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

Milan Korda[1], Didier Henrion[2,3,4], Colin N. Jones[1]

Draft of September 8, 2016

Abstract

We study the convergence rate of moment-sum-of-squares hierarchies of semidefinite programs for optimal control problems with polynomial data. It is known that these hierarchies generate 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.

Keywords: optimal control, moment relaxations, polynomial sums of squares, semidefinite programming, approximation theory.

1 Introduction

The moment-sum-of-squares hierarchy (also know as Lasserre hierarchy) of semidefinite programs was originally introduced in [10] 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. [13] for an introductory survey, [11] for a comprehensive overview and [3] for control applications.

This hierarchy was extended in [12] 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

---

[1]Laboratoire d'Automatique, École Polytechnique Fédérale de Lausanne, Station 9, CH-1015, Lausanne, Switzerland. {`milan.korda,colin.jones`}`@epfl.ch`

[2]CNRS; LAAS; 7 avenue du colonel Roche, F-31400 Toulouse; France. `henrion@laas.fr`

[3]Université de Toulouse; LAAS; F-31400 Toulouse; France.

[4]Faculty of Electrical Engineering, Czech Technical University in Prague, Technická 2, CZ-16626 Prague, Czech Republic.

then transposed and adapted to an approximation hierarchy for transcendental sets relevant for systems control [2], such as regions of attraction [7] and maximal invariant sets for controlled polynomial differential and difference equations [9], still with rigourous analytic convergence guarantees.

Central to the moment-sum-of-squares hierarchies of [12, 7, 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 [12] 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 [7], global convergence (i.e. in L^1 norm on compact domains, or equivalently, almost uniformly) was established in [8].

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 [14]. 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.

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 $\text{Int}(X)$.

2 Problem setup

Consider the discounted infinite-horizon optimal control problem

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

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

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

$$
U = \{u \in \mathbb{R}^m \mid g_i^U(u) \geq 0, i = 1, \ldots, n_U\},
$$

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

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

is the truncated quadratic module associated with the sets X and U (see [13] or [11]), 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 hieararchy in the static polynomial optimization case [10] which is based on Putinar's Positivstellensatz [15]. We recall that some polynomials of degree $d + d_f$ non-negative on $X \times U$ may not belong to $Q_{d+d_f}(X \times U)$ [11]. 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) $0 \in \mathrm{Int}(X)$ *and* $0 \in \mathrm{Int}(U)$.

(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* $v \mapsto \inf_{u \in U} \{l(x,u) \mid v = f(x,u)\}$ *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 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 Archimedianity 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., [18]) 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. [8]:

Theorem 1 *There exists a* $d_0 \geq 0$ *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 $V^\star$.

---

[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 provovided that the set defined by g_i's is compact and an estimate of its dimeter is known. For a precise definition of this condition see Section 3.6.2 of [13].

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

3   Convergence rate

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

Theorem 2 (Bagby et al. [1]) *If $h : X \to \mathbb{R}$ is a function such that $\nabla h \in C^1(X)$, then there exists a sequence of polynomials $(p_d)_{d=1}^{\infty}$ satisfying $\deg(p_d) \leq d$ such that $\|h - p_d\|_{C^1(X)} \leq c_1/d$ for some constant $c_1 \geq 0$ depending on h and X only.*

Now we turn to the second fundamental result. Given a polynomial $p \in \mathbb{R}[x]_d$ expressed in a multivariate monomial basis as

$$p(x) = \sum_{\substack{\alpha \in \mathbb{N}^n \\ |\alpha| \leq d}} \beta_\alpha x^\alpha$$

with $|\alpha| = \sum_{i=1}^{n} \alpha_i$ and $x^\alpha = \prod_{i=1}^{n} x^{\alpha_i}$, we define

$$\|p\|_{\mathbb{R}[x]} = \max_\alpha \frac{|\beta_\alpha|}{\binom{|\alpha|}{\alpha}}, \tag{5}$$

where the multinomial coefficient $\binom{|\alpha|}{\alpha}$ is defined by

$$\binom{|\alpha|}{\alpha} := \frac{|\alpha|!}{\alpha_1! \cdot \ldots \cdot \alpha_n!}.$$

Theorem 3 (Nie & Schweighofer [14]) *Let $p \in \mathbb{R}[x, u]_{d_p}$ and let*

$$p_{\min} := \min_{(x,u) \in X \times U} p(x, u) \quad \text{with} \quad p_{\min} > 0.$$

Then $p \in Q_d(X \times U)$ provided that

$$d \geq c_2 \exp\left(d_p^2 (n+m)^{d_p} \frac{\|p\|_{\mathbb{R}[x,u]}}{p_{\min}} \right)^{c_2}, \tag{6}$$

where the constant c_2 depends only on the sets X and U.

