

Linear conic optimization for nonlinear optimal control

Didier Henrion^{1,2,3}, Edouard Pauwels^{1,2}

Draft of July 7, 2014

Abstract

Infinite-dimensional linear conic formulations are described for nonlinear optimal control problems. The primal linear problem consists of finding occupation measures supported on optimal relaxed controlled trajectories, whereas the dual linear problem consists of finding the largest lower bound on the value function of the optimal control problem. Various approximation results relating the original optimal control problem and its linear conic formulations are developed. As illustrated by a couple of simple examples, these results are relevant in the context of finite-dimensional semidefinite programming relaxations used to approximate numerically the solutions of the infinite-dimensional linear conic problems.

1 Motivation

In [8, 9], J.-B. Lasserre described a hierarchy of convex semidefinite programming (SDP) problems allowing to compute bounds and find global solutions for finite-dimensional nonconvex polynomial optimization problems. Each step in the hierarchy consists of solving a primal moment SDP problem and a dual polynomial sum-of-squares (SOS) SDP problem corresponding to discretizations of infinite-dimensional linear conic problems, namely a primal linear programming (LP) problem on the cone of nonnegative measures, and a dual LP problem on the cone of nonnegative continuous functions. The number of variables (number of moments in the primal SDP, degree of the SOS certificates in the dual SDP) increases when progressing in the hierarchy, global optimality can be ensured by checking rank conditions on the moment matrices, and global optimizers can be extracted by numerical linear algebra. For more information on the moment-SOS hierarchy and its applications, see [11].

This approach was then extended to polynomial optimal control in [10]. Whereas the key idea in [8, 9] was to reformulate a (finite-dimensional) nonconvex polynomial optimization on a compact semi-algebraic set into an LP in the (infinite-dimensional) space of

¹CNRS, LAAS, 7 avenue du colonel Roche, F-31400 Toulouse, France.

²Université de Toulouse, LAAS, F-31400 Toulouse, France.

³Faculty of Electrical Engineering, Czech Technical University in Prague, Technická 2, CZ-16626 Prague, Czech Republic

probability measures supported on this set, the key idea in [10], also developed in [5], was to reformulate an (infinite-dimensional) nonconvex polynomial optimal control problem with compact constraint set into an LP in the (infinite-dimensional) space of occupation measures supported on this set. Note that LP formulations of optimal control problems (on ordinary differential equations and partial differential equations) are classical, and can be traced back to the work by L. C. Young, Filippov, as well as Warga and Gamkrelidze, amongst many others. For more details and a historical survey, see e.g. [4, Part III]. We believe that what is innovative in [10] is the observation that the infinite-dimensional linear formulations for optimal control problems can be solved numerically with a moment-SOS hierarchy of the same kind as those used in polynomial optimization.

The objective of this contribution is to revisit the approach of [10] and to survey the use of occupation measures to linearize polynomial optimal control problems. This is an opportunity to describe duality in infinite-dimensional conic problems, as well as various approximation results on the value function of the optimal control problem. The primal LP consists of finding occupation measures supported on optimal relaxed controlled trajectories, whereas the dual LP consists of finding the largest lower bound on the value function of the optimal control problem. The value function is the solution (in a suitably defined weak sense) of a nonlinear partial differential equation called the Hamilton-Jacobi-Bellman equation, see e.g. [13, Chapters 8 and 9] and [3, Chapters 19 and 24]. It is traditionally used for verification of optimality, and for explicit computation of optimal control laws, but we do not describe these applications here.

2 Polynomial optimal control

We consider polynomial optimal control problems (POCPs) of the form

$$\begin{aligned}
 v^*(t_0, x_0) &:= \inf_u \int_{t_0}^T l(x(t), u(t)) dt + l_T(x(T)) \\
 \text{s.t.} \quad &\dot{x}(t) = f(x(t), u(t)), \quad x(t_0) = x_0 \\
 &x(t) \in X, \quad t \in [t_0, T] \\
 &u(t) \in U, \quad t \in [t_0, T] \\
 &x(T) \in X_T
 \end{aligned} \tag{1}$$

where the dot denotes time derivative, $l \in \mathbb{R}[x, u]$ is a given Lagrangian (integral cost), $l_T \in \mathbb{R}[x]$ is a given terminal cost, $f \in \mathbb{R}[x, u]^n$ is a given dynamics (vector field), $X \subset \mathbb{R}^n$ is a given compact state constraint set, $U \subset \mathbb{R}^m$ is a given compact control constraint set, $X_T \subset X$ is a given compact terminal state constraint set. Also given are the terminal time $T \geq 0$, the initial time $t_0 \in [0, T]$ and the initial condition $x_0 \in X$. In POCP (1), the minimum is with respect to all control laws $u \in \mathcal{L}^\infty([t_0, T]; U)$ which are bounded functions of time t with values in U , and the resulting state trajectories $x \in \mathcal{L}^\infty([t_0, T]; X)$ which are bounded functions of time t with values in X .

