf{u}_G) = 0$ . Define

$$\bar{\mathbf{x}} = \mathbf{x} - \mathbf{x}_G, \quad \bar{\mathbf{u}} = \mathbf{u} - \mathbf{u}_G. \quad (3)$$

Now, linearize the system around  $(\mathbf{x}_G, \mathbf{u}_G)$  to yield the dynamics:

$$\dot{\bar{\mathbf{x}}}(t) \approx \mathbf{A}\bar{\mathbf{x}}(t) + \mathbf{B}\bar{\mathbf{u}}(t). \quad (4)$$

We assume that this linearization of  $\mathbf{f}$  about the goal is controllable. Define the quadratic regulator cost-to-go function as

$$J(\bar{\mathbf{x}}') = \int_0^\infty [\bar{\mathbf{x}}^T(t)\mathbf{Q}\bar{\mathbf{x}}(t) + \bar{\mathbf{u}}^T(t)\mathbf{R}\bar{\mathbf{u}}(t)] dt, \quad (5)$$

$$\mathbf{Q} = \mathbf{Q}^T \geq 0, \mathbf{R} = \mathbf{R}^T > 0, \bar{\mathbf{x}}(0) = \bar{\mathbf{x}}'.$$

The optimal cost-to-go function for the linear system is given by

$$J^*(\bar{\mathbf{x}}) = \bar{\mathbf{x}}^T \mathbf{S} \bar{\mathbf{x}}, \quad (6)$$

where  $\mathbf{S}$  is the positive-definite solution to the equation:

$$0 = \mathbf{Q} - \mathbf{SBR}^{-1}\mathbf{B}^T\mathbf{S} + \mathbf{SA} + \mathbf{A}^T\mathbf{S} \quad (7)$$

(given by the MATLAB `lqr` function). The optimal feedback policy for the linear system is given by

$$\bar{\mathbf{u}}^* = -\mathbf{R}^{-1}\mathbf{B}^T\mathbf{S}\bar{\mathbf{x}} = -\mathbf{K}\bar{\mathbf{x}}. \quad (8)$$

### 3.1.1 Time-Invariant LQR Verification

We acquire an estimate of the region of attraction of the LQR controller by constructing a function,  $V(\mathbf{x})$ , which is a valid Lyapunov function for the nonlinear system over some sub-level set:

$$\mathcal{B}_G(\rho) = \{\mathbf{x} | 0 \leq V(\mathbf{x}) \leq \rho\}. \quad (9)$$

To demonstrate asymptotic stability, we require:

- $V(\mathbf{x})$  is positive definite in  $\mathcal{B}_G(\rho)$ ,
- $\dot{V}(\mathbf{x})$  is negative definite in  $\mathcal{B}_G(\rho)$ .

Such a function demonstrates the goal state is stabilized, and furthermore, all initial conditions in  $\mathcal{B}_G(\rho)$  will converge to  $\mathbf{x}_G$  [33].

Here we use  $V(\mathbf{x}) = J^*(\bar{\mathbf{x}})$ , the linear optimal cost-to-go function, which is a Lyapunov function for the linear system and hence serves well as a local Lyapunov function candidate for the nonlinear system. By construction, this choice satisfies the positive definite constraint on  $V(\mathbf{x})$ . For the second condition, first observe that

$$\dot{V}(\mathbf{x}) = \dot{J}(\bar{\mathbf{x}}) = 2\bar{\mathbf{x}}^T \mathbf{Sf}(\mathbf{x}_G + \bar{\mathbf{x}}, \mathbf{u}_G - \mathbf{K}\bar{\mathbf{x}}). \quad (10)$$

SOS programs certify the *global* non-negativity (or non-positivity) of a polynomial, however our condition on  $\dot{J}^*$  is:

$$\dot{J}^*(\bar{\mathbf{x}}) < 0 \quad \forall \bar{\mathbf{x}} \neq 0 \in \mathcal{B}_G(\rho), \quad \dot{J}^*(\mathbf{0}) = 0. \quad (11)$$

We take two steps to make our condition compatible with SOS, which are common in the SOS literature [25]. First, we transform the negativity constraint into a non-positivity constraint by requiring:

$$\dot{J}^*(\bar{\mathbf{x}}) \leq -\epsilon \|\bar{\mathbf{x}}\|_2^2 \quad \forall \bar{\mathbf{x}} \in \mathcal{B}_G(\rho), \quad (12)$$

for some small positive constant  $\epsilon$ . Next, we remove the restriction of  $\bar{\mathbf{x}} \in \mathcal{B}_G(\rho)$  using a Lagrange multiplier. Take any non-negative polynomial  $h(\bar{\mathbf{x}})$ . We examine the inequality:

$$\dot{J}^*(\bar{\mathbf{x}}) + h(\bar{\mathbf{x}})(\rho - J^*(\bar{\mathbf{x}})) \leq -\epsilon \|\bar{\mathbf{x}}\|_2^2. \quad (13)$$

Note that the sign of  $(\rho - J^*(\bar{\mathbf{x}}))$  is positive for all  $\bar{\mathbf{x}} \in \mathcal{B}_G(\rho)$  by definition, and negative outside this ball. As a result, if (13) holds for any  $h(\bar{\mathbf{x}}) \geq 0$  it immediately implies (12). This allows us to search for such an  $h(\bar{\mathbf{x}})$ :

$$\begin{aligned} & \text{find} && h(\bar{\mathbf{x}}), && (14) \\ & \text{subject to} && \hat{J}^*(\bar{\mathbf{x}}) + h(\bar{\mathbf{x}}) (\rho - J^*(\bar{\mathbf{x}})) \leq -\epsilon \|\bar{\mathbf{x}}\|_2^2, \\ & && h(\bar{\mathbf{x}}) \geq 0. \end{aligned}$$

The polynomial  $h$  must be of sufficiently large order to counteract higher order terms in  $\hat{J}^*(\bar{\mathbf{x}})$ . We will make frequent use of Lagrange multipliers such as  $h(\cdot)$  to handle regional non-negativity conditions throughout this paper. In the case where  $\mathbf{f}$  is polynomial in  $\mathbf{x}$  and  $\mathbf{u}$ , demonstrating a feasible  $h(\cdot)$  is sufficient to provide an exact verification of our desired condition, though it is important to note that not all non-negative polynomials are SOS[25].

For the case where  $\mathbf{f}$  is not polynomial, we can approximate this condition using a Taylor expansion of  $\mathbf{f}$ , with order  $> 1$ . Let us denote the closed-loop dynamics as

$$\mathbf{f}^{(\text{cl})}(\mathbf{x}) = \mathbf{f}(\mathbf{x}, \mathbf{u}_G - \mathbf{K}(\mathbf{x} - \mathbf{x}_G)). \quad (15)$$

We denote the truncated Taylor expansion of  $\mathbf{f}^{(\text{cl})}$  to order  $N_f > 1$ :

$$\mathbf{f}^{(\text{cl})}(\mathbf{x}_G + \bar{\mathbf{x}}) \approx \hat{\mathbf{f}}^{(\text{cl})}(\bar{\mathbf{x}}). \quad (16)$$

For example, when  $N_f = 2$ , the  $k$ -th component of the closed-loop dynamics is approximated by:

$$\hat{\mathbf{f}}_k^{(\text{cl})}(\bar{\mathbf{x}}) = \sum_i \bar{x}_i \left[ \frac{\partial \mathbf{f}_k^{(\text{cl})}(\mathbf{x})}{\partial x_i} \right]_{\mathbf{x}=\mathbf{x}_G} + \frac{1}{2} \sum_{i,j} \bar{x}_i \bar{x}_j \left[ \frac{\partial^2 \mathbf{f}_k^{(\text{cl})}(\mathbf{x})}{\partial x_i \partial x_j} \right]_{\mathbf{x}=\mathbf{x}_G}, \quad (17)$$

where we also have that  $\mathbf{f}_k^{(\text{cl})}(\mathbf{x}_G) = 0$ . When the polynomial approximation of the dynamics is used, we replace the  $\hat{J}^*(\bar{\mathbf{x}})$  in the SOS program above with  $\hat{\hat{J}}^*(\bar{\mathbf{x}})$ :

$$\hat{\hat{J}}(\bar{\mathbf{x}}) = 2\bar{\mathbf{x}}^T \mathbf{S} \hat{\mathbf{f}}^{(\text{cl})}(\bar{\mathbf{x}}). \quad (18)$$

Figure 2 presents a numerical exploration of the estimated region of attraction for the simple pendulum. For increasing values of  $\rho$ , the maximum value of  $\hat{J}^*(\bar{\mathbf{x}})$  in  $\mathcal{B}_G$  is plotted. Also plotted is the maximum value achieved by  $\hat{\hat{J}}(\bar{\mathbf{x}})$  sampling the boundary of this ellipse, and the maximum value of the true (non-polynomial)  $J^*(\bar{\mathbf{x}})$  sampled similarly. One cannot generally rely on the Taylor expansion to be conservative as in Figure 2. Since we are interested in the Taylor expansion over a bounded region ( $\mathcal{B}_G(\rho)$ ), one can bound the error of the expansion if desired and ensure the region of attraction estimate is conservative, as in [6].






