On the asymptotic behaviour of the value function in optimal control problems

V. Ašković

Sorbonne University, Paris

Work in collaboration with Emmanuel Trélat (Sorbonne University, Paris) and Hasnaa Zidani (INSA, Rouen)

PGMO Days, EDF, Saclay, December 2021

December 1, 2021

Optimal control problem

Consider the following optimal control problem (OCP):

$$\left(\mathcal{P}_{[0,T]}\right)_{x,z} \begin{cases} v(T,x,z) := \inf_{u\in\mathcal{U}} \int_0^T f^0(y(t),u(t))dt \\ \dot{y}(t) = f(y(t),u(t)) \ \ \forall t \in [0,T] \\ y(0) = x, \ \ y(T) = z \end{cases} \tag{1}$$

where f is $C^1(\mathbb{R}^n \times \mathbb{R}^m, \mathbb{R}^n)$ and f^0 is $C^1(\mathbb{R}^n \times \mathbb{R}^m, \mathbb{R})$.

Assume the existence of the optimal triple $(\widehat{y}_T(.), \widehat{\lambda}_T(.), \widehat{u}_T(.))$ coming from the extremal equations of the Pontryagin Maximum Principle (PMP).

Turnpike property

When T is large, $(\widehat{y}_T(.), \widehat{\lambda}_T(.), \widehat{u}_T(.))$ remains "close" to the turnpike $(\bar{y}, \bar{\lambda}, \bar{u})$ except around $t = 0$ and $t = T$.

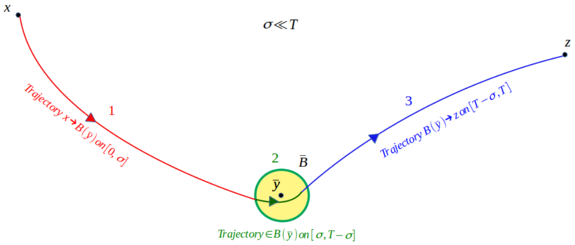

Static (steady) optimization problem (SOP)

$$(\bar{y}, \bar{u}) = \arg\min_{f(y,u)=0} f^0(y, u) \tag{2}$$

$\bar{\lambda}$ is the associated Lagrange multiplier.

1-D Example

$$\min_u \left((u-1)^2 + y^2 \right), \quad \dot{y} = y + 2u, \text{ s.t. } y(0) = x, \quad y(T) = z \quad (3)$$

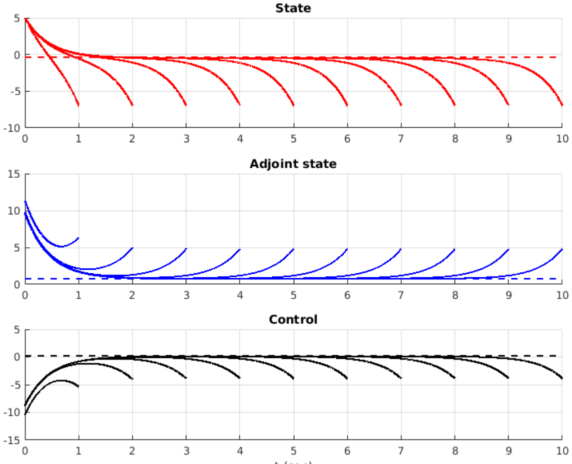

Bibliography

Turnpike property introduced in the 1920 & 1930s. Observed in different types of (OCP) amongst which economical systems.

Optimal control problems over large time intervals, B. Anderson, P. Kokotovic, 1987 $\rightarrow$ the turnpike phenomenon qualitatively observed in the LQ case

Turnpike property in finite-dimensional nonlinear optimal control, E. Trélat, E. Zuazua, 2014 $\rightarrow$ turnpike property locally characterized through an exponential inequality.

Linear turnpike theorem, E. Trélat, 2020 $\rightarrow$ exponential inequality extended globally under the strict dissipativity assumption.

