

Controller Synthesis for Discrete-Time Polynomial Systems via Occupation Measures

Wei-qiao Han and Russ Tedrake

Abstract—In this paper, we design nonlinear state feedback controllers for discrete-time polynomial dynamical systems via the occupation measure approach. We propose the discrete-time controlled Liouville equation, and use it to formulate the controller synthesis problem as an infinite-dimensional linear programming problem on measures, which is then relaxed as finite-dimensional semidefinite programming problems on moments of measures and their duals on sums-of-squares polynomials. Nonlinear controllers can be extracted from the solutions to the relaxed problems. The advantage of the occupation measure approach is that we solve convex problems instead of generally non-convex problems, and the computational complexity is polynomial in the state and input dimensions, and hence the approach is more scalable. In addition, we show that the approach can be applied to over-approximating the backward reachable set of discrete-time autonomous polynomial systems and the controllable set of discrete-time polynomial systems under known state feedback control laws. We illustrate our approach on several dynamical systems.

I. INTRODUCTION

Given a discrete-time polynomial dynamical system and a target set in state space, we are interested in designing controllers that steer the system to the target set without violating state or control input constraints. Controller synthesis for polynomial systems is a challenging problem in robotics and control. Traditional approaches include designing a linear quadratic regulator (LQR) based on linearized dynamics in a neighborhood of the fixed point, model predictive control (MPC), feedback linearization, dynamic programming, and Lyapunov-based approaches. These approaches each have their limitations. LQR control and linear MPC only work for a small region around the fixed point. To plan for the entire state space, the LQR-Trees method [23] and the approximate explicit-MPC method [15] have been invented. Feedback linearization does not work if there are limits on the inputs. Dynamic programming only works for systems with small dimensionality. Lyapunov-based approaches are generally non-convex, but can be convexified by incorporating the integrator into the controller structure [16] or adding delayed states in the Lyapunov function [17].

Recently the area has seen the development of the occupation measure approach [11] (also known as the Lasserre hierarchy strategy on occupation measures [5]). The general framework of the approach is to first formulate the problem as an infinite-dimensional LP on measures and its dual on continuous functions, and to then approximate the LP by a hierarchy of finite-dimensional semidefinite programming

(SDP) programs on moments of measures and their duals on sums-of-squares (SOS) polynomials. The earliest notable application of the approach is the outer approximation of the region of attraction of continuous-time polynomial systems [6]. The advantage of the approach is that the problem is formulated as a series of convex optimization problems instead of general non-convex problems, and theoretically the approximation to the real set can be made arbitrarily close. Since then the occupation measure approach has been attracting increasing attention and study. It has been applied to the approximation of the region of attraction, the backward reachable set, and the maximum controllable set for continuous-time polynomial systems [8]–[10], [21]. It has also been applied to controller synthesis for continuous-time nonhybrid/hybrid polynomial systems [9], [14], [25].

Studies on discrete-time polynomial systems, however, are relatively sparse compared to those on continuous-time polynomial systems. In [20], the authors considered the discrete-time nonlinear stochastic optimal control problem, which can be interpreted in terms of the Bellman equation. In [13], the authors proposed the discrete-time Liouville equation and used it to formulate an optimization problem that approximates the forward reachable set of discrete-time autonomous polynomial systems.

We are particularly interested in discrete-time systems. One reason is that any physical system simulated by a digital computer is discrete in time, and the control input sent by the digital computer is also discrete in time. When modeling robots making and breaking contact with the environment, the continuous-time systems using some contact models need to handle measure differential inclusions for impacts [19], while the discrete-time models equally capture the complexity of the constrained hybrid dynamics without worrying about impulsive events and event detection [4], [15]. As another related example, in N -step capturability analysis used to study balancing in legged robots, the decision-making is discrete on a footstep-to-footstep level, and the entire problem formulation is asking about the viability kernel, also known as the backward reachable set [18].

In this paper, we propose a controller synthesis method for discrete-time polynomial systems via the occupation measure approach. We propose the discrete-time controlled Liouville equation, and use it to formulate the problem as an infinite-dimensional LP, approximated by a family of finite-dimensional SDP's. By solving SDP's of certain degrees, we are able to extract controllers as polynomials of the corresponding degrees. Unlike Lyapunov-based approaches, our controller synthesis process does not simultaneously return

the controllable region, and hence the stability of the closed-loop system has to be checked a posteriori. Nevertheless, we show that our approach can be applied to over-approximating both the backward reachable set of discrete-time autonomous polynomial systems and the controllable set of discrete-time polynomial systems given any polynomial state feedback control law. We illustrate our approach on several dynamical systems. Our work can be viewed as the discrete-time counterpart of [14], and a pathway towards controller synthesis for discrete-time hybrid polynomial systems.

II. PROBLEM FORMULATION

A. Problem statement

Let $n, m \in \mathbb{N}$. Consider the discrete-time control-affine polynomial system

$$x_{t+1} = \phi(x_t, u_t) := f(x_t) + g(x_t)u_t.$$

The sets $X \subseteq \mathbb{R}^n$ and $U \subseteq \mathbb{R}^m$ are state and control input constraint sets, respectively. The vectors $x_t \in X$ and $u_t \in U$ represent states and control inputs at time $t \in \mathbb{N}$, respectively. $f(x)$ and $g(x)$ are polynomial maps. Denote the target set by $Z \subseteq X$. Our goal is to design a polynomial state feedback controller $u_t = u(x_t) \in U$ that steers the system to the target set Z without violating state and control input constraints.

$\mathbb{R}[x]$ (resp. $\mathbb{R}[u]$) stands for the set of polynomials in the variable $x = (x_1, \dots, x_n)$ (resp. $u = (u_1, \dots, u_m)$). $\mathbb{R}_{2r}[x]$ (resp. $\mathbb{R}_{2r}[u]$) stands for the set of polynomials in the variable x (resp. u) of degree at most $2r$.

