Lyapunov Functions and Feedback in Nonlinear Control

Francis Clarke

Professeur à l'Institut Universitaire de France. Institut Desargues, Université de Lyon 1, 69622 Villeurbanne, France. clarke@igd.univ-lyon1.fr

Summary. The method of Lyapunov functions plays a central role in the study of the controllability and stabilizability of control systems. For nonlinear systems, it turns out to be essential to consider nonsmooth Lyapunov functions, even if the underlying control dynamics are themselves smooth. We synthesize in this article a number of recent developments bearing upon the regularity properties of Lyapunov functions. A novel feature of our approach is that the guidability and stability issues are decoupled. For each of these issues, we identify various regularity classes of Lyapunov functions and the system properties to which they correspond. We show how such regularity properties are relevant to the construction of stabilizing feedbacks. Such feedbacks, which must be discontinuous in general, are implemented in the sample-and-hold sense. We discuss the equivalence between open-loop controllability, feedback stabilizability, and the existence of Lyapunov functions with appropriate regularity properties. The extent of the equivalence confirms the cogency of the new approach summarized here.

1 Introduction

We consider a system governed by the standard control dynamics

$$\dot{x}(t) = f(x(t), u(t)) \text{ a.e., } u(t) \in \mathcal{U} \text{ a.e.}$$

or equivalently (under mild conditions) by the differential inclusion

$$\dot{x}(t) \in F(x(t)) \text{ a.e.}$$

The issue under consideration is that of guiding the state x to the origin. (The use of more general target sets presents no difficulties in the results presented here.)

A century ago, for the uncontrolled case in which the multifunction F is given by a (smooth) single-valued function (that is, $F(x) = \{f(x)\}$), Lyapunov introduced a criterion for the *stability* of the system, a property whereby all the trajectories $x(t)$ of the system tend to the origin (in a certain sense which we gloss over for now). This criterion involves the existence of a certain function V, now known as a Lyapunov function. Later, in the classical works of Massera, Barbashin and Krasovskii, and Kurzweil, this sufficient condition for stability was also shown to be necessary (under various sets of hypotheses).

In extending the technique of Lyapunov functions to control systems, a number of new issues arise. To begin with, we can distinguish two cases: we may require that *all* trajectories go to the origin (strong stability) or that (for a suitable choice of the control function) *some* trajectory goes to zero (weak stability, or controllability). In the latter case, unlike the former, it turns out that characterizing stability in terms of smooth Lyapunov functions is not possible; thus elements of *nonsmooth analysis* become essential. Finally, the issue of stabilizing *feedback design* must be considered, for this is one of the main reasons to introduce control Lyapunov functions. Here again regularity intervenes: in general, such feedbacks must be discontinuous, so that a method of implementing them must be devised, and new issues such as robustness addressed.

While these issues have been considered for decades, they have only recently been resolved in a unified and (we believe) satisfactory way. Several new tools have contributed to the analysis, notably: proximal analysis and attendant Hamilton-Jacobi characterizations of monotonicity properties of trajectories, semiconcavity, and sample-and-hold implementation of discontinuous feedbacks. The point of view in which the issues of guidability and stability are decoupled is also very recent. Our purpose here is to sketch the complete picture of these related developments for the first time, thereby synthesizing a guide for their comprehension. The principal results being summarized here appear in the half-dozen joint articles of Clarke, Ledyaev, Rifford and Stern cited in the references, and in the several works by Rifford; the article [8] of Clarke, Ledyaev, Sontag and Subbotin is also called upon. The necessary background in nonsmooth analysis is provided by the monograph of Clarke, Ledyaev, Stern and Wolenski [10].

Of course there is an extensive literature on the issues discussed here, with contributions by Ancona, Artstein, Bressan, Brockett, Coron, Kellett, Kokotovic, Praly, Rosier, Ryan, Sontag, Sussmann, Teel, and many others; these are discussed and cited in the introductions of the articles mentioned above. General references for Lyapunov functions in control include [2] and [14].

2 Strong Stability

We shall say that the control system $\dot{x}(t) \in F(x(t))$ a.e. is *strongly asymptotically stable* if every trajectory $x(t)$ is defined for all $t \geq 0$ and satisfies $\lim_{t\to+\infty} x(t) = 0$, and if in addition the origin has the familiar local property known as 'Lyapunov stability'. The following result, which unifies and extends several classical theorems dealing with the uncontrolled case, is due to Clarke, Ledyaev and Stern [9]:

Theorem 1. *Let F have compact convex nonempty values and closed graph. Then the system is strongly asymptotically stable if and only if there exists a pair of C^∞ functions*

$$V : \mathbb{R}^n \to \mathbb{R}, \quad W : \mathbb{R}^n \backslash \{0\} \to \mathbb{R}$$

satisfying the following conditions:

1. Positive Definiteness:

$$V(x) > 0 \text{ and } W(x) > 0 \quad \forall x \neq 0, \text{ and } V(0) \geq 0.$$