On turnpike and dissipativity properties of continuous time optimal control problems, T. Faulwasser, M. Korda, N. Jones, D. Bonvin, 2021 $\rightarrow$ relation between the dissipativity and the measure turnpike

"Intuition" of the value function expansion

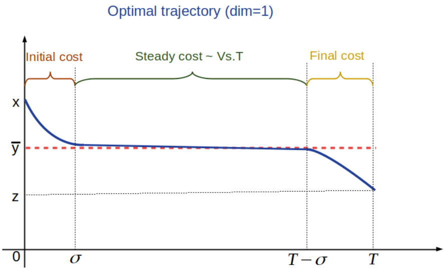

$$v(T,x,z) \underset{T\to+\infty}{\approx} V_s.T + \text{Initial cost}(x) + \text{Final cost}(z) \qquad (4)$$

What has been done so far and our contribution

On the turnpike property and the long time behavior of the HJ equation, C. Esteve, H. Kouhkouh, D. Pighin, E. Zuazua, June 2020 $\rightarrow$ asymptotic expansion of the value associated to the LTI dynamics with running quadratic cost and terminal cost. The proof relies on the exponential turnpike inequality.

What has been done so far and our contribution

On the turnpike property and the long time behavior of the HJ equation, C. Esteve, H. Kouhkouh, D. Pighin, E. Zuazua, June 2020 $\rightarrow$ asymptotic expansion of the value associated to the LTI dynamics with running quadratic cost and terminal cost. The proof relies on the exponential turnpike inequality.

We precise the expansion in the LQ case *intrinsically* that's to say without relying on the exponential turnpike inequality.

What has been done so far and our contribution

On the turnpike property and the long time behavior of the HJ equation, C. Esteve, H. Kouhkouh, D. Pighin, E. Zuazua, June 2020 $\rightarrow$ asymptotic expansion of the value associated to the LTI dynamics with running quadratic cost and terminal cost. The proof relies on the exponential turnpike inequality.

We precise the expansion in the LQ case *intrinsically* that's to say without relying on the exponential turnpike inequality.

We generalize the result to the nonlinear case for a general cost, relying on the strict dissipativity property, which is a priori, a weaker assumption than the exponential turnpike.

Main result:

Under appropriate assumptions*, we have the following expansion:

$$v(T,x,z) \underset{T\to+\infty}{=} V_s.T + v_f(x) + v_b(z) + o(1) \tag{5}$$

where $V_s := f^0(\bar{y},\bar{u})$, $w(y,u) := f^0(y,u) - f^0(\bar{y},\bar{u})$ and:

$$(\mathcal{P}_{\infty f})_x \begin{cases} v_f(x) := \min\limits_{u(.)} \displaystyle\int_0^{+\infty} w(y(t),u(t))dt \\ \dot{y}(t) = f(y(t),u(t)),\ y(0) = x \end{cases} \tag{6}$$

$$(\mathcal{P}_{\infty b})_z \begin{cases} v_b(z) := \min\limits_{u(.)} \displaystyle\int_0^{+\infty} w(y(t),u(t))dt \\ \dot{y}(t) = -f(y(t),u(t)),\ y(0) = z \end{cases} \tag{7}$$

Strict dissipativity

The dissipativity concept was introduced by Willems in "Dissipative Dynamical Systems -Part I and II", 1972

Definition:

The family of (OCP) $\left(\mathcal{P}_{[0,T]}\right)_{x,z}$ indexed by T is strictly dissipative at $(\bar{y},\bar{u})$ with respect to the supply rate function $w(.)$ if there exists a bounded function $S:\mathbb{R}^n \mapsto \mathbb{R}$ and a class $\mathcal{K}$ function $\alpha(.)$ s.t. for any admissible couple $(y(.),u(.))$ one has:

$$S(x) + \int_0^T w(y(t),u(t))dt \geqslant S(z) + \int_0^T \alpha\left(\left\| \begin{array}{c} y(t)-\bar{y} \\ u(t)-\bar{u} \end{array} \right\|\right) dt \tag{8}$$

for any T large enough.

Physical interpretation & Example

Systems that check (8) are called *(strict) dissipative* systems.

They're of particular interest in engeenering and physics.

The main idea behind this is that dissipative systems have certain input-output properties related to the conservation, dissipation and transport of energy.