Assume

$$X := \{x \in \mathbb{R}^n : h_i^X(x) \geq 0, h_i^X(x) \in \mathbb{R}[x], i = 1, \dots, n_X\},$$

is a compact basic semi-algebraic set. Furthermore, assume that the moments of the Lebesgue measure on X are available. For example, if X is an n -dimensional ball or box, then it satisfies this assumption.

Assume

$$U := \{u \in \mathbb{R}^m : h_i^U(u) \geq 0, h_i^U(u) \in \mathbb{R}[u], i = 1, \dots, n_U\} \\ = [a_1, b_1] \times \dots \times [a_m, b_m],$$

where $a_i, b_i \in \mathbb{R}, i = 1, \dots, m$ are upper and lower limits on control inputs. Furthermore, without loss of generality, assume

$$U := [-1, 1]^m,$$

because the dynamics equation can be scaled and shifted.

Assume

$$Z := \{x \in \mathbb{R}^n : h_i^Z(x) \geq 0, h_i^Z(x) \in \mathbb{R}[x], i = 1, \dots, n_Z\},$$

is a compact basic semi-algebraic set. In practice, we may choose Z to be a small ball or box around the origin.

B. Notations

In this subsection, we introduce some notations in real analysis, functional analysis, and polynomial optimization. For an introduction to these three subjects, please refer to [3], [1], and [12], respectively.

Let $X \subseteq \mathbb{R}^n$ be a compact set. $\mathcal{C}(X)$ denotes the Banach space of continuous functions on X equipped with the sup-norm. Its topological dual, denoted by $\mathcal{C}'(X)$, is the set of all continuous linear functionals on $\mathcal{C}(X)$. $\mathcal{M}(X)$ denotes the Banach space of finite signed Radon measures on the Borel σ -algebra $\mathcal{B}(X)$ equipped with the total variation norm. By Riesz Representation Theorem, $\mathcal{M}(X)$ is isometrically isomorphic to $\mathcal{C}'(X)$. $\mathcal{C}_+(X)$ (resp. $\mathcal{M}_+(X)$) denotes the cone of non-negative elements of $\mathcal{C}(X)$ (resp. $\mathcal{M}(X)$). The topology in $\mathcal{C}_+(X)$ is the strong topology of uniform convergence while the topology in $\mathcal{M}_+(X)$ is the weak-star topology. For any $A \in \mathcal{B}(X)$, λ_A denotes the restriction of the Lebesgue measure on A . For $\mu, \nu \in \mathcal{M}(X)$, we say μ is dominated by ν , denoted by $\mu \leq \nu$, if $\nu - \mu \in \mathcal{M}_+(X)$.

Define $r_i^X := \lceil \deg h_i^X / 2 \rceil, i = 1, \dots, n_X$, $r_i^U := \lceil \deg h_i^U / 2 \rceil, i = 1, \dots, n_U$, and $r_i^Z := \lceil \deg h_i^Z / 2 \rceil, i = 1, \dots, n_Z$. $\Sigma[x]$ (resp. $\Sigma_r[x]$) denotes the cone of SOS polynomials (resp. SOS polynomials of degree up to $2r$) in the variable x . \mathbf{Q}_r^X (resp. $\mathbf{Q}_r^{XU}, \mathbf{Q}_r^Z$) denotes the r -truncated quadratic module generated by the defining polynomials of X (resp. $X \times U, Z$), assuming $h_0^X(x) = 1$ (resp. $h_0^U(u) = 1, h_0^Z(x) = 1$):

$$\mathbf{Q}_r^X := \left\{ \sum_{i=0}^{n_X} \sigma_i(x) h_i^X(x) : \sigma_i \in \Sigma_{r-r_i^X}[x], i = 0, \dots, n_X \right\},$$

$$\mathbf{Q}_r^{XU} := \left\{ \sum_{i=0}^{n_X} \sigma_i^X(x, u) h_i^X(x) + \sum_{i=0}^{n_U} \sigma_i^U(x, u) h_i^U(u) : \right. \\ \left. \sigma_i^X \in \Sigma_{r-r_i^X}[x, u], \sigma_j^U \in \Sigma_{r-r_j^U}[x, u], \right. \\ \left. i = 0, \dots, n_X, j = 0, \dots, n_U \right\},$$

$$\mathbf{Q}_r^Z := \left\{ \sum_{i=0}^{n_Z} \sigma_i(x) h_i^Z(x) : \sigma_i \in \Sigma_{r-r_i^Z}[x], i = 0, \dots, n_Z \right\}.$$

A set $\Omega = \{x \in \mathbb{R}^n : h_i(x) \geq 0, h_i(x) \in \mathbb{R}[x], i = 1, \dots, n_\Omega\} \subseteq \mathbb{R}^n$ is said to satisfy Putinar's condition if there exists $\sigma \in \mathbb{R}[x]$ such that $\sigma = \sigma_0 + \sum_{i=1}^{n_\Omega} \sigma_i h_i$ for some $\{\sigma_i\}_{i=0}^{n_\Omega} \subset \Sigma[x]$, and the level set $\{x \in \mathbb{R}^n : \sigma(x) \geq 0\}$ is compact. Putinar's condition can be satisfied by including the polynomial $N - \|x\|_2^2$, where N is a sufficiently large real number, in the defining polynomials $\{h_i\}$. Here x represents a general variable with n components, so for the set $X \times U$, we consider the field \mathbb{R}^{m+n} , the polynomial ring $\mathbb{R}[x, u]$, and the cone $\Sigma[x, u]$.

III. OPTIMIZATION FORMULATION

A. Discrete-time controlled Liouville equation

The Liouville equation for continuous-time systems is a partial differential equation describing the evolution of the system state over time. The discrete-time analogue of the Liouville equation was studied in Markov decision process, and was incorporated into the occupation measure approach

in [10], [13], [20]. For our controller synthesis purpose, we are going to propose a new form of the Liouville equation, which we call the *discrete-time controlled Liouville equation*.