In the following developments it will be crucial to bound the norm $\|\cdot\|_{\mathbb{R}[x]}$ of a polynomial by its supremum norm $\|\cdot\|_{C(X)}$. We remark that such a bound is possible only for a "generic" set X such that any polynomial vanishing on X necessarily vanishes everywhere. A sufficient condition for this is $\mathrm{Int}(X) \neq \emptyset$. This is the reason for Assumption 1 (c).

Lemma 1 *If $p \in \mathbb{R}[x]_d$, $x \in \mathbb{R}^n$, then*

$$\|p\|_{\mathbb{R}[x]} \leq 3^{d+1} \|p\|_{C([-1,1]^n)} \tag{7}$$

for all $d \geq 0$.

Proof: The idea is to use a multivariate Markov inequality to bound the derivatives of the polynomial at zero (and hence its coefficients) in terms of its supremum norm on $[-1,1]^n$.

Let $p = \sum_\alpha \beta_\alpha x^\alpha \in \mathbb{R}[x]_d$. From [17, Theorem 6], we have

$$\left|\frac{\partial^{|\alpha|} p}{\partial x^\alpha}(0)\right| \leq |T_d^{(|\alpha|)}(0) + i S_d^{(|\alpha|)}(0)| \cdot \|p\|_{C([-1,1]^n)}$$

for all multiindices α satisfying $|\alpha| \leq d$, where $i = \sqrt{-1}$, $T_d(y) = \cos(d \arccos(y))$, $y \in [-1,1]$, denotes the d-th univariate Chebyshev polynomial of the first kind, $S_d(y) = \sin(d \arccos(y)) = d^{-1}\sqrt{1-x^2}T_d'(y)$, $y \in [-1,1]$, and $h^{(k)}$ signifies the k-th derivative of a function $h : \mathbb{R} \to \mathbb{R}$. It is easy to see that $S_d^{(k)}(0) = d^{-1} T_d^{(k+1)}(0)$ and hence

$$|T_d^{(|\alpha|)}(0) + i S_d^{(|\alpha|)}(0)| \leq |T_d^{(|\alpha|)}(0)| + \frac{1}{d}|T_d^{(|\alpha|+1)}(0)| = |\alpha|! \cdot |t_{d,|\alpha|}| + \frac{(|\alpha|+1)!}{d} \cdot |t_{d,|\alpha|+1}|$$

$$\leq \left[|\alpha|! + \frac{(|\alpha|+1)!}{d}\right]\bar{t}_d,$$

where $t_{d,k}$ denotes the k-th coefficient of T_d when expressed in the monomial basis (i.e., $T_d(y) = \sum_{k=0}^d t_{d,k}y^k$) and $\bar{t}_d = \max_{k \in \{0,\dots,d\}} |t_{d,k}|$.

Since $\beta_\alpha = (\alpha_1! \cdot \ldots \cdot \alpha_n!)^{-1}\frac{\partial^{|\alpha|} p}{\partial x^\alpha}(0)$, we get

$$\frac{|\beta_\alpha|}{\binom{|\alpha|}{\alpha}} = \frac{(\alpha_1! \cdot \ldots \cdot \alpha_n!)|\beta_\alpha|}{|\alpha|!} = \frac{1}{|\alpha|!}\left|\frac{\partial^{|\alpha|} p}{\partial x^\alpha}(0)\right| \leq \left[1 + \frac{|\alpha|+1}{d}\right]\bar{t}_d\|p\|_{C([-1,1]^n)}.$$

In view of (5) and since $|\alpha| \leq d$ we get

$$\|p\|_{\mathbb{R}[x]} \leq \left[2 + \frac{1}{d}\right]\bar{t}_d\|p\|_{C([-1,1]^n)}.$$

