On Infinite Linear Programming and the Moment Approach to Deterministic Infinite Horizon Discounted Optimal Control Problems

Journal Article

Author(s):

Kamoutsis, Angeliki; Sutter, Tobias; Esfahani, Peyman M.; Lygeros, John

Publication date:

2017-07

Permanent link:

<https://doi.org/10.3929/ethz-b-000232554>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

IEEE Control Systems Letters 1(1), <https://doi.org/10.1109/LCSYS.2017.2710234>

On Infinite Linear Programming and the Moment Approach to Deterministic Infinite Horizon Discounted Optimal Control Problems

Angeliki Kamoutsi, Tobias Sutter, Peyman Mohajerin Esfahani, and John Lygeros

Abstract—We revisit the linear programming approach to deterministic, continuous time, infinite horizon discounted optimal control problems. In the first part, we relax the original problem to an infinite-dimensional linear program over a measure space and prove equivalence of the two formulations under mild assumptions, significantly weaker than those found in the literature until now. The proof is based on duality theory and mollification techniques for constructing approximate smooth subsolutions to the associated Hamilton-Jacobi-Bellman equation. In the second part, we assume polynomial data and use Lasserre's hierarchy of primal-dual moment-sum-of-squares semidefinite relaxations to approximate the value function and design an approximate optimal feedback controller. We conclude with an illustrative example.

Index Terms—optimal control, discounted occupation measures, moments, sum-of-squares, infinite linear programming

I. INTRODUCTION

WE study the infinite horizon optimal control problem (OCP) with discounted payoff subject to state constraints. A comprehensive theoretical framework has been developed over the years to tackle such OCPs, e.g., via the Pontryagin's maximum principle and the associated Hamilton-Jacobi-Bellman equation (HJB) [1, Sections III.2, IV.5]. However, these sophisticated and mathematically elegant methods are not always easy to implement in practice. Several approaches, including shooting methods [2], model predictive control [3], direct methods [4] and neural-network-based algorithms [5], have been proposed in the literature to alleviate this difficulty.

In this work, we study an approximate dynamic programming approach based on infinite-dimensional linear programming (LP). The LP formulation of deterministic and stochastic optimal control problems has been studied extensively in the literature [6], [7], [8], [9], [10]. The main idea is to embed the set of admissible controls for the original problem in a new space, a space of measures. In our particular setting, we identify each admissible control with a discounted occupation measure and observe that it satisfies a linear equation that provides a measure-theoretic interpretation of the system

dynamics. In this way, we introduce an infinite-dimensional LP (primal) over an appropriate space of measures. The optimal value of the primal LP is no greater than the optimal value of the OCP and, under mild assumptions described in Theorem III.1, the two are equal. Using duality theory, we introduce the dual LP, which involves finding the supremum of all smooth subsolutions to the associated Hamilton-Jacobi-Bellman (HJB) equation, and prove that there is no duality gap. As a result, we derive a characterization of the value function as the upper envelope of the smooth subsolutions to the HJB equation. It is worth mentioning that the LP approach is particularly appealing for dealing with unconventional problems involving additional constraints or secondary costs, where traditional dynamic programming techniques are not applicable. In particular, in the OCP of our interest, the state and input constraints are directly incorporated into the measure space associated with the primal program and the function space associated with the dual program.

We then, use finite-dimensional approximations of the primal LP and its dual to approximate the optimal value function and extract a near optimal feedback control. In the literature, several works propose finite LP approximations based on state-and-control-space gridding [7], [11]. On the other hand, more recently, Lasserre et al. in [12], [13], used convex optimization methods and proposed a hierarchy of finite-dimensional semidefinite programming (SDP) relaxations to tackle nonlinear finite horizon OCPs with polynomial data. We extend the latter approach to the infinite horizon case.

Contribution. This paper is inspired and extends the techniques discussed in [12], [13] and [8]. We show equivalence of the LP formulation to the OCP (Theorem III.1) under weaker assumptions than those known in the literature [7]; in particular, our assumptions do not require Lipschitz continuity of $V_C^{Y_\delta}$ with respect to δ [7, Theorem 4.4], a technical assumption that usually cannot be checked even for simple OCPs. Moreover Assumption 1 in [7], needs to hold only for the cost function of the OCP and not for all continuous functions. Our proof of equivalence is based on duality theory and mollification techniques for constructing approximate smooth subsolutions to the HJB equation, also proposed in [8] for finite horizon differential inclusions and in [10] for controlled diffusions. Our finite-dimensional optimization approximations are based on Lasserre's hierarchy of semidefinite relaxations [12] instead of the grid-based LP approximations proposed in [7]. However, unlike [12], we treat the infinite horizon polynomial OCP

Research was supported by the European Union 7th Framework Program "Scalable Proactive Event-Driven Decision-making (SPEEDD)".

The authors are with the Automatic Control Laboratory, ETH Zürich, Switzerland and the Delft Center for Systems and Control, TU Delft, Netherlands; Emails: {kamoutsa, sutter, lygeros}@control.ee.ethz.ch, P.MohajerinEsfahani@tudelft.nl

by introducing the notion of discounted occupation measures. Finally, recently Korda et al. [14] have presented a new approach, where the primal LP is tightened and optimization is carried out over polynomial densities of measures. In this way, they are able to provide a controller design with strong convergence guarantees. However, their initial assumptions restrict significantly the class of OCPs for which their method is applicable. Although in our controller extraction, which follows closely the reasoning in [13], we do not provide convergence guarantees, our simulation results show better performance in the presented numerical example.