Example in mechanics
Consider a 1-D mechanical system with a mass, a spring and a damper. The equation of motion is:

$$m\ddot{x} + d\dot{x} + kx = f(t) \text{ with } x(0) = x_0,\ \dot{x}(0) = \dot{x}_0$$

Example in mechanics

m is the mass, d the damper constant, k the stiffness of the spring and f the force acting on the mass.

The energy of the system $V(x,\dot{x}) := \frac{1}{2}kx^2 + \frac{1}{2}m\dot{x}^2$ checks:

$$\underbrace{V(x(T),\dot{x}(T))}_{\text{energy at time T}} = \underbrace{V(x(0),\dot{x}(0))}_{\text{initial energy}} + \underbrace{\int_0^T f(t).\dot{x}(t)dt}_{\text{externally supplied energy}} - \underbrace{\int_0^T d.\dot{x}^2(t)dt}_{\text{dissipated energy}} \tag{9}$$

Idea of the proof

From strict dissipativity assumption:

$$\frac{1}{T}\int_0^T \alpha\left(\left\| \begin{matrix} \widehat{y}_T(t)-\bar{y} \\ \widehat{u}_T(t)-\bar{u} \end{matrix} \right\|\right) dt \leqslant \frac{S(x)-S(z)}{T} + \frac{1}{T}\int_0^T w(\widehat{y}_T(t),\widehat{y}_T(t))dt \longrightarrow 0 \text{ as } T\to+\infty \tag{10}$$

Idea of the proof

From strict dissipativity assumption:

$$\frac{1}{T}\int_0^T \alpha\left(\left\| \begin{array}{c} \widehat{y}_T(t)-\bar{y} \\ \widehat{u}_T(t)-\bar{u} \end{array} \right\|\right) dt \leqslant \frac{S(x)-S(z)}{T} + \frac{1}{T}\int_0^T w(\widehat{y}_T(t),\widehat{y}_T(t))dt \longrightarrow 0 \text{ as } T\to+\infty \tag{10}$$

$$\exists t(T)\in[0,T] \text{ s.t. } \begin{pmatrix} \widehat{y}_T(t(T)) \\ \widehat{u}_T(t(T)) \end{pmatrix} \longrightarrow \begin{pmatrix} \bar{y} \\ \bar{u} \end{pmatrix} \tag{11}$$

Idea of the proof

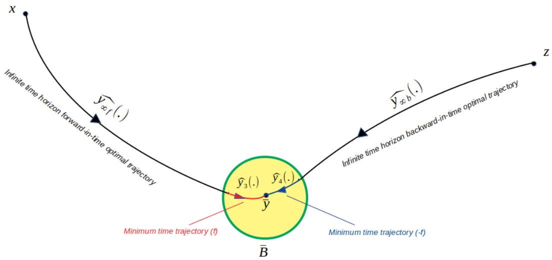

Idea of the proof

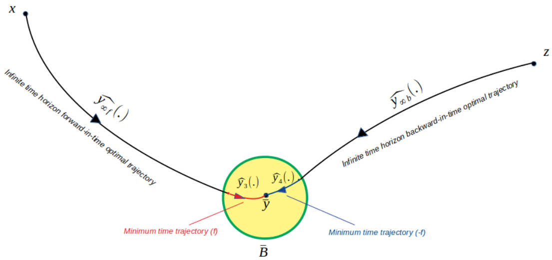

The optimal trajectory is a concatenation of three trajectories:

$$\underbrace{\widehat{y}_T(t)}_{\text{cost}\to v(T,x,z)-V_s.T} \underset{T\to+\infty}{\approx} \underbrace{\widehat{y}_{\infty f}(t)}_{\text{cost}\to v_f(x)} + \underbrace{\widehat{y}_{\infty b}(T-t)}_{\text{cost}\to v_b(z)} + \underbrace{\widehat{y}_3(t)+\widehat{y}_4(T-t)}_{\text{cost}\to\epsilon<<1} \tag{12}$$

Assumptions for the main result

(H_1): C^1 regularity of the cost f^0 and the dynamics f.

