Automatica 67 (2016) 54–66

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Controller design and value function approximation for nonlinear dynamical systems☆

Milan Korda [a], Didier Henrion [b,c,d], Colin N. Jones [a]

[a] Laboratoire d'Automatique, École Polytechnique Fédérale de Lausanne, Station 9, CH-1015, Lausanne, Switzerland
[b] CNRS, LAAS, 7 avenue du colonel Roche, F-31400 Toulouse, France
[c] Université de Toulouse, LAAS, F-31400 Toulouse, France
[d] Faculty of Electrical Engineering, Czech Technical University in Prague, Technická 2, CZ-16626 Prague, Czech Republic

ARTICLE INFO

Article history:
Received 23 March 2015
Received in revised form
25 November 2015
Accepted 15 December 2015
Available online 5 February 2016

Keywords:
Optimal control
Nonlinear control
Sum-of-squares
Semidefinite programming
Occupation measures
Value function approximation

ABSTRACT

This work considers the infinite-time discounted optimal control problem for continuous time input-affine polynomial dynamical systems subject to polynomial state and box input constraints. We propose a sequence of sum-of-squares (SOS) approximations of this problem obtained by first lifting the original problem into the space of measures with continuous densities and then restricting these densities to polynomials. These approximations are tightenings, rather than relaxations, of the original problem and provide a sequence of rational controllers with value functions associated to these controllers converging (under some technical assumptions) to the value function of the original problem. In addition, we describe a method to obtain polynomial approximations from above and from below to the value function of the extracted rational controllers, and a method to obtain approximations from below to the optimal value function of the original problem, thereby obtaining a sequence of asymptotically optimal rational controllers with explicit estimates of suboptimality. Numerical examples demonstrate the approach.

1. Introduction

This paper considers the infinite-time discounted optimal control problem for continuous-time input-affine polynomial dynamical systems subject to polynomial state constraints and box input constraints. This problem has a long history in both control and economics literature. Various methods to tackle this problem have been developed, often based on the analysis of the associated Hamilton–Jacobi–Bellman equation.

In this work we take a different approach: We first lift the problem into an infinite-dimensional space of measures with continuous densities where this problem becomes convex; in fact a linear program (LP). This lifting is a *tightening*, i.e., its optimal value is greater than or equal to the optimal value of the original problem, and under suitable technical conditions the two optimal values coincide. This infinite-dimensional LP is then further tightened by restricting the class of functions to polynomials of a prescribed degree and replacing nonnegativity constraints by sufficient sum-of-squares (SOS) constraints. This leads to a hierarchy of semidefinite programming (SDP) tightenings of the original problem indexed by the degree of the polynomials. The solutions to the SDPs yield immediately a sequence of *rational* controllers, and we prove that, under suitable technical assumptions, the value functions associated to these controllers converge *from above* to the value function of the original problem.

We also describe how to obtain a sequence of polynomial approximations converging from above and from below to the value function associated to each rational controller. Combined with existing techniques to obtain polynomial under approximations of the value function of the original problem (adapted to our setting), this method can be viewed as a design tool providing a sequence of rational controllers asymptotically optimal in the original problem with explicit estimates of suboptimality in each step.

The idea of lifting a nonlinear problem to an infinite-dimensional space dates back at least to the work of Young (1969) and subsequent works of Rubio (1985), Vinter and Lewis (1978), Warga (1972) and many others, both in deterministic and stochastic settings. These works typically lift the original problem into the space of measures and this lifting is a *relaxation* (i.e., its optimal value is less than or equal to the optimal value of the original problem) and under suitable conditions the two values coincide.

☆ The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Kok Lay Teo under the direction of Editor Ian R. Petersen.