Figure 2: Polynomial verification of LTI feedback on the damped simple pendulum ( $m = 1$  kg,  $l = 0.5$  m,  $b = 0.1$  m<sup>2</sup> kg s<sup>-1</sup>,  $g = 9.8$  m s<sup>-2</sup>,  $\mathbf{Q} = \text{diag}([10, 1])$ ,  $\mathbf{R} = 15$ ,  $N_f = 3$ ,  $N_m = 2$ )

Finally, we formulate a convex optimization to find the largest region  $\mathcal{B}_G(\rho)$  for which we can verify the second condition is also satisfied:

$$\begin{aligned}
 & \text{maximize} && \rho, && (19) \\
 & \text{subject to} && \hat{J}^*(\bar{\mathbf{x}}) + h(\bar{\mathbf{x}}) \left( \rho - \hat{J}^*(\bar{\mathbf{x}}) \right) \leq -\epsilon \|\bar{\mathbf{x}}\|_2^2, \\
 & && \rho > 0, \\
 & && h(\bar{\mathbf{x}}) \geq 0.
 \end{aligned}$$

This optimization can be performed via a line-search on  $\rho$ . At each step of the line-search the SOS program (14) is executed. If the program is feasible then  $\rho$  can be increased, otherwise it must be decreased. Thus, the maximal  $\rho$  which can be verified is approached by a sequence of SOS programs.

### 3.2 Trajectory Optimization

In order to increase the set of states which can reach the goal beyond the time-invariant LQR design, we will design and stabilize a feasible trajectory for the system with an initial condition outside of the verified region of attraction of the time-invariant controller. Trajectory design can be accomplished readily by trajectory optimization algorithms including shooting methods, multiple shooting methods, and direct collocation methods [3], all of which are quite mature, and can perform well on even very complicated nonlinear systems. Given the nonlinear system  $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$ , we solve for a feasible trajectory of the system  $\mathbf{x}_0(t)$ ,  $\mathbf{u}_0(t)$  over the finite time interval  $[t_0, t_f]$  which (locally) optimizes a cost function of the form

$$J = \int_{t_0}^{t_f} [1 + \mathbf{u}_0^T \mathbf{R} \mathbf{u}_0] dt, \quad (20)$$

often subject to a final value constraint (for instance,  $\mathbf{x}_0(t_f) = \mathbf{x}_G$ ). Throughout the paper, we perform trajectory optimization with a multiple shooting method with the initial trajectory specified by a random duration “tape” of random

control actions; this constrained nonlinear optimization is implemented with sequential quadratic programming (SQP), using SNOPT [9].

### 3.3 Time-Varying LQR

Given a nominal trajectory,  $\mathbf{x}_0(t)$ ,  $\mathbf{u}_0(t)$ , (a solution of equation 2) over a finite time interval,  $t \in [t_0, t_f]$ , we stabilize the trajectory using a time-varying LQR controller. Linearizing the system around the trajectory, we obtain:

$$\bar{\mathbf{x}}(t) = \mathbf{x}(t) - \mathbf{x}_0(t), \quad \bar{\mathbf{u}}(t) = \mathbf{u}(t) - \mathbf{u}_0(t), \quad (21)$$

$$\dot{\bar{\mathbf{x}}}(t) \approx \mathbf{A}(t)\bar{\mathbf{x}}(t) + \mathbf{B}(t)\bar{\mathbf{u}}(t). \quad (22)$$

and define a quadratic regulator (tracking) cost function:

$$J(\bar{\mathbf{x}}', t') = \bar{\mathbf{x}}^T(t_f)\mathbf{Q}_f\bar{\mathbf{x}}(t_f) + \int_{t'}^{t_f} [\bar{\mathbf{x}}^T(t)\mathbf{Q}\bar{\mathbf{x}}(t) + \bar{\mathbf{u}}^T(t)\mathbf{R}\bar{\mathbf{u}}(t)] dt, \quad (23)$$

$$\mathbf{Q}_f = \mathbf{Q}_f^T > \mathbf{0}, \mathbf{Q} = \mathbf{Q}^T \geq \mathbf{0}, \mathbf{R} = \mathbf{R}^T > \mathbf{0}, \bar{\mathbf{x}}(t') = \bar{\mathbf{x}}'.$$

In general,  $\mathbf{Q}$  and  $\mathbf{R}$  could easily be made a function of time as well. With time-varying dynamics, the resulting cost-to-go is time-varying. It can be shown that the optimal cost-to-go,  $J^*$ , is given by

$$J^*(\bar{\mathbf{x}}, t) = \bar{\mathbf{x}}^T\mathbf{S}(t)\bar{\mathbf{x}}, \quad \mathbf{S}(t) = \mathbf{S}^T(t) > \mathbf{0}. \quad (24)$$

where  $\mathbf{S}(t)$  is the solution to

$$-\dot{\mathbf{S}} = \mathbf{Q} - \mathbf{S}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^T\mathbf{S} + \mathbf{S}\mathbf{A} + \mathbf{A}^T\mathbf{S}, \quad \mathbf{S}(t_f) = \mathbf{Q}_f, \quad (25)$$

and the optimal feedback policy is given by

$$\bar{\mathbf{u}}^*(t) = -\mathbf{R}^{-1}\mathbf{B}^T(t)\mathbf{S}(t)\bar{\mathbf{x}}(t) = -\mathbf{K}(t)\bar{\mathbf{x}}(t). \quad (26)$$

#### 3.3.1 Time-Varying LQR Verification

In order to describe the stabilization of a finite-length nominal trajectory, rather than asymptotic stability, we specify a bounded goal region of state space,  $\mathcal{B}_f$ , and search for a time-varying region,  $\mathcal{B}(\rho(\cdot), t)$ , for which the closed-loop system obeys:

$$\mathbf{x}(t) \in \mathcal{B}(\rho(\cdot), t) \implies \mathbf{x}(t_f) \in \mathcal{B}_f \quad \forall t \in [t_0, t_f]. \quad (27)$$

We search for such a region in terms of a (now time-varying) positive definite, differentiable, radially unbounded function  $V(\mathbf{x}, t)$ . At a moment in time, the region is determined as a sublevel set:

$$\mathcal{B}(\rho(\cdot), t) = \{\mathbf{x} | 0 \leq V(\mathbf{x}, t) \leq \rho(t)\}. \quad (28)$$

Where  $\rho(t)$  is chosen to ensure (27) holds. We also define the goal region similarly:

$$\mathcal{B}_f = \{\mathbf{x} | 0 \leq V(\mathbf{x}, t_f) \leq \rho_f\}. \quad (29)$$

To ensure (27), we choose the function  $\rho(t) : [t_0, t_f] \mapsto \mathbb{R}^+$  with the following conservative properties:

- $\rho(t_f) \leq \rho_f$ .
- $\rho(t)$  is discontinuous at finitely many points,  $\tau_1, \dots, \tau_M$ , and:

$$\lim_{t \uparrow \tau_m} \rho(t) \leq \lim_{t \downarrow \tau_m} \rho(t). \quad (30)$$

- The right derivative,  $\frac{d}{dt}_+ \rho(t)$ , exists and:

$$V(\mathbf{x}, t) = \rho(t) \implies \dot{V}(\mathbf{x}, t) \leq \frac{d}{dt}_+ \rho(t). \quad (31)$$

This ensures that the value  $V(\mathbf{x}, t)$  decreases faster along trajectories than our level-set  $\rho(t)$ . It is useful to note that given two such functions,  $\rho_1(t)$  and  $\rho_2(t)$ , their point-wise maximum  $\rho(t) = \max\{\rho_1(t), \rho_2(t)\}$  also satisfies these conditions, which verifies that  $\mathcal{B}(\rho_1(\cdot), t) \cup \mathcal{B}(\rho_2(\cdot), t)$  satisfies (27).

