2013 European Control Conference (ECC)
July 17-19, 2013, Zürich, Switzerland.

Convex computation of the region of attraction of polynomial control systems*

Didier Henrion[a,b,c], Milan Korda[d]

Abstract— **We address the long-standing problem of computing the region of attraction (ROA) of a target set (typically a neighborhood of an equilibrium point) of a controlled nonlinear system with polynomial dynamics and semialgebraic state and input constraints. We show that the ROA can be computed by solving a convex linear programming (LP) problem over the space of measures. In turn, this problem can be solved approximately via a classical converging hierarchy of convex finite-dimensional linear matrix inequalities (LMIs). Our approach is genuinely primal in the sense that convexity of the problem of computing the ROA is an outcome of optimizing directly over system trajectories. The dual LP on nonnegative continuous functions (approximated by polynomial sum-of-squares) allows us to generate a hierarchy of semialgebraic outer approximations of the ROA at the price of solving a sequence of LMI problems with asymptotically vanishing conservatism. This sharply contrasts with the existing literature which follows an exclusively dual Lyapunov approach yielding either nonconvex bilinear matrix inequalities or conservative LMI conditions. The approach is simple and readily applicable as the outer approximations are the outcome of a single semidefinite program with no additional data required besides the problem description.**

I. Introduction

Given a nonlinear control system, a state-constraint set and a target set (e.g. a neighborhood of an attracting orbit or an equilibrium point), the constrained controlled region of attraction (ROA) is the set of all initial states that can be steered with an admissible control to the target set without leaving the state-constraint set. The target set can be required to be reached at a given time or at any time before a given time[1]. The problem of computing the ROA (and variations thereof) lies at the heart of viability theory (see, e.g., [2]) and goes by many other names, e.g., the reach-avoid or target-hitting problem (see, e.g., [19]); in the language of viability theory the ROA itself is sometimes referred to as the capture basin [2].

We show that, in the case of polynomial dynamics, semi-algebraic state-constraint, input-constraint and target sets, the computation of the ROA boils down to solving an infinite-dimensional linear programming (LP) problem in the cone of nonnegative Borel measures. Our approach is genuinely primal in the sense that we optimize over state-space system trajectories modeled with occupation measures [17], [7].

In turn, this LP can be solved approximately by a classical hierarchy of finite-dimensional convex linear matrix inequality (LMI) relaxations. The dual LP on nonnegative continuous functions and its LMI relaxations on polynomial sum-of-squares provide explicitly an asymptotically converging sequence of nested semialgebraic outer approximations of the ROA.

Most of the existing literature on ROA computation follows Zubov's approach [20], [28], [10] and uses a dual Lyapunov certificate; see [25], the survey [8], Section 3.4 in [9], and more recently [24] and [4] and the references therein. These approaches either enforce convexity with conservative LMI conditions (whose conservatism is difficult if not impossible to evaluate systematically) or they rely on nonconvex bilinear matrix inequalities (BMIs), with all their inherent numerical difficulties. In contrast, we show in this paper that the problem of computing the ROA has actually a convex infinite-dimensional LP formulation, and that this LP can be solved with a hierarchy of convex finite-dimensional LMIs with asymptotically vanishing conservatism.

We believe that our approach is closer in spirit to set-oriented approaches [5], level-set and Hamilton-Jacobi approaches [15], [19], [21] or transfer operator approaches [26], even though we do not discretize w.r.t. time and/or space. In our approach, we model a measure with a finite number of its moments, which can be interpreted as a frequency-domain discretization (by analogy with Fourier coefficients which are moments w.r.t. the unit circle).

Another way to evaluate the contribution of our paper is to compare it with the recent works [13], [11] which deal with polynomial approximations of semialgebraic sets. In these references, the sets to be approximated are given a priori (as a polynomial sublevel set, or as a feasibility region of a polynomial matrix inequality). In contrast, in the current paper the set to be approximated (namely the ROA of a nonlinear dynamical system) is not known in advance, and our contribution can be understood as an application and extension of the techniques of references [13], [11] to sets defined implicitly by differential equations.

The benefits of our occupation measure approach are overall the convexity of the problem of finding the ROA, and the availability of publicly available software to implement and solve the hierarchy of LMI relaxations.

Our primary focus in this paper is the computation of the constrained finite-time controlled region of attraction of a given set. This problem is formally stated in Section II and solved using occupation measures in Section IV; the occupation measures themselves are introduced in Section III. A dual problem on the space of continuous functions is discussed in Section V. The hierarchy of finite-dimensional LMI relaxations of the infinite dimensional LP is described in Section VI. Numerical examples are presented in Section VII, and we conclude in Section VIII.

For space reasons, the more technical proofs are omitted in this conference version of the paper. Please refer to the

*An extended version of this work, [14], has been submitted for a journal publication in August 2012.

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

[b]Université de Toulouse; LAAS; F-31400 Toulouse; France.

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

[d]Laboratoire d'Automatique, École Polytechnique Fédérale de Lausanne, Station 9, CH-1015, Lausanne, Switzerland. milan.korda@epfl.ch

[1]The cases of an arbitrarily long but finite time and of asymptotic convergence are not addressed in this paper.

extended version [14] for detailed proofs.

