

Inner approximations of the region of attraction for polynomial dynamical systems

Milan Korda* Didier Henrion** Colin N. Jones*

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

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

Abstract: In a previous work we developed a convex infinite dimensional linear programming (LP) approach to approximating the region of attraction (ROA) of polynomial dynamical systems subject to compact basic semialgebraic state constraints. Finite dimensional relaxations to the infinite-dimensional LP lead to a truncated moment problem in the primal and a polynomial sum-of-squares problem in the dual. This primal-dual linear matrix inequality (LMI) problem can be solved numerically with standard semidefinite programming solvers, producing a hierarchy of outer (i.e. exterior) approximations of the ROA by polynomial sublevel sets, with a guarantee of almost uniform and set-wise convergence. In this companion paper, we show that our approach is flexible enough to be modified so as to generate a hierarchy of polynomial inner (i.e., interior) approximations of the ROA with similar convergence guarantees.

1. INTRODUCTION

Given an autonomous nonlinear system and a target set, the region of attraction (ROA) is the set of all states that end in the target set at a given time without leaving the state constraint set⁵. The ROA is one of the principal sets associated to any dynamical system and goes by many other names in the literature (e.g., backward reachable set or capture basin [4]).

In [6] we showed (in a controlled setting) that there is a genuinely primal convex characterization of the ROA. Optimization over system trajectories is formulated as optimization over occupation measures, leading to an infinite dimensional linear programming (LP) problem in the cone of nonnegative measures. Finite dimensional relaxations of the dual of this problem then provide a converging sequence of *outer* approximations to the ROA. For a description of alternative techniques for numerical approximations of the ROA, please consult [5] or [6] and the many references therein.

In this paper we show, within the same measure-theoretic framework, that there exists an infinite dimensional LP whose finite-dimensional relaxations provide a converging sequence of *inner* approximations to the ROA. This paper can therefore be seen as a complement to [6]. To simplify our developments and to emphasize our contribution, we focus only on the uncontrolled setting. The main idea is to construct a converging sequence of outer approximations to the *complement* of the ROA. There are certain difficulties, topological in nature, associated with this approach. A careful distinction had to be made between trajectories leaving the constraint set and trajectories hitting its boundary. This then translates to a (sometimes subtle, but necessary) distinction between open and closed semialge-

⁵ There are various modifications on this setup (e.g., one may consider asymptotic convergence instead of finite-time reachability, with or without constraints, in the presence of disturbances and/or uncertainty, or in a controlled setting); most of these modifications are amenable to the methods presented in this paper, sometimes with different qualitative results.

braic sets. Fortunately, the LP formulation proposed in [6] was flexible enough to allow for these modifications.

Generally speaking, and consistently with our previous work [6], we believe that the main virtues of our approach are overall convexity, conceptual simplicity and compactness. Both primal and dual finite-dimensional relaxations turn out to be linear matrix inequalities (LMI), also called semidefinite programming (SDP) problems, with no tuning parameters besides the relaxation order and no initialization data besides the defining ingredients of the problem. In addition, the inner approximations obtained are particularly simple – they are given by a sublevel set of a single polynomial of a predefined degree. Therefore, an ROA approximation in analytic form can be readily obtained by solving a single LMI using freely available software (e.g., SeDuMi [12]).

2. PROBLEM STATEMENT

Consider the autonomous system

$$\dot{x}(t) = f(t, x(t)), \quad x(t) \in X \subset \mathbb{R}^n, \quad t \in [0, T] \quad (1)$$

with a given vector field f with polynomial entries $f_i \in \mathbb{R}[t, x]$, $i = 1, \dots, n$, final time $T > 0$. The state trajectory $x(\cdot)$ is constrained to a nonempty open⁶ basic semialgebraic set⁷

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

where the polynomial $g_X \in \mathbb{R}[x]$ is such that the set

$$\bar{X} := \{x \in \mathbb{R}^n : g_X(x) \geq 0\} \supset X$$

is compact⁸.

⁶ The requirement of the constraint set being open is merely technical, for this considerably simplifies the developments and the proofs.

⁷ For clarity of exposition we consider the constraint set given by a single superlevel set of a polynomial. The approach can, however, be straightforwardly extended to constraint sets defined by the intersection of finitely many polynomial superlevel sets.

⁸ Note that the closed semialgebraic set $\bar{X} = \{x : g_X(x) \geq 0\}$ can be strictly larger than the closure of the open semialgebraic set

The vector field f is polynomial and therefore Lipschitz on the compact set \bar{X} . As a result, for any $x_0 \in \bar{X}$ there exists a unique maximal solution $x(\cdot)$ to ODE (1). The time interval on which this solution is defined contains the time interval on which $x(t) \in \bar{X}$.

2.1 Region of attraction (ROA)

Given a final time T and an open bounded basic semialgebraic target set

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

the region of attraction (ROA) is defined as

$$X_0 := \left\{ x_0 \in X : \exists x(\cdot) \text{ s.t. } \dot{x}(t) = f(t, x(t)), x(0) = x_0, \right. \\ \left. x(T) \in X_T, x(t) \in X, \forall t \in [0, T] \right\}. \quad (3)$$

In words, the ROA is the set of all initial states from X for which the unique solution to (1) stays in X for all $t \in [0, T]$ and ends in the target set at time T .