2. Properness: *The sublevel sets* $\{x : V(x) \leq c\}$ *are bounded* $\forall c$.
3. Strong Infinitesimal Decrease:

$$\max_{v \in F(x)} \langle \nabla V(x), v \rangle \; \leq \; -W(x) \quad x \neq 0.$$

We refer to the function (V, W) as a *strong Lyapunov function* for the system. Note that in this result, whose somewhat technical proof we shall not revisit here, the system multifunction F itself need not even be continuous, yet strong stability is equivalent to the existence of a smooth Lyapunov function: this is a surprising aspect of these results. As we shall see, this is in sharp contrast to the case of weak stability, where stronger hypotheses on the underlying system are required. In fact, in addition to the hypotheses of Theorem 1, we shall suppose henceforth that F is locally Lipschitz with linear growth. Even so, Lyapunov functions will need to be nondifferentiable in the controllability context.

Finally, we remark that in the positive definiteness condition, the inequality $V(0) \geq 0$ is superfluous when V is continuous (which will not be the case later); also, it could be replaced by the more traditional condition $V(0) = 0$ in the present context.

3 Guidability and Controllability

The Case for Less Regular Lyapunov Functions

Strong stability is most often of interest when F arises from a perturbation of an ordinary (uncontrolled) differential equation. In most control settings, it is weak (open loop) stability that is of interest: the possibility of guiding *some* trajectory to 0 in a suitable fashion. It is possible to distinguish two distinct aspects of the question: on the one hand, the possibility of guiding the state from any prescribed initial condition to 0 (or to an arbitrary neighborhood of 0), and on the other hand, that of keeping the state close to 0 when the initial condition is already near 0. In a departure from the usual route, we choose to decouple these two issues, introducing the term 'guidability' for the first. We believe that in so doing, a new level of clarity emerges in connection with Lyapunov theory.

A point α is *asymptotically guidable* to the origin if there is a trajectory x satisfying $x(0) = \alpha$ and $\lim_{t\to\infty} x(t) = 0$. When every point has this property, and when additionally the origin has the familiar local stability property known as *Lyapunov stability*, it is said in the literature to be *GAC*: (open loop) *globally asymptotically controllable* (to 0). A well-known *sufficient* condition for this property is the existence of a smooth (C^1, say) pair (V, W) of functions satisfying the positive definiteness and properness conditions of Theorem 1, together with *weak infinitesimal decrease:*

$$\min_{v \in F(x)} \langle \nabla V(x), v \rangle \ \leq \ -W(x) \quad x \neq 0.$$

Note the presence of a minimum in this expression rather than a maximum.

It is a fact, however, that as demonstrated by simple examples (see [6] or [23]), the existence of a smooth function V with the above properties fails to be a necessary condition for global asymptotic controllability; that is, the familiar converse Lyapunov theorems of Massera, Barbashin and Krasovskii, and Kurzweil do not extend to this weak controllability setting, at least not in smooth terms.

It is natural therefore to seek to weaken the smoothness requirement on V so as to obtain a necessary (and still sufficient) condition for a system to be GAC. This necessitates the use of some construct of nonsmooth analysis to replace the gradient of V that appears in the infinitesimal decrease condition. In this connection we use the *proximal subgradient* $\partial_P V(x)$, which requires only that the (extended-valued) function V be lower semicontinuous. In proximal terms, the weak infinitesimal decrease condition becomes

$$\sup_{\zeta \in \partial_P V(x)} \min_{v \in F(x)} \langle \zeta, v \rangle \ \leq \ -W(x) \quad x \neq 0.$$

Note that this last condition is trivially satisfied when x is such that $\partial_P V(x)$ is empty, in particular when $V(x) = +\infty$. (The supremum over the empty set is $-\infty$.) Henceforth, a *general Lyapunov pair* (V, W) refers to extended-valued lower semicontinuous functions $V : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ and $W : \mathbb{R}^n \backslash \{0\} \to \mathbb{R} \cup \{+\infty\}$ satisfying the positive definiteness and properness conditions of Theorem 1, together with proximal weak infinitesimal decrease.

The following is proved in [10]:

Theorem 2. *Let (V, W) be a general Lyapunov pair for the system. Then any $\alpha \in \operatorname{dom} V$ is asymptotically guidable to 0.*

We proceed to make some comments on the proof. To show that any initial condition can be steered towards zero (in the presence of a Lyapunov function), one can invoke the infinitesimal decrease condition to deduce that the function $V(x)+y$ is *weakly decreasing* for the augmented dynamics $F(x) \times \{W(x)\}$ (see pp. 213-214 of [10] for details); this implies the existence of a trajectory x such that the function

$$t \mapsto V(x(t)) + \int_0^t W(x(\tau))\, d\tau$$

is nonincreasing, which in turn implies that $x(t) \to 0$. We remark that viability theory can also be used in this type of argument; see for example [1].