II. Problem statement

Consider the control system

$$\dot{x}(t) = f(t, x(t), u(t)),\ x(t) \in X,\ u(t) \in U,\ t \in [0, T] \quad (1)$$

with a given polynomial vector field f with entries $f_i \in \mathbb{R}[t, x, u]$, $i \in \mathbb{Z}_{[1,n]}$, given final time $T > 0$, and given compact semialgebraic state and input constraints

$$\begin{array}{l} x(t) \in X := \{x \in \mathbb{R}^n : g_{X\,i}(x) \geq 0, i \in \mathbb{Z}_{[1,n_X]}\} \\ u(t) \in U := \{u \in \mathbb{R}^m : g_{U\,i}(u) \geq 0, i \in \mathbb{Z}_{[1,n_U]}\} \end{array} \quad (2)$$

for $t \in [0, T]$, where $g_{X\,i} \in \mathbb{R}[x]$, $g_{U\,i} \in \mathbb{R}[u]$, and $\mathbb{Z}_{[i,j]}$ denotes the set of consecutive integers $\{i, i{+}1, \ldots, j\}$. Given a compact semialgebraic target set

$$X_T := \{x \in \mathbb{R}^n \ :\ g_{T\,i}(x) \geq 0, i \in \mathbb{Z}_{[1,n_T]}\} \subset X,$$

with $g_{T\,i} \in \mathbb{R}[x]$, let

$$\begin{array}{l} \mathcal{X}(x_0) := \big\{x(\cdot)\ :\ x(t) = x_0 + \int_0^t f(\tau, x(\tau), u(\tau))d\tau, \\ \quad u(t) \in U,\, x(t) \in X,\, x(T) \in X_T,\, \forall t \in [0, T]\big\} \end{array} \quad (3)$$

denote the set of all absolutely continuous admissible controlled trajectories $x(\cdot)$ starting from x_0, generated by an admissible control $u(\cdot) \in L^1([0, T]; \mathbb{R}^m)$.

The constrained controlled region of attraction (ROA) is then defined as

$$X_0 := \{x_0 \in X \ :\ \mathcal{X}(x_0) \neq \emptyset\}. \quad (4)$$

In words, the ROA is the set of all initial conditions for which there exists an admissible controlled trajectory. By construction the set X_0 is bounded and unique.

In the sequel we propose an infinite-dimensional LP approach to computing ROA X_0 and show how this reformulation can be approximated by a sequence of LMI problems converging to the solution to the LP.

III. Occupation measures

In the paper we use the following notations:

- $I_A(\cdot)$ is the indicator function of a set A, i.e., a function equal to 1 on A and 0 elsewhere;
- λ denotes the Lebesgue measure on $X \subset \mathbb{R}^n$ such that
$$\lambda(A) = \int_X I_A(x)\, d\lambda(x) = \int_X I_A(x) dx = \int_A dx$$
is the standard n-dimensional volume of a set $A \subset X$;
- $\operatorname{spt}\mu$ denotes the support of a measure μ, that is, the closed set of all points x such that $\mu(A) > 0$ for every neighborhood A of x.

A. Liouville's equation

Given an initial condition x_0 and an admissible trajectory $x(\cdot\,|\,x_0) \in \mathcal{X}(x_0)$ with its corresponding control $u(\cdot\,|\,x_0) \in L^1([0, T]; \mathbb{R}^m)$ that we assume to be a measurable function of x_0, define the *occupation measure*

$$\mu(A \times B \times C\,|\,x_0) := \int_0^T I_{A\times B\times C}(t, x(t\,|\,x_0), u(t\,|\,x_0))\, dt$$

for all subsets $A \times B \times C$ in the Borel σ-algebra of subsets of $[0, T] \times X \times U$. Next, for a set K let $M(K)$ denote the Banach space of signed Borel measures supported on K, so that a measure $\nu \in M(K)$ can be interpreted as a function that takes any subset of K and returns a number in $\mathbb{R}$. Alternatively, elements of $M(K)$ can be interpreted as linear functionals acting on the Banach space of continuous functions $C(K)$, that is, as elements of the dual space $C(K)'$. The action of a measure $\nu \in M(K)$ on a test function $v \in C(K)$ can be modeled with the duality pairing

$$\langle \nu,\ v\rangle\ := \int_K v(z)\, d\nu(z).$$

Define further the linear operator $\mathcal{L} : C^1([0, T] \times X) \to C([0, T] \times X \times U)$ by

$$v \mapsto \mathcal{L}v := \frac{\partial v}{\partial t} + \sum_{i=1}^n \frac{\partial v}{\partial x_i} f_i(t, x, u) = \frac{\partial v}{\partial t} + \operatorname{grad} v \cdot f$$

and its adjoint operator $\mathcal{L}' : C([0, T] \times X \times U)' \to C^1([0, T] \times X)'$ by the adjoint relation

$$\langle \mathcal{L}'\nu, v\rangle := \langle \nu, \mathcal{L}v\rangle = \int_{[0,T]\times X\times U} \mathcal{L}v(t, x, u)\, d\nu(t, x, u)$$

for all $\nu \in M([0, T] \times X \times U) = C([0, T] \times X \times U)'$ and $v \in C^1([0, T] \times X)$.