2.2 Complement ROA

The idea to get *inner* approximations of the ROA X_0 is to construct *outer* approximations of the complement ROA $X_0^c := X \setminus X_0$. By continuity of solutions to (1), the set X_0^c is equal to

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

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 words, X_0^c is the set of initial states that give rise to trajectories which do not end up in X_T at time T and/or violate the state constraint at some point between 0 and T .

3. OCCUPATION MEASURES

In this section we introduce the concept of occupation measures and show how the nonlinear system dynamics can be equivalently described by a linear equation on measures.

3.0.0.1. Notation We will use the following notation. The vector space of all signed Borel measures with support contained in a Borel set K is denoted by $M(K)$. The support (i.e., the smallest closed set whose complement has a zero measure) of a measure μ is denoted by $\text{spt } \mu$. The space of continuous functions on K is denoted by $C(K)$ and likewise the space of continuously differentiable functions is $C^1(K)$. The indicator function of a set K (i.e., the function equal to one on K and zero otherwise) is denoted by $I_K(\cdot)$. The symbol λ denotes the n -dimensional Lebesgue measure (i.e., the standard n -dimensional volume). The integral of a function v w.r.t a measure μ over a set K is denoted by $\int_K v(x) d\mu(x)$. Sometimes we for simplicity omit the integration variable and/or the set over which we integrate if they are obvious from the context.

$X = \{x : g_X(x) > 0\}$, consider in \mathbb{R} e.g. $g_X(x) = (1 - x^2)(2 + x)^2$. For a similar reason, note also that X bounded does not imply \bar{X} bounded. Indeed, in \mathbb{R}^2 with $g_X(x) = (1 - x_1^2 - x_2^2)(2 + x_1)^2$ we have $X = \{x : \|x\| < 1\}$ and $\bar{X} = \{x : \|x\| \leq 1\} \cup \{x : x_1 = -2\}$.

Now assume $x_0 \in \bar{X}$ and define the first hitting time of X_∂ as

$$\tau(x_0) := \min \{T, \inf \{t \geq 0 : x(t | x_0) \in X_\partial\}\}, \quad (4)$$

where $x(\cdot | x_0)$ denotes the unique trajectory starting from x_0 (which is well defined on the time interval $[0, \tau(x_0)]$). Then we define the *occupation measure* associated to the trajectory starting from x_0 by

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

for all Borel⁹ sets $A \times B \subset [0, T] \times \bar{X}$. The interpretation is that the occupation measure measures the time spent by the trajectory $x(\cdot | x_0)$ in subsets of the state space.

The occupation measure enjoys the following important property: for any measurable function $v(t, x)$ the equality

$$\int_0^{\tau(x_0)} v(t, x(t)) dt = \int_{[0, T] \times \bar{X}} v(t, x) d\mu(t, x | x_0) \quad (5)$$

holds. In words, the time integral of a function evaluated along the trajectory $x(\cdot | x_0)$ is equal to the integral of the function w.r.t. the occupation measure associated to x_0 . Therefore, loosely speaking, all information about the trajectory $x(\cdot | x_0)$ is encoded by the occupation measure $\mu(\cdot | x_0)$.

Now suppose that the initial state is not a single point but that its spatial distribution is given by an *initial measure* $\mu_0 \in M(\bar{X})$. Then we define the *average occupation measure* $\mu \in M([0, T] \times \bar{X})$ as

$$\mu(A \times B) := \int_{\bar{X}} \mu(A \times B | x_0) d\mu_0(x_0).$$

Lastly, we define the final measure $\mu_T \in M([0, T] \times \bar{X})$ by

$$\mu_T(B) := \int_{\bar{X}} I_B(x(T | x_0)) d\mu_0(x_0).$$

To derive an equation linking the three principal measures, consider a test function $v \in C^1([0, T] \times \bar{X})$ evaluated along a trajectory. Using the chain rule and equation (5) we obtain

$$v(\tau(x_0), x(\tau(x_0) | x_0)) - v(0, x_0) = \int_0^{\tau(x_0)} \frac{d}{dt} v(t, x(t | x_0)) dt \\ = \int_0^{\tau(x_0)} \left(\frac{\partial v}{\partial t} + \text{grad } v \cdot f(t, x(t | x_0)) \right) dt \\ = \int_{[0, T] \times \bar{X}} \left(\frac{\partial v}{\partial t} + \text{grad } v \cdot f(t, x) \right) d\mu(t, x | x_0) \\ = \int_{[0, T] \times \bar{X}} \mathcal{L}v(t, x) d\mu(t, x | x_0)$$

where the linear operator $\mathcal{L} : C^1([0, T] \times \bar{X}) \rightarrow C([0, T] \times \bar{X})$ is defined by

$$v \mapsto \mathcal{L}v := \frac{\partial v}{\partial t} + \text{grad } v \cdot f.$$

Integrating the above equation w.r.t. μ_0 leads to the equation

$$\int_{[0, T] \times \bar{X}} v(t, x) d\mu_T(t, x) - \int_{\bar{X}} v(0, x) d\mu_0(x) = \\ \int_{[0, T] \times \bar{X}} \mathcal{L}v(t, x) d\mu(t, x) \quad \forall v \in C^1([0, T] \times \bar{X}), \quad (6)$$

⁹ For brevity we drop the adjective ‘‘Borel’’ in the sequel.

which is a linear equation linking the measures μ_0 , μ and μ_T . Equation (6) is sometimes referred to as Liouville's equation.

