Stability and Feedback Stabilization

EDUARDO D. SONTAG
Department of Mathematics, Rutgers University,
New Brunswick, USA

Article Outline

Glossary

Stability A globally asymptotically stable equilibrium is a state with the property that all solutions converge to this state, with no large excursions.

Stabilization A system is stabilizable (with respect to a given state) if it is possible to find a feedback law that renders that state a globally asymptotically stable equilibrium.

Lyapunov and control-Lyapunov functions A control-Lyapunov functions is a scalar function which decreases along trajectories, if appropriate control actions are taken. For systems with no controls, one has a Lyapunov function.

Definition of the Subject

The problem of stabilization of equilibria is one of the central issues in control. In addition to its intrinsic interest, it represents a first step towards the solution of more complicated problems, such as the stabilization of periodic orbits or general invariant sets, or the attainment of other control objectives, such as tracking, disturbance rejection, or output feedback, all of which may be interpreted as requiring the stabilization of some quantity (typically, some sort of "error" signal). A very special case, when there are no inputs, is that of stability.

Introduction

This article discusses the problem of stabilization of an equilibrium, which we take without loss of generality to be the origin, for a finite-dimensional system $\dot{x} = f(x, u)$. The objective is to find a *feedback law* $u = k(x)$ which renders the origin of the "closed-loop" system $\dot{x} = f(x, k(x))$ globally asymptotically stable. The problem of stabilization of equilibria is one of the central issues in control. In addition to its intrinsic interest, it represents a first step towards the solution of more complicated problems, such as the stabilization of periodic orbits or general invariant sets, or the attainment of other control objectives, such as tracking, disturbance rejection, or output feedback, all of which may be interpreted as requiring the stabilization of some quantity (typically, some sort of "error" signal). A very special case (when there are no inputs u) is that of stability.

After setting up the basic definitions, we consider *linear* systems. Linear systems are widely used as models for physical processes, and they also play a major role in the general theory of local stabilization. We briefly review pole assignment and linear-quadratic optimization as approaches to obtaining feedback stabilizers.

In general, there is a close connection between the existence of continuous stabilizing feedbacks and smooth *control-Lyapunov functions*, (cfl's), which constitute an extension of the classical concept of Lyapunov functions from dynamical system theory. We discuss the role of clf's in design methods and "universal" formulas for feedback controls.

For nonlinear systems, it has been known since the late 1970s that, in general, there are topological obstructions to the existence of even continuous stabilizers. We review these obstructions, using tools from degree theory.

Finally, we turn to discontinuous stabilization and the associated issue of defining precisely a "solution" for a differential equation with discontinuous right-hand side. We introduce techniques from nonsmooth analysis and differential games, in order to deal with discontinuous controllers. In particular, we discuss the effect of measurement errors on the performance of such controllers.

Preliminaries

In this article, we restrict attention to continuous-time deterministic systems whose states evolve in finite-dimensional Euclidean spaces $\mathbb{R}^n$. (This excludes many other equally important objects of study in control theory: systems which evolve on infinite dimensional spaces and are described by PDE's, systems evolving on manifolds which serve to model state constraints, discrete-time systems described by difference equations, and stochastic systems, among others.) In order to streamline the presentation, we suppose throughout that controls take values in $\mathcal{U} = \mathbb{R}^m$ (constraints in controls would lead to proper subsets $\mathcal{U}$). A *control* (other names: *input*, *forcing function*) is any measurable locally essentially bounded map $u(\cdot)\colon [0,\infty) \to \mathcal{U} = \mathbb{R}^m$. In general, we use the notation $|x|$ for Euclidean norms, and use $\|u\|$ to indicate the essential supremum of a function $u(\cdot)$. For basic terminology and facts about control systems, see [25].

Given a map $f\colon \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ which is locally Lipschitz and satisfies $f(0,0) = 0$, we consider the associated forced system of ordinary differential equations

$$\dot{x}(t) = f(x(t), u(t))\,. \tag{1}$$

The maximal solution $x(\cdot)$ of (1) which corresponds to a given initial state $x(0) = x^0$ and to a given control u is defined on some maximal interval $[0, t_{\max}(x^0, u))$, and is denoted by $x(t, x^0, u)$. In the special case when f does not depend on u, we have an unforced system, or system with no inputs

$$\dot{x}(t) = f(x(t))\,. \tag{2}$$

Unforced systems are associated to a controlled system (1) in two different ways. The first is when one substitutes a feedback law $u = k(x)$ in (1) to obtain a "closed-loop" system $\dot{x} = f(x, k(x))$. The second is when one considers instead the autonomous system $\dot{x} = f(x, 0)$ which models the behavior of (1) in the absence of any controls. All definitions stated for unforced systems are implicitly applied also to systems with inputs (1) by setting $u \equiv 0$; for instance, we define the global asymptotic stability (GAS) property for (2), but we say that (1) is GAS if $\dot{x} = f(x, 0)$ is. For systems with no inputs (2) we write just $x(t, x^0)$ instead of $x(t, x^0, u)$.

Stability and Asymptotic Controllability Stability is one of the most important objectives in control theory, because a great variety of problems can be recast in stability terms. This includes questions of driving a system to a desired configuration (e. g., an inverted pendulum on a cart, to its upwards position), or the problem of tracking a reference signal (such as a pilot's command to an aircraft). We focus in this talk on global asymptotic stabilization.

Recall that the class of $\mathcal{K}_\infty$ functions consists of all $\alpha\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ which are continuous, strictly increasing, unbounded, and satisfy $\alpha(0) = 0$. The class of $\mathcal{KL}$ functions consists of those $\beta\colon \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ with the properties that

(1) $\beta(\cdot, t) \in \mathcal{K}_\infty$ for all t, and
(2) $\beta(r, t)$ decreases to zero as $t \to \infty$.

We will also use $\mathcal{N}$ to denote the set of all nondecreasing functions $\sigma\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Expressed in terms of such comparison functions, the property of *global asymptotic stability (GAS)* of the origin for a system with no inputs (2) becomes:

$$(\exists \beta \in \mathcal{KL}) \quad |x(t, x^0)| \leq \beta\left(|x^0|, t\right) \quad \forall x^0\,, \forall t \geq 0\,.$$

This definition is equivalent to a more classical "ε-δ" definition usually provided in textbooks, which defines GAS as the combination of stability and global attractivity. For one implication, simply observe that

$$|x(t, x^0)| \leq \beta\left(|x^0|, 0\right)$$

provides the stability (or "small overshoot") property, while

$$|x(t, x^0)| \leq \beta\left(|x^0|, t\right) \underset{t\to\infty}{\longrightarrow} 0$$

gives attractivity. The converse implication is an easy exercise.

More generally, we define what it means for the system with inputs (1) to be (open loop, globally) *asymptotically controllable (AC)* (to the origin). The definition amounts to requiring that for each initial state x^0 there exists some control $u = u_{x^0}(\cdot)$ defined on $[0, \infty)$, such that the corresponding solution $x(t, x^0, u)$ is defined for all $t \geq 0$, and converges to zero as $t \to \infty$, with "small" overshoot. Moreover, we wish to rule out the possibility that $u(t)$ becomes unbounded for x near zero. The precise formulation is as follows.

$$(\exists \beta \in \mathcal{KL})(\exists \sigma \in \mathcal{N}) \quad \forall x^0 \in \mathbb{R}^n \exists u(\cdot),\ \|u\| \leq \sigma(\left|x^0\right|),$$

$$|x(t, x^0, u)| \leq \beta\left(|x^0|, t\right) \quad \forall\, t \geq 0\,.$$

In particular, (global) asymptotic stability amounts to the existence of $\beta \in \mathcal{KL}$ such that $|x(t, x^0, u)| \leq \beta\,(|x^0|, t)$ holds for all $t \geq 0$. A very special case is that of *exponential* stability, in which an estimate of the type $|x(t, x^0, u)| \leq M e^{-\lambda t}|x^0|$ holds. For linear systems (see below), asymptotic stability and exponential stability coincide. It is a puzzling fact that for general systems, one can find continuous coordinate changes that make asymptotically stable systems exponentially stable [8,13] (a fact of little practical utility, since finding such coordinate changes is as hard as establishing stability to being with).

Feedback Stabilization A map $k\colon \mathbb{R}^n \to \mathcal{U}$ is a *feedback stabilizer* for the system with inputs (1) if k is locally bounded (that is, k is bounded on each bounded subset of $\mathbb{R}$), $k(0) = 0$, and the closed-loop system

$$\dot{x} = f(x, k(x)) \tag{3}$$

is GAS, i. e. there is some $\beta \in \mathcal{KL}$ so that $|x(t)| \leq \beta(|x(0)|, t)$ for all solutions and all $t \geq 0$. (A technical difficulty with this definition lies the possible lack of solutions of (1) when k is not regular enough. We ignore this for now, but will most definitely return to this issue later.)