Again, we choose here to use  $V(\mathbf{x}, t) = J^*(\mathbf{x}, t)$ , which is positive definite as the LQR derivation ensures  $\mathbf{S}(t)$  is uniformly positive definite. In particular, this gives us:

$$\mathcal{B}_f = \{\mathbf{x} | 0 \leq \bar{\mathbf{x}}^T \mathbf{Q}_f \bar{\mathbf{x}} \leq \rho_f\}. \quad (32)$$

When  $\mathbf{x}(t_f) = \mathbf{x}_G$ , a stabilizable goal, one can choose  $\mathbf{Q}_f = \mathbf{S}_{(l_{ti})}$  and  $\rho_f = \rho_{(l_{ti})}$  to ensure trajectories land in an infinite-horizon LQR controller's region of attraction.

Now we have:

$$\hat{J}^*(\bar{\mathbf{x}}, t) = \bar{\mathbf{x}}^T \dot{\mathbf{S}}(t) \bar{\mathbf{x}} + 2\bar{\mathbf{x}}^T \mathbf{S}(t) \mathbf{f}(\hat{\mathbf{x}}_0(t) + \bar{\mathbf{x}}, \hat{\mathbf{u}}_0(t) - \mathbf{K}(t)\bar{\mathbf{x}}). \quad (33)$$

Here, even if  $\mathbf{f}$  is polynomial in  $\mathbf{x}$  and  $\mathbf{u}$  and the input tape  $\mathbf{u}_o(t)$  was polynomial, our analysis must make use of  $\mathbf{x}_0(t)$ ,  $\mathbf{S}(t)$ , and  $\mathbf{K}(t)$  which are the result of numerical integration (e.g., with `ode45` in Matlab). We will approximate this temporal dependence with (elementwise) piecewise polynomials using splines of order  $N_t$ , where  $N_t$  is often chosen to be 3 (cubic splines), with the knot points at the timesteps output by the variable step integration, which we denote  $t_0, t_1, \dots, t_N$ , with  $t_N = t_f$ , e.g.:

$$\forall t \in [t_k, t_k + 1], \quad S_{ij}(t) \approx \sum_{m=0}^{N_t} \alpha_{ijm} (t - t_k)^m = \hat{S}_{ij}(t), \quad (34)$$

$$\hat{J}^*(\bar{\mathbf{x}}, t) = \sum_i \sum_j \bar{x}_i \bar{x}_j \hat{S}_{ij}(t). \quad (35)$$

If  $\mathbf{f}$  is non-polynomial, we approximate  $\hat{J}^*(\bar{\mathbf{x}}, t)$  by first taking the Taylor expansion in  $\bar{\mathbf{x}}$  to arrive at a polynomial at each knot,  $t_k$ , and then interpolate these coefficients of the resulting polynomials as piecewise polynomials of time.

### 3.3.2 Optimizing for $\rho(t)$

In general, we would desire the largest possible  $\rho(t)$ , perhaps measured by the volume of state-space included in  $\mathcal{B}(\rho(\cdot), t)$ . However, searching over a suitably parameterized  $\rho(t)$  over all time-segments while verifying invariance appears intractable. Instead, we find a piecewise-linear<sup>3</sup>  $\rho(t)$  over the intervals  $[t_k, t_{k+1})$ :

$$\rho_k(t) = \beta_{1k}t + \beta_{0k}, \quad (36)$$

$$\rho(t) = \begin{cases} \rho_k(t), & \forall t \in [t_k, t_{k+1}), \\ \rho_f, & t = t_f. \end{cases} \quad (37)$$

which we construct backwards in time, starting with  $k = N - 1$ . For a given  $\rho_k$ , it is easy to test if  $\rho_k(t_{k+1}) = \beta_{1k}t_{k+1} + \beta_{0k} \leq \rho(t_{k+1})$ . The more complicated step of verification requires:

$$\hat{J}^*(\bar{\mathbf{x}}, t) = \rho_k(t) \implies \hat{J}^*(\bar{\mathbf{x}}, t) \leq \dot{\rho}_k(t) = \beta_{1k} \quad \forall t \in [t_k, t_{k+1}). \quad (38)$$

which can be tested by the SOS feasibility program:

$$\begin{aligned} & \text{find } h_1(\bar{\mathbf{x}}), h_2(\bar{\mathbf{x}}), h_3(\bar{\mathbf{x}}), \\ & \text{subject to } \hat{J}^*(\bar{\mathbf{x}}, t) - \dot{\rho}_k(t) + h_1(\bar{\mathbf{x}}, t)(\rho_k(t) - \hat{J}^*(\bar{\mathbf{x}}, t)) \\ & \quad + h_2(\bar{\mathbf{x}}, t)(t - t_k) + h_3(\bar{\mathbf{x}}, t)(t_{k+1} - t) \leq 0, \\ & \quad h_2(\bar{\mathbf{x}}, t) \geq 0, \\ & \quad h_3(\bar{\mathbf{x}}, t) \geq 0. \end{aligned} \quad (39)$$

Choosing the three Lagrange multipliers,  $h_1, h_2, h_3$  to be polynomials of sufficient order. Note that  $h_1(\bar{\mathbf{x}})$  is used to restrict our attention to an equality constraint, and thus does not need to be non-negative. This feasibility test enables a search for  $\beta_{.k}$ . The feasible region is generally non-convex, but there are some guarantees as to the local minima.

Our assumptions about the system dynamics ensure that, if the polynomial approximations of  $\mathbf{x}_0(t)$  and  $\hat{J}^*(\bar{\mathbf{x}}, t)$  are sufficiently close to true solutions of the dynamics and Riccati equations:

$$\exists \epsilon_k > 0 \text{ s.t. } \forall t \in [t_k, t_{k+1}) J^*(\bar{\mathbf{x}}, t) \leq \epsilon_k \implies \hat{J}^*(\bar{\mathbf{x}}, t) \leq 0. \quad (40)$$

We formulate the optimization:

$$\begin{aligned} & \underset{\alpha_k}{\text{maximize}} \quad \rho_k(t) = \alpha_k, \\ & \text{subject to} \quad \alpha_k \leq \rho(t_{k+1}), \\ & \quad \text{SOS Program (39)}. \end{aligned} \quad (41)$$

---

<sup>3</sup>Higher order polynomials are of course possible, but for small time-steps, piecewise linear appears to perform well.

This program is amenable to line-search, similar to the LTI region of attraction. A piecewise constant  $\rho(t)$  would be defined by these values, but has the disadvantage that  $\rho(t)$  will monotonically increasing.

This value obtained by (39),  $\alpha_k^*$ , can be used as an initial guess to optimizing  $\rho_k(t)$  over a class of polynomials. A particularly successful optimization strategy has been the following optimization:

$$\begin{aligned} & \underset{\beta_k}{\text{maximize}} && \rho_k(t_k) = \beta_{1k}t_k + \beta_{0k}, && (42) \\ & \text{subject to} && \rho_k(t_{k+1}) \leq \rho(t_{k+1}), \\ & && \text{SOS Program (39)}. \end{aligned}$$

Provides coefficients  $\beta^*$  such that:

$$\rho_k(t) = \max \{ \alpha_k^*, \beta_{1k}t + \beta_{0k} \} \quad k = N - 1, \dots, 1. \quad (43)$$

form a valid, piecewise linear certificate, which avoids locally maximizing  $\rho_k$  at the expense of later steps.

### 3.4 Verification with Input Saturation

This section describes a modified condition for certificates covering systems with saturated inputs. We examine the single-input case where the control law  $\mathbf{u}(t) = \mathbf{u}_0(t) - \mathbf{K}(t)\bar{\mathbf{x}}(t)$  is mapped through:

$$g(\mathbf{u}(t)) = \begin{cases} \mathbf{u}_+ & \mathbf{u}(t) \geq \mathbf{u}_+, \\ \mathbf{u}_- & \mathbf{u}(t) \leq \mathbf{u}_-, \\ \mathbf{u}(t) & \text{o.w.} \end{cases}$$

We begin by calculating the smallest level-set  $\rho$  for which either constraint becomes active. e.g.:

$$\underset{\bar{\mathbf{x}}'}{\text{minimize}} \quad J^*(\bar{\mathbf{x}}') \quad (44)$$

$$\text{subject to} \quad \mathbf{u}_+ = \mathbf{u}_0 - \mathbf{K}\bar{\mathbf{x}}' \quad (45)$$

Let  $\rho_{\max}$  be the solution to this convex quadratic program. Similarly,  $\rho_{\min}$  can be computed as the smallest level-set for which  $\mathbf{u}^*(t) = \mathbf{u}_-$ . These values allow for a case-wise analysis of the saturation.