Basic definitions and notations. Let $\mathbb{X}$ be a normed vector space. The *weak* topology* on its topological dual $\mathbb{X}^*$ is the smallest topology with respect to which the linear functionals in $\{\mathbb{X}^* \ni x^* \rightarrow \langle x^*, x \rangle := x^*(x) : x \in \mathbb{X}\}$ are continuous. Let $\mathcal{X}$ be a metrizable topological space and let $\mathcal{B}(\mathcal{X})$ be its Borel σ -algebra. Let $C(\mathcal{X})$ be the Banach space of real-valued bounded continuous functions on $\mathcal{X}$ together with the sup-norm $\|\cdot\|_\infty$. We denote by $\mathcal{M}(\mathcal{X})$ the Banach space of finite signed Borel measures on $\mathcal{X}$ equipped with the total variation norm and by $\mathcal{P}(\mathcal{X})$ the set of Borel probability measures. Let $\delta_x \in \mathcal{P}(\mathcal{X})$ be the Dirac measure centred on $x \in \mathcal{X}$. When $\mathcal{X}$ is compact, $C(\mathcal{X})^* \simeq \mathcal{M}(\mathcal{X})$. In particular, we identify each $\mu \in \mathcal{M}(\mathcal{X})$ with the bounded linear functional $\langle \mu, \cdot \rangle : C(\mathcal{X}) \rightarrow \mathbb{R}$, $\langle \mu, l \rangle := \int_{\mathcal{X}} l \, d\mu$, for all $l \in C(\mathcal{X})$. Moreover, if $\mathcal{M}(\mathcal{X})^+$ is the convex cone of finite nonnegative Borel measures on $\mathcal{X}$, then its dual convex cone is the set $C(\mathcal{X})^+$ of nonnegative continuous functions on $\mathcal{X}$. The *support* of a measure $\mu \in \mathcal{M}(\mathcal{X})^+$ is the smallest closed set whose complement has zero measure and is denoted by $\text{spt } \mu$. If $X \subset \mathbb{R}^n$, we consider the Banach space $C^1(X) = \{\phi \in C(X) : \frac{\partial \phi}{\partial x_i} \in C(X), \text{ for all } i = 1, \dots, n\}$, with $\|\phi\|_1 = \|\phi\|_\infty + \sum_{i=1}^n \left\| \frac{\partial \phi}{\partial x_i} \right\|_\infty$. Let $\mathbb{R}[x]$ be the space of polynomials of the variable $x \in \mathbb{R}^n$. Let $\alpha = (\alpha_1, \dots, \alpha_n)^\top \in \mathbb{N}^n$ be a multi-index. A monomial is defined as $x^\alpha := x_1^{\alpha_1} x_2^{\alpha_2} \dots x_n^{\alpha_n}$ and its degree is $|\alpha| := \sum_{i=1}^n \alpha_i$. A polynomial $p \in \mathbb{R}[x]$ of degree $\deg(p) = d$ is given by $p(x) := \sum_{|\alpha| \leq d} p_\alpha x^\alpha$. Similarly, let $\mathbb{R}[x, u]$ be the space of polynomials of the variable $(x, u) \in \mathbb{R}^n \times \mathbb{R}^m$. Then a polynomial $q \in \mathbb{R}[x, u]$ of degree d has the form $q(x, u) := \sum_{\substack{\gamma \in \mathbb{N}^n \times \mathbb{N}^m \\ |\gamma| \leq d}} q_\gamma (x, u)^\gamma = \sum_{\substack{\alpha \in \mathbb{N}^n, \beta \in \mathbb{N}^m \\ |\alpha| + |\beta| \leq d}} q_{\alpha\beta} x^\alpha u^\beta$. Let $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^m$ be compact subsets. Given a measure $\mu \in \mathcal{M}(X \times U)^+$, the real number $z_\gamma := \int_{X \times U} (x, u)^\gamma \, d\mu(x, u) = \int_{X \times U} x^\alpha u^\beta \, d\mu(x, u)$, is called its *moment* of order $\gamma = (\alpha, \beta) \in \mathbb{N}^n \times \mathbb{N}^m$. Given a real sequence $z = (z_\gamma)_{\gamma \in \mathbb{N}^n \times \mathbb{N}^m}$, the *Riesz linear functional* $L_z : \mathbb{R}[x, u] \rightarrow \mathbb{R}$ is defined by $L_z(q) := \sum_{\gamma \in \mathbb{N}^n \times \mathbb{N}^m} q_\gamma z_\gamma$, for each polynomial $q(x, u) = \sum_{\gamma \in \mathbb{N}^n \times \mathbb{N}^m} q_\gamma (x, u)^\gamma$. Let $\nu_d(x, u) := ((x, u)^\gamma)_{|\gamma| \leq d}$. The *moment matrix* of order d is given by $M_d(z) := L_z(\nu_d(x, u) \nu_d(x, u)^\top)$, where the latter notation means that L_z is applied entrywise on the input matrix. Similarly, given $h \in \mathbb{R}[x, u]$ the *localizing matrix* of order d w.r.t. z and h is defined by $M_d(hz) := L_z(h(x, u) \nu_d(x, u) \nu_d(x, u)^\top)$. A polynomial $h \in \mathbb{R}[x, u]$ is a *sum of squares (s.o.s.)* if it can be written as $h(x, u) = \sum_{j \in J} h_j^2(x, u)$, for some finite index set J , and $\{h_j : j \in$

$J\} \subset \mathbb{R}[x, u]$. We denote by $\Sigma[x, u] \subset \mathbb{R}[x, u]$, the space of s.o.s. polynomials.

II. PROBLEM STATEMENT

Let $X \subset \mathbb{R}^n$ and $U \subset \mathbb{R}^m$. We consider the following infinite horizon control problem:

$$\begin{aligned} \dot{x}(t) &= f(x(t), u(t)), \quad t > 0 \\ x(0) &= x_0, \end{aligned} \quad (1)$$

where $x_0 \in X$ and $f : \mathbb{R}^n \times U \rightarrow \mathbb{R}^n$. A *control* is a Borel measurable function $u(\cdot) : [0, \infty) \rightarrow U$. The set of controls is denoted by $\mathcal{U}^0$. Under Assumption II.1 (H1) and (H2) below, for each $u(\cdot) \in \mathcal{U}^0$, the control system (1) admits a unique solution $x(\cdot|x_0, u(\cdot))$. We impose the state constraint

$$x(t|x_0, u(\cdot)) \in X, \quad \text{for all } t \geq 0. \quad (2)$$

The set of *admissible controls* is given by $\mathcal{U}_X(x_0) := \{u(\cdot) \in \mathcal{U}^0 : (2) \text{ holds}\}$. The OCP is described by

$$V_X(x_0) := \inf_{u(\cdot) \in \mathcal{U}_X(x_0)} \int_0^\infty e^{-\lambda t} g(x(t|x_0, u(\cdot)), u(t)) \, dt, \quad (3)$$

where $\lambda > 0$ is the *discount factor* and $g : \mathbb{R}^n \times U \rightarrow \mathbb{R}$. Consider the following conditions:

Assumption II.1 (Control model).