4. PRIMAL LP

In this section we follow the approach developed in [6] and derive an infinite-dimensional linear programming (LP) characterization of the complement ROA X_0^c . Certain sublevel sets of feasible solutions to the dual of this LP then yield inner approximations to the ROA X_0 .

The basic idea is to maximize the mass of the initial measure μ_0 under the constraint that it is dominated by the Lebesgue measure, i.e., $\mu_0 \leq \lambda$. System dynamics is captured by Liouville's equation (6) and state and terminal constraints are handled by suitable constraints on the support of the measures. The key idea is then to split the final measure in two measures such that each measure is supported on a suitable compact basic semialgebraic set. More explicitly, we let

$$\mu_T := \mu_T^1 + \mu_T^2$$

with $\mu_T^1 \in M([0, T] \times X_\partial)$ and $\mu_T^2 \in M(\{T\} \times X_T^c)$. That is, we require that $\text{spt } \mu_T^1 \subset [0, T] \times X_\partial$ and $\text{spt } \mu_T^2 \subset \{T\} \times X_T^c$. The interpretation is that measure μ_T^1 models the trajectories that leave X , whereas measure μ_T^2 models the trajectories that end in X_T^c (i.e., not in X_T). These support constraints on the final measure(s) along with system dynamics enforce that the support of the initial measure μ_0 must be contained in X_0^c . Since there are no other constraints on μ_0 besides $\mu_0 \leq \lambda$, maximization of its mass should yield the restriction of the Lebesgue measure λ to X_0^c .

The constraint $\mu_0 \leq \lambda$ can be rewritten equivalently as $\mu_0 + \hat{\mu}_0 = \lambda$ for some nonnegative slack measure $\hat{\mu}_0 \in M(X)$. This is equivalent to requiring that $\int w d\mu_0 + \int w d\hat{\mu}_0 = \int w d\lambda$ for all test functions $w \in C(\bar{X})$. In addition, we can drop the time argument from the definition of μ_T^2 since its time component is supported on a singleton.

This leads to the following optimization problem:

$$\begin{aligned} p^* &= \sup \int 1 d\mu_0 \\ \text{s.t. } & \int v d\mu_T^1 + \int v(T, \cdot) d\mu_T^2 - \int v(0, \cdot) d\mu_0 = \int \mathcal{L}v d\mu \quad \forall v \in C^1 \\ & \int w d\mu_0 + \int w d\hat{\mu}_0 = \int w d\lambda \quad \forall w \in C \\ & \mu_0 \geq 0, \mu \geq 0, \mu_T^1 \geq 0, \mu_T^2 \geq 0, \hat{\mu}_0 \geq 0 \\ & \text{spt } \mu \subset [0, T] \times \bar{X}, \text{spt } \mu_0 \subset \bar{X}, \text{spt } \hat{\mu}_0 \subset \bar{X} \\ & \text{spt } \mu_T^1 \subset [0, T] \times X_\partial, \text{spt } \mu_T^2 \subset X_T^c, \end{aligned} \tag{7}$$

where the supremum is over the vector of nonnegative measures

$$(\mu_0, \mu, \mu_T^1, \mu_T^2, \hat{\mu}_0) \in M(\bar{X}) \times M([0, T] \times \bar{X}) \times M([0, T] \times X_\partial) \times M(X_T^c) \times M(\bar{X}).$$

Problem (7) is an infinite-dimensional LP in the cone of nonnegative measures. Indeed, the objective is linear, the first two constraints are linear equality constraints and the remaining constraints are conic constraints (the set of nonnegative measures supported on a given set is a positive cone in the vector space of all measures supported on the same set).

The discussion leading to problem (7) is formalized in the following result.

Theorem 1. The optimal value of LP problem (7) is equal to the volume of the complement ROA X_0^c , that is, $p^* = \lambda(X_0^c)$. Moreover, the supremum is attained by the restriction of the Lebesgue measure to the complement ROA X_0^c .

Proof. Closely follows arguments in [6]. By definition of the relaxed complement ROA, the unique trajectory $x(\cdot)$ associated to any initial condition $x_0 \in X_0^c$ either hits X_∂ at some $t \in [0, T]$ or ends in X_T^c . Therefore for any initial measure $\mu_0 \leq \lambda$ with $\text{spt } \mu_0 \subset X \subset \bar{X}$ there exist an occupation measure μ , final measures μ_T^1, μ_T^2 and a slack measure $\hat{\mu}_0$ such that the constraints of problem (7) are satisfied. One such measure μ_0 is the restriction of the Lebesgue measure to X_0^c , and therefore $p^* \geq \lambda(X_0^c)$.