Let $\mathcal{A} \subset [0, T] \times X$ denote the set of values (t_0, x_0) for which there is a controlled trajectory $(x, u) \in \mathcal{L}^\infty([t_0, T]; X \times U)$ starting at $x(t_0) = x_0$ and admissible for POCP (1). The function $(t_0, x_0) \mapsto v^*(t_0, x_0)$ defined in (1) is called the value function, and its domain is \mathcal{A} .

3 LP formulation

As explained in the introduction, to derive an LP formulation of POCP (1) we have to introduce measures on trajectories, the so-called occupation measures. The first step is to replace classical controls with probability measures, and for this we have to define additional notations.

Given a compact set $X \subset \mathbb{R}^n$, let $\mathcal{C}(X)$ denote the space of continuous functions supported on X , and let $\mathcal{C}_+(X)$ denote its nonnegative elements, the cone of nonnegative continuous functions on X . Let $\mathcal{M}_+(X) = \mathcal{C}_+(X)'$ denote its topological dual, the set of all nonnegative continuous linear functional on $\mathcal{C}(X)$. By a Riesz representation theorem, these are nonnegative Borel-regular measures, or Borel measures, supported on X . The topology in $\mathcal{C}_+(X)$ is the strong topology of uniform convergence, whereas the topology in $\mathcal{M}_+(X)$ is the weak-star topology. The duality bracket

$$\langle v, \mu \rangle := \int_X v(x) \mu(dx)$$

denotes the integration of a function $v \in \mathcal{C}_+(X)$ against a measure $\mu \in \mathcal{M}_+(X)$. For background on weak-star topology see e.g. [12, Section 5.10] or [2, Chapter IV]. Finally, let us denote by $\mathcal{P}(X)$ the set of probability measures supported on X , consisting of Borel measures $\mu \in \mathcal{M}_+(X)$ such that $\langle 1, \mu \rangle = 1$.

3.1 Relaxed controls

In POCP (1), given $(t_0, x_0) \in \mathcal{A}$, let $(x_k, u_k)_{k \in \mathbb{N}} \in \mathcal{L}^\infty([t_0, T]; X \times U)$ denote a minimizing sequence of admissible controlled trajectories, i.e. it holds

$$\dot{x}_k(t) = f(x_k(t), u_k(t)), \quad x(t_0) = x_0$$

and

$$\lim_{k \rightarrow \infty} \int_{t_0}^T l(x_k(t), u_k(t)) dt + l_T(x_k(T)) = v^*(t_0, x_0).$$

In general the infimum in POCP (1) is not attained, so our next step is to assume that, at each time $t \in [t_0, T]$, the control is not a vector $u(t) \in U$, but a time-dependent probability measure $\omega(du | t) \in \mathcal{P}(U)$ which rules the distribution of the control in U . We use the notation $\omega_t := \omega(\cdot | t)$ to emphasize the dependence on time. This is called a relaxed control, or stochastic control, or Young measure in the functional analysis literature. POCP (1) is then relaxed to

$$\begin{aligned} v_R^*(t_0, x_0) &:= \min_{\omega} \int_{t_0}^T \langle l(x(t), \cdot), \omega_t \rangle dt + l_T(x(T)) \\ \text{s.t.} & \quad \dot{x}(t) = \langle f(x(t), \cdot), \omega_t \rangle, \quad x(t_0) = x_0 \\ & \quad x(t) \in X, \quad t \in [t_0, T] \\ & \quad \omega_t \in \mathcal{P}(U), \quad t \in [t_0, T] \\ & \quad x(T) \in X_T \end{aligned} \tag{2}$$

where the minimization is w.r.t. a relaxed control. Note that we replaced the infimum in POCP (1) with a minimum in relaxed POCP (2). Indeed, it can be proved that this minimum is always attained using (weak-star) compactness of the space of probability measures with compact support.

Since classical controls $u \in \mathcal{L}^\infty([t_0, T]; U)$ are a particular case of relaxed controls $\omega_t \in \mathcal{P}(U)$ corresponding to the choice $\omega_t = \delta_{u(t)}$ for a.e. $t \in [t_0, T]$, the minimum in relaxed POCP (2) is smaller than the infimum in classical POCP (1), i.e.

$$v^*(t_0, x_0) \geq v_R^*(t_0, x_0).$$

Contrived optimal control problems (e.g. with overly stringent state constraints) can be cooked up such that the inequality is strict, i.e. $v^*(t_0, x_0) > v_R^*(t_0, x_0)$, see e.g. the examples in [7, Appendix C]. We do not consider that these examples are practically relevant, and hence the following assumption will be made.

Assumption 1 (No relaxation gap) *For any relaxed controlled trajectory (x, ω_t) admissible for relaxed POCP (2), there is a sequence of controlled trajectories $(x_k, u_k)_{k \in \mathbb{N}}$ admissible for POCP (1) such that*

$$\lim_{k \rightarrow \infty} \int_{t_0}^T v(x_k(t), u_k(t)) dt = \int_{t_0}^T \langle v(x(t), \cdot), \omega_t \rangle dt$$

for every function $v \in \mathcal{C}(X \times U)$. Then it holds

$$v^*(t_0, x_0) = v_R^*(t_0, x_0)$$

for every $(t_0, x_0) \in \mathcal{A}$.