Given measurable spaces (X_1, \mathcal{A}_1) and (X_2, \mathcal{A}_2) , a measurable function $p : X_1 \rightarrow X_2$ and a measure $\nu : \mathcal{A}_1 \rightarrow [0, +\infty]$, the pushforward measure of ν is defined to be $p_*\nu : \mathcal{A}_2 \rightarrow [0, +\infty]$

$$p_*\nu(A) := \nu(p^{-1}(A))$$

for all $A \in \mathcal{A}_2$. Define π to be the projection map from $X \times U$ to X , i.e., $\pi : X \times U \rightarrow X, (x, u) \mapsto x$. The system dynamics $\phi : X \times U \rightarrow X$ is as defined in the previous section. Let $X_0, X_T \subseteq X$ be the measurable sets containing all possible initial states and final states of the system, respectively. The discrete-time controlled Liouville equation is

$$\mu + \pi_*\nu = \phi_*\nu + \mu_0, \quad (1)$$

where $\mu_0 \in \mathcal{M}_+(X_0)$, $\mu \in \mathcal{M}_+(X_T)$ and $\nu \in \mathcal{M}_+(X \times U)$.

We can view the initial measure μ_0 as the distribution of the mass of the initial states of the system trajectories (not necessarily normalized to 1), the occupation measure ν as describing the volume occupied by the trajectories, and the final measure μ as the distribution of the mass of the final states of the system trajectories. For example, $\mu_0 = \delta_{x_0}$, $\nu = \delta_{(x_0, u_0)} + \dots + \delta_{(x_{T-1}, u_{T-1})}$, and $\mu = \delta_{x_T}$ is a solution to the controlled Liouville equation, describing the system trajectory $\{x_0, x_1 = \phi(x_0, u_0), \dots, x_T = \phi(x_{T-1}, u_{T-1})\}$, where δ_x is the Dirac measure centered at x . It is possible that the measure ν can be disintegrated as $\nu_1(du|x)\nu_2(dx)$ for some measure ν_2 on X and some probability measure $\nu_1(du|x)$ on $U(x)$ for every $x \in X$, as noted in [20].

B. Primal-dual infinite-dimensional LP

We formulate the infinite-dimensional LP on measures as follows:

$$\begin{aligned} \sup \quad & \int_X 1 d\mu_0 \\ \text{s.t.} \quad & \mu + \pi_*\nu = \phi_*\nu + \mu_0, \\ & \mu_0 + \hat{\mu}_0 = \lambda_X, \\ & \mu_0, \hat{\mu}_0 \in \mathcal{M}_+(X), \mu \in \mathcal{M}_+(Z), \\ & \nu \in \mathcal{M}_+(X \times U). \end{aligned} \quad (2)$$

The objective is to maximize the mass of the initial measure. The first constraint is the controlled Liouville equation. Notice that we require the final measure μ to be supported on Z . This constraint, together with the objective, means that we want as many system trajectories as possible to land in Z . The second constraint ensures that the initial measure is dominated by the Lebesgue measure on X , and if the optimal solution is achieved, then the initial measure would be the Lebesgue measure on a set of initial states whose trajectories end in Z and the optimal value is the volume of the set (similar to the idea in Theorem 3.1 in [7]).

The dual LP on continuous functions is given by

$$\begin{aligned} \inf \quad & \int_X w(x) d\lambda_X \\ \text{s.t.} \quad & v(x) - v(\phi(x, u)) \geq 0, \forall x \in X, \forall u \in U, \\ & w(x) - v(x) - 1 \geq 0, \forall x \in X, \\ & w(x) \geq 0, \forall x \in X, \\ & v(x) \geq 0, \forall x \in Z, \\ & v, w \in \mathcal{C}(X). \end{aligned} \quad (3)$$

IV. SEMIDEFINITE RELAXATIONS

We have formulated the infinite-dimensional LP on measures and its dual on continuous functions, but we cannot solve them directly. A practical solution is to approximate the original LP by a family of finite-dimensional SDP's. This relaxation is based on the idea that measures can be characterized by their moments, just as signals can be characterized by their Fourier coefficients. By solving the relaxed SDP's of certain degrees, we can extract controllers in the form of polynomials of corresponding degrees. In this section, we first introduce some background knowledge on moments of measures. For more detailed treatments, please refer to [12]. Next we formulate the relaxed SDP's on moments of measures and their dual on SOS polynomials. Finally, we show how to extract controllers from the SDP solutions.

A. Preliminaries

Any polynomial $p(x) \in \mathbb{R}[x]$ can be expressed in the monomial basis as

$$p(x) = \sum_{\alpha} p_{\alpha} x^{\alpha},$$

where $\alpha \in \mathbb{N}^n$, and $p(x)$ can be identified with its vector of coefficients $p := (p_{\alpha})$ indexed by α . Any measure μ is characterized by its sequence of moments, defined by

$$\int x^{\alpha} d\mu, \alpha \in \mathbb{N}^n.$$

Given a sequence of real numbers $y := (y_{\alpha})$, we define the linear functional $\ell_y : \mathbb{R}[x] \rightarrow \mathbb{R}$ by

$$\ell_y(p(x)) := p^{\top} y = \sum_{\alpha} p_{\alpha} y_{\alpha}.$$

If $y = (y_{\alpha})$ is a sequence of moments for some measure μ , i.e.,

$$y_{\alpha} = \int x^{\alpha} d\mu,$$

then μ is called a representing measure for y . If y has a representing measure μ , then the linear function ℓ_y is the same as integration with respect to μ :

$$\int p d\mu = \int \sum_{\alpha} p_{\alpha} x^{\alpha} d\mu = \sum_{\alpha} p_{\alpha} \int x^{\alpha} d\mu = \ell_y(p(x)).$$

Given $r \in \mathbb{N}$, define $\mathbb{N}_r^n = \{\beta \in \mathbb{N}^n : |\beta| := \sum_i \beta_i \leq r\}$. Define the moment matrix $M_r(y)$ of order r with entries indexed by multi-indices α (rows) and β (columns)

$$[M_r(y)]_{\alpha, \beta} := \ell_y(x^{\alpha} x^{\beta}) = y_{\alpha + \beta}, \forall \alpha, \beta \in \mathbb{N}_r^n.$$

If y has a representing measure, then $M_r(y) \succeq 0, \forall r \in \mathbb{N}$. However, the converse is generally not true.