E-mail addresses: milan.korda@epfl.ch (M. Korda), henrion@laas.fr (D. Henrion), colin.jones@epfl.ch (C.N. Jones).

http://dx.doi.org/10.1016/j.automatica.2016.01.022

More recently, this infinite-dimensional lifting was utilized numerically by *relaxing* the infinite-dimensional LP into a finite-dimensional SDP (Lasserre, Henrion, Prieur, & Trélat, 2008) or finite-dimensional LP (Gaitsgory & Quincampoix, 2009). Whereas the LP relaxations are obtained by classical state- and control-space gridding, the SDP relaxations are obtained by optimizing over truncated moment sequences (i.e., involving only finitely many moments) of the measures and imposing conditions *necessary* for these truncated moment sequences to be feasible in the infinite-dimensional lifted problem. These finite-dimensional relaxations provide *lower bounds* on the value function of the optimal control problem and seem to be difficult to use for control design with strong convergence guarantees; a controller extraction from the relaxations is possible although no convergence (e.g., Gaitsgory & Quincampoix, 2009; Henrion, Lasserre, & Savorgnan, 2008) or only very weak convergence can be established (e.g., Korda, Henrion, & Jones, 2014b; Majumdar, Vasudevan, Tobenkin, & Tedrake, 2014 in the related context of region of attraction approximation).

Contrary to these works, in this approach we tighten the infinite-dimensional LP by optimizing over polynomial densities of the measures and imposing conditions *sufficient* for these densities to be feasible in the infinite-dimensional lifted problem, thereby obtaining upper bounds as opposed to lower bounds. Crucially, to ensure that polynomial densities of arbitrarily low degrees exist for our problem (and therefore the resulting SDP tightenings are feasible), we work with free initial and final measures and set up the cost function and constraints such that this additional freedom does not affect optimality. Importantly, we do *not* assume that the state constraint set is control invariant, a requirement that is often imposed in the existing literature (e.g., Rantzer & Hedlund, 2003) but rarely met in practice.

The presented approach bears some similarity with the density approach of Prajna, Parrilo, and Rantzer (2004) for global stabilization later extended to optimal control (in a purely theoretical setting) in Rantzer and Hedlund (2003) and recently generalized to optimal stabilization of a given invariant set in Raghunathan and Vaidya (2014) (providing both theoretical results and a practical computation method). However, contrary to Prajna et al. (2004) we consider the problem of optimal control, not stabilization and moreover we work under state constraints. Contrary to Raghunathan and Vaidya (2014) we work in continuous time, consider a more general problem (optimal control, not optimal stabilization of a given set) and our approach of finite-dimensional approximation is completely different in the sense that it is based purely on convex optimization and it does not rely on state-space discretization. Moreover, and importantly, our approach comes with convergence guarantees.

Finally, let us mention that this work is inspired by Lasserre (2011), where a converging sequence of upper bounds on *static* polynomial optimization problems was proposed, as opposed to a converging sequence of lower bounds as originally developed in Lasserre (2001).

2. Preliminaries

2.1. Notation

We use $L(X;Y)$ to denote the space of all Lebesgue measurable functions defined on a set $X \subset \mathbb{R}^n$ and taking values in the set $Y \subset \mathbb{R}^m$. If the space Y is not specified it is understood to be $\mathbb{R}$. The spaces of integrable functions and essentially bounded functions are denoted by $L^1(X;Y)$ and $L^\infty(X;Y)$, respectively. The spaces of continuous respectively k-times continuously differentiable functions are denoted by $C(X;Y)$ respectively $C^k(X;Y)$. By a (Borel) measure we understand a countably additive mapping from (Borel) sets to nonnegative real numbers. Integration of a continuous function v with respect to a measure μ on a set X is denoted by $\int_X v(x)\,d\mu(x)$ or also $\int v\,d\mu$ when the variable and domain of integration are clear from the context. A probability measure is a measure with unit mass (i.e., $\int 1 d\mu = 1$). The support of a measure μ, defined as the smallest closed set whose complement has zero measure, is denoted by $\operatorname{spt}\mu$. The ring of all multivariate polynomials in a variable x is denoted by $\mathbb{R}[x]$, the vector space of all polynomials of degree no more than d is denoted by $\mathbb{R}[x]_d$, and the vector space of m-dimensional polynomial vectors is denoted by $\mathbb{R}[x]^m$. The boundary of a set X is denoted by ∂X, the interior by X° and the closure by $\bar{X}$. The Euclidean distance of a point x from a set X is denoted by $\mathrm{dist}_X(x)$. For a possibly matrix-valued function $f \in C(X;\mathbb{R}^{n\times m})$ we define $\|f\|_{C^0(X)} := \sup_{x\in X}\max_{i,j}|f_{i,j}(x)|$ and for a vector-valued function $g \in C^1(X;\mathbb{R}^n)$ we define $\|g\|_{C^1(X)} := \|g\|_{C^0(X)} + \|\frac{\partial g}{\partial x}\|_{C^0(X)}$, where $\frac{\partial g}{\partial x}$ denotes the Jacobian of g. If clear from the context we write $\|\cdot\|_{C^0}$ for $\|\cdot\|_{C^0(X)}$ and similarly for the C^1 norm. The set of consecutive integers $i, i+1, \ldots, j$ is denoted by $\mathbb{Z}_{i,j}$.