Given a test function $v \in C^1([0, T] \times X)$, it follows from the above definition of the occupation measure μ that

$$\begin{aligned} v(T, x(T))\ &=\ v(0, x(0)) + \int_0^T \dot{v}(t, x(t\,|\,x_0))\, dt \\ &=\ v(0, x(0)) + \int_0^T \mathcal{L}v(t, x(t\,|\,x_0), u(t\,|\,x_0))\, dt \\ &=\ v(0, x(0)) + \int_{[0,T]\times X\times U} \mathcal{L}v(t, x, u)\, d\mu(t, x, u\,|\,x_0). \end{aligned} \quad (5)$$

Now consider that the initial state is not a single point but that its distribution in space is modeled by an *initial measure* $\mu_0 \in M(X)$, and that to each initial state x_0 an admissible control function $u(\cdot \mid x_0) \in L^1([0, T]; \mathbb{R}^m)$ is assigned in such a way that $x(\cdot\,|\,x_0)$ is admissible[2]. Then we can define the average occupation measure $\mu \in M([0, T] \times X \times U)$ by

$$\mu(A \times B \times C) := \int_X \mu(A \times B \times C\,|\,x_0)\, d\mu_0(x_0), \quad (6)$$

and the *final measure* $\mu_T \in M(X_T)$ by

$$\mu_T(B) := \int_X I_B(x(T\,|\,x_0))\, d\mu_0(x_0). \quad (7)$$

It follows by integrating (5) with respect to μ_0 that

$$\int_{X_T} v(T, x)\, d\mu_T(x) = \int_X v(0, x)\, d\mu_0(x) + \int_{[0,T]\times X\times U} \mathcal{L}v(t, x, u)\, d\mu(t, x, u)$$

for all $v \in C^1([0, T] \times X)$, or more concisely

$$\langle \mu_T, v(T, \cdot)\rangle\ =\ \langle \mu_0, v(0, \cdot)\rangle + \langle \mu, \mathcal{L}v\rangle\ \forall\, v \in C^1([0, T] \times X), \quad (8)$$

which is a linear equation linking the nonnegative measures μ_T, μ_0 and μ. Denoting δ_t the Dirac measure at a point t and $\otimes$ the product of measures, we can write $\langle \mu_0, v(0, \cdot)\rangle = \langle \delta_0 \otimes$

[2] The measure μ_0 can be thought of as the probability distribution of x_0 although we do not require that its mass be normalized to one.

$\mu_0,\, v\rangle$ and $\langle \mu_T,\, v(T,\cdot)\rangle = \langle \delta_T \otimes \mu_T,\, v\rangle$. Then, Eq. (8) can be rewritten equivalently using the adjoint $\mathcal{L}'$ as $\langle \mathcal{L}'\mu,\, v\rangle = \langle \delta_T \otimes \mu_T, v\rangle - \langle \delta_0 \otimes \mu_0,\, v\rangle \quad \forall\, v \in C^1([0,T]\times X)$, and since this equation is required to hold for all test functions v, we obtain a linear operator equation

$$\mathcal{L}'\mu = \delta_T \otimes \mu_T - \delta_0 \otimes \mu_0. \tag{9}$$

This equation is classical in fluid mechanics and statistical physics, where $\mathcal{L}'$ is usually written using distributional derivatives of measures (see [14]); then the equation is referred to as Liouville's partial differential equation (PDE).

Each family of admissible trajectories starting from a given initial distribution $\mu_0 \in M(X)$ satisfies Liouville's equation (9). The converse may not hold in general although for the computation of the ROA the two formulations can be considered equivalent, at least from a practical viewpoint. Let us briefly elaborate more on this subtle point now.

B. Relaxed ROA

The control system $\dot{x}(t) = f(t, x(t), u(t))$, $u(t) \in U$, can be viewed as a differential inclusion

$$\dot{x}(t) \in f(t, x(t), U) := \{f(t, x(t), u) \;:\; u \in U\}. \tag{10}$$

We show in [14, Appendix A, Lemma 1] that any triplet of measures satisfying Liouville's equation (9) corresponds to a family of trajectories of the convexified inclusion

$$\dot{x}(t) \in \operatorname{conv} f(t, x(t), U) \tag{11}$$

starting from the initial distribution μ_0, where conv denotes the convex hull. Let us denote the set of absolutely continuous admissible trajectories of (11) by

$$\begin{aligned}\bar{\mathcal{X}}(x_0) := \{x(\cdot) : &\dot{x}(t) \in \operatorname{conv} f(t, x(t), U) \text{ a.e., } x(0) = x_0,\\ & x(T) \in X_T,\ x(t) \in X\ \forall\, t \in [0,T]\},\end{aligned}$$

where a.e. stands for "almost everywhere" w.r.t. the Lebesgue measure on $[0,T]$. Given a family[3] of admissible trajectories of the convexified inclusion starting from an initial distribution μ_0, the occupation and final measures can be defined in a complete analogy via (6) and (7), but now there are only the time and space arguments in the occupation measure, not the control argument.

