Systems & Control Letters 100 (2017) 1–5

Contents lists available at ScienceDirect

Systems & Control Letters

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

Convergence rates of moment-sum-of-squares hierarchies for optimal control problems

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

[a] Department of Mechanical Engineering, University of California, Santa Barbara, CA 93106-5070, United States
[b] CNRS, LAAS, 7 avenue du colonel Roche, F-31400 Toulouse, France
[c] Université de Toulouse, F-31400 Toulouse, France
[d] Faculty of Electrical Engineering, Czech Technical University in Prague, Technická 2, CZ-16626 Prague, Czech Republic
[e] Laboratoire d'Automatique, École Polytechnique Fédérale de Lausanne, Station 9, CH-1015, Lausanne, Switzerland

ARTICLE INFO

Article history:
Received 7 September 2016
Received in revised form 15 November 2016
Accepted 26 November 2016

Keywords:
Optimal control
Moment relaxations
Polynomial sums of squares
Convergence rate
Semidefinite programming
Approximation theory

ABSTRACT

We study the convergence rate of the moment-sum-of-squares hierarchy of semidefinite programs for optimal control problems with polynomial data. It is known that this hierarchy generates polynomial under-approximations to the value function of the optimal control problem and that these under-approximations converge in the L^1 norm to the value function as their degree d tends to infinity. We show that the rate of this convergence is $O(1/\log\log d)$. We treat in detail the continuous-time infinite-horizon discounted problem and describe in brief how the same rate can be obtained for the finite-horizon continuous-time problem and for the discrete-time counterparts of both problems.

1. Introduction

The moment-sum-of-squares hierarchy (also known as the Lasserre hierarchy) of semidefinite programs was originally introduced in [1] in the context of polynomial optimization. It allows one to solve globally non-convex optimization problems at the price of solving a sequence, or hierarchy, of convex semidefinite programming problems, with convergence guarantees; see e.g. [2] for an introductory survey, [3] for a comprehensive overview and [4] for control applications.

This hierarchy was extended in [5] to polynomial optimal control, and later on in [6] to global approximations of semi-algebraic sets, originally motivated by volume and integral estimation problems. The approximation hierarchy for semi-algebraic sets derived in [6] was then transposed and adapted to an approximation hierarchy for transcendental sets relevant for systems control [7], such as regions of attraction [8] and maximal invariant sets for controlled polynomial differential and difference equations [9], still with rigorous analytic convergence guarantees.

Central to the moment-sum-of-squares hierarchies of [5,8,9] are polynomial subsolutions of the Hamilton–Jacobi–Bellman equation, providing certified lower bounds, or under-approximations, of the value function of the optimal control problem. It was first shown in [5] that the hierarchy of polynomial subsolutions of increasing degree converges locally (i.e. pointwise) to the value function on its domain. Later on, as an outcome of the results of [8], global convergence (i.e. in L^1 norm on compact domains, or equivalently, almost uniformly) was established in [10].

The current paper is motivated by the analysis of the rate of convergence of the moment-sum-of-squares hierarchy for static polynomial optimization achieved in [11]; see also [12] and references therein for latest developments. We show that a similar analysis can be carried out in the dynamic case, i.e. for assessing the rate of convergence of the moment-sum-of-squares hierarchy for polynomial optimal control. For ease of exposition, we focus on the discounted infinite-horizon continuous-time optimal control problem and briefly describe (in Section 5) how the same convergence rate can be obtained for the finite-time continuous version of the problem and for the discrete counterparts of both problems.

Our main Theorem 4 gives estimates on the rate of convergence of the polynomial under-approximations to the value function in the L^1 norm. As a direct outcome of this result, we derive in Corollary 2 that the rate of convergence is in $O(1/\log\log d)$, where d is the degree of the polynomial approximation. As far as we know, this is the first estimate of this kind in the context of moment-sum-of-squares hierarchies for polynomial optimal control.