- (H1) f, g are bounded and continuous, Lipschitz in x uniformly in u ;
- (H2) X and U are compact;
- (H3) $\lambda > L$, where L is the Lipschitz constant for f ;
- (H4) $\mathcal{U}_X(x_0) \neq \emptyset$.

Under suitable assumptions [1, Thm. 5.10], the value function V_X is the unique constrained viscosity solution of the *Hamilton-Jacobi-Bellman (HJB) equation*

$$\lambda \phi(x) + \sup_{u \in U} \{-f(x, u) \cdot \nabla_x \phi(x) - g(x, u)\} = 0, \quad x \in \text{Int} X.$$

In general, the infimum in (3) is not attained, so our next step is to consider a relaxation of the OCP. A *relaxed control* or a *Young measure* is a measurable map $m(\cdot) : [0, \infty) \rightarrow \mathcal{P}(U)$, i.e., for each $h \in C(U)$, the scalar-valued function $[0, +\infty) \ni t \rightarrow \int_U h(u) \, dm_t(du)$ is Lebesgue measurable. The set of relaxed controls is denoted by $\mathcal{V}^0$. We consider the relaxed control system

$$\begin{aligned} \dot{x}(t) &= \int_U f(x(t), u) \, dm_t(u), \quad t \geq 0, \\ x(0) &= x_0. \end{aligned} \quad (4)$$

Given a relaxed control $m(\cdot)$, we denote by $x(\cdot|x_0, m(\cdot))$ the solution of the system (4). We impose the state constraint

$$x(t|x_0, m(\cdot)) \in X, \quad \text{for all } t \geq 0. \quad (5)$$

The set of admissible relaxed controls is given by $\mathcal{V}_X(x_0) := \{m(\cdot) \in \mathcal{V}^0 : (5) \text{ holds}\}$. Note that $\mathcal{U}_X(x_0) \subset \mathcal{V}_X(x_0)$ by

identifying each control $u(\cdot)$ with the Dirac measure-valued function $t \rightarrow \delta_{u(t)}$. The relaxed OCP is described by

$$V_X^R(x_0) := \min_{m(\cdot) \in \mathcal{V}(x_0)} \int_0^\infty e^{-\lambda t} \int_U g(x(t), u) dm_t(u) dt. \quad (6)$$

By using the weak* compactness of $\mathcal{P}(X)$ one can prove that the minimum is indeed attained. Moreover, $V_X^R(x_0) \leq V_X(x_0)$. We add the following assumption.

Assumption II.2 (No relaxation gap). $V_X(x_0) = V_X^R(x_0)$, for all $x_0 \in X$ such that $\mathcal{U}_X(x_0) \neq \emptyset$.

Sufficient conditions under which Assumption II.2 is satisfied are given e.g., by the assumptions of Filipov-Wazewski type theorems with state constraints [15] or by conditions described in [16, Prop. 4.3]. In particular, Assumption II.2 is satisfied in case of invariance of X w.r.t. the solutions of the OCP or in case of uncontrolled dynamics. Moreover in the case of input-affine dynamics, Assumption II.2 is satisfied under Assumption II.1 and if U is convex. For a related discussion on sufficient conditions see also the paragraph following Assumption 1 in [16].

III. LINEAR PROGRAMMING FORMULATION

For each $u(\cdot) \in \mathcal{U}_X(x_0)$, we define the discounted occupation measure $M^u \in \mathcal{M}(X \times U)^+$,

$$M^u(E) := \int_0^\infty e^{-\lambda t} \delta_{(x(t|x_0, u(\cdot)), u(t))}(E) dt, \quad (7)$$

for all $E \in \mathcal{B}(X \times U)$. Moreover, for each $m(\cdot) \in \mathcal{V}_X(x_0)$, we define the discounted relaxed occupation measure $N^m \in \mathcal{M}(X \times U)^+$,

$$N^m(E) := \int_0^\infty e^{-\lambda t} \int_U \delta_{(x(t|x_0, m(\cdot)), u)}(E) dm_t(u) dt, \quad (8)$$

for all $E \in \mathcal{B}(X \times U)$. Note that for every $l \in C(X \times U)$,

$$\langle M^u, l \rangle = \int_0^\infty e^{-\lambda t} l(x(t|x_0, u(\cdot)), u(t)) dt \quad (9)$$

$$\langle N^m, l \rangle = \int_0^\infty e^{-\lambda t} \int_U l(x(t|x_0, m(\cdot)), u) dm_t(u) dt. \quad (10)$$

By setting $\mathcal{M}_X := \{M^u : u(\cdot) \in \mathcal{U}_X(x_0)\}$ and $\mathcal{R}_X := \{N^m : m(\cdot) \in \mathcal{V}_X(x_0)\}$, we can rewrite the OCP (3) as $V_X(x_0) = \inf\{\langle \mu, g \rangle : \mu \in \mathcal{M}_X\}$, and the relaxed OCP (6) as $V_X^R(x_0) = \inf\{\langle \mu, g \rangle : \mu \in \mathcal{R}_X\}$.

In both cases, the cost is linear over μ but $\mathcal{R}_X$ and $\mathcal{M}_X$ are nonconvex sets. To transform the OCP into a convex program, we define the linear operator $A : C^1(X) \rightarrow C(X \times U)$ by $(A\phi)(x, u) := \lambda\phi(x) - \nabla_x \phi(x) \cdot f(x, u)$, for all $(x, u) \in X \times U$, $\phi \in C^1(X)$. Since $\|A\phi\|_\infty \leq (\lambda + \sup_{(x, u) \in X \times U} |f(x, u)|) \|\phi\|_\infty$, for all $\phi \in C^1(X)$, A is bounded and in particular it is weakly continuous. Let $A^* : \mathcal{M}(X \times U) \rightarrow C^1(X)^*$ be the adjoint map of A , defined by $\langle A^* \mu, \phi \rangle = \langle \mu, A\phi \rangle = \int_{X \times U} A\phi d\mu$, for all $\mu \in \mathcal{M}(X \times U)$ and $\phi \in C^1(X)$. By the weak continuity of

A , A^* is well defined and weakly* continuous. We introduce the linear program:

$$\mathcal{J}(x_0) := \inf\{\langle \mu, g \rangle : \mu \in \mathcal{F}_X\}, \quad (11)$$

where, $\mathcal{F}_X := \{\mu \in \mathcal{M}(X \times U)^+ : A^* \mu = \delta_{x_0}\}$ and δ_{x_0} is identified as an element of $C^1(X)^*$. The following proposition establishes the relationship between the above programs and asserts that the infinite LP (11) admits a minimizer.