(H_2): Local controlabillity at $(\bar{y}, \bar{u})$ and reachability of the turnpike.

(H_3): Uniqueness of the turnpike $(\bar{y}, \bar{\lambda}, \bar{u})$.

(H_4): Boundedness of the admissible trajectories and controls uniformly with respect to T ("compact world")

(H_5): Finiteness of the costs associated to $(\mathcal{P}_{\infty f})_x$ and $(\mathcal{P}_{\infty b})_z$

(H_6): Uniqueness of the optimal trajectories to $\left(\mathcal{P}_{[0,T]}\right)_{x,z}$, $(\mathcal{P}_{\infty f})_x$ and $(\mathcal{P}_{\infty b})_z$

(H_7): Strict dissipativity property

The Linear Quadratic case

Under the Kalman condition on (A, B):

$$v(T, x, z) \underset{T\to+\infty}{=} V_s.T + \underbrace{F(x) + \langle\bar{\lambda}, x - \bar{y}\rangle}_{v_f(x)} + \underbrace{B(z) + \langle\bar{\lambda}, \bar{y} - z\rangle}_{v_b(z)} + o(1)$$

where:

$$(\mathcal{P}^{lin}_{\infty f}) \begin{cases} F(x) := \min\limits_{u_1(.)} \dfrac{1}{2}\displaystyle\int_0^{+\infty} \left(\|u_1(t)\|^2 + \|y_1(t)\|^2\right) dt \\ \dot{y}_1 = Ay_1 + Bu_1, \quad y_1(0) = x - \bar{y} \end{cases}$$

$$(\mathcal{P}^{lin}_{\infty b}) \begin{cases} B(z) := \min\limits_{u_2(.)} \dfrac{1}{2}\displaystyle\int_0^{+\infty} \left(\|u_2(t)\|^2 + \|y_2(t)\|^2\right) dt \\ \dot{y}_2 = -Ay_2 - Bu_2, \quad y_2(0) = z - \bar{y} \end{cases}$$

Remarks & Perspectives

In practice, the strict dissipativity is the less intuitive assumption.

Remarks & Perspectives

In practice, the strict dissipativity is the less intuitive assumption.

In the LQ case, it is automatically and globally verified. This is a consequence of the quadratic nature of the Lagrangian associated to the steady optimization problem.

Remarks & Perspectives

In practice, the strict dissipativity is the less intuitive assumption.

In the LQ case, it is automatically and globally verified. This is a consequence of the quadratic nature of the Lagrangian associated to the steady optimization problem.

In the general nonlinear case, the local strict dissipativity can be seen as a consequence of some properties of the augmented Lagrangian around the turnpike.

Remarks & Perspectives

In practice, the strict dissipativity is the less intuitive assumption.

In the LQ case, it is automatically and globally verified. This is a consequence of the quadratic nature of the Lagrangian associated to the steady optimization problem.

In the general nonlinear case, the local strict dissipativity can be seen as a consequence of some properties of the augmented Lagrangian around the turnpike.

We continue our work in order to obtain convergence results at the control & trajectory level.

Remarks & Perspectives

In practice, the strict dissipativity is the less intuitive assumption.

In the LQ case, it is automatically and globally verified. This is a consequence of the quadratic nature of the Lagrangian associated to the steady optimization problem.

In the general nonlinear case, the local strict dissipativity can be seen as a consequence of some properties of the augmented Lagrangian around the turnpike.

We continue our work in order to obtain convergence results at the control & trajectory level.

To be followed...

Thank you for your attention!

Exponential turnpike inequality

Under appropriate assumptions, there exists $C > 0$, $\nu > 0$ such that if $T > T_0$, then:

$$\|\widehat{y}_T(t) - \bar{y}\| + \|\widehat{\lambda}_T(t) - \bar{\lambda}\| + \|\widehat{u}_T(t) - \bar{u}\| \leqslant C\left(e^{-\nu.t} + e^{-\nu.(T-t)}\right) \tag{13}$$

Turnpike property in finite-dimensional nonlinear optimal control, E. Trélat, E. Zuazua, 2014

Linear turnpike theorem, E. Trélat, 2020