2.2. SOS programming

Crucial to the material presented in the paper is the ability to decide whether a polynomial $p \in \mathbb{R}[x]$ is nonnegative on a set

$$X = \{x \in \mathbb{R}^n \mid g_i(x) \geq 0,\ i = 1,\ldots,n_g\},$$

with $g_i \in \mathbb{R}[x]$. A *sufficient* condition for p to be nonnegative on X is that it belongs to the truncated quadratic module of degree d associated to X,

$$Q_d(X) := \left\{ s_0 + \sum_{i=1}^{n_g} g_i(x)s_i(x) \;\middle|\; s_0 \in \Sigma_{2\lfloor \frac{d}{2}\rfloor}, s_i \in \Sigma_{2\left\lfloor \frac{(d-\deg g_i)}{2}\right\rfloor} \right\},$$

where Σ_{2k} is the set of all polynomial sum-of-squares (SOS) of degree at most $2k$. Note in particular that $Q_{d+1}(X) \supset Q_d(X)$. If $p \in Q_d(X)$ for some $d \geq 0$ then clearly p is nonnegative on X, and the following fundamental result shows that a certain converse result holds.

Proposition 1 (Putinar, 1993). *Let $N - \|x\|^2 \in Q_d(X)$ for some $d > 0$ and $N \geq 0$ and let $p \in \mathbb{R}[x]$ be strictly positive on X. Then $p \in Q_d(X)$ for some $d \geq 0$.*

Combining with the Stone–Weierstrass Theorem, as an immediate corollary we get:

Corollary 1. *Let $f \in C(X)$ be nonnegative on X and let $N - \|x\|^2 \in Q_d(X)$ for some $d > 0$ and $N \geq 0$. Then for every $\epsilon \geq 0$ there exists $d \geq 0$ and $p_d \in Q_d(X)$ such that $\|f - p_d\|_{C^0} < \epsilon$.*

Corollary 1 says that polynomials in $\bigcup_{d=0}^{\infty} Q_d(X)$ are dense (with respect to the C^0 norm) in the space of continuous functions nonnegative on X.

In the rest of the text we use standard algebraic operations on sets. For instance if we write that $p \in gQ_d(X) + h\mathbb{R}[x]_d$, then it means that $p = gq + hr$ with $q \in Q_d(X)$ and $r \in \mathbb{R}[x]_d$.

The inclusion of $p \in Q_d(X)$ for a fixed d is equivalent to the existence of a positive semidefinite matrix W such that $p(x) = b(x)^\top W b(x)$, where $b(x)$ is a basis of $\mathbb{R}[x]_{d/2}$, the vector space of polynomials of degree at most $d/2$. Comparing coefficients leads to a set of affine constraints on the coefficients of p and the entries of W. Deciding whether $p \in Q_d(X)$ therefore translates to the feasibility of a semidefinite programming problem with the coefficients of p entering affinely. As a result, optimization of a linear function of the coefficients of p subject to the constraint $p \in Q_d(X)$ translates to a semidefinite programming problem (SDP) and hence to a well-understood and widely studied class of convex optimization problems for which powerful algorithms and off-the-shelf software are available. See, e.g., Lasserre (2009) and the references therein for more details.