

Controller design and region of attraction estimation for nonlinear dynamical systems

Milan Korda¹, Didier Henrion^{2,3,4}, Colin N. Jones¹

Draft of March 20, 2014

Abstract

This work presents a method to obtain inner and outer approximations of the region of attraction of a given target set as well as an admissible controller generating the inner approximation. The method is applicable to constrained polynomial dynamical systems and extends to trigonometric and rational systems. The method consists of three steps: compute outer approximations, extract a polynomial controller while guaranteeing the satisfaction of the input constraints, compute inner approximations with respect to the closed-loop system with this controller. Each step of the method is a convex optimization problem, in fact a semidefinite program consisting of minimizing a linear function subject to linear matrix inequality (LMI) constraints. The inner approximations are positively invariant provided that the target set is included in the inner approximation and/or is itself invariant.

Keywords: Region of attraction, polynomial control systems, occupation measures, linear matrix inequalities (LMIs), convex optimization, viability theory, reachable set, capture basin.

1 Introduction

In this paper, the region of attraction (ROA) of a given target set is defined as the set of all states that can be steered to the target set at any time while satisfying state and control input constraints. The problem of ROA characterization and computation with its many variations have a long tradition in both control and viability theory (where the ROA is typically called the capture basin [1]). Computational methods start with the seminal work of Zubov [21] and are surveyed in, e.g., [7, 11, 3] and the book [2].

This work proposes a computationally tractable method for obtaining both inner and outer approximations of the ROA and, importantly, an approximate polynomial controller. The

¹Laboratoire d'Automatique, École Polytechnique Fédérale de Lausanne, Station 9, CH-1015, Lausanne, Switzerland. {milan.korda,colin.jones}@epfl.ch

²CNRS; LAAS; 7 avenue du colonel Roche, F-31400 Toulouse; France. henrion@laas.fr

³Université de Toulouse; LAAS; F-31400 Toulouse; France.

⁴Faculty of Electrical Engineering, Czech Technical University in Prague, Technická 2, CZ-16626 Prague, Czech Republic.

approach consists of three steps: compute outer approximations, extract a polynomial controller satisfying the input constraints, compute inner approximations of the closed-loop system with the extracted controller. Thus, the approach can also be viewed as a design tool providing a polynomial controller with an analytically known inner and outer approximations of its ROA.

Computing both inner and outer approximations, compared to just one or the other, enables assessing the tightness of the approximations obtained and provides a valuable insight into achievable performance and/or safety of a given constrained control system. For instance, a natural application for outer approximations is in collision avoidance, whereas a typical application of a (positively invariant) inner approximation is as a terminal constraint of a model predictive controller ensuring recursive feasibility of the underlying optimization problem [17].

The approach builds on and extends the ideas of [7] and [11] and the controller extraction procedure of [16] which was also sketched earlier in [6] in the context of switching sequence design. The main contributions with respect these works are:

- Contrary to [7], we treat the infinite time version of the ROA computation problem. The approach uses discounting similarly to our previous work [11] but here we treat the problem of computing the ROA, not maximum controlled invariant set as in [11].
- Contrary to [10] we compute inner approximations for *controlled* systems. This significantly extends the applicability of the approach but brings additional practical and theoretical challenges. In addition, under mild conditions, the inner approximations obtained are controlled invariant.
- Contrary to [16] the extracted controller is guaranteed to satisfy the input constraints and comes with an explicit estimate of its ROA, both from inside and outside.
- The formulation providing outer approximations of the ROA is based on a different idea than that of [7, 11, 16] and provides tighter estimates on the numerical examples investigated.

As in previous works [7, 10, 11], the method presented in this paper studies how whole ensembles of trajectories evolve through time using the concept of occupation measures. To obtain the outer approximations, we first characterize the ROA as a value function of a certain non-linear optimal control problem which we then relax using measures in the spirit of [14]. This leads to a primal infinite-dimensional linear program (LP) which is then relaxed using a hierarchy of finite-dimensional semidefinite programming problems (SDPs) whose solutions can be used to extract approximate polynomial controllers. Finite-dimensional approximations of the dual infinite-dimensional LP in the space of continuous functions are sum-of-squares (SOS) problems and provide outer approximations to the ROA. To obtain the inner approximations, we characterize directly using measures the *complement* of the ROA associated with the closed-loop system with the extracted polynomial controller. This leads to an infinite-dimensional primal LP in the space of measures. Finite-dimensional approximations of the dual LP on continuous functions are SOS problems and provide outer approximations to the complement of the ROA and hence inner approximations to the ROA itself.

Note in passing that the use of occupation measures has a long tradition both in deterministic and stochastic control; see, e.g., [11] for a historical account with the emphasis on applications to ROA and MCI set computation.

The paper is organized as follows. Section 2 defines the problem to be solved; Section 3 introduces the occupation measures; Section 4 presents the outer approximation formulation; Section 5 describes the controller extraction procedure; Section 6 presents the inner approximations; and Section 7 demonstrates the whole procedure on numerical examples.

1.1 Notation

Throughout the paper we work with standard Euclidean spaces; all subsets of these spaces we refer to are automatically assumed Borel measurable. The spaces of continuous and once continuously differentiable functions on a set X are denoted by $C(X)$ and $C^1(X)$, respectively. By a measure we understand a countably-additive mapping from sets to real numbers. Integration of a function $v(\cdot)$ with respect to a measure μ over a set X is denoted by $\int_X v(x) d\mu(x)$; often we omit the integration variable or the domain of integration and write $\int_X v d\mu$ or $\int v d\mu$. The support of a measure μ (i.e., the smallest closed set whose complement has a zero measure) is denoted by $\text{spt } \mu$. A moment sequence $\{y_\alpha\}_{\alpha \in \mathbb{N}^n}$ of a measure μ on \mathbb{R}^n is defined by $y_\alpha = \int_{\mathbb{R}^n} x^\alpha d\mu(x) = \int_{\mathbb{R}^n} x_1^{\alpha_1} \cdot \dots \cdot x_n^{\alpha_n} d\mu(x)$. The indicator function of a set A , i.e., the function equal to one on the set and zero elsewhere, is denoted by $\mathbb{I}_A(\cdot)$.

2 Problem description

Consider the polynomial input-affine dynamical system

$$\dot{x}(t) = f(x(t)) + G(x(t))u(t), \quad (1)$$

where the vector- and matrix-valued functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $G : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ have polynomial entries. The system is subject to a basic semialgebraic state constraint¹

$$x(t) \in X := \{x \in \mathbb{R}^n : g_X(x) > 0\}, \quad (2)$$

where g_X is a polynomial, and box input constraints

$$u(t) \in U := [0, \bar{u}]^m, \quad \bar{u} \geq 0. \quad (3)$$

The assumption that the input constraint is of the form (3) is made without loss of generality since any box in \mathbb{R}^m can be affinely transformed² to $[0, \bar{u}]^m$. It is also worth mentioning that

¹The assumption that X is given by a super-level set of single polynomial is made for ease of exposition; all results extend immediately to arbitrary basic semialgebraic sets. This extension is briefly described in Appendix A.

²Any box in \mathbb{R}^m can, of course, be also affinely transformed to $[0, 1]^m$. However, we decided to consider the more general form $[0, \bar{u}]^m$ so that it is immediately apparent where the upper bound \bar{u} comes into play in the optimization problems providing the region of attraction estimates.

arbitrary polynomial dynamical systems of the form $\dot{x} = f(x, u)$ can also be handled by considering the dynamic extension

$$\begin{bmatrix} \dot{x} \\ \dot{u} \end{bmatrix} = \begin{bmatrix} f(x, u) \\ v \end{bmatrix}$$

where the real control input u is treated as a state and v is a new, unconstrained, control input. Some of our convergence results hinge on the compactness of the input constraint set and therefore one may want to impose additional bounds on the new control input v , which correspond to slew-rate constraints on the true control input u , a requirement often encountered in practice.