Given a polynomial $u(x) \in \mathbb{R}[x]$ with coefficient vector $u = (u_\gamma)$, define the localizing matrix w.r.t. y and u to be the matrix indexed by multi-indices α (rows) and β (columns)

$$[M_r(uy)]_{\alpha,\beta} := \ell_y(u(x)x^\alpha x^\beta) \\ = \sum_{\gamma} u_\gamma y_{\gamma+\alpha+\beta}, \forall \alpha, \beta \in \mathbb{N}_r^n.$$

If y has a representing measure μ , then $M_r(uy) \succeq 0$ whenever the support of μ is contained in $\{x \in \mathbb{R}^n : u(x) \geq 0\}$. Conversely, if X is a compact semi-algebraic set as defined in Section II, if X satisfies Putinar's condition, and if $M_r(h_j^X y) \succeq 0, j = 0, \dots, n_X, \forall r$, then y has a finite Borel representing measure with support contained in X (Theorem 3.8(b) in [12]).

B. Primal-dual finite-dimensional SDP

For each $r \geq r_{\min} := \max_{i,j,k} \{r_i^X, r_j^U, r_k^Z\}$, let $y_0 = (y_{0,\beta}), \beta \in \mathbb{N}_{2r}^n$, be the finite sequence of moments up to degree $2r$ of the measure μ_0 . Similarly, y_1, \hat{y}_0, y^X , and z are finite sequences of moments up to degree $2r$ associated with measures $\mu, \hat{\mu}_0, \lambda_X$, and ν , respectively. Let $d := \text{degree } \phi$. The infinite-dimensional LP on measures (2) can be relaxed with the following semidefinite program on moments of measures:

$$\begin{aligned} \sup \quad & y_{0,0} \\ \text{s.t.} \quad & y_{1,\beta} + \ell_z(x^\beta) = \ell_z(\phi(x, u)^\beta) + y_{0,\beta}, \forall \beta \in \mathbb{N}_{2r}^n, \\ & y_{0,\beta} + \hat{y}_{0,\beta} = y_\beta^X, \forall \beta \in \mathbb{N}_{2r}^n, \\ & \mathbf{M}_{r-r_j^X}(h_j^X y_0) \succeq 0, j = 1, \dots, n_X, \\ & \mathbf{M}_{r-r_j^X}(h_j^X \hat{y}_0) \succeq 0, j = 1, \dots, n_X, \\ & \mathbf{M}_{r-d-r_j^X}(h_j^X z) \succeq 0, j = 1, \dots, n_X, \\ & \mathbf{M}_{r-d-r_j^U}(h_j^U z) \succeq 0, j = 1, \dots, n_U, \\ & \mathbf{M}_{r-r_j^Z}(h_j^Z y_1) \succeq 0, j = 1, \dots, n_Z. \end{aligned} \quad (4)$$

The dual of (4) is the following SDP on polynomials of degrees up to $2r$:

$$\begin{aligned} \inf_{v,w} \quad & \sum_{\beta \in \mathbb{N}_{2r}^n} w_\beta y_\beta^X \\ \text{s.t.} \quad & v - v \circ \phi \in \mathbf{Q}_{rd}^{XU}, \\ & w - v - 1 \in \mathbf{Q}_r^X, \\ & w \in \mathbf{Q}_r^X, v \in \mathbf{Q}_r^Z, \\ & v, w \in \mathbb{R}_{2r}[x], \end{aligned} \quad (5)$$

where \circ denotes function composition. The dual SDP (5) is a strengthening of the dual LP (3) by requiring nonnegative polynomials in (3) to be SOS polynomials up to certain degrees.

C. Controller extraction

The controllers can be extracted from the primal SDP (4) as in [9], [14]. We describe the procedure in detail in the following.

Fix $r \in \mathbb{N}$ in the SDP's (4) and (5). Let each u_i be a degree- r polynomial in x , $i = 1, \dots, m$. Identify u_i with its vector of coefficients $(u_{i,\alpha})$. ν is a measure supported on $X \times U$. By solving the primal SDP (4), we obtain the moments of ν (as subsequences of z):

$$\tau_{i,\alpha} := \int x^\alpha u_i d\nu, \forall \alpha \in \mathbb{N}_r^n, \\ \rho_\alpha := \int x^\alpha d\nu, \forall \alpha \in \mathbb{N}_r^n.$$

Then

$$M_r(\rho) \cdot (u_{i,\alpha})_\alpha = (\tau_{i,\alpha})_\alpha,$$

where $(u_{i,\alpha})_\alpha$ is the column vector of coefficients of the polynomial $u_i(x)$ indexed by α , and $(\tau_{i,\alpha})_\alpha$ is the column vector consisting of $\tau_{i,\alpha}$'s indexed by α . The controller $u_i(x)$ can be approximated by taking the pseudo-inverse of the moment matrix $M_r(\rho)$:

$$(u_{i,\alpha})_\alpha = [M_r(\rho)]^+ \cdot (\tau_{i,\alpha})_\alpha.$$

As noted in [9], the approximated controller does not always satisfy the control input constraints. The easiest remedy is to limit the control input to be the boundary values, ± 1 , if the constraints are violated. For all the examples in the Examples Section, we used this method. Most of the time, the control input constraints were not violated. Another method is to solve an SOS optimization problem as in [9].

In general, our controller synthesis method is heuristic. The controllable region needs to be checked a posteriori. In the next section, we show that we can over-approximate the controllable region using a simplified form of our optimization formulation.