Note that this assumption is satisfied under the classical controllability and/or convexity conditions used in the Filippov-Ważewski Theorem with state constraints, see [6] and the discussions around Assumption I in [5] and Assumption 2 in [7]. However, let us point out that Assumption 1 does not imply that the infimum is attained in POCP (1). Conversely, if the infimum is attained, the values of POCP (1) and relaxed POCP (2) coincide, and Assumption 1 is satisfied.

3.2 Occupation measure

Given initial data $(t_0, x_0) \in \mathcal{A}$, and given a relaxed control $\omega_t \in \mathcal{P}(U)$, the unique solution of the ODE

$$\dot{x}(t) = \langle f(x(t), \cdot), \omega_t \rangle, \quad x(t_0) = x_0 \quad (3)$$

in relaxed POCP (2) is given by

$$x(t) = x_0 + \int_{t_0}^t \langle f(x(s), \cdot), \omega_s \rangle ds \quad (4)$$

for every $t \in [t_0, T]$. Let us then define

$$\mu(dt, dx, du) := dt \delta_{x(t)}(dx) \omega_t(du) \in \mathcal{M}_+([t_0, T] \times X \times U) \quad (5)$$

as the occupation measure concentrated uniformly in time on the state trajectory starting at x_0 at time t_0 , for the given relaxed control ω_t . An analytic interpretation is that integration w.r.t. the occupation measure is equivalent to time-integration along system trajectories, i.e.

$$\int_{t_0}^T v(t, x(t)) dt = \int_{t_0}^T \int_X \int_U v(t, x) \mu(dt, dx, du) = \langle v, \mu \rangle$$

given any test function $v \in \mathcal{C}([t_0, T] \times X)$.

Let us define the linear operator $\mathcal{L} : \mathcal{C}^1([t_0, T] \times X) \rightarrow \mathcal{C}([t_0, T] \times X \times U)$ by

$$v \mapsto \mathcal{L}v := \frac{\partial v}{\partial t} + \sum_{i=1}^n \frac{\partial v}{\partial x_i} f_i = \frac{\partial v}{\partial t} + \text{grad } v \cdot f.$$

Given a continuously differentiable test function $v \in \mathcal{C}^1([t_0, T] \times X)$, notice that

$$\begin{aligned} v(T, x(T)) - v(t_0, x(t_0)) &= \int_{t_0}^T dv(t, x(t)) = \int_{t_0}^T \dot{v}(t, x(t)) dt \\ &= \int_{t_0}^T \mathcal{L}v(t, x(t)) dt = \langle \mathcal{L}v, \mu \rangle \end{aligned}$$

which can be written more concisely as

$$\langle v, \mu_T \rangle - \langle v, \mu_0 \rangle = \langle \mathcal{L}v, \mu \rangle \quad (6)$$

upon defining respectively the initial and terminal occupation measures

$$\mu_0(dt, dx) := \delta_{t_0}(dt) \delta_{x(t_0)}(dx), \quad \mu_T(dt, dx) := \delta_T(dt) \delta_{x(T)}(dx). \quad (7)$$

Let us define the adjoint linear operator $\mathcal{L}' : \mathcal{C}([t_0, T] \times X)' \rightarrow \mathcal{C}^1([t_0, T] \times X \times U)'$ by the relation $\langle v, \mathcal{L}'\mu \rangle := \langle \mathcal{L}v, \mu \rangle$ for all $\mu \in \mathcal{M}([t_0, T] \times X)$ and $v \in \mathcal{C}^1([t_0, T] \times X)$. More explicitly, this operator can be expressed as

$$\mu \mapsto \mathcal{L}'\mu = -\frac{\partial \mu}{\partial t} - \sum_{i=1}^n \frac{\partial (f_i \mu)}{\partial x_i} = -\frac{\partial \mu}{\partial t} - \text{div } f \mu$$

where the derivatives of measures are understood in the weak sense, i.e. via their action on smooth test functions, and the change of sign comes from integration by parts. Equation (6) can be rewritten equivalently as $\langle v, \mu_T \rangle - \langle v, \mu_0 \rangle = \langle v, \mathcal{L}'\mu \rangle$ and since this equation should hold for all test functions $v \in \mathcal{C}^1([t_0, T] \times X)$, we obtain a linear partial differential equation (PDE) on measures $\mathcal{L}'\mu = \mu_T - \mu_0$ that we write

$$\frac{\partial \mu}{\partial t} + \text{div } f \mu + \mu_T = \mu_0. \quad (8)$$

This linear transport equation is classical in fluid mechanics, statistical physics and analysis of PDEs. It is called the equation of conservation of mass, or the continuity equation, or the advection equation, or Liouville's equation. Under the assumption that the initial data $(t_0, x_0) \in \mathcal{A}$ and the control law $\omega_t \in \mathcal{P}(U)$ are given, the following result can be found e.g. in [14, Theorem 5.34] or [1].