Now we show that $p^* \leq \lambda(X_0^c)$. Take a vector of measures $(\mu_0, \mu, \mu_T^1, \mu_T^2, \hat{\mu}_0)$ feasible in (7) and suppose that $\lambda(\text{spt } \mu_0 \setminus X_0^c) > 0$. Since any level set of a polynomial has a zero Lebesgue measure we have $\lambda(X_\partial) = 0$ and

$$\lambda(\text{spt } \mu_0 \setminus (X_0^c \cup X_\partial)) = \lambda(\text{spt } \mu_0 \setminus X_0^c) > 0.$$

By a superposition principle [1, Theorem 3.2] using arguments of [6, Appendix A, Lemma 4], there exists a family of admissible trajectories of the ODE (1) starting from μ_0 generating the occupation measure μ and the final measure $\mu_T = \mu_T^1 + \mu_T^2$. However, this is a contradiction since $\text{spt } \mu_0 \setminus (X_0^c \cup X_\partial) \subset X_0$, which means that all trajectories starting from $\text{spt } \mu_0 \setminus (X_0^c \cup X_\partial)$ neither hit X_∂ nor end in X_T^c . Thus, $\lambda(\text{spt } \mu_0 \setminus X_0^c) = 0$ and so $\lambda(\text{spt } \mu_0) \leq \lambda(X_0^c)$. Combining this with the constraint $\mu_0 \leq \lambda$ we get $\mu_0(X) = \mu_0(\text{spt } \mu_0) \leq \lambda(\text{spt } \mu_0) \leq \lambda(X_0^c)$ for any feasible μ_0 . Therefore $p^* \leq \lambda(X_0^c)$ and thus in fact $p^* = \lambda(X_0^c)$.

5. DUAL LP

In this section we derive a dual LP on continuous functions, prove the absence of a duality gap between the primal and dual LPs and relate feasible solutions to the dual to the indicator function of the complement ROA X_0^c .

By standard infinite-dimensional LP theory (see, e.g., [2]), the dual to LP (7) reads

$$\begin{aligned} d^* &= \inf \int_X w(x) d\lambda(x) \\ \text{s.t. } & \mathcal{L}v(t, x) \leq 0, \quad \forall (t, x, u) \in [0, T] \times \bar{X} \\ & w(x) \geq v(0, x) + 1, \quad \forall x \in \bar{X} \\ & v(T, x) \geq 0, \quad \forall x \in X_T^c \\ & v(t, x) \geq 0, \quad \forall (t, x) \in [0, T] \times X_\partial \\ & w(x) \geq 0, \quad \forall x \in \bar{X}, \end{aligned} \tag{8}$$

where the infimum is over $(v, w) \in C^1([0, T] \times \bar{X}) \times C(\bar{X})$.

The intuition is that given $x_0 \in X_0^c$ the constraint $\mathcal{L}v \leq 0$ forces v to decrease along trajectories as long as it does not hit X_∂ or end in X_T^c . Because of the constraint $v \geq 0$ on $[0, T] \times X_\partial \cup \{T\} \times X_T^c$ we must have $v(0, \cdot) \geq 0$ on X_0^c . Consequently, $w(x) \geq 1$ on X_0^c . This instrumental observation is formalized in the following Lemma.

Lemma 2. If $\mathcal{L}v \leq 0$ on $[0, T] \times \bar{X}$, $v \geq 0$ on $([0, T] \times X_\partial) \cup (\{T\} \times X_T^c)$ and $w \geq v(0, \cdot) + 1$ on \bar{X} , then $w \geq 1$ on X_0^c .

Proof. Take $x_0 \in X_0^c$ and consider the first hitting time of X_∂ , $\tau := \tau(x_0)$, defined by (4). By definition of X_0^c the trajectory starting from x_0 will either hit X_∂ or end in X_T^c .

Therefore $x(\tau) \in ([0, T] \times X_\partial) \cup (\{T\} \times X_T^c)$ and $x(t) \in X$ for $t \in [0, \tau]$. Therefore $v(\tau, x(\tau)) \geq 0$, $\mathcal{L}v(t, x(t)) \leq 0$, $\forall t \in [0, \tau]$ and so

$$\begin{aligned} 0 \leq v(\tau, x(\tau)) &= v(0, x_0) + \int_0^\tau \mathcal{L}v(t, x(t)) dt \leq v(0, x_0) \\ &\leq w(x_0) - 1. \end{aligned}$$

The following result is of key importance for subsequent developments.

Theorem 3. There is no duality gap between primal LP problems (7) on measures and dual LP problem (8) on functions, that is, $p^* = d^*$.

