

Controller Synthesis for Discrete-time Hybrid Polynomial Systems via Occupation Measures

Wei-qiao Han and Russ Tedrake

Abstract—We consider the feedback design for stabilizing a rigid body system by making and breaking multiple contacts with the environment without prespecifying the timing or the number of occurrence of the contacts. We model such a system as a discrete-time hybrid polynomial system, where the state-input space is partitioned into several polytopic regions with each region associated with a different polynomial dynamics equation. Based on the notion of occupation measures, we present a novel controller synthesis approach that solves finite-dimensional semidefinite programs as approximations to an infinite-dimensional linear program to stabilize the system. The optimization formulation is simple and convex, and for any fixed degree of approximations the computational complexity is polynomial in the state and control input dimensions. We illustrate our approach on some robotics examples.

I. INTRODUCTION

In robot locomotion and manipulation, it is common that a legged robot balances itself or a robotic hand manipulates an object by making and breaking multiple contacts with the environment. However, there are no simple rules to stabilize such a system. Local stabilization methods, such as linear-quadratic regulator (LQR), are unable to reason about the change of system dynamics. In this paper, we model the system as a hybrid system, or more specifically a piecewise polynomial system, i.e., a system whose state-input space is partitioned into several polytopic regions, with each region associated with a different polynomial dynamics equation. Such a hybrid system modeling approach is a natural fit for many problems in robotics. For example, a system with contacts is hybrid, with each hybrid mode associated with a particular contact mode.

Controller synthesis for hybrid systems is a long-standing challenging problem. The earliest work [1] dates back to 1960s, where the hybrid model and the optimization formulation for controlling the system were introduced. More recently, Branicky et al. considered several algorithms for optimal control of hybrid systems [2], [3] and derived necessary conditions for the existence of optimal control law [4]. Approaches based on dynamic programming [4]–[6] and the maximum principle [7]–[10] were proposed by several authors to solve the hybrid optimal control problem. (More reviews on hybrid systems can be found in [11]–[15].) Approaches based on dynamic programming design

Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA. weiqiaoh, russt@mit.edu

This work was supported by Air Force/Lincoln Laboratory Award No. 7000374874, and Department of the Navy, Office of Naval Research, Award No. N00014-18-1-2210. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Office of Naval Research.




Fig. 1. Examples used in this work. (A) Linear inverted pendulum with an elastic wall. (B) Variable height inverted pendulum (legged robot).

control policies for the whole discretized state space and the computational complexity grows exponentially with respect to the dimension of the state. Approaches based on the maximum principle are trajectory optimization methods, which design a trajectory for a specific initial state, and they usually assume mode switch sequences are fixed, while we are more interested in methods that automatically find mode switch sequences. Recent works based on the hybrid minimum principle and gradients [16], [17] do drop the assumption of fixed mode switch sequences. Posa et al. [18] avoided the mode switch problem by formulating the linear complementarity problem for physical contacts. However, we are interested in designing control policies that cover most of the state space rather than designing a trajectory for a fixed initial state. Sampling based approaches [19]–[21] achieve a compromise between covering the state space with control policies and designing a single trajectory for an initial state, and are promising for high-dimensional systems.

Piecewise-affine (PWA) systems, a special class of hybrid systems, where the state-input space is divided into polytopic regions with each region associated with a different affine dynamics equation, have been extensively studied. PWA systems can arise by linearizing the dynamics of a nonlinear/hybrid system at a few points in the state space [22], [23]. The optimal control policy for PWA systems can be computed offline by (i) dynamic programming and multi-parametric convex programming [24]–[31], or (ii) representing the system as a mixed logical dynamical (MLD) system and then applying multi-parametric mixed integer convex programming [32]–[34], or (iii) enumerating mode switch sequences and applying multi-parametric convex programming [35]. In general, the computational time of these methods grows exponentially as the number of time steps grow until, if possible, convergence. Lyapunov-function-based approaches have also been used for the feedback control of PWA systems, where linear or bilinear matrix inequalities are formed to search for a PWA controller

and a piecewise-quadratic Lyapunov function that proves the stability of the closed-loop system [36]–[38]. The main drawbacks of these approaches are that the formulations are usually too conservative and hence may not be able to find feasible solutions for practical problems, and that bilinear matrix inequalities are non-convex and often intractable. We think that piecewise polynomial functions, which include PWA functions as a special case, are better characterizations of the true dynamics of many systems.

Recent years has seen the development of the occupation measure approach [39]. The general framework of the approach is to first formulate the problem as an infinite-dimensional linear programming (LP) problem 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 approach has been applied to approximating the region of attraction, the maximum controllable set, or the forward/backward reachable set for discrete-time/continuous-time autonomous/controlled hybrid/non-hybrid polynomial systems [40]–[45]. It has also been applied to controller synthesis for those systems [42], [46]–[49].

The closest work to this paper is [49], in which the authors used occupation measures to design controllers for continuous-time hybrid systems. The main difference is that we focus on discrete-time systems as digital computers control robots via discrete signals. Also, we usually test our closed-loop systems in simulators [50] that use time-stepping schemes and complementarity problems [51], [52] to model contacts. Previous works [42], [47] on controller synthesis for discrete-time non-hybrid systems using occupation measures lack the notion of “the state finally landing in the target set”, as is commonly used in continuous-time systems, and hence enforcing the system to reach a target set depends completely on the objective function, which can be hard to tune. Our previous work [48] solved this problem by proposing a new form of the Liouville equation.

In this paper, based on our previous work [48], we find what we call “the one-step transition sets” and formulate an optimization problem for discrete-time hybrid systems with the goal of steering as many initial states to the target set as possible in finite time. As far as we know, this is the first time that occupation measure approach has been applied to controller synthesis for discrete-time hybrid polynomial systems. The advantages of our approach include the following: (i) The approach only needs to solve SDPs, which are convex. For a fixed degree of SDP approximations, the computational complexity is polynomial in the state and control dimensions, and hence in principle our approach is more scalable than those based on dynamic programming. (ii) The approach does not fix any mode switch sequences beforehand, and it finds them automatically. (iii) The approach does controller synthesis for all initial states simultaneously instead of only one initial state. (iv) The approach does not fix the number of time steps. (v) As will be shown by the examples, the controller can take a simple form and is efficient to be applied

online, implying low degrees of the SDP approximations work well in some cases. In the last section, we will discuss the limitations of our approach. For example, the controllable set has to be checked a posteriori.

