

LINEAR-CONVEX CONTROL AND DUALITY

R.T. Rockafellar

*Department of Mathematics, University of Washington
Seattle, WA 98195-4350, USA
Email: rtr@math.washington.edu*

R. Goebel

*3518 NE 42 St., Seattle, WA 98105, USA
Email: goebel@math.washington.edu*

An optimal control problem with linear dynamics and convex but not necessarily quadratic and possibly infinite-valued or nonsmooth costs can be analyzed in an appropriately formulated duality framework. The paper presents key elements of such a framework, including a construction of a dual optimal control problem, optimality conditions in open loop and feedback forms, and a relationship between primal and dual optimal values. Some results on differentiability and local Lipschitz continuity of the gradient of the optimal value function associated with a convex optimal control problem are included.

Keywords: Optimal control, generalized problem of Bolza, convex value function, duality, Hamilton-Jacobi theory

1. Introduction

The classical Linear Quadratic Regulator problem (\mathcal{LQR}) involves minimizing a quadratic cost subject to linear dynamics. See Ref. 1 for a detailed exposition. When the system is subject to constraints or actuator saturation, or when nonquadratic costs or barrier functions are involved, the linear and quadratic techniques applicable to \mathcal{LQR} are no longer adequate. Often though, the optimal control problem one needs to solve is convex.

Much like what is appreciated in convex optimization, convex structure of an optimal control problem makes available techniques based on an appropriately formulated duality framework. Existence results, optimality conditions, and sensitivity to perturbations can be analyzed in such a framework even if constraints or nonsmooth costs are involved.

Following the contributions made over the years by the first author and

his students and former students, we outline in this paper some key aspects of a duality framework for the study of general finite-horizon and infinite-horizon convex optimal control problems. For simplicity of presentation, but also with control engineering applications in mind, we specialize the key results to a Linear-Convex Regulator (\mathcal{LCR}): a problem in which the quadratic costs as in the \mathcal{LQR} are replaced by general convex functions.

Besides expository purposes, the choice of the duality related results we state is motivated by the task of analyzing the regularity properties of optimal value functions associated with a control problem. The properties of our interest include finiteness, coercivity, differentiability, strict convexity, and finally, continuity of the gradient. In particular, based on some recent results of the authors on the duality between local strong convexity of a convex function and the local Lipschitz continuity of the gradient of its conjugate, we identify some mild assumptions on a convex optimal control problem that guarantee that both the primal value function and the dual value function have locally Lipschitz continuous gradients.

2. Duality in finite horizon optimal control

The full power of convex duality techniques was brought to the calculus of variations problems on finite time horizons by Rockafellar in Ref. 2, and led, in particular, to very general existence results in Ref 3. We point the reader to Ref. 2 for an overview of other early developments relying, to an extent, on convex duality. Regarding linear-quadratic optimal control problems with constraints, a duality framework was suggested in Ref. 4 and led, among other things, to simple optimality conditions for problems where some state constraints are present; see Ref 5. The duality tools were specialized to the study of value functions, in the Hamilton-Jacobi framework, by Rockafellar and Wolenski in Refs. 6,7. Comparison to related Hamilton-Jacobi developments can be found in Refs. 6,8.

2.1. *The (finite horizon) Linear-Convex Regulator*

We will focus on the following optimal control problem:

$$\text{minimize } \int_{\tau}^T q(y(t)) + r(u(t)) dt + g(x(T)) \quad (1)$$

over all integrable control functions $u : [\tau, T] \rightarrow \mathbb{R}^k$, subject to linear dynamics

$$\dot{x}(t) = Ax(t) + Bu(t), \quad y(t) = Cx(t), \quad (2)$$

and the initial condition

$$x(\tau) = \xi. \quad (3)$$

The absolutely continuous arc $x : [\tau, T] \rightarrow \mathbb{R}^n$ denotes the state and $y : [\tau, T] \rightarrow \mathbb{R}^m$ denotes the output of the linear system (2). The matrices A , B , and C are any matrices of appropriate dimensions; we will explicitly make assumptions of observability or controllability when needed.

When the functions q , r , and g are quadratic and given by symmetric and positive definite matrices the problem described above is the classical Linear-Quadratic Regulator; see Ref. 1 for background and detailed analysis. Here, we will allow q , r , and g to be more general convex functions. Consequently, we will call the problem defined by (1), (2), (3) a Linear-Convex Regulator, or just \mathcal{LCR} . When an explicit reference to initial conditions is needed, we will speak of the $\mathcal{LCR}(\tau, \xi)$ problem.

Below, $\overline{\mathbb{R}} = [-\infty, +\infty]$ and a function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is called *coercive* if $\lim_{|x| \rightarrow \infty} f(x)/|x| = \infty$.

Standing Assumptions. *The functions $q : \mathbb{R}^m \rightarrow \overline{\mathbb{R}}$, $r : \mathbb{R}^k \rightarrow \overline{\mathbb{R}}$, and $g : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ are elements of the class \mathcal{C} , where*

$$\mathcal{C} = \bigcup_{i=1}^{\infty} \left\{ f : \mathbb{R}^i \rightarrow \overline{\mathbb{R}} \mid \begin{array}{l} f \text{ is convex, lower semicontinuous,} \\ f(0) = 0, f(x) \geq 0 \text{ for all } x \in \mathbb{R}^i \end{array} \right\}. \quad (4)$$

Furthermore, $q(y)$ is finite for all $y \in \mathbb{R}^m$ while r is coercive.

In what follows, we will say that $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is proper if it never equals $-\infty$ and is finite somewhere.

2.2. Convex conjugate functions

For any convex, lower semicontinuous (lsc.), and proper function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$, its *conjugate function* $f^* : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is defined by

$$f^*(p) = \sup_{x \in \mathbb{R}^n} \{p \cdot x - f(x)\}.$$

The function f^* is itself a convex, lsc., and proper function, and the function conjugate to it is f . It is also immediate from the definition of f^* that $f \in \mathcal{C}$ if and only $f^* \in \mathcal{C}$; this will in particular imply that our assumptions on \mathcal{LCR} and on the problem dual to \mathcal{LCR} are symmetric.

A basic example of a pair of conjugate functions is

$$f(x) = \frac{1}{2}x \cdot Qx, \quad f^*(p) = \frac{1}{2}p \cdot Q^{-1}p, \quad (5)$$

where Q is a symmetric and positive definite matrix. Another is

$$f(x) = \begin{cases} \frac{1}{2}x \cdot Qx & \text{if } x \in X \\ +\infty & \text{if } x \notin X \end{cases}, \quad f^*(p) = \sup_{x \in X} \left\{ p \cdot x - \frac{1}{2}x \cdot Qx \right\} \quad (6)$$

for Q as above and a closed convex and nonempty set $X \subset \mathbb{R}^n$. Requiring that $f(0) = 0$ implies that $0 \in X$. If $0 \in \text{int } X$, then f^* is quadratic on some neighborhood of 0, and given there by $\frac{1}{2}p \cdot Q^{-1}p$. In general, f^* is differentiable and ∇f^* is Lipschitz continuous; related properties of not necessarily quadratic convex functions will be discussed in Section 3.