We then compute a separate SOS condition for each active constraint, and choose the largest  $\rho$  which satisfies all. Let

$$\dot{J}_-(\bar{\mathbf{x}}) = \frac{\partial}{\partial \bar{\mathbf{x}}} J(\bar{\mathbf{x}})\mathbf{f}(\bar{\mathbf{x}}, \mathbf{u}_-) + \frac{\partial}{\partial t} J(\bar{\mathbf{x}}), \quad (46)$$

$$\dot{J}_K(\bar{\mathbf{x}}) = \frac{\partial}{\partial \bar{\mathbf{x}}} J(\bar{\mathbf{x}})\mathbf{f}(\bar{\mathbf{x}}, \mathbf{u}_0 - \mathbf{K}\bar{\mathbf{x}}) + \frac{\partial}{\partial t} J(\bar{\mathbf{x}}), \quad (47)$$

$$\dot{J}_+(\bar{\mathbf{x}}) = \frac{\partial}{\partial \bar{\mathbf{x}}} J(\bar{\mathbf{x}})\mathbf{f}(\bar{\mathbf{x}}, \mathbf{u}_+) + \frac{\partial}{\partial t} J(\bar{\mathbf{x}}). \quad (48)$$

Then the verification inequality (14) can be replaced by

$$\dot{J}_-(\bar{\mathbf{x}}) + h_1(\bar{\mathbf{x}})(\rho - J(\bar{\mathbf{x}})) + h_2(\bar{\mathbf{x}})(\mathbf{u}_- - \mathbf{u}_0 + \mathbf{K}\bar{\mathbf{x}}) \leq -\epsilon\|\bar{\mathbf{x}}\|_2^2, \quad (49)$$

$$\begin{aligned} \dot{J}_K(\bar{\mathbf{x}}) + h_3(\bar{\mathbf{x}})(\rho - J(\bar{\mathbf{x}})) + h_4(\bar{\mathbf{x}})(\mathbf{u}_0 - \mathbf{K}\bar{\mathbf{x}} - \mathbf{u}_-) \\ + h_5(\bar{\mathbf{x}})(\mathbf{u}_+ - \mathbf{u}_0 + \mathbf{K}\bar{\mathbf{x}}) \leq -\epsilon\|\bar{\mathbf{x}}\|_2^2, \end{aligned} \quad (50)$$

$$\dot{J}_+(\bar{\mathbf{x}}) + h_6(\bar{\mathbf{x}})(\rho - J(\bar{\mathbf{x}})) + h_7(\bar{\mathbf{x}})(\mathbf{u}_0 - \mathbf{K}\bar{\mathbf{x}} - \mathbf{u}_+) \leq -\epsilon\|\bar{\mathbf{x}}\|_2^2, \quad (51)$$

with each  $h_i(\bar{\mathbf{x}})$  sum of squares. Again, these Lagrange multipliers serve to restrict the verification to the regions of interest. Note that (50) is always evaluated, as it corresponds to the unsaturated region. Inequalities (49) and (51) verify stability with  $\mathbf{u}$  saturated at its minimum and maximum value, respectively, and need only be evaluated if  $\rho > \rho_{\min}$  and  $\rho > \rho_{\max}$ , respectively.

Analogous additions can be made to the LTV verification inequality (39). It is straightforward to extend this for higher input dimensions, however in general there will be  $3^m$  different combinations of saturated and unsaturated inputs, where  $m$  is the control dimension, which could quickly become computationally prohibitive for systems with many inputs.

## 4 The LQR-Tree Algorithm

Given the ability to design, stabilize, and verify a trajectory from an initial condition to the goal, the LQR-Tree algorithm proceeds by growing a randomized tree of stabilizing controllers to reach the goal. Because both the feedback design and verification work backwards in time, we grow the tree backwards, starting from the goal. The result is that the tree becomes a large web of stabilizing controllers which grab initial conditions and pull them towards the goal (with formal certificates of stability for the nonlinear, continuous state and action system).

The goal of the algorithm is to cover an entire region of interest, the set of points from which the goal state is reachable, or a specified bounded subset of this region, with this stabilizing controller. In Section 6, we define this property as “probabilistic feedback coverage”. To achieve this, we grow our tree in the fashion of an RRT, where new subgoals are chosen at random from a uniform distribution over the state space. Unlike the RRTs, we have additional information from the estimated regions of stability (funnel), and we can immediately discard sample points which are sampled inside the previously verified region. Define  $\mathcal{C}_k$  as the estimated region of stability for the entire tree after  $k$  iterations of the algorithm. Discarding random samples from  $\mathcal{C}_k$  occurs rapidly, due in part to our choice of a simple ellipsoidal estimate for the funnels, and causes a dramatic improvement in the sparsity of the randomized tree.

One consequence of this sampling strategy is that the algorithm which attempts to connect the random sample back to the goal must be capable of establishing this connection over some finite distance through state space. For this reason, we use the trajectory optimization algorithms described briefly in Section 3.2 to grow a trajectory from the sampled point forward until it connects

to the existing tree. These algorithms are local, and are subject to local minima, but we prove in Section 6 that they are sufficient to obtain the desired coverage property of the algorithm. If the trajectory optimization fails to connect to the tree, the sampled point is discarded in order to allow for the possibility that the tree is not reachable from that sampled point; if the tree is reachable then this region will be sampled again in the future with a more extensive tree available to connect to. The trajectory optimization is set to timeout with a failure after some relatively small number of major iterations of the SQP algorithm in order to facilitate fast sampling and discarding of new samples.

The most important component of the trajectory optimization which attempts to connect to the tree is the “final tree constraint”, which specifies that the optimized trajectory must connect back to the tree, but is free to connect to any part of the continuous manifold of points which form the tree. On the  $k$ th iteration of the algorithm, given a candidate trajectory terminating in  $\mathbf{x}_0(t_f)$ , we enforce the vector constraint that

$$\min_{\mathbf{x}' \in \mathcal{T}_{k-1}} \|\mathbf{x}_0(t_f) - \mathbf{x}'\|_2 = \mathbf{0}, \quad (52)$$

where  $\mathcal{T}_k$  is the set of all points in the tree after  $k$  iterations of the algorithm. A well implemented optimization algorithm can satisfy this constraint by designing a trajectory which connects to the tree, then continuing to “walk” the terminal condition along the constraint manifold of the tree to maximize the objective.

## 4.1 The Algorithm

The algorithm proceeds by producing a tree,  $T$ , with nodes containing the tuples,  $\{\mathbf{x}, \mathbf{u}, \mathbf{S}, \mathbf{K}, \rho_c, i\}$ , where  $J^*(\bar{\mathbf{x}}, t) = \bar{\mathbf{x}}^T \mathbf{S} \bar{\mathbf{x}}$  is the local quadratic approximation of the value function,  $\bar{\mathbf{u}}^* = -\mathbf{K} \bar{\mathbf{x}}$  is the feedback controller,  $J^*(\bar{\mathbf{x}}, t) \leq \rho(t)$  is the funnel where  $\rho(t)$  is described by the vector of polynomial coefficients  $\rho_c$ , and  $i$  is a pointer to the parent node.