A. Notations

Let \mathbb{N} (resp. \mathbb{R}) be the set of nonnegative integers (resp. real numbers). Let $n \in \mathbb{N}$. Let $\mathbb{R}[x] := \mathbb{R}[x_1, \dots, x_n]$ be the ring of polynomials in $x := (x_1, \dots, x_n)$, and let $\mathbb{R}_r[x]$ be the set of polynomials in $\mathbb{R}[x]$ with degree at most r . Let $\Sigma[x]$ denote the cone of SOS polynomials in x , and $\Sigma_r[x]$ the cone of SOS polynomials of degree up to $2r$.

Let $X \subseteq \mathbb{R}^n$ be a compact set. Let $\mathcal{C}(X)$ denote the space of all continuous functions on X . It is a Banach space equipped with the sup-norm. Its topological dual $\mathcal{C}'(X)$ is the set of all continuous linear functionals on $\mathcal{C}(X)$. $\mathcal{C}_+(X)$ denotes the cone of non-negative elements of $\mathcal{C}(X)$, i.e., it is the cone of non-negative functions on X . If μ is a finite measure on the Borel σ -algebra $\mathcal{B}(X)$, then μ is a Radon measure [53]. Let $\mathcal{M}(X)$ denote the Banach space of finite signed Radon measures on $\mathcal{B}(X)$ equipped with the total variation norm, and $\mathcal{M}_+(X)$ is the cone of (unsigned) Radon measures. By Riesz Representation Theorem, $\mathcal{M}(X)$ is isometrically isomorphic to $\mathcal{C}'(X)$. The topology in $\mathcal{M}_+(X)$ is the weak-star topology, while the topology in $\mathcal{C}_+(X)$ is the strong topology of uniform convergence. For $\mu, \nu \in \mathcal{M}(X)$, if $\nu - \mu \in \mathcal{M}_+(X)$, then we say μ is dominated by ν , and denote it by $\mu \leq \nu$. For a measure space (X, \mathcal{A}, μ) and $V \in \mathcal{A}$, let $\mathcal{A}|_V = \{A \subseteq V | A \in \mathcal{A}\}$ and let $\mu_V : \mathcal{A}|_V \rightarrow [0, +\infty]$ denote the restriction of μ on $\mathcal{A}|_V$, i.e., $\mu_V(A) = \mu(A)$, for $A \in \mathcal{A}|_V$. More reviews on real and functional analysis can be found in [54], [55].

II. PROBLEM FORMULATION

Let $n, m, s \in \mathbb{N}$. Consider the discrete-time hybrid control-affine polynomial system $\mathcal{H} = (X, U, \mathcal{I}, \mathcal{S}, \phi)$:

- $\mathcal{I} = \{1, \dots, s\}$ is the set of indices of the state space cells. Each (discrete state) $i \in \mathcal{I}$ is called a “mode”.
- The state constraint set $X \subseteq \mathbb{R}^n$ is the union of all state space cells X_i 's, $X = \bigcup_{i \in \mathcal{I}} X_i$. The state space cell $X_i \subseteq \mathbb{R}^n$ is a compact set with nonempty interior, $\forall i \in \mathcal{I}$. $X_i \cap X_j$ has empty interior, $\forall i \neq j \in \mathcal{I}$.
- $U \subseteq \mathbb{R}^m$ is the control input constraint set.
- The dynamics in the cell X_i is given by $x_{t+1} = \phi_i(x_t, u_t) := f_i(x_t) + g_i(x_t)u_t$, for $(x_t, u_t) \in X_i \times U$, where f_i and g_i are polynomials, $i \in \mathcal{I}$, and the vectors x_t and u_t represent the state and the control input at time t , respectively. The dynamics of the system at time t is completely determined by the cell in which the state x_t resides. The dynamics at the boundary $X_i \cap X_j$ is defined to be ϕ_k , where $k = \min\{i, j\}$.
- $\mathcal{S} = \{(i, j) \in \mathcal{I} \times \mathcal{I} : \exists x_k \in X_i, u_k \in U, \text{ s.t. } x_{k+1} = \phi_i(x_k, u_k) \in X_j\}$ is the set of ordered pairs (i, j) of indices denoting possible switches from cell i to cell j .

Problem 1. Given a discrete-time hybrid control-affine polynomial system \mathcal{H} and a target set $Z \subseteq X$, design a

state feedback controller $u : X \rightarrow U$ that maximizes the volume of the set X_0 where the trajectory starting from X_0 under u reaches Z in finite time.

To solve this problem, we are going to design piecewise polynomial controllers expressed as $u = u^i(x) \in U$ if $x \in X_i$, where $u^i(x)$ is a polynomial in x , $i \in \mathcal{I}$. In the end, our controller would be an approximation of the optimal controller, and we show by examples that the approximate controller is good enough for practical purposes.

Assume $X := \{x \in \mathbb{R}^n | h_j^X(x) \geq 0, h_j^X(x) \in \mathbb{R}[x], j = 1, \dots, n_X\}$, and $X_i := \{x \in \mathbb{R}^n | h_j^{X_i}(x) \geq 0, h_j^{X_i}(x) \in \mathbb{R}[x], j = 1, \dots, n_{X_i}\}, i \in \mathcal{I}$, are compact basic semi-algebraic sets with nonempty interior. Furthermore, assume that the moments of the Lebesgue measure on X_i 's are available. For example, X_i can be an n -dimensional ball or box. Assume that the origin is in the interior of X_1 . This assumption excludes systems where the origin is at the boundaries of multiple cells. By properly modifying the optimization formulation, we can deal with other classes of hybrid systems, but in this paper we only show how to work with this kind of hybrid systems.

Assume $U := \{u \in \mathbb{R}^m | h_j^U(u) \geq 0, h_j^U(u) \in \mathbb{R}[u], j = 1, \dots, n_U\} = [a_1, b_1] \times \dots \times [a_m, b_m]$. In the sequel, without loss of generality, we assume $U := [-1, 1]^m$, because the dynamics equations can be scaled and shifted arbitrarily.

Assume $Z := \{x \in \mathbb{R}^n | h_j^Z(x) \geq 0, h_j^Z(x) \in \mathbb{R}[x], j = 1, \dots, n_Z\}$ is a compact basic semi-algebraic set with nonempty interior. In practice, we may choose Z to be a small ball or box around the origin so that after the system enters Z , we may turn on LQR or some other local control methods to regulate the system to the origin. For this purpose, assume that the origin is in the interior of Z , and that $Z \subseteq X_1 \subseteq X$.

