Annual Reviews in Control 35 (2011) 13–33

Contents lists available at ScienceDirect

Annual Reviews in Control

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

Lyapunov functions and discontinuous stabilizing feedback ☆

Francis Clarke

Université de Lyon, Institut Camille Jordan, 69622 Villeurbanne, France

ARTICLE INFO

Article history:
Received 30 September 2010
Accepted 12 February 2011

Keywords:
Controllability
Discontinuous control
Feedback
Nonlinear theory
Stabilization

ABSTRACT

We study the controllability and stability of control systems that are nonlinear, and for which, for whatever reason, linearization fails. We begin by motivating the need for two seemingly exotic tools: nonsmooth control-Lyapunov functions, and discontinuous feedbacks. With the aid of nonsmooth analysis, we build a theory around these tools. We proceed to apply it in various contexts, focusing principally on the design of discontinuous stabilizing feedbacks.

1. Introduction

Our interest centers throughout on the standard control system

$$x'(t) = f(x(t), u(t)) \text{ a.e.}, \quad u(t) \in U \text{ a.e.}, \tag{$*$}$$

where the dynamics function $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ and the control set $U \subset \mathbb{R}^m$ are given, and 'a.e.' is the abbreviation of 'almost everywhere'. A *control* on some interval $[a, b]$ of interest refers to a measurable function $u(\cdot)$ defined on $[a, b]$ and having values in U. By a *trajectory* of the system $(*)$ we mean (as usual) an absolutely continuous state function $x : [a, b] \to \mathbb{R}^n$ corresponding to some choice of control $u(\cdot)$.

It is assumed throughout that f is continuous, U is compact, and f is locally Lipschitz with respect to the state variable in the following sense: for every bounded subset $S \subset \mathbb{R}^n$, there exists $K = K_f(S)$ such that

$$|f(x, u) - f(y, u)| \leqslant K_f|x - y| \quad \forall x, y \in S, \ u \in U. \tag{1}$$

We remark that this Lipschitz behavior is automatically present if f is continuously differentiable, but differentiability of f is irrelevant to our discussion. Much more to the point are the assumptions that are *not* being made: f is not linear, U is not simply 'large enough' to be effectively ignored.

The central issue under discussion will be the convergence of state trajectories $x(t)$ to an equilibrium, which we take to be the origin: stability, controllability, and feedback stabilization. Stabilization to the origin is a simple representative of various other objectives that can be treated by the techniques that we shall describe. (Stabilization to other target sets will also be involved later.)

One way to steer trajectories to zero is to invent a cost whose minimization will have that effect. (Indeed, in a certain sense, this is rather close to being the only effective strategy that we know.) The positive features of such an approach, as well as certain inherent difficulties which arise when we employ it, are well illustrated by what is called the *dynamic programming* technique in optimal control. It will furnish us with valuable insight into our stabilization problem, and provide guidance about the mathematical tools needed.

1.1. Dynamic programming and minimal time

The *minimal-time problem* refers to finding a trajectory of $(*)$ that reaches the origin as quickly as possible from a given initial point α. Thus we seek the least $T \geqslant 0$ admitting a control function $u(\cdot)$ on $[0, T]$ having the property that the resulting trajectory x with $x(0) = \alpha$ satisfies $x(T) = 0$. The dynamic programming approach centers upon the *minimal-time function* $T(\cdot)$, defined on $\mathbb{R}^n$ as follows: $T(\alpha)$ is the least time T defined above.

The *principle of optimality* makes two observations about $T(\cdot)$. The first of these is that, for any trajectory $x(\cdot)$ beginning at α, for any two times s,t with $0 \leqslant s < t$, we have

$$T(x(s)) \leqslant T(x(t)) + t - s. \tag{2}$$

This reflects the fact that, starting at the point $x(s)$, we may choose the two-step strategy of following the trajectory x until time t, and then proceeding optimally from the point $x(t)$ to the origin. The time required for this two-step strategy is the right side of (2); the inequality holds because there may be a better strategy beginning from $x(s)$.