It follows from the theorem that the existence of a lower semicontinuous Lyapunov pair (V, W) with V everywhere finite-valued implies the global asymptotic guidability to 0 of the system. This does not imply Lyapunov stability at the origin, however, so it cannot characterize global asymptotic controllability. An early and seminal result due to Sontag [22] considers *continuous* functions V, with the infinitesimal decrease condition expressed in terms of Dini derivates. Here is a version of it in proximal subdifferential terms:

Theorem 3. *The system is GAC if and only if there exists a* continuous *Lyapunov pair* (V, W).

For the sufficiency, the requisite guidability evidently follows from the previous theorem. The continuity of V provides the required local stability: roughly speaking, once $V(x(t))$ is small, its value cannot take an upward jump, so $x(t)$ remains near 0.

The proof of the converse theorem (that a continuous Lyapunov function must exist when the system is globally asymptotically controllable) is more challenging. One route is as follows: In [7] it was shown that certain locally Lipschitz *value functions* give rise to *practical* Lyapunov functions (that is, assuring stable controllability to arbitrary neighborhoods of 0, as in Theorem 4 below). Building upon this, Rifford [18, 19] was able to combine a countable family of such functions in order to construct a global *locally Lipschitz* Lyapunov function. This answered a long-standing open question in the subject. Rifford also went on to show the existence of a *semiconcave* Lyapunov function, a property whose relevance to feedback construction will be seen in the following sections.

Finally, we remark that the equivalence of the Dini derivate and of the proximal subdifferential forms of the infinitesimal decrease condition is a consequence of Subbotin's Theorem (see [10]).

Practical guidability.

The system is said to be (open-loop) *globally practically guidable* (to the origin) if for each initial condition α and for every $\varepsilon > 0$ there exists a trajectory x and a time T (both depending on α and ε) such that $|x(T)| \le \varepsilon$. We wish to characterize this property in Lyapunov terms. For this purpose we need an extension of the Lyapunov function concept.

ε-Lyapunov functions.

An *ε-Lyapunov pair* for the system refers to lower semicontinuous functions $V : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ and $W : \mathbb{R}^n \backslash \overline{B}(0, \varepsilon) \to \mathbb{R} \cup \{+\infty\}$ satisfying the usual properties of a Lyapunov pair, but with the role of the origin replaced by the closed ball $\overline{B}(0, \varepsilon)$:

1. *Positive Definiteness:*

$$V(x) > 0 \text{ and } W(x) > 0 \quad \forall x \notin \overline{B}(0, \varepsilon), \text{ and } V \geq 0 \text{ on } \overline{B}(0, \varepsilon).$$

2. *Properness:* The sublevel sets $\{x : V(x) \leq c\}$ are bounded $\forall c$.
3. *Weak Infinitesimal Decrease:*

$$\min_{v \in F(x)} \langle \nabla V(x), v \rangle \ \leq \ -W(x) \quad x \notin \overline{B}(0, \varepsilon).$$

The results in [7] imply:

Theorem 4. *The system is globally practically guidable to the origin if and only if there exists a locally Lipschitz ε-Lyapunov function for each $\varepsilon > 0$.*

We do not know whether global asymptotic guidability can be characterized in analogous terms, or whether practical guidability can be characterized by means of a *single* Lyapunov function. However, it is possible to do so for *finite-time* guidability (see Section 6 below).

4 Feedback

The Case for More Regular Lyapunov Functions

The need to consider discontinuous feedback in nonlinear control is now well established, together with the attendant need to define an appropriate solution concept for a differential equation in which the dynamics fail to be continuous in the state. The best-known solution concept in this regard is that of Filippov. For the stabilization issue, and using the standard formulation

$$\dot{x}(t) = f(x(t), u(t)) \text{ a.e.}, \ u(t) \in \mathcal{U} \text{ a.e.}$$

rather than the differential inclusion, the issue becomes that of finding a feedback control function $k(x)$ (having values in $\mathcal{U}$) such that the ensuing differential equation

$$\dot{x} = g(x), \quad \text{where} \quad g(x) := f(x, k(x))$$

has the required stability. The *central question* in the subject has long been: If the system is open loop globally asymptotically controllable (to the origin), is there a feedback k such that the resulting g exhibits global asymptotic stability (of the origin)? It has long been known that continuous feedback laws cannot suffice for this to be the case; it also turns out that admitting discontinuous feedbacks interpreted in the Filippov sense is also inadequate.

The question was settled by Clarke, Ledyaev, Sontag and Subbotin [8], who used the proximal aiming method (see also [11]) to show that the answer is positive if the (discontinuous) feedbacks are implemented in the *closed-loop system sampling* sense (also referred to as *sample-and-hold*). We proceed now to describe the sample-and-hold implementation of a feedback.