Proof. Here we only outline the basic steps; for a detailed argument in a similar setting see [6, Theorem 2]. Since the supports of all measures are compact, the initial measure is dominated by the Lebesgue measure and the final time is finite, we have $\mu_0(\bar{X}) \leq \lambda(\bar{X}) < \infty$, $\mu_T([0, T] \times \bar{X}) = \mu_0(\bar{X}) < \infty$ and $\mu([0, T] \times X) \leq T\mu_T([0, T] \times \bar{X}) < \infty$, where the last two inequalities follow by plugging in $v(t, x) = 1$ and $v(t, x) = t$ in Liouville's equation (6). Therefore $p^* < \infty$ and the feasible set of problem (7) is weakly-* bounded. Furthermore, the feasible set of (7) is nonempty since $(\mu_0, \mu, \mu_T^1, \mu_T^2, \hat{\mu}_0) = (0, 0, 0, 0, \lambda)$ is a trivial feasible point; therefore $0 \leq p^* < \infty$. The absence of a duality gap then follows from [2, Theorem 3.10] using Alaoglu's theorem (see, e.g., [10, Chapter 5]) and the weak-* continuity of the adjoint of the operator \mathcal{L} .

Next, we establish our first convergence¹⁰ result.

Theorem 4. There is a sequence of feasible solutions to problem (8) such that its w -component converges from above to I_{X_0} in L^1 norm and almost uniformly.

Proof. Follows by the same arguments as Theorem 3 in [6].

6. LMI RELAXATIONS

In this section we derive finite dimensional semidefinite programming (SDP) or linear matrix inequality (LMI) relaxations to the infinite dimensional LPs (7) and (8) and establish several convergence results relating these relaxations to the infinite dimensional LPs and to the ROA.

In what follows, $\mathbb{R}_k[\cdot]$ denotes the vector space of real multivariate polynomials of total degree less than or equal to k .

Derivation of the finite dimensional relaxations is standard and the reader is referred to [6, Section 5] or to the comprehensive reference [9]; therefore we only highlight the main ideas. First of all, since the supports of all measures feasible in (7) are compact, these measures are determined by their moments, i.e., by integrals of all monomials (which is a sequence of real numbers when indexed in, e.g., the canonical monomial basis). Therefore, it suffices to restrict the test functions $w(x)$ and $v(t, x)$ in (7) to all monomials, reducing the linear equality constraints of (7) to linear equality constraints on the moments. Next, by the celebrated Putinar Positivstellensatz (see [9, 11]), the constraint that the support of a measure is included in a given compact basic semialgebraic set is equivalent to the

¹⁰Please refer to [6] or, e.g., [3] for definitions of the various types of convergence relevant in this context.

feasibility of an infinite sequence of LMIs involving the so-called moment and localizing matrices, which are linear in the coefficients of the moment sequence. By truncating the moment sequence and taking only the moments corresponding to monomials of total degree less than or equal to $2k$ we obtain a necessary condition for this truncated moment sequence to be the first part of a moment sequence of a measure with the desired support.

This procedure leads to the primal SDP relaxation of order k

$$\begin{aligned} p_k^* &= \max (y_0)_0 \\ \text{s.t. } &A_k(y, y_0, y_T^1, y_T^2, \hat{y}_0) = b_k \\ &M_k(y) \succeq 0, M_{k-d_{X_i}}(g_{X_i}, y) \succeq 0 \\ &M_k(y_0) \succeq 0, M_{k-d_X}(g_X, y_0) \succeq 0 \\ &M_k(y_T^1) \succeq 0, M_{k-d_T}(g_X, y_T^1) \succeq 0, M_{k-d_T}(-g_X, y_T^1) \succeq 0 \\ &M_k(y_T^2) \succeq 0, M_{k-d_T}(g_X, y_T^2) \succeq 0, M_{k-d_T}(-g_T, y_T^2) \succeq 0 \\ &M_k(\hat{y}_0) \succeq 0, M_{k-d_X}(g_X, \hat{y}_0) \succeq 0 \\ &M_{k-1}(t(T-t), y) \succeq 0, M_{k-1}(t(T-t), y_T^1) \succeq 0 \end{aligned} \quad (9)$$

where the notation $\succeq 0$ stands for positive semidefinite and the minimum is over moment sequences $(y, y_0, y_T^1, y_T^2, \hat{y}_0)$ truncated to degree $2k$ corresponding to measures $\mu, \mu_0, \mu_T^1, \mu_T^2$ and $\hat{\mu}_0$. The linear equality constraint captures the two linear equality constraints of (7) 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 $M_k(\cdot)$ are the moment and localizing matrices, following the notations of [9] or [6]. In problem (9), a linear objective is minimized subject to linear equality constraints and LMI constraints; therefore problem (9) is an SDP problem.