For example, if (1) is a model of an undamped spring/mass system, where u represents the net effect of external forces, one obvious way to asymptotically stabilize the system is to introduce damping. In control-theoretic terms, this means that we choose $u(t) = k(x(t))$ to be a negative linear function of the velocity. Physically, one may implement a feedback controller by means of an analog device. In the example of the spring/mass system, one could achieve this by adding friction or connecting a dashpot. Alternatively, in modern control technology, one uses a digital computer to measure the state x and compute the appropriate control action to be applied. (There are many implementation issues which arise in digital control and are ignored in our theoretical formulation "$u = k(x)$", among them the effect of delays in the actual computation of the control $u(t)$ to be applied at time t, and the effect of quantization due to the finite precision of measuring devices and the digital nature of the computer. These issues are addressed in the literature, although a comprehensive theoretical framework is still lacking.)

Observe that, obviously, if there exists a feedback stabilizer for (1), then (1) is also AC (we just use $u(t) := k(x(t, x^0))$ as u_{x^0}). Thus, it is very natural to ask whether the converse holds: *is every asymptotically controllable system also feedback stabilizable?*

Linear Systems

A *linear system* is a system (1) for which the map f is linear. In other words, there are two linear transformations $A\colon \mathbb{R}^n \to \mathbb{R}^n$ and $B\colon \mathbb{R}^m \to \mathbb{R}^n$ so that the equations take the form

$$\dot{x} = Ax + Bu\,. \qquad (4)$$

Such a system is completely specified once that we are given A and B, which we identify by abuse of notation with their respective matrices $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$ with respect to the canonical bases in $\mathbb{R}^n$ and $\mathbb{R}^m$. We also say "the system (A, B)" when referring to (4).

It is natural to look specifically for *linear* feedbacks $k\colon \mathbb{R}^n \to \mathbb{R}^m$ which stabilize a linear system (just as a linear term, inversely proportional to velocity, stabilizes an undamped harmonic oscillator). (In fact, this is no loss of generality, since it can be easily proved for linear systems [25] that if a feedback stabilizer $u = k(x)$ exists, then there also exists a linear feedback stabilizer.) We write $u = Fx$, when expressing $k(x) = Fx$ in matrix terms with respect to the canonical bases. Substituting this control law into (4) results in the equation $\dot{x} = (A + BF)x$. Thus, the mathematical problem reduces to:

> given $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$, find $F \in \mathbb{R}^{m\times n}$ such that $A + BF$ is Hurwitz.

(Recall that a Hurwitz matrix is one all whose eigenvalues have negative real parts, and that the origin of the system $\dot{x} = Hx$ is globally asymptotically stable if and only if H is a Hurwitz matrix.) The fundamental stabilization result for linear systems is as follows [25]:

Theorem 1 *A linear system is asymptotically controllable if and only if it admits a linear feedback stabilizer.*

A Remark on Linearization

If the dynamics map f in (1) is continuously differentiable, we may expand to first order $f(x, u) = Ax + Bu + o(x, u)$. Let us suppose that the linearized system (A, B) is AC, and pick a linear feedback stabilizer $u = Fx$, whose existence is guaranteed by Theorem 1. Then, the same feedback law $k(x) = Fx$, when fed back into the original system (1), results in $\dot{x} = f(x, Fx) = (A + BF)x + o(x)$. Thus, k locally stabilizes the origin for the nonlinear system. Of course, the assumption that the linearization is AC is not always satisfied. Systems in which inputs enter multiplicatively, such as those controlling reaction rates in chemical problems, lead to degenerate linearizations. In addition, even if the linearized system (A, B) is AC, in general a linear stabilizer $u = Fx$ will not work as a global stabilizer. For example, the system $\dot{x} = x + x^2 + u$ can be locally stabilized with $u := -2x$, but any linear feedback $u = -fx$ $(f > 1)$ results in an additional equilibrium away from the origin (at $x = f - 1$). Nonlinear feedback must be used (obviously, in this, example, $u = -2x - x^2$ works for global stabilization).

Returning to linear systems, let us note that Theorem 1 is of great interest because (1) there is a simple algebraic test to check the AC property, and (2) there are several practically useful algorithms for obtaining a stabilizing F, including pole placement and optimization, which we sketch next.

Pole Placement

The first technique for stabilization is purely algebraic. In order to simplify this exposition, we will suppose that the system (4) is not just AC but is in fact *controllable*, meaning that every state can be steered, in finite time, to every other state. (Any AC system (4) can be decomposed into two components, of which one is already GAS and the other one is controllable, cf. [25], so this represents no loss of generality.) Controllability is characterized by the property – generically satisfied for pairs (A, B) – that

$$\text{rank}\left[B\middle|AB\middle|A^2B\middle|\dots\middle|A^{n-1}B\right] = n$$

(note that the composite matrix shown has n rows and nm columns).

Two pairs (A, B) and $(\widetilde{A}, \widetilde{B})$ are said to be *feedback equivalent* if there exist $T \in GL(n, \mathbb{R})$, $F_0 \in \mathbb{R}^{m\times n}$, and $V \in GL(m, \mathbb{R})$ so that

$$\begin{aligned}(A + BF_0)T &= T\widetilde{A} \quad \text{and} \\ BV &= T\widetilde{B}\,.\end{aligned} \tag{5}$$

Feedback equivalence corresponds to changes of basis in the state and control-value spaces (invertible matrices T and V, respectively) and feedback transformations $u = F_0x + u'$, where u' is a new control. (An equivalent way to describe feedback equivalence is by the requirement that two pairs should be in the same orbit under the action of a "feedback group" which is obtained as a suitable semidirect product of $GL(n, \mathbb{R})$, $(\mathbb{R}^{m\times n}, +)$, and $GL(m, \mathbb{R})$.) Controllability is preserved under feedback equivalence. Moreover, if (5) holds and if one finds a matrix $\widetilde{F}$ so that $\widetilde{A} + \widetilde{B}\widetilde{F}$ is Hurwitz, then

$$A + BF = T(\widetilde{A} + \widetilde{B}\widetilde{F})T^{-1}$$

is also Hurwitz, where $F := F_0 + V\widetilde{F}T^{-1}$. Thus, the task of finding a stabilizing feedback F can be reduced to the same problem for any pair $(\widetilde{A}, \widetilde{B})$ which is feedback equivalent to the given pair (A, B).

One then proceeds to show that there always exists an equivalent pair $(\widetilde{A}, \widetilde{B})$ which has a form simple enough that the existence of $\widetilde{F}$ is trivial to establish. In order to find such a pair, it is useful to study the classification of controllable pairs under feedback equivalence. This classification is closely related to Kronecker's theory of "matrix pencils" applied to polynomial matrices $[\lambda I - A, B] = \lambda[I, 0] + [-A, B]$ modulo matrix equivalence, cf. [25]. The orbits under feedback equivalence are in one-to-one correspondence with the possible partitions of $n = \kappa_1 + \ldots + \kappa_r$ into the sum of r positive integers, $r \leq m$, and in each orbit one can find a pair $(\widetilde{A}, \widetilde{B})$ which is in "controller canonical form", for which F can be trivially found. For simplicity, let is just discuss here the very special case of single-input systems ($m = 1$). For this case, the action of the feedback group is transitive, and each controllable system is feedback equivalent to the following special system:

$$A := \begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix} \qquad B := \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix}.$$

For this system, a stabilizing feedback is trivial to obtain. Indeed, take the polynomial $p(\lambda) = (\lambda + 1)^n = \lambda^n - \alpha_n\lambda^{n-1} - \ldots - \alpha_2\lambda - \alpha_1$. With $F = (\alpha_1, \alpha_2, \alpha_3, \ldots, \alpha_n)$,

$$A + BF := \begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ \alpha_1 & \alpha_2 & \alpha_3 & \dots & \alpha_n \end{pmatrix}$$

has characteristic polynomial $(\lambda+1)^n$, and hence is a Hurwitz matrix, as required for stabilization.

Observe that, instead of the particular $p(\lambda)$ which we used, we could have picked any polynomial all whose roots have negative real parts, and the same argument applies. The conclusion is that, not only can we make $A + BF$ Hurwitz, but we can assign to it any desired set of n eigenvalues (as long as they form a set closed under conjugation). This is the reason that the technique is called *eigenvalue placement* (or "pole placement" because the eigenvalues of A are the poles of the resolvent $(\lambda I - A)^{-1}$). See Chap. 5 in [25] for a detailed treatment of the pole placement problem.