In the remainder of the text we make the following standing assumption:

Assumption 1 *The set $\bar{X} := \{x \in \mathbb{R}^n : g_X(x) \geq 0\}$ is compact.*

This assumption is of a technical nature, required for the convergence results of Section 6.

Given a target set

$$X_T := \{x \in \mathbb{R}^n : g_T(x) > 0\} \subset X,$$

where the function g_T is a polynomial, the goal of the paper is to compute inner approximations of the *region of attraction* (ROA)

$$X_0 = \left\{ x_0 \in \mathbb{R}^n : \exists u(\cdot), x(\cdot), \tau \in [0, \infty) \text{ s.t. } \begin{aligned} &\dot{x}(t) = f(x(t)) + G(x(t))u(t) \text{ a.e.,} \\ &x(t) \in X, u(t) \in U, \forall t \in [0, \tau], x(0) = x_0, x(\tau) \in X_T \end{aligned} \right\},$$

where a.e. means “almost everywhere” with respect to the Lebesgue measure on $[0, \tau]$, $x(\cdot)$ is absolutely continuous and $u(\cdot)$ is measurable.

In words, the region of attraction X_0 is the set of all initial states that can be steered to the target set X_T at any time $\tau \in [0, \infty)$ in an admissible way, i.e., without violating the state or input constraints.

Our approach to compute the inner and outer approximations to the ROA consists of three steps:

1. Compute an outer approximation to the ROA,
2. Extract a polynomial controller out of this outer approximation,
3. Compute an inner approximation for the closed-loop system with this controller.

These steps are detailed in the rest of the paper.

3 Occupation measures

The key ingredient of our approach is the use of measures to capture the evolution of a family of the trajectories of a dynamical system starting from a given initial distribution.

Assume therefore that the initial state is not a single point but that its spatial distribution is given by an *initial measure* μ_0 and that to each initial condition a stopping time $\tau(x_0) \in [0, \infty]$ is assigned. Assume that the support of μ_0 and the stopping time $\tau(\cdot)$ are chosen such that there exists a controller $u(x)$ such that all closed-loop trajectories $x(\cdot | x_0)$ starting from initial conditions $x_0 \in \text{spt } \mu_0$ remain in X for all $t \in [0, \tau(x_0))$.

Then we can define the (*average*) *discounted occupation measure* as

$$\mu(A) = \int_X \int_0^{\tau(x_0)} e^{-\beta t} \mathbb{I}_A(x(t | x_0)) dt d\mu_0(x_0), \quad A \subset X, \quad (4)$$

where $\beta > 0$ is a discount factor. This measure measures the average (where the averaging is over the distribution of the initial state) discounted time spent in subsets of the state-space in the time interval $[0, \tau(x_0))$.

The *discounted final measure* μ_T is defined by

$$\mu_T(A) = \int_X e^{-\beta\tau(x_0)} \mathbb{I}_A(x(\tau(x_0) | x_0)) d\mu_0(x_0), \quad A \subset X, \quad (5)$$

where we define $e^{-\beta\tau(x_0)} := 0$ whenever $\tau(x_0) = +\infty$. The discounted final measure captures the time-discounted spatial distribution of the state at the stopping time $\tau(x_0)$.

The equation linking the three measures is a variant of the discounted Liouville equation

$$\int_X v(x) d\mu_T(x) + \beta \int_X v(x) d\mu(x) = \int_X v(x) d\mu_0(x) + \int_X \nabla v(x) \cdot [f(x) + G(x)u(x)] d\mu(x), \quad (6)$$

which holds for all $v \in C^1(X)$. This equation will replace the system dynamics (16) when studying the evolution of trajectories starting from the initial distribution μ_0 over the possibly infinite time intervals $[0, \tau(x_0))$. The equation is derived in Appendix B.

4 Outer approximations

In this section we formulate an infinite-dimensional linear program (LP) in the space of measures characterizing the ROA (or more precisely a set closely related to it) and whose finite-dimensional relaxations allow for the extraction of an approximate controller. Finite-dimensional relaxations of the dual to this LP then provide outer approximations to the ROA.

Here we take an approach different³ from [7, 11, 16] and formulate the ROA computation problem as an optimal control problem whose value function characterizes the ROA; this problem is then relaxed using measures. The optimal control problem reads:

³The works [7, 11, 16] do not use the value function of an optimal control problem and argue directly in the space of trajectories modelled by measures. This approach can also be used to characterize the infinite-time ROA sought here; however, on our class of numerical examples, the tightness of the finite-dimensional relaxations and the quality of the extracted controllers seem to be inferior to those of the approach presented here. Note that this is something peculiar to the infinite-time ROA problem and has not been observed for the finite-time ROA problem of [7, 16] and the maximum controlled invariant set problem of [11].

$$\begin{aligned}
V(x_0) &:= \sup_{\tau \in [0, \infty], u(\cdot)} e^{-\beta\tau} [\mathbb{I}_{X_T}(x(\tau)) - \mathbb{I}_{X \setminus X_T}(x(\tau))] \\
&\text{s.t.} \quad \dot{x}(t) = f(x(t)) + G(x(t))u(t), \quad x(0) = x_0 \\
&\quad \quad \quad x(t) \in X, \quad u(t) \in U \quad \forall t \in [0, \tau],
\end{aligned} \tag{7}$$

where the objective function is zero if $\tau = +\infty$. Clearly, the initial conditions that can be steered to X_T achieve a strictly positive objective; the initial conditions that necessarily leave X a strictly negative objective; and the initial conditions that can be kept within X forever but do not enter X_T achieve a zero objective. Therefore $X_0 = \{x \mid V(x) > 0\}$.

Now we formulate an infinite-dimensional LP relaxation of the problem (7) in the space of measures. First, and this is the key insight due to [16] ensuring that an approximate controller can be extracted, we view each component $u_i(x)$ of the controller $u(x)$ in (6) as a density of a *control measure*⁴ σ_i defined by

$$\sigma_i(A) = \int_A u_i(x) d\mu(x), \quad i = 1, 2, \dots, m, \quad A \subset X. \tag{8}$$

Defining the linear differential operators $\mathcal{L}_f : C^1(X) \rightarrow C(X)$ and $\mathcal{L}_{G_i} : C^1(X) \rightarrow C(X)$, $i = 1, \dots, m$, by

$$\mathcal{L}_f v := \nabla v \cdot f, \quad \mathcal{L}_{G_i} v := \nabla v \cdot G_i,$$

where G_i denotes the i^{th} column of G , the Liouville equation (6) can be rewritten as

$$\int_X v d\mu_T + \beta \int_X v d\mu = \int_X v d\mu_0 + \int_X \mathcal{L}_f v d\mu + \sum_{i=1}^m \int_X (\mathcal{L}_{G_i} v) d\sigma_i, \tag{9}$$

where we used the fact that $d\sigma_i(x) = u_i(x)d\mu(x)$ in view of (8).