The SDP problem dual to problem (9) turns out to be the sum-of-squares problem

$$\begin{aligned} d_k^* &= \inf w \cdot l \\ \text{s.t. } &-\mathcal{L}v(t, x) = p(t, x) + q_1(t, x)t(T-t) + q_2(t, x)g_X(x) \\ &w(x) - v(0, x) - 1 = p_0(x) + q_{01}(x)g_X(x) \\ &v(t, x) = p_{T1}(x) + q_{T1}(t, x)t(T-t) + r(x)g_X(x) \\ &v(T, x) = p_{T2}(x) + q_{T2}(x)g_X(x) - q_{T3}(x)g_T(x) \\ &w(x) = s_0(x) + s_1(x)g_X(x), \end{aligned} \quad (10)$$

where l is the vector of Lebesgue moments over X indexed in the same basis in which the polynomial $w(x)$ with coefficients w is expressed. The minimum is over polynomials $v(t, x) \in \mathbb{R}_{2k}[t, x]$, $w(x) \in \mathbb{R}_{2k}[x]$, over the polynomial $r(x)$ and polynomial sum-of-squares $p(t, x)$, $q_1(t, x)$, $q_2(t, x)$, $q_{01}(x)$, $p_{T1}(x)$, $p_{T2}(x)$, $q_{T1}(x)$, $q_{T2}(x)$, $q_{T3}(x)$, $s_0(x)$, $s_1(x)$ of appropriate degrees. The constraints that polynomials are sum-of-squares can be written explicitly as LMI constraints (see, e.g., [9]), and the objective is linear in the coefficients of the polynomial $w(x)$; therefore problem (10) can be formulated as an SDP problem.

Theorem 5. There is no duality gap between primal LMI problem (9) and dual LMI problem (10), i.e. $p_k^* = d_k^*$.

Proof. Follows by the same arguments based on standard SDP duality theory as Theorem 4 in [6].

Now we prove our main convergence results.

Theorem 6. Let $w_k \in \mathbb{R}_{2k}[x]$ denote the w -component of a solution to the dual LMI problem (10) and let $\bar{w}_k(x) = \min_{i \leq k} w_i(x)$. Then $1 - w_k$ converges from below to I_{X_0} in L^1 norm and $1 - \bar{w}_k$ converges from below to I_{X_0} in L^1 norm and almost uniformly.

Proof. It follows from Theorem 4 and from the density of polynomials in the space of continuous functions on compact sets (for a detailed argument in a similar setting

see [6, Theorem 5]) that w_k and \bar{w}_k converge from above to $I_{X_0^c}$ in L_1 and almost uniformly on X , respectively. Therefore $1 - w_k$ and $1 - \bar{w}_k$ converge from below to $I_{X_0} = 1 - I_{X_0^c}$ on X in the same manner.

The next Corollary follows immediately from Theorem 6.

Corollary 7. The sequence of infima of LMI problems (10) converges monotonically from above to the supremum of the LP problem (8), i.e., $d^* \leq d_{k+1}^* \leq d_k^*$ and $\lim_{k \rightarrow \infty} d_k^* = d^*$. Similarly, the sequence of maxima of LMI problems (9) converges monotonically from above to the maximum of the LP problem (7), i.e., $p^* \leq p_{k+1}^* \leq p_k^*$ and $\lim_{k \rightarrow \infty} p_k^* = p^*$.

Proof: Follows the proof of Corollary 1 in [6]. \square

Our last results establishes set-wise convergence of inner approximations to the ROA.

Theorem 8. Let $w_k \in \mathbb{R}_{2k}[x]$ denote the w -component of a solution to the dual LMI problem (10) and let $X_{0k} := \{x \in X : w_k(x) < 1\}$. Then $X_{0k} \subset X_0$,

$$\lim_{k \rightarrow \infty} \lambda(X_0 \setminus X_{0k}) = 0 \quad \text{and} \quad \lambda(X_0 \setminus \cup_{k=1}^{\infty} X_{0k}) = 0.$$

Proof. Follows the proof of Theorem 6 in [6]. From Lemma 2 we have $w_k(x) \leq 1 \Rightarrow x \in X_0$ for all $x \in X$, and therefore $I_{X_{0k}} \leq I_{X_0}$. Since $w_k \geq 0$ on X we also have $1 - w_k \leq I_{X_{0k}}$ on X and therefore $1 - w_k \leq I_{X_{0k}} \leq I_{X_0}$ on X . From Theorem 6, we have $1 - w_k \rightarrow I_{X_0}$ in L^1 norm on X . Consequently,

$$\begin{aligned} \lambda(X_0) &= \int_X I_{X_0} d\lambda = \lim_{k \rightarrow \infty} \int_X 1 - w_k d\lambda \\ &\leq \lim_{k \rightarrow \infty} \int_X I_{X_{0k}} d\lambda = \lim_{k \rightarrow \infty} \lambda(X_{0k}) \\ &\leq \lim_{k \rightarrow \infty} \lambda(\cup_{i=1}^k X_{0i}) = \lambda(\cup_{k=1}^{\infty} X_{0k}). \end{aligned}$$

But since $X_{0k} \subset X_0$ for all k , we must have

$$\lim_{k \rightarrow \infty} \lambda(X_{0k}) = \lambda(X_0) \quad \text{and} \quad \lambda(\cup_{k=1}^{\infty} X_{0k}) = \lambda(X_0),$$

which proves the theorem.