Variational Approach

A second technique for stabilization is based on optimal control techniques. We first pick any two symmetric positive definite matrices $R \in \mathbb{R}^{m\times m}$ and $Q \in \mathbb{R}^{n\times n}$ (for instance the identity matrices of the respective sizes). Next, we consider the problem of minimizing, for each initial state x^0 at time $t = 0$, the infinite-horizon cost

$$\mathcal{J}_{x^0}(u) := \int_0^\infty \left\{u(t)'Ru(t) + x(t)'Qx(t)\right\} \mathrm{d}t$$

over all controls $u\colon [0,\infty) \to \mathbb{R}^m$ which make $\mathcal{J}_{x^0}(u) < \infty$, where $x(t) = x(t, x^0, u)$ and prime indicates transpose. The main result from linear-quadratic optimal control (cf. Sect. 8.4 in [25]) implies that, for AC systems, there is a unique solution u to this problem, which is given in the following form: there is a matrix $F \in \mathbb{R}^{m\times n}$ such that solving $\dot{x} = (A+BF)x$ with $x(0) = x^0$ gives that $u(t) := Fx(t)$ minimizes $\mathcal{J}_{x^0}(\cdot)$. Moreover, this F stabilizes the system (which is intuitively to be expected, since $\mathcal{J}_{x^0}(u) < \infty$

implies that solutions $x(t)$ must be in L^2), and F can be computed by the formula

$$F := -R^{-1}B'P\,, \tag{6}$$

where P is a symmetric and positive definite solution of the *Matrix Algebraic Riccati Equation*

$$PBR^{-1}B'P - PA - A'P - Q = 0\,. \tag{7}$$

A Sufficient Nonlinear Condition

Although of limited applicability, it is worth remarking that there is a partial extension to nonlinear systems of the stabilization method which was just described. For simplicity, we specialize our discussion to control-affine systems, i. e., those for which the input appears only in an affine form. This class is sufficient for the study of most forced mechanical systems. The Eq. (1) becomes:

$$\dot{x} = g_0(x) + \sum_{i=1}^{m} u_i g_i(x) = g_0(x) + G(x)u\,.$$

(See Sect. 8.5 in [25] for general f, and for proofs). We now pick two continuous functions $Q\colon \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ and $R\colon \mathbb{R}^n \to \mathbb{R}^{n\times n}$, so that $R(x)$ is a symmetric positive definite matrix for each x.

In general, we say that a continuous function

$$V\colon \mathbb{R}^n \to \mathbb{R}_{\geq 0}$$

is *positive definite* if $V(x) = 0$ only if $x = 0$, and it is *proper* (or "weakly coercive") if for each $a \geq 0$ the set $\{x | V(x) \leq a\}$ is compact, or, equivalently, $V(x) \to \infty$ as $|x| \to \infty$ (radial unboundedness). Given any such V which is also differentiable, we denote the vector function whose components are the directional derivatives of V in the directions of the various control vector fields g_i, $i \geq 1$ by:

$$L_G V(x) := \nabla V(x) G(x) = \left(L_{g_1} V(x), \ldots, L_{g_m} V(x)\right)\,,$$

and also write $L_{g_0} V(x) := \nabla V(x) g_0(x)$.

We consider the following PDE on such functions V:

$$\forall x \quad Q(x) + L_{g_0} V(x) - \frac{1}{4} L_G V(x) R(x)^{-1} \left(L_G V(x)\right)' = 0\,. \tag{8}$$

This reduces to the Algebraic Riccati Eq. (7) in the special case of linear systems, quadratic $V(x) = x'Px$ and $Q(x) = x'Qx$, and constant matrices $R(x) \equiv R$. We also take the following generalization of the feedback law (6):

$$k(x) := -\frac{1}{2} R(x)^{-1} \left(L_G V(x)\right)'\,. \tag{9}$$

Finally, we assume that Q is a positive definite function. One then has [25]:

Theorem 2 *Suppose that V is a twice continuously differentiable, positive definite, and proper solution of the PDE (8). Then, k defined by (9) stabilizes the system.*

This theorem arises from the following optimization problem: for each state $x^0 \in \mathbb{R}^n$, minimize the cost

$$\mathcal{J}_{x^0}(u) := \int_0^\infty u(t)' R(x(t)) u(t) + Q(x(t))\mathrm{d}t\,,$$

where $x(t) = x(t, x^0, u)$, over all those controls $u\colon [0,\infty) \to \mathcal{U}$ for which the solution $x(t, x^0, u)$ of (1) is defined for all $t \geq 0$ and satisfies $\lim_{t\to\infty} x(t) = 0$. Under the above assumptions, and as for linear systems, one also concludes that for each state x^0 the solution of $\dot{x} = f(x, k(x))$ with initial state $x(0) = x^0$ exists for all $t \geq 0$, the control $u(t) = k(x(t))$ is optimal, and $V(x^0)$ is the optimal cost from initial state x^0. Moreover, the formula for k arises from the Hamilton–Jacobi–Bellman equation of optimal control theory, because

$$k(x) = \operatorname*{argmin}_u \left\{\nabla V(x) \cdot f(x,u) + u' R(x) u + Q(x)\right\}$$

when $f(x,u) = g_0(x) + G(x)u$.

There are applications where this method has proven useful. Unfortunately, however, and in contrast to the linear case, in general there exists no positive definite, proper, and C^2 solution V of the above PDE. On the other hand, the formula (9) does appear, with variations, in other contexts, including generalizations of the idea of adding damping to systems, cf. Sect. 5.9 in [25], and, more generally, the use of auxiliary positive definite and proper functions V, in similar roles, will be central to the control-Lyapunov ideas discussed later.

Nonlinear Systems: Continuous Feedback

One of the central topics which we will address here concerns possibly discontinuous feedback laws k. Before turning to that subject, however, we study continuous feedback. When dealing with linear systems, linear feedback is natural, and indeed sufficient from a theoretical standpoint, as shown by the results just reviewed. However, for our general study, major technical questions arise in even deciding on just what degree of regularity should be imposed on the feedback maps k.

It turns out that the precise requirements away from 0, say asking whether k is merely continuous or smooth, are not very critical; it is often the case that one can "smooth out" a continuous feedback (or, even, make it real-analytic, via Grauert's Theorem) away from the origin. So, in order to avoid unnecessary complications in exposition due to nonuniqueness, let us call a feedback k *regular* if it is locally Lipschitz on $\mathbb{R}^n \setminus \{0\}$. For such k, solutions of initial value problems $\dot{x} = f(x, k(x))$, $x(0) = x^0$, are well defined (at least for small time intervals $[0, \varepsilon)$) and, provided k is a stabilizing feedback, are unique (cf. Exercise 5.9.9 in [25]).

On the other hand, behavior at the origin cannot be "smoothed out" and, at zero, the precise degree of smoothness plays a central role in the theory [12]. For instance, consider the system

$$\dot{x} = x + u^3 .$$

The continuous (and, in fact, smooth away from zero) feedback $u = k(x) := -\sqrt[3]{2x}$ globally stabilizes the system (the closed-loop system becomes $\dot{x} = -x$). However, there is no possible stabilizing feedback which is differentiable at the origin, since $u = k(x) = O(x)$ implies that

$$\dot{x} = x + O\left(x^3\right)$$

about $x = 0$, which means that the solution starting at any positive and small point moves to the right, instead of towards the origin. (A general result, assuming that A has no purely imaginary eigenvalues, cf. [25], Section 5.8, is that if – and only if – $\dot{x} = Ax + Bu + o(x, u)$ can be locally asymptotically stabilized using a feedback which is differentiable at the origin, the linearization $\dot{x} = Ax + Bu$ must be AC itself. In the example that we gave, this linearization is just $\dot{x} = x$, which is not AC.)

We now turn to the question of existence of regular feedback stabilizers. We first study a comparatively trivial case, namely systems with one state variable and one input. After that, we turn to multidimensional systems.

The Special Case $n = m = 1$

There are algebraic obstructions to the stabilization of $\dot{x} = f(x, u)$ if the input u appears nonlinearly in f. Ignoring the requirement that there be a $\sigma \in \mathcal{N}$ so that controls can be picked with $\|u\| \leq \sigma(|x^0|)$, asymptotic controllability is, for $n = m = 1$, equivalent to:

$$(\forall x \neq 0)(\exists u)\ x f(x, u) < 0 \qquad (10)$$

(this is proved in [28]; it is fairly obvious, but some care must be taken to deal with the fact that one is allowing arbitrary measurable controls; the argument proceeds by first approximating such controls by piecewise constant ones). Let us introduce the following set:

$$\mathcal{O} := \{(x, u) | x f(x, u) < 0\} ,$$

and let $\pi: (x, u) \mapsto x$ be the projection into the first coordinate in $\mathbb{R}$. Then, (10) is equivalent to:

$$\pi\mathcal{O} = \mathbb{R} \setminus \{0\} .$$

(One can easily include the requirement "$\|u\| \leq \sigma(|x^0|)$" by asking that for each interval $[-K, K] \subset \mathbb{R}$ there must be some compact set $C_K \subset \mathbb{R}^2$ so that $[-K, K] \subseteq \pi(C_K)$. For simplicity, we ignore this technicality.) In these terms, a stabilizing feedback is nothing else than a locally bounded map $k: \mathbb{R} \to \mathbb{R}$ such that $k(0) = 0$ and so that k is a section of π on $\mathbb{R} \setminus \{0\}$:

$$(x, k(x)) \in \mathcal{O} \quad \forall x \neq 0 .$$

For a regular feedback, we ask that k be locally Lipschitz on $\mathbb{R} \setminus \{0\}$.