Lemma 1 (Liouville PDE = Cauchy ODE) *There exists a unique solution to the Liouville PDE (8) which is concentrated on the solution of the Cauchy ODE (3), i.e. such that (5) and (7) hold.*

In our context of conic optimization, the relevance of the Liouville PDE (8) is its linearity in the occupation measures μ , μ_0 and μ_T , whereas the Cauchy ODE (3) is nonlinear in the state trajectory $x(t)$.

3.3 Primal LP on measures

The cost in relaxed POCP (2) can therefore be written

$$\int_{t_0}^T \langle l(x(t), \cdot), \omega_t \rangle dt + l_T(x(T)) = \langle l, \mu \rangle + \langle l_T, \mu_T \rangle$$

and we can now define a relaxed optimal control problem as an LP in the cone of non-negative measures:

$$\begin{aligned} p^*(t_0, x_0) &:= \min_{\mu, \mu_T} \langle l, \mu \rangle + \langle l_T, \mu_T \rangle \\ \text{s.t.} & \quad \frac{\partial \mu}{\partial t} + \operatorname{div} f \mu + \mu_T = \delta_{t_0} \delta_{x_0} \\ & \quad \mu \in \mathcal{M}_+([t_0, T] \times X \times U) \\ & \quad \mu_T \in \mathcal{M}_+(\{T\} \times X_T) \end{aligned} \quad (9)$$

where the minimization is w.r.t. the occupation measure μ (which includes the relaxed control ω_t , see (5)) and the terminal measure μ_T , for a given initial measure $\mu_0 = \delta_{t_0} \delta_{x_0}$ which is the right-hand side in the Liouville equation constraint.

Note that in LP (9) the infimum is always attained since the admissible set is (weak-star) compact and the functional is linear. However, since classical trajectories are a particular case of relaxed trajectories corresponding to the choice (5), the minimum in LP (9) is smaller than the minimum in relaxed POCP (2) (this latter one being equal to the infimum in POCP (1), recall Assumption 1), i.e.

$$v^*(t_0, x_0) \geq p^*(t_0, x_0). \quad (10)$$

The following result, due to [15], essentially based on convex duality, shows that there is no gap occurring when considering more general occupation measures than those concentrated on solutions of the ODE.

Lemma 2 *It holds $v^*(t_0, x_0) = p^*(t_0, x_0)$ for all $(t_0, x_0) \in \mathcal{A}$.*

3.4 Dual LP on functions

Primal measure LP (9) has a dual LP in the cone of nonnegative continuous functions:

$$\begin{aligned} d^*(t_0, x_0) &:= \sup_v v(t_0, x_0) \\ \text{s.t.} & \quad l + \frac{\partial v}{\partial t} + \operatorname{grad} v \cdot f \in \mathcal{C}_+([t_0, T] \times X \times U) \\ & \quad l_T - v(T, \cdot) \in \mathcal{C}_+(X_T) \end{aligned} \quad (11)$$

where maximization is with respect to a continuously differentiable function $v \in \mathcal{C}^1([t_0, T] \times X)$ which can be interpreted as a Lagrange multiplier of the Liouville equation in (9).

In general the supremum in dual LP (11) is not attained, and weak duality with the primal LP (9) holds

$$p^*(t_0, x_0) \geq d^*(t_0, x_0)$$

but it can be shown that there is actually no duality gap:

Lemma 3 (No duality gap) *It holds $p^*(t_0, x_0) = d^*(t_0, x_0)$ for all $(t_0, x_0) \in \mathcal{A}$.*

Proof : The proof follows along the same lines as the proof of [7, Theorem 2]. First we observe that $(t_0, x_0) \in \mathcal{A}$ and Assumption 1 imply that $p^*(t_0, x_0)$ is finite. Second, we use the condition that the cone

$$\{(\langle l, \mu \rangle + \langle l_T, \mu_T \rangle, \frac{\partial \mu}{\partial t} + \operatorname{div} f \mu + \mu_T) : \mu \in \mathcal{M}_+([t_0, T] \times X \times U), \mu_T \in \mathcal{M}_+(\{T\} \times X_T)\}$$

is closed in the weak-star topology. This is a classical sufficient condition for the absence of a duality gap in infinite-dimensional LPs, see e.g. [2, Chapter IV, Theorem 7.2]. \square

4 Approximation results

Primal LP (9) and dual LP (11) are infinite-dimensional conic problems. If we want to solve them with a computer, we invariably have to use discretization and approximation schemes. The aim of this section is to derive various approximation results that prove useful when designing numerical methods based on moment-SOS hierarchies.