7. NUMERICAL EXAMPLES

In this section we present two numerical examples. The primal problems on measures were modeled using Glopoly 3 [7] interfaced with the SDP solver SeDuMi [12]; this solver also returns the solution to the dual SDP relaxation. In Section 7.3 we then investigate how tight low order approximations can be obtained.

In general, for high order relaxations to be reached, the approach requires that the data are properly scaled (e.g. to unit boxes); in particular the time interval $[0, T]$ should be scaled to $[0, 1]$ by multiplying the vector field f by T .

7.1 Univariate cubic dynamics

Consider the system given by

$$\dot{x} = x(x - 0.5)(x + 0.5),$$

the constraint set $X = [-1, 1]$, the final time $T = 10$ and the target set $X_T = [-0.3, 0.3]$. The ROA in this setup is $X_0 = [-0.5, 0.5]$. Polynomial approximations to the complement ROA for degrees $d \in \{16, 24\}$ are shown in Figure 1. As expected, the functional convergence of the polynomial to the discontinuous indicator function is rather slow. A slightly better convergence is observed in the volume error of the sublevel set approximation to the

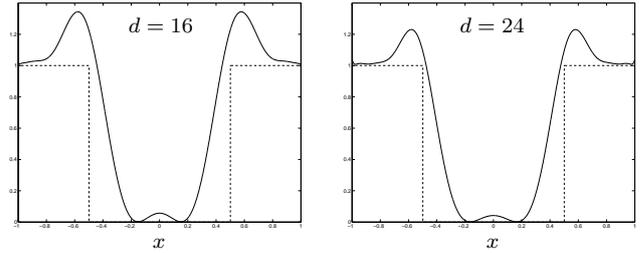


Fig. 1. Univariate cubic dynamics – polynomial approximations (solid line) to the complement ROA indicator function $I_{X_0^c} = I_{[-1, -0.5]} + I_{[0.5, 1]}$ (dashed line) for degrees $d \in \{16, 24\}$.

ROA documented in Table 1. The relatively slow convergence could be significantly improved if a tighter constraint set X was employed; see Section 7.3 below. Alternative polynomial bases (e.g. Chebyshev polynomials) would also allow tighter higher order approximations; see [8] for more details.

Table 1. Univariate cubic dynamics – relative error of the inner approximations to the ROA $X_0 = [-0.5, 0.5]$ as a function of the approximating polynomial degree.

degree	16	20	24	28
error	11.4%	6.4%	4.84%	4.54%

7.2 Van der Pol oscillator

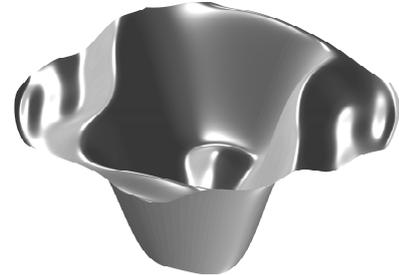


Fig. 2. Van der Pol oscillator – degree 18 polynomial approximation to the indicator function of the complement ROA.

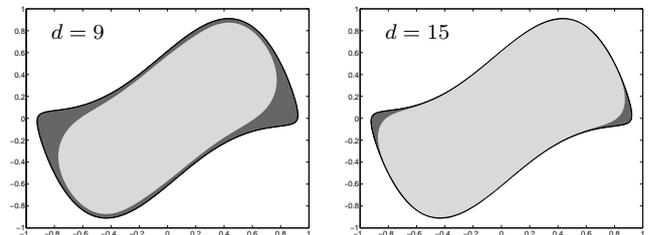


Fig. 3. Van der Pol oscillator – polynomial inner approximations (light gray) to the ROA (dark gray) for degrees $d \in \{9, 15\}$.

As a second example consider a scaled¹¹ version of the uncontrolled reversed-time Van der Pol oscillator given by

¹¹The coefficients were chosen so that the ROA fits within the box $[-1, 1]^2$.

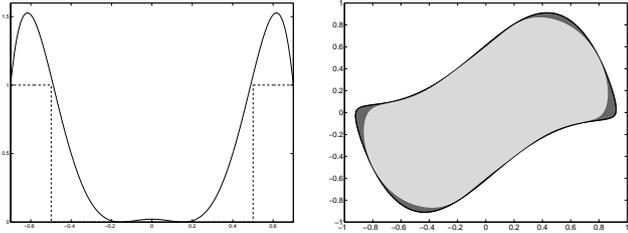


Fig. 4. Low order approximations to the ROA. Top: tighter constraint set X and low order w (compare with Figure 1). Bottom: low order w only (compare with Figure 3).

$$\begin{aligned}\dot{x}_1 &= -2.0x_2, \\ \dot{x}_2 &= 0.80x_1 + 10(x_1^2 - 0.21)x_2.\end{aligned}$$

We take $T = 1$ and $X_T = \{x \in \mathbb{R}^n : \|x\|_2 \leq 0.50\}$ and $X := \{x \in \mathbb{R}^n : \|x\| \leq 1.1\}$. The ROA is bounded, having the characteristic Van der Pol shape. Plots of polynomial sublevel set approximations of degrees $d \in \{9, 15\}$ are shown in Figure 3. We observe a relatively fast convergence to the ROA, which is also documented by the relative volume errors reported in Table 2. Figure 2 then shows a degree 18 polynomial approximation to the indicator function of the complement ROA.