Proposition III.1. $\mathcal{M}_X \subset \mathcal{R}_X \subset \mathcal{F}_X$ and so $\mathcal{J}(x_0) \leq V_X^R(x_0) \leq V_X(x_0)$. Moreover, if $\mathcal{F}_X \neq \emptyset$, then the infimum in (11) is attained.

Proof. Since $\mathcal{U}_X(x_0) \subset \mathcal{V}_X(x_0)$, we get $\mathcal{M}_X \subset \mathcal{R}_X$ and $V_X^R(x_0) \leq V_X(x_0)$. Moreover, if $m(\cdot) \in \mathcal{V}(x_0)$, then for each $\phi \in C^1(X)$ it holds that $\nabla_x \phi(x(t|x_0, m(\cdot))) \cdot \int_U f(x(t|x_0, u)) dm_t(u) = \frac{d}{dt} \phi(x(t|x_0, m(\cdot)))$, for almost all $t \geq 0$. Integration by parts gives $\langle N^u, A\phi \rangle = \phi(x_0)$, for all $\phi \in C^1(X)$. Equivalently, $A^* N^u = \delta_{x_0}$ and thus $\mathcal{R}_X \subset \mathcal{F}_X$ and $\mathcal{J}(x_0) \leq V_X^R(x_0)$. Next, note that $\mathcal{F}_X = \mathcal{M}(X \times U)^+ \cap (A^*)^{-1}(\{\delta_{x_0}\}) \subset \{\mu \in \mathcal{M}(X \times U) : \|\mu\| \leq \frac{1}{\lambda}\}$. By the Banach-Alaoglu Theorem and weak* continuity of A^* , we get that $\mathcal{F}_X \neq \emptyset$ is weakly* compact in $\mathcal{M}(X \times U)$. The solvability of (11) follows by the weak* continuity of $\mathcal{M}(X \times U) \ni \mu \rightarrow \langle \mu, g \rangle$. $\square$

The dual linear program of (11) is

$$\mathcal{J}^*(x_0) := \sup_{\phi \in C^1(X)} \{\phi(x_0) : A\phi \leq g \text{ on } X \times U\}. \quad (12)$$

The following notion will be valuable for the interpretation of the feasible solutions of (12).

Definition III.1. A function $\phi \in C^1(X)$ is called a smooth subsolution to the HJB equation if $A\phi \leq g$ on $X \times U$.

Notice that the dual program (12) is always feasible. Indeed, the function $\phi(x) := C$, $x \in X$ with $C \leq \frac{-\sup_{(x, u) \in X \times U} |g(x, u)|}{\lambda}$ is feasible for the dual program. Moreover, any smooth subsolution to the HJB equation, gives a global lower bound for the value function in (3).

Lemma III.1 (Lower bound). If $\phi \in C^1(X)$ is a feasible solution for the dual program (12), then $\phi \leq V_X$ on X .

Proof. Let $x \in X$ s.t. $\mathcal{U}_X(x) \neq \emptyset$. Since $\langle M^u, A\phi \rangle = \phi(x)$, for all $u(\cdot) \in \mathcal{U}_X(x)$ and $A\phi \leq g$, we get $\phi(x) \leq \langle M^u, g \rangle$, for all $u(\cdot) \in \mathcal{U}_X(x)$. Thus, $\phi(x) \leq V_X(x)$. $\square$

Theorem III.1 (Equivalence). Under Assumptions II.1 and II.2, $V_X(x_0) = \mathcal{J}(x_0) = \mathcal{J}^*(x_0)$.

Note that the assumptions in Theorem III.1 are weaker than those in [7, Theorem 4.4], as indicated in the introduction.

The proof of Theorem III.1, is divided in several parts. The first step is to prove that strong duality holds.

Lemma III.2 (Absence of duality gap). Under Assumption II.1 (H1) and (H2), if (11) is feasible, then there is no duality gap, i.e., $\mathcal{J}^*(x_0) = \mathcal{J}(x_0)$.

Proof. By virtue of [17, Th. 3.10], it suffices to prove that the set $D = \{(A^*\mu, \langle \mu, g \rangle) : \mu \in \mathcal{M}(X \times U)^+\}$ is $\sigma(C^1(X)^* \times \mathbb{R}, C^1(X) \times \mathbb{R})$ -closed [18, Def. 8.2.1.]. Similarly to [12, Theorem 2.3 (ii)], this follows by the Banach-Alaoglu Theorem and the weak* continuity of A^* . $\square$

Next, we consider the case of unconstrained state space. Let $C_0(\mathbb{R}^n \times U)$ be the Banach space of bounded continuous functions that vanish at infinity with the sup-norm. Then $\mathcal{M}(\mathbb{R}^n \times U)$ is isometrically isomorphic to $C_0(\mathbb{R}^n \times U)^*$ and $\mathcal{M}(\mathbb{R}^n \times U) \subset C(\mathbb{R}^n \times U)^*$. Let

$$V(x_0) := \inf_{u(\cdot) \in \mathcal{U}^0} \int_0^\infty e^{-\lambda t} g(x(t|x_0, u(\cdot)), u(t)) dt. \quad (13)$$