V. UNCONTROLLED CASE: OUTER APPROXIMATION OF THE BACKWARD REACHABLE SET

In this section, we consider a special case – the discrete-time autonomous polynomial system

$$x_{t+1} = f(x_t),$$

where $X, f(x)$, and the target set Z are defined as before. Given a time step $T \in \mathbb{N}$, define the T -step backward reachable set

$$X_0^T := \{x_0 \in X : x_t = f(x_{t-1}) = \dots = f^t(x_0) \in Z \\ \text{for some } 0 \leq t \leq T, \text{ and } x_{t_i} \in X, \forall 0 \leq t_i \leq t\}.$$

This is the set of points in X that enter the target region Z within T time steps and whose trajectories do not leave X before entering Z . Once a point enters Z , what happens to it next is not our concern. We are going to over-approximate the backward reachable set

$$X_0^\infty := \bigcup_{T=0}^{\infty} X_0^T,$$

which is the union of all points in X that enter Z in finite time. Denote by \bar{X}_0^∞ the closure of X_0^∞ .

The primal LP is obtained from LP (2) by modifying the Liouville equation to be the same as the one in [13] and modifying the support of the occupation measure ν to be X . The primal and dual LP's are formulated as follows

$$\begin{aligned} p := \sup \quad & \int_X 1 d\mu_0 \\ \text{s.t.} \quad & \nu + \mu = f_*\nu + \mu_0, \\ & \mu_0 + \hat{\mu}_0 = \lambda_X, \\ & \mu_0, \hat{\mu}_0, \nu \in \mathcal{M}_+(X), \\ & \mu \in \mathcal{M}_+(Z). \end{aligned} \quad (6)$$

$$\begin{aligned} d := \inf \quad & \int_X w(x) d\lambda_X \\ \text{s.t.} \quad & v(x) - v(f(x)) \geq 0, \forall x \in X, \\ & w(x) - v(x) - 1 \geq 0, \forall x \in X, \\ & w(x) \geq 0, \forall x \in X, \\ & v(x) \geq 0, \forall x \in Z, \\ & v, w \in \mathcal{C}(X). \end{aligned} \quad (7)$$

Proposition 1. Suppose there exists a constant $M > 0$ such that for any feasible solution $(\mu_0, \hat{\mu}_0, \nu, \mu)$ of the LP (6), the mass of ν is bounded by M , i.e., $\int_X 1 d\nu < M$.

(a) If $f(X) \subseteq X$, then LP (6) admits an optimal solution $(\mu_0^*, \hat{\mu}_0^*, \nu^*, \mu^*)$ such that $\mu_0^* = \lambda_{X_0^\infty}$ and $p^* = \text{vol} X_0^\infty$.

(b) There is no duality gap between the primal LP (6) and the dual LP (7). \square

The semidefinite relaxations can be obtained similarly. The primal is

$$\begin{aligned} p_r := \sup_{y_0, \hat{y}_0, z, a} \quad & y_{0,0} \\ \text{s.t.} \quad & y_{1,\beta} + z_\beta = \ell_z(f(x)^\beta) + y_{0,\beta}, \forall \beta \in \mathbb{N}_{2r}^n, \\ & y_{0,\beta} + \hat{y}_{0,\beta} = y_\beta^X, \forall \beta \in \mathbb{N}_{2r}^n, \\ & \mathbf{M}_{r-r_j^X}(h_j^X y_0) \succeq 0, j = 1, \dots, n_X, \\ & \mathbf{M}_{r-r_j^X}(h_j^X \hat{y}_0) \succeq 0, j = 1, \dots, n_X, \\ & \mathbf{M}_{r-d-r_j^X}(h_j^X z) \succeq 0, j = 1, \dots, n_X, \\ & \mathbf{M}_{r-r_j^Z}(h_j^Z y_1) \succeq 0, j = 1, \dots, n_Z. \end{aligned} \quad (8)$$

The dual is

$$\begin{aligned} d_r := \inf_{v, w} \quad & \sum_{\beta \in \mathbb{N}_{2r}^n} w_\beta y_\beta^X \\ \text{s.t.} \quad & v - v \circ f \in \mathbf{Q}_{rd}^X, \\ & w - v - 1 \in \mathbf{Q}_r^X, \\ & w \in \mathbf{Q}_r^X, v \in \mathbf{Q}_r^Z, \\ & v, w \in \mathbb{R}_{2r}[x]. \end{aligned} \quad (9)$$

Proposition 2. Let $r \geq r_{\min}$.

(a) The primal SDP (8) and the dual SDP (9) are both feasible. If the primal SDP (8) has a strictly feasible solution, then there is no duality gap between the primal SDP (8)

and the dual SDP (9), and the optimal value of SDP (9) is attained.

(b) Let (v_r, w_r) be a feasible solution to SDP (9). Define

$$X_{0r} = \{x \in X | w_r(x) - 1 \geq 0\}.$$

Then $X_{0r} \supseteq \bar{X}_0^\infty \supseteq X_0^\infty$. Suppose the conditions in Part (a) hold. In addition, if there exists a sequence of polynomials $(u_k)_{k \in \mathbb{N}}$ satisfying (i) $u_k > 1_{\bar{X}_0^\infty}$ on X , (ii) (u_k) converges to $1_{\bar{X}_0^\infty}$ in L^1 norm, and (iii) $u_k(x) - u_k(f(x)) > 0, \forall x \in X$, then SDP (9) has an optimal solution (v_r, w_r) such that

$$\lim_{r \rightarrow \infty} \int_X |w_r(x) - 1_{\bar{X}_0^\infty}(x)| d\lambda_X = 0.$$