Let $\pi = \{t_i\}_{i \geq 0}$ be a partition of $[0, \infty)$, by which we mean a countable, strictly increasing sequence t_i with $t_0 = 0$ such that $t_i \to +\infty$ as $i \to \infty$. The *diameter* of π, denoted diam (π), is defined as $sup_{i \geq 0}(t_{i+1} - t_i)$. Given an initial condition x_0, the π-trajectory $x(\cdot)$ corresponding to π and an arbitrary feedback law $k : \mathbb{R}^n \to \mathcal{U}$ is defined in a step-by-step fashion as follows. Between t_0 and t_1, x is a classical solution of the differential equation

$$\dot{x}(t) = f(x(t), k(x_0)), \quad x(0) = x_0, \quad t_0 \leq t \leq t_1.$$

(Of course in general we do not have uniqueness of the solution, nor is there necessarily even one solution, although nonexistence can be ruled out when blow-up of the solution in finite time cannot occur, as is the case in the stabilization problem.) We then set $x_1 := x(t_1)$ and restart the system at $t = t_1$ with control value $k(x_1)$:

$$\dot{x}(t) = f(x(t), k(x_1)), \quad x(t_1) = x_1, \quad t_1 \leq t \leq t_2,$$

and so on in this fashion. The trajectory x that results from this procedure is an actual state trajectory corresponding to a piecewise constant open-loop control; thus it is a physically meaningful one. When results are couched in terms of π-trajectories, the issue of defining a solution concept for discontinuous differential equations is effectively sidestepped. Making the diameter of the partition smaller corresponds to increasing the sampling rate in the implementation.

We remark that the use of possibly discontinuous feedback has arisen in other contexts. In linear time-optimal control, one can find discontinuous feedback syntheses as far back as the classical book of Pontryagin et al [17]; in these cases the feedback is invariably piecewise constant relative to certain partitions of state space, and solutions either follow the switching surfaces or cross them transversally, so the issue of defining the solution in other than a classical sense does not arise. Somewhat related to this is the approach that defines a multivalued feedback law [4]. In stochastic control, discontinuous feedbacks are the norm, with the solution understood in terms of stochastic differential equations. In a similar vein, in the control of certain linear partial differential equations, discontinuous feedbacks can be interpreted in a distributional sense. These cases are all unrelated to the one under discussion. We remark too that the use of discontinuous pursuit strategies in differential games [15] is well-known, together with examples to show that, in general, it is not possible to achieve the result of a discontinuous optimal strategy to within any tolerance by means of a continuous strategy (so there can be a positive unbridgeable gap between the performance of continuous and discontinuous feedbacks).

We can use the π-trajectory formulation to implement feedbacks for either guidability or stabilization (see [12]); we limit attention here to the latter issue.

It is natural to say that a feedback $k(x)$ (continuous or not) *stabilizes* the system in the sample-and-hold sense provided that for every initial value x_0, for all $\epsilon > 0$, there exists $\delta > 0$ and $T > 0$ such that if the diameter of the partition π is less than δ, then the corresponding π-trajectory x beginning at x_0 satisfies

$$\|x(t)\| \leq \epsilon \quad \forall\, t \geq T.$$

The following theorem is proven in [8]:

Theorem 5. *The system is open loop globally asymptotically controllable if and only if there exists a (possibly discontinuous) feedback $k : \mathbb{R}^n \to \mathcal{U}$ which stabilizes it in the sample-and-hold sense.*

The proof of the theorem actually yields precise estimates regarding how small the step size diam (π) must be for a prescribed stabilization tolerance to ensue, and of the resulting stabilization time, in terms of the given data. These estimates are uniform on bounded sets of initial conditions, and are a consequence of the method of proximal aiming. The latter, which can be viewed as a geometric version of the Lyapunov technique, appears to be difficult to implement in practice, however. One of our principal goals is to show how stabilizing feedbacks can be defined much more conveniently if one has at hand a sufficiently regular Lyapunov function.

The Smooth Case

We begin with the case in which a C^1 smooth Lyapunov function exists, and show how the natural 'pointwise feedback' described below stabilizes the system (in the sample-and-hold sense). For $x \neq 0$, we define $k(x)$ to be any element $u \in \mathcal{U}$ satisfying

$$\langle \nabla V(x), f(x,u) \rangle \leq -W(x).$$

Note that at least one such u does exist, in light of the infinitesimal decrease condition. We mention two more definitions that work: take u to be the element minimizing the inner product above over $\mathcal{U}$, or take any $u \in \mathcal{U}$ satisfying $\langle \nabla V(x), f(x,u) \rangle \leq -W(x)/2$.

Theorem 6. *The pointwise feedback k described above stabilizes the system in the sense of closed-loop system sampling.*

We proceed to sketch the elementary proof of this theorem, which we deem to be a basic result in the theory of control Lyapunov functions.

We begin with a remark: for any $R > 0$, there exists $\delta_R > 0$ such that for all $\alpha \in B(0,R)$ and for all $u \in \mathcal{U}$, any solution x of $\dot{x} = f(x,u),\ x(0) = \alpha$ satisfies

$$|x(t)| \leq R + 1 \quad \forall t \in [0, \delta_R]$$

(this is a simple consequence of the linear growth hypothesis and Gronwall's Lemma).