For each $u(\cdot)$ in $\mathcal{U}^0$, we define the corresponding discounted occupation measure $M^u \in \mathcal{M}(\mathbb{R}^n \times U)^+$ by (7), for all $E \in \mathcal{B}(\mathbb{R}^n \times U)$. Then, (9) holds for every $l \in C(\mathbb{R}^n \times U)$. We set $\mathcal{M} := \{M^u : u(\cdot) \in \mathcal{U}^0(x_0)\}$. Then the OCP (13) can be written as $V(x_0) = \inf\{\langle \mu, g \rangle : \mu \in \mathcal{M}\}$. As before, we define the linear operator $B : C^1(\mathbb{R}^n) \rightarrow C(\mathbb{R}^n \times U)$ by $(B\phi)(x, u) := \lambda\phi(x) - \nabla_x \phi(x) \cdot f(x, u)$, for all $(x, u) \in \mathbb{R}^n \times U$, $\phi \in C^1(\mathbb{R}^n)$. B is linear, bounded and weakly continuous. Moreover, its adjoint $B^* : C(\mathbb{R}^n \times U)^* \rightarrow C^1(\mathbb{R}^n)^*$ is well defined and weakly* continuous. We consider the convex optimization problem

$$P(x_0) := \inf\{\langle \mu, g \rangle : \mu \in \mathcal{F}\},$$

where, $\mathcal{F} := \{\mu \in \mathcal{M}(\mathbb{R}^n \times U)^+ : B^*\mu = \delta_{x_0}\}$. Similarly as before, $\mathcal{M} \subset \mathcal{F}$ and $\mathcal{F}$ is convex and weakly* compact in $\mathcal{M}(\mathbb{R}^n \times U)$. In particular, the following is true.

Lemma III.3. *Under Assumptions (H1) and (H3), $\mathcal{F}$ is the weak* convex closure of $\mathcal{M}$.*

Proof. (H1) and (H3) imply that the value function V is Lipschitz continuous [1, Prop. 2.1]. By Rademacher's Theorem, V is almost everywhere differentiable. Therefore, $BV \leq g$ for almost all $x \in \mathbb{R}^n$ and all $u \in U$. W.l.o.g., we assume that $g \in C_0(\mathbb{R}^n \times U)$. This assumption can be lifted after the proof of Theorem. By using standard mollification techniques as in [19, Lemma 3.2], we can construct a sequence of approximate smooth subsolutions to the HJB equation for the unconstrained OCP (13). That is, there exist $\{V_n\}_n \subset C^1(\mathbb{R}^n)$ and $\{\varepsilon_n\}_n \subset [0, \infty)$, such that $\lim_{n \rightarrow \infty} \|V - V_n\|_\infty = 0$, $BV_n \leq g + \varepsilon_n$ on $\mathbb{R}^n \times U$ and $\lim_{n \rightarrow \infty} \varepsilon_n = 0$. We will use this result to prove that $V(x_0) = P(x_0)$. Clearly, $P(x_0) \leq V(x_0)$. Assume for the sake of contradiction that $P(x_0) < V(x_0)$. Then, there exists $\mu \in \mathcal{M}(\mathbb{R}^n \times U)^+$, s.t. $B^*\mu = \delta_{x_0}$ and $\langle \mu, g \rangle < V(x_0)$. So, for all $n \in \mathbb{N}$, $V_n(x_0) = \langle \delta_{x_0}, V_n \rangle = \langle B^*\mu, V_n \rangle = \langle \mu, BV_n \rangle \leq \langle \mu, g + \varepsilon_n \rangle$. By taking $n \rightarrow \infty$, we get $V(x_0) \leq \langle \mu, g \rangle$, which is a contradiction. Therefore, we have proved that $V(x_0) = P(x_0)$. Finally, assume for the sake of contradiction that $\text{conv}(\mathcal{M})^* \neq \mathcal{F}$. Let $\nu \in \mathcal{F}/\text{conv}(\mathcal{M})^*$. By the separation theorem, there exists $g \in C_0(\mathbb{R}^n \times U)$ and $s \in \mathbb{R}$ such that $\langle M^u, g \rangle \geq s$, for all $u \in \mathcal{U}^0$ and $\langle \nu, g \rangle < s$. Since a function in $C_0(\mathbb{R}^n \times U)$ can be uniformly approximated by a Lipschitz continuous function in $C_0(\mathbb{R}^n \times U)$ [19, pg. 124], we can replace the function

$g \in C_0(\mathbb{R}^n \times U)$ with a new Lipschitz continuous function in $C_0(\mathbb{R}^n \times U)$ and adjust s so that the separation property still holds. Assume that this specific g is the cost function of the OCP (13). Then, $P(x_0) < s \leq V(x_0)$, which is a contradiction. $\square$

Next, we consider for each $m(\cdot) \in \mathcal{V}^0$, the corresponding relaxed occupation measure $N^m \in \mathcal{M}(\mathbb{R}^n \times U)^+$ defined by (8), for each $E \in \mathcal{B}(\mathbb{R}^n \times U)$ and satisfying (10), for each $l \in C(\mathbb{R}^n \times U)$. Set $\mathcal{R} := \{N^m : m(\cdot) \in \mathcal{V}^0\}$. It is a known result that $\mathcal{R}$ is the weak* closure of $\mathcal{M}$.

Corollary III.1. *For each $\mu \in \mathcal{F}$, there exists $\Lambda \in \mathcal{P}(\mathcal{R})$ such that $\langle \mu, l \rangle = \int_{\mathcal{R}} \langle \gamma, l \rangle d\Lambda(\gamma)$, for all $l \in C(\mathbb{R}^n \times U)$.*

Proof. Similar arguments as in [8, Cor. 1.4] prove the result for all $l \in C_0(\mathbb{R}^n \times U)$. Since a bounded continuous function can be approximated pointwise by an increasing sequence of continuous functions with compact support, the application of the monotone convergence theorem completes the proof. $\square$

We will now return to the state constrained OCP (3). The following Lemma indicates the relation between all the previously defined sets.

Lemma III.4. *We make the convention that each measure $\mu \in \mathcal{M}(X \times U)^+$ is a measure on the whole space $\mathbb{R}^n \times U$ with $\text{spt } \mu \subset X \times U$. The following hold,*

$$\begin{aligned} \mathcal{M}_X &= \{\mu \in \mathcal{M} : \text{spt } \mu \subset X \times U\}, \\ \mathcal{R}_X &= \{\mu \in \mathcal{R} : \text{spt } \mu \subset X \times U\}, \\ \mathcal{F}_X &= \{\mu \in \mathcal{F} : \text{spt } \mu \subset X \times U\}. \end{aligned}$$

Proof. We will prove the inclusion $\supset$ in the first assertion. Let $M^u \in \mathcal{M}$, for some $u(\cdot) \in \mathcal{U}^0$, such that $\text{spt } M^u \subset X \times U$. Note that $\|M^u\| = \frac{1}{\lambda}$. Therefore, $\int_0^{+\infty} e^{-\lambda t} dt = \frac{1}{\lambda} = M^u(X \times U) = \int_0^{+\infty} e^{-\lambda t} \delta_{(x(t), u(t))}(X \times U) dt$. Thus, $(x(t), u(t)) \in X \times U$, for almost all $t \geq 0$. So, $u(\cdot) \in \mathcal{U}_X(x_0)$. $\square$

We are now ready to prove Theorem III.1.