Define the one-step transition sets Y_{ij} , for any $(i, j) \in \mathcal{S}$, to be the preimage of X_j under f_i in $X_i \times U$, i.e., $Y_{ij} := \{(x, u) \in X_i \times U | f_i(x, u) \in X_j\} = f_i^{-1}(X_j) \cap X_i \times U = \{(x, u) \in X_i \times U | h_k^{X_j}(f_i(x, u)) \geq 0, k = 1, \dots, n_{X_j}\}$. The last equality implies that Y_{ij} is a basic semi-algebraic set. So we can write $Y_{ij} = \{(x, u) \in \mathbb{R}^{n+m} | h_k^{Y_{ij}}(x, u) \geq 0, h_k^{Y_{ij}} \in \mathbb{R}[x, u], k = 1, \dots, n_{Y_{ij}}\}, \forall (i, j) \in \mathcal{S}$. Since X_j is compact and hence closed, and since f_i is continuous, $f_i^{-1}(X_j)$ is closed. Since $X_i \times U$ is compact, Y_{ij} , being a closed subset of a compact set, is compact.

Let $r_j^{X_i} := \lceil \deg h_j^{X_i} / 2 \rceil, i \in \mathcal{I}, j = 0, \dots, n_{X_i}$. Let $\mathbf{Q}_r^{X_i} := \{\sum_{j=0}^{n_{X_i}} \sigma_j(x) h_j^{X_i}(x) : \sigma_j \in \Sigma_{r-r_j^{X_i}}[x], j = 0, \dots, n_{X_i}\}$ denote the r -truncated quadratic module generated by the defining polynomials of X_i , assuming $h_0^{X_i}(x) = 1$. Analogously, we define $\mathbf{Q}_r^{X_i U}$, \mathbf{Q}_r^Z , and $\mathbf{Q}_r^{Y_{ij}}$. Note that $\mathbf{Q}_r^{X_i U}$ is generated by defining polynomials of X_i and U .

(Putinar's condition) A compact basic semialgebraic set defined by $\Omega := \{x \in \mathbb{R}^n | h_j^\Omega(x) \geq 0, j = 1, \dots, n_\Omega\}$ satisfies Putinar's condition if there exists $h \in \mathbb{R}[x]$ such that $h = \sigma_0 + \sum_{j=1}^{n_\Omega} \sigma_j h_j$ for some $\{\sigma_j\}_{j=0}^{n_\Omega} \subset \Sigma[x]$ and the level set $\{x \in \mathbb{R}^n | h(x) \geq 0\}$ is compact.

Putinar's condition can be satisfied by including $N - \|x\|_2^2$

in the defining polynomials of Ω for some large $N \in \mathbb{R}$. If Ω satisfies Putinar's condition, then Putinar's Positivstellensatz [56] says that any positive polynomial on Ω has an SOS representation.

III. OPTIMIZATION FORMULATION

A. Discrete-time controlled Liouville equation

The Liouville equation describes the evolution of the state distribution of the system over time. For continuous-time systems, it takes the form of partial differential equations (see e.g. [40]). For discrete-time systems, it is realized by the push forward operator [45], [48]. In the following, we introduce the discrete-time controlled Liouville equation in the non-hybrid setting as proposed in [48], and in the next subsection we shall see how it enables us to formulate optimization problems for hybrid systems.

Given measurable spaces (S_1, \mathcal{A}_1) and (S_2, \mathcal{A}_2) , a measurable function $f : S_1 \rightarrow S_2$, and a measure $\nu : \mathcal{A}_1 \rightarrow [0, +\infty]$, the pushforward measure $f_*\nu : \mathcal{A}_2 \rightarrow [0, +\infty]$ is defined to be $f_*\nu(A) := \nu(f^{-1}(A))$, for all $A \in \mathcal{A}_2$.

Assume $s = 1$, i.e. $X = X_1$ (just for the rest of this subsection). Let $X_0, X_T \subseteq X$ be the measurable sets of all possible initial states and all possible final states of system trajectories, respectively. Define π to be the projection $\pi : X \times U \rightarrow X, (x, u) \mapsto x$. It simply extracts the state from a state-input pair. $\phi : X \times U \rightarrow X$ describes the system dynamics as defined in the previous section. 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 $\nu \in \mathcal{M}_+(X \times U)$.

The initial measure μ_0 can be viewed as the distribution of the mass of the initial states of the system trajectories (not necessarily normalized to 1). The occupation measure ν describes the volume occupied by the trajectories. The final measure μ can be viewed as the distribution of the mass of the final states of the system trajectories. The measures μ_0, ν , and μ will be decision variables in our optimization, which enables us to search over trajectories with certain properties.

B. Primal-dual infinite-dimensional LP

Define the projections $\pi_i : X_i \times U \rightarrow X_i, (x, u) \mapsto x, i \in \mathcal{I}$. The infinite-dimensional LP on measures is formulated as follows:

$$p := \sup_{\mu_0^i, \hat{\mu}_0^i, \mu, \nu_i, \mu_{ij}} \sum_{i \in \mathcal{I}} \int_X 1 d\mu_0^i \quad \text{subject to:} \quad (2)$$

$$\sum_{j:(i,j) \in \mathcal{S}} \pi_{i*}\mu_{ij} + \pi_{i*}\nu_i = \phi_{i*}\nu_i + \mu_0^i + \sum_{j:(j,i) \in \mathcal{S}} \phi_{j*}\mu_{ji},$$

$$\text{for } i \neq 1, i \in \mathcal{I},$$

$$\mu + \pi_{i*}\nu_i = \phi_{i*}\nu_i + \mu_0^i + \sum_{j:(j,i) \in \mathcal{S}} \phi_{j*}\mu_{ji}, \text{ for } i = 1,$$

$$\mu_0^i + \hat{\mu}_0^i = \lambda_{X_i}, \mu_0^i, \hat{\mu}_0^i \in \mathcal{M}_+(X_i), \forall i \in \mathcal{I},$$

$$\mu \in \mathcal{M}_+(Z), \nu_i \in \mathcal{M}_+(X_i \times U), \forall i \in \mathcal{I},$$