Let $\bar{\mu}(t,x)$ denote the (t,x)-marginal of the occupation measure μ defined through (6), that is,

$$\bar{\mu}(A \times B) := \mu(A \times B \times U) \quad \forall\, A \subset [0,T],\ B \subset X.$$

The correspondence between the convexified inclusion (11) and the measures satisfying the Liouville equation (9) is captured by the following result:

Lemma 1. *Let (μ_0, μ, μ_T) be a triplet of measures satisfying the Liouville equation (9) such that $\operatorname{spt}\mu_0 \subset X$, $\operatorname{spt}\mu \subset [0,T]\times X \times U$ and $\operatorname{spt}\mu_T \subset X_T$. Then there exists a family of absolutely continuous admissible trajectories of (11) starting from μ_0 (i.e., trajectories in $\bar{\mathcal{X}}(x_0)$) such that the occupation measure and the terminal measure generated by this family of trajectories are equal to $\bar{\mu}$ and μ_T, respectively.*

[3]Each such family can be described by a measure on $C([0,T];\mathbb{R}^n)$ which is supported on the absolutely continuous solutions to (11). Note that there may be more than one trajectory corresponding to a single initial condition x_0 since the inclusion (11) may admit multiple solutions.

Define now the relaxed region of attraction as

$$\bar{X}_0 := \{x_0 \in X \;:\; \bar{\mathcal{X}}(x_0) \neq \emptyset\}.$$

Clearly $X_0 \subset \bar{X}_0$ and the inclusion can be strict; see [14, Appendix C] for concrete examples. However, by the Filippov-Ważewski relaxation Theorem [3], the trajectories of the original inclusion (10) are dense (w.r.t. the metric of uniform convergence of absolutely continuous functions of time) in the set of trajectories of the convexified inclusion (11). This implies that the relaxed region of attraction $\bar{X}_0$ corresponds to the region of attraction of the original system but with infinitesimally dilated constraint sets X and X_T; see [14, Appendix B] for more details. Therefore, we argue that there is little difference between the two ROAs from a practical point of view. Nevertheless, because of this subtle distinction we make the following standing assumption in the remainder of the paper.

Assumption 1. *Control system (1) is such that $\lambda(X_0) = \lambda(\bar{X}_0)$.*

In other words, the volume of the classical ROA X_0 is assumed to be equal to the volume of the relaxed ROA $\bar{X}_0$. Obviously, this is satisfied if $X_0 = \bar{X}_0$, but otherwise these sets may differ by a set of zero Lebesgue measure. Any of the following conditions on control system (1) is *sufficient* for Assumption 1 to hold:

- $\dot{x}(t) \in f(t, x(t), U)$ with $f(t, x, U)$ convex for all t, x,
- $\dot{x}(t) = f(t, x(t)) + g(t, x(t))u(t)$, $u(t) \in U$ with U convex,
- uncontrolled dynamics $\dot{x}(t) = f(t, x(t))$,

as well as all controllability assumptions allowing the application of the constrained Filippov-Ważewski Theorem; see, e.g., [6] and the discussion around Assumption I in [7].

IV. Primal LP on measures

A natural way to compute the ROA is to maximize the support of the initial measure μ_0 subject to the constraint (9) and suitable support constraints on μ_0, μ and μ_T. However, a direct maximization of the volume of the support leads to a non-convex problem. The key idea behind the presented approach then consists in replacing the optimization over the support of μ_0 by the maximization of its mass under the constraint that μ_0 is dominated by the Lebesgue measure, i.e., $\mu_0 \le \lambda$. This constraint can be equivalently rewritten as $\mu_0 + \hat{\mu}_0 = \lambda$ for some nonnegative measure (a slack variable) $\hat{\mu}_0 \in M(X)$. This leads to the following LP:

$$\begin{array}{rcll} p^* & = & \sup & \mu_0(X) \\ & & \text{s.t.} & \mathcal{L}'\mu = \delta_T \otimes \mu_T - \delta_0 \otimes \mu_0 \\ & & & \mu_0 + \hat{\mu}_0 = \lambda \\ & & & \mu \ge 0,\ \mu_0 \ge 0,\ \mu_T \ge 0,\ \hat{\mu}_0 \ge 0 \\ & & & \operatorname{spt}\mu \subset [0,T]\times X \times U, \\ & & & \operatorname{spt}\mu_0 \subset X,\ \operatorname{spt}\hat{\mu}_0 \subset X \\ & & & \operatorname{spt}\mu_T \subset X_T. \end{array} \tag{12}$$

The following theorem is a cornerstone of our approach.

Theorem 1. *The optimal value of LP problem (12) is equal to the volume of the ROA X_0, that is, $p^* = \lambda(X_0)$. Moreover, the supremum is attained by the restriction of the Lebesgue measure to the ROA X_0.*

V. DUAL LP ON FUNCTIONS

In this section, we derive a dual formulation of problem (12) on the space of continuous functions. A certain super-level set of any feasible solution to the dual problem yields an outer approximation to the ROA X_0.

Consider the LP problem

$$\begin{array}{rcll} d^* & = & \inf & \int_X w(x)\,d\lambda(x) \\ & & \text{s.t.} & \mathcal{L}v(t,x,u) \le 0,\ \forall (t,x,u) \in [0,T]\times X \times U \\ & & & w(x) \ge v(0,x)+1,\ \forall x \in X \\ & & & v(T,x) \ge 0,\ \forall x \in X_T \\ & & & w(x) \ge 0,\ \forall x \in X, \end{array} \tag{13}$$

where the infimum is over $(v,w) \in C^1([0,T]\times X)\times C(X)$. The interpretation of the dual is intuitive: the constraint $\mathcal{L}v \le 0$ forces v to decrease along trajectories and hence necessarily $v(0,x) \ge 0$ on X_0 because of the constraint $v(T,x) \ge 0$ on X_T. Consequently, $w(x) \ge 1$ on X_0. This instrumental observation is formalized in the following Lemma.