Clearly, there is no reason for Lipschitz, or for that matter, just continuous, sections of π to exist. As an illustration, take the system

$$\dot{x} = x\left[(u-1)^2 - (x-1)\right]\left[(u+1)^2 + (x-2)\right] .$$

Let

$$\mathcal{O}_1 = \left\{(u+1)^2 < (2-x)\right\} \text{ and } \mathcal{O}_2 = \{(u-1)^2 < (x-1)\}$$

(these are the interiors of two disjoint parabolas). Here, $\mathcal{O}$ has three connected components, namely $\mathcal{O}_2$, $\mathcal{O}_1$ intersected with $x < 0$, and $\mathcal{O}_1$ intersected with $x < 0$. It is clear that, even though $\pi\mathcal{O} = \mathbb{R}$, there is no continuous curve (graph of $u = k(x)$) which is always in $\mathcal{O}$ and projects onto $\mathbb{R} \setminus \{0\}$. On the other hand, there exist many possible feedback stabilizers provided that we allow one discontinuity. It is also possible to provide examples, even with $f(x, u)$ smooth, for which an infinite number of switches are needed in any possible stabilizing feedback law. Finally, it may even be possible to stabilize semiglobally with a regular feedback, meaning that for each compact subset K of the state-space there is a regular, even smooth, feedback law $u = k_K(x)$ such that all states in K get driven asymptotically to the origin, but yet it may be impossible to find a single $u = k(x)$ which works globally. See [26] for details.

When feedback laws are required to be continuous at the origin, new obstructions arise. The case of systems with $n = m = 1$ is also a good way to introduce this subject. The first observation is that stabilization about the origin (even if just local) means that we must have, near zero:

$$f(x, k(x)) \begin{cases} > 0 & \text{if} \quad x < 0 \\ < 0 & \text{if} \quad x > 0 \\ = 0 & \text{if} \quad x = 0 \end{cases} .$$

In fact, all that we need is that $f(x_1, k(x_1)) < 0$ for some $x_1 > 0$ and $f(x_2, k(x_2)) > 0$ for some $x_2 < 0$. This guarantees, via the intermediate-value theorem that, if k is continuous, the projection

$$(-\varepsilon, \varepsilon) \to \mathbb{R}\,, \quad x \mapsto f(x, k(x))$$

is onto a neighborhood of zero, for each $\varepsilon > 0$. It follows, in particular, that the image of

$$(-\varepsilon, \varepsilon) \times (-\varepsilon, \varepsilon) \to \mathbb{R}\,, \quad (x, u) \mapsto f(x, u)$$

also contains a neighborhood of zero, for any $\varepsilon > 0$ (that is, the map $(x, u) \mapsto f(x, u)$ is open at zero). This last property is intrinsic, being stated in terms of the original data $f(x, u)$ and not depending upon the feedback k. Brockett's condition, to be described next, is a far-reaching generalization of this argument; in its proof, degree theory replaces the use of the intermediate value theorem.

Obstructions and Necessary Degree Conditions

If there are "obstacles" in the state-space, or more precisely if the state-space is a proper subset of $\mathbb{R}^n$, discontinuities in feedback laws cannot in general be avoided, since the domain of attraction of an asymptotically stable vector field must be diffeomorphic to Euclidean space. But even if states evolve in Euclidean spaces, similar obstructions may arise. These are due not to the topology of the state space, but to "virtual obstacles" implicit in the form of the system equations. These obstacles occur when it is impossible to move instantaneously in certain directions, even if it is possible to move eventually in every direction, the phenomenon of "nonholonomy". As an illustration, let us consider a model for the "shopping cart" shown in Fig. 1 ("knife-edge" or "unicycle" are other names for this example). The state is given by the orientation θ, together with the coordinates x_1, x_2 of the midpoint between the back wheels. The front wheel is a castor, free to rotate.

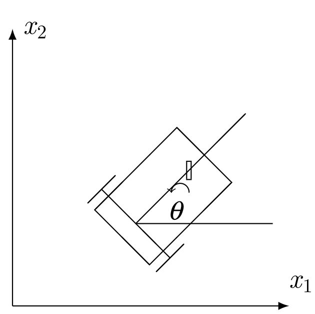

Stability and Feedback Stabilization, Figure 1
Shopping cart

There is a non-slipping constraint on movement: the velocity $(\dot{x}_1, \dot{x}_2)'$ must be parallel to the vector $(\cos\theta, \sin\theta)'$. This leads to the following equations:

$$\begin{aligned} \dot{x}_1 &= u_1 \cos\theta \\ \dot{x}_2 &= u_1 \sin\theta \\ \dot{\theta} &= u_2 \end{aligned}$$

where we may view u_1 as a "drive" command and u_2 as a steering control. (In practice, one would implement these controls by means of differential forces on the two back corners of the cart.) The feedback transformation $z_1 := \theta, z_2 := x_1 \cos\theta + x_2 \sin\theta, z_3 := x_1 \sin\theta - x_2 \cos\theta$, $v_1 := u_2$, and $v_2 := u_1 - u_2 z_3$ brings the system into the system with equations $\dot{z}_1 = v_1, \dot{z}_2 = v_2, \dot{z}_3 = z_2 v_1$ known as "Brockett's example" or "nonholonomic integrator" (yet another change can bring the third equation into the form $\dot{z}_3 = z_1 v_2 - z_2 v_1$). We view the system as having state space $\mathbb{R}^3$. Although a physically more accurate state space would be the manifold $\mathbb{R}^2 \times \mathbb{S}^1$, the necessary condition to be given is of a local nature, so the global structure is unimportant.

This system is (obviously) completely controllable (formally, controllability can be checked using the Lie algebra rank condition, as in e. g. [25], Exercise 4.3.16), and in particular is AC. However, we may expect that discontinuities are unavoidable due to the non-slip constraint, which does not allow moving from, for example the position $x_1 = 0, \theta = 0, x_2 = 1$ in a straight line towards the origin. Indeed, one then has [3]:

Theorem 3 *If there is a stabilizing feedback which is regular and continuous at zero, then the map $(x, u) \mapsto f(x, u)$ is open at zero.*

The test fails here, since no points of the form $(0, \varepsilon, *)$ belong to the image of the map

$$\begin{aligned} \mathbb{R}^5 \to \mathbb{R}^3 \colon (x_1, x_2, \theta, u_1, u_2)' &\mapsto f(x, u) \\ &= (u_1 \cos\theta, u_1 \sin\theta, u_2)' \end{aligned}$$

for $\theta \in (-\pi/2, \pi/2)$, unless $\varepsilon = 0$.

More generally, it is impossible to continuously stabilize any system without drift

$$\dot{x} = u_1 g_1(x) + \ldots + u_m g_m(x) = G(x)u$$

if $m < n$ and $\operatorname{rank}[g_1(0), \ldots, g_m(0)] = m$ (this includes all totally nonholonomic mechanical systems). Indeed, under these conditions, the map $(x, u) \mapsto G(x)u$ cannot contain a neighborhood of zero in its image, when restricted to a small enough neighborhood of zero. Indeed, let us first rearrange the rows of G:

$$G(x) \rightsquigarrow \begin{pmatrix} G_1(x) \\ G_2(x) \end{pmatrix}$$

so that $G_1(x)$ is of size $m \times m$ and is nonsingular for all states x that belong to some neighborhood N of the origin. Then,

$$\begin{pmatrix} 0 \\ a \end{pmatrix} \in \operatorname{Im}\left[N \times \mathbb{R}^m \to \mathbb{R}^n : (x,u) \mapsto G(x)u\right] \Rightarrow a = 0$$

(since $G_1(x)u = 0 \Rightarrow u = 0 \Rightarrow G_2(x)u = 0$ too).

If the condition $\operatorname{rank}[g_1(0), \ldots, g_m(0)] = m$ is violated, we cannot conclude a negative result. For instance, the system $\dot{x}_1 = x_1 u$, $\dot{x}_2 = x_2 u$ has $m = 1 < 2 = n$ but it can be stabilized by means of the feedback law $u = -(x_1^2 + x_2^2)$.

Observe that for linear systems, Brockett's condition says that

$$\operatorname{rank}[A, B] = n$$

which is the Hautus controllability condition (see e. g. [25], Lemma 3.3.7) at the zero mode.