The input constraints $u_i(x) \in [0, \bar{u}]$ are then enforced by requiring that $0 \leq \sigma_i \leq \bar{u}\mu$, which is equivalent to saying that σ_i is a nonnegative measure absolutely continuous with respect to μ with density (i.e., Radon-Nikodým derivative) taking values in $[0, \bar{u}]$. The constraint $0 \leq \sigma_i \leq \bar{u}\mu$ can be written equivalently as $\sigma_i \geq 0$, $\sigma_i + \hat{\sigma}_i = \bar{u}\mu$ for some non-negative slack measure $\hat{\sigma}_i \geq 0$. The state constraint $x(t) \in X$ is handled by requiring that $\text{spt } \mu \subset X$. Further we decompose the final measure as $\mu_T = \mu_T^1 + \mu_T^2$ with⁵ $\text{spt } \mu_T^1 \subset X_T$ and $\text{spt } \mu_T^2 \subset X$. Finally we set the initial measure equal to the Lebesgue measure (i.e., $d\mu_0(x) = dx$).

⁴In [16] the authors use signed measures which they then decompose using the Jordan decomposition – this step is avoided here since our control input is, without loss of generality, non-negative; this reduces the number of control measures by half and therefore makes the subsequent SDP relaxations more tractable.

⁵We could have constrained the support of μ_T^2 to $X \setminus X_T$ instead of X ; however, this is unnecessary by virtue of optimality in (10) and avoids using the difference of two basic semialgebraic sets which may not be basic semialgebraic.

This leads to the following infinite-dimensional primal LP on measures:

$$\begin{aligned}
p_o^* &= \sup \int 1 d\mu_T^1 - \int 1 d\mu_T^2 \\
\text{s.t. } & \beta \int_X v d\mu + \int_{X_T} v d\mu_T^1 + \int_X v d\mu_T^2 = \int_X v dx + \int_X \mathcal{L}_f v d\mu + \sum_{i=1}^m \int_X (\mathcal{L}_{G_i} v) d\sigma_i \quad \forall v \in C^1(X) \\
& \int p_i d\sigma_i + \int p_i d\hat{\sigma}_i - \bar{u} \int p_i d\mu = 0 \quad \forall i \in \{1, \dots, m\} \quad \forall p_i \in C(X) \\
& \text{spt } \mu \subset X, \text{ spt } \mu_T^1 \subset X_T, \text{ spt } \mu_T^2 \subset X, \\
& \text{spt } \sigma_i \subset X, \text{ spt } \hat{\sigma}_i \subset X, \quad \forall i \in \{1, \dots, m\} \\
& \mu \geq 0, \mu_T^1 \geq 0, \mu_T^2 \geq 0, \\
& \sigma_i \geq 0, \hat{\sigma}_i \geq 0, \quad \forall i \in \{1, \dots, m\},
\end{aligned} \tag{10}$$

where the supremum is over

$$(\mu, \mu_T^1, \mu_T^2, \sigma_1, \dots, \sigma_m, \hat{\sigma}_1, \dots, \hat{\sigma}_m) \in M(X) \times M(X_T) \times M(X) \times M(X)^m \times M(X)^m.$$

The dual LP on continuous functions provides approximations from above to the value function and therefore outer approximations of the ROA. The dual LP reads

$$\begin{aligned}
d_o^* &= \inf \int_X v(x) dx \\
\text{s.t. } & \mathcal{L}_f v(x) + \bar{u} \sum_{i=1}^m p_i(x) \leq \beta v(x), \quad \forall x \in X \\
& p_i(x) \geq \mathcal{L}_{G_i} v(x), \quad \forall x \in X, \quad i \in \{1, \dots, m\} \\
& p_i(x) \geq 0, \quad \forall x \in X, \quad i \in \{1, \dots, m\} \\
& v(x) \geq 1, \quad \forall x \in X_T \\
& v(x) \geq -1, \quad \forall x \in X
\end{aligned} \tag{11}$$

where the infimum is over $(v, p_1, \dots, p_m) \in C^1(X) \times C(X)^m$.

The following Lemma shows that the zero super-level set of any function $v \in C^1$ feasible in (11) is an outer approximation to the ROA.

Lemma 1 *If $v \in C^1(X)$ is feasible in (11), then $X_0 \subset \{x : v(x) > 0\}$.*

Proof: Fix an $x_0 \in X_0$. Then by definition there exists m control functions $u_i(\cdot | x_0) \in [0, \bar{u}]$ and a time $\tau \in [0, \infty)$ such that $x(\tau | x_0) \in X_T$. Therefore $x(t | x_0) \in X$ for all $t \in [0, \tau]$ and consequently, using the constraints of (11),

$$\begin{aligned}
\frac{d}{dt} v(x(t | x_0)) &= \mathcal{L}_f v(x(t | x_0)) + \sum_{i=1}^m \mathcal{L}_{G_i} v(x(t | x_0)) u_i(x(t | x_0)) \leq \mathcal{L}_f v + \sum_{i=1}^m p_i u_i \\
&\leq \mathcal{L}_f v + \bar{u} \sum_{i=1}^m p_i \leq \beta v(x(t | x_0))
\end{aligned}$$

for all $t \in [0, \tau)$. Using the Gronwall's inequality we have $v(x(\tau | x_0)) \leq e^{\beta\tau} v(x_0)$ and therefore

$$v(x_0) \geq e^{-\beta\tau} v(x(\tau | x_0)) > 0$$

as desired. Here the last inequality follows from the fact that $v(x) \geq 1$ on X_T . \square

4.1 Finite dimensional relaxations

The infinite dimensional LPs (10) and (11) can be solved only approximately; a systematic way of approximating them is the so-called Lasserre hierarchy of semidefinite programming (SDP) relaxations [13], originally introduced for static polynomial optimization and later extended to the dynamic case [14].

Instead of optimizing over measures, this hierarchy takes only finitely many moments of the measures for the primal problem (10) while imposing conditions necessary for these truncated moment sequences to be feasible in the primal LP via the so-called moment and localizing matrices. On the dual side, the function space is restricted to polynomials of a given degree while imposing sufficient conditions for the non-negativity via sum-of-squares conditions. One then refers to the relaxation of order k when the first $2k$ moments of the measures are taken on the primal side and polynomials of total degree up to $2k$ on the dual side. Please refer to [7] or [16] for more details on how to construct the relaxations for this particular problem or to [13] for a general treatment.

Let $v_k(\cdot)$ be a polynomial of degree $2k$ solving the k^{th} order relaxation of the dual SDP (11). Then this polynomial is feasible in (11) and therefore, in view of Lemma 1, we can define the k^{th} order outer approximation of X_0 by

$$X_{0k}^{\mathcal{O}} := \{x \in X : v_k(x) > 0\}. \quad (12)$$

The following theorem states that the running intersection $\cap_{i=1}^k X_{0i}^{\mathcal{O}}$ converges monotonically to the set no smaller than $\{x \mid V(x) \geq 0\}$, which is the union of the ROA X_0 and the set of all states which can be kept within X forever.

Theorem 1 *The following statements are true: $\cap_{i=1}^k X_{0i}^{\mathcal{O}} \supset X_0$ for all $k \geq 1$ and*

$$\lim_{k \rightarrow \infty} \text{vol}(\cap_{i=1}^k X_{0i}^{\mathcal{O}} \setminus \{x \mid V(x) \geq 0\}) = 0.$$

Proof: The proof follows the convergence of $v_k(\cdot)$ to $V(\cdot)$ in L_1 norm which can be established using similar reasoning as in [11, Theorem 6]; details are omitted for brevity. \square

Note that in the case where the volume of the set of states which can be kept in X forever but cannot be steered to X_T is positive, the set to which the running intersection of $X_{0k}^{\mathcal{O}}$ converges can be strictly larger than X_0 (in the sense of positive volume difference); nevertheless by virtue of optimality in (10) the controller attaining the infimum in (10) generates X_0 (i.e., admissibly steers any state in X_0 to X_T).