Lemma 2. *If $\mathcal{L}v \le 0$ on $[0,T]\times X\times U$, $v(T,\cdot) \ge 0$ on X_T and $w \ge v(0,\cdot)+1$ on X, then $w \ge 1$ on X_0.*

We have the following salient result:

Theorem 2. *There is no duality gap between primal LP problems (12) on measures and dual LP problem (13) on functions, in the sense that $p^* = d^*$.*

It follows from Theorem 2 that the supremum in LPs (12) is attained (by the restriction of the Lebesgue measure to X_0), a statement already proved by a different means in Theorem 1. In contrast, the infimum in problem (13) is not attained in $C^1([0,T]\times X)\times C(X)$, but the w-component of feasible solutions of (13) converges to the discontinuous indicator I_{X_0} as we show next.

Before we state our convergence results, we recall the following types of convergence of a sequence of functions $w_k : X \to \mathbb{R}$ to a function $w : X \to \mathbb{R}$ on a compact set $X \subset \mathbb{R}^n$. As $k \to \infty$, the functions w_k converge to w:

- in L^1 norm if $\int_X |w_k - w|\,d\lambda \to 0$,
- almost uniformly if $\forall\,\epsilon > 0, \exists\, B \subset X,\ \lambda(B) < \epsilon$, such that $w_k \to w$ uniformly on $X \backslash B$.

We also recall that convergence in L^1 norm implies convergence in Lebesgue measure and that almost uniform convergence implies convergence almost everywhere (see [1, Theorems 2.5.2 and 2.5.3]). Our results are stated in terms of the stronger notions of L^1 norm and almost uniform convergence.

Theorem 3. *There is a sequence of feasible solutions to problem (13) such that its w-component converges from above to I_{X_0} in L^1 norm and almost uniformly.*

VI. LMI RELAXATIONS

In this section we show how the infinite dimensional LP problem (12) can be approximated by a hierarchy of LMI problems with the approximation error vanishing as the relaxation order tends to infinity. The dual LMI problem then yields a converging sequence of outer approximations to the ROA. The derivation of the finite-dimensional relaxations is omitted for space reasons; see [14] for a detailed derivation.

Let $\mathbb{R}_k[x]$ denote the vector space of real multivariate polynomials of total degree less than or equal to k and let

$$d_{X\,i} := \left\lceil \frac{\deg g_{X\,i}}{2} \right\rceil,\ d_{U\,i} := \left\lceil \frac{\deg g_{U\,i}}{2} \right\rceil,\ d_{T\,i} := \left\lceil \frac{\deg g_{T\,i}}{2} \right\rceil.$$

The primal semidefinite programming problem then reads

$$\begin{array}{lll} p_k^* \;=\; \max & (y_0)_0 & \\ \quad\quad \text{s.t.} & A_k(y,y_0,y_T,\hat{y}_0) = b_k & \\ M_k(y) \succeq 0, & M_{k-d_{X\,i}}(g_{X\,i},y) \succeq 0, i \in \mathbb{Z}_{[1,n_X]} \\ M_{k-1}(t(T-t),y) \succeq 0, & M_{k-d_{U\,i}}(g_{U\,i},y) \succeq 0, i \in \mathbb{Z}_{[1,n_U]} \\ M_k(y_0) \succeq 0, & M_{k-d_{X\,i}}(g_{X\,i},y_0) \succeq 0, i \in \mathbb{Z}_{[1,n_X]} \\ M_k(y_T) \succeq 0, & M_{k-d_{T\,i}}(g_{T\,i},y_T) \succeq 0, i \in \mathbb{Z}_{[1,n_T]} \\ M_k(\hat{y}_0) \succeq 0, & M_{k-d_{X\,i}}(g_{X\,i},\hat{y}_0) \succeq 0, i \in \mathbb{Z}_{[1,n_X]}, \end{array} \tag{14}$$

where the notation $\succeq 0$ stands for positive semidefinite and the minimum is over moment sequences $(y,y_0,y_T,\hat{y}_0)$ truncated to degree $2k$. The linear equality constraint captures the two linear equality constraints of (12) with $v(t,x) \in \mathbb{R}_{2k}[t,x]$ and $w(x) \in \mathbb{R}_{2k}[x]$ being monomials of total degree less than or equal to $2k$. The matrices $M_k(\cdot)$ and $M_k(\cdot,\cdot)$ are the moment and localizing matrices, respectively (see [16] for details). Problem (14) is a semidefinite program (SDP), where a linear function is minimized subject to convex linear matrix inequality (LMI) constraints, or equivalently a finite-dimensional LP in the cone of positive semidefinite matrices.

For the remainder of the section we make the following standard assumption.

Assumption 2. *One of the polynomials modeling the sets X, U resp. X_T, is equal to $g_{X\,i}(x) = R_X^2 - \|x\|_2^2$, $g_{U\,i}(u) = R_U^2 - \|u\|_2^2$ resp. $g_{T\,i}(x) = R_T^2 - \|x\|_2^2$, with R_X, R_U resp. R_T sufficiently large constants.*

Assumption 2 is made without loss of generality since X, U and X_T are bounded, and polynomials modeling ball constraints can be added to the semialgebraic descriptions of these sets.