$$\mu_{ij} \in \mathcal{M}_+(Y_{ij}), \forall (i, j) \in \mathcal{S}.$$

Remember that we assume $Z \subseteq X_1$. The first constraint is a direct application of the Liouville equation. In this constraint, we are considering any mode i whose state space cell does not contain the origin, i.e., $1 \neq i \in \mathcal{I}$. The measure ν_i is the occupation measure describing the volume occupied by system trajectories in X_i . The final measure $\sum_{j:(i,j) \in \mathcal{S}} \pi_{i*} \mu_{ij}$ consists of the measures supported on the one-step transition sets Y_{ij} for all j such that $(i,j) \in \mathcal{S}$, meaning that if f_i is applied again to those final states, the system will leave X_i and enter other modes. The initial measure $\mu_0^i + \sum_{j:(j,i) \in \mathcal{S}} \phi_{j*} \mu_{ji}$ consists of two parts, μ_0^i describing the distribution of initial states originated from X_i , and $\sum_{j:(j,i) \in \mathcal{S}} \phi_{j*} \mu_{ji}$ describing the distribution of states coming from other modes. The second constraint is the Liouville equation for mode 1. It differs from the first constraint in the final measure, which is a measure μ supported on the target set Z . This is crucial to our goal of controlling the system to the target set, because adding the term $\sum_{j:(i,j) \in \mathcal{S}} \pi_{i*} \mu_{ij}$ to the final measure in the second constraint would cause chattering effects – for many initial states, the system would bounce between two modes and finally stop somewhere at the boundary of two modes. The third constraint ensures that any initial measure originated from mode i is dominated by the Lebesgue measure on X_i , $\forall i$. The objective is to maximize the sum of the mass of initial measures originated from all modes. The combination of the third constraint and the objective means maximizing the volume of the set of initial states that can be controlled to the target set. The whole optimization design is specific to solving Problem 1. For other optimal control purposes, the objective can be modified to be some other reward or cost functions, and the constraints can be modified accordingly.

The dual LP on continuous functions is given by

$$\begin{aligned} \inf_{w_i, v_i, i \in \mathcal{I}} \sum_{i \in \mathcal{I}} \int_{X_i} w_i(x) d\lambda_{X_i} \quad \text{subject to:} \quad (3) \\ v_i(x) - v_i(\phi_i(x, u)) \geq 0, \forall x \in X_i, i \in \mathcal{I}, \forall u \in U, \\ v_i(x) - v_j(\phi_i(x, u)) \geq 0, \forall (x, u) \in Y_{ij}, i \neq 1, (i, j) \in \mathcal{S}, \\ -v_j(\phi_i(x, u)) \geq 0, \forall (x, u) \in Y_{ij}, i = 1, (i, j) \in \mathcal{S}, \\ w_i(x) - v_i(x) - 1 \geq 0, w_i(x) \geq 0, \forall x \in X_i, i \in \mathcal{I}, \\ v_1(x) \geq 0, \forall x \in Z, \quad v_i, w_i \in \mathcal{C}(X_i), i \in \mathcal{I}. \end{aligned}$$

IV. SEMIDEFINITE RELAXATIONS

Although we cannot directly solve the infinite-dimensional LP's, we can approximate them by finite-dimensional SDP's and extract the controller from the solution to the SDP's. Semidefinite approximation to the infinite-dimensional LP is quite standard in the literature (see e.g. [39]). It is based on the idea that a measure can be characterized by its sequence of moments, just as a signal can be characterized by its sequence of Fourier coefficients. In this section, we first introduce some background knowledge about moments of measures, which can also be found in [57]. Next we formulate the relaxed SDP's on moments of measures and their duals on SOS polynomials. Finally, we briefly mention how to extract controllers.

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 . The (sequence of) moments $y := (y_{\alpha})$ for a measure μ is defined to be $y_{\alpha} := \int x^{\alpha} d\mu, \alpha \in \mathbb{N}^n$. Given a sequence of real numbers $y = (y_{\alpha})$, y is not necessarily a sequence of moments for some measure μ . If it is, then μ is called a representing measure for y . Given any y , we define the linear functional $\ell_y : \mathbb{R}[x] \rightarrow \mathbb{R}$ by $\ell_y(p(x)) := p^{\top} y = \sum_{\alpha} p_{\alpha} y_{\alpha}$. Integration of a polynomial p against a measure μ with the sequence of moments y can be expressed as a linear functional $\ell_y: \int p d\mu = \int \sum_{\alpha} p_{\alpha} x^{\alpha} d\mu = \sum_{\alpha} p_{\alpha} \int x^{\alpha} d\mu = \sum_{\alpha} p_{\alpha} y_{\alpha}$.

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) to be $[M_r(y)]_{\alpha, \beta} := \ell_y(x^{\alpha} x^{\beta}) = y_{\alpha+\beta}, \forall \alpha, \beta \in \mathbb{N}_r^n$. The moment matrix can be expressed as a bilinear form $\langle \cdot, \cdot \rangle_y$ on $\mathbb{R}[x]_r$: $\langle p, q \rangle_y := \ell_y(pq) = p^{\top} M_r(y) q, \forall p, q \in \mathbb{R}[x]_r$. 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) $[M_r(uy)]_{\alpha, \beta} := \ell_y(u(x) x^{\alpha} x^{\beta}) = \sum_{\gamma} u_{\gamma} y_{\gamma+\alpha+\beta}, \forall \alpha, \beta \in \mathbb{N}_r^n$. Similar to the moment matrix, we have $\langle p, M_r(uy) q \rangle = \ell_y(upq) = p^{\top} M_r(uy) q, \forall p, q \in \mathbb{R}[x]_r$. 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 Ω is a compact basic semi-algebraic set as defined in Section II, if Putinar's condition holds, and if $M_r(h_j^{\Omega} y) \succeq 0, j = 0, \dots, n_{\Omega}, \forall r$, then y has a finite Borel representing measure with support contained in Ω (Theorem 3.8(b) in [57]).