5 Controller extraction

In this section we describe how a polynomial controller approximately feasible in the primal LP (10) can be extracted from the solution to the finite-dimensional relaxations of (10). Given a truncated moment sequence solving the k^{th} primal relaxation, the idea is to find, component-by-component, polynomial controllers $u_i^k(x)$, $i = 1, \dots, m$, of a predefined total

degree $\deg(u_i^k) \leq k$ that approximately satisfy the relation (8). Details of this procedure are described below.

First, note that satisfying relation (8) is equivalent to satisfying

$$\int_X v(x) u_i^k(x) d\mu(x) = \int_X v(x) d\sigma_i(x)$$

for all polynomials⁶ $v(\cdot)$ and therefore, by linearity, it is equivalent to the linear equation

$$\int_X x^\alpha u_i^k(x) d\mu(x) = \int_X x^\alpha d\sigma_i(x), \quad (13)$$

where the multindex $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{N}^n$ runs over all nonnegative integer n -tuples. The data available to us after solving the k^{th} order primal relaxation are the first $2k$ moments⁷ of the measures μ and σ_i . The approach for controller extraction of [16] consists of setting $\deg(u_i^k) := k$ and satisfying the equation (13) exactly for all moments up to degree k , i.e., for all $\alpha \in \mathbb{N}^n$ such that $\sum_{i=1}^n \alpha_i \leq k$. For this choice of $\deg(u_i^k)$, the linear equation (13) takes a particularly simple form

$$M_k(\mathbf{y}^{2k}) \mathbf{u}_i^k = \boldsymbol{\sigma}_i^k, \quad (14)$$

where $M_k(\mathbf{y}^{2k})$ is the $\binom{n+k}{n} \times \binom{n+k}{n}$ moment matrix (see [13] for definition) associated to \mathbf{y}^{2k} , the vector of the first $2k$ moments of μ , $\boldsymbol{\sigma}_i^k$ is the vector of the first k moments of σ_i , and \mathbf{u}_i^k is the vector of the coefficients of the polynomial $u_i^k(\cdot)$.

However, this extraction procedure does not ensure the satisfaction of the input constraints and indeed typically leads to controllers violating the input constraints. One remedy is to pose input constraints on the polynomial $u_i^k(x)$ as sum-of-squares (SOS) constraints and minimize over \mathbf{u}_i^k the moment mismatch $\|M_k(\mathbf{y}^{2k}) \mathbf{u}_i^k - \boldsymbol{\sigma}_i^k\|_2$ subject to these input constraints; this immediately translates to an SDP problem, as proposed originally in [9]. Denoting by \hat{u}_i^k the “true” controller, coefficients of which satisfy the equation (14), this approach is equivalent to minimizing the $L^2(\mu)$ error⁸ $\int_X (u_i^k(x) - \hat{u}_i^k(x))^2 d\mu(x)$. The problem of this extraction procedure is now apparent – the $L^2(\mu)$ criterion weights subsets of X according to the occupation measure μ and therefore will penalize more those subsets of X where the trajectories spend a large amount of time and penalize less those subsets where the trajectories spend little time. This is clearly undesirable and is likely to lead to a poor closed-loop performance of the extracted controller. What we would like to have is a uniform penalization over the constraint set X or even better over the region of attraction X_0 (since any control is futile outside X_0 anyway). This leads to the following two-step procedure that we propose:

1. Extract the controller polynomial \hat{u}_i^k satisfying exactly (14) but violating the input constraints.

⁶This follows from the compactness of the constraint set X and the fact that the polynomials are dense in $C(X)$ w.r.t. the supremum norm on compact sets.

⁷The first $2k$ moments belong to a measure feasible in (10) only asymptotically, i.e. when $k \rightarrow \infty$; at a finite relaxation order k the first $2k$ moments may not belong to a measure feasible in (10) (or any nonnegative measure at all).

⁸The integral $\int_X (u_i^k(x) - \hat{u}_i^k(x))^2 d\mu(x)$ should be understood symbolically, replacing moments of the “true” occupation measure μ by \mathbf{y}^{2k} .

2. Minimize the $L^2(\nu)$ error $\int_X (u_i^k(x) - \hat{u}_i^k(x))^2 d\nu(x)$ subject to the input constraints,

where ν is a reference measure, ideally equal to the uniform measure on the ROA X_0 . The second step can be carried out only approximately since the ROA, let alone the uniform distribution on it, are not known in advance. Here we use the first $2k$ moments of the uniform measure on X as a rough proxy for the uniform measure on X_0 . This leads to the following SOS optimization problem for controller extraction:

$$\begin{aligned} \min_{u_i^k, s_1^i, s_2^i} \quad & \mathbf{u}_i^k M_1 \mathbf{u}_i^k - 2\mathbf{u}_i^k M_2 \hat{\mathbf{u}}_i^k \\ \text{s.t.} \quad & u_i^k - g_X s_1^i = \text{SOS} \\ & \bar{u} - u_i^k - g_X s_2^i = \text{SOS} \\ & s_1^i = \text{SOS}, \quad s_2^i = \text{SOS}, \end{aligned} \tag{15}$$

where M_1 and M_2 are respectively the $\binom{n+\deg(u_i^k)}{n} \times \binom{n+\deg(u_i^k)}{n}$ and $\binom{n+\deg(u_i^k)}{n} \times \binom{n+k}{n}$ top-left sub-blocks of $M_k(\mathbf{y}_\nu^{2k})$, the moment matrix associated to the first $2k$ moments of the reference measure ν . The optimization is over the coefficients \mathbf{u}_i^k of the controller polynomial $u_i^k(x)$ and the coefficients of the SOS polynomial multipliers $s_1^i(x)$ and $s_2^i(x)$.

It is worth noting that if a low-complexity controller is desired one can enforce sparsity on the coefficient vector of the polynomial u_i^k by adding an l_1 -regularization term $\gamma \|\mathbf{u}_i^k\|_1$ for some $\gamma > 0$ to the objective of (15) and/or fix an a priori sparsity pattern of \mathbf{u}_i^k .

To state a convergence result for the extracted controllers, we abuse notation and set $\int v(x) d\mu_k := \sum_\alpha v_\alpha \mathbf{y}^{2k}(\alpha)$ for a polynomial $v(x) = \sum_\alpha v_\alpha x^\alpha$ of total degree less than $2k$.

The following convergence result for the controllers \hat{u}_i^k was established in [16]:

Lemma 2 *There exists a subsequence $\{k_j\}_{j=1}^\infty$ such that*

$$\lim_{j \rightarrow \infty} \int_X v(x) \hat{u}_i^{k_j}(x) d\mu_{k_j}(x) = \int_X v(x) \hat{u}_i^*(x) d\mu^*(x),$$

for all polynomials $v(x)$, where $u_i^*(x) d\mu^*(x) = d\sigma_i^*(x)$ and (μ^*, σ_i^*) is a part of an optimal solution to the primal LP (10).

Establishing a stronger notion of convergence and extending the convergence result to the controllers $u_i^k(x)$ satisfying the input constraints obtained from (15) is currently investigated by the authors.