Table 2. Van der Pol oscillator – relative error of the inner approximation to the ROA X_0 as a function of the approximating polynomial degree.

degree	9	12	15	18
error	18.3%	8.4%	3.8%	3.1%

7.3 Low order approximations

In the examples above, relatively high order polynomials had to be used to obtain tight approximations, which can limit subsequent applicability of the approximations. There are several ways to obtain low order approximations of similar quality. First of all, since the integral of a polynomial w is minimized over the constraint set X , it is desirable that X be a good outer approximation of the ROA. Of course, selecting X is possible only if it is an artificially specified outer approximation of the ROA, not a constraint set coming from physical requirements on the system. More importantly, notice that in problem (10) the system dynamics enters the constraints on the polynomial $v(t, x)$, whereas the polynomial $w(x)$ is only upper-bounding $v(t, x) + 1$ for $t = 0$. Since the inner approximations are given by sublevel sets of w , it is possible and plausible to choose different degrees of w and v – low for w and higher for v . Both techniques are illustrated in Figure 4; in Figure 4 (a) we consider the univariate cubic dynamics and we both shrink the constraint set X and choose low order w while keeping v of higher order. In Figure 4 (b) we consider the Van Der Pol oscillator, keeping the constraint set X unchanged and only selecting low order w . The inner approximations obtained are indeed significantly tighter for the given degrees (compare with Figures 1 and 3).

8. CONCLUSION

This paper presented an infinite dimensional convex characterization of the region of attraction (ROA) for uncontrolled polynomial systems, following the approach initiated in our previous work [6]. Finite dimensional dual

relaxations yield a converging sequence of inner approximations to the ROA, thereby complementing the outer approximations of [6]. One of the virtues of the approach is its conceptual simplicity – the resulting approximation is the outcome of a single SDP or LMI problem with no free parameters except for the relaxation order. The approximations itself are also simple, given by sublevel sets of polynomials of predefined degrees.

Nevertheless, this approach does not escape the curse of dimensionality – indeed, whereas the number of variables of the LMI relaxations grows polynomially with the relaxation order, this number grows exponentially with the state dimension. Tailored structure-exploiting SDP solvers could enable this approach to reach higher dimensions. In addition, a different choice of basis functions (e.g., Chebyshev polynomials rather than monomials) would improve numerical conditioning of the LMIs, allowing higher order relaxations to be computed.

Future research directions include inner approximations in a controlled setting and the related problem of robust region of attraction / reachable set computation with either unknown but constant uncertainty or a time-varying disturbance. Finally, approximation of the asymptotic region of attraction and the maximum (controlled) invariant set are also amenable to similar tools.

REFERENCES

- [1] L. Ambrosio. Transport equation and Cauchy problem for non-smooth vector fields. In L. Ambrosio et al. (eds.), *Calculus of variations and nonlinear partial differential equations. Lecture Notes in Mathematics*, Vol. 1927, Springer-Verlag, Berlin, 2008.
- [2] E. J. Anderson, P. Nash. *Linear programming in infinite-dimensional spaces: theory and applications*. Wiley, New York, NY, 1987.
- [3] R. B. Ash. *Real analysis and probability*. Academic Press, San Diego, CA, 1972.
- [4] J. P. Aubin, A. M. Bayen, P. Saint-Pierre. *Viability theory: new directions*. Springer-Verlag, Berlin, 2011.
- [5] G. Chesi. *Domain of attraction; analysis and control via SOS programming. Lecture Notes in Control and Information Sciences*, Vol. 415, Springer-Verlag, Berlin, 2011.
- [6] D. Henrion, M. Korda. Convex computation of the region of attraction of polynomial control systems. LAAS-CNRS Research Report 12488, [arXiv:1208.1751](https://arxiv.org/abs/1208.1751), August 2012.
- [7] D. Henrion, J. B. Lasserre, and J. Löfberg. *Gloptipoly 3: moments, optimization and semidefinite programming. Optimization Methods and Software*, 24:761–779, 2009.
- [8] D. Henrion, J. B. Lasserre, C. Savorgnan. Approximate volume and integration for basic semialgebraic sets. *SIAM Review*, 51:722–743, 2009.
- [9] J. B. Lasserre. *Moments, positive polynomials and their applications*. Imperial College Press, London, UK, 2009.
- [10] D. G. Luenberger. *Optimization by vector space methods*. Wiley, New York, NY, 1969.
- [11] M. Putinar. Positive polynomials on compact semialgebraic sets. *Indiana University Mathematics Journal*, 42:969–984, 1993.
- [12] I. Pólik, T. Terlaky, and Y. Zinchenko. SeDuMi: a package for conic optimization. IMA workshop on Optimization and Control, Univ. Minnesota, Minneapolis, Jan. 2007.