It remains to bound $\bar{t}_d$. From the generating recurrence of $T_{d+1}(y) = 2yT_d(y) - T_{d-1}(y)$ starting from $T_0 = 1$ and $T_1 = y$, it follows that $\bar{t}_d \leq \tilde{t}_d$, where $\tilde{t}_d$ solves the linear difference equation

$$\tilde{t}_{d+1} = 2\tilde{t}_d + \tilde{t}_d$$

with the initial condition $\tilde{t}_0 = 1$ and $\tilde{t}_1 = 1$. The solution to this equation is

$$\tilde{t}_d = (1+\sqrt{2})^d\left(\frac{\sqrt{2}}{2} + \frac{1}{2}\right) + (1-\sqrt{2})^d\left(\frac{1}{2} - \frac{\sqrt{2}}{2}\right) \leq 3^d, \quad d \geq 1.$$

Therefore $\bar{t}_d \leq 3^d$ for $d \geq 1$ and hence

$$\|p\|_{\mathbb{R}[x]} \leq \left[2 + \frac{1}{d}\right]3^d\|p\|_{C([-1,1]^n)} \leq 3^{d+1}\|p\|_{C([-1,1]^n)}, \quad d \geq 1.$$

Since $\|p\|_{\mathbb{R}[x]} = \|p\|_{C([-1,1]^n)}$ for $d = 0$, the result follows. $\qquad\square$

In order to state an immediate corollary of this result, crucial for subsequent developments, we define

$$r := \frac{1}{\sup\{s > 0 \mid [-s,s]^{n+m} \subset X \times U\}}, \tag{8}$$

which is the reciprocal value of the length of the side of the largest box centered around the origin included in $X \times U$. By Assumption 1 (a) and (c), we have $r \in [1, \infty)$.

Corollary 1 *If $p \in \mathbb{R}[x,u]_d$, then*

$$\|p\|_{\mathbb{R}[x,u]} \leq k(d)\|p\|_{C(X \times U)}, \tag{9}$$

where $k(d) = 3^{d+1}r^d$ with r defined in (8).

Proof: Set $\tilde{p}((x,u)) := p(r^{-1}(x,u))$. Then we have

$$\|\tilde{p}\|_{C([-1,1]^{n+m})} = \|p\|_{C([-1/r,1/r]^{n+m})} \leq \|p\|_{C(X \times U)} \tag{10}$$

since $[-1/r, 1/r]^{n+m} \subset X \times U$ by definition of r (8). In addition

$$\|\tilde{p}\|_{\mathbb{R}[x,u]} = \max_{\alpha} r^{-|\alpha|} \frac{|\beta_\alpha|}{\binom{|\alpha|}{\alpha}} = r^{-d} \max_{\alpha} r^{d-|\alpha|} \frac{|\beta_\alpha|}{\binom{|\alpha|}{\alpha}} \geq r^{-d}\|p\|_{\mathbb{R}[x,u]}. \tag{11}$$

Combining (10), (11) and Lemma 1 we get

$$\|p\|_{\mathbb{R}[x,u]} \leq r^d\|\tilde{p}\|_{\mathbb{R}[x,u]} \leq 3^{d+1}r^d\|\tilde{p}\|_{C([-1,1]^{n+m})} \leq 3^{d+1}r^d\|p\|_{C(X \times U)} = k(d)\|p\|_{C(X \times U)}.$$

as desired. $\qquad\square$

Now we turn to analyzing the Bellman inequality (2). The following immediate property of this inequality will be of importance:

Lemma 2 *Let V satisfy (2) and let $a \in \mathbb{R}$. Then $\tilde{V} := V - a$ satisfies*

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

Proof: We have $l - \beta\tilde{V} + \nabla\tilde{V} \cdot f = l - \beta V + \nabla V \cdot f + \beta a \geq \beta a$, since V satisfies (2). $\quad\square$

We will also need a result which estimates the distance between the best polynomial approximation of a given degree to the value function and polynomials of the same degree satisfying Bellman's inequality. A similar result in discrete time and with discrete state and control spaces can be found in [5].