☆ Chaire en théorie mathématique du contrôle, Institut universitaire de France. Plenary talk given at the IFAC Conference on Nonlinear Control Systems (NOLCOS), Bologna, September 2010.

E-mail address: clarke@math.univ-lyon1.fr

1367-5788/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.arcontrol.2011.03.001

The second observation is that equality holds in (2) if x is a trajectory that joins α to the origin in minimal time; that is, if $x(T) = 0$ for $T = T(\alpha)$. This reflects the fact that when x is a minimal-time trajectory, there is no better strategy than the two-step one described above. Combining these two observations, we find that, for any trajectory $x(\cdot)$, the function $t \mapsto T(x(t)) + t$ is nondecreasing; it is constant when x is a minimal-time trajectory.

Since $t \mapsto T(x(t)) + t$ is nondecreasing, we expect to have

$$\langle \nabla T(x(t)), x'(t)\rangle + 1 \ \geqslant \ 0,$$

with equality when $x(\cdot)$ is an optimal trajectory. The possible values of $x'(t)$ for a trajectory being precisely the elements of the set $f(x(t), U)$, we arrive at

$$\min_{u\in U} \langle \nabla T(x), f(x,u)\rangle + 1 = 0. \tag{3}$$

We define the (lower) *Hamiltonian function h* as follows:

$$h(x,p) := \min_{u\in U}\langle p, f(x,u)\rangle. \tag{4}$$

In terms of h, the partial differential Eq. (3) above reads

$$h(x, \nabla T(x)) + 1 = 0, \tag{5}$$

a special case of the *Hamilton–Jacobi equation*.

We have now reached the first stage in the dynamic programming approach: solve the Hamilton–Jacobi equation (5), together with the boundary condition $T(0) = 0$, to find $T(\cdot)$. How will this help us find minimal-time trajectories?

To answer this question, we recall that an optimal trajectory is such that equality holds in (3). This suggests the following procedure: for each x, let $k(x)$ be a point in U satisfying

$$\min_{u\in U} \langle \nabla T(x), f(x,u)\rangle = \langle \nabla T(x), f(x, k(x))\rangle = -1. \tag{6}$$

Then, if we construct $x(\cdot)$ via the initial-value problem

$$x'(t) = f(x(t), k(x(t))), \quad x(0) = \alpha, \tag{7}$$

we obtain a minimum-time trajectory (from α).

Let us see why this so: if $x(\cdot)$ satisfies (7), then, in light of (6), we have

$$\begin{aligned}(d/dt)T(x(t)) &= \langle \nabla T(x(t)), x'(t)\rangle \\ &= \langle \nabla T(x(t)), f(x(t), k(x(t)))\rangle = -1.\end{aligned}$$

Integrating, we find

$$T(x(t)) = T(\alpha) - t,$$

which implies that at $\tau = T(\alpha)$, we have $T(x(\tau)) = 0$, whence $x(\tau) = 0$ (since T is zero only at the origin). Therefore $x(\cdot)$ is a minimal-time trajectory.

This second stage of the dynamic programming approach has provided a feedback $k(\cdot)$ which, from *any* initial value α, generates via (7) a minimal-time trajectory; k constitutes what can be considered the ultimate solution to our problem: an *optimal feedback synthesis*.

We remark that the Hamilton–Jacobi equation (5) has another use, when we know that it has a unique solution $T(\cdot)$ satisfying $T(0) = 0$ (namely, the minimal-time function). We refer to the *verification method* in optimal control (see for example Clarke (1989)). It would work here as follows: suppose we have formulated a conjecture that, for each α, a certain trajectory x_α is a minimal-time one from the initial condition α. We proceed to calculate $T(\alpha)$ (provisionally) based on this conjecture; that is, by setting $T(\alpha)$ equal to the time required for x_α to join α to 0. Then, if the resulting function T satisfies (5), our conjecture is verified (since, by uniqueness, T must then coincide with the minimal-time function). If T fails to satisfy (5), then our conjecture is certainly false (and the way in which (5) fails may help us amend it).