Proof. By Proposition III.1, there exists a minimizer $\mu_0 \in \mathcal{F}_X$, such that $\langle \mu_0, g \rangle = \mathcal{J}_X(x_0) \leq V_X^R(x_0)$. Then, by Corollary III.1, there exists a Borel probability measure Λ on $\mathcal{R}$ such that $\langle \mu_0, l \rangle = \int_{\mathcal{R}} \langle \gamma, l \rangle d\Lambda(\gamma)$, for all $l \in C(\mathbb{R}^n \times U)$. We set $l(x, u) = d(x, X)$, where $d(x, X)$ is the euclidean distance of x from X . Since, $\text{spt } \mu_0 \subset X \times U$, we have $\langle \mu_0, d(\cdot, X) \rangle = 0$. Therefore, $\langle \gamma, d(\cdot, X) \rangle = 0$, for Λ -almost all $\gamma \in \mathcal{R}$. Equivalently, $\Lambda(\mathcal{R}_X) = 1$. So, $\mathcal{J}_X(x_0) = \langle \mu_0, g \rangle = \int_{\mathcal{R}_X} \langle \gamma, g \rangle d\Lambda(\gamma)$. We have that $\langle \mu_0, g \rangle \leq \langle \gamma, g \rangle$, for all $\gamma \in \mathcal{R}_X$. Assume for the sake of contradiction that $\langle \mu_0, g \rangle < \langle \gamma, g \rangle$, for all $\gamma \in \mathcal{R}_X$. Then, $\langle \mu_0, g \rangle < \int_{\mathcal{R}_X} \langle \gamma, g \rangle d\Lambda(\gamma)$ which is a contradiction. Therefore, there exists $\gamma_0 \in \mathcal{R}_X$ such that $\mathcal{J}(x_0) = \langle \mu_0, g \rangle = \langle \gamma_0, g \rangle \geq V_X^R(x_0)$. So, we have proved that $\mathcal{J}(x_0) = V_X^R(x_0)$ and that the relaxed admissible control associated to γ_0 is optimal for the relaxed OCP (6). Finally, under Assumption II.2 we have $V_X(x_0) = V_X^R(x_0) = \mathcal{J}(x_0)$. $\square$

IV. PRIMAL-DUAL MOMENT-S.O.S. LMIS, OPTIMAL CONTROL SYNTHESIS AND NUMERICAL EXAMPLE

Throughout this section we make the following assumptions on the data of the OCP (3).

Assumption IV.1 (Polynomial Data).

- (A1) $g \in \mathbb{R}[x, u]$ and $\{f_k\}_{k=1}^n \subset \mathbb{R}[x, u]$;
(A2) $X \times U$ is a compact basic semi-algebraic set of the form $X \times U = \{(x, u) \in \mathbb{R}^n \times \mathbb{R}^m : q_i(x, u) \geq 0, i = 1, \dots, N\}$, for some $N \in \mathbb{N}$ and $\{q_i\}_{i=1}^N \subset \mathbb{R}[x, u]$;
(A3) There exists $\xi \in \mathbb{R}[x, u]$ such that $\xi \in \left\{ \sigma_0 + \sum_{i=1}^N \sigma_i q_i : \{\sigma_i\}_{i=0}^N \subset \Sigma[x, u] \right\}$ and the level set $\{(x, u) \in \mathbb{R}^n \times \mathbb{R}^m : \xi(x, u) \geq 0\}$ is compact.

Assumption (A3) is not restrictive. Indeed, if necessary, we can add in the representation of $X \times U$ the polynomial inequality $q_{N+1}(x, u) := K - \|(u, x)\|_2^2 \geq 0$, where K is an upper bound of $X \times U$. Then, the new equivalent representation satisfies Assumption (A3).

We now formulate Lasserre's hierarchy of primal-dual moment-s.o.s. semidefinite relaxations [20] for our specific problem. Under Assumption (A3), we are able to apply Putinar's Postivstellensatz [21] and characterize in a computationally tractable way the moment-sequences of a finite Borel measure on the compact basic semi-algebraic set $X \times U$. So, we conclude that the primal LP (11) is equivalent to an infinite SDP problem. By optimizing over finite truncated sequences of moments, we obtain a hierarchy of finite SDP relaxations. In particular, let $\deg(g) = 2v_0$ or $2v_0 - 1$ and $\deg(q_i) = 2v_i$ or $2v_i - 1$, for all $i = 1, \dots, N$. For each $r \geq \max_{i=0, \dots, N} v_i$, we consider the LMI-relaxation of order r ,

$$\mathcal{J}_r(x_0) := \begin{cases} \inf_{z=(z_\gamma)_{|\gamma| \leq 2r}} & L_z(g) \\ \text{s.t.} & L_z(h^{(\alpha)}) = b^{(\alpha)}, \text{ for all } \alpha \in \mathcal{A}_r \\ & M_r(z) \succeq 0, \\ & M_{r-v_i}(q_i z) \succeq 0, i = 1, \dots, N, \end{cases} \quad (14)$$

where $h^{(\alpha)}(x, u) := \lambda x^\alpha - \nabla x^\alpha \cdot f(x, u)$ and $b^{(\alpha)} := x_0^\alpha$, for each $\alpha \in \mathbb{N}^n$ and $\mathcal{A}_r := \{\alpha \in \mathbb{N}^n : \deg(h^{(\alpha)}) \leq 2r\}$.