Lemma 3 *Let*

$$\hat{V}_d \in \underset{V \in \mathbb{R}[x]_d}{\arg\min} \|V - V^\star\|_{C^1(X)}.$$

Then there exists a polynomial $\tilde{V}_d \in \mathbb{R}[x]_d$ satisfying (2) and such that

$$\|\tilde{V}_d - V^\star\|_{C^1(X)} \leq \|\hat{V}_d - V^\star\|_{C^1(X)}\left(2 + \frac{\|f\|_{C^0(X)}}{\beta}\right). \tag{12}$$

Proof: Let $\tilde{V}_d := \hat{V}_d - a$. We will find an $a \geq 0$ such that $\tilde{V}_d$ satisfies the Bellman inequality. We have

$$
\begin{aligned}
l - \beta\tilde{V}_d + \nabla\tilde{V}_d f &= l - \beta\hat{V}_d + \nabla\hat{V}_d f + \beta a \\
&= l - \beta V^\star + \nabla V^\star f + \beta(V^\star - \hat{V}_d) + (\nabla\hat{V}_d - \nabla V^\star)f + \beta a \\
&\geq \beta(V^\star - \hat{V}_d) + (\nabla\hat{V}_d - \nabla V^\star)f + \beta a \\
&\geq -\beta\|\hat{V}_d - V^\star\|_{C^1(X)} - \|\hat{V}_d - V^\star\|_{C^1(X)}\|f\|_{C^0(X)} + \beta a,
\end{aligned}
$$

7

and hence if

$$a := \left(1 + \frac{\|f\|_{C^0(X)}}{\beta}\right)\|\hat{V}_d - V^\star\|_{C^1(X)},$$

then $\tilde{V}_d$ satisfies Bellman's inequality (2) and estimate (12) holds. $\qquad\square$

Now we are in position to prove our main result which bounds the gap, in L^1 norm, between the value function $V^\star$ of the optimal control problem (1) and any optimal solution $V_d^\star$ of the sum-of-squares program (4):

Theorem 4 *It holds that* $\|V^\star - V_d^\star\|_{L_1(\mu_0)} < \epsilon$ *for all integer*

$$d \geq c_2 \cdot \exp\left[\left(6d_p(\epsilon)^2(3r(n+m))^{d_p(\epsilon)}\frac{M + \beta\epsilon + \delta_1\|f\|_{C^0}}{\beta\epsilon}\right)^{c_2}\right] \tag{13}$$

$$= O\left(\exp\left[\frac{1}{\epsilon^{3c_2}}(3(n+m)r)^{\frac{c_3}{\epsilon}}\right]\right), \tag{14}$$

where $d_p = \left\lceil \frac{2c_1}{\epsilon}\left(2 + \frac{\|f\|_{C^0}}{\beta}\right) + d_f\right\rceil$, $M = \|l - \beta V^\star + \nabla V^\star \cdot f\|_{C^0(X \times U)} < \infty$, r *is defined in (8),* $c_3 = 2c_1c_2(2\beta + \|f\|_{C^0})/\beta$, *the constant* c_1 *depends only on* $V^\star$ *and* X *and* U, *whereas the constant* c_2 *depends only on sets* X *and* U.

Proof: According to Theorem 2 and Lemma 3 we can find a polynomial $\tilde{V}_{\tilde{d}}$ of degree no more than

$$\tilde{d} = \left\lceil \frac{2c_1}{\epsilon}\left(2 + \frac{\|f\|_{C^0}}{\beta}\right)\right\rceil$$

such that $\|V^\star - \tilde{V}_{\tilde{d}}\|_{C^1} \leq \frac{\epsilon}{2}$ and such that $\tilde{V}_{\tilde{d}}$ satisfies the Bellman inequality (2). Let V be an arbitrary polynomial feasible in (4) for some $d \geq 0$. Then

$$\|V^\star - V_d^\star\|_{L_1} \leq \|V^\star - V\|_{L_1} \leq \|V^\star - V\|_{C^0} \leq \|V^\star - \tilde{V}_{\tilde{d}}\|_{C^0} + \|V - \tilde{V}_{\tilde{d}}\|_{C^0} \leq \frac{\epsilon}{2} + \|V - \tilde{V}_{\tilde{d}}\|_{C^0}. \tag{15}$$

Hence, the goal is to find a degree $d \geq 0$ and a polynomial V feasible in (4) for that d satisfying $\|V - \tilde{V}_{\tilde{d}}\|_{C^0} \leq \epsilon/2$. Setting $V := \tilde{V}_{\tilde{d}} - \epsilon/2$, we clearly have $\|V - \tilde{V}_{\tilde{d}}\|_{C^0} \leq \epsilon/2$; in addition, using Lemma 2 we know that