\square

Remark. Part (a) is a standard strong duality theorem for SDP's. Part (b) indicates that X_{0r} is an outer approximation of the closure of the backward reachable set. If we define

$$\tilde{X}_{0r} = \bigcap_{k=r_{\min}}^r X_{0k},$$

then the approximation by the sequence of sets $\{\tilde{X}_{0r}\}_r$ is monotone. The last technical condition in Part (b) can be understood as follows. Since \bar{X}_0^∞ is closed, the indicator function $1_{\bar{X}_0^\infty}$ is upper semi-continuous. So there exists a decreasing sequence of bounded continuous functions $(u_k)_{k \in \mathbb{N}}$ converging pointwise to $1_{\bar{X}_0^\infty}$ on X . By the Dominated Convergence Theorem, $(u_k)_{k \in \mathbb{N}}$ converges to $1_{\bar{X}_0^\infty}$ in L^1 norm. By the Stone-Weierstrass Theorem, each u_k can be approximated uniformly arbitrarily well by polynomials. Therefore, there exists a sequence of polynomials $(\tilde{u}_k)_{k \in \mathbb{N}}$ satisfying conditions (i) and (ii). So (iii) is an additional constraint. If (iii) holds, then Putinar's Positivstellensatz implies that SDP (9) has a feasible solution whose w -component resembles (\tilde{u}_k) . This establishes the vanishing error of the hierarchical SDP approximations. \square

In practice, however, given a system it is not known a priori if condition (iii) holds or not. Even if it is known, current numerical solvers can only handle SDP's up to a certain degree. Whether the approximation up to that degree is good or not is not known.

While the approach approximates the backward reachable set of autonomous systems, it can also approximate the backward controllable set of systems subject to polynomial state feedback control inputs. This is immediately seen by plugging the polynomial control law $u_t = u(x_t)$ into the control affine polynomial system $x_{t+1} = f(x_t) + g(x_t)u_t = f(x_t) + g(x_t)u(x_t)$, yielding a polynomial closed-loop dynamical system.

VI. EXAMPLES

We illustrate our methods on five discrete-time polynomial systems. All computations are done using MATLAB 2016b, the SDP solver MOSEK 8, and the polynomial optimization toolbox Spotless [24].

A. Van der Pol oscillator

In this example, we are going to over approximate the backward reachable set of the uncontrolled reversed-time Van der Pol oscillator (Example 9.2 in [6]) given by

$$\begin{aligned}\dot{x}_1 &= -2x_2, \\ \dot{x}_2 &= 0.8x_1 + 10(x_1^2 - 0.21)x_2.\end{aligned}$$

Discretizing the model with the explicit Euler scheme with a sampling time $\delta t = 0.01$, the discrete-time system is

$$\begin{aligned}x_1^+ &= (-2x_2)\delta t + x_1, \\ x_2^+ &= (0.8x_1 + 10(x_1^2 - 0.21)x_2)\delta t + x_2.\end{aligned}$$

Choose $X = \{x \in \mathbb{R}^2 : |x_1|^2 \leq 1.5^2, |x_2|^2 \leq 1.5^2\}$ and $Z = \{x \in \mathbb{R}^2 : |x_1|^2 \leq 0.1^2, |x_2|^2 \leq 0.1^2\}$.

We approximate the backward reachable set by degree-14 and 16 polynomials. As show in Figure 1, the gray areas are the approximate backward reachable sets. The areas enclosed by the red lines are the true backward reachable set, which was obtained analytically by integrating backwards in time.




Fig. 1. Degree-14 and degree-16 outer approximations to the backward reachable set (or region of attraction).

B. Double integrator

Consider a double integrator discretized by the explicit Euler scheme with a sampling time $\delta t = 0.01$. The discrete-time dynamics equations are

$$\begin{aligned}x_1^+ &= x_1 + 0.01x_2, \\ x_2^+ &= x_2 + 0.01u.\end{aligned}$$

We are going to design controllers and then approximate the backward reachable set of the closed loop system. We consider the state constraint set $X = \{x \in \mathbb{R}^2 : |x_1| \leq 1, |x_2| \leq 1\}$, and the target set $Z = \{x \in \mathbb{R}^2 : \|x\|_2^2 \leq 0.05^2\}$. We search for a degree-1 controller.

As shown in the left plot of Figure 2, the green area is a degree-10 approximation of the backward reachable set of the closed loop system. We cover X by a uniform 20×20 grid, and compute the trajectories of the grid vertices under the extracted controller. The red markers represent the vertices that can be steered to Z under the extracted controller in $T = 10^4$ time steps without violating state or control input constraints. In the right plot of Figure 2, we plotted the trajectories of four initial states, $(-0.8, 0.8)$, $(-0.6, -0.6)$, $(0.6, 0.4)$, and $(0.5, -0.68)$, under the extracted controller.




Fig. 2. Left: The green area is a degree-10 approximation of the backward reachable set of the closed loop system. The red markers represent the controllable sample points. Right: Trajectories of four initial states under the extracted controller.

C. Dubin's car

Consider the Dubin's car model (Example 2 in [14])

$$\dot{a} = v \cos(\theta), \dot{b} = v \sin(\theta), \dot{\theta} = \omega,$$

or by a change of coordinates, the Brockett integrator

$$\dot{x}_1 = u_1, \dot{x}_2 = u_2, \dot{x}_3 = x_2 u_1 - x_1 u_2.$$

The system has an uncontrollable linearization and does not admit any continuous time-invariant control law that makes the origin asymptotically stable [2]. We are going to design a polynomial control law for the system.

Discretize the system using the explicit Euler scheme with a sampling time $\delta t = 0.01$. Choose $X = \{x \in \mathbb{R}^3 : \|x\|_\infty \leq 1\}$, and $Z = \{x \in \mathbb{R}^3 : \|x\|_2^2 \leq 0.1^2\}$. We search for a degree-4 controller. We sample the 2D sections $\{x \in X : x_3 = 0\}$ and $\{x \in X : x_2 = 0\}$ uniformly, and compute whether the grid vertices can be steered to Z under the extracted controller in 10^4 time steps. In the left two plots of Figure 3, the red vertices represent the initial states that can be regulated to the target set under the extracted controller, while the blue vertices are the rest. The right plot of Figure 3 shows the trajectories of the eight initial states $(\pm 0.9, \pm 0.9, \pm 0.5)$ under the extracted controller. They all reach the target set Z , represented by a red ball. Some other initial states that cannot reach the target set actually end up somewhere very close to the target set. For example the initial state $(0.8, -0.6, 0.7)$ ends up at $(0, 0, 0.1224)$.