A fundamental property of the pair f and f^* , that is the basis for Euler-Lagrange and Hamiltonian optimality conditions for convex optimal control problems, is that for any $x, p \in \mathbb{R}^n$,

$$f(x) + f^*(p) \geq x \cdot p, \quad (7)$$

while $f(x) + f^*(p) = x \cdot p$ if and only if $p \in \partial f(x)$ if and only if $p \in \partial f^*(x)$. Here, ∂f is the *subdifferential* of the convex function f :

$$\partial f(x) = \{p \in \mathbb{R}^n \mid f(x') \geq f(x) + p \cdot (x' - x) \text{ for all } x' \in \mathbb{R}^n\}.$$

A standard reference for the material just presented is Ref. 9.

2.3. General duality framework

Many of the results we will state or use were developed not in an optimal control setting, but rather, in the framework of duality for calculus of variations problems. We briefly outline this framework.

Given convex, lsc., and proper functions $L, l : \mathbb{R}^{2n} \rightarrow \overline{\mathbb{R}}$, consider

$$\tilde{L}(p, w) = L^*(w, p), \quad \tilde{l}(p_\tau, p_T) = l^*(p_\tau, -p_T),$$

and two generalized problems of Bolza type: the primal problem \mathcal{P} of minimizing, over all absolutely continuous $x : [\tau, T] \rightarrow \mathbb{R}^n$, the cost functional

$$\int_\tau^T L(x(t), \dot{x}(t)) dt + l(x(\tau), x(T)), \quad (8)$$

and the dual problem $\tilde{\mathcal{P}}$ of minimizing, over all absolutely continuous $p : [\tau, T] \rightarrow \mathbb{R}^n$, the (dual) cost functional

$$\int_\tau^T \tilde{L}(p(t), \dot{p}(t)) dt + \tilde{l}(p(\tau), p(T)). \quad (9)$$

Note that the problem dual to $\tilde{\mathcal{P}}$ is the original \mathcal{P} .

Directly from (7), it follows that it is always the case that $\inf(\mathcal{P}) \geq -\inf(\tilde{\mathcal{P}})$, and that any absolutely continuous $p : [\tau, T] \rightarrow \mathbb{R}^n$ provides a

lower bound on $\inf(\mathcal{P})$ through (9), while any absolutely continuous $x : [\tau, T] \rightarrow \mathbb{R}^n$ provides a lower bound on $\inf(\tilde{\mathcal{P}})$ through (8). Under some mild assumptions, given in Refs. 2,3 for a more general time-dependant case and specialized by Rockafellar and Wolenski in Ref. 6 to the autonomous case, the following holds:

$$-\infty < \inf(\mathcal{P}) = -\inf(\tilde{\mathcal{P}}) < +\infty, \quad (10)$$

and moreover, optimal solutions for both the primal and the dual problem exist. Regarding L , Ref. 6 required that: (i) the set $F(x) = \{v \mid L(x, v) < \infty\}$ be nonempty for all x , and there exist a constant ρ such that $\text{dist}(0, F(x)) \leq \rho(1 + |x|)$ for all x , and (ii) there exist constants α and β and a coercive, proper, nondecreasing function θ on $[0, \infty)$ such that $L(x, v) \geq \theta(\max\{0, |v| - \alpha|x|\}) - \beta|x|$ for all x and v . A particularly attractive feature of these assumptions, besides their generality, is that L satisfies them if and only if \tilde{L} does.

The Linear-Convex Regulator can be reformulated in the generalized Bolza framework. (For this, and other equivalences between optimal control problem formulations, see Chapter 1.3 in Ref 10.) To this end, for each fixed $\tau \leq T$, $\xi \in \mathbb{R}^n$, one considers

$$\begin{aligned} L(x, v) &= q(Cx) + \min_u \{r(u) \mid Ax + Bu = v\}, \\ l(x_\tau, x_T) &= \delta_\xi(x_\tau) + g(x_T), \end{aligned} \quad (11)$$

where δ_ξ is the indicator of ξ : $\delta_\xi(\xi) = 0$ while $\delta_\xi(x) = \infty$ if $x \neq \xi$. If a solution $x : [\tau, T] \rightarrow \mathbb{R}^n$ to the resulting Bolza problem $\mathcal{P}(\tau, \xi)$ is found, an optimal control $u : [\tau, T] \rightarrow \mathbb{R}^k$ for \mathcal{LCR} can be then recovered, at almost every $t \in [\tau, T]$, as the minimizer of $r(u)$ over all u such that $Ax(t) + Bu = \dot{x}(t)$. Given such L and l , the dual Bolza problem $\tilde{\mathcal{P}}(\tau, \xi)$ is defined by

$$\begin{aligned} \tilde{L}(p, w) &= r^*(B^*p) + \min_z \{q^*(z) \mid -A^*p + C^*z = w\}, \\ \tilde{l}(p_\tau, p_T) &= \xi \cdot p_\tau + g^*(-p_T). \end{aligned}$$

Our Standing Assumptions on q , r , and g guarantee that L and l as above (and equivalently, \tilde{L} and \tilde{l}) meet the growth conditions of Ref. 6 and consequently, that (10) holds. We note that in most cases, and even if q and r are quadratic functions, L in (11) does not satisfy the classical coercivity conditions: $L(x, v)$ is not bounded below by a coercive function of v . This can easily be seen in the case of B being an identity: then $L(x, v) = q(Cx) + r(v - Ax)$.

2.4. The primal and the dual value functions

The (primal) optimal value function $V : (-\infty, T] \times \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is defined, for each $(\tau, \xi) \in (-\infty, T] \times \mathbb{R}^n$, as the infimum in the problem $\mathcal{LCR}(\tau, \xi)$. For the quadratic case (of q , r , and g quadratic), $V(\tau, \cdot)$ is quadratic for each $\tau \leq T$. In general, it is immediate that $V(\tau, \cdot)$ is a convex function. More precisely, Theorem 2.1 and Corollary 7.6. in Ref. 6 state the following:

Theorem 2.1 (value function – basic regularity).

- For each $\tau \leq T$, $V(\tau, \cdot)$ is an element of \mathcal{C} . If g is finite-valued, then so is $V(\tau, \cdot)$.
- $V(\tau, \cdot)$ depends epi-continuously on $\tau \in (-\infty, T]$.

Epi-continuity above means that the epigraphs of $V(\tau, \cdot)$, i.e., the sets $\{(\xi, \alpha) \in \mathbb{R}^{n+1} \mid \alpha \geq V(\tau, \xi)\}$, depend continuously on τ . Such concept of continuity of V takes into account the fact that the effective domains of $V(\tau, \cdot)$ may depend on τ . (See Chapters 4 and 5 in Ref. 11 for details on set convergence and continuity of set-valued mappings.)