$$l - \beta V + \nabla V \cdot f \geq \frac{1}{2}\beta\epsilon > 0 \tag{16}$$

and hence V strictly satisfies the Bellman inequality and as a consequence of the Putinar's Positivstellensatz [15] there exists a degree $d \geq 0$ such that V is feasible in (4). To bound the degree d we apply the bound of Theorem 3 on $p := l - \beta V + \nabla V \cdot f$. From (16) we know that $p_{\min} \geq \frac{1}{2}\beta\epsilon$. Next, we need to bound $\|p\|_{\mathbb{R}[x,u]}$ by bounding $\|p\|_{C^0(X \times U)}$ and using Corollary 1. We have

$$\|p\|_{C^0} = \|l - \beta V + \nabla V \cdot f\|_{C^0} = \|l - \beta\tilde{V}_{\tilde{d}} + \nabla\tilde{V}_{\tilde{d}} \cdot f + \frac{1}{2}\beta\epsilon\|_{C^0}$$

$$\leq \|l - \beta V^\star + \nabla V^\star \cdot f + \|_{C^0} + \beta\|V^\star - \tilde{V}_{\tilde{d}}\|_{C^0} + \|V^\star - \tilde{V}_{\tilde{d}}\|_{C^1}\|f\|_{C^0} + \frac{1}{2}\beta\epsilon$$

$$\leq M + \frac{1}{2}\beta\epsilon + \frac{1}{2}\epsilon\|f\|_{C^0} + \frac{1}{2}\beta\epsilon = M + \beta\epsilon + \frac{1}{2}\epsilon\|f\|_{C^0}.$$

Finally, we need to estimate the degree of p. We have (assuming without loss of generality that $\tilde{d} + d_f - 1 \geq \deg(l)$)

$$\deg(p) = \deg\big(l - \beta V + \nabla V \cdot f\big) \leq \tilde{d} + d_f - 1 \leq \left\lceil \frac{2c_1}{\epsilon}\Big(2 + \frac{\|f\|_{C^0}}{\beta}\Big) + d_f \right\rceil$$

Setting $d_p := \left\lceil \frac{2c_1}{\epsilon}\Big(2 + \frac{\|f\|_{C^0}}{\beta}\Big) + d_f \right\rceil$ and using Theorem 3 and Corollary 1, we conclude that for

$$d \geq c_2 \cdot \exp\Big[\Big(2d_p^2(n+m)^{d_p}\frac{k(d_p)(M + \beta\epsilon + \delta_1\|f\|_{C^0})}{\beta\epsilon}\Big)^{c_2}\Big],$$

the polynomial V is feasible in (4). Since $\|\tilde{V}_{\tilde{d}} - V\|_{C^0} \leq \frac{\epsilon}{2}$, we conclude from (15) that $\|V^\star - V_d^\star\|_{L_1} \leq \epsilon$. Inserting the expression for $k(d) = 3^{d+1}r^d$ from Corollary 1 yields (13) and carrying out asymptotic analysis for $\epsilon \to 0$ yields

$$d \geq O\Big(\exp\big[\frac{1}{\epsilon^{3c_2}}(3(n+m)r)^{\frac{c_3}{\epsilon}}\big]\Big),$$

which is (14).

$\square$

Corollary 2 *It holds $\|V^\star - V_d^\star\|_{L_1(\mu_0)} = O(1/\log\log d)$.*

Proof: Follows by inverting the asymptotic expression (14) using the fact that

$$(3(n+m)r)^{\frac{2c_3}{\epsilon}} \geq \frac{1}{\epsilon^{3c_2}}(3(n+m)r)^{\frac{c_3}{\epsilon}}$$

for small ϵ.

$\square$

4 Discussion

The bound on the convergence rate $O(1/\log\log d)$ should be compared with the bound $O(1/\sqrt[c_2]{\log d})$ derived in [14] for *static* polynomial optimization problems (here $c_2 \geq 1$ is the, in general, unknown constant from Theorem 3). The additional logarithm appearing in our bound seems to be unavoidable due to fundamental results of approximation theory (known as Bernstein inequalities) implying that Lipschitz continuous functions cannot be approximated by polynomials with rate faster than $1/d$ (in the sense that there exists a Lipschitz continuous function whose best degree-d approximation converges to f with the rate exactly C/d, $C > 0$, in the supremum norm on $[-1,1]^n$); this implies the $1/\epsilon$ dependence of d_p from Theorem 4 which then propagates to doubly exponential dependence on $1/\epsilon$ through Theorem 3.

