Hamiltonian Dynamical Systems for Convex Problems of Optimal Control: Implications for the Value Function

Rafal Goebel[1]

Abstract

For fully convex problems of optimal control, the Hamiltonian dynamical system provides a global description of evolution of the subdifferentials of the value function from those of the initial cost. We employ this description and the convex structure of the problem to investigate the differentiability properties of the value function. Motivation is provided by questions of regularity of optimal feedback, the key ingredients of which the the value function, and by the fact that the Hamiltonian may lead to a reasonable dynamical system, even if the underlying control problem involves various constraints and penalties.

1 Introduction

The Hamiltonian function and the associated Hamiltonian dynamical system are important objects in the analysis of optimal control problems. The dynamical system is involved in optimality conditions, and the Hamiltonian itself characterizes the value function – a key ingredient in the optimal feedback – through the Hamilton-Jacobi equation.

In this note we discuss how, for fully convex problems, properties of the Hamiltonian dynamical system influence the regularity of the value function. Results in this direction are taken mainly from Goebel [8], [9]. A key tool is provided by a global description of the evolution of subdifferentials of the value function in the Hamiltonian dynamical system, given by Rockafellar and Wolenski [14]. In other words, this description globally validates – for the fully convex setting – a generalized method of characteristics for the appropriately understood Hamilton-Jacobi equation. We describe how the Hamiltonian dynamical system "stores" all the information about the initial cost of the problem (Theorem 2.3), show when "shocks" introducing nonsmoothness of the value function cannot appear (Theorem 3.1), and state how Hamiltonians with Lipschitz continuous gradients lead to Lipschitz continuity of the optimal feedback (Theorem 3.3). Through examples we illustrate that problems with abundant nonsmoothness in the cost can still lead to reasonable Hamiltonians, satisfying in particular the assumptions of mentioned results.

Problems of our interest are the fully convex generalized problems of Bolza: given any $\tau \in [0,+\infty)$ and any $\xi \in I\!R^n$ we consider the problem

$$\mathcal{P}(\tau,\xi): \quad \begin{array}{c} \text{minimize} \;\; g(x(0)) + \displaystyle\int_0^\tau L(x(t),\dot{x}(t))\,dt \\ \text{subject to} \;\; x(\tau)=\xi, \end{array} \tag{1}$$

where both the *initial cost* $g : I\!R^n \mapsto \overline{I\!R}$, and the *Lagrangian* $L : I\!R^n \times I\!R^n \mapsto \overline{I\!R}$, are convex functions. Minimization is carried out over all absolutely continuous arcs $x : [0,\tau] \mapsto I\!R^n$ satisfying the terminal condition $x(\tau