B. Primal-dual finite-dimensional SDP

For each $r \geq r_{min} := \max_{k_1, k_2, k_3, k_4} \{r_{k_1}^{X_i} (i \in \mathcal{I}), r_{k_2}^U, r_{k_3}^{Y_{ij}} ((i, j) \in \mathcal{S}), r_{k_4}^Z\}$, let $y_0^i = (y_{0, \beta}^i), \beta \in \mathbb{N}_{2r}^n$, be the finite sequence of moments up to degree $2r$ of the measure μ_0^i . Similarly, $y_1, \hat{y}_0^i, y^{X_i}, z_i, y_{ij}$ are finite sequences of moments up to degree $2r$ associated with measures $\mu, \hat{\mu}_0^i, \lambda_{X_i}, \nu_i$, and μ_{ij} , respectively. Let $d_i := \text{degree } \phi_i$. The infinite-dimensional LP on measures (2) can be relaxed with the following semidefinite program on moments of measures:

$$\begin{aligned} p_r := \sup_{y_0^i, y_1, \hat{y}_0^i, z_i, y_{ij}} \sum_{i \in \mathcal{I}} y_{0,0}^i \quad \text{subject to:} \quad (4) \\ \sum_{j:(i,j) \in \mathcal{S}} \ell_{y_{ij}}(x^{\beta}) + \ell_{z_i}(x^{\beta}) = \ell_{z_i}(\phi_i(x, u)^{\beta}) + y_{0, \beta}^i \\ + \sum_{j:(j,i) \in \mathcal{S}} \ell_{y_{ji}}(\phi_j(x, u)^{\beta}), \forall \beta \in \mathbb{N}_{2r}^n, 1 \neq i \in \mathcal{I}, \\ y_{1, \beta} + \ell_{z_i}(x^{\beta}) = \ell_{z_i}(\phi_i(x, u)^{\beta}) + y_{0, \beta}^i \\ + \sum_{j:(j,i) \in \mathcal{S}} \ell_{y_{ji}}(\phi_j(x, u)^{\beta}), \forall \beta \in \mathbb{N}_{2r}^n, i = 1, \\ y_{0, \beta}^i + \hat{y}_{0, \beta}^i = y_{\beta}^{X_i}, \forall \beta \in \mathbb{N}_{2r}^n, i \in \mathcal{I}, \end{aligned}$$

$$\begin{aligned}
\mathbf{M}_{r-r_k^{X_i}}(h_k^{X_i} y_0^i) &\succeq 0, k = 1, \dots, n_{X_i}, i \in \mathcal{I}, \\
\mathbf{M}_{r-r_k^{X_i}}(h_k^{X_i} \hat{y}_0^i) &\succeq 0, k = 1, \dots, n_{X_i}, i \in \mathcal{I}, \\
\mathbf{M}_{r d_i - r_k^{X_i}}(h_k^{X_i} z_i) &\succeq 0, k = 1, \dots, n_{X_i}, i \in \mathcal{I}, \\
\mathbf{M}_{r d_i - r_k^U}(h_k^U z_i) &\succeq 0, k = 1, \dots, n_U, i \in \mathcal{I}, \\
\mathbf{M}_{r d_i - r_k^{Y_{ij}}}(h_k^{Y_{ij}} y_{ij}) &\succeq 0, k = 1, \dots, n_{Y_{ij}}, (i, j) \in \mathcal{S}, \\
\mathbf{M}_{r-r_k^Z}(h_k^Z y_1) &\succeq 0, k = 1, \dots, n_Z.
\end{aligned}$$

Proposition 1. Assume all assumptions in Section II are satisfied. Assume the sets X_i for $i \in \mathcal{I}$, Y_{ij} for $(i, j) \in \mathcal{S}$, and Z all satisfy Putinar's condition. Suppose there exists $r_0 \geq r_{min}$ such that (i) for every $r \geq r_0, r \in \mathbb{N}$, SDP (4) is feasible and the optimal solution to SDP (4) exists, and (ii) there exists a ball of constant radius (independent of r) in the solution space such that it contains the feasible set of the SDP (4) for all $r \geq r_0, r \in \mathbb{N}$. Then $\lim_{r \rightarrow \infty} p_r = p$.

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

$$\begin{aligned}
\inf_{v, w} \sum_{i \in \mathcal{I}} \sum_{\beta \in \mathbb{N}_{2r}^{X_i}} (w_i)_\beta y_\beta^{X_i} \quad &\text{subject to:} \quad (5) \\
v_i - v_i \circ \phi_i &\in \mathbf{Q}_{r d_i}^{X_i U}, i \in \mathcal{I}, \\
v_i - v_j \circ \phi_i &\in \mathbf{Q}_{r d_i}^{Y_{ij}}, i \neq 1, (i, j) \in \mathcal{S}, \\
-v_j \circ \phi_i &\in \mathbf{Q}_{r d_i}^{Y_{ij}}, i = 1, (i, j) \in \mathcal{S}, \\
w_i - v_i - 1 &\in \mathbf{Q}_r^{X_i}, w_i \in \mathbf{Q}_r^{X_i}, i \in \mathcal{I}, \\
v_1 &\in \mathbf{Q}_r^Z, v_i, w_i \in \mathbb{R}_{2r}[x], i \in \mathcal{I}.
\end{aligned}$$

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. In practice, we use off-the-shelf numerical solvers to solve the dual SDP (5) and simultaneously get information about the primal SDP (4). Usually r has to be less than 10 due to limitations on current SDP solvers, but such r already provides good controllers for many practical problems. For fixed r or for r less than a fixed constant, the computational complexity is polynomial in the state and control input dimensions.

C. Controller extraction

We extract a polynomial controller for each mode. Combining these controllers, we have a piecewise polynomial controller on the entire state space. The procedure of extracting a controller for mode i from the moments of the measure ν_i is the same as in Section IV.C. of [48].

However, the extracted controller may 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 both examples in the next section, we used this method, and we noticed that most of the time, the control input constraints were not violated. Another method is to solve an SOS optimization problem as in [42].

V. EXAMPLES

We illustrate our controller synthesis method on some discrete-time hybrid polynomial systems. All computations

are done using MATLAB 2016b, the SDP solver MOSEK 8, and the polynomial optimization toolbox Spotless [58].

A. Linear inverted pendulum with a wall

Consider the problem of balancing a linear inverted pendulum to its upright position with the existence of a nearby elastic wall, as depicted in Figure 1(A) (Example A in [22]). The system can be modelled as a hybrid system with two modes: pendulum not in contact with the wall (mode 1), and pendulum in contact with the wall (mode 2). In this example, we linearize the dynamics and work with the PWA system as in [22]. (We could have approximated the system with higher order polynomials, synthesized the controller, and run on the real system as in Example E in [48], but we are more interested in the comparison with the traditional model predictive control (MPC) approach.)