Therefore the primary point of improvement of the bound from Theorem 4 and Corollary 2 is the fundamental bound of Theorem 3 derived in [14]. As the authors of [14] remark, this bound is far from tight, at least in two special cases: the univariate case (i.e., $n+m = 1$ in our setting) or the case of a single constraint defining $X \times U$. In these cases the exponential in (6) can be dropped, which results in $O(1/\log d)$ asymptotic rate of convergence in Corollary 2. In the general case, however, it is unknown whether the exponential in (6) can be dropped or whether the bound (6) can be improved otherwise [14].

5   Extensions

The approach for deriving this bound can be extended to other settings. In particular, similar bounds, with identical $O(1/\log\log d)$ asymptotics, hold for the finite-horizon version of the problem, both in continuous and discrete time, as well as for the discounted discrete-time infinite-horizon variant (the former was treated using the moment-sum-of-squares approach, in continuous time, in [12] and the latter was treated in [16]). The derivation in discrete-time is completely analogous and the results hold under milder assumptions (Assumption 1 (d) can be replaced by $V^\star$ Lipschitz and Assumption 1 (e) can be dropped completely). For the finite-horizon continuous-time problem, the only difference is in Lemma 3, where the constant shift $\tilde{V}(x) = \hat{V}(x) - a$, is replaced by the affine shift $\tilde{V}(t, x) = \hat{V}(t, x) - a - b(T - t)$ for suitable $a > 0$, $b > 0$ ensuring that $\tilde{V}$ satisfies the corresponding finite-time Bellman inequality and its boundary condition (hence the two degrees of freedom).

References

[1] T. Bagby, L. Bos, N. Levenberg. Multivariate simultaneous approximation. Constructive approximation, 18:569-577, 2002.

[2] G. Chesi. Domain of attraction; analysis and control via SOS programming. Lecture Notes in Control and Information Sciences, Vol. 415, Springer-Verlag, Berlin, 2011.

[3] G. Chesi. LMI techniques for optimization over polynomials in control: a survey. IEEE Transactions on Automatic Control, 55: 2500-2510, 2010.

[4] E. de Klerk, R. Hess, M. Laurent. Improved convergence rates for Lasserre-type hierarchies of upper bounds for box-constrained polynomial optimization. arXiv:1603.03329, March 2016.

[5] D. P. de Farias, B. Van Roy. The Linear Programming Approach to Approximate Dynamic Programming. Operations Research, 51:850-865.

[6] D. Henrion, J. B. Lasserre, C. Savorgnan. Approximate volume and integration for basic semialgebraic sets. SIAM Review 51(4):722-743, 2009.

[7] D. Henrion, M. Korda. Convex computation of the region of attraction of polynomial control systems. IEEE Transactions on Automatic Control, 59:297-312, 2014.

[8] D. Henrion, E. Pauwels. Linear conic optimization for nonlinear optimal control. arXiv:1407.1650, July 2014.

[9] M. Korda, D. Henrion, C. N. Jones. 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.

[10] J. B. Lasserre. Optimisation globale et théorie des moments. C. R. Acad. Sci. Paris. Série I, 331:929934, 2000.

[11] J. B. Lasserre. Moments, positive polynomials and their applications. Imperial College Press, London, UK, 2010.

[12] J. B. Lasserre, D. Henrion, C. Prieur, E. Trélat. Nonlinear optimal control via occupation measures and LMI relaxations. SIAM Journal on Control and Optimization, 47:1643-1666, 2008.

[13] M. Laurent. Sums of squares, moment matrices and polynomial optimization. In M. Putinar, S. Sullivan (eds.). Emerging applications of algebraic geometry, Vol. 149 of IMA Volumes in Mathematics and its Applications, Springer, Berlin, 2009.

[14] J. Nie, M. Schweighofer. On the complexity of Putinar's Positivstellensatz. Journal of Complexity 23:135-150, 2007.

[15] M. Putinar. Positive polynomials on compact semi-algebraic sets. Indiana University Mathematics Journal, 42:969-984, 1993.

[16] C. Savorgnan, J. B. Lasserre, M. Diehl. Discrete-time stochastic optimal control via occupation measures and moment relaxations. Proceedings of the 48th IEEE Conference on Decision and Control, 2009.

[17] V. I. Skalyga. Analogs of the Markov and Schaeffer-Duffin inequalities for convex bodies. Mathematical Notes 68:130-134, 2000.

[18] R. Vinter. Convex duality and nonlinear optimal control. SIAM Journal on Control and Optimization, 31:518:538, 1993.