Now let positive numbers r and ε be given; we show that for any $\alpha \in B(0,r)$ there is a trajectory x beginning at α that enters the ball $B(0,\varepsilon)$ in finite time.

Let $R \geq r$ be chosen so that

$$V(x) \leq \max_{B(0,r)} V \Longrightarrow x \in B(0,R).$$

For simplicity, let us assume that ∇V is locally Lipschitz (as otherwise, the argument is carried out with a modulus of continuity). We proceed to choose $K > 0$ such that for every $u \in \mathcal{U}$, the function

$$x \mapsto \langle \nabla V(x), f(x,u) \rangle$$

is Lipschitz on $B(0, R+1)$ with constant K, together with positive numbers M and m satisfying

$$|f(x,u)| \leq M \quad \forall x \in B(0, R+1),\ \forall u \in \mathcal{U},$$

and

$$W(x) \geq\ m \quad \forall x \in B(0, R+1) \backslash B(0, \varepsilon).$$

Now let $\pi = \{t_i\}_{0 \leq i \leq N}$ be a partition (taken to be uniform for simplicity) of step size $\delta \leq \delta_R$, of an interval $[0,T]$, where $t_0 = 0, t_N = T, T = N\delta$. We apply the pointwise feedback k relative to this partition, and with initial condition $x(0) = x_0 := \alpha$. We proceed to compare the values of V at the first two nodes:

$$V(x(t_1)) - V(x(t_0)) = \langle \nabla V(x(t^*)), \dot{x}(t^*) \rangle\, (t_1 - t_0)$$

(by the mean value theorem, for some $t^* \in (0, \delta)$)

$$\begin{aligned} &= \langle \nabla V(x(t^*)), f(x(t^*), k(t_0)) \rangle\, (t_1 - t_0) \\ &\leq \langle \nabla V(x(t_0)), f(x(t_0), k(t_0)) \rangle\, (t_1 - t_0) \\ &\qquad + K\, |x(t^*) - x(t_0)|\, (t_1 - t_0) \end{aligned}$$

(by the Lipschitz condition)

$$\leq -W(x(t_0))\delta + KM\delta^2$$

(by the way k is defined)

$$\leq -m\delta + KM\delta^2.$$

Note that these estimates apply because $x(t_1)$ and $x(t^*)$ remain in $B(0, R + 1)$, and, in the case of the last step, provided that x_0 does *not* lie in the ball $B(0, \varepsilon)$. Inspection of the final term above shows that if δ is taken less than $m/(2KM)$, then the value of V between the two nodes has decreased by at least $m\delta/2$. It follows from the definition of R that $x(t_1) \in B(0, R)$. Consequently, the same argument as above can be applied to the next partition subinterval, and so on. Iteration then yields:

$$V(x(N\delta)) - V(x_0) \leq -mN/2.$$

This will contradict the nonnegativity of V when N exceeds $2V(x_0)/m$, so it follows that the argument must fail at some point, which it can only do when a node $x(t_i)$ lies in $B(0, \varepsilon)$.

This proves that any sample-and-hold trajectory generated by the feedback k enters $B(0, \varepsilon)$ in a time that is bounded above in a way that depends only upon ε and $|\alpha|$ (and V), provided only that the step size is sufficiently small, as measured in a way that depends only on $|\alpha|$.

That k stabilizes the system in the sense of closed-loop system sampling now follows.

Remark.

Rifford [20] has shown that the existence of a smooth Lyapunov pair is equivalent to the existence of a locally Lipschitz one satisfying weak decrease in the sense of generalized gradients (rather than proximal subgradients), which in turn is equivalent to the existence of a stabilizing feedback in the Filippov (rather than sample-and-hold) sense.

5 Semiconcavity

The 'Right' Regularity for Lyapunov Functions

We have seen that a smooth Lyapunov function generates a stabilizing feedback in a very simple and natural way. But since a smooth Lyapunov function does not necessarily exist, we still require a way to handle the general case. It turns out that the two issues can be reconciled through the notion of *semiconcavity.* This is a certain regularity property (not implying smoothness) which can always be guaranteed to hold for some Lyapunov function (if the system is globally asymptotically controllable, of course), and which permits a natural extension of the pointwise definition of a stabilizing feedback.

A function $\phi : \mathbb{R}^n \to \mathbb{R}$ is said to be (globally) semiconcave provided that for every ball $B(0, r)$ there exists $\gamma = \gamma(r) \geq 0$ such that the function $x \mapsto \phi(x) - \gamma|x|^2$ is (finite and) concave on $B(0, r)$. (Hence ϕ is locally the sum of a concave function and a quadratic one.) Observe that any function of class C^2 is semiconcave; also, any semiconcave function is locally Lipschitz, since both concave functions and smooth functions have that property. (There is a local definition of semiconcavity that we omit for present purposes.)