Fig. 3. Top left: The 2D section $\{x \in X : x_3 = 0\}$. Bottom left: The 2D section $\{x \in X : x_2 = 0\}$. The red vertices represent the initial states that can be regulated to the target set under the extracted controller. Right: Trajectories of the eight initial states $(\pm 0.9, \pm 0.9, \pm 0.5)$ under the extracted controller. The red ball in the center is the target set Z .

D. Controlled 3D Van der Pol oscillator

Consider the controlled 3D Van der Pol oscillator (Example 2 in [9]) discretized by the explicit Euler scheme with a sampling time $\delta t = 0.01$. The dynamics are given by

$$\begin{aligned} x_1^+ &= x_1 - 2x_2\delta t \\ x_2^+ &= x_2 + (0.8x_1 - 2.1x_2 + x_3 + 10x_1^2x_2)\delta t \\ x_3^+ &= x_3 + (-x_3 + x_3^3 + 0.5u)\delta t \end{aligned}$$

Let the state constraint set be the unit ball $X = \{x \in \mathbb{R}^3 : \|x\|_2^2 \leq 1\}$ and the target set be $Z = \{x \in \mathbb{R}^3 : \|x\|_2^2 \leq 0.1^2\}$. We search for a degree-1 controller, i.e., an affine controller.




Fig. 4. Left: Controllable sample points in red and uncontrollable points in blue. Right: Trajectories of six initial states under the extracted controller. The red ball in the center is the target set Z .

We choose as our sample points the uniform $5 \times 5 \times 5$ grid vertices that are inside the unit ball X . As shown in the left plot in figure 4, the red dots represent the sample points that can be controlled to the target set under the extracted controller in 10^4 time steps. The blue dots represent those cannot. In the right plot, we show the trajectories of six initial states $(0.6, -0.6, -0.2)$, $(-0.6, -0.6, 0.2)$, $(0.6, 0.2, 0.6)$, $(0.6, -0.2, 0.6)$, $(-0.2, 0.6, -0.6)$, and $(-0.2, -0.6, 0.6)$ under the extracted controller. The red ball in the center represents the target set.

E. Cart-pole system




Fig. 5. The cart-pole system.

Consider balancing the cart-pole system [22], shown in Figure 5, to its upright position, an unstable equilibrium. We are allowed to apply only horizontal force on the cart, so the system is underactuated. The equations of motion are given by

$$\begin{aligned} (m_c + m_p)\ddot{x} + m_p l \ddot{\theta} \cos \theta - m_p l \dot{\theta}^2 \sin \theta &= f, \\ m_p l \ddot{x} \cos \theta + m_p l^2 \ddot{\theta} + m_p g l \sin \theta &= 0. \end{aligned}$$

Let $\mathbf{x} = [x, \theta, \dot{x}, \dot{\theta}]^\top$ and $\mathbf{u} = f$. Choose $m_c = 10, m_p = 1, l = 0.5, g = 9.81, X = \{\mathbf{x} \in \mathbb{R}^4 : |x| \leq 4, |\theta| \leq \pi/6, |\dot{x}| \leq 4, |\dot{\theta}| \leq 2\}$, $f \in [-40, 40]$, and $Z = \{\mathbf{x} \in \mathbb{R}^4 : \|\mathbf{x}\|_\infty \leq 0.5\}$. We Taylor-expand the equation of motion to the third order around the unstable equilibrium $\mathbf{x} = [0, \pi, 0, 0]^\top$, and synthesize a third degree polynomial controller. We sample points uniformly in six 2D sections and compute the controllable points under our controller (represented by red circles in Fig 6) using the true equations of motion. Each section is obtained by setting two variables to be 0. For example, the section in the $x - \theta$ plane is $\{\mathbf{x} \in X : \dot{x} = 0, \dot{\theta} = 0\}$. As a comparison, we also compute the controllable points under the infinite-horizon LQR controller (represented by blue dots in Fig 6) with Q and R being identity matrices.

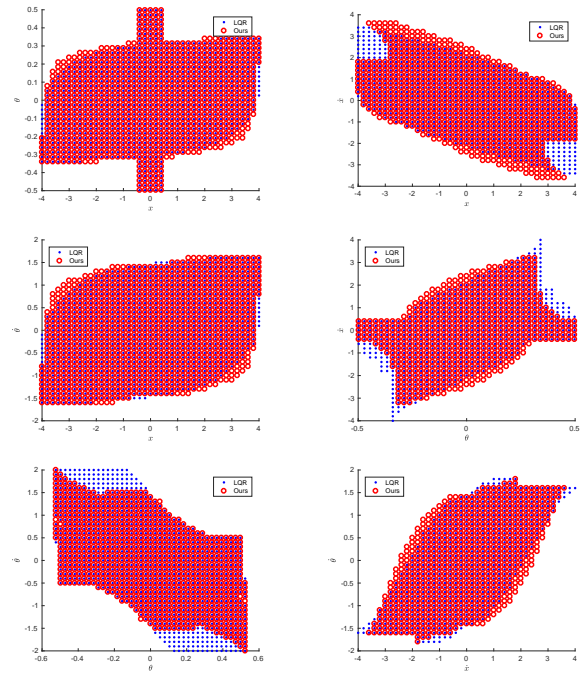


Fig. 6. Controllable sample points in six 2D sections. The red circles represent the controllable points using our controller, while the blue dots represent the controllable points using the infinite-horizon LQR controller.