The dual of (14) is

$$\mathcal{J}_r^*(x_0) := \begin{cases} \sup_{\lambda_{\alpha}, \sigma_i} & \phi(x_0) \\ \text{s.t.} & \phi(x) = \sum_{\deg h^{(\alpha)} \leq 2r} \lambda_{\alpha} x^{\alpha}, \\ & g - A\phi = \sigma_0 + \sum_{i=1}^N \sigma_i q_i, \\ & \lambda_{\alpha} \in \mathbb{R}, \sigma_i \in \Sigma[x, u], \\ & \deg \sigma_i q_i \leq 2r, i = 0, \dots, N. \end{cases} \quad (15)$$

We highlight that the dual SDP (15) is a tightening of the dual linear program (12), since its feasible solutions are polynomial subsolutions to the HJB equation. Moreover, note that the s.o.s. representation that $g - A\phi$ is required to satisfy can be expressed via SDP feasibility tests, since the degree of the involved s.o.s. polynomials is fixed.

Note that the number of moments, i.e., the size of the vector $z = (z_\gamma)_{|\gamma| \leq 2r}$, in the LMI relaxation (14) of order r is $R := \binom{n+m+2r}{2r}$. Therefore, for a fixed state space with

dimension n and control space with dimension m , R grows as $O(r^{n+m})$. If the relaxation order r is fixed, then R grows as $O((n+m)^r)$, that is polynomially in the size of the problem. This means that taking into account the current performance of general-purpose SDP solvers, the moment approach is appealing for small to medium size OCPs. One can exploit structure, (e.g., sparsity) of the specific instance OCP under consideration to improve scalability and overcome this computational limit [20, Sec. 4.6], [22]. Moreover, in our work, we have chosen the monomial basis to represent polynomials. However, other bases (e.g., Chebyshev) may be more efficient from a computational point of view [14]. One other possible strategy is to develop alternative positivity certificates, which should be less computationally demanding [23].

Theorem IV.1 ([20, Thm. 4.3], [24, Lemma 5]). *Under Assumption IV.1 and if (11) is feasible, then $\mathcal{J}_r^*(x_0) = \mathcal{J}_r(x_0) \uparrow \mathcal{J}(x_0) = \mathcal{J}^*(x_0)$, as $r \rightarrow \infty$. If in addition, Assumptions II.1 and II.2 hold, then $\mathcal{J}_r^*(x_0) = \mathcal{J}_r(x_0) \uparrow V_X(x_0)$, as $r \rightarrow \infty$.*

Some preliminary results regarding the convergence rate, are available in the recent work [25].

We next discuss how the resulting dual s.o.s. SDP relaxations can be used for the computation of an approximate feedback controller. Numerical examples in [13] for the finite horizon case, have shown that the optimal solution ϕ^* of (15) approximates well the value function $V(\cdot)$ along optimal trajectories starting from x_0 , but it gives an unsatisfactory approximation for the other points. This observation was formalized in [26, Theorem 1]. We would like to enlarge this region and have a good approximation of the value function V on a given set $X_0 \subset X$. Assume that $U_X(x_0) \neq \emptyset$, for all $x_0 \in X_0$. Let μ_0 be the uniform probability measure on X_0 and consider the average value $\bar{V}(\mu_0) := \langle V, \mu_0 \rangle = \int_{X_0} V(x) d\mu_0(x)$. Consider the primal averaged linear program

$$\mathcal{J}(\mu_0) := \inf_{\mu \in \mathcal{M}(X \times U)^+} \{ \langle \mu, g \rangle : A^* \mu = \mu_0 \}, \quad (16)$$

with corresponding dual

$$\mathcal{J}^*(\mu_0) := \sup_{\phi \in C^1(X)} \{ \langle \mu_0, \phi \rangle : A\phi \leq g \text{ on } X \times U \}. \quad (17)$$

Intuitively, the primal averaged LP (16) describes a superposition of a possibly uncountable set of OCPs. Under the assumptions of Theorem III.1 and by linearity, we have $\bar{V}(\mu_0) = \mathcal{J}(\mu_0) = \mathcal{J}^*(\mu_0)$.

Assume that a relaxation of order r of the averaged dual program (17) has an optimal solution ϕ . By routine modifications in [26, Theorem 2], we get that ϕ is a good approximation of the value function V along all optimal trajectories which start from X_0 . A natural candidate for feedback control is derived by the following minimization problem: For fixed $x \in X$

$$u^*(x) = \arg \min_{u \in U} (g(x, u) - A\phi(x, u)). \quad (18)$$

For each $x \in X$, consider a closed neighborhood $\mathcal{S}_x$ of x . Similarly to [13], we propose the following heuristic iterative algorithm.

- 1) Set $\bar{x} = x(t)$.

- 2) Solve the semidefinite relaxation of the averaged dual program (17) (where $X_0 = \mathcal{S}_{\bar{x}}$).
- 3) Apply the control law derived by (18) until $x(t) \notin \mathcal{S}_{\bar{x}}$.
- 4) Go to step 1.

We conclude with an illustrative numerical example. Consider the nonlinear double integrator

$$\inf_{u(\cdot)} \int_0^\infty e^{-0.1t} (x_1(t)^2 + x_2(t)^2) dt$$

$$\begin{aligned} \text{s.t. } \quad & \dot{x}_1(t) = x_2(t) + 0.1x_1(t)^3, \\ & \dot{x}_2(t) = -0.3u(t), \\ & x(0) = (0, 0.7)^T, \\ & x(t) \in X := \{x : \|x\|_2 \leq 1\}, \quad u(t) \in U := [-1, 1]. \end{aligned}$$

A problem with similar data was considered in [14, Section 7.1.]. The following table shows the approximation of the optimal value for different relaxation orders of the moment SDP (14). It seems that the sequence $(J_r(x_0))_r$ converges to 2.2043, as $r \rightarrow \infty$. Note that 2.2043 is a lower bound for $V_X(x_0)$.