The dual to the SDP problem (14) is given by

$$\begin{array}{l} d_k^* \;=\; \inf \quad \mathbf{w}'l \\ \quad\quad\quad \text{s.t.} \;\; -\mathcal{L}v(t,x,u) = p(t,x,u) + q_0(t,x,u)t(T-t) \\ \quad\quad + \sum_{i=1}^{n_X} q_i(t,x,u)g_{X\,i}(x) + \sum_{i=1}^{n_U} r_i(t,x,u)g_{U\,i}(x) \\ \quad w(x) - v(0,x) - 1 = p_0(x) + \sum_{i=1}^{n_X} q_{0i}(x)g_{X\,i}(x) \\ \quad v(T,x) = p_T(x) + \sum_{i=1}^{n_T} q_{T\,i}(x)g_{T\,i}(x) \\ \quad w(x) = s_0(x) + \sum_{i=1}^{n_X} s_{0i}(x)g_{X\,i}(x), \end{array} \tag{15}$$

where l is the vector of Lebesgue moments over X indexed in the same basis in which the polynomial $w(x)$ with coefficients $\mathbf{w}$ is expressed. The minimum is over polynomials $v(t,x) \in \mathbb{R}_{2k}[t,x]$ and $w \in \mathbb{R}_{2k}[x]$, and polynomial sum-of-squares $p(t,x,u)$, $q_i(t,x,u)$, $i \in \mathbb{Z}_{[0,n_X]}$, $r_i(t,x,u)$, $i \in \mathbb{Z}_{[1,n_U]}$, $p_0(x)$, $p_T(x)$, $q_{0i}(x)$, $q_{T\,i}(x)$, $s_0(x)$, $s_{0i}(x)$, $i \in \mathbb{Z}_{[1,n_X]}$ of appropriate degrees. The constraints that polynomials are sum-of-squares can be written explicitly as LMI constraints (see, e.g., [16]), and the objective is linear in the coefficients of the polynomial $w(x)$; therefore problem (15) can be formulated as an SDP.

Theorem 4. *There is no duality gap between primal LMI problem (14) and dual LMI problem (15), i.e. $p_k^* = d_k^*$.*

Now we state convergence results analogous to those of Theorem 3 as well as a set-wise convergence to the ROA

X_0 of certain super-level sets of the polynomial solutions to (15).

Theorem 5. *Let $w_k \in \mathbb{R}_{2k}[x]$ denote the w-component of a solution to the dual LMI problem (15) and let $\bar{w}_k(x) = \min_{i\leq k} w_i(x)$. Then w_k converges from above to I_{X_0} in L^1 norm and $\bar{w}_k$ converges from above to I_{X_0} in L^1 norm and almost uniformly.*

The following Corollary follows immediately from Theorem 5.

Corollary 1. *The sequence of infima of LMI problems (15) converges monotonically from above to the supremum of the LP problem (13), i.e., $d^* \leq d^*_{k+1} \leq d^*_k$ and $\lim_{k\to\infty} d^*_k = d^*$. Similarly, the sequence of maxima of LMI problems (14) converges monotonically from above to the maximum of the LP problem (12), i.e., $p^* \leq p^*_{k+1} \leq p^*_k$ and $\lim_{k\to\infty} p^*_k = p^*$.*

Theorem 5 establishes a functional convergence of w_k to I_{X_0} and Corollary 1 a convergence of the primal and dual optima p^*_k and d^*_k to the volume of the ROA $\lambda(X_0) = p^* = d^*$. Finally, the following theorem establishes a set-wise convergence of the unit super-level sets of w_k to X_0.

Theorem 6. *Let $w_k \in \mathbb{R}_{2k}[x]$ denote the w-component of a solution to the dual LMI problem (15) and let $X_{0k} := \{x \in \mathbb{R}^n \,:\, w_k(x) \geq 1\}$. Then $X_0 \subset X_{0k}$,*

$$\lim_{k\to\infty} \lambda(X_{0k} \setminus X_0) = 0 \;\; \text{and} \;\; \lambda(\cap_{k=1}^{\infty} X_{0k} \setminus X_0) = 0.$$

VII. Numerical examples

In this section we present two examples to illustrate our approach: an uncontrolled Van der Pol oscillator and a minimum-time control of a double integrator. For numerical implementation, one can either use Gloptipoly 3 [12] to formulate the primal problem on measures and then extract the dual solution provided by SeDuMi [22] or formulate directly the dual SOS problem using, e.g., YALMIP [27] or SOSTOOLS [23].

A. Van der Pol oscillator

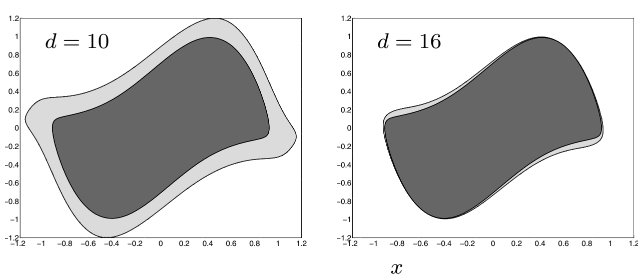

Fig. 1. Van der Pol oscillator – polynomial outer approximations (light gray) to the ROA (dark gray) for degrees $d \in \{10, 16\}$.