4.1 Lower bound on value function

First, observe that there always exists an admissible solution for dual LP (11). For example, choose $v(t, x) := a + b(T - t)$ with $a \in \mathbb{R}$ such that $l_T(x) \geq a$ on X_T and $b \in \mathbb{R}$ such that $l(x, u) \geq b$ on $X \times U$. Moreover, by construction, any admissible function for dual LP (11) gives a global lower bound on the value function:

Lemma 4 (Lower bound on value function) *If $v \in \mathcal{C}^1([t_0, T] \times X)$ is admissible for dual LP (11), then $v^* \geq v$ on $[t_0, T] \times X$.*

Proof : If $(t_1, x_1) \notin \mathcal{A}$, then $v^*(t_1, x_1)$ is unbounded and the statement holds because $v(t_1, x_1)$ must be finite. Let $(t_1, x_1) \in \mathcal{A}$ be given. If v is admissible for dual LP (11), then for any admissible trajectory $(x, u) \in \mathcal{L}^\infty([t_0, T]; X \times U)$, starting at $x(t_1) = x_1$, it holds $\int_{t_1}^T (l(x(t), u(t)) + \mathcal{L}v(t, x(t), u(t))) dt = \int_{t_1}^T l(x(t), u(t)) dt + v(T, x(T)) - v(t_1, x_1) \geq 0$ since $l + \mathcal{L} \geq 0$ on $[t_0, T] \times X \times U$. Moreover $\int_{t_1}^T l(x(t), u(t)) dt + l_T(x(T)) \geq \int_{t_1}^T l(x(t), u(t)) dt + v(T, x(T))$ since $l_T - v(T, \cdot) \geq 0$ on X_T . Combining the two inequalities yields $\int_{t_1}^T l(x(t), u(t)) dt + l_T(x(T)) \geq v(t_1, x_1)$ and the expected inequality follows by taking the infimum over admissible trajectories. \square

The relation between the dual LP (11) and the original POCP (1) is given by the following result:

Lemma 5 (Maximizing sequence) *Given $(t_0, x_0) \in \mathcal{A}$, there is a sequence $(v_k)_{k \in \mathbb{N}}$ admissible for the dual LP (11) such that $v^*(t_0, x_0) \geq v_k(t_0, x_0)$ and $\lim_{k \rightarrow \infty} v_k(t_0, x_0) = v^*(t_0, x_0)$.*

Proof : From Lemma 4, it holds $v^*(t_0, x_0) \geq v_k(t_0, x_0)$ for every function v_k admissible for dual LP (11). By Assumption 1 and Lemma 3, it holds $d^*(t_0, x_0) = p^*(t_0, x_0) = v^*(t_0, x_0)$ and hence there exists a maximizing sequence $v_k \in \mathcal{C}^1([t_0, T] \times X)$ for LP (11). \square

4.2 Uniform approximation of the value function along trajectories

In this section, we investigate the properties of maximizing sequences given by Lemma 5, and in particular their convergence to the value function of POCP (1). We first demonstrate the lower semicontinuity of the value of LP (9). This leads to the lower semicontinuity of the value of POCP (1), by considering Assumption 1 and Lemma 2. Note that lower semicontinuity is readily ensured when the set $\{(f(x, u), l(x, u) + a) : u \in U, a \geq 0\}$ is convex in \mathbb{R}^{n+1} for all x , with U compact, see e.g. [13, Section 6.2]. Indeed, in this case, the infimum is attained in POCP (1), and Assumption 1 is readily satisfied.

Lemma 6 (Lower semi-continuity of the value of the measure LP) *The function $(t_0, x_0) \rightarrow p^*(t_0, x_0)$ is lower semicontinuous.*

Proof : We need to show that given a sequence $(t_k, x_k)_{k \in \mathbb{N}}$ such that $\lim_{k \rightarrow \infty} (t_k, x_k) = (t, x) \in \mathbb{R}^{n+1}$, it holds that $\liminf_{k \rightarrow \infty} p^*(t_k, x_k) \geq p^*(t, x)$. Suppose that (t, x) is such that measure LP (9) is feasible. If the left-hand side is not finite, the result holds. If the left-hand side is finite, we can consider, up to taking a subsequence, that $\liminf_{k \rightarrow \infty} p^*(t_k, x_k) = \lim_{k \rightarrow \infty} p^*(t_k, x_k) < \infty$. Since the infimum is attained in measure LP (9), we have a sequence of measures $(\mu_k, \mu_{T_k})_{k \in \mathbb{N}}$ such that $p^*(t_k, x_k) = \langle \mu_k, l \rangle + \langle \mu_{T_k}, l_T \rangle$ and $\frac{\partial \mu_k}{\partial t} + \operatorname{div} f \mu_k + \mu_{T_k} = \delta_{t_k} \delta_{x_k}$. Convergence of (t_k, x_k) to (t, x) implies weak-star convergence of $\delta_{t_k} \delta_{x_k}$ to $\delta_t \delta_x$. Using the same closedness argument as in the proof of Lemma 3, we can consider that, up to a subsequence, μ_k and μ_{T_k} converge to some measures μ and μ_T in the weak-star topology and that $\frac{\partial \mu}{\partial t} + \operatorname{div} f \mu + \mu_T = \delta_t \delta_x$. Hence, we have $\liminf_{k \rightarrow \infty} p^*(t_k, x_k) = \langle \mu, l \rangle + \langle \mu_T, l_T \rangle$ and the pair (μ, μ_T) is feasible for problem $p^*(t, x)$. Therefore $\liminf_{k \rightarrow \infty} p^*(t_k, x_k) \geq p^*(t, x)$ which proves the result when LP (9) is feasible for (t, x) . Using similar arguments, one can show that if (t, x) is such that LP (9) is not feasible, there cannot be infinitely many k such that LP (9) is feasible for (t_k, x_k) . \square

The following result extends the convergence properties of the maximizing sequence.