Remark 1 *The method of controller extraction ensuring the satisfaction of input constraints using sum-of-squares programming is not the only one possible but is particularly convenient for subsequent inner approximation computation since the extracted controller is polynomial. For instance, other viable approach is to simply clip the controller $\hat{u}_i^k(x)$ on the input constraint set; in this case the closed-loop dynamics is piecewise polynomial defined over a semi-algebraic partition and it is still possible to compute inner approximations by a straightforward modification of the approach of Section 6.*

6 Inner approximations

Given a controller $u(x)$ satisfying input constraints extracted from the outer approximations, all we need to do to obtain *inner approximations* to the ROA X_0 is to compute inner approximations of the ROA for the closed-loop system

$$\dot{x} = \bar{f}(x) := f(x) + G(x)u(x). \quad (16)$$

Remark 2 *All results of this section apply to uncontrolled systems $\dot{x} = \bar{f}(x)$ with an arbitrary polynomial vector field \bar{f} , not only those of the special form (16).*

In order to compute the inner approximations we combine ideas of our two previous works [10] and [11] for computation of inner approximations to the ROA in a finite-time setting and outer approximations to the maximum controlled invariant set, respectively. This combination of the two approaches retains strong theoretical guarantees of both and seems to exhibit faster convergence of the finite-dimensional SDP relaxations.

The key idea of [10] that we adopt here is to characterize the *complement* of the ROA

$$X_0^c := X \setminus X_0.$$

By continuity of solutions to (16), this set is equal to

$$X_0^c = \left\{ x_0 \in X : \exists x(\cdot) \text{ s.t. } \dot{x}(t) = \bar{f}(x(t)) \text{ and} \right. \\ \left. \exists \tau \in [0, \infty) \text{ s.t. } x(\tau) \in X_\partial \text{ and/or } x(t) \in X_T^c \forall t \in [0, \infty) \right\},$$

where

$$X_T^c := \{x \in \mathbb{R}^n : g_X(x) \geq 0, g_T(x) \leq 0\}$$

is the complement of X_T in X and

$$X_\partial := \{x \in \mathbb{R}^n : g_X(x) = 0\}.$$

In order to compute outer approximations of the complement ROA X_0^c we study families of trajectories starting from an initial distribution μ_0 . The time evolution is captured by the discounted occupation measure (6) and the distribution at the stopping time $\tau(\cdot) \in [0, \infty]$ by the final measure μ_T (5). The three measures are again linked by the discounted Liouville's equation (6).

The complement ROA X_0^c (and hence also the ROA X_0) is then obtained by maximizing the mass of the initial measure subject to the discounted Liouville equation (6), support constraints $\text{spt } \mu \subset X_T^c$, $\text{spt } \mu_0 \subset X_T^c$, $\mu_T \subset X_\partial$, and subject to the constraint that the initial measure is dominated by the Lebesgue measure (which is equivalent to saying that the density of the initial measure is below one). The last constraint is equivalent to the existence of a nonnegative slack measure $\hat{\mu}_0$ such that $\int_X w d\mu_0 + \int_X w d\hat{\mu}_0 = \int_X w dx$ for all $w \in C(X)$. This optimization procedure yields an optimal initial measure with density (w.r.t. the Lebesgue measure) equal to one on X_0^c and zero otherwise.

Writing the above formally leads to the following primal infinite-dimensional LP on measures

$$\begin{aligned}
p_I^* &= \sup \int 1 d\mu_0 \\
\text{s.t. } & \int_X v d\mu_T + \beta \int_X v d\mu = \int_X v d\mu_0 + \int_X \nabla v \cdot \bar{f} d\mu, & \forall v \in C^1(X), \\
& \int_X w d\mu_0 + \int_X w d\hat{\mu}_0 = \int_X w dx, & \forall w \in C(X), \\
& \text{spt } \mu \subset X_T^c, \text{ spt } \mu_0 \subset X_T^c, \text{ spt } \mu_T \subset X_\partial, \text{ spt } \hat{\mu}_0 \subset X_T^c, \\
& \mu_0 \geq 0, \mu \geq 0, \mu_T \geq 0, \hat{\mu}_0 \geq 0,
\end{aligned} \tag{17}$$

where the supremum is over the vector of nonnegative measures

$$(\mu_0, \mu, \mu_T, \hat{\mu}_0) \in M(X_T^c) \times M(X_T^c) \times M(X_\partial) \times M(X_T^c).$$

The dual infinite-dimensional linear program on continuous functions reads

$$\begin{aligned}
d_I^* &= \inf \int_X w(x) dx \\
\text{s.t. } & \nabla v(x) \cdot \bar{f}(x) \leq \beta v(x), & \forall x \in X_T^c, \\
& w(x) \geq v(x) + 1, & \forall x \in X_T^c, \\
& w(x) \geq 0, & \forall x \in X_T^c, \\
& v(x) \geq 0, & \forall x \in X_\partial,
\end{aligned} \tag{18}$$

where the infimum is over the pair of functions $(v, w) \in C^1(X) \times C(X)$.

The following lemma establishes that the set $\{x \in X : v(x) < 0\}$ for any function $v \in C^1(X)$ feasible in (18) provides an inner approximation to the ROA X_0 .

Lemma 3 *If $v \in C^1(X)$ is feasible in (18), then $\{x \in X : v(x) < 0\} \subset X_0$.*

Proof: We will prove the contrapositive, i.e., that whenever $x_0 \in X_0^c$, then $v(x) \geq 0$. For that we distinguish two cases.

First, assume that $x_0 \in X_0^c$ and that $x(t) \in X$ for all $t \in [0, \infty)$. In that case necessarily also $x(t) \in X_T^c$ for all $t \in [0, \infty)$ and the first constraint of (18) implies that $\frac{d}{dt}v(x(t)) \leq \beta v(x(t))$ for all $t \in [0, \infty)$. Using Gronwall's inequality, this implies that $v(x(t)) \leq v(x_0)e^{\beta t}$ or $v(x_0) \geq e^{-\beta t}v(x(t))$ for all $t \in [0, \infty)$. Since X_T^c is compact and v continuous this implies that $v(x(t))$ is bounded and therefore necessarily $v(x_0) \geq 0$.

Second, assume that there exists a time $\tau \in [0, \infty)$ such that $x(\tau) \in X_\partial$. Then the second constraint of (18) implies that $v(x(\tau)) \geq 0$, and therefore, using again the first constraint of (18) and Gronwall's inequality, we get $0 \leq v(x(\tau)) \leq e^{\beta\tau}v(x_0)$ and therefore $v(x_0) \geq e^{-\beta\tau}v(x(\tau)) \geq 0$ as desired. \square

6.1 Choice of the discount factor β

The LPs (17) and (18) depend on the discount factor $\beta > 0$ which is a free parameter. Theoretical results pertaining to the infinite-dimensional LPs (17) and (18) and convergence guarantees of their finite-dimensional relaxations do not depend on the value of β as long as

it is strictly positive. However, the *speed of convergence* and the *quality* (i.e., the tightness) of the ROA estimates coming out of the finite-dimensional relaxations does depend on β .