## 4.2 Executing LQR-Tree feedback

The resulting LQR-Tree policy is a function,

$$\mathbf{u} = \pi_T(\mathbf{x}, t_T, b), \quad (53)$$

with internal controller state  $t_T$  - the time corresponding to the current execution - and  $b$  - a unique identifier to the currently active branch of the tree. At the onset of control from an initial condition  $\mathbf{x}_i$ , these controller variables are initialized as

$$[t_T, b] = \operatorname{argmax}_{b \in T, t \in [t_0^b, t_f^b]} c(\mathbf{x}_i, t, b), \quad (54)$$

$$c(\mathbf{x}, t, b) = \rho^b(t) - (\mathbf{x}_i - \mathbf{x}_0^b(t))^T \mathbf{S}^b(t) (\mathbf{x} - \mathbf{x}_0^b(t)), \quad (55)$$

---

**Algorithm 1** LQR-Tree ( $\mathbf{f}, \mathbf{x}_G, \mathbf{u}_G, \mathbf{Q}, \mathbf{R}$ )

---

```
1:  $[\mathbf{A}, \mathbf{B}] \leftarrow$  linearization of  $\mathbf{f}(\mathbf{x}, \mathbf{u})$  around  $\mathbf{x}_G, \mathbf{u}_G$ 
2:  $[\mathbf{K}, \mathbf{S}] \leftarrow$  LQR( $\mathbf{A}, \mathbf{B}, \mathbf{Q}, \mathbf{R}$ )
3:  $\rho_c \leftarrow$  level-set computed as described in section 3.1.1
4:  $T.\text{init}(\{\mathbf{x}_g, \mathbf{u}_g, \mathbf{S}, \mathbf{K}, \rho_c, \text{NULL}\})$ 
5: for  $k = 1$  to  $K$  do
6:    $\mathbf{x}_{rand} \leftarrow$  random sample
7:   if  $\mathbf{x}_{rand} \in \mathcal{C}_k$  then
8:     continue
9:   end if
10:   $[t, \mathbf{x}_0(t), \mathbf{u}_0(t)]$  from trajectory optimization with a “final tree constraint”
11:  if  $\mathbf{x}_0(t_f) \notin \mathcal{T}_k$  then
12:    continue
13:  end if
14:   $[\mathbf{K}(t), \mathbf{S}(t)]$  from time-varying LQR
15:   $\rho_c \leftarrow$  level-set computed as in section 3.3.1
16:   $i \leftarrow$  pointer to branch in  $T$  containing  $\mathbf{x}_0(t_f)$ 
17:   $T.\text{add-branch}(\mathbf{x}_0(t), \mathbf{u}_0(t), \mathbf{S}(t), \mathbf{K}(t), \rho_c, i)$ 
18: end for
```

---

where superscript  $b$  denotes the nominal times, trajectories, and funnels on branch  $b$  of the tree (where  $b$  is from an index set of branches). We refer to the quantity,  $c(t)$ , as the *confidence*. A positive value of  $c(t)$  corresponds to a verification of stability to the goal and larger values of  $c(t)$  mean the current state is further inside the region of attraction. Although it is tempting to continually search over  $t_T$  as the policy is executed to maximize  $c(t)$ , this can lead to undesirable properties such as chattering. Therefore, in practice, we execute the policy with time evolving naturally ( $\dot{t}_T = 1$ ) unless a disturbance occurs which takes the state outside of the funnel, in which case we re-evaluate  $t_T$  and  $b$ . Similarly, branch index  $b$  is held constant until the end of its tape, at which point  $b$  is set to the parent trajectory of the previous branch, closer to the goal.

## 5 Simulations

We illustrate the operation of the algorithm on a two-dimensional toy problem, for which we can carefully plot the entire state space and covered set,  $\mathcal{C}_k$ , on each iteration of the algorithm. Here we use the swing-up task for a torque-limited simple pendulum,  $I\ddot{\theta} + b\dot{\theta} + mgl \sin \theta = \tau$ , with  $m = 1, l = .5, b = .1, I = ml^2, g = 9.8, |\tau| \leq 3$ . Here  $\mathbf{x} = [\theta, \dot{\theta}]^T$  and  $\mathbf{u} = \tau$ . The parameters of the LQR-tree algorithm were  $\mathbf{x}_G = [\pi, 0]^T$ ,  $\mathbf{u}_G = 0$ ,  $\mathbf{Q} = \text{diag}([10, 1])$ ,  $\mathbf{R} = 20$ ,  $N_f = 3$ ,  $N_t = 3$ .

Figure 3(a) shows the region of attraction (gray shaded region) after computing the linear time-invariant (LTI) LQR solution around the unstable equilib-



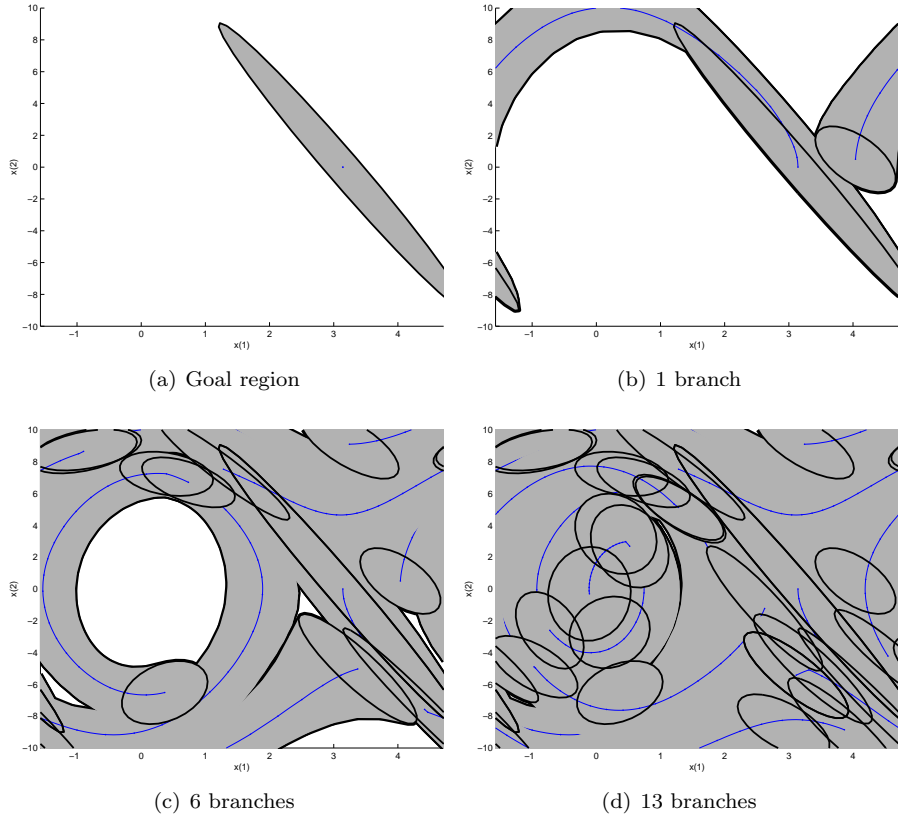


Figure 3: An LQR-tree for the simple pendulum. The  $x$ -axis is  $\theta \in [-\pi/2, 3\pi/2]$  (note that the state wraps around this axis), and the  $y$ -axis is  $\dot{\theta} \in [-10, 10]$ . The gray shaded region represents the covered set,  $\mathcal{C}_k$ , (aka, the “funnels”), after  $k$  branches have been added to the tree.

rium. Figure 3(b) shows the entire nominal trajectory (blue) to the first random sample point, and the funnels that have been computed for this trajectory. Note that the state-space of the pendulum lives on a cylinder, and that the trajectory (and region of attraction) wraps around from the left to the right. Plots (c-d) show the region of attraction as it grows to fill the state space, with the black ellipses denoting the inlet and outlet of each funnel in the tree.

In this example, the entire space is probabilistically covered (1000 random points chosen sequentially were all in the region of attraction) after the tree contained just 13 branches. On average, the algorithm terminates after 13.25 branches have been added for the simple pendulum with these parameters. For contrast, [4] shows a well-tuned single-directional RRT for the simple pendulum which has 5600 nodes. However the cost of adding each node is considerably greater here than in the traditional RRT, dominated by the line search used to

maximize the estimated region of stability.

## 6 LQR-Trees achieve “Probabilistic Feedback Coverage”

Consider a smooth system

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)), \quad (56)$$

with  $\mathbf{x}(t) \in \mathcal{X}$  and  $\mathbf{u}(t) \in \mathcal{U}$  where  $\mathcal{X}$  and  $\mathcal{U}$  are open subsets of  $\mathbb{R}^n$  and  $\mathbb{R}^m$ , respectively. We assume that for all piecewise continuous functions  $\mathbf{u} : [0, \infty) \rightarrow \mathcal{U}$ , a unique solution of (56) exists, and solutions have continuous dependence on controls and initial conditions. Note that there are many polynomial differential equations for which this is not true, however for well-modelled physically-motivated systems these are natural assumptions.

We wish to consider the performance of an iterative algorithm for the task of planning motions to a particular goal state  $\mathbf{x}_G \in \mathcal{X}$ .