Theorem 1 (Uniform convergence along relaxed trajectories) *For any sequence $(v_k)_{k \in \mathbb{N}}$ admissible for the dual LP (11), for any solution (x, ω_t) of relaxed POCP (2), and for any $t \in [t_0, T]$, it holds*

$$0 \leq v^*(t, x(t)) - v_k(t, x(t)) \leq v^*(t_0, x_0) - v_k(t_0, x_0) \xrightarrow[k \rightarrow \infty]{} 0.$$

Proof : Let $(x_j, u_j)_{j \in \mathbb{N}}$ be an approximating sequence for (x, ω_t) , whose existence is guaranteed by Assumption 1. For any $j \in \mathbb{N}$, $k \in \mathbb{N}$, and $t \in [t_0, T]$, we have $l_T(x_j(T)) - v_k(T, x_j(T)) + \int_t^T (l(x_j(s), u_j(s)) + \mathcal{L}(s, x_j(s), u_j(s))) ds = l_T(x_j(T)) + \int_t^T l(x_j(s), u_j(s)) ds - v_k(t, x_j(t))$. Both the first term and the integrand are positive in the left-hand side. Therefore, the right-hand side is a decreasing function of t . Moreover, the trajectory is suboptimal, and v_k is a lower bound on the value function. It holds that $0 \leq v^*(t, x_j(t)) - v_k(t, x_j(t)) \leq l_T(x_j(T)) + \int_t^T l(x_j(s), u_j(s)) ds - v_k(t, x_j(t)) \leq l_T(x_j(T)) + \int_{t_0}^T l(x_j(s), u_j(s)) ds - v_k(t_0, x_0)$. Letting j tend to infinity, using the lower semicontinuity of v^* , we conclude that $0 \leq v^*(t, x(t)) - v_k(t, x(t)) \leq \liminf_{j \rightarrow \infty} v^*(t, x_j(t)) - v_k(t, x_j(t)) \leq \lim_{j \rightarrow \infty} (l_T(x_j(T)) + \int_{t_0}^T l(x_j(s), u_j(s)) ds - v_k(t_0, x_0)) = v^*(t_0, x_0) - v_k(t_0, x_0)$. \square

It is important to notice that Theorem 1 holds for any trajectory realizing the minimum of POCP (2) and therefore, for all of them simultaneously. In addition, these trajectories are identified with limiting trajectories of POCP (1) by Assumption 1.

5 Optimal control over a set of initial conditions

Liouville equation (8) is used as a linear equality constraint in POCP (9) with a Dirac right-hand side as an initial condition. However, this right-hand side can be replaced by more general probability measures. The linearity of the constraint allows to extend most of the results of the previous section to this setting. It leads to similar convergence guarantees regarding a (possibly uncountable) set of optimal control problems. These guarantees hold for solutions of a single infinite-dimensional LP.

Suppose that we are given a set of initial conditions $X_0 \subset X$, such that $(t_0, x_0) \in \mathcal{A}$ for every $x_0 \in X_0$. Given a probability measure $\xi_0 \in \mathcal{P}(X_0)$, let

$$\mu_0(dt, dx) = \delta_{t_0}(dt)\xi_0(dx)$$

and consider the following average value

$$\bar{v}^*(\mu_0) := \int_{X_0} v^*(t, x) \mu_0(dt, dx) = \langle v^*, \mu_0 \rangle \quad (12)$$

where v^* is the value of POCP (1). Under Assumption 1, by linearity this value is equal to the value of POCP (9) with μ_0 as the right-hand side of the equality constraint, namely the primal averaged LP

$$\begin{aligned} \bar{p}^*(\mu_0) &:= \min_{\mu, \mu_T} \langle l, \mu \rangle + \langle l_T, \mu_T \rangle \\ \text{s.t.} & \quad \frac{\partial \mu}{\partial t} + \text{div } f \mu + \mu_T = \mu_0 \\ & \quad \mu \in \mathcal{M}_+([t_0, T] \times X \times U) \\ & \quad \mu_T \in \mathcal{M}_+(\{T\} \times X_T) \end{aligned} \quad (13)$$

with dual averaged LP

$$\begin{aligned} \bar{d}^*(\mu_0) &:= \sup_v \langle v, \mu_0 \rangle \\ \text{s.t.} & \quad l + \frac{\partial v}{\partial t} + \text{grad } v \cdot f \in \mathcal{C}_+([t_0, T] \times X \times U) \\ & \quad l_T - v(T, \cdot) \in \mathcal{C}_+(X_T). \end{aligned} \quad (14)$$