Idea of the Proof One may prove Brockett's condition in several ways. A proof based on degree theory is probably easiest, and proceeds as follows (for details see for instance [25], Sect 5.9). The basic fact, due to Krasnosel'ski, is that if the system $\dot{x} = F(x) = f(x, k(x))$ has the origin as an asymptotically stable point and F is regular (since k is), then the degree (index) of F with respect to zero is $(-1)^n$, where n is the system dimension. In particular, the degree is also nonzero with respect to points p in a neighborhood of 0, which means that the equation $F(x) = p$ can be solved for small p, and hence $f(x,u) = p$ can be solved as well. The proof that the degree is $(-1)^n$ follows by exhibiting a homotopy, namely

$$F_t(x^0) = \frac{1}{t}\left[x\left(\frac{t}{1-t}, x^0\right) - x^0\right],$$

between $F_0 = F$ and $F_1(x) = -x$, and noting that the degree of the latter is obviously $(-1)^n$. An alternative proof uses Lyapunov functions. Asymptotic stability implies the existence of a smooth Lyapunov function V for $\dot{x} = F(x) = f(x, k(x))$, so, on the boundary ∂B of a sublevel set $B = \{x | V(x) \le c\}$ we have that F points towards the interior of B. Thus, for p small, $F(x) - p$ still points to the interior, which means that B is invariant with respect to the perturbed vector field $\dot{x} = F(x) - p$. Provided that a fixed-point theorem applies to continuous maps $B \to B$, this implies that $F(x) - p$ must vanish somewhere in B, that is, the equation $F(x) = p$ can be solved. (Because, for each small $h > 0$, the time-h flow ϕ of $F - p$ has a fixed point $x_h \in B$, i. e. $\phi(h, x_h) = x_h$, so picking a convergent subsequence $x_h \to \bar{x}$ gives that $0 = (\phi(h, x_h) - x_h)/(h) \to F(\bar{x}) - p$.) A fixed point theorem can indeed be applied, because B is a retract of $\mathbb{R}^n$ (use the flow itself); note that this argument gives a weaker conclusion than the degree condition.

Control-Lyapunov Functions

The method of control-Lyapunov functions ("clf's") provides a powerful tool for studying stabilization problems, both as a basis of theoretical developments and as a method for actual feedback design.

Before discussing clf's, let us quickly review the classical concept of Lyapunov functions, through a simple example. Consider first a damped spring-mass system $\ddot{y} + \dot{y} + y = 0$, or, in state-space form with $x_1 = y$ and $x_2 = \dot{y}$, $\dot{x}_1 = x_2$, $\dot{x}_2 = -x_1 - x_2$. One way to verify global asymptotic stability of the equilibrium $x = 0$ is to pick the (Lyapunov) function $V(x_1, x_2) := \frac{3}{2} x_1^2 + x_1 x_2 + x_2^2$, and observe that $\nabla V(x).f(x) = -|x|^2 < 0$ if $x \neq 0$, which means that

$$\frac{\mathrm{d}V(x(t))}{\mathrm{d}t} = -\left|x(t)^2\right| < 0$$

along all nonzero solutions, and thus the energy-like function V decreases along all trajectories, which, since V is a nondegenerate quadratic form, implies that $x(t)$ decreases, and in fact $x(t) \to 0$. Of course, in this case one could compute solutions explicitly, or simply note that the characteristic equation has all roots with negative real part, but Lyapunov functions are a general technique. (In fact, the classical converse theorems of Massera and Kurzweil [17,19] show that, whenever a system is GAS, there always exists a smooth Lyapunov function V.)

Now let us modify this example to deal with a control system, and consider a forced (but undamped) harmonic oscillator $\ddot{x} + x = u$, i. e. $\dot{x}_1 = x_2$, $\dot{x}_2 = -x_1 + u$. The damping feedback $u = -x_2$ stabilizes the system, but let us pretend that we do not know that. If we take the same V as before, now the derivatives along trajectories are, using "$\dot{V}(x,u)$" to denote $\nabla V(x).f(x,u)$ and omitting arguments t in $x(t)$ and $u(t)$:

$$\dot{V}(x,u) = -x_1^2 + x_1 x_2 + x_2^2 - (x_1 + 2x_2)u\,.$$

This expression is affine in u. Thus, if x is a state such that $x_1 + 2x_2 \neq 0$, then we may pick a control value u (which depends on this current state x) such that $\dot{V} < 0$. On the other hand, if $x_1 + 2x_2 = 0$, then the expression reduces to $\dot{V} = -5x_2^2$ (for any u), which is negative unless x_2 (and hence also $x_1 = -2x_2$) vanishes.

In conclusion, for each $x \neq 0$ there is some u so that $\dot{V}(x, u) < 0$. This is, except for some technicalities to be discussed, the characterizing property of control-Lyapunov functions. For any given compact subset B in $\mathbb{R}^n$, we now pick some compact subset $\mathcal{U}_0 \subset \mathcal{U}$ so that

$$\forall x \in B, x \neq 0\,, \ \exists u \in \mathcal{U}_0 \quad \text{such that } \dot{V}(x, u) < 0\,. \tag{11}$$

In principle, then, we could then stabilize the system, for states in B, by using the steepest descent feedback law:

$$k(x) := \operatorname*{argmin}_{u \in \mathcal{U}_0} \nabla V(x) \cdot f(x, u) \tag{12}$$

("argmin" means "pick any u at which the min is attained"; we restricted $\mathcal{U}$ to be assured that $\dot{V}(x, u)$ attains a minimum). Note that the stabilization problem becomes, in these terms, a set of static nonlinear programming problems: minimize a function of u, for each x. Global stabilization is also possible, by appropriately picking $\mathcal{U}_0$ as a function of the norm of x; later we discuss a precise formulation.

Control–Lyapunov functions, if understood non-technically as the basic paradigm "look for a function $V(x)$ with the properties that $V(x) \approx 0$ if and only if $x \approx 0$, and so that for each $x \neq 0$ it is possible to decrease $V(x)$ by some control action," constitute a very general approach to control (sometimes expressed in a dual fashion, as maximization of some measure of success). They appear in such disparate areas as A.I. game-playing programs (position evaluations), energy arguments for dissipative systems, program termination (Floyd/Dijkstra "variant"), and learning control ("critics" implemented by neural-networks). More relevantly to this paper, the idea underlies much of modern feedback control design, as illustrated for instance by the books [7,11,15,16,25].

Differentiable clf's: Precise Definition A *differentiable control-Lyapunov function* (clf) is a differentiable function $V\colon \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ which is proper, positive definite, and *infinitesimally decreasing*, meaning that there exists a positive definite continuous function $W\colon \mathbb{R}^n \to \mathbb{R}_{\geq 0}$, and there is some $\sigma \in \mathcal{N}$, so that

$$\sup_{x \in \mathbb{R}^n} \min_{|u| \leq \sigma(|x|)} \nabla V(x) \cdot f(x, u) + W(x) \leq 0\,. \tag{13}$$

This is basically the same as condition (11), with $\mathcal{U}_0 =$ the ball of radius $\sigma(|x|)$ picked as a function of x. The main difference is that, instead of saying "$\nabla V(x) \cdot f(x, u) < 0$ for $x \neq 0$" we write $\nabla V(x) \cdot f(x, u) \leq -W(x)$, where W is negative when $x \neq 0$. The two definitions are equivalent, but the "Hamiltonian" version used here is the correct one for the generalizations to be given, to nonsmooth V.

The basic result is due to Artstein [2]:

Theorem 4 *A control-affine system $\dot{x} = g_0(x) + \sum u_i g_i(x)$ admits a differentiable clf if and only if it admits a regular stabilizing feedback.*

The proof of sufficiency is easy: if there is such a k, then the converse Lyapunov theorem, applied to the closed-loop system $F(x) = f(x, k(x))$, provides a smooth V such that

$$L_F V(x) = \nabla V(x) F(x) < 0 \quad \forall x \neq 0\,.$$

This gives that for all nonzero x there is some u (bounded on bounded sets, because k is locally bounded by definition of feedback) so that $\dot{V}(x, u) < 0$; and one can put this in the form (13).

The necessity is more interesting. The original proof in [2] proceeds by a nonconstructive argument involving partitions of unity, but it is also possible [24,25] to exhibit explicitly a feedback, written as a function:

$$k\left(\nabla V(x) \cdot g_0(x), \ldots, \nabla V(x) \cdot g_m(x)\right)$$

of the directional derivatives of V along the vector fields defining the system (*universal formulas* for stabilization). Taking for simplicity $m = 1$, one such formula is:

$$k(x) := -\frac{a(x) + \sqrt{a(x)^2 + b(x)^4}}{b(x)} \quad (0 \text{ if } b = 0)$$

where $a(x) := \nabla V(x) \cdot g_0(x)$ and $b(x) := \nabla V(x) \cdot g_1(x)$. The expression for k is analytic in a,b when $x \neq 0$, because the clf property means that $a(x) < 0$ whenever $b(x) = 0$ [24,25].