Let $m = 1, l = 1, d = 0.1, g = 10$, and $k = 1000$. Let the state be $x = (\theta, \dot{\theta})$. Linearizing the dynamics around the upright position $\theta = 0$, and discretizing the model with the explicit Euler scheme with the sampling time $\delta t = 0.01$, the system becomes a piecewise-affine system $x^+ = A_i x + B_i u + a_i, x \in X_i, i \in \mathcal{I} = \{1, 2\}$, where the dynamics and the state space are the same as given in Example A of [22]. The target set is $Z = \{x \in \mathbb{R}^2 \mid |x_1| \leq 0.03, |x_2| \leq 0.1\}$, which is inside the maximum LQR-control invariant set.

We search for a piecewise-affine controller, and the extracted controller is

$$\begin{aligned}
u^1(x) &= 0.10336 - 6.7202x_1 - 1.6978x_2, \text{ for } x \in X_1, \\
u^2(x) &= -0.62962 + 5.4774x_1 - 0.60315x_2, \text{ for } x \in X_2.
\end{aligned}$$

The first two plots in Figure 2 show trajectories of five initial states $(0.08, 0.2 + i0.1)$ for $i = 0, 1, \dots, 4$, under the extracted controller and the uniformly sampled initial states with controllable states in red and uncontrollable states in blue. The mode switch sequences are not specified beforehand, and it turns out that small differences in the initial states can result in very different mode switch sequences as shown in the trajectory plot.

We compare our method with the standard MPC approach. With the MPC approach, we use a binary variable for each time step to indicate whether the system mode at that time step is mode 1 or mode 2. Therefore, the MPC approach amounts to solving online a mixed integer quadratic programming (MIQP) problem with the terminal set being the maximum LQR-control invariant set and the terminal function being the solution to the Riccati equation. Fixing a mode sequence over the time horizon, the MIQP problem becomes a QP problem, and the feasible set of the QP is a polytope. The union of the feasible sets of the QP problems over all possible mode sequences is the feasible set of the MIPQ problem. In this case (with the terminal set and the terminal cost as described above), the feasible set is the same as the controllable set. The feasible set of the MIQP problem with time horizon $T = 10$ is depicted in the right plot (taken from Figure 5 of [22]) of Figure 2. We choose the time horizon to be 10, because on an Intel i5, 2.3 GHz machine using Gurobi 8.0.0, the worst computation time for




Fig. 2. Linear inverted pendulum with a wall. Left: trajectories of five initial states under the occupation-measure-derived controller. Middle: controllable states under the occupation-measure-derived controller plotted in red and uncontrollable states plotted in blue. Right: feasible sets of the MPC approach.

solving an MIQP with the time horizon of 10 already exceeds the sampling time 0.01 s. Each polytope in the right plot represents a feasible set of a QP problem with a certain mode sequence. For example, the large blue region on the left represents the feasible set of the QP problem with the mode sequence being all 1's for the 10 time steps.

Our approach has three advantages. First, we have a simpler controller. Our controller is piecewise-affine with only two pieces, while the MPC controller is much more complicated. Second, we have a larger controllable set (noticing our controllable set includes the white region pointed by an arrow in the right plot of Figure 2). This is because the MPC approach, restricted by the sampling time, has a limited time horizon, while our approach computes solutions for all finite time horizons. Third, the online computation time of our approach is negligible, while the MPC approach needs to solve an MIQP problem online, which is > 0.01 s in the worst case for a time horizon of 10. It is amazing that our approach produces such a simple controller that controls such a large region.

B. Variable height inverted pendulum (legged robot)

Consider balancing a legged robot modelled as a variable height inverted pendulum [59], [60], as depicted in Figure 1(B). In this example, we assume there is one massless base (foot), and the robot can place it on two fixed places. The center of pressure b can range from anywhere in the base. The center of mass (CoM) q is constrained in a box area, independent of the distance to the center of base (CoB), the middle point of base. The force exerted by the ground is $F = m(q-b)u$, where u controls the magnitude of the force. The dynamics of the system is given by $m\ddot{q} = -mg + m(q-b)u$. Both b and u are control inputs, and b is hybrid with two possible steps to take on. The system is a hybrid polynomial system with 2 modes, 4 states, and 2 control inputs. The goal is to balance the CoM to the upright position, which is 1 m above the CoB. Assume the length of the base is 0.2 m. We search for two 3rd order piecewise polynomial controllers b and u . The computation time on an Intel i7 3.3 GHz, 32 GB RAM machine was about 10 minutes. Figure 3 shows the controllable points in the 2-dimensional $q_x - q_y$ plane sections of the 4-dimensional state space, where q_x and q_y are the horizontal and vertical displacement of the CoM from the desired position. Projections of three trajectories onto the $q_x - q_y$ plane are also plotted. In the first experiment (left), the base can be placed on the intervals $[-0.1, 0.1]$ and




Fig. 3. The $q_x - q_y$ sections of the state space, where other states are 0. Controllable points are plotted in green. The black box in each plot indicates the target set Z . The vertical solid line is the boundary of two modes. The vertical dashed line indicates a boundary of a base. q_y is the deviation from the desired height (1 m above the ground). Three curves in each plot are the projections of the trajectories of three initial states $x_0 = [q_x, q_y, \dot{q}_x, \dot{q}_y] = [0.2, 0.08, 0, 0]$, $[0.08, 0.05, 0, 0]$, and $[0.04, -0.08, 0, 0]$ onto the plane.

$[0.1, 0.3]$, i.e., for mode 1, $b \in [-0.1, 0.1]$, and for mode 2, $b \in [0.1, 0.3]$. In the second experiment (right), the base can be placed on the intervals $[-0.1, 0.1]$ and $[0.15, 0.35]$. In both experiments, $q_x = 0.1$ is the boundary of two modes, i.e., if $q_x < 0.1$, then $b \in [-0.1, 0.1]$, otherwise b is in the other interval ($[0.1, 0.3]$ or $[0.15, 0.35]$).