TABLE I
NUMERICAL RESULTS FOR VARYING RELAXATION ORDER r .

r	2	3	4	5	6	7	11
$J_r(x_0)$	0.1121	1.6465	2.1978	2.2042	2.2042	2.2043	2.2043
CPU time	0.79	0.92	1.24	2.65	6.77	17.73	677.73

For the design of an approximate feedback control law, we use the averaged LP formulations (16) and (17), where the initial condition is described by the uniform probability measure μ_0 on the whole state constraint set X . If an optimal solution ϕ_r^* of the averaged dual SDP relaxation of order r is given, then by using (18) a feedback control is obtained by $u_r(x) = \text{sign} \frac{\partial}{\partial x_2} \phi_r^*(x)$, $x \in X$. Let V^{u_r} be its corresponding cost. Then V^{u_r} is an upper bound for $V_X(x_0)$. Figure 1 displays the trajectories obtained by the dual SDP relaxations of order $r = 3$ and $r = 4$. The corresponding costs are $V^{u_3} = 2.2479$ and $V^{u_4} = 2.2582$. Since the OCP of our example has no analytical solution, we evaluate the performance of the extracted feedback controller by evaluating the gap $G_r = 100 \frac{V^{u_r} - 2.2043}{2.2043}$. We have $G_3 = 1.98\%$ and $G_4 = 2.45\%$. This is a strong indication that the extracted feedback controllers are near optimal. In fact, one can verify that in the presented example, Assumptions II.1, II.2 and IV.1 are satisfied (input-affine polynomials dynamics with convex control space). Therefore, Theorems III.1 and IV.1 hold.

REFERENCES

- [1] M. Bardi and I. Capuzzo-Dolcetta, *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*. Birkhäuser Boston, Inc., Boston, MA, 1997.
- [2] H. J. Pesch, "A practical guide to the solution of real-life optimal control problems," *Control and Cybernetics*, vol. 23, no. 1/2, pp. 7–60, 1994.
- [3] L. Grüne and J. Pannek, *Nonlinear Model Predictive Control: Theory and Algorithms*. Springer Publishing Company, Incorporated, 2013.
- [4] O. von Stryk and R. Bulirsch, "Direct and indirect methods for trajectory optimization," *Annals of Operations Research*, vol. 37, no. 1, pp. 357–373, 1992.
- [5] Y. Tassa and T. Erez, "Least squares solutions of the HJB equation with neural network value-function approximators," *Trans. Neur. Netw.*, vol. 18, no. 4, pp. 1031–1041, Jul. 2007.
- [6] D. Hernández-Hernández, O. Hernández-Lerma, and M. Taksar, "The linear programming approach to deterministic optimal control problems," *Appl. Math. (Warsaw)*, vol. 24, no. 1, pp. 17–33, 1996.
- [7] V. Gaitsgory and M. Quincampoix, "Linear programming approach to deterministic infinite horizon optimal control problems with discounting," *SIAM J. Control Optim.*, vol. 48, no. 4, pp. 2480–2512, 2009.

Fig. 1. Trajectories obtained from relaxation order $r = 3$ and $r = 4$.

- [8] R. Vinter, "Convex duality and nonlinear optimal control," *SIAM J. Control Optim.*, vol. 31, no. 2, pp. 518–538, 1993.
- [9] A. G. Bhatt and V. S. Borkar, "Occupation measures for controlled Markov processes: characterization and optimality," *Ann. Probab.*, vol. 24, no. 3, pp. 1531–1562, 1996.
- [10] W. H. Fleming and D. Vermes, "Convex duality approach to the optimal control of diffusions," *SIAM J. Control Optim.*, vol. 27, no. 5, pp. 1136–1155, 1989.
- [11] P. Mohajerin Esfahani, T. Sutter, D. Kuhn, and J. Lygeros, "From Infinite to Finite Programs: Explicit Error Bounds with Applications to Approximate Dynamic Programming," <https://arxiv.org/abs/1701.06379>, Jan. 2017.
- [12] J. B. Lasserre, D. Henrion, C. Prieur, and E. Trélat, "Nonlinear optimal control via occupation measures and LMI-relaxations," *SIAM J. Control Optim.*, vol. 47, no. 4, pp. 1643–1666, 2008.
- [13] D. Henrion, J. B. Lasserre, and C. Savorgnan, "Nonlinear optimal control synthesis via occupation measures," in *2008 47th IEEE Conference on Decision and Control*, Dec 2008, pp. 4749–4754.
- [14] M. Korda, D. Henrion, and C. N. Jones, "Controller design and value function approximation for nonlinear dynamical systems," *Automatica J. IFAC*, vol. 67, pp. 54–66, 2016.
- [15] H. Frankowska and F. Rampazzo, "Filippov's and Filippov-Ważewski's theorems on closed domains," *Journal of Differential Equations*, vol. 161, no. 2, pp. 449–478, 2000.
- [16] V. Gaitsgory, "Averaging and near viability of singularly perturbed control systems," *Journal of Convex Analysis*, vol. 13, no. 2, pp. 329–352, 2006.
- [17] E. J. Anderson and P. Nash, *Linear programming in infinite-dimensional spaces*. John Wiley & Sons, Ltd., Chichester, 1987.
- [18] L. Narici and E. Beckenstein, *Topological vector spaces*, 2nd ed. CRC Press, Boca Raton, FL, 2011.
- [19] W. H. Fleming and D. Vermes, *Generalized Solutions in the Optimal Control of Diffusions*. New York, NY: Springer New York, 1988, pp. 119–127.
- [20] J. B. Lasserre, *Moments, Positive Polynomials and Their Applications*. London: Imperial College Press, 2010.
- [21] M. Putinar, "Positive polynomials on compact semi-algebraic sets," *Indiana Univ. Math. J.*, vol. 42, no. 3, pp. 969–984, 1993.
- [22] Z. Yang, G. Fantuzzi, and A. Papachristodoulou, "Exploiting Sparsity of Coefficient Matching Conditions in Sum-of-Squares Programming using ADMM," *ArXiv e-prints*, 2017.
- [23] A. A. Ahmadi and A. Majumdar, "DSOS and SDSOS optimization: LP and SOCP-based alternatives to sum of squares optimization," in *2014 48th Annual Conference on Information Sciences and Systems (CISS)*, March 2014, pp. 1–5.
- [24] M. Korda, D. Henrion, and 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, Oct. 2014.
- [25] —, "Convergence rates of moment-sum-of-squares hierarchies for optimal control problems," *Systems & Control Letters*, vol. 100, pp. 1–5, 2017.
- [26] D. Henrion and E. Pauwels, "Linear conic optimization for nonlinear optimal control," 2014.