We define the following sets:

- $\mathcal{R}(\mathbf{x}_G)$ : the set of points from which a state  $\mathbf{x}_G$  is reachable, possibly in infinite time. I.e.  $\mathbf{x}_0 \in \mathcal{R}(\mathbf{x}_G)$  if and only if there exists a piecewise continuous control signal  $\mathbf{u}^* : [0, \infty) \rightarrow \mathcal{U}$  such that the system (56) with initial condition  $\mathbf{x}_0$  and input  $\mathbf{u}^*$  asymptotically approaches  $\mathbf{x}_G$ .
- $\mathcal{C}_k$ : The set of points covered (e.g., inside the region of attraction of some trajectory) by the feedback motion planning algorithm after  $k$  iterations. This is the union of all funnels so far added, so we have  $\mathcal{C}_k \subseteq \mathcal{C}_{k+1}$ .
- $\mathcal{C}_\infty$ : the limit set of the coverage sets  $\mathcal{C}_k$ :

$$\mathcal{C}_\infty = \lim_{k \rightarrow \infty} \mathcal{C}_k. \quad (57)$$

That is,  $\mathbf{x} \in \mathcal{C}_\infty$  if and only if there exists a  $k_0$  such that  $\forall k > k_0 . \mathbf{x} \in \mathcal{C}_k$ .

Note that the limit  $\mathcal{C}_\infty$  is well-defined since  $\mathcal{C}_k \subseteq \mathcal{C}_{k+1}$  for all  $k$ , which implies that  $\limsup \mathcal{C}_k = \liminf \mathcal{C}_k$ .

We will be interested in proving the following property, which plays an analogous role in feedback-motion-planning as probabilistic completeness does in open-loop motion planning [16]:

**Definition 1** *A feedback motion planning algorithm achieves probabilistic feedback coverage for the goal state  $\mathbf{x}_G$  if  $\text{cl}(\mathcal{C}_\infty) = \text{cl}(\mathcal{R}(\mathbf{x}_G))$  with probability one<sup>4</sup>.*

The LQR-Trees algorithm works by randomly sampling states, and attempting to link them back to the tree. We make the following assumption:

<sup>4</sup>Here  $\text{cl}(\cdot)$  denotes the closure of a set.

**Assumption 1** *Each new state to be linked back to the tree is sampled from a distribution with non-zero probability density everywhere on  $\mathcal{X}$ .*

For bounded  $\mathcal{X}$  a uniform probability density is adequate, however for unbounded  $\mathcal{X}$  more care will need to be taken to ensure the density is integrable.

**Assumption 2** *(a) For the goal point  $\mathbf{x}_G$ , the LTI system is completely controllable, (b) for any feasible trajectory  $(\mathbf{x}^*, \mathbf{u}^*)$  of (56), the associated LTV Riccati differential equation has a positive-definite stabilizing solution from any positive-definite final-time condition.*

Assumption 2(b) is an assumption on local exponential stabilizability of the system dynamics by linear feedback. Taken together, these imply that  $\mathcal{C}_k$  has non-empty interior for every  $k$  and hence  $\mathcal{R}(\mathbf{x}_G)$  has non-empty interior, since each  $\mathcal{C}_k \subset \mathcal{R}(\mathbf{x}_G)$ .

Let  $\mathcal{T}_k$  be the set of states on the branches of the LQR-tree after  $k$  nodes have been added. We make the following assumption on the availability of an open-loop motion-planning algorithm:

**Assumption 3** *The open-loop motion-planner used to reconnect to the tree has the following property: given any point  $\mathbf{x}_0 \in \mathcal{R}(\mathbf{x}_G)$ , then the motion-planner will, with non-zero probability, find a finite time  $t_f$  and a control signal  $\mathbf{u}^*(t)$  defined on  $[0, t_f]$  such that the solution of  $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}^*)$  with  $\mathbf{x}(0) = \mathbf{x}_0$  achieves  $\mathbf{x}(t_f) \in \mathcal{T}_k$ .*

In Subsection 6.1 we will discuss this assumption, and how it can be satisfied in practice.

We are now ready to state the main result of this section:

**Theorem 1** *For any system and goal state satisfying Assumptions 1, 2, 3, the LQR-Trees algorithm achieves probabilistic feedback coverage.*

**Proof:** For a particular instance of the randomized tree, define

$$\mathcal{F} = \bigcap_{k=1}^{\infty} \{\mathcal{R}(\mathbf{x}_G)/\mathcal{C}_k\}, \quad (58)$$

i.e. the set of all states from which the goal is reachable that fail to be added to the tree. Note that  $\mathcal{F}$  is the complement of  $\mathcal{C}_\infty$  relative to  $\mathcal{R}(\mathbf{x}_G)$ . Now,  $\text{cl}(\mathcal{C}_\infty)$  equals  $\text{cl}(\mathcal{R}(\mathbf{x}_G))$  if and only if  $\mathcal{F}$  has empty interior. We will prove that this must be the case by contradiction.

Suppose  $\mathcal{F}$  has non-empty interior, with the interior denoted by  $\text{int}(\mathcal{F})$ . In this case the Lebesgue measure of  $\text{int}(\mathcal{F}) > 0$  and hence by Assumption 1, there is a non-zero probability of sampling from  $\text{int}(\mathcal{F})$ . Furthermore, by Assumption 3, for any such sample there is a non-zero probability that it can be connected back to the tree. Since all the samplings and trajectory searches are independent events, this implies that with probability one a sample  $\mathbf{x}_i$  will eventually be found in  $\text{int}(\mathcal{F})$  which connects back to the tree with a feasible trajectory  $\mathbf{x}^*(t)$ , i.e.  $\mathbf{x}^*(t_i) = \mathbf{x}_i$ ,  $\mathbf{x}^*(t_f) \in \mathcal{T}_k$  for some  $t_i, t_f$ .

Now, if such an  $\mathbf{x}_i$  were found, then by Assumption 2 the solution of Riccati equation associated with  $\mathbf{x}^*(t)$  is positive-definite for all  $t \in [t_i, t_f]$ . Since  $\mathcal{X}$  and  $\mathcal{U}$  are open sets, and  $f(\mathbf{x}, \mathbf{u})$  is smooth, this implies that around each point on the trajectory  $x^*(t)$  there is an ellipsoid of positive radius which will be added to  $\mathcal{C}_{k+1}$ . Since  $\mathbf{x}_i \in \mathcal{F}$ , some of these ellipsoids will have non-zero measure intersection with  $\text{int}(\mathcal{F})$ . However, this would contradict the definition  $\mathcal{F}$ , hence there is zero probability that  $\mathcal{F}$  has non-empty interior. Therefore with probability one  $\text{cl}(\mathcal{C}_\infty)$  equals  $\text{cl}(\mathcal{R}(\mathbf{x}_G))$ .

□

Note that we have not assumed anywhere that  $\mathcal{X}, \mathcal{U}$ , or  $\mathcal{R}(\mathbf{x}_G)$  are bounded, however unbounded sets will require careful selection of sampling density, and in practice may be problematic for satisfaction of Assumption 3.

## 6.1 Discussion on the Motion Planning Assumption

We briefly argue that Assumption 3 is reasonable in practice. Firstly, due to the smoothness of (56) and the complete controllability of the linearized system about  $\mathbf{x}_G$ , if from a sampled point the system can be made to asymptotically approach  $\mathbf{x}_G$  then it can be made to reach  $\mathbf{x}_G$  in finite time.

In the present implementation of LQR-Trees, a path back to the tree is sought by randomly sampling a piecewise-constant control signal  $\mathbf{u}$ , where the time horizon  $t_f$  is randomly sampled from a distribution which has non-zero density everywhere on  $[0, T]$  for some large  $T$  and the sampling interval is a random variable with non-zero density on  $[0, t_f]$ . Therefore, for sufficiently large  $T$ , arbitrarily fine approximations of any finite-length piece-wise continuous control signal can be sampled.

Using the sampled  $\mathbf{u}$  as an initial guess, a search is performed using over input signals with the cost function (20) subject to the constraints that  $\mathbf{x}(t_2) \in \mathcal{T}_k$  and  $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$ . The optimization is posed in direct collocation or multiple-shooting form, and is solved numerically via an SQP method in which nonlinear equality constraints are handled by  $l_1$  regularized elastic variables [9].

