Systems & Control Letters 106 (2017) 79–86

Contents lists available at ScienceDirect

Systems & Control Letters

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

Cost-to-travel functions: A new perspective on optimal and model predictive control

Boris Houska [a,*], Matthias A. Müller [b]

[a] School of Information Science and Technology, ShanghaiTech University, China
[b] Institute for Systems Theory and Automatic Control, University of Stuttgart, Germany

ARTICLE INFO

Article history:
Received 1 October 2016
Received in revised form 13 April 2017
Accepted 13 June 2017

Keywords:
Optimal control
Model predictive control
Dissipativity

ABSTRACT

This paper concerns a class of functions, named cost-to-travel functions, which find applications in model-based control. For a given (potentially nonlinear) control system, the cost-to-travel function associates with any given start and end point in the state space and any given travel duration the minimum economic cost of the associated point-to-point motion. Cost-to-travel functions are a generalization of cost-to-go functions, which are often used in the context of dynamic programming as well as model predictive control. We discuss the properties of cost-to-travel functions, their relations to existing concepts in control such as dissipativity, but also a variety of control-theoretic applications of this function class. In particular, we discuss how cost-to-travel functions can be used to analyze the properties of economic model predictive control with return constraints.

1. Introduction

Optimal control as well as model predictive control theory are well-established fields of research that are covered by many modern textbooks [1–3]. The foundations of optimal control theory go back to the work of Bellman [4], Pontryagin [5], and Filippov [6]. Bellman's principle of optimality is the basis for dynamic programming algorithms, which construct the optimal value function, also called "cost-to-go function", of an optimal control problem by a backward recursion [4,7]. Cost-to-go functions are the solution of the Hamilton–Jacobi–Bellman equation [7] and can be used as a starting point for deriving Pontryagin's maximum principle [1,5].

Cost-to-go functions can also be considered as an important tool for analyzing closed-loop systems. For example, quadratic cost-to-go functions can be used to derive closed-loop optimal control laws for linear systems leading to the famous linear quadratic regulator (LQR) [2]. Moreover, modern model predictive control theory [3] frequently uses cost-to-go functions as Lyapunov function candidates that can be used to derive a number of stability results for model predictive control with tracking objectives under additional controllability assumptions with or without terminal costs and constraints [8,9].

In recent years, so-called *economic* MPC schemes have received significant attention. Here, the primary control objective is not the stabilization of an a priori given setpoint or trajectory as in tracking MPC, but the optimization of a general performance criterion [10–12]. In such economic MPC schemes, certain dissipativity conditions have turned out to play a crucial role, both for classifying the optimal operating behavior for a given system as well as for establishing closed-loop convergence and performance guarantees [10,11,13,14]. The notion of dissipativity goes back to Willems [15] and it was already shown in these early results [15] (compare also [16]) that dissipativity can – at least in principle – be characterized via cost-to-go functions of suitably defined optimal control problems.

The goal of this letter is to propose a new type of functions, named "cost-to-travel functions". Cost-to-travel functions can be viewed as a generalization of cost-to-go functions, which open new perspectives on the above mentioned existing control theory results. The main contributions can be outlined as follows.

- Section 3.1 introduces a functional equation for cost-to-travel functions, which is summarized in Proposition 1 and which can be used to construct double-sided dynamic programming algorithms for solving time-autonomous optimal control problems to global optimality. The run-time complexity of double-sided dynamic programming is proportional to the logarithm of the time horizon of the given optimal control problem.
- Theorem 1 exploits the properties of cost-to-travel functions in order to provide an alternative, elegant proof of a result that has originally been established in [17], namely that for any linear system with strictly convex stage-cost and convex constraints every optimal periodic orbit is a steady-state.

* Corresponding author.
E-mail addresses: borish@shanghaitech.edu.cn (B. Houska), matthias.mueller@ist.uni-stuttgart.de (M.A. Müller).

http://dx.doi.org/10.1016/j.sysconle.2017.06.005

- Theorem 2 establishes an alternative characterization of a certain dissipativity condition for general (potentially nonlinear) control systems, which is crucial in economic MPC.
- Theorem 3 and the associated Lemma 1 establish (under a mild controllability assumption) necessary and sufficient conditions to check whether a given optimal m-periodic orbit of a time-autonomous control system is optimal among all N-periodic orbits with $N \in \mathbb{N}$.
- Theorem 4 establishes conditions under which local asymptotic orbital stability of economic MPC with return constraints can be verified. The corresponding proof uses the fact that the closed-loop state-trajectories of economic MPC with return constraints are equal to selected iterates of a block-coordinate descent method applied to minimize a particularly defined "cost-of-a-round-trip function".

Notation.

We use the symbols $\mathbb{R}$ and $\mathbb{N}$ to denote the set of real numbers and strictly positive integers, respectively. The notation $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ denotes the set of positive integers including 0.