Semiconcavity is an important regularity property in partial differential equations; see for example [5]. The fact that the semiconcavity of a Lyapunov function V turns out to be useful in stabilization is a new observation, and may be counterintuitive: V often has an interpretation in terms of energy, and it may seem more appropriate to seek a *convex* Lyapunov function V. We proceed now to explain why semiconcavity is a highly desirable property, and why a convex V would be of less interest (unless it were smooth, but then it would be semiconcave too).

Recall the ideal case discussed above, in which (for a smooth V) we select a function $k(x)$ such that

$$\langle \nabla V(x), f(x, k(x)) \rangle \leq -W(x).$$

How might this appealing idea be adapted to the case in which V is nonsmooth? We cannot use the proximal subdifferential $\partial_P V(x)$ directly, since it may be empty for 'many' x. We are led to consider the *limiting subdifferential* $\partial_L V(x)$, which, when V is continuous, is defined by applying a natural limiting operation to $\partial_P V$:

$$\partial_L V(x) := \left\{ \zeta = \lim_{i \to \infty} \zeta_i : \zeta_i \in \partial_P V(x_i), \lim_{i \to \infty} x_i = x \right\}.$$

It follows readily that, when V is locally Lipschitz, $\partial_L V(x)$ is nonempty for all x. By passing to the limit, the weak infinitesimal decrease condition for proximal subgradients implies the following:

$$\min_{u \in U} \langle f(x,u), \zeta \rangle \leq -W(x) \; \forall \zeta \in \partial_L V(x), \forall x \neq 0.$$

Accordingly, let us consider the following idea: for each $x \neq 0$, choose some element $\zeta \in \partial_L V(x)$, then choose $k(x) \in U$ such that

$$\langle f(x, k(x)), \zeta \rangle \leq -W(x).$$

Does this lead to a stabilizing feedback, when (of course) the discontinuous differential equation is interpreted in the sample-and-hold sense? When V is smooth, the answer is 'yes', as we have seen. But when V is merely locally Lipschitz, a certain 'dithering' phenomenon may arise to prevent k from being stabilizing. However, if V is semiconcave (on $\mathbb{R}^n \backslash \{0\}$), this does not occur, and stabilization is guaranteed. This accounts in part for the desirability of a semiconcave Lyapunov function, and the importance of knowing one always exists.

The proof that the pointwise feedback defined above is stabilizing hinges upon the following fact in nonsmooth analysis:

Lemma.

Suppose that $V(x) = g(x) + \gamma |x|^2$, where g is a concave function. Then for any $\zeta \in \partial_L V(x)$, we have

$$V(y) - V(x) \le \langle \zeta, y - x \rangle + \gamma |y - x|^2 \ \forall\, y.$$

The proof of Theorem 6 can be mimicked when V is semiconcave rather than smooth, by invoking the 'decrease property' described in the lemma at a certain point. The essential step remains the comparison of the values of V at successive nodes; for the first two, for example, we have

$$V(x(t_1)) - V(x(t_0)) \le \langle \zeta, x(t_1) - x(t_0) \rangle (t_1 - t_0) + \gamma |x(t_1) - x(t_0)|^2$$

(where $\zeta \in \partial_L V(x(t_0)$, by the lemma)

$$= \langle \zeta, f(x(t^*), k(t_0)) \rangle (t_1 - t_0) + \gamma |x(t_1) - x(t_0)|^2$$

(for some $t^* \in (t_0, t_1)$, by the mean value theorem)

$$\le \langle \zeta, f(x(t_0), k(t_0)) \rangle (t_1 - t_0) \\ + K_V K_f |x(t^*) - x(t_0)| (t_1 - t_0) + \gamma M^2 \delta^2$$

(where K_V and K_f are suitable Lipschitz constants for V and f)

$$\le -W(x(t_0))\delta + K_V K_f M(1 + \gamma M)\delta^2,$$

by the way k is defined. Then, as before, a decrease in the value of V can be guaranteed by taking δ sufficiently small, and the proof proceeds as before. (The detailed argument must take account of the fact that V is only semiconcave away from the origin, and that a parameter γ as used above is available only on bounded subsets of $\mathbb{R}^n \backslash \{0\}$.)

6 Finite-Time Guidability

So far we have been concerned with possibly asymptotic approach to the origin. There is interest in being able to assert that the origin can be reached in finite time. If such is the case from any initial condition, then we say that the system is *globally guidable* in finite time (to 0). There is a well-studied local version of this property that bears the name *small-time local controllability* (STLC for short). A number of verifiable criteria exist which imply that the system has property STLC, which is stronger than Lyapunov stability; see [3].