Thus, the question of existence of regular feedback, for control-affine systems, reduces to the search for differentiable clf's, and this gives rise to a vast literature dealing with the construction of such V's, see [7,15,16,25] and references therein. Many other theoretical issues are also answered by Artstein's theorem. For example, via Kurzweil's converse theorem one has that the existence of k merely continuous on $\mathbb{R}^n \setminus \{0\}$ suffices for the existence of smooth (infinitely differentiable) V, and from here one may in turn find a k which is smooth on $\mathbb{R}^n \setminus \{0\}$. In addition, one may easily characterize the existence of k continuous at zero as well as regular: this is equivalent to the *small control property*: for each $\varepsilon > 0$ there is some $\delta > 0$ so that $0 < |x| < \delta$ implies that $\min_{|u| \leq \varepsilon} \nabla V(x) \cdot f(x, u) < 0$ (if this property holds, the universal formula automatically provides such a k). We should note that Artstein provided a result valid for general, not necessarily control-affine systems $\dot{x} = f(x,u)$; however, the obtained "feedback" has values in sets of relaxed controls, and is not a feedback law in the classical sense. Later, we discuss a different generalization.

Differentiable clf's will in general not exist, because of obstructions to regular feedback stabilization. This leads us naturally into the twin subjects of discontinuous feedbacks and non-differentiable clf's.

Discontinuous Feedback

The previous results and examples show that, in order to develop a satisfactory general theory of stabilization, one in which one proves the implication "asymptotic controllability implies feedback stabilizability," we must allow discontinuous feedback laws $u = k(x)$. But then, a major technical difficulty arises: solutions of the initial-value problem $\dot{x} = f(x, k(x))$, $x(0) = x^0$, interpreted in the classical sense of differentiable functions or even as (absolutely) continuous solutions of the integral equation $x(t) = x^0 + \int_0^t f(x(s), k(x(s)))\mathrm{d}s$, do not exist in general. The only general theorems apply to systems $\dot{x} = F(x)$ with continuous F. For example, there is no solution to $\dot{x} = -\operatorname{sign} x$, $x(0) = 0$, where $\operatorname{sign} x = -1$ for $x < 0$ and $\operatorname{sign} x = 1$ for $x \geq 0$. So one cannot even pose the stabilization problem in a mathematically consistent sense.

There is, of course, an extensive literature addressing the question of discontinuous feedback laws for control systems and, more generally, differential equations with discontinuous right-hand sides. One of the best-known candidates for the concept of solution of (3) is that of a *Filippov solution* [6,9], which is defined as the solution of a certain differential inclusion with a multivalued right-hand side which is built from $f(x, k(x))$. Unfortunately, there is no hope of obtaining the implication "asymptotic controllability implies feedback stabilizability" if one interprets solutions of (3) as Filippov solutions. This is a consequence of results in [5,22], which established that the existence of a discontinuous stabilizing feedback in the Filippov sense implies the Brockett necessary conditions, and, moreover, for systems affine in controls it also implies the existence of regular feedback (which we know is in general impossible).

A different concept of solution originates with the theory of discontinuous positional control developed by Krasovskii and Subbotin in the context of differential games in [14], and it is the basis of the new approach to discontinuous stabilization proposed in [4], to which we now turn.

Limits of High-Frequency Sampling

By a *sampling schedule* or *partition* $\pi = \{t_i\}_{i\geq 0}$ of $0, +\infty$ we mean an infinite sequence

$$0 = t_0 < t_1 < t_2 < \dots$$

with $\lim_{i\to\infty} t_i = \infty$. We call

$$\mathbf{d}(\pi) := \sup_{i\geq 0}(t_{i+1} - t_i)$$

the *diameter* of π. Suppose that k is a given feedback law for system (1). For each π, the *π-trajectory starting from x^0* of system (3) is defined recursively on the intervals $[t_i, t_{i+1})$, $i = 0, 1, \dots$, as follows. On each interval $t_i, t_{i+1})$, the initial state is measured, the control value $u_i = k(x(t_i))$ is computed, and the constant control $u \equiv u_i$ is applied until time t_{i+1}; the process is then iterated. That is, we start with $x(t_0) = x^0$ and solve recursively

$$\dot{x}(t) = f(x(t), k(x(t_i)))\,,\ t \in t_i, t_{i+1})\,, \quad i = 0, 1, 2, \dots$$

using as initial value $x(t_i)$ the endpoint of the solution on the preceding interval. The ensuing π-trajectory, which we denote as $x_\pi(\cdot, x^0)$, is defined on some maximal nontrivial interval; it may fail to exist on the entire interval $[0, +\infty)$ due to a blow-up on one of the subintervals $t_i, t_{i+1})$. We say that it is *well defined* if $x_\pi(t, x^0)$ is defined on all of $[0, +\infty)$.

Definition The feedback $k \colon \mathbb{R}^n \to \mathcal{U}$ *stabilizes* the system (1) if there exists a function $\beta \in \mathcal{KL}$ so that the following property holds: For each

$$0 < \varepsilon < K$$

there exists a $\delta = \delta(\varepsilon, K) > 0$ such that, for every sampling schedule π with $\mathbf{d}(\pi) < \delta$, and for each initial state x^0 with $|x^0| \leq K$, the corresponding π-trajectory of (3) is well-defined and satisfies

$$\left|x_\pi(t, x^0)\right| \leq \max\{\beta(K, t), \varepsilon\} \quad \forall t \geq 0\,. \tag{14}$$

In particular, we have

$$\left|x_\pi(t, x^0)\right| \leq \max\left\{\beta\left(\left|x^0\right|, t\right), \varepsilon\right\} \quad \forall t \geq 0 \tag{15}$$

whenever $0 < \varepsilon < |x^0|$ and $\mathbf{d}(\pi) < \delta(\varepsilon, |x^0|)$ (just take $K := |x^0|$).

Observe that the role of δ is to specify a lower bound on intersampling times. Roughly, one is requiring that

$$t_{i+1} \leq t_i + \theta(|x(t_i)|)$$

for each i, where θ is an appropriate positive function.

Our definition of stabilization is physically meaningful, and is very natural in the context of sampled-data (computer control) systems. It says in essence that a feedback k stabilizes the system if it drives all states asymptotically to the origin and with small overshoot when using *any fast enough sampling schedule*. A high enough sampling frequency is generally required when close to the origin, in order to guarantee small displacements, and also at infinity, so as to preclude large excursions or even blow-ups in finite time. This is the reason for making δ depend on ε and K.

This concept of stabilization can be reinterpreted in various ways. One is as follows. Pick any initial state x^0, and consider any sequence of sampling schedules π_ℓ whose diameters $\mathbf{d}(\pi_\ell)$ converge to zero as $\ell \to \infty$ (for instance, constant sampling rates with $t_i = i/\ell$, $i = 0, 1, 2, \ldots$). Note that the functions $x_\ell := x_{\pi_\ell}(\cdot, x^0)$ remain in a bounded set, namely the ball of radius $\beta(|x^0|, 0)$ (at least for ℓ large enough, for instance, any ℓ so that $\mathbf{d}(\pi_\ell)\delta(|x^0|/2, |x^0|)$). Because $f(x, k(x))$ is bounded on this ball, these functions are equicontinuous, and (Arzela–Ascoli's Theorem) we may take a subsequence, which we denote again as $\{x_\ell\}$, so that $x_\ell \to x$ as $\ell \to \infty$ (uniformly on compact time intervals) for some absolutely continuous (even Lipschitz) function $x\colon [0, \infty) \to \mathbb{R}^n$. *We may think of any limit function $x(\cdot)$ that arises in this fashion as a generalized solution of the closed-loop Eq. (3).* That is, generalized solutions are the limits of trajectories arising from arbitrarily high-frequency sampling when using the feedback law $u = k(x)$. Generalized solutions, for a given initial state x^0, may not be unique – just as may happen with continuous but non-Lipschitz feedback – but there is always existence, and, moreover, for any generalized solution, $|x(t)| \leq \beta(|x^0|, t)$ for all $t \geq 0$. This is precisely the defining estimate for the GAS property. Moreover, if k happens to be regular, then the unique solution of $\dot{x} = f(x, k(x))$ in the classical sense is also the unique generalized solution, so we have a reasonable extension of the concept of solution. (This type of interpretation is somewhat analogous, at least in spirit, to the way in which "relaxed" controls are interpreted in optimal trajectory calculations, namely through high-frequency switching of approximating regular controls.) The definition of stabilization was given in [4] in a slightly different form; see [26] for a discussion of the equivalence.

Stabilizing Feedbacks Exist

The main result is [4]:

Theorem 5 *The system (1) admits a stabilizing feedback if and only if it is asymptotically controllable.*

Necessity is clear. The sufficiency statement is proved by construction of k, and is based on the following ingredients:

- Existence of a nonsmooth control-Lyapunov function V.
- Regularization on shells of V.
- Pointwise minimization of a Hamiltonian for the regularized V.