This dependence can be exploited to speed-up the convergence of the finite-dimensional relaxations by observing that if a vector of measures $(\mu_0, \mu, \mu_T, \hat{\mu}_0)$ is feasible in the primal LP (17) with a given value of the discount factor $\beta > 0$, then for any other value of $\beta' > 0$ there must exist discounted occupation and terminal measures $\mu^{\beta'}$ and $\mu_T^{\beta'}$ such that the vector of measures $(\mu_0, \mu^{\beta'}, \mu_T^{\beta'}, \hat{\mu}_0)$ is also feasible in (17). Therefore, instead of considering a single value of β we can define a vector $\boldsymbol{\beta} = (\beta_1, \dots, \beta_{n_\beta})$, $\beta_i > 0$, and optimize over vectors of discounted occupation and terminal measures $\boldsymbol{\mu} = (\mu^1, \dots, \mu^{n_\beta})$ and $\boldsymbol{\mu}_T = (\mu_T^1, \dots, \mu_T^{n_\beta})$, components of which have to satisfy the discounted Liouville equation with *one common* initial measure μ_0 . This leads to the following modified primal LP

$$\begin{aligned}
p_I^* &= \sup \int 1 d\mu_0 \\
\text{s.t.} \quad & \int_X v d\mu_T^i + \beta_i \int_X v d\mu^i = \int_X v d\mu_0 + \int_X \nabla v \cdot \bar{f} d\mu^i, \quad i \in \{1, \dots, n_\beta\}, \quad \forall v \in C^1(X), \\
& \int_X w d\mu_0 + \int_X w d\hat{\mu}_0 = \int_X w dx, \quad \forall w \in C(X), \\
& \text{spt } \mu_0 \subset X_T^c, \quad \text{spt } \hat{\mu}_0 \subset X_T^c, \\
& \text{spt } \mu^i \subset X_T^c, \quad \text{spt } \mu_T^i \subset X_\partial, \quad i \in \{1, \dots, n_\beta\}, \\
& \mu_0 \geq 0, \quad \hat{\mu}_0 \geq 0 \\
& \mu^i \geq 0, \quad \mu_T^i \geq 0, \quad i \in \{1, \dots, n_\beta\}
\end{aligned} \tag{19}$$

where the supremum is over the vector of nonnegative measures

$$(\mu_0, \boldsymbol{\mu}, \boldsymbol{\mu}_T, \hat{\mu}_0) \in M(X_T^c) \times M(X_T^c)^{n_\beta} \times M(X_\partial)^{n_\beta} \times M(X_T^c).$$

As already remarked the optimal values of the infinite-dimensional primal LPs (17) and (19) are the same; however, the finite-dimensional SDP relaxations of (19) are likely to converge faster than those of (17) (and will never converge slower provided that the β used in (17) is among the components of the vector $\boldsymbol{\beta}$ used in (19)).

For completeness we also state the dual infinite-dimensional LP on continuous functions where the function v from (18) is replaced by a vector of functions $\boldsymbol{v} = (v_1, \dots, v_{n_\beta})$:

$$\begin{aligned}
d_I^* &= \inf \int_X w(x) dx \\
\text{s.t.} \quad & \nabla v_i(x) \cdot \bar{f}(x) \leq \beta_i v_i(x), \quad \forall x \in X_T^c, \quad i \in \{1, \dots, n_\beta\}, \\
& w(x) \geq 1 + \sum_{i=1}^{n_\beta} v_i(x), \quad \forall x \in X_T^c, \\
& w(x) \geq 0, \quad \forall x \in X_T^c, \\
& v_i(x) \geq 0, \quad \forall x \in X_\partial, \quad i \in \{1, \dots, n_\beta\},
\end{aligned} \tag{20}$$

where the infimum is over the vector of functions $(\boldsymbol{v}, w) \in C^1(X)^{n_\beta} \times C(X)$.

Remark 3 For any functions v_1, \dots, v_{n_β} feasible in (20) the results of Lemmata 3 and 4 and Corollaries 1 and 2 hold with the function v replaced by $\sum_{i=1}^{n_\beta} v_i$.

6.2 Finite-dimensional relaxations

The infinite dimensional primal and dual LPs (17) and (18) give rise to finite-dimensional SDP relaxations in exactly the same fashion as outlined in Section 4.1 or described in more detail in, for instance, [7, Section VI]. Further details are omitted for brevity.

Let v_i^k , $i = 1, \dots, n_\beta$, and w_k denote the polynomials of degree $2k$ solving the k^{th} order SDP relaxation of the dual LP (20) and let

$$X_{0k}^{\mathcal{I}} := \{x \in X : \sum_{i=1}^{n_\beta} v_i(x) < 0\} \quad (21)$$

denote the k^{th} order inner approximation (since $X_{0k}^{\mathcal{I}} \subset X_0$ by Lemma 3 and Remark 3).

The following theorem summarizes convergence properties of the finite-dimensional relaxations. Let $X_0^{\text{cl}} \subset X_0$ be the *closed-loop* ROA associated to the closed-loop system (16) (defined analogously as ROA X_0).

Theorem 2 *The following convergence properties hold:*

- *The optimal values of the primal and dual SDP relaxations converge to the volume of $X \setminus X_0^{\text{cl}}$ as the relaxation order k tends to infinity.*
- *The functions w_k converge in L^1 to the indicator function of the set $X \setminus X_0^{\text{cl}}$.*
- *The sets $X_{0k}^{\mathcal{I}}$ converge from inside to X_0^{cl} in the sense that the volume of $X_0^{\text{cl}} \setminus X_{0k}^{\mathcal{I}}$ tends to zero as the relaxation order k tends to infinity.*

Proof: The proof follows by the same arguments as Theorem 6, Corollary 7 and Theorem 8 of [10] using Theorems 2 and 4 of [11] in place of Theorems 1 and 5 of [10]. \square

6.3 Invariance of the inner approximations

In this section we investigate under what conditions the inner approximations $X_{0k}^{\mathcal{I}}$ are controlled invariant for the system (1) and positively invariant for the system (16). Recall that a subset of \mathbb{R}^n is called *positively invariant* for an uncontrolled ODE if trajectories starting in the set remain in the set forever. Similarly a subset of \mathbb{R}^n is called *controlled invariant* for a controlled ODE if there exists an admissible control input such that the trajectories starting in the set remain in the set forever.

The following Lemma leads almost immediately to the characterization of the invariance of $X_{0k}^{\mathcal{I}}$; it says that trajectories starting in $X_{0k}^{\mathcal{I}}$ stay there until they reach the target set X_T for the closed-loop system (16).

Lemma 4 *If $x(0) \in X_{0k}^{\mathcal{I}}$, then $x(t) \in X_{0k}^{\mathcal{I}}$ for all $t \in [0, \tau]$, where $\tau = \inf\{s \in [0, \infty) \mid x(s) \in X_T\} < \infty$ is the first time that $x(t)$ reaches X_T , and $x(t)$ is the solution to (16).*

Proof: If $x(0) \in X_T$ then there is nothing to prove. Assume therefore $x(0) \in X \setminus X_T \subset X_T^c$. By Lemma 3 (and its analogy for the problem (20)) if $x(0) \in X_{0k}^I$, then, by the definition of the ROA X_0 , $\tau = \inf\{s \in [0, \infty) \mid x(s) \in X_T\}$ is finite and $x(t) \in X$ for all $t \in [0, \tau]$. This implies that $x(t) \in X_T^c$ for all $t \in [0, \tau)$ and therefore $\sum_{i=1}^{n_\beta} v_i(x(t)) < \sum_{i=1}^{n_\beta} v_i(x(0)) < 0$, where the first inequality follows from the first constraint of (20) and the Gronwall's inequality and the second one from the definition of X_{0k}^I (21). This proves the claim. \square

The following two immediate Corollaries give conditions under which X_{0k}^I is invariant.