The absence of duality gap is justified in the same way as in Lemma 3. Moreover, Lemma 4 also holds, and, as in Lemma 5, we have the existence of maximizing lower bounds v_k such that

$$\lim_{k \rightarrow \infty} \langle v_k, \mu_0 \rangle = \bar{v}^*(\mu_0) = \bar{p}^*(\mu_0) = \bar{d}^*(\mu_0).$$

Intuitively, primal LP (13) models a superposition of optimal control problems. The LP formulation allows to express it as a single program over measures satisfying a transport equation. A relevant question here is the relation between solutions of averaged measure LP (13) and optimal trajectories of the original problem POCP (1). The intuition is that measure solutions of LP (13) represent a superposition of optimal trajectories of the relaxed POCP (2). These trajectories are themselves limiting trajectories of the original POCP (1). The superposition principle of [1, Theorem 3.2] allows to formalize this intuition and to extend the result of Theorem 1 to this setting.

Theorem 2 (Uniform convergence on support of optimal measure) *For any solution (μ, μ_T) of primal averaged LP (13), there are parametrized measures $\xi_t \in \mathcal{P}(X)$ (for the state) and $\omega_t \in \mathcal{P}(U)$ (for the control) such that $\mu(dt, dx, du) = dt \xi_t(dx) \omega_t(du)$, $\mu_0(dt, dx) = \delta_{t_0}(dt) \xi_{t_0}(dx)$ and $\mu_T(dt, dx) = \delta_T(dt) \xi_T(dx)$. In addition, if $(v_k)_{k \in \mathbb{N}}$ is a maximizing sequence for dual averaged LP (14), for any $t \in [t_0, T]$, it holds*

$$0 \leq \int_X (v^*(t, x) - v_k(t, x)) \xi_t(dx) \leq \int_X (v^*(t_0, x_0) - v_k(t_0, x_0)) \xi_0(dx_0) \xrightarrow[k \rightarrow \infty]{} 0.$$

Proof : The decomposition is given by Lemma 3 in [7]. It asserts the existence of a measure $\sigma \in \mathcal{M}(\mathcal{C}([t_0, T], X))$ supported on trajectories admissible for relaxed POCP (2) and such that for any measurable function $w : X \rightarrow \mathbb{R}$, it holds $\int_X w(x) \xi_t(dx) = \int_{\mathcal{C}([t_0, T], X)} w(x(t)) \sigma(dx(\cdot))$. By Assumption 1, all trajectories of the support of σ are pointwise limits of sequences of feasible trajectories of POCP (1). Hence σ -almost all of these sequences must be minimizing sequences for POCP (1), otherwise, that would contradict optimality of (μ, μ_T) . The result follows by discarding the trajectories which are not limits of minimizing sequences. This does not change σ or ξ_t . Theorem 1 applies to σ -almost all these trajectories and we have $0 \leq \int_{\mathcal{C}([t_0, T], X)} (v^*(t, x(t)) - v_k(t, x(t))) \sigma(dx(\cdot)) = \int_X (v^*(t, x) - v_k(t, x)) \xi_t(dx) \leq \int_X (v^*(t_0, x_0) - v_k(t_0, x_0)) \xi_0(dx_0)$. \square

A remarkable practical implication of this result is that maximizing sequences of averaged dual LP (14) provide an approximation to the value function of POCP (1) that is uniform in time and almost uniform in space along limits of optimal trajectories starting from X_0 .

6 Numerical illustrations

In Sections 3 and 5 we reformulated nonlinear optimal control problems as abstract linear conic optimization problems that involve manipulations of measures and continuous functions in their full generality. The results presented in Section 4 are related to properties of minimizing or maximizing elements, or sequences of elements for these problems. From a practical point of view, it is possible to construct these sequences using the same numerical

tools as in static polynomial optimization. On the primal side, this allows to approximate the minimizing elements of measure LP problems with a converging hierarchy of moment SDP problems. On the dual side, we can construct numerically maximizing sequences of polynomial SOS certificates for the continuous function LP problems. The convergence properties that we investigated hold in particular for these solutions of the moment-SOS hierarchy.

This section illustrates convergence properties of the sequence of approximations of value functions computed using moment-SOS hierarchies. We consider simple, but largely spread, optimal control problems for which the value function (or optimal trajectories) are known.

6.1 Uniform approximation along an optimal trajectory

Consider the one-dimensional turnpike POCP analyzed in Section 22.2 of [3]:

$$\begin{aligned}
 v^*(t_0, x_0) := \inf_u & \int_{t_0}^2 (x(t) + u(t)) dt \\
 \text{s.t.} & \dot{x}(t) = 1 + x(t) - x(t)u(t) \\
 & x(t_0) = x_0 \\
 & u(t) \in [0, 3].
 \end{aligned} \tag{15}$$




Figure 1: Optimal trajectory $t \mapsto x^*(t)$ starting at $(t_0, x_0) = (0, 0)$ for the turnpike POCP (15).