In order to sketch this construction, we start by quickly reviewing a basic concept from nonsmooth analysis.

Proximal Subgradients Let V be any continuous function $\mathbb{R}^n \to \mathbb{R}$ (or even, just lower semicontinuous and with extended real values). A *proximal subgradient* of V at the point $x \in \mathbb{R}^n$ is any vector $\zeta \in \mathbb{R}^n$ such that, for some $\sigma > 0$ and some neighborhood $\mathcal{O}$ of x,

$$V(y) \geq V(x) + \zeta \cdot (y - x) - \sigma^2 \left|y - x^2\right| \quad \forall y \in \mathcal{O}\,.$$

In other words, proximal subgradients are the possible gradients of supporting quadratics at the point x. The set of all proximal subgradients at x is denoted $\partial_p V(x)$.

Nonsmooth Control–Lyapunov Functions A continuous (but not necessarily differentiable) $V\colon \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is a *control-Lyapunov function* (clf) if it is proper, positive definite, and infinitesimally decreasing in the following generalized sense: there exist a positive definite continuous $W\colon \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ and a $\sigma \in \mathcal{N}$ so that

$$\sup_{x \in \mathbb{R}^n} \max_{\zeta \in \partial_p V(x)} \min_{|u| \leq \sigma(|x|)} \zeta \cdot f(x, u) + W(x) \leq 0\,. \qquad (16)$$

This is the obvious generalization of the differentiable case in (13); we are still asking that one should be able to make $\nabla V(x) \cdot f(x, u) < 0$ by an appropriate choice of $u = u_x$, for each $x \neq 0$, except that now we replace $\nabla V(x)$ by the proximal subgradient set $\partial_p V(x)$. An equivalent property is to ask that V be a viscosity supersolution of the corresponding Hamilton–Jacobi–Bellman equation.

For nonsmooth clf's, the main basic result is [23,26]:

Theorem 6 *The system (1) is asymptotically controllable if and only if it admits a continuous clf.*

The proof is based on first constructing an appropriate W, and then letting V be the optimal cost (Bellman function) for the problem $\min \int_0^\infty W(x(s))\mathrm{d}s$. However, some care has to be taken to insure that V is continuous, and the cost has to be adjusted in order to deal with possibly unbounded minimizers. An important and very useful refinement of this result is the fact that a locally Lipschitz clf can also be shown to exist [21].

Regularization Once V is known to exist, the next step in the construction of a stabilizing feedback is to obtain Lipschitz approximations of V. For this purpose, one considers the Iosida–Moreau inf-convolution of V with a quadratic function:

$$V_\alpha(x) := \inf_{y\in\mathbb{R}^n}\left[V(y) + \frac{1}{2\alpha^2}|y-x|^2\right]$$

where the number $\alpha > 0$ is picked constant on appropriate regions. One has that $V_\alpha(x) \nearrow V(x)$, uniformly on compacts. Since V_α is locally Lipschitz, Rademacher's Theorem insures that V_α is differentiable almost everywhere. The feedback k is then made equal to a pointwise minimizer k_α of the Hamiltonian, at the points of differentiability (compare with (12) for the case of differentiable V):

$$k_\alpha(x) := \operatorname*{argmin}_{u\in\mathcal{U}_0} \nabla V_\alpha(x)\cdot f(x,u)\,,$$

where α and the compact $\mathcal{U}_0 = \mathcal{U}_0(\alpha)$ are chosen constant on certain compacts and this choice is made in between level curves. The critical fact is that V_α is itself a clf for the original system, at least when restricted to the region where it is needed. More precisely, on each shell of the form

$$C = \{x\in\mathbb{R}^n | r \le |x| \le R\}\,,$$

there are positive numbers m and α_0 and a compact subset $\mathcal{U}_0$ such that, for each $0 < \alpha \le \alpha_0$, each $x \in C$, and every $\zeta \in \partial_{\mathrm{p}} V_\alpha(x)$,

$$\min_{u\in\mathcal{U}_0} \zeta\cdot f(x,u) + m \le 0\,.$$

See [26]. (Actually, this description is oversimplified, and the proof is a bit more delicate. One must define, on appropriate compact sets

$$k(x) := \operatorname*{argmin}_{u\in\mathcal{U}_0} \zeta_\alpha(x)\cdot f(x,u)\,,$$

where $\zeta_\alpha(x)$ is carefully chosen. At points x of nondifferentiability, $\zeta_\alpha(x)$ is not a proximal subgradient of V_α, since $\partial_{\mathrm{p}} V_\alpha(x)$ may well be empty. One uses, instead, the fact that $\zeta_\alpha(x)$ happens to be in $\partial_{\mathrm{p}} V(x')$ for some $x' \approx x$. See [4] for details.)

Sensitivity to Small Measurement Errors

We have seen that every asymptotically controllable system admits a feedback stabilizer k, generally discontinuous, which renders the closed-loop system $\dot{x} = f(x,k(x))$ GAS. On the other hand, one of the main motivations for the use of feedback is in order to deal with uncertainty, and one possible source of uncertainty are measurement errors in state estimation. The use of discontinuous feedback means that undesirable behavior – chattering – may arise. In fact, one of the main reasons for the focus on continuous feedback is precisely in order to avoid such behaviors. Thus, we turn now to an analysis of the effect of measurement errors. Suppose first that k is a continuous function of x. Then, if the error e is small, using the control $u' = k(x+e)$ instead of $u = k(x)$ results in behavior which remains close to the intended one, since $k(x+e) \approx k(x)$; moreover, if $e \ll x$ then stability is preserved. This property of robustness to small errors when k is continuous can be rigorously established by means of a Lyapunov proof, based on the observation that, if V is a Lyapunov function for the closed-loop system, then continuity of $f(x,k(x+e))$ on e means that

$$\nabla V(x)\cdot f(x,k(x+e)) \approx \nabla V(x)\cdot f(x,k(x)) < 0\,.$$

Unfortunately, when k is not continuous, this argument breaks down. However, it can be modified so as to avoid invoking continuity of k. Assuming that V is continuously differentiable, one can argue that

$$\nabla V(x)\cdot f(x,k(x+e)) \approx \nabla V(x+e)\cdot f(x,k(x+e)) < 0$$

(using the Lyapunov property at the point $x+e$ instead of at x). This observation leads to a theorem, formulated below, which says that a discontinuous feedback stabilizer, robust with respect to small observation errors, can be found provided that there is a C^1 clf. In general, as there are no C^1, but only continuous, clf's, one may not be able to find any feedback law that is robust in this sense.

There are many well-known techniques for avoiding chattering, and a very common one is the introduction of deadzones where no action is taken. The feedback constructed in [4], with no modifications needed, can always be used in a manner robust with respect to small observation errors, using such an approach. Roughly speaking, the general idea is as follows. Suppose that the true current state, let us say at time $t = t_i$, is x, but that the controller uses $u = k(\tilde{x})$, where $\tilde{x} = x + e$, and e is small. Call x' the state that results at the next sampling time, $t = t_{i+1}$. By continuity of solutions on initial conditions, $|x' - \tilde{x}'|$ is also small, where $\tilde{x}'$ is the state that would have resulted from applying the control u if the true state had been $\tilde{x}$. By continuity, it follows that $V_\alpha(x) \approx V_\alpha(\tilde{x})$ and also $V_\alpha(x') \approx V_\alpha(\tilde{x}')$. On the other hand, the construction in [4] provides that $V_\alpha(\tilde{x}') < V_\alpha(\tilde{x}) - \mathrm{d}(t_{i+1} - t_i)$, where d is some positive constant (this is valid while we are far from the origin). Hence, if e is sufficiently small compared to the intersample time $t_{i+1} - t_i$, it will necessarily be the case that $V_\alpha(x')$ must also be smaller than $V_\alpha(x)$. This discussion may be formalized in several ways; see [26] for a precise statement.

If we insist upon fast sampling, a necessary condition arises, as was proved in the recent paper [18] (which, in turn, represented an extension of the work by Hermes [10] for classical solutions under observation error). We next discuss the main result from that paper. We consider systems

$$\dot{x}(t) = f(x(t), k(x(t) + \mathrm{e}(t)) + \mathrm{d}(t)) \quad (17)$$

in which there are observation errors as well as, now, possible actuator errors $d(\cdot)$. Actuator errors $d(\cdot)\colon [0,\infty) \to \mathcal{U}$ are Lebesgue measurable and locally essentially bounded, and observation errors $\mathrm{e}(\cdot)\colon [0,\infty) \to \mathbb{R}^n$ are locally bounded. We define solutions of (17), for each sampling schedule π, in the usual manner, i. e., solving recursively on the intervals $t_i, t_{i+1})$, $i = 0, 1, \ldots$, the differential equation

$$\dot{x}(t) = f\big(x(t), k(x(t_i) + \mathrm{e}(t_i)) + \mathrm{d}(t)\big) \quad (18)$$

with $x(0) = x^0$. We write $x(t) = x_\pi(t, x^0, \mathrm{d}, \mathrm{e})$ for the solution, and say it is *well-defined* if it is defined for all $t \geq 0$.