Corollary 1 *If the target set X_T is controlled / positively invariant for the system (1) / (16), then so is the set $X_{0k}^I \cup X_T$.*

Corollary 2 *If $\text{closure}(X_T) \subset X_{0k}^I$, then X_{0k}^I is controlled / positively invariant for the system (1) / (16).*

7 Numerical examples

In this section we present numerical examples illustrating the approach. As a modeling tool we used Gloptipoly [8] which allows to model directly the primal problems on measures. The outer and inner approximations are then extracted from the dual variables provided by a primal-dual SDP solver (in our case MOSEK). Equivalently one can model the dual SOS problems (in our case using YALMIP [15] or SOSOPT [19]) and extract the primal moment vector (which is needed to obtain the approximate control law) as a part of the dual variables associated with the constraints of the dual SDP.

7.1 Nonlinear double integrator

As our first example we consider the nonlinear double integrator

$$\begin{aligned}\dot{x}_1 &= x_2 + 0.1x_1^3 \\ \dot{x}_2 &= 0.3u\end{aligned}$$

subject to the constraints $u \in [-1, 1]$ and $x \in X := \{x : \|x\|_2 < 1.2\}$. As a terminal set we take a small neighbourhood of the origin $X_T := \{x : \|x\|_2 < 0.1\}$. First we obtain two approximate polynomial controllers of degree four by solving the fourth order (i.e., considering moments up to total degree eight) SDP relaxation of the two primal LPs (10). Then we obtain inner approximations from the eighth order (corresponding to degrees of w and v_i , $i = 1, \dots, 5$, equal to 16) SDP relaxation of the dual LP (20) where we choose the vector $\beta = (10, 1, 0.1, 0.01, 0.001)$; the inner approximations given by (21) are compared in Figure 1 with an outer approximation (12) obtained from solving an SDP relaxation of the dual LP (11) with degrees of w and v equal to 16. We can see that, for this example, both the inner and outer estimates are almost tight. Computation times of the SDP relaxations are reported in Table 1.

Table 1: Nonlinear double integrator – Computation time comparison for different degrees $d = 2k$ of the polynomials in the SDP relaxations of the outer dual LP (11) and the inner dual LP (20). Reported is pure solver time of the MOSEK SDP solver (excluding Gloptipoly and Yalmip parsing time). Larger solve time for the inner approximations is because of the larger number of decision variables and constraints in (20) since in (20) there is one polynomial v_i associated to each of the five values of the discount factor β_i .

d	6	8	10	12	16
Inner	0.42 s	0.67 s	1.03 s	1.61 s	4.41 s
Outer	0.17 s	0.23 s	0.37 s	0.66 s	1.02 s

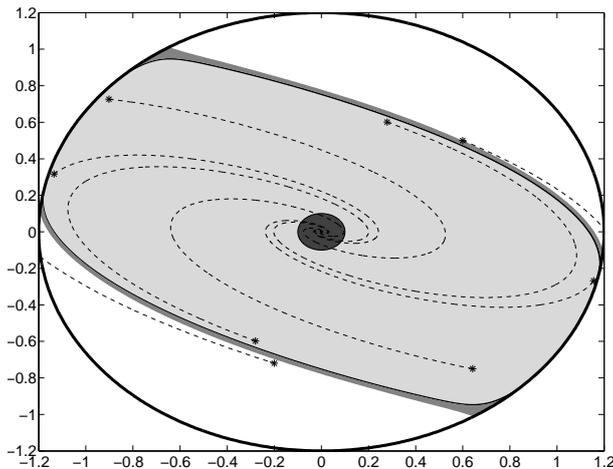


Figure 1: Nonlinear double integrator – light: inner approximation (21) with v_i of 16; darker: outer approximations (12) with v of degree 16; dark, small: target set; thick line: constraint set boundary; dashed thin lines: closed-loop trajectories.

7.2 Controlled 3D Van der Pol oscillator

As our second example we consider a controlled Van der Pol oscillator in three dimensions given by

$$\begin{aligned}
 \dot{x}_1 &= -2x_2 \\
 \dot{x}_2 &= 0.8x_1 - 2.1x_2 + x_3 + 10x_1^2x_2 \\
 \dot{x}_3 &= -x_3 + x_3^3 + 0.5u
 \end{aligned}$$

subject to the constraints $u \in [-1, 1]$ and $x \in X = \{x : \|x\|_2 < 1\}$. As a terminal set we take a small neighbourhood of the origin $X_T = \{x : \|x\|_2 < 0.1\}$. First we extract an approximate controller $u(x)$ of degree four by solving an SDP relaxation of fourth order (i.e., considering moments up to total degree eight) of the primal LP (10). After that we compute an inner approximation by solving an SDP relaxation of the dual LP (20) with $\beta = (1, 0.1, 0.01, 0.001)$ and polynomials w and v_i , $i = 1, \dots, 4$, of degree 10. To assess the tightness of the inner approximation we compute an outer approximation by solving an SDP relaxation of the dual LP (11) with w and v polynomials of degree 16. Figure 2 shows the comparison and also several trajectories of the closed-loop system which, as expected, converge to the target

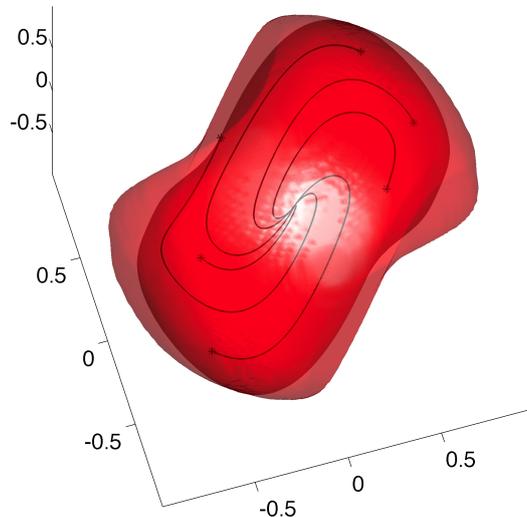


Figure 2: Controlled 3D Van der Pol oscillator – inner approximation of degree 10 (dark, smaller); outer approximation of degree 16 (light, larger); closed-loop trajectories (black). The degree of the polynomial controller generating the inner approximation is four.

set (not shown) whenever starting in the inner approximation. We observe a relatively good tightness of the inner approximation. Computation times are reported in Table 2.

Table 2: Controlled 3D Van der Pol oscillator – Computation time comparison for different degrees $d = 2k$ of the polynomials in the SDP relaxations of the outer dual LP (11) and the inner dual LP (20). The same comments as for Table 1 apply.

d	6	8	10	12	16
Inner	1.16 s	3.6 s	12.3 s	32.7 s	213 s
Outer	0.48 s	0.88 s	2.8 s	7.4 s	54.4 s

8 Conclusion

We have presented a method for computing inner and outer approximations of the region of attraction for input-affine polynomial dynamical systems subject to box input and semi-algebraic state constraints. The method combines and extends the ideas of our previous works [7, 10, 11] and the controller extraction idea of [16]. The inner approximations are controlled invariant provided that the target set is included in the inner approximation and/or itself controlled invariant.

The approach is based purely on convex optimization, in fact semidefinite programming (SDP), and is tractable for systems of moderate state-space dimension (say up to n equal to 6 or 8) using interior point solvers, in our case MOSEK; the number of control inputs

m plays a secondary role in the computation complexity since the number of variables and constraints in the SDP relaxations grows only linearly with m . The linear growth is due to the fact that one measure on \mathbb{R}^n is associated to each control input, where n is the state dimension. The total number of variables therefore grows as $O(mn^d)$ when the polynomial degree is held fixed and as $O(md^n)$ when m and n are held fixed; this is a significantly more favourable growth rate than $O((m+n)^d)$ or $O(d^{n+m})$ of [7, 10, 11]. Larger systems could be tackled using first-order methods; for instance SDPNAL [20] shows promising results on our problem class (after suitable preconditioning of problem data).