We now rain on this parade by pointing out that there are serious obstacles to rigorously justifying the route that we have just outlined. There is, to begin with, the issue of *controllability*: Is it always possible to steer α to 0 in finite time? And if this holds, do minimal-time trajectories exist? Even if this is true, how do we know that $T(\cdot)$ is differentiable? If this fails to be the case, then we shall need to replace the gradient ∇T used above by some suitably generalized derivative. Next, we would have to examine anew the argument that led to the Hamilton–Jacobi equation (5), which itself will require reformulation in some way that allows for nonsmooth solutions. Will the Hamilton–Jacobi equation generalized in such a way admit T as the unique solution?

Assuming that all this can be done, the second stage above offers fresh difficulties of its own. Even if T is smooth, there is in general no *continuous* function $k(\cdot)$ satisfying (6) for each x. When k is discontinuous, the classical concept of 'solution' to (7) is inappropriate; what solution concept should we use instead? Would optimal trajectories still result?

That these difficulties are real, and indeed that they arise in the simplest problems, can be illustrated by the following example, familiar from any introductory text in optimal control.

The double integrator. This refers to the system $x'' = u$, or, in terms of the standard formulation $(*)$:

$$x'(t) = y(t), \ \ y'(t) = u(t), \ \ u(t) \in [-1, +1]. \tag{8}$$

Thus $n = 2$, $m = 1$, and the dynamics are linear. It is not difficult to show that all initial points $(x(0), y(0)) = (\alpha, \beta)$ are controllable to the origin in finite time; existence theory tells us that minimal-time trajectories exist. The Maximum Principle helps us to identify them: they turn out to be bang–bang with at most one switch between +1 and -1. We can then calculate the minimal-time function $T(\cdot)$:

$$T(\alpha,\beta) = \begin{cases} -\beta + \sqrt{2\beta^2 - 4\alpha} & \text{when } (\alpha,\beta) \text{ is left of } S \\ +\beta + \sqrt{2\beta^2 + 4\alpha} & \text{when } (\alpha,\beta) \text{ is right of } S, \end{cases}$$

where the *switching curve* S in the $x - y$ plane is given by $y^2 = 2|x|$; see Fig. 1. The resulting function $T(\cdot)$ is seen to be continuous, but it fails to be differentiable or even locally Lipschitz along the switching curve. The optimal feedback synthesis consists of taking ($k = -1$ to the right or on the upper branch) of S, and $k = +1$ otherwise.

We see therefore that our doubts correspond to real difficulties, and they explain why the dynamic programming approach to optimal control, very prominent in the 1950s and 60s, is now frequently ignored in engineering texts, or else relegated to a heuristic role, perhaps in exercises. In fact, however, the difficulties have now been successfully and rigorously resolved, through the use of nonsmooth analysis, viscosity solutions, and discontinuous feedbacks. These very same tools will play a central role in the stabilization issue, which we turn to now.

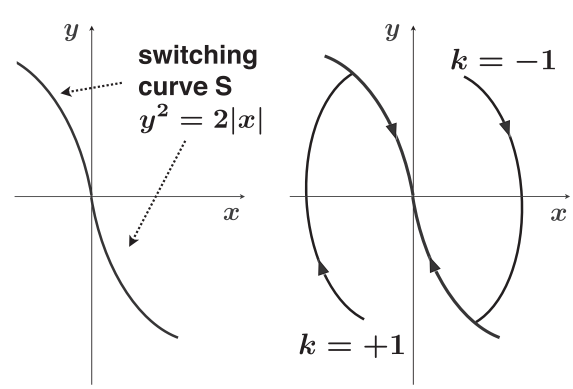

Fig. 1. The double integrator: switching curve and optimal synthesis.