Definition The feedback $k\colon \mathbb{R}^n \to \mathcal{U}$ *stabilizes* the system (17) if there exists a function $\beta \in \mathcal{KL}$ so that the following property holds: For each

$$0 < \varepsilon < K$$

there exist $\delta = \delta(\varepsilon, K) > 0$ and $\eta = \eta(\varepsilon, K)$ such that, for every sampling schedule π with $\mathbf{d}(\pi) < \delta$, each initial state x^0 with $|x^0| \leq K$, and each e, d such that $|\mathrm{e}(t)| \leq \eta$ for all $t \geq 0$ and $|\mathrm{d}(t)| \leq \eta$ for almost all $t \geq 0$, the corresponding π-trajectory of (17) is well-defined and satisfies

$$\left|x_\pi(t, x^0, \mathrm{d}, \mathrm{e})\right| \leq \max\left\{\beta\left(K, t\right), \varepsilon\right\} \quad \forall t \geq 0\,. \quad (19)$$

In particular, taking $K := |x^0|$, one has that

$$\left|x_\pi(t, x^0, d, e)\right| \leq \max\left\{\beta\left(\left|x^0\right|, t\right), \varepsilon\right\} \quad \forall t \geq 0$$

whenever $0 < \varepsilon < |x^0|$, $\mathbf{d}(\pi) < \delta(\varepsilon, |x^0|)$, and for all t, $|\mathrm{e}(t)| \leq \eta(\varepsilon, |x^0|)$, and $|\mathrm{d}(t)| \leq \eta(\varepsilon, |x^0|)$.

The main result in [18] is as follows.

Theorem 7 *There is a feedback which stabilizes the system (17) if and only if there is a C^1 clf for the unperturbed system (1).*

It is interesting to note that, as a corollary of Artstein's Theorem, for control-affine systems $\dot{x} = g_0(x) + \sum u_i g_i(x)$ we may conclude that if there is a discontinuous feedback stabilizer that is robust with respect to small noise, then there is also a regular one, and even one that is smooth on $\mathbb{R}^n \setminus \{0\}$. For non control-affine systems, however, there may exist a discontinuous feedback stabilizer that is robust with respect to small noise, yet there is no regular feedback.

Briefly, the sufficiency part of Theorem 7 proceeds by taking a pointwise minimization of the Hamiltonian, for a given C^1 clf, i. e. $k(x)$ is defined as any u with $|u| \leq \sigma(|x|)$ which minimizes $\nabla V(x) \cdot f(x,u)$. The necessity part is based on the following technical fact: if the perturbed system can be stabilized, then the differential inclusion

$$\dot{x} \in F(x) := \bigcap_{\varepsilon > 0} \overline{\mathrm{co}}\, f(x, k(x + \varepsilon B))$$

(where B denotes the unit ball in $\mathbb{R}^n$) is strongly asymptotically stable. One may then apply converse Lyapunov theorems for upper semicontinuous compact convex differential inclusions to deduce the existence of V.

We now summarize exactly which implications hold, writing "robust" to mean stabilization of the system subject to observation and actuator noise:

$$\begin{array}{ccccc} C^1 V & \Longleftrightarrow & \exists \text{ robust } k & & \\ \Downarrow & & \Downarrow & & \\ C^0 V & \Longleftrightarrow & \exists k & \Longleftrightarrow & \mathrm{AC}\,. \end{array}$$

Future Directions

There are several alternative approaches to feedback stabilization, notably the very appealing approach to discontinuous stabilization throgh *patchy feedbacks* [1], as well as other related "hybrid" approaches [20]. It is also extremely important to understand the effect of "large" disturbances on the behavior of feedback systems. This study leads one to the very active area of input to state stability (ISS) and related notions (output to state stability as a model of detectability, input to output stability for the study of regulation problems, and so forth), see [27].

Bibliography

Primary Literature

1. Ancona F, Bressan A (2003) Flow stability of patchy vector fields and robust feedback stabilization. SIAM J Control Optim 41:1455–1476
2. Artstein Z (1983) Stabilization with relaxed controls. Nonlinear Anal Theory Methods Appl 7:1163–1173
3. Brockett R (1983) Asymptotic stability and feedback stabilization. In: Differential geometric control theory. Birkhauser, Boston, pp 181–191

4. Clarke FH, Ledyaev Y, Sontag E, Subbotin A (1997) Asymptotic controllability implies feedback stabilization. IEEE Trans Autom Control 42(10):1394–1407
5. Coron J-M, Rosier L (1994) A relation between continuous time-varying and discontinuous feedback stabilization. J Math Syst Estim Control 4:67–84
6. Filippov A (1985) Differential equations with discontinuous right-hand side. Nauka, Moscow. English edition: Kluwer, Dordrecht (1988)
7. Freeman R, Kokotović P (1996) Robust nonlinear control design. Birkhäuser, Boston
8. Grüne L, Sontag E, Wirth F (1999) Asymptotic stability equals exponential stability, and ISS equals finite energy gain—if you twist your eyes. Syst Control Lett 38(2):127–134
9. Hájek O (1979) Discontinuous differential equations, i & ii. J Diff Equ 32(2):149–170, 171–185
10. Hermes H (1967) Discontinuous vector fields and feedback control. In: Differential equations and dynamical systems. Proc Internat Sympos, Mayaguez 1965. Academic Press, New York, pp 155–165
11. Isidori A (1995) Nonlinear control systems: An introduction. Springer, Berlin
12. Kawski M (1989) Stabilization of nonlinear systems in the plane. Syst Control Lett 12:169–175
13. Kawski M (1995) Geometric homogeneity and applications to stabilization. In: Krener A, Mayne D (eds) Nonlinear control system design symposium (NOLCOS), Lake Tahoe, June 1995. Elsevier
14. Krasovskii N, Subbotin A (1988) Game-theoretical control problems. Springer, New York
15. Krstić M, Deng H (1998) Stabilization of uncertain nonlinear systems. Springer, London
16. Krstić M, Kanellakopoulos I, Kokotović PV (1995) Nonlinear and adaptive control design. Wiley, New York
17. Kurzweil J (1956) On the inversion of Ljapunov's second theorem on stability of motion. Am Math Soc Transl 2:19–77
18. Ledyaev Y, Sontag E (1999) A Lyapunov characterization of robust stabilization. Nonlinear Anal 37(7):813–840
19. Massera JL (1956) Contributions to stability theory. Ann Math 64:182–206
20. Prieur C (2006) Robust stabilization of nonlinear control systems by means of hybrid feedbacks. Rend Sem Mat Univ Pol Torino 64:25–38
21. Rifford L (2002) Semiconcave control-Lyapunov functions and stabilizing feedbacks. SIAM J Control Optim 41:659–681
22. Ryan E (1994) On Brockett's condition for smooth stabilizability and its necessity in a context of nonsmooth feedback. SIAM J Control Optim 32:1597–1604
23. Sontag E (1983) A Lyapunov-like characterization of asymptotic controllability. SIAM J Control Optim 21(3):462–471
24. Sontag E (1989) A "universal" construction of Artstein's theorem on nonlinear stabilization. Syst Control Lett 13(2):117–123
25. Sontag E (1998) Mathematical control theory. In: Deterministic finite-dimensional systems. Texts in Applied Mathematics, vol 6. Springer, New York
26. Sontag E (1999) Stability and stabilization: Discontinuities and the effect of disturbances. In : Nonlinear analysis, differential equations and control (Montreal, 1998). NATO Sci Ser C Math Phys Sci, vol 528. Kluwer, Dordrecht, pp 551–598
27. Sontag E (2006) Input to state stability: Basic concepts and results. In: Nistri P, Stefani G (eds) Nonlinear and optimal control theory. Springer, Berlin, pp 163–220
28. Sontag E, Sussmann H (1980) Remarks on continuous feedback. In: Proc IEEE Conf Decision and Control, Albuquerque, Dec 1980. pp 916–921

Books and Reviews

Clarke FH, Ledyaev YS, Stern RJ, Wolenski P (1998) Nonsmooth analysis and control theory. Springer, New York

Freeman R, Kokotović PV (1996) Robust nonlinear control design. Birkhäuser, Boston

Isidori A (1995) Nonlinear control systems, 3rd edn. Springer, London

Isidori A (1999) Nonlinear control systems II. Springer, London

Khalil HK (1996) Nonlinear systems, 2nd edn. Prentice-Hall, Upper Saddle River

Krstić M, Deng H (1998) Stabilization of uncertain nonlinear systems. Springer, London

Sepulchre R, Jankovic M, Kokotović PV (1997) Constructive nonlinear control. Springer, New York