VI. CONCLUSION AND DISCUSSION

We have presented a controller synthesis method for discrete-time hybrid polynomial systems via the notion of occupation measures. We noticed that controllers of certain degrees work better for certain systems. For example, for PWA systems we found PWA controllers generally work better than controllers of higher degrees, while for higher degree systems, PWA controllers do not work well. There are some limitations. First, the controller synthesis process is heuristic, providing no stability guarantees for the closed-loop system. Controllable regions have to be computed a posteriori. Second, besides the degree of the controller, there is not much more room for parameter tuning. One possible way to tune parameters is to change the state space constraints that are not hard imposed. For example, increasing or decreasing the limit on the maximum velocity of a rigid body would result in different controllers. Third, if the truncated moments in the objective of the SDP (5) are large, the SDP solver might run into numerical issues, due to the immaturity of the current SDP solvers and the Spotless software. A possible solution is to rescale the state space constraints together with the system dynamics.

REFERENCES

- [1] H. Witsenhausen, "A class of hybrid-state continuous-time dynamic systems," *IEEE Transactions on Automatic Control*, vol. 11, no. 2, pp. 161–167, 1966.
- [2] M. S. Branicky and S. K. Mitter, "Algorithms for optimal hybrid control," in *Proceedings of 1995 34th IEEE Conference on Decision and Control*, vol. 3. IEEE, 1995, pp. 2661–2666.
- [3] M. S. Branicky, "Studies in hybrid systems: Modeling, analysis, and control," MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR INFORMATION AND DECISION SYSTEMS, Tech. Rep., 1995.
- [4] M. S. Branicky, V. S. Borkar, and S. K. Mitter, "A unified framework for hybrid control: Model and optimal control theory," *IEEE transactions on automatic control*, vol. 43, no. 1, pp. 31–45, 1998.
- [5] S. Hedlund and A. Rantzer, "Convex dynamic programming for hybrid systems," *IEEE Transactions on Automatic Control*, vol. 47, no. 9, pp. 1536–1540, 2002.
- [6] J. Lygeros, C. Tomlin, and S. Sastry, "Hybrid systems: modeling, analysis and control," *preprint*, 2008.
- [7] H. J. Sussmann, "Set-valued differentials and the hybrid maximum principle," in *Proceedings of the 39th IEEE Conference on Decision and Control (Cat. No. 00CH37187)*, vol. 1. IEEE, 2000, pp. 558–563.
- [8] B. Piccoli, "Necessary conditions for hybrid optimization," in *Proceedings of the 38th IEEE Conference on Decision and Control (Cat. No. 99CH36304)*, vol. 1. IEEE, 1999, pp. 410–415.
- [9] M. S. Shaikh and P. E. Caines, "On trajectory optimization for hybrid systems: Theory and algorithms for fixed schedules," in *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, vol. 2. IEEE, 2002, pp. 1997–1998.
- [10] —, "On the hybrid optimal control problem: theory and algorithms," *IEEE Transactions on Automatic Control*, vol. 52, no. 9, pp. 1587–1603, 2007.
- [11] P. J. Antsaklis and X. D. Koutsoukos, "Hybrid systems: Review and recent progress," *Software-Enabled Control: Information Technology for Dynamical Systems*, vol. 273, p. 298, 2003.
- [12] R. Goebel, R. G. Sanfelice, and A. R. Teel, "Hybrid dynamical systems," *IEEE Control Systems Magazine*, vol. 29, no. 2, pp. 28–93, 2009.
- [13] F. Zhu and P. J. Antsaklis, "Optimal control of hybrid switched systems: A brief survey," *Discrete Event Dynamic Systems*, vol. 25, no. 3, pp. 345–364, 2015.
- [14] J. Lunze and F. Lamnabhi-Lagarigue, *Handbook of hybrid systems control: theory, tools, applications*. Cambridge University Press, 2009.
- [15] P. Tabuada, *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.
- [16] B. Passenberg, P. E. Caines, M. Sobotka, O. Stursberg, and M. Buss, "The minimum principle for hybrid systems with partitioned state space and unspecified discrete state sequence," in *49th IEEE Conference on Decision and Control (CDC)*. IEEE, 2010, pp. 6666–6673.
- [17] B. Passenberg, M. Sobotka, O. Stursberg, M. Buss, and P. E. Caines, "An algorithm for discrete state sequence and trajectory optimization for hybrid systems with partitioned state space," in *49th IEEE Conference on Decision and Control (CDC)*. IEEE, 2010, pp. 4223–4229.
- [18] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 69–81, 2014.
- [19] M. S. Branicky, M. M. Curtiss, J. A. Levine, and S. B. Morgan, "Rrts for nonlinear, discrete, and hybrid planning and control," in *42nd IEEE International Conference on Decision and Control (IEEE Cat. No. 03CH37475)*, vol. 1. IEEE, 2003, pp. 657–663.
- [20] M. S. Branicky, M. M. Curtiss, J. Levine, and S. B. Morgan, "Sampling-based planning, control and verification of hybrid systems," *IEE Proceedings-Control Theory and Applications*, vol. 153, no. 5, pp. 575–590, 2006.
- [21] S. Sadraddini and R. Tedrake, "Sampling-based polytopic trees for approximate optimal control of piecewise affine systems," *arXiv preprint arXiv:1809.09716*, 2018.
- [22] T. Marcucci, R. Deits, M. Gabiccini, A. Biechi, and R. 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*. IEEE, 2017, pp. 31–38.
- [23] R. Deits, T. Koolen, and R. Tedrake, "Lvis: Learning from value function intervals for contact-aware robot controllers," *arXiv preprint arXiv:1809.05802*, 2018.
- [24] M. Baoti, F. J. Christophersen, and M. Morari, "Constrained optimal control of hybrid systems with a linear performance index," *IEEE Transactions on Automatic Control*, vol. 51, no. 12, pp. 1903–1919, 2006.
- [25] M. Baotic, F. J. Christophersen, and M. Morari, "Infinite time optimal control of hybrid systems with a linear performance index," in *42nd IEEE International Conference on Decision and Control (IEEE Cat. No. 03CH37475)*, vol. 3. IEEE, 2003, pp. 3191–3196.
- [26] M. Baotic, F. J. Christophersen, and M. Morari, "A new algorithm for constrained finite time optimal control of hybrid systems with a linear performance index," in *2003 European Control Conference (ECC)*. IEEE, 2003, pp. 3323–3328.
- [27] F. Borrelli, M. Baotic, A. Bemporad, and M. Morari, "An efficient algorithm for computing the state feedback optimal control law for discrete time hybrid systems," in *Proceedings of the 2003 American Control Conference, 2003.*, vol. 6. IEEE, 2003, pp. 4717–4722.
- [28] M. Barić, P. Grieder, M. Baotic, and M. Morari, "Optimal control of pwa systems by exploiting problem structure," *IFAC Proceedings Volumes*, vol. 38, no. 1, pp. 307–312, 2005.
- [29] F. J. Christophersen, M. Baotic, and M. Morari, "Optimal control of piecewise affine systems: A dynamic programming approach," in *Control and Observer Design for Nonlinear Finite and Infinite Dimensional Systems*. Springer, 2005, pp. 183–198.
- [30] F. Borrelli, M. Baotic, A. Bemporad, and M. Morari, "Dynamic programming for constrained optimal control of discrete-time linear hybrid systems," *Automatica*, vol. 41, no. 10, pp. 1709–1721, 2005.
- [31] M. Barić, P. Grieder, M. Baotic, and M. Morari, "An efficient algorithm for optimal control of pwa systems with polyhedral performance indices," *Automatica*, vol. 44, no. 1, pp. 296–301, 2008.
- [32] A. Bemporad, F. Borrelli, and M. Morari, "Optimal controllers for hybrid systems: Stability and piecewise linear explicit form," in *Proceedings of the 39th IEEE Conference on Decision and Control (Cat. No. 00CH37187)*, vol. 2. IEEE, 2000, pp. 1810–1815.
- [33] —, "Piecewise linear optimal controllers for hybrid systems," in *Proceedings of the 2000 American Control Conference. ACC (IEEE Cat. No. 00CH36334)*, vol. 2. IEEE, 2000, pp. 1190–1194.
- [34] —, "On the optimal control law for linear discrete time hybrid systems," in *International workshop on hybrid systems: computation and control*. Springer, 2002, pp. 105–119.
- [35] D. Mayne and S. Rakovic, "Optimal control of constrained piecewise affine discrete time systems using reverse transformation," in *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, vol. 2. IEEE, 2002, pp. 1546–1551.
- [36] L. Rodrigues, "Dynamic output feedback controller synthesis for piecewise-affine systems," Ph.D. dissertation, Stanford University, 2002.
- [37] M. Lazar, "Model predictive control of hybrid systems: Stability and robustness," Ph.D. dissertation, Eindhoven: Technische Universiteit Eindhoven, 2006.
- [38] W. Han and R. 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*. IEEE, 2017, pp. 842–849.
- [39] J. B. Lasserre, D. Henrion, C. Prieur, and E. Trélat, "Nonlinear optimal control via occupation measures and lmi-relaxations," *SIAM journal on control and optimization*, vol. 47, no. 4, pp. 1643–1666, 2008.
- [40] D. Henrion and M. Korda, "Convex computation of the region of attraction of polynomial control systems," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 297–312, 2014.
- [41] M. Korda, D. Henrion, and C. N. Jones, "Inner approximations of the region of attraction for polynomial dynamical systems," *IFAC Proceedings Volumes*, vol. 46, no. 23, pp. 534–539, 2013.
- [42] —, "Controller design and region of attraction estimation for nonlinear dynamical systems," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 2310–2316, 2014.
- [43] —, "Convex computation of the maximum controlled invariant set for polynomial control systems," *SIAM Journal on Control and Optimization*, vol. 52, no. 5, pp. 2944–2969, 2014.
- [44] V. Shia, R. Vasudevan, R. Bajcsy, and R. Tedrake, "Convex computation of the reachable set for controlled polynomial hybrid systems," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014, pp. 1499–1506.