Theorem 7. *The system is globally guidable in finite time if and only if there exists a general Lyapunov pair (V, W) with V finite-valued and $W \equiv 1$. If the system has the property STLC, then it is globally guidable in finite time iff there exists a Lyapunov pair (V, W) with V continuous and $W \equiv 1$.*

The proof of the theorem revolves around the much-studied minimal time function $T(\cdot)$. If the system is globally guidable in finite time, then $(T, 1)$ is the required Lyapunov pair: positive definiteness and properness are easily checked, and weak infinitesimal decrease follows from the (now well-known) fact that T satisfies the proximal Hamilton-Jacobi equation

$$h(x, \partial_P T(x)) + 1 = 0, \quad x \neq 0.$$

This is equivalent to the assertion that T is a viscosity solution of a related equation; see [10].

The sufficiency in the first part of the theorem follows much as in the proof of Theorem 2: we deduce the existence of a trajectory x for which $V(x(t)) + t$ is nonincreasing as long as $x(t) \neq 0$; this implies that $x(\tau)$ equals 0 for some $\tau \leq V(x(0))$. As for the second part of the theorem, it follows from the fact that, in the presence of STLC, the minimal time function is continuous.

7 An Equivalence Theorem

The following result combines and summarizes many of the ones given above concerning the regularity of Lyapunov functions and the presence of certain system properties.

Theorem 8. *The following are equivalent:*

1. *The system is open-loop globally asymptotically controllable.*
2. *There exists a continuous Lyapunov pair (V, W).*
3. *There exists a locally Lipschitz Lyapunov pair (V, W) with V semiconcave on $\mathbb{R}^n \backslash \{0\}$.*
4. *There exists a globally stabilizing sample-and-hold feedback.*

If, a priori, the system has Lyapunov stability at 0, then the following item may be added to the list:

5. *There exists for each positive ε a locally Lipschitz ε-Lyapunov function.*

If, a priori, the system has the property STLC, the following further item may be added to the list:

6. *There exists a continuous Lyapunov pair (V, W) with $W \equiv 1$.*

In this last case, the system is globally guidable in finite time.

8 Some Related Issues

Robustness.

It may be thought in view of the above that there is no advantage in having a smooth Lyapunov function, except the greater ease of dealing with derivatives rather than subdifferentials. In any case, stabilizing feedbacks will be discontinuous; and they can be conveniently defined in a pointwise fashion if the Lyapunov function is semiconcave. In fact, however, there is a robustness consequence to the existence of a smooth Lyapunov function.

The robustness of which we speak here is with respect to possible error e in state measurement when the feedback law is implemented: we are at x, but measure the state as $x + e$, and therefore apply the control $k(x + e)$ instead of the correct value $k(x)$. When k is continuous, then for e small enough this error will have only a small effect: the state may not approach the origin, but will remain in a neighborhood of it, a neighborhood that shrinks to the origin as e goes to zero; that is, we get practical stabilization. This feature of continuous feedback laws is highly desirable,and in some sense essential, since some imprecision seems inevitable in practice. One might worry that a discontinuous feedback law might not have this robustness property, since an arbitrarily small but nonzero e could cause $k(x)$ and $k(x + e)$ to differ significantly.

It is a fact that the (generally discontinuous) feedback laws constructed above do possess a *relative robustness* property: if, in the sample-and-hold implementation, the measurement error is at most of the same order of magnitude as the partition diameter, then practical stabilization is obtained. To put this another way, the step size may have to be big enough relative to the potential errors (to avoid dithering, for example). At the same time, the step size must be sufficiently small for stabilization to take place, so there is here a conflict that may or may not be reconcilable. It appears to us to be a great virtue of the sample-and-hold method that it allows, apparently for the first time, a precise error analysis of this type.

There is another, stronger type of robustness (called absolute robustness), in which the presence of small errors preserves practical stabilization independently of the step size. Ledyaev and Sontag [16] have shown that there exists an absolutely robust stabilizing feedback if and only if there exists a smooth Lyapunov pair. This, then, is an advantage that such systems have. Recall that the nonholonomic integrator, though stabilizable, does not admit a smooth Lyapunov function and hence fails to admit an absolutely robust stabilizing feedback.

State constraints.

There are situations in which the state x is naturally constrained to lie in a given closed set S, so that in steering the state to the origin, we must respect the condition $x(t) \in S$. The same questions arise as in the unconstrained case: is the possibility of doing this in the open-loop sense characterized by some kind of Lyapunov function, and would such a function lead to the definition of a stabilizing feedback that respects the state constraint? The more challenging case is that in which the origin lies on the boundary of S, but the case in which 0 lies in the interior of S is also of interest, since it localizes around the origin the global and constraint-free situation that has been the focus of this article.

An important consideration in dealing with state constraints is to identify a class of sets S for which meaningful results can be obtained. Recently Clarke and Stern [13, 12], for what appears to have been the first time, have extended many of the Lyapunov and stabilization methods discussed above to the case of state constraints specified by a set S which is *wedged* (see [10]). This rather large class of sets includes smooth manifolds with boundaries and convex bodies (as well as their closed complements). A set is wedged (or epi-Lipschitz) when its (Clarke) tangent cone at each point has nonempty interior, which is equivalent to the condition that locally (and after a change of coordinates), it is the epigraph of a Lipschitz function.

A further hypothesis is made regarding the consistency of the state constraint with the dynamics of the system: for every nonzero vector ζ in the (Clarke) normal cone to a point $x \in \operatorname{bdry} S$, there exists $u \in \mathcal{U}$ such that $\langle f(x,u), \zeta \rangle < 0$. Thus an 'inward-pointing' velocity vector is always available.

Under these conditions, and in terms of suitably defined extensions to the state-constrained case of the underlying definitions, one can prove an equivalence between open-loop controllability, closed-loop stabilization, and the existence of more or less regular (and in particular semiconcave) Lyapunov functions.

Regular and essentially stabilizing feedbacks.

In view of the fact that a GAC system need not admit a continuous stabilizing feedback, the question arises of the extent to which the discontinuities can be minimized. Ludovic Rifford has exploited the existence of a semiconcave Lyapunov function, together with both proximal and generalized gradient calculus, to show that when the system is affine in the control, there exists a stabilizing feedback whose discontinuities form a set of measure zero. Moreover, the discontinuity set is *repulsive* for the trajectories generated by the feedback: the trajectories lie in that set at most initially. This means that in applying the feedback, the solutions can be understood in the usual Carathéodory sense; robustness ensues as well. In the case of planar systems, Rifford has gone on to settle an open problem of Bressan by classifying the types of discontinuity that must occur in stabilizing feedbacks.

More recently, Rifford [21] has introduced the concept of *stratified* semiconcave Lyapunov functions, and has shown that every GAC system must admit one. Building upon this, he proves that there then exists a *smooth* feedback which *almost* stabilizes the system (that is, from almost all initial values). This highly interesting result is presented in Rifford's article in the present collection.

References

1. Aubin J (1991) Viability theory. Birkhäuser Boston

2. Bacciotti A, Rosier L (2001) Lyapunov functions and stability in control theory. Springer-Verlag London
3. Bardi M, Capuzzo-Dolcetta I (1997) Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations. Birkhäuser Boston
4. Berkovitz L (1989) Optimal feedback controls, SIAM J Control Optim 27:991–1006
5. Cannarsa P, Sinestrari C (2003) Semiconcave functions Hamilton-Jacobi equations and optimal control. Birkhäuser Boston to appear
6. Clarke F (2001) Nonsmooth analysis in control theory a survey, European J Control 7:63–78
7. Clarke F, Ledyaev Y, Rifford L, Stern R (2000) Feedback stabilization and Lyapunov functions, SIAM J Control Optim 39:25–48
8. Clarke F, Ledyaev Y, Sontag E, Subbotin A (1997) Asymptotic controllability implies feedback stabilization, IEEE Trans Aut Control 42:1394–1407
9. Clarke F, Ledyaev Y, Stern R (1998) Asymptotic stability and smooth Lyapunov functions, J Differential Equations 149:69–114
10. Clarke F, Ledyaev Y, Stern R, Wolenski P (1998) Nonsmooth analysis and control theory. Springer-Verlag New York
11. Clarke F, Ledyaev Y, Subbotin A (1997) The synthesis of universal feedback pursuit strategies in differential games, SIAM J Control Optim 35:552–561
12. Clarke F, Stern R (2003) CLF and feedback characterizations of state constrained controllability and stabilization, preprint
13. Clarke F, Stern R (2003) State constrained feedback stabilization, SIAM J Control Optim 42:422-441
14. Freeman R, Kokotović P (1996) Robust nonlinear control design state space and Lyapunov techniques. Birkhäuser Boston
15. Krasovskii N, Subbotin A (1988) Game-theoretical control problems. Springer-Verlag New York
16. Ledyaev Y, Sontag E (1999) A Lyapunov characterization of robust stabilization, Nonlinear Analysis 37:813–840
17. Pontryagin L, Boltyanskii R, Gamkrelidze R, Mischenko E (1962) The mathematical theory of optimal processes. Wiley-Interscience New York
18. Rifford L (2000) Existence of Lipschitz and semiconcave control-Lyapunov functions, SIAM J Control Optim 39:1043–1064
19. Rifford L (2000) Problèmes de stabilisation en théorie du contrôle. PhD Thesis Université Claude Bernard Lyon I
20. Rifford L (2001) On the existence of nonsmooth control-Lyapunov functions in the sense of generalized gradients, ESAIM Control Optim Calc Var 6:593–611
21. Rifford L (2003) Stratified semiconcave control-Lyapunov functions and the stabilization problem, preprint
22. Sontag E (1983) A Lyapunov-like characterization of asymptotic controllability, SIAM J Control Optim 21:462–471
23. Sontag E (1999) Stability and stabilization discontinuities and the effect of disturbances. In: Clarke F, Stern R (eds) Nonlinear analysis differential equations and control. Kluwer Acad Publ Dordrecht