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

http://dx.doi.org/10.1016/j.sysconle.2016.11.010
0167-6911/

1.1. Notation

The set of all continuous functions on a set $X \subset \mathbb{R}^n$ is denoted by $C(X)$; the set of all k-times continuously differentiable functions is denoted by $C^k(X)$. For $h \in C(X)$, we denote $\|h\|_{C^0(X)} := \max_{x\in X}|h(x)|$ and for $h \in C^1(X)$ we denote $\|h\|_{C^1(X)} := \max_{x\in X} |h(x)| + \max_{x\in X}\|\nabla h(x)\|_2$ where ∇h is the gradient of h. The L^1 norm with respect to a measure μ_0 of a measurable function $h : \mathbb{R}^n \to \mathbb{R}$ is denoted by $\|h\|_{L_1(\mu_0)} := \int_{\mathbb{R}^n} h(x)\mu_0(dx)$. The set of all multivariate polynomials in a variable x of total degree no more than d is denoted by $\mathbb{R}[x]_d$. The symbol $\mathbb{R}[x]_d^n$ denotes the n-fold cartesian product of this set, i.e., the set of all vectors with n entries, where each entry is a polynomial from $\mathbb{R}[x]_d$. The interior of a set $X \subset \mathbb{R}^n$ is denoted by $\operatorname{int} X$.

2. Problem setup

Consider the discounted infinite-horizon optimal control problem

$$\begin{aligned} V^\star(x_0) \quad := \quad & \inf_{u(\cdot),\, x(\cdot)} \quad \int_0^\infty e^{-\beta t} l(x(t), u(t))\, dt \\ & \text{s.t.} \quad x(t) = x_0 + \int_0^t f(x(s), u(s))\, ds \\ & \qquad\; x(t) \in X,\; u(t) \in U \qquad \forall t \geq 0 \end{aligned} \tag{1}$$

where $\beta > 0$ is a given discount factor, $f \in \mathbb{R}[x, u]_{d_f}^n$ and $l \in \mathbb{R}[x, u]_{d_l}$ are given multivariate polynomials and the state and input constraint sets X and U are of the form

$$X = \{x \in \mathbb{R}^n : g_i^X(x) \geq 0,\, i = 1, \ldots, n_X\},$$
$$U = \{u \in \mathbb{R}^m : g_i^U(u) \geq 0,\, i = 1, \ldots, n_U\},$$

where $g_i^X \in \mathbb{R}[x]_{d_i^X}$ and $g_i^U \in \mathbb{R}[u]_{d_i^U}$ are multivariate polynomials. The function V^* in (1) is called the *value function* of the optimal control problem (1).

Let us recall the Hamilton–Jacobi–Bellman inequality

$$l(x, u) - \beta V(x, u) + \nabla V(x, u) \cdot f(x, u) \geq 0 \quad \forall (x, u) \in X \times U \tag{2}$$

which plays a crucial role in the derivation of the convergence rates. In particular, for any function $V \in C^1(X)$ that satisfies (2) it holds

$$V(x) \leq V^\star(x) \quad \forall x \in X. \tag{3}$$

The following polynomial sum-of-squares optimization problem provides a sequence of lower bounds to the value function indexed by the degree d:

$$\begin{aligned} \max_{V\in\mathbb{R}[x]_d} \quad & \int_X V(x)\, d\mu_0(x) \\ \text{s.t.} \quad & l - \beta V + \nabla V \cdot f \in Q_{d+d_f}(X \times U), \end{aligned} \tag{4}$$

where μ_0 is a given probability measure supported on X (e.g., the uniform distribution), and

$$Q_{d+d_f}(X \times U) := \Big\{ s_0 + \sum_{i=1}^{n_X} g_i^X s_X^i + \sum_{i=1}^{n_U} g_i^U s_U^i \;: \; s_0 \in \Sigma_{\lfloor (d+d_f)/2 \rfloor},\, s_X^i \in \Sigma_{\lfloor (d+d_f-d_X^i)/2 \rfloor},\, s_U^i \in \Sigma_{\lfloor (d+d_f-d_U^i)/2 \rfloor} \Big\},$$

is the truncated quadratic module associated with the sets X and U (see [2] or [3]), where Σ_d is the cone of sums of squares of polynomials of degree up to d. Note that whenever V is feasible in (4), then V satisfies Bellman's inequality (2), because polynomials in $Q_{d+d_f}(X \times U)$ are non-negative on $X \times U$ by construction. Therefore any polynomial V feasible in (4) satisfies also (3) and hence is an under-approximation of $V^\star$ on X.