- [45] V. Magron, P.-L. Garoche, D. Henrion, and X. Thirioux, “Semidefinite approximations of reachable sets for discrete-time polynomial systems,” *arXiv preprint arXiv:1703.05085*, 2017.
- [46] A. Majumdar, R. Vasudevan, M. M. Tobenkin, and R. Tedrake, “Convex optimization of nonlinear feedback controllers via occupation measures,” *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1209–1230, 2014.
- [47] C. Savorgnan, J. B. Lasserre, and M. 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*. IEEE, 2009, pp. 519–524.
- [48] W. Han and R. Tedrake, “Controller synthesis for discrete-time polynomial systems via occupation measures,” in *Intelligent Robots and Systems (IROS), 2018 IEEE/RSJ International Conference on*. IEEE, 2018.
- [49] P. Zhao, S. Mohan, and R. Vasudevan, “Control synthesis for nonlinear optimal control via convex relaxations,” in *American Control Conference (ACC), 2017*. IEEE, 2017, pp. 2654–2661.
- [50] R. Tedrake and the Drake Development Team, “Drake: A planning, control, and analysis toolbox for nonlinear dynamical systems,” 2016. [Online]. Available: <https://drake.mit.edu>
- [51] D. Stewart and J. C. Trinkle, “An implicit time-stepping scheme for rigid body dynamics with coulomb friction,” in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 1. IEEE, 2000, pp. 162–169.
- [52] N. Chakraborty, S. Berard, S. Akella, and J. C. Trinkle, “An implicit time-stepping method for multibody systems with intermittent contact,” in *Robotics: Science and Systems*, 2007.
- [53] H. L. Royden and P. Fitzpatrick, *Real analysis*. Macmillan New York, 1988, vol. 32.
- [54] G. B. Folland, *Real analysis: modern techniques and their applications*. John Wiley & Sons, 2013.
- [55] J. B. Conway, *A course in functional analysis*. Springer Science & Business Media, 2013, vol. 96.
- [56] M. Putinar, “Positive polynomials on compact semi-algebraic sets,” *Indiana University Mathematics Journal*, vol. 42, no. 3, pp. 969–984, 1993.
- [57] J.-B. Lasserre, *Moments, positive polynomials and their applications*. World Scientific, 2010, vol. 1.
- [58] M. M. Tobenkin, F. Permenter, and A. Megretski, “Spotless polynomial and conic optimization,” 2013.
- [59] T. Koolen, M. Posa, and R. Tedrake, “Balance control using center of mass height variation: limitations imposed by unilateral contact,” in *Humanoid Robots (Humanoids), 2016 IEEE-RAS 16th International Conference on*. IEEE, 2016, pp. 8–15.
- [60] M. Posa, T. Koolen, and R. Tedrake, “Balancing and step recovery capturability via sums-of-squares optimization,” in *Robotics: Science and Systems*, 2017.