Now, suppose a sampled  $\mathbf{x}_0 \in \mathcal{R}(\mathbf{x}_G)$ , i.e., there exists a piecewise continuous function  $\mathbf{u}_0 : [0, t_1] \rightarrow \mathcal{U}$  which drives the state from  $\mathbf{x}_0$  to a point on the tree  $\mathbf{x}(t_f) \in \mathcal{T}_k$ . Since  $\mathcal{U}$  is open and solutions of (56) are continuous with respect to controls, there exists an open neighborhood of piecewise continuous functions  $[0, t_2] \rightarrow \mathcal{U}$ ,  $t_2 > t_1$  containing  $\mathbf{u}_0$ , which is the region of convergence to a local optimum for the optimization (20) considered over piecewise-continuous functions. Since our sampling can produce arbitrarily fine approximations of such functions, it can be argued that there is a non-zero probability of numerically finding a feasible solution of (20).

In practice, if the optimization fails to find a solution given the initial seed, the sampled point is discarded as temporarily uncontrollable, however after the tree has grown a nearby point will eventually be sampled from which the search is successful. Another choice for open-loop motion-planner could be a randomized-sampling based planner such as the RRT [16], run for a finite (but random) number of steps.




Figure 4: Controllability-limited example

## 6.2 Controllability-limited example

The performance of the algorithm when searching near a boundary of controllability in state space is an important consideration, particularly when applying it to complicated, underactuated systems in which the location and form of such boundaries is unknown. Figure 4 demonstrates the performance of the algorithm when applied to the following second-order polynomial dynamic system:

$$\ddot{x} = x + u^2 - u - 3/4. \quad (59)$$

This system has a known controllable set defined by  $\dot{x} < 1 - x$  (whose boundary is shown as a red line in the figure). Intuitively, the force applied to the system is bounded to be in the set  $[-1, \infty)$ . From this realization, energy considerations clearly show that the boundary of the controllable set is defined by the line  $\dot{x} = 1 - x$ . The algorithm will provably fill the sampled space (shown in light blue) which is capable of reaching the goal. As the figure demonstrates, this is achieved by placing ever smaller funnels closer and closer to the boundary, until the boundary is approximated by the surfaces of the funnels.

## 7 Discussion

In this section, we briefly discuss the limitations of LQR-Trees, some possible variations of the main algorithm, and some implementation details.

## 7.1 Limitations of LQR-Trees

While LQR-Trees has broad applicability in the realm of nonlinear systems, it does have limitations with respect to the dimension of the systems that can be solved. Its scaling behavior can best be understood by considering the four components that make up the algorithm: 1) a local motion planner (e.g., multiple shooting) 2) a global motion planner which connects the local plans (RRTs) 3) local control design (time-varying LQR) and 4) local verification (SOS).

The local motion planner and the local control design (components 1 and 3) both scale well with dimension. The global motion planner (component 2) has poor asymptotic behavior, but has been shown to work well on systems with 30 or more dimensions [15], and, for particular tasks, on systems with over a 1000 dimensions [32]. The bottleneck in addressing systems of increasing dimensionality is component 4, verification using SOS. We have had success applying LQR-Trees to systems with up to 5 dimensions, and the literature has examples involving rigid body dynamics in 6 dimensions, e.g. [11].

The reliance on linearized controller synthesis means that LQR-Trees is not directly applicable to nonlinear systems which are controllable but have uncontrollable linearizations, such as nonholonomic systems.

## 7.2 Variations of the algorithm

- **Discrete-time formulation.** Although the presentation here focused on the continuous time case, the approach has a natural discrete-time formulation. In many ways, the discrete time formulation is simpler, mitigating many of the numerical subtleties to be described in Section 7.3, and may in fact be preferable for implementation on a sampled-data control system. The primary drawbacks are that (1) it is more difficult to do a careful time-discretization of a polynomial system than a linear one, and (2) the tree will now consist of a collection of points instead of a smooth manifold, complicating any efforts to create a final value constraint which can “walk along the tree”.
- **Compatible with optimal trajectories.** The LQR-tree algorithm provides a relatively efficient way to fill the controllable state space with funnels, but does not make any claim on the optimality of the resulting trajectories. If tracking particular trajectories, or optimal trajectories, is important for a given problem, then it is quite natural to seed the LQR-tree with one or more locally optimal trajectories, then use the random exploration to fill in any missing regions.
- **Early termination.** For higher dimensional problems, covering the controllable state space may be unnecessary or impractical. Based on the RRTs, the LQR-trees can easily be steered towards a region of state space (e.g., by sampling from that region with slightly higher probability) containing important initial conditions. Termination could then occur when some important subspace is covered by the tree.

- **Bidirectional trees.** Although LQR-trees only grow backwards from the goal, a partial covering tree (from an early termination) could also serve as a powerful tool for real-time planning. Given a new initial condition, a forward RRT simply has to grow until it intersects with the *volume*,  $C_k$ , defined by the region of attraction of the backwards tree.
- **Other local stabilizers.** In this paper we have proposed using LQR as a local stabilizer for points and trajectories. The main reason for this choice is simplicity of computation, and the fact that the “cost-to-go” automatically provides a reasonable candidate Lyapunov function. Other popular control design techniques, such as  $H^\infty$  or LMI-based methods also supply candidate Lyapunov functions, and might be more appropriate for certain systems. In particular, some information about the nonlinearities could be encoded in an uncertainty structure. These methods require more computational effort than LQR, but if they stabilize much larger regions that may be a useful trade-off.
- **Other regional stability verifiers.** Similarly, in this paper we have proposed SOS as a regional stability verifier. This allows a very large class of smooth nonlinear systems to be handled under a single framework for which reliable computational tools are available. However, for systems of high dimension, SOS programs can result in very large semidefinite programs, which become a computational bottleneck. An alternative version which approximates the verification stage via sampling can be performed with much lower computational cost, at the expense of losing the rigorous stability guarantees [31].

Many interesting systems in robotics can be decomposed into “easy” subsystems (e.g. actuated links with linear or nearly-linear dynamics) and “hard” subsystems (e.g. unactuated links with highly nonlinear dynamics). Taking advantage of this structure could reap large benefits in computation of stability regions. Taking this line even further, there appear to be deep connections between SOS methods and classical methods for nonlinear stability analysis such as integral quadratic constraints [20], which we intend to investigate.

- **Multi-query algorithms.** Another very interesting question is the question of reusing the previous computational work when the goal state is changed. In the pendulum example, consider having a new goal state,  $\mathbf{x}_G = [\pi + 0.1, 0]^T$  - this would of course require a non-zero torque to stabilize. To what extent could the tree generated for stabilizing  $\mathbf{x}_G = [\pi, 0]^T$  be used to stabilize this new fixed point? If one can find a trajectory to connect up the new goal state near the root of the tree, then the geometry of the tree can be preserved, but naively, one would think that all of the stabilizing controllers and the verification would have to be re-calculated. Interestingly, there is also a middle-road, in which the existing feedback policy is kept for the original tree, and the estimated funnels are not re-computed, but simply scaled down to make sure that the funnels from the

old tree transition completely into the funnel for the new tree. This could be accomplished very efficiently, by just propagating a new  $\rho_{max}$  through the tree, but might come at the cost of losing coverage.

### 7.3 Implementation details

The LQR-Tree algorithm presented here, which relies on a continuous time formulation and optimization over polynomials, requires the practitioner to address a number of potential numerical issues. We list a few of these here:

- **Balancing coordinates.** One drawback of working with polynomials is that their associated optimization problems are often poorly conditioned (high-order polynomials blow up quickly), and solutions to the Riccati differential equation can easily have a large range of eigenvalues. An essential step for getting good performance out of the semidefinite programming was finding a coordinate transformation:

$$\mathbf{x}_b = \mathbf{T}\mathbf{x}, \tag{60}$$

where  $\mathbf{T}$  was selected to numerically condition the problem by making the matrices  $\mathbf{T}'\mathbf{S}\mathbf{T}$  and  $\mathbf{T}'\left(2\mathbf{S}\frac{\partial\mathbf{f}(\mathbf{x}_0+\bar{\mathbf{x}},\mathbf{u}_0-\mathbf{K}\bar{\mathbf{x}})}{\partial\mathbf{x}}+\dot{\mathbf{S}}\right)\mathbf{T}$  are as close as possible to the identity matrix. Without this balancing, the semidefinite solvers often return errors with infeasible programs.