The truncated quadratic module is essential to the proof of convergence of the moment-sum-of-squares hierarchy in the static polynomial optimization case [1] which is based on Putinar's Positivstellensatz [13]. We recall that some polynomials of degree $d + d_f$ non-negative on $X \times U$ may not belong to $Q_{d+d_f}(X \times U)$ [3]. On the other hand, optimizing over the polynomials belonging to $Q_{d+d_f}(X \times U)$ is "simple" (it translates to semidefinite programming) while optimizing over the cone of non-negative polynomials is very difficult in general. In particular, the optimization problem (4) translates to a finite-dimensional semidefinite programming problem (SDP). The fact that the truncated quadratic module has an explicit SDP representation and hence can be tractably optimized over is one of the main reasons for the popularity of the moment-sum-of-squares hierarchies across many fields of science.

Throughout the paper we impose the following standing assumptions.

Assumption 1. The following conditions hold:

(a) $X \subset [-1, 1]^n$ and $U \subset [-1, 1]^m$.
(b) The sets of polynomials $(g_i^X)_{i=1}^{n_X}$ and $(g_i^U)_{i=1}^{n_U}$ both satisfy the Archimedian condition.[1]
(c) $0 \in \operatorname{int} X$ and $0 \in \operatorname{int} U$.
(d) The function $\nabla V^\star$ is Lipschitz continuous on X.
(e) The set $f(x, U)$ is convex for all $x \in X$ and the function $v \mapsto \inf_{u\in U}\{l(x, u) : v = f(x, u)\}$ is convex for all $x \in X$.

The Assumption (a) and (b) are made without loss of generality since the sets X and U are assumed to be compact and hence they can be scaled such that they are included in the unit ball; adding redundant ball constraints $1 - \|x\|^2$ and $1 - \|u\|^2$ in the description of X and U then implies the Archimedian condition. Assumption (c) essentially requires that the sets X and U have nonempty interiors (a mild assumption) since then a change of coordinates can always be carried out such that the origin is in the interior of these sets. Assumption (d) is an important regularity assumption necessary for the subsequent developments. Assumption (e) is a standard assumption ensuring that the value function of the so-called relaxed formulation of the problem (4) coincides with $V^\star$ (see, e.g., [14]) and is satisfied, e.g., for input-affine[2] systems with input-affine cost function provided that U is convex. This class of problems is by far the largest and practically most relevant for which this assumption holds although other problems exist that satisfy this assumption as well.[3]

Under Assumption 1, the hierarchy of lower bounds generated by problem (4) converges from below in the L^1 norm to the value function $V^\star$; see e.g. [10]:

Theorem 1. *There exists $d_0 \geq 0$ such that the problem (4) is feasible for all $d \geq d_0$. In addition $V \leq V^\star$ for any V feasible in (4) and $\lim_{d\to\infty}\|V^\star - V_d^\star\|_{L_1(\mu_0)} = 0$, where $V_d^\star$ is an optimal solution to (4).*

The goal of this paper is to derive bounds on the convergence rate of $V_d^\star$ to $V^\star$.

3. Convergence rate

The convergence rate is a consequence of the following fundamental results from approximation theory and polynomial optimization.

[1] A sufficient condition for a set of polynomials $(g_i)_{i=1}^n$ to satisfy the Archimedian condition is $g_i = N - \|x\|_2^2$ for some i and some $N \geq 0$, which is a non-restrictive condition provided that the set defined by the g_i is compact and an estimate of its diameter is known. For a precise definition of this condition see Section 3.6.2 of [2].

[2] A system is input-affine if $f(x, u) = f_x(x) + f_u(x)u$ for some functions f_x and f_u.

[3] For example, consider $l(x, u) = x^2, f(x, u) = x + u^2, U = [-1, 1]$.