Consider a scaled version of the uncontrolled reversed-time Van der Pol oscillator given by $\dot{x}_1 = -2x_2$, $\dot{x}_2 = 0.8x_1 + 10(x_1^2 - 0.21)x_2$. The system has one stable equilibrium at the origin with a bounded region of attraction $X_0 \subset X := [-1.2, 1.2]^2$. In order to compute an outer approximation to this region we take $T = 100$ and $X_T = \{x \,:\, \|x\|_2 \leq 0.01\}$. Plots of polynomial super-level set approximations of degree $d \in \{10, 16\}$ are shown in Figure 1. We observe a relatively fast convergence of the

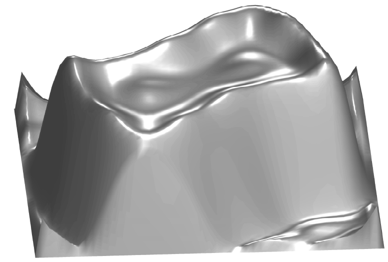

Fig. 2. Van der Pol oscillator – a polynomial approximation of degree 18 of the ROA indicator function I_{X_0}.

super-level sets to the ROA – this is confirmed by the relative volume error[4] summarized in Table I. Numerically, a better behavior is expected when using alternative polynomial bases (e.g., Chebyshev polynomials) instead of the monomials; see the conclusion for a discussion.

TABLE I
Van der Pol oscillator – relative error of the outer approximation to the ROA X_0 as a function of the approximating polynomial degree.

degree	10	12	14	16
error	49.3 %	19.7 %	11.1 %	5.7 %

B. Double integrator

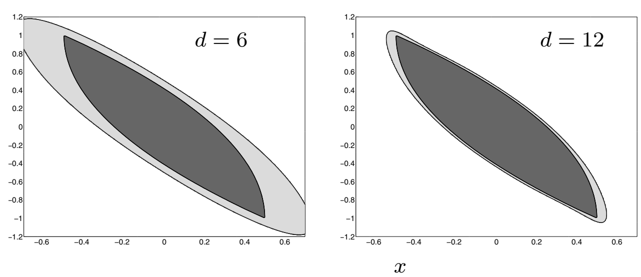

Fig. 3. Double integrator – polynomial outer approximations (light gray) to the ROA (dark gray) for degrees $d \in \{6, 12\}$.

Consider a minimum time control of a double integrator $\dot{x}_1 = x_2$, $\dot{x}_2 = u$. The goal is to find an approximation of the set of initial states X_0 that can be steered to the origin in $T = 1$. Therefore we set $X_T = \{0\}$ and the constraint set such that $X_0 \subset X$, e.g., $X = [-0.7, 0.7] \times [-1.2, 1.2]$. The solution to this problem can be computed analytically as $X_0 = \{x \,:\, V(x) \leq 1\}$, where

$$V(x) = \begin{cases} x_2 + 2\sqrt{x_1 + \frac{1}{2}x_2^2} & \text{if} \quad x_1 + \frac{1}{2}x_2|x_2| > 0, \\ -x_2 + 2\sqrt{-x_1 + \frac{1}{2}x_2^2} & \text{otherwise.} \end{cases}$$

The computation results are depicted in Figure 3; again we observe a relatively fast convergence of the super-level set approximations, which is confirmed by the relative volume errors in Table II.

[4]The relative volume error was computed approximately by Monte Carlo integration.

TABLE II

Double integrator – relative error of the outer approximation to the ROA X_0 as a function of the approximating polynomial degree.

degree	6	8	10	12
error	75.7 %	32.6 %	21.2 %	16.0 %

VIII. Conclusion

The main contributions of this paper can be summarized as follows:

- contrary to most of the existing systems control literature, we propose a convex formulation for the problem of computing the controlled region of attraction;
- our approach is constructive in the sense that we rely on standard hierarchies of finite-dimensional LMI relaxations whose convergence can be guaranteed theoretically and for which public-domain interfaces and solvers are available;
- we deal with polynomial dynamics and semialgebraic input and state constraints, therefore covering a broad class of nonlinear control systems;
- additional properties (e.g., convexity) of the approximations can be enforced by additional constraints on the polynomial (e.g., Hessian being negative definite).
- the approach extremely simple to use – the outer approximations are obtained from the solution of a single semidefinite program with *no additional data* required besides the problem description.

The problem of computing the reachability set, i.e. the set of all states that can be reached from a given set of initial conditions under input and state constraints, can be addressed with the same techniques. Basically, the initial and final measures exchange roles. Computation of maximum (controlled) invariant sets should also be amenable to our approach. Furthermore, there is a straightforward extension to piecewise polynomial dynamics defined over a semialgebraic partition of the state and input spaces – one measure is then defined for each region of the partition. Our approach should also allow for extensions to discrete-time controlled systems, stochastic systems (either discrete-time controlled Markov processes or controlled SDEs) and/or uncertain systems.

The hierarchy of LMI relaxations described in this paper generates a sequence of nested outer approximations of the ROA, but it should also be possible, using a similar approach, to compute valid inner approximations.