The approach can be readily extended to trigonometric, rational or rational-trigonometric dynamics with the same theoretical guarantees since a variant of the Putinar Positivstellensatz [18] (which Theorem 2 hinges on) holds for trigonometric polynomials as well [5].

At present no proof of convergence of the inner approximations X_{0k}^I to the ROA X_0 is available. This is since Lemma 2 provides only very weak convergence guarantees of the extracted controller to a controller generating the ROA X_0 . Strengthening Lemma 2 or the possibility of proving the convergence by other means is currently investigated by the authors.

Appendix A

Here we describe how the presented approach can be extended to handle the situation where the state constraint set X and/or the target set X_T are given by multiple polynomial inequalities. Assume therefore for this section that

$$X = \{x : g_X^i > 0, i = 1, \dots, n_X\}, \quad X_T = \{x : g_T^i > 0, i = 1, \dots, n_T\} \subset X,$$

and that the set

$$\bar{X} = \{x : g_X^i \geq 0, i = 1, \dots, n_X\}$$

is compact. The primal and dual LPs (10) and (11) providing the outer approximations stay the same since the set \bar{X} is compact basic semialgebraic. A slight modification is needed for the primal and dual LPs (17) and (18) providing the inner approximations. There, with multiple constraints, we define the sets X_∂ and X_T^c as

$$X_\partial = \bigcup_{i=1}^{n_X} X_\partial^i := \bigcup_{i=1}^{n_X} \{x : g_X^i = 0, g_X^j \geq 0, j \in \{1, \dots, n_X\} \setminus \{i\}\},$$

$$X_T^c = \bigcup_{i=1}^{n_T} X_T^{c_i} := \bigcup_{i=1}^{n_T} \{x : g_T^i \leq 0, g_X^j \geq 0, j \in \{1, \dots, n_X\}\}.$$

Since \bar{X} is compact, so are X_∂^i and $X_T^{c_i}$. The sets X_∂ and X_T^c are therefore unions of compact basic semialgebraic sets. In the primal LP the measures $\mu_0, \hat{\mu}_0$ and μ with the supports in X_T^c are decomposed as the sum of n_T measures each with the support in $X_T^{c_i}$ and analogously the terminal measure with the support in X_∂ is decomposed as the sum of n_X measures with the supports in X_∂^i . In the dual LP (18) this translates to imposing the inequalities for each $X_T^{c_i}$ and X_∂^i ; for instance, the first inequality of (18) now translates to the n_T inequalities $\nabla v \cdot \bar{f}(x) \leq \beta v(x) \forall x \in X_T^{c_i}, i = 1, \dots, n_T$.

Appendix B

This Appendix derives the discounted Liouville equation (6). For any test function $v \in C^1(X)$ we have

$$\begin{aligned}
 \int_X \nabla v(x) \cdot \bar{f}(x) d\mu(x) &= \int_X \int_0^{\tau(x_0)} e^{-\beta t} \nabla v(t|x_0) \cdot \bar{f}(x(t|x_0)) dt d\mu_0(x_0) \\
 &= \int_X \int_0^{\tau(x_0)} e^{-\beta t} \frac{d}{dt} v(x(t|x_0)) dt d\mu_0(x_0) \\
 &= \beta \int_X \int_0^{\tau(x_0)} e^{-\beta t} v(x(t|x_0)) dt d\mu_0(x_0) + \int_X e^{-\beta \tau} v(x(\tau(x_0)|x_0)) \mu_0(x_0) - \int_X v(x_0) \mu_0(x_0) \\
 &= \beta \int_X v(x) d\mu(x) + \int_X v(x) d\mu_T(x) - \int_X v(x) d\mu_0(x),
 \end{aligned}$$

which is exactly (6). Here we have used integration by parts in the third equality, and the definition of the initial, discounted occupation and terminal measures in the fourth.

References

- [1] J. P. Aubin, A. M. Bayen, P. Saint-Pierre. Viability theory: new directions. Springer-Verlag, Berlin, 2011.
- [2] G. Chesi. Domain of attraction; analysis and control via SOS programming. Lecture Notes in Control and Information Sciences, Vol. 415, Springer-Verlag, Berlin, 2011.
- [3] G. Chesi. LMI techniques for optimization over polynomials in control: a survey, IEEE Transactions on Automatic Control, 55: 2500-2510, 2010.
- [4] M. Claeys, D. Arzelier, D. Henrion, J.-B. Lasserre. Measures and LMI for impulsive optimal control with applications to space rendezvous problems. Proceedings of the American Control Conference, Montreal, 2012.
- [5] B. Dumitrescu. Positivstellensatz for trigonometric polynomials and multidimensional stability tests. IEEE Transactions on Circuits and Systems, 54:353-356, 2007.
- [6] D. Henrion, J. Daafouz, M. Claeys. Optimal switching control design for polynomial systems: an LMI approach. IEEE Conference on Decision and Control, Florence, Italy, 2013.
- [7] D. Henrion, M. Korda. Convex computation of the region of attraction of polynomial control systems. IEEE Transactions on Automatic Control, 59:297-312, 2014.
- [8] D. Henrion, J. B. Lasserre, and J. Löfberg. Gloptipoly 3: moments, optimization and semidefinite programming. Optimization Methods and Software, 24:761-779, 2009.
- [9] D. Henrion, J. B. Lasserre, M. Mevissen. Mean squared error minimization for inverse moment problems. Applied Mathematics & Optimization, 2014.

- [10] M. Korda, D. Henrion, C. N. Jones. Inner approximations of the region of attraction for polynomial dynamical systems, IFAC Symposium on Nonlinear Control Systems (NOLCOS), Toulouse, July 2013.
- [11] M. Korda, D. Henrion, C. N. Jones. Convex computation of the maximum controlled invariant set for polynomial control systems. [arXiv:1303.6469](#). A shorter version focusing only on the discrete-time was presented at the IEEE Conference on Decision and Control, Florence, Italy, 2013.
- [12] M. Korda, D. Henrion, C. N. Jones. Controller design and region of attraction estimation for nonlinear dynamical systems, 2014. [arXiv:1310.2213](#)
- [13] J. B. Lasserre. Moments, positive polynomials and their applications. Imperial College Press, London, UK, 2009.
- [14] J. B. Lasserre, D. Henrion, C. Prieur, E. Trélat. Nonlinear optimal control via occupation measures and LMI relaxations. *SIAM Journal on Control and Optimization*, 47:1643-1666, 2008.
- [15] J. Löfberg. YALMIP : A toolbox for modeling and optimization in MATLAB. In *Proc. IEEE CCA/ISIC/CACSD Conference*, Taipei, Taiwan, 2004.
- [16] A. Majumdar, R. Vasudevan, M. M. Tobenkin, R. Tedrake. Convex optimization of nonlinear feedback controllers via occupation measures. *Proceedings of Robotics Science and Systems (RSS)*, 2013. [arXiv: 1305.7484](#)
- [17] D. Q. Mayne, J. B. Rawlings, C.V. Rao, P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789-814, 2000.
- [18] M. Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42:969-984, 1993.
- [19] P. Seiler. SOSOPT: A toolbox for polynomial optimization. University of Minnesota, 2010.
- [20] X. Zhao, D. Sun, K. C. Toh. A Newton-CG Augmented Lagrangian Method for Semidefinite Programming, *SIAM J. Optimization*, 20:1737-1765, 2010.
- [21] V. I. Zubov. *Methods of A. M. Lyapunov and their application*. Noordhoff, Groningen, 1964.