For this problem, the infimum is attained at a unique optimal control which is piecewise constant. The optimal trajectory $t \mapsto x^*(t)$ starting at $(t_0, x_0) = (0, 0)$ is presented in Figure 1. The uniform convergence of approximate value functions $t \mapsto v_k(t, x^*(t))$ to




Figure 2: Differences $t \mapsto v^*(t, x^*(t)) - v_k(t, x^*(t))$ between the actual value function and its polynomial approximations of increasing degrees $k = 3, 5, 7, 9$ along the optimal trajectory $t \mapsto x^*(t)$ starting at $(t_0, x_0) = (0, 0)$ for the turnpike POCP (15). We observe uniform convergence along this trajectory, as well as time decrease of the difference, as predicted by the theory.

the true value function $t \mapsto v^*(t, x^*(t))$ along this optimal trajectory, stated by Theorem 1, is illustrated in Figure 2. Moreover, the difference $t \mapsto v^*(t, x^*(t)) - v_k(t, x^*(t))$ is a decreasing function of time as we observed in the proof of Theorem 1.

6.2 Uniform approximation over a set

Consider the classical linear quadratic regulator problem:

$$\begin{aligned}
 v^*(t_0, x_0) := & \inf_u \int_0^1 (10x(t)^2 + u(t)^2) dt \\
 \text{s.t.} & \quad \dot{x}(t) = x(t) + u(t) \\
 & \quad x(t_0) = x_0.
 \end{aligned} \tag{16}$$

For each (t_0, x_0) , the infimum is attained and the value of the problem can be computed by solving a Riccati differential equation. To illustrate Theorem 2 we are interested in the average value (12) for an initial measure μ_0 concentrated at time 0 and uniformly distributed in space in $X_0 = [-1, 1]$. We approximate this value with primal and dual solutions of LPs (13) and (14). The contour lines (in decimal logarithmic scale) of the difference between the true value function $(t, x) \mapsto v^*(t, x)$ and a polynomial approximation of degree 6 $(t, x) \mapsto v_6(t, x)$ is represented in Figure 3. We also show the support of optimal trajectories starting from X_0 . This illustrates the fact that the approximation of the value function is correct in this region, as stated by Theorem 2. It is noticeable




Figure 3: Countour lines (at $0, -1, -2, -3, \dots$) of the decimal logarithm of the difference $(t, x) \mapsto v^*(t, x) - v_6(t, x)$ between the actual value function and its polynomial approximation of degree 6 for LQR POCP (16). The dark area represents the set of optimal trajectories starting from $x_0 \in X_0 = [-1, 1]$ at time $t_0 = 0$. We observe that the difference is smaller in this area, as predicted by the theory.

that this is computed by a single linear program and provides approximation guarantees uniformly over an uncountable set of optimal control problems.

References

- [1] L. Ambrosio. Transport equation and Cauchy problem for non-smooth vector fields. In L. Ambrosio et al. (eds.), *Calculus of variations and nonlinear partial differential equations*. Lecture Notes in Mathematics, Vol. 1927, Springer-Verlag, Berlin, 2008.
- [2] A. Barvinok. *A Course in Convexity*. American Mathematical Society, Providence, NJ, 2002.
- [3] F. Clarke. *Functional analysis, calculus of variations and optimal control*. Springer-Verlag, London, UK, 2013.
- [4] H. O. Fattorini. *Infinite dimensional optimization and control theory*. Cambridge Univ. Press, Cambridge, UK, 1999.
- [5] V. Gaitsgory, M. Quincampoix. Linear programming approach to deterministic infinite horizon optimal control problems with discounting. *SIAM Journal on Control and Optimization*, 48:2480-2512, 2009.

- [6] H. Frankowska, F. Rampazzo. Filippov's and Filippov-Ważewski's theorems on closed domains. *Journal of Differential Equations*, 161(2):449-478, 2000
- [7] D. Henrion, M. Korda. Convex computation of the region of attraction of polynomial control systems. *IEEE Transactions on Automatic Control*, 59(2):297-312, 2014.
- [8] J. B. Lasserre, *Optimisation globale et théorie des moments*. *Comptes Rendus de l'Académie des Sciences Paris, Série I, Mathématique*, 331(11):929-934, 2000.
- [9] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization* 11(3):796–817, 2001.
- [10] 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.
- [11] J. B. Lasserre. *Moments, positive polynomials and their applications*. Imperial College Press, London, UK, 2010.
- [12] D. G. Luenberger. *Optimization by vector space methods*. Wiley, New York, NY, 1969.
- [13] E. Trélat. *Contrôle optimal : théorie et applications*. Vuibert, Paris, 2005.
- [14] C. Villani. *Topics in optimal transportation*. AMS, Providence, RI, 2003.
- [15] R. B. Vinter. Convex duality and nonlinear optimal control. *SIAM Journal on Control and Optimization* 31(2):518-538, 1993.