Numerical examples indicate that the choice of monomials as a dense basis for the set of continuous functions on compact sets, while mathematically appropriate (and notationally convenient), is not always satisfactory regarding convergence and quality of the approximations. However, this is not peculiar to ROA computation problems – a similar behavior was already observed when computing the volume (and moments) of semialgebraic sets in [13]. To achieve better performance, we recommend the use of alternative polynomial bases such as (appropriately scaled tensor products of) Chebyshev polynomials.

Acknowledgments

The authors would like to thank Mathieu Claeys and Jean-Bernard Lasserre for helpful discussions, Masakazu Kojima for valuable insights into computational aspects and the anonymous referees for useful comments.

References

[1] R. B. Ash. Real analysis and probability. Academic Press, San Diego, CA, 1972.
[2] J. P. Aubin, A. M. Bayen, P. Saint-Pierre. Viability Theory: New Directions. Springer-Verlag, Berlin, 2011.
[3] J. P. Aubin, H. Frankowska. Set-valued analysis. Springer-Verlag, Berlin, 1990.
[4] G. Chesi. Domain of attraction; analysis and control via SOS programming. Lecture Notes in Control and Information Sciences, Vol. 415, Springer-Verlag, Berlin, 2011.
[5] M. Dellnitz, G. Froyland, O. Junge. The algorithms behind GAIO – Set oriented numerical methods for dynamical systems. In B. Fiedler (ed.), Ergodic theory, analysis, and efficient simulation of dynamical systems, Springer-Verlag, Berlin, 2001.
[6] H. Frankowska, F. Rampazzo. Filippov's and Filippov-Ważewski's theorems on closed domains. Journal of Differential Equations, 161:449-478, 2000.
[7] V. Gaitsgory, M. Quincampoix. Linear programming approach to deterministic infinite horizon optimal control problems with discounting. SIAM Journal on Control and Optimization, 48:2480-2512, 2009.
[8] R. Genesio, M. Tartaglia, A. Vicino. On the estimation of asymptotic stability regions: state of the art and new proposals. IEEE Transactions on Automatic Control, 30:747–755, 1985.
[9] W. M. Haddad, V. Chellaboina. Nonlinear dynamical systems and control, a Lyapunov-based approach. Princeton Univ. Press, Princeton, NJ, 2008.
[10] W. Hahn. Stability of motion. Springer-Verlag, Berlin, 1967.
[11] D. Henrion, J. B. Lasserre. Inner approximations for polynomial matrix inequalities and robust stability regions. IEEE Transactions on Automatic Control, 57:1456-1467, 2012.
[12] D. Henrion, J. B. Lasserre, and J. Löfberg. Gloptipoly 3: moments, optimization and semidefinite programming. Optimization Methods and Software, 24:761–779, 2009.
[13] D. Henrion, J. B. Lasserre, C. Savorgnan. Approximate volume and integration for basic semialgebraic sets. SIAM Review, 51:722-743, 2009.
[14] D. Henrion, M. Korda. Convex computation of the region of attraction of polynomial control systems. Manuscript, arXiv number arXiv:1208.1751, August 2012.
[15] A. B. Kurzhanski, P. Varaiya. Ellipsoidal techniques for reachability analysis. In Hybrid systems: computation and control (B. Krogh and N. Lynch, eds.), LNCIS 1790, Springer-Verlag, Berlin, 2000.
[16] J. B. Lasserre. Moments, positive polynomials and their applications. Imperial College Press, London, UK, 2009.
[17] 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.
[18] D. G. Luenberger. Optimization by vector space methods. Wiley, New York, NY, 1969.
[19] K. Margellos, J. Lygeros. Hamilton-Jacobi formulation for reach-avoid differential games. IEEE Transactions on Automatic Control, 56:1849-1861, 2011.
[20] S. G. Margolis, W. G. Vogt. Control engineering applications of V. I. Zubov's construction procedure for Lyapunov functions. IEEE Transactions on Automatic Control, 8:101-113, 1963.
[21] I. Mitchell, C. Tomlin. Overapproximating reachable sets by Hamilton-Jacobi projections. Journal of Scientific Computing, 19:323-346, 2003.
[22] I. Pólik, T. Terlaky, and Y. Zinchenko. SeDuMi: a package for conic optimization. IMA workshop on Optimization and Control, Univ. Minnesota, Minneapolis, Jan. 2007.
[23] S. Prajna, A. Papachristodoulou, P. Seiler, P. A. Parrilo. SOSTOOLS: sum of squares optimization toolbox for MATLAB, 2004.
[24] U. Topcu, A. K. Packard, P. Seiler, G. J. Balas. Robust region-of-attraction estimation. IEEE Transactions on Automatic Control, 55:137-142, 2010.
[25] A. Vannelli, M. Vidyasagar. Maximal Lyapunov functions and domains of attraction for autonomous nonlinear systems. Automatica, 21:69-80, 1985.
[26] K. Wang, U. Vaidya. Transfer operator approach for computing domain of attraction. IEEE Conference on Decision and Control (CDC), Atlanta, GA, 2010.
[27] J. Löfberg. YALMIP : A toolbox for modeling and optimization in MATLAB. In Proc. IEEE CCA/ISIC/CACSD Conference, Taipei, Taiwan, 2004.
[28] V. I. Zubov. Methods of A. M. Lyapunov and their application. Noordhoff, Groningen, 1964.