VII. CONCLUSION

We have presented a controller synthesis method for discrete-time polynomial systems via the occupation measure approach. We have also showed how to over approximate the backward reachable set of a discrete-time autonomous polynomial system and the backward controllable set of a discrete-time polynomial system under state feedback control laws. The advantage of our approach is that we solve convex optimization problems instead of generally non-convex problems, and the computational complexity is polynomial in the state and input dimensions. However, for controller synthesis, our method is heuristic – stability is not guaranteed in any region. In our future work, we will consider the discrete-time hybrid systems.

ACKNOWLEDGMENT

This work was supported by Air Force/Lincoln Laboratory Award No. 7000374874 and Army Research Office Award No. W911NF-15-1-0166.

REFERENCES

- [1] John B Conway. *A course in functional analysis*, volume 96. Springer Science & Business Media, 2013.
- [2] David DeVon and Timothy Bretl. Kinematic and dynamic control of a wheeled mobile robot. In *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, pages 4065–4070. IEEE, 2007.
- [3] Gerald B Folland. *Real analysis: modern techniques and their applications*. John Wiley & Sons, 2013.
- [4] Weiqiao Han and Russ Tedrake. Feedback design for multi-contact push recovery via lmi approximation of the piecewise-affine quadratic regulator. In *Humanoid Robotics (Humanoids), 2017 IEEE-RAS 17th International Conference on*, pages 842–849. IEEE, 2017.
- [5] Didier Henrion. The lasserre hierarchy in robotics. http://webdav.tuebingen.mpg.de/robust_mpc_legged_robots/henrion_slides.pdf, May 2016. Accessed: 2018-02-28.
- [6] Didier Henrion and Milan Korda. Convex computation of the region of attraction of polynomial control systems. *IEEE Transactions on Automatic Control*, 59(2):297–312, 2014.
- [7] Didier Henrion, Jean B Lasserre, and Carlo Savorgnan. Approximate volume and integration for basic semialgebraic sets. *SIAM review*, 51(4):722–743, 2009.
- [8] Milan Korda, Didier Henrion, and Colin N Jones. Inner approximations of the region of attraction for polynomial dynamical systems. *IFAC Proceedings Volumes*, 46(23):534–539, 2013.
- [9] Milan Korda, Didier Henrion, and Colin N Jones. Controller design and region of attraction estimation for nonlinear dynamical systems. *IFAC Proceedings Volumes*, 47(3):2310–2316, 2014.
- [10] Milan Korda, Didier Henrion, and Colin N Jones. Convex computation of the maximum controlled invariant set for polynomial control systems. *SIAM Journal on Control and Optimization*, 52(5):2944–2969, 2014.
- [11] Jean B Lasserre, Didier Henrion, Christophe Prieur, and Emmanuel Trélat. Nonlinear optimal control via occupation measures and lmi-relaxations. *SIAM journal on control and optimization*, 47(4):1643–1666, 2008.
- [12] Jean-Bernard Lasserre. *Moments, positive polynomials and their applications*, volume 1. World Scientific, 2010.
- [13] Victor Magron, Pierre-Loïc Garoche, Didier Henrion, and Xavier Thiriaux. Semidefinite approximations of reachable sets for discrete-time polynomial systems. *arXiv preprint arXiv:1703.05085*, 2017.
- [14] Anirudha Majumdar, Ram Vasudevan, Mark M Tobenkin, and Russ Tedrake. Convex optimization of nonlinear feedback controllers via occupation measures. *The International Journal of Robotics Research*, 33(9):1209–1230, 2014.
- [15] Tobia Marcucci, Robin Deits, Marco Gabiccini, Antonio Biechi, and Russ Tedrake. Approximate hybrid model predictive control for multi-contact push recovery in complex environments. In *Humanoid Robotics (Humanoids), 2017 IEEE-RAS 17th International Conference on*, pages 31–38. IEEE, 2017.
- [16] Mohd Md Saat. *Controller synthesis for polynomial discrete-time systems*. PhD thesis, ResearchSpace@ Auckland, 2013.
- [17] Jose Luis Pitarch, Antonio Sala, Jimmy Lauber, and Thierry-Marie Guerra. Control synthesis for polynomial discrete-time systems under input constraints via delayed-state lyapunov functions. *International Journal of Systems Science*, 47(5):1176–1184, 2016.
- [18] Michael Posa, Twan Koolen, and Russ Tedrake. Balancing and step recovery capturability via sums-of-squares optimization. In *Robotics: Science and Systems*, 2017.
- [19] Michael Posa, Mark Tobenkin, and Russ Tedrake. Stability analysis and control of rigid-body systems with impacts and friction. *IEEE Transactions on Automatic Control*, 61(6):1423–1437, 2016.
- [20] Carlo Savorgnan, Jean B Lasserre, and Moritz Diehl. Discrete-time stochastic optimal control via occupation measures and moment relaxations. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 519–524. IEEE, 2009.
- [21] Victor Shia, Ram Vasudevan, Ruzena Bajcsy, and Russ Tedrake. Convex computation of the reachable set for controlled polynomial hybrid systems. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 1499–1506. IEEE, 2014.
- [22] Russ Tedrake. Underactuated robotics: Algorithms for walking, running, swimming, flying, and manipulation (course notes for mit 6.832). *Downloaded in Fall*, 2014.
- [23] Russ Tedrake, Ian R Manchester, Mark Tobenkin, and John W Roberts. Lqr-trees: Feedback motion planning via sums-of-squares verification. *The International Journal of Robotics Research*, 29(8):1038–1052, 2010.
- [24] Mark M Tobenkin, Frank Permenter, and Alexandre Megretski. Spot-less polynomial and conic optimization, 2013.
- [25] Pengcheng Zhao, Shankar Mohan, and Ram Vasudevan. Optimal control for nonlinear hybrid systems via convex relaxations. *arXiv preprint arXiv:1702.04310*, 2017.