- **Piecewise polynomial approximations of numerical integration.** For the continuous time formulation, it was essential to carefully manage the piecewise polynomial interpolations of solutions from trajectory optimization and from the Riccati differential equation. Simple cubic spline representations often caused even the nominal trajectory to appear to be unstable. In the current implementation, we represent the output of the trajectory optimization with a cubic piecewise polynomial that, for every segment, matches the knot points and the derivatives of the knot points which are known from the dynamics function. The Riccati differential equation is solved using MATLAB `ode45`, interpolations of that solution are done using `deval`, and calculations of  $\dot{\mathbf{S}}$  are done by explicitly re-evaluating the Riccati equation.
- **Time dependence on  $\dot{J}$ .** Finally, one must be careful when performing verification on the time-varying Lyapunov candidate. Naively, if one used a cubic spline representation of the nominal trajectory and of  $\mathbf{S}$ , then  $\dot{\mathbf{J}}$  quickly becomes an unnaturally high order. Instead, we fit the best cubic spline directly to  $\dot{J}$ .

### 7.4 A Software Distribution

Due to the potential complexity of getting a robust and high performance implementation of the LQR-Tree algorithm, the authors are making a MATLAB toolbox available at <http://groups.csail.mit.edu/locomotion/software.html>



## 8 Summary and Conclusions

Advances in direct computation of Lyapunov functions have enabled a new class of feedback motion planning algorithms for complicated dynamical systems. This paper presented the LQR-Tree algorithm which uses Lyapunov computations to evaluate the regions of attraction of randomized trees stabilized with LQR feedback. We prove, and demonstrate through simulation examples, that this algorithm has the property of “probabilistic feedback coverage”. Furthermore, initial experimental results suggest that this coverage occurs efficiently, requiring only a small number of stabilized trajectories to cover the controllable space.

Further investigation of this algorithm will likely result in a covering motion planning strategy for underactuated systems with dimensionality greater than what is accessible by discretization algorithms like dynamic programming, and early termination strategies which provide targeted coverage of state space in much higher dimensional systems. The resulting policies will have certificates guaranteeing their performance on the system model.

## Acknowledgements

The authors wish to thank Alexandre Megretski and Pablo Parrilo for their invaluable assistance on the formulations and numerical implementations of the sums-of-squares methods. We also wish to thank Philipp Reist and Steve LaValle for valuable discussions. This work was supported by the National Science Foundation [Awards 0746194 and 0915148], and the Air Force Research Laboratory [Contract FA8650-05-C- 7262].

## 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] Christopher G. Atkeson and Benjamin Stephens. Random sampling of states in dynamic programming. In *Advances in Neural Information Processing Systems*, 2008.
- [3] John T. Betts. *Practical Methods for Optimal Control Using Nonlinear Programming*. SIAM Advances in Design and Control. Society for Industrial and Applied Mathematics, 2001.
- [4] Michael Branicky and Michael Curtiss. Nonlinear and hybrid control via RRTs. *Proc. Intl. Symp. on Mathematical Theory of Networks and Systems*, 2002.
- [5] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek. Sequential composition of dynamically dexterous robot behaviors. *International Journal of Robotics Research*, 18(6):534–555, June 1999.

- [6] Graziano Chesi. Estimating the domain of attraction for non-polynomial systems via LMI optimizations. *Automatica*, 45(6):1536 – 1541, 2009.
- [7] G. Franzè, A. Casavola, D. Famularo, and E. Garone. An off-line MPC strategy for nonlinear systems based on SOS programming. In L. Magni, D. M. Raimondo, and F. Allgower, editors, *Nonlinear Model Predictive Control*, volume 384 of *Lecture Notes in Control and Information Sciences*, pages 491–499. Springer-Verlag, 2009.
- [8] Emilio Frazzoli, Munther A. Dahleh, and Eric Feron. Real-Time Motion Planning for Agile Autonomous Vehicles. *Journal of Guidance, Control, and Dynamics*, 25(1):116–129, JanuaryFebruary 2002.
- [9] Philip E. Gill, Walter Murray, and Michael A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131, 2005.
- [10] Rafal Goebel, Christophe Prieur, and Andrew R. Teel. Smooth patchy control lyapunov functions. *Automatica*, 45:675–683, 2009.
- [11] Gollu, N., Rodrigues, and L. Control of large angle attitude maneuvers for rigid bodies using sum of squares. pages 3156 –3161, july 2007.
- [12] Tor A. Johansen. Computation of Lyapunov functions for smooth nonlinear systems using convex optimization. *Automatica*, 36(11):1617 – 1626, 2000.
- [13] L.E. Kavraki, P. Svestka, JC Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, August 1996.
- [14] Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, 3rd edition, December 2001.
- [15] J. Kuffner, K. Nishiwaki, S. Kagami, M. Inaba, and H. Inoue. Motion planning for humanoid robots under obstacle and dynamic balance constraints. *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, 1:692–698 vol.1, 2001.
- [16] S. LaValle and J. Kuffner. Rapidly-exploring random trees: Progress and prospects. In *Proceedings of the Workshop on the Algorithmic Foundations of Robotics*, 2000.
- [17] Steven M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.
- [18] M.T. Mason. The mechanics of manipulation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 544–548. IEEE, 1985.
- [19] A. Megretski. Systems polynomial optimization tools (SPOT), available online: <http://web.mit.edu/ameg/www/>.

- [20] A. Megretski and A. Rantzer. System analysis via integral quadratic constraints. *IEEE Transactions on Automatic Control*, 42(6):819–830, Jun 1997.
- [21] Ian M. Mitchell and Jeremy A. Templeton. A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems. In *Proceedings of Hybrid Systems Computation and Control*, March 2005.
- [22] Y. Nesterov and A. Nemirovskii. *Interior-Point Polynomial Algorithms in Convex Programming*. Studies in Applied Mathematics. Society for Industrial Mathematics, 1995.
- [23] A. Papachristodoulou, M. M. Peet, and S. Lall. Analysis of polynomial systems with time delays via the sum of squares decomposition. *IEEE Transactions on Automatic Control*, 54(5):1058–1064, May 2009.
- [24] A. Papachristodoulou and S. Prajna. On the construction of Lyapunov functions using the sum of squares decomposition. In *Proceedings of the 41st IEEE Conference on Decision and Control*, Dec 2002.
- [25] Pablo A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, May 18 2000.
- [26] N. Peterfreund and Y. Baram. Convergence analysis of nonlinear dynamical systems by nested Lyapunov functions. *IEEE Transactions on Automatic Control*, 43(8):1179–1184, Aug 1998.
- [27] S. Prajna, P.A. Parrilo, and A. Rantzer. Nonlinear control synthesis by convex optimization. *Automatic Control, IEEE Transactions on*, 49(2):310–314, Feb. 2004.
- [28] Stephen Prajna, Antonis Papachristodoulou, Peter Seiler, and Pablo A. Parrilo. *SOSTOOLS: Sum of Squares Optimization Toolbox for MATLAB Users guide*, 2.00 edition, June 1 2004.
- [29] Stephen Prajna and Anders Rantzer. Convex programs for temporal verification of nonlinear dynamical systems. *SIAM Journal of Control and Optimization*, 46(3):999–1021, 2007.
- [30] Anders Rantzer. A dual to Lyapunov’s stability theorem. *Systems & Control Letters*, 42(3):161 – 168, 2001.
- [31] Philipp Reist and Russ Tedrake. Simulation-based LQR-trees with input and state constraints. In *To appear in the Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2010.
- [32] Alexander Shkolnik and Russ Tedrake. Path planning in 1000+ dimensions using a task-space Voronoi bias. In *Proceedings of the IEEE/RAS International Conference on Robotics and Automation (ICRA)*. IEEE/RAS, 2009.

- [33] Jean-Jacques E. Slotine and Weiping Li. *Applied Nonlinear Control*. Prentice Hall, October 1990.
- [34] W. Tan and A. Packard. Stability region analysis using polynomial and composite polynomial Lyapunov functions and sum-of-squares programming. *IEEE Transactions on Automatic Control*, 53(2):565–571, March 2008.
- [35] Russ Tedrake. LQR-Trees: Feedback motion planning on sparse randomized trees. In *Proceedings of Robotics: Science and Systems (RSS)*, page 8, 2009.
- [36] Claire J. Tomlin, Ian M. Mitchell, Alexandre M. Bayen, and Meeko K. M. Oishi. Computational techniques for the verification and control of hybrid systems. In *Multidisciplinary Methods for Analysis Optimization and Control of Complex Systems*, Mathematics in Industry, pages 151–175. Springer Berlin Heidelberg, 2005.
- [37] Ufuk Topcu, Andrew Packard, and Peter Seiler. Local stability analysis using simulations and sum-of-squares programming. *Automatica*, 44(10):2669 – 2675, 2008.