2. Cost-to-travel function

Let $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$ be a given continuous right-hand side function, $l : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}$ a given continuous stage cost, $\mathbb{X} \subseteq \mathbb{R}^{n_x}$ a closed set modeling state constraints, and $\mathbb{U} \subseteq \mathbb{R}^{n_u}$ a compact set modeling control constraints. The cost-to-travel function

$$V : \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \times \mathbb{N}_0 \to \mathbb{R} \cup \{\infty\}$$

is defined as the optimal value of the optimization problem

$$V(a,b,N) = \min_{x,u} \sum_{k=0}^{N-1} l(x_k,u_k) \quad \text{s.t.} \begin{cases} \forall k \in \{0,1,\ldots,N-1\}, \\ x_{k+1} = f(x_k,u_k) \\ x_k \in \mathbb{X}, \ u_k \in \mathbb{U} \\ x_0 = a, \ b = x_N \in \mathbb{X} \end{cases} \tag{1}$$

for all $a, b \in \mathbb{R}^{n_x}$ and all strictly positive integers $N \in \mathbb{N}$ as well as

$$V(a,b,0) = \begin{cases} 0 & \text{if } a = b \in \mathbb{X} \\ \infty & \text{otherwise.} \end{cases}$$

Here, x_k and u_k denote the state and control of the discrete-time system f at time k. Since the functions f and l are continuous, the set $\mathbb{X}$ closed, and the set $\mathbb{U}$ compact, Problem (1) either has a minimizer or is infeasible. If Problem (1) is infeasible, we set $V(a,b,N) = \infty$ such that V is well-defined for all points $a, b \in \mathbb{R}^{n_x}$ and all $N \in \mathbb{N}_0$. Notice that this implies that N-step controllability (under constraints) from a to b is equivalent to requiring $V(a,b,N) < \infty$. The function $V(a,b,N)$ can be interpreted as the minimum cost that is needed to bring the system from state a to state b in time N.

2.1. Cost-to-go function

The cost-to-travel function V is closely related to the so-called cost-to-go function $J : \mathbb{R}^{n_x} \times \mathbb{N}_0 \to \mathbb{R}$. For the case that $m : \mathbb{R}^{n_x} \to \mathbb{R}$ is a given continuous terminal cost, the cost-to-go function is defined as

$$J(a,N) = \min_b V(a,b,N) + m(b)$$

for all $a \in \mathbb{R}^{n_x}$ and all $N \in \mathbb{N}_0$. Here, the motivation for introducing the end cost m depends on the context. In standard optimal control m is used to model cost contributions that depend on the terminal state only, while in model predictive control the motivation for introducing m is to approximate the infinite horizon cost and to enforce stability of the resulting closed-loop system [3,18]. Here, the infinite horizon cost-to-go function is typically defined as

$$J_\infty(a) = \lim_{N\to\infty} \min_b V(a,b,N), \tag{2}$$

if this limit exists.

2.2. Guiding example

Throughout this paper, the guiding example

$$\begin{aligned} f(x,u) &= x + u, & \mathbb{X} &= [-2,2], \\ l(x,u) &= x^2 - |u|, & \mathbb{U} &= [-4,4] \end{aligned} \tag{3}$$

is used in order to illustrate some of the theoretical developments. Notice that the one-step cost-to-travel function has the explicit form

$$V(a,b,1) = \begin{cases} a^2 - |a-b| & \text{if } a,b \in [-2,2] \\ \infty & \text{otherwise.} \end{cases} \tag{4}$$

Moreover, we have

$$\lim_{N\to\infty} \min_b V(a,b,N) = \begin{cases} -\infty & \text{if } a \in [-2,2] \\ +\infty & \text{otherwise} \end{cases}$$

for this example, i.e., the standard cost-to-go function does not take finite values without further modification.

3. Properties of cost-to-travel functions

3.1. Functional equation

The cost-to-travel function V satisfies "Bellman's principle of optimality" [4], which yields a double-sided dynamic programming recursion:

Proposition 1. *The cost-to-travel function V satisfies a functional recursion of the form*

$$V(a,b,N+M) = \min_c V(a,c,N) + V(c,b,M) \tag{5}$$

for all $a, b \in \mathbb{R}^{n_x}$ and all $N, M \in \mathbb{N}_0$.

Proof. The cost of traveling from a to b via c in $N+M$ steps is always larger than or equal to the cost to travel from a to b without necessarily visiting c, but we can always find a way-point c for which equality holds [4]. □

Example 1. For the guiding example (3), the functional equation (5) yields

$$\begin{aligned} &V(a,b,2) \\ &= \min_c V(a,c,1) + V(c,b,1) \\ &= \min_c a^2 + c^2 - |a-c| - |b-c| \\ &= a^2 + \min\{-|a-1| - |b-1| + 1, -|a| - |b|, \\ &\quad -|a+1| - |b+1| + 1\} \end{aligned} \tag{6}$$

for all $a, b \in [-2,2]$.

Notice, that if the time horizon of the discrete time optimal control problem

$$\min_{x,u} \sum_{k=0}^{N-1} l(x_k,u_k) + m(x_N) \quad \text{s.t.} \begin{cases} \forall k \in \{0,1,\ldots,N-1\}, \\ x_{k+1} = f(x_k,u_k) \\ x_0 = \hat{x}_0 \\ x_k \in \mathbb{X}, \ u_k \in \mathbb{U} \end{cases} \tag{7}$$