The value function V can be equivalently defined, at each (τ, ξ) , as the infimum in the generalized problem of Bolza $\mathcal{P}(\tau, \xi)$ related to $\mathcal{LCR}(\tau, \xi)$ through (11). From (10) it follows that $V(\tau, \xi) = \inf(\mathcal{P}(\tau, \xi)) = -\inf(\tilde{\mathcal{P}}(\tau, \xi))$. Since $l(p_\tau, p_T) = \inf_{\pi \in \mathbb{R}^n} \{p_\tau \cdot \pi + \delta_{p_\tau}(\pi)\}$ and $\delta_{p_\tau}(\pi) = \delta_\pi(p_\tau)$, one obtains that $V(\tau, \xi)$ equals

$$\begin{aligned} & -\inf \left\{ \inf_{\pi \in \mathbb{R}^n} \{ \xi \cdot \pi + \delta_{p(\tau)}(\pi) \} + g^*(-p(T)) + \int_\tau^T \tilde{L}(p(t), \dot{p}(t)) dt \right\} \\ & = \sup_{\pi \in \mathbb{R}^n} \left\{ -\xi \cdot \pi - \inf \left\{ \delta_\pi(p(\tau)) + g^*(-p(T)) + \int_\tau^T \tilde{L}(p(t), \dot{p}(t)) dt \right\} \right\}. \end{aligned}$$

The first and the last infimum above are taken over all arcs $p : [\tau, T] \rightarrow \mathbb{R}^n$. The presence of the term $\delta_\pi(p(\tau))$ ensures that the last infimum can be taken only over those arcs $p(\cdot)$ for which $p(\tau) = \pi$. This suggests that the last infimum defines a dual value function, parameterized by τ and π . We now make the formal definitions, following the ideas of Ref. 4 regarding the structure of a dual optimal control problem, and of Ref. 6 regarding the dual value function.

The optimal control problem dual to \mathcal{LCR} is as follows:

$$\text{minimize} \quad \int_\tau^T r^*(s(t)) + q^*(z(t)) dt + g^*(-p(T)) \quad (12)$$

over all integrable control functions $z : [\tau, T] \rightarrow \mathbb{R}^m$, subject to linear dynamics

$$\dot{p}(t) = -A^*p(t) + C^*z(t), \quad s(t) = B^*p(t), \quad (13)$$

(here, A^* denotes the transpose of A , etc.) and the initial condition

$$p(\tau) = \pi. \quad (14)$$

The absolutely continuous arc $p : [\tau, T] \rightarrow \mathbb{R}^n$ denotes the state and $s : [\tau, T] \rightarrow \mathbb{R}^k$ denotes the output of the linear system (13). We will denote this problem by $\widetilde{\mathcal{LCR}}$, and when a direct reference to initial conditions is needed, by $\mathcal{LCR}(\tau, \pi)$.

Note that the functions r^* , q^* , and $g^*(-\cdot)$ are in the class \mathcal{C} (recall (4)). Furthermore, r^* is finite-valued (since r is coercive) and q^* is coercive (since q is finite-valued). Thus $\widetilde{\mathcal{LCR}}$ has the same growth properties that \mathcal{LCR} has, based on the Standing Assumptions. Note also that if \mathcal{LCR} is in fact a linear quadratic regulator, with q , r , and g quadratic, then so is $\widetilde{\mathcal{LCR}}$, thanks to (5). If \mathcal{LCR} is a linear quadratic regulator with control constraints $u \in U$ for some closed convex set U (usually with $0 \in \text{int } U$) then the dual problem is unconstrained, but the state cost r^* is not quadratic (recall (6)).

The (dual) optimal value function $\widetilde{V} : (-\infty, T] \times \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is defined, at a given (τ, π) , as the infimum in the problem $\mathcal{LCR}(\tau, \pi)$. By the symmetry of the Standing Assumptions, the properties attributed to V in Theorem 2.1 are present also for \widetilde{V} .

The idea of the computation that followed Theorem 2.1 can be now summarized. For details, see Theorem 5.1 in Ref. 6.

Theorem 2.2 (value function duality). *For each $\tau \leq T$, the value functions $V(\tau, \cdot)$ and $\widetilde{V}(\tau, \cdot)$ are convex conjugates of each other, up to a minus sign. That is,*

$$\widetilde{V}(\tau, \pi) = \sup_{\xi} \{-\pi \cdot \xi - V(\tau, \xi)\}, \quad V(\tau, \xi) = \sup_{\pi} \{-\xi \cdot \pi - \widetilde{V}(\tau, \pi)\}.$$

This result allows for studying properties of $V(\tau, \cdot)$ by analyzing corresponding properties of $\widetilde{V}(\tau, \cdot)$. For example, it is a general property of convex functions that a function is coercive if and only if its conjugate is finite-valued. Thus, if g is coercive, then g^* is finite-valued, then by Theorem 2.1 the function $\widetilde{V}(\tau, \cdot)$ is finite-valued for each τ , and so $V(\tau, \cdot)$ is coercive for each τ . Further correspondences, between differentiability and strict or strong convexity, will be explored in Section 3.

2.5. The Hamiltonian and open-loop optimality conditions

With the generalized problems of Bolza \mathcal{P} as described in Subsection 2.3, one can associate the (maximized) Hamiltonian $H : \mathbb{R}^{2n} \rightarrow \overline{\mathbb{R}}$:

$$H(x, p) = \sup_v \{p \cdot v - L(x, v)\}.$$

Convexity of L implies that $H(x, p)$ is convex in p for a fixed x and concave in x for a fixed p . In presence of the growth properties of L mentioned following (10), it is also finite-valued. (The said growth properties have equivalent characterizations in terms of H , see Theorem 2.3 in Ref. 6.) The finiteness of H in turn guarantees that

$$-H(x, p) = \sup_w \{x \cdot w - \tilde{L}(p, w)\}.$$

In other words, the Hamiltonian associated with the dual problem $\tilde{\mathcal{P}}$ is exactly $-H$. This further underlines the symmetry between \mathcal{P} and $\tilde{\mathcal{P}}$.

A Hamiltonian trajectory on $[\tau, T]$ is a pair of arcs $x, p : [\tau, T] \rightarrow \mathbb{R}^n$ such that

$$\dot{x}(t) \in \partial_p H(x(t), p(t)), \quad -\dot{p}(t) \in \hat{\partial}_x H(x(t), p(t)) \quad (15)$$

for almost all $t \in [\tau, T]$. In (15), $\partial_p H(x, p)$ is the subdifferential, in the sense of convex analysis, of the convex function $H(x, \cdot)$, while $\hat{\partial}_x H(x, p) = -\partial_x (-H(x, p))$ where the second subdifferential is in the sense convex analysis. Under suitable assumptions (which also guarantee (10)) the following are equivalent:

- $x(\cdot)$ solves \mathcal{P} , $p(\cdot)$ solves $\tilde{\mathcal{P}}$;
- $x(\cdot), p(\cdot)$ form a Hamiltonian trajectory and

$$(p(\tau), -p(T)) \in \partial l(x(\tau), x(T)).$$

In the convex setting, the key behind this equivalence is the inequality (7) and the relationship between subdifferentials of L , \tilde{L} and H , rather than any regularity properties of these functions. For details, see Ref. 2.

When the problem \mathcal{P} comes from reformulating \mathcal{LCR} in the Bolza framework as done in (11), one has

$$H(x, p) = p \cdot Ax - q(Cx) + r^*(B^*p). \quad (16)$$

The Hamiltonian system (15) turns into

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Bu(t), & u(t) &\in \partial r^*(s(t)), & s(t) &= B^*p(t), \\ \dot{p}(t) &= -A^*p(t) + C^*z(t), & z(t) &\in \partial q(y(t)), & y(t) &= Cx(t), \end{aligned} \quad (17)$$

if one additionally requires that $u(t)$ be a minimizer of $r(u)$ over all u such that $\dot{x}(t) = Ax(t) + Bu$ and that $z(t)$ be a minimizer of $q^*(z)$ over all z such that $\dot{p}(t) = -A^*p(t) + C^*z$. Details are given in Lemma 3.8 of Ref. 12; see also Theorem 4.7 in Ref. 5.

Theorem 2.3 (open-loop optimality). *For a control $u : [\tau, T] \rightarrow \mathbb{R}^k$, the following statements are equivalent:*

- $u(\cdot)$ is an optimal control for $\mathcal{LCR}(\tau, \xi)$ and $-\pi \in \partial_\xi V(\tau, \xi)$;
- there exists an integrable control $z : [\tau, T] \rightarrow \mathbb{R}^m$ so that arcs $x, p : [\tau, T] \rightarrow \mathbb{R}^n$ corresponding to controls $u(\cdot), z(\cdot)$ through the initial condition $(x(\tau), p(\tau)) = (\xi, \pi)$ and (17) satisfy

$$-p(T) \in \partial g(x(T)).$$

It can be added that if $z(\cdot)$ corresponds to an optimal $u(\cdot)$ as described in the theorem above, then $z(\cdot)$ is optimal for the problem defining $\tilde{V}(\tau, \pi)$, and also $-\xi \in \partial_\pi \tilde{V}(\tau, \pi)$.

2.6. Hamilton-Jacobi results

The Hamiltonian H and the associated Hamiltonian dynamical system (15) is involved in two different characterizations of the value function. One is the (generalized) Hamilton-Jacobi partial differential equation, as written in Theorem 2.5 of Ref. 6. It says that, for all $\tau < T$, we have $\sigma + H(x, p) = 0$ for each subgradient (σ, p) (in the sense of Definition 8.3 of Ref. 11) of $V(\tau, \xi)$. We stress that the subgradient is taken with respect to both time and space variables, and thus it can not be understood in the sense of convex analysis. Subject to a boundary condition, V is in fact the unique solution of this generalized Hamilton-Jacobi equation, as shown by Galbraith in Ref. 8.

The other characterization will involve the flow S_τ generated by the Hamiltonian dynamical system (15):

$$S_\tau(\xi_T, \pi_T) = \left\{ (\xi, \pi) \mid \begin{array}{l} \text{there exists a Hamiltonian traj. } (x(\cdot), p(\cdot)) \text{ on } [\tau, T] \\ \text{with } (x(\tau), p(\tau)) = (\xi, \pi), (x(T), p(T)) = (\xi_T, \pi_T) \end{array} \right\}.$$

We have the following result, as stated in Theorem 2.4 of Ref. 6:

Theorem 2.4 (Hamiltonian flow of $\partial_\xi V(\tau, \cdot)$). *For any $\tau \leq T$,*

$$\text{gph}(-\partial_\xi V(\tau, \cdot)) = S_\tau(\text{gph}(-\partial g)).$$

In other words, the graph of the subdifferential of $V(\tau, \cdot)$ is essentially obtained from the graph of the subdifferential of g by flowing backwards

along Hamiltonian trajectories. This result can also be thought of as saying that the (generalized) method of characteristics works globally in this convex setting. Thanks to Theorem 2.2 and the fact that subdifferential mappings of a pair of convex conjugate functions are inverses of one another, Theorem 2.4 also gives a description of $\partial_\pi \tilde{V}(\tau, \cdot)$.

Since the subdifferential of a convex function is a monotone mapping, the graph of $-\partial g$ is an “anti-monotone set”, that is, for any $(x_1, p_1), (x_2, p_2) \in \text{gph}(-\partial g)$, we have $(x_1 - x_2) \cdot (p_1 - p_2) \leq 0$. Thanks to the concavity in x , convexity in p , of $H(x, p)$, the flow S_τ has a certain monotonicity, or rather anti-monotonicity preserving property:

$$(x_1(\tau) - x_2(\tau)) \cdot (p_1(\tau) - p_2(\tau)) \leq (x_1(T) - x_2(T)) \cdot (p_1(T) - p_2(T)) \quad (18)$$

for any pair of Hamiltonian trajectories $(x_i(\cdot), p_i(\cdot))$, $i = 1, 2$, on $[\tau, T]$. This was noted in Theorem 4 of Ref. 13. In what follows, we will need this result specified to the case of \mathcal{LCR} .

Lemma 2.1 (preservation of monotonicity). *Let $(x_i(\cdot), p_i(\cdot))$, $i = 1, 2$, be Hamiltonian trajectories on $[\tau, T]$ for the Hamiltonian (16). Then*

$$\begin{aligned} & \frac{d}{dt}(x_1(t) - x_2(t)) \cdot (p_1(t) - p_2(t)) \\ & \in (\partial q(Cx_1(t)) - \partial q(Cx_2(t))) \cdot (Cx_1(t) - Cx_2(t)) \\ & \quad + (\partial r^*(B^*p_1(t)) - \partial r^*(B^*p_2(t))) \cdot (B^*p_1(t) - B^*p_2(t)) \geq 0 \end{aligned}$$

The inequality above can be easily shown directly from the definition of a Hamiltonian trajectory. It also obviously implies (18).

2.7. Feedback optimality conditions

Open-loop optimal solutions of optimal control problems are not well-suited to cope with uncertainty or disturbance. This, and the classical control engineering problem of feedback stabilization, motivate the pursuit of an optimal feedback: a mapping that at each state produces the set of optimal controls to be applied.

Theorem 2.3, Theorem 2.4, and the Hamiltonian system (15) suggest that optimal arcs $x(\cdot)$ satisfy the inclusion $\dot{x}(t) \in \partial_y H(x(t), -\partial_\xi V(t, x(t)))$. Indeed, if a Hamiltonian trajectory $(x(\cdot), p(\cdot))$ on $[\tau, T]$ is such that $p(T) \in -\partial g(x(T))$ then, by Theorem 2.4, $p(t) \in -\partial_\xi V(t, x(t))$ for almost all $t \in [\tau, T]$. A more difficult – and more important – issue is to what extent the stated inclusion is sufficient for optimality. To an extent, this was addressed in Ref. 12. We state the relevant result in Theorem 2.5.

The *optimal feedback* mapping $\Phi : (-\infty, T] \times \mathbb{R}^n$ for the problem \mathcal{LCR} is defined, at each $t \leq T$, $x \in \mathbb{R}^n$, by

$$\Phi(t, x) = \partial r^*(-B^* \partial_\xi V(t, x)) = \bigcup_{p \in \partial_\xi V(t, x)} \partial r^*(-B^* p).$$

Theorem 2.5 (optimal feedback). *Fix $(\tau, \xi) \in (-\infty, T] \times \mathbb{R}^n$. If $u : [\tau, T] \rightarrow \mathbb{R}^k$ is an optimal control for $\mathcal{LCR}(\tau, \xi)$ and $\partial_\xi V(\tau, \xi) \neq \emptyset$, then*

$$u(t) \in \Phi(t, x(t)) \text{ for almost all } t \in [\tau, T]. \quad (19)$$

On the other hand, if an integrable control $u : [\tau, T] \rightarrow \mathbb{R}^k$ satisfies (19), and the absolutely continuous $x : [\tau, T] \rightarrow \mathbb{R}^n$ satisfying (2) and $x(\tau) = \xi$ is such that $x(t) \in \text{int dom } V(\tau, \cdot)$ for almost all $t \in [\tau, T]$, then $u(\cdot)$ is optimal for $\mathcal{LCR}(\tau, \xi)$.

The interior $\text{int dom } V$ of the effective domain of V , i.e., of the set $\text{dom } V$ of all points where V is finite, is nonempty if $\text{int dom } g$ is nonempty, and $(\tau, \xi) \in \text{int dom } V$ if and only if $\tau < T$ and $\xi \in \text{int dom } V(\tau, \cdot)$; see Proposition 7.2 and Corollary 7.5 in Ref. 6. In general, the feedback mapping Φ is outer semicontinuous and locally bounded on $\text{int dom } V$, but it need not have convex values; see Example 3.7 in Ref. 12.

3. Regularity of the value function

We now give some results on how regularity of the data of \mathcal{LCR} may be reflected in the properties of V , and thus, of the optimal feedback mapping Φ . The central role played by the Hamiltonian dynamical system (15) in the analysis of convex optimal control problems is reflected below in the fact that the only tools from the previous chapter we use are Theorem 2.4, which describes how $\nabla_\xi V(\tau, \cdot)$ evolves in the Hamiltonian system, and Lemma 2.1, about the monotonicity-preserving properties of the Hamiltonian flow.

We note that similar regularity properties of the value function in a nonconvex setting require stronger assumptions than the ones posed here; see for example Refs. 14,15. To an extent though, regularity properties (or singularities) of the value function are propagated via the Hamiltonian system; see Ref. 16.

3.1. Convex-valued and single-valued optimal feedback

A natural condition to guarantee that the optimal feedback Φ has convex values is that $\partial_\xi V(\tau, \xi)$ be single-valued, which is equivalent to essential

strict differentiability of $V(\tau, \cdot)$. Theorem 3.1 will show that this property is in a sense inherited by V from g .

We need some preliminary material. Let $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ be a proper, lsc., and convex function. The function f is *essentially differentiable* if $D = \text{int dom } f$ is nonempty, f is differentiable on D , and $\lim_{i \rightarrow \infty} |\nabla f(x_i)| = \infty$ for any sequence $\{x_i\}_{i=1}^{\infty}$ converging to a boundary point of D . The function f is *essentially strictly convex* if f is convex on every convex subset of $\text{dom } \partial f = \{x \in \mathbb{R}^n \mid \partial f(x) \neq \emptyset\}$; that is, if for each convex $S \subset \text{dom } \partial f$, we have $(1 - \lambda)f(x_1) + \lambda f(x_2) > f((1 - \lambda)x_1 + \lambda x_2)$ for all $\lambda \in (0, 1)$, all $x_1, x_2 \in S$. The subdifferential mapping is a *monotone mapping*: for all $x_1, x_2 \in \text{dom } \partial f$, all $y_1 \in \partial f(x_1)$, $y_2 \in \partial f(x_2)$, $(x_1 - x_2) \cdot (y_1 - y_2) \geq 0$. It is *strictly monotone* if $(x_1 - x_2) \cdot (y_1 - y_2) > 0$ as long as $x_1 \neq x_2$.

The following statements are equivalent: (i) f is essentially differentiable; (ii) ∂f is single-valued on $\text{dom } \partial f$; (iii) ∂f^* is strictly monotone; and (iv) f^* is essentially strictly convex.

The result below is a special case of Theorem 4.4 in Ref. 12. Note that no regularity is assumed about r .

Theorem 3.1 (value function – essential differentiability).

Suppose that q and g are essentially differentiable. Then, for all $\tau \leq T$, $V(\tau, \cdot)$ is essentially differentiable.

Proof. By Theorem 2.4, any two points in $\text{gph } \partial_\xi V(\tau, \cdot)$ can be expressed as $(x_i(\tau), p_i(\tau))$, $i = 1, 2$, where $(x_i(\cdot), p_i(\cdot))$ are Hamiltonian trajectories on $[\tau, T]$ with $p_i(T) = -\nabla g(x_i(T))$, $i = 1, 2$. If $x_1(\tau) = x_2(\tau)$, then by Lemma 2.1 one obtains

$$0 = (x_1(\tau) - x_2(\tau)) \cdot (p_1(\tau) - p_2(\tau)) \leq (x_1(T) - x_2(T)) \cdot (p_1(T) - p_2(T)) \leq 0,$$

with the last inequality above coming from monotonicity of $\nabla g(\cdot)$. Hence

$$(x_1(T) - x_2(T)) \cdot (p_1(T) - p_2(T)) = 0$$

which by essential strict convexity of g^* (a property dual to essential differentiability of g) implies that $p_1(T) = p_2(T)$. The fact that $(x_1(t) - x_2(t)) \cdot (p_1(t) - p_2(t))$ is constant on $[\tau, T]$ and Theorem 4 of Ref. 13 implies that $(x_1(\cdot), p_2(\cdot))$ is a Hamiltonian trajectory. In particular,

$$\begin{aligned} \dot{p}_1(t) - \dot{p}_2(t) &= -A^*(p_1(t) - p_2(t)) + C^*(\nabla q(Cx_1(t)) - \nabla q(Cx_2(t))) \\ &= -A^*(p_1(t) - p_2(t)), \end{aligned}$$

and since $p_1(T) = p_2(T)$, it must be the case that $p_1(\tau) = p_2(\tau)$.

We have shown that, for any two points $(x_i(\tau), p_i(\tau))$, $i = 1, 2$, in $\text{gph } \partial_\xi V(\tau, \cdot)$, if $x_1(\tau) = x_2(\tau)$ then also $p_1(\tau) = p_2(\tau)$. This amounts to single-valuedness of $\partial_\xi V(\tau, \cdot)$ which in turn is equivalent to essential differentiability of $V(\tau, \cdot)$. \square

Under the assumptions of Theorem 3.1, the feedback Φ has convex values, since ∂r^* has convex values. If furthermore ∂r^* is single-valued, and so if r^* is differentiable (as r^* is finite, its essential differentiability is just differentiability), then Φ is single-valued and in fact continuous. Thus, if q and g are essentially differentiable and r is strictly convex, then the optimal feedback Φ is continuous on $\text{int dom } V$. (C.f. Proposition 3.5 in Ref. 12.)

A different set of assumptions leading to essential differentiability of $V(\tau, \cdot)$ is strict convexity of both q and r^* and observability of (A, C) and $(-A^*, B^*)$. This combines arguments that apply to the case of strictly concave, strictly convex Hamiltonian, as in Ref. 17 and basic properties of observable linear systems (c.f. Theorem 3.1 in Ref. 18). We note though that strict convexity of r^* may be absent in simple cases, like when r represents a quadratic cost and a convex control constraint set (recall (6)).

The next natural property of the optimal feedback to pursue, with the uniqueness of solutions to the closed loop equation

$$\dot{x}(t) = Ax(t) + B\Phi(t, x(t)) \quad (20)$$

in mind, is the local Lipschitz continuity of $\Phi(t, \cdot)$. This motivates the developments in the next section.

3.2. Convex functions with locally Lipschitz gradients

Let $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ be a proper, lsc., and convex function. The property of f being differentiable with ∇f Lipschitz continuous on \mathbb{R}^n with constant $\rho > 0$ is equivalent to f^* being *strongly convex* with constant $\sigma = \rho^{-1}$:

$$(1-\lambda)f^*(p_1) + \lambda f^*(p_2) \geq f^*((1-\lambda)p_1 + \lambda p_2) + \frac{1}{2}\sigma\lambda(1-\lambda)|p_1 - p_2|^2 \quad (21)$$

for all $\lambda \in (0, 1)$, all $p_1, p_2 \in \mathbb{R}^n$. A manipulation of (21) shows that strong convexity of f^* with constant σ is equivalent to $f^* - \frac{1}{2}\sigma|\cdot|^2$ being convex. Furthermore, strong convexity of f^* is equivalent to its subdifferential being *strongly monotone* with constant σ : $(p_1 - p_2) \cdot (x_1 - x_2) \geq \sigma|p_1 - p_2|^2$ for all $(p_1, x_1), (p_2, x_2) \in \text{gph } \partial f^*$.

We now report some recent results from Ref. 19 regarding localization of the properties just mentioned. In what follows, we will say that a proper,

lsc., and convex function $g : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is *essentially locally strongly convex* if for any compact and convex $K \subset \text{dom } \partial g$, g is strongly convex on K , in the sense that there exists $\sigma > 0$ such that (21) is satisfied for all $p_1, p_2 \in K$ and all $\lambda \in (0, 1)$.

Theorem 3.2 (strong convexity and gradient continuity). *For a proper, lsc., and convex function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$, the following conditions are equivalent:*

- (a) f is essentially differentiable and ∇f is locally Lipschitz continuous on $\text{int dom } f$;
- (b) the mapping ∂f is single-valued on $\text{dom } \partial f$ and locally Lipschitz continuous relative to $\text{gph } \partial f$ in the sense that for each compact subset $S \subset \text{gph } \partial f$ there exists $\rho > 0$ such that for any $(x_1, y_1), (x_2, y_2) \in S$ one has

$$|y_1 - y_2| \leq \rho |x_1 - x_2|;$$

- (c) the mapping ∂f^* is locally strongly monotone relative to $\text{gph } \partial f^*$ in the sense that for each compact subset $S \subset \text{gph } \partial f^*$ there exists $\sigma > 0$ such that for any $(y_1, x_1), (y_2, x_2) \in S$ one has

$$(y_1 - y_2) \cdot (x_1 - x_2) \geq \sigma |y_1 - y_2|^2.$$

If f is such that $\text{dom } \partial f^*$ is convex and open, then each of the conditions above is equivalent to:

- (d) f^* is essentially locally strongly convex.

This in particular says that for a coercive (and proper and lsc) convex function f , essential differentiability of f and local Lipschitz continuity of ∇f is equivalent to (naturally understood) local strong convexity of f^* .

For convenience, we define the following classes of convex functions:

$$\mathcal{C}^+ = \left\{ f \in \mathcal{C} \mid \begin{array}{l} f \text{ is essentially differentiable and} \\ \nabla f \text{ is locally Lipschitz continuous on } \text{int dom } f \end{array} \right\},$$

$$\mathcal{C}^{++} = \{ f \in \mathcal{C}^+ \mid f \text{ is essentially locally strongly convex} \}.$$

Since for any essentially differentiable f , $\text{dom } \partial f = \text{int dom } f$ is open and convex, for such f the condition that f is essentially locally strongly convex is equivalent, by Theorem 3.2, to $f^* \in \mathcal{C}^+$. Thus, $f \in \mathcal{C}^{++}$ if and only if $f, f^* \in \mathcal{C}^+$. This immediately implies that $f \in \mathcal{C}^{++}$ if and only if $f^* \in \mathcal{C}^{++}$.

3.3. Locally Lipschitz continuous optimal feedback

The result below is a local version of Theorem 3.6 of Ref. 20, where global Lipschitz continuity of ∇H and ∇g was assumed. Here though, we make stronger symmetric assumptions on g . This simplifies the proof.

Theorem 3.3 (value function – strong local regularity).

Suppose that $q, r^* \in C^+$ while $g \in C^{++}$. Then, for all $\tau \leq T$, $V(\tau, \cdot)$ is an element of C^{++} .

Proof. By Theorem 2.4, any point in $\text{gph } \partial_\xi V(\tau, \cdot)$ can be expressed as $(x(\tau), p(\tau))$ where $(x(\cdot), p(\cdot))$ is a Hamiltonian trajectories on $[\tau, T]$ with $p(T) = -\nabla g(x(T))$. By local Lipschitz continuity of $\nabla H(\cdot, \cdot)$, one can find a compact neighborhood $X_\tau \times P_\tau$ of $(x(\tau), p(\tau))$ and a compact neighborhood $X_T \times P_T$ of $(x(T), p(T))$ such that:

- (i) $X_\tau \times P_\tau \subset S_\tau(X_T \times P_T)$;
- (ii) no Hamiltonian trajectory from $X_\tau \times P_\tau$ blows up in time less or equal to $T - \tau$ (and so in particular, there exists a compact set $K \subset \mathbb{R}^{2n}$ so that all Hamiltonian trajectories on $[\tau, T]$ from $X_\tau \times P_\tau$ are in K);
- (iii) ∇H is Lipschitz continuous with constant M_1 on K ;
- (iv) $X_T \subset \text{int dom } g$, $P_T \subset \text{int dom } g^*$;
- (v) $\nabla g(\cdot)$ is Lipschitz continuous with constant M_2 and strongly monotone with constant $1/M_2$ relative to $X_T \times P_T$.

Now, pick any two different points $(x_i(\tau), p_i(\tau))$, $i = 1, 2$, in $\text{gph } \partial_\xi V(\tau, \cdot)$ with $x_i(\tau) \in X_\tau$, $p_i(\tau) \in P_\tau$, and let $(x_i(\cdot), p_i(\cdot))$, $i = 1, 2$, to be the associated Hamiltonian trajectories on $[\tau, T]$ with $p_i(T) = -\nabla g(x_i(T))$. Note that by Theorem 3.1 $x_1(\tau) \neq x_2(\tau)$ and $p_1(\tau) \neq p_2(\tau)$. To shorten the notation, we will write $\alpha(t) = x_1(t) - x_2(t)$, $\beta(t) = p_1(t) - p_2(t)$. We have

$$-\|\alpha(\tau)\|\|\beta(\tau)\| \leq \alpha(\tau) \cdot \beta(\tau) \leq \alpha(T) \cdot \beta(T) \leq -\frac{1}{M_2} \|\beta(T)\|^2$$

where the second inequality comes from Lemma 2.1, and the last inequality comes the strong monotonicity of ∇g (as in (v) above). In particular,

$$\frac{1}{\|\alpha(\tau)\|\|\beta(\tau)\|} \leq \frac{M_2}{\|\beta(T)\|^2}.$$

Lipschitz continuity of ∇H yields ^a

$$\|\beta(\tau)\| \leq \|\beta(T)\| + \left(e^{M_1(T-\tau)} - 1\right) \|(\alpha(T), \beta(T))\|.$$

Squaring both sides and multiplying by the previous displayed inequality yields

$$\begin{aligned} \frac{\|\beta(\tau)\|}{\|\alpha(\tau)\|} &\leq M_2 \left(1 + \left(e^{M_1(T-\tau)} - 1\right) \sqrt{1 + \frac{\|\alpha(T)\|^2}{\|\beta(T)\|^2}}\right)^2 \\ &\leq M_2 \left(1 + \left(e^{M_1(T-\tau)} - 1\right) \sqrt{1 + M_2^2}\right)^2. \end{aligned}$$

This shows that $\nabla_\xi V(\tau, \cdot)$ is Lipschitz continuous relative to $X_\tau \times P_\tau$. Strong monotonicity of $\nabla_\xi V(\tau, \cdot)$ can be shown similarly, by using a bound on $\|\alpha(\tau)\|$ in place of the bound on $\|\beta(\tau)\|$ in the arguments above. \square

Minor modifications of the proof above do show that the Lipschitz constant for $\nabla_\xi V(\tau, \cdot)$ is locally bounded in both τ and ξ . This, and local Lipschitz continuity of ∇r^* , implies that the solutions to (20) are unique.

4. Infinite horizon problems

The original control engineering motivation for considering convex optimal control problem – construction of stabilizing feedback – calls for posing problems on an infinite time horizon. (Finite time horizon problems can then be used to approximate the infinite horizon ones, or, as seen in Receding Horizon control (also referred to as Model Predictive Control), to obtain stabilizing feedbacks without necessarily having them optimal; see Ref. 21 for a survey.) Besides control engineering, infinite-horizon problems have seen treatment in theoretical economics; see the book by Carlson, Haurie, and Leizarowitz, Ref. 22, for an overview, Ref. 23 for an example of duality

^aThis bound comes from reversing time, i.e., considering the backward Hamiltonian flow, with time 0 corresponding to T and $T - \tau$ corresponding to τ . By Lipschitz continuity of ∇H with constant M_1 one gets

$$\left\| \frac{d}{dt}(\alpha(t), \beta(t)) \right\| \leq M_1 \|(\alpha(t), \beta(t))\|$$

and thus $\|(\alpha(t), \beta(t))\| \leq \|(\alpha(0), \beta(0))\| e^{M_1 t}$. One also gets

$$\|\dot{\beta}(t)\| \leq M_1 \|(\alpha(t), \beta(t))\| \leq M_1 \|(\alpha(0), \beta(0))\| e^{M_1 t}$$

which then yields $\|\beta(t)\| \leq \|\beta(0)\| + \|(\alpha(0), \beta(0))\| (e^{M_1 t} - 1)$.

techniques, and Ref. 17 for an illustration of the role played by the Hamiltonian dynamical system. For an infinite-horizon \mathcal{LCR} , some convexity tools were used in Ref. 24. A detailed, and largely self-contained analysis of a Linear-Quadratic Regulator with control constraints was carried out in Ref. 25. General calculus of variations problems on the infinite-time horizon were treated in Ref. 26, based to an extent on the duality and Hamilton-Jacobi developments of Ref. 6.

The (primal) infinite horizon optimal value function $W : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is defined as the infimum of

$$\int_0^\infty q(y(t)) + r(u(t)) dt,$$

taken over all locally integrable controls $u : [0, \infty) \rightarrow \mathbb{R}^k$, subject to dynamics (2) and the initial condition

$$x(0) = \xi.$$

The (dual) infinite horizon optimal value function $\widetilde{W} : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is defined as the infimum of

$$\int_0^\infty r^*(s(t)) + q^*(z(t)) dt,$$

taken over all locally integrable controls $z : [0, \infty) \rightarrow \mathbb{R}^m$, subject to dynamics (13) and the initial condition

$$p(0) = \pi.$$

Addition to the Standing Assumptions. *The functions q, q^*, r, r^* are positive definite (0 at 0, positive elsewhere). The pairs of matrices (A, C) and $(-A^T, B^T)$ are observable.*

Of course, the value functions W and \widetilde{W} are convex. Results regarding their lower semicontinuity and the existence of optimal arcs in the problems defining them are in Ref. 26. The functions W and \widetilde{W} are both positive definite, and the optimal arcs, for both W and \widetilde{W} , tend to 0 as $t \rightarrow \infty$. Indeed, for W , note that $W(\xi) = 0$ only if the optimal arc x satisfies $Cx(t) = 0$ while the optimal control is $u(t) = 0$ for all $t \geq 0$. But then $\dot{x}(t) = Ax(t)$, and observability of (A, C) gives $x(t) = 0$ for all $t \geq 0$, in particular $\xi = x(0) = 0$. Given an optimal arc $x(\cdot)$ with $W(x(0)) < \infty$, we have $W(x(T)) = \int_0^T q(Cx(t)) + r(u(t)) dt$ and by finiteness of $W(\xi)$, the integral tends to 0 as $T \rightarrow \infty$. Positive definiteness of $W(x(T))$ now implies that $x(T) \rightarrow 0$. Arguments for \widetilde{W} are symmetric.

Now, [26, Corollary 3.5] says that:

- $W(\xi) = \widetilde{W}^*(-\xi)$, equivalently, $\widetilde{W}(\pi) = W^*(-\pi)$;
- $f = W$ is the unique $f \in \mathcal{C}$ such that $H(x, -\partial f(x)) = 0$ for all $x \in \mathbb{R}^n$;
- $f = \widetilde{W}$ is the unique $f \in \mathcal{C}$ such that $H(-\partial f(p), p) = 0$ for all $p \in \mathbb{R}^n$.

In other words, the infinite horizon value functions are conjugate to each other, and are the unique solutions to the generalized stationary Hamilton-Jacobi equations. (The equation $H(x, -\partial f(x)) = 0$ is to be understood as $H(x, -p) = 0$ for all $p \in \partial f(x)$.) It is interesting to note that with such a concept of a solution, it is no longer true that under standard assumptions of (A, B) being stabilizable and (A, C) being detectable, and with q and r quadratic, the (quadratic) value function W is the unique solution to the Hamilton-Jacobi equation. (This occurs even though the matrix defining W is the unique solution to the matrix Riccati equation, to which the Hamilton-Jacobi equation simplifies if one only looks for quadratic solutions.) See Section 4.1 in Ref. 26 for details.

Another consequence of the fact that optimal arcs for both W and \widetilde{W} converge to 0 is a result relating these two value functions to the finite horizon value functions discussed earlier. It turns out that both W and \widetilde{W} can be simultaneously approximated by a pair of conjugate finite horizon value functions $V(\tau, \cdot)$ and $\widetilde{V}(\tau, \cdot)$. Indeed, as the proof of Theorem 3.4 in Ref. 26 suggests, we have the following:

Theorem 4.1 (finite-horizon approximation). *Consider the functions $W^T, \widetilde{W}^T : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ defined at points $\xi, \pi \in \mathbb{R}^n$ as the optimal values in $\mathcal{LCR}(0, \xi)$, $\widetilde{\mathcal{LCR}}(0, \pi)$. Suppose that g and g^* are finite-valued on some neighborhood of 0. Then W^T converges epi-graphically to W while \widetilde{W}^T converges epi-graphically to \widetilde{W} as $T \rightarrow \infty$.*

For the infinite horizon case, the Hamiltonian dynamical system (15) again yields a description of the subdifferential of W (and of \widetilde{W}). Indeed, $\text{gph } -\partial W$ consists of all points (ξ, π) for which there exists a Hamiltonian trajectory x, p on $[0, \infty)$ with $x(0) = \xi$, $p(0) = \pi$, and $x(t) \rightarrow 0$, $p(t) \rightarrow 0$ as $t \rightarrow \infty$. See Proposition 3.7 in Ref. 26 which generalizes the results given for strictly concave / strictly convex Hamiltonians in Ref. 17. We add that for such a Hamiltonian trajectory, x is optimal for $W(\xi)$ while p is optimal for $\widetilde{W}(\pi)$.

The said description of ∂W leads to the following result.

Proposition 4.1. *Suppose that there exists a neighborhood of 0 on which*

both q and r^* are strictly convex. Then W is both essentially differentiable and essentially strictly convex.

The assumption is equivalent to strict convexity of q and essential differentiability of r . By duality, the same assumptions imply that \widetilde{W} is also essentially differentiable and essentially strictly convex. The result can be shown via arguments combining Lemma 2.1 and controllability assumptions we have. (Similar results, shown via slightly different methods, are in Theorem 4.1 of Ref. 17 and Theorem 3.1 of Ref. 18.)

An interesting application of Proposition 4.1 lies in showing that for a wide class of linear systems with saturation nonlinearities there does exist a continuous and stabilizing feedback. (Such an issue was partially addressed in Refs. 27,28 via tools not related to optimality.) Consider a system

$$\dot{x}(t) = Ax(t) + B\sigma(u(t)) \quad (22)$$

where $\sigma : \mathbb{R}^k \rightarrow \mathbb{R}^k$ is a saturation nonlinearity. Often, for single input systems, $\sigma(u) = \arctan(u)$ or $\sigma(u) = u$ for $u \in [-1, 1]$, $\sigma(u) = -1$ for $u < -1$, $\sigma(u) = 1$ for $u > 1$. For these, and for many other commonly encountered in the control literature saturation functions, $\sigma = \nabla s$ for a convex function $s \in \mathcal{C}$ which is strictly convex near 0. Taking $r = s^*$ and, for simplicity, any positive definite quadratic q , leads to an essentially differentiable W . The optimal feedback $\Phi(x) = \nabla s(-B^*\nabla W(x))$ leads to asymptotic stability for the closed loop system $\dot{x}(t) = Ax(t) + B\Phi(x(t))$. This asymptotic stability can in fact be verified with W as a Lyapunov function. But since $\nabla s = \sigma$, this means that choosing $u(x) = -B^*\nabla W(x)$ in (22) also leads to asymptotic stability. As ∇W is continuous, there does exist a continuous stabilizing feedback for the nonlinear system (22). For details, see Ref. 18.

We conclude by stating a result on local Lipschitz continuity of ∇W . Since optimal trajectories for problems defining both W and \widetilde{W} converge to 0, local properties around 0 of these functions to an extent determine global properties. Another way to say this is that, on each compact subset of $\text{dom } W$, W equals to a finite horizon value function defined by \mathcal{LCR} with $g = W$ in a neighborhood of 0, as long as T is large enough. This and Theorem 3.3 lead to the following.

Proposition 4.2. *Suppose that there exists a neighborhood of 0 on which W is differentiable and strongly convex and ∇W is Lipschitz continuous. Then $W \in \mathcal{C}^{++}$.*

An important example of problems where W is differentiable and

strongly convex around 0 is provided by a \mathcal{LQR} problem with control constraints: u is restricted to some closed convex set U and $0 \in \text{int } U$. Indeed, then on some neighborhood of 0, W is quadratic and equals to the value function of the unconstrained \mathcal{LQR} problem.

Acknowledgments

Research by R.T. Rockafellar was partially supported by the NSF Grant DMS-0104055.

References

1. B. Anderson and J. Moore, *Optimal Control – Linear Quadratic Methods* (Prentice-Hall, 1990).
2. R. Rockafellar, *J. Math. Anal. Appl.* **32**, 174 (1970).
3. R. Rockafellar, *Trans. Amer. Math. Soc.* **159**, 1 (1971).
4. R. Rockafellar, *SIAM J. Control Optim.* **25**, 781 (1987).
5. R. Rockafellar, *SIAM J. Control Optim.* **27**, 1007 (1989).
6. R. Rockafellar and P. Wolenski, *SIAM J. Control Optim.* **39**, 1323 (2000).
7. R. Rockafellar and P. Wolenski, *SIAM J. Control Optim.* **39**, 1351 (2000).
8. G. Galbraith, *SIAM J. Control Optim.* **39**, 281 (2000).
9. R. Rockafellar, *Convex Analysis* (Princeton University Press, 1970).
10. F. Clarke, *Optimization and Nonsmooth Analysis* (Wiley, 1983).
11. R. Rockafellar and R. J.-B. Wets, *Variational Analysis* (Springer, 1998).
12. R. Goebel, *Set-Valued Analysis* **12**, 127 (2004).
13. R. Rockafellar, *Pacific J. Math.* **33**, 411 (1970).
14. C. Byrnes and H. Frankowska, *C.R. Acad. Sci. Paris* **315**, 427 (1992).
15. N. Caroff and H. Frankowska, *Trans. Amer. Math. Soc.* **348**, 3133 (1996).
16. H. Frankowska and A. Ochal, *J. Math. Anal. Appl.* **306**, 714 (2005).
17. R. Rockafellar, *J. Optim. Theory Appl.* **12** (1973).
18. R. Goebel, *IEEE Trans. Automat. Contr.* **50**, 650 (2005).
19. R. Goebel and R. Rockafellar, Local strong convexity and local Lipschitz continuity of the gradient of convex functions, in preparation.
20. R. Goebel, *SIAM J. Control Optim.* **43**, 1781 (2005).
21. D. Mayne, J. Rawlings, C. Rao and P. Scokaert, *Automatica* **36**, 789 (2000).
22. D. Carlson, A. Haurie and A. Leizarowitz, *Infinite Horizon Optimal Control: Deterministic and Stochastic Systems* (Springer-Verlag, 1991).
23. L. Benveniste and J. Scheinkman, *J. Economic Theory* **27**, 1 (1982).
24. G. D. Blasio, *SIAM J. Control Opt.* **29**, 909 (1991).
25. R. Goebel and M. Subbotin, *IEEE Trans. Automat. Contr.* (2007), to appear.
26. R. Goebel, *Trans. Amer. Math. Soc.* **357**, 2187 (2005).
27. E. Sontag and H. Sussmann, Nonlinear output feedback design for linear systems with saturating controls, in *Proc. 29 IEEE Conf. Decision and Control, Honolulu, 1990*, 1990.
28. E. Sontag, H. Sussmann and Y. Yang, *IEEE Trans. Automat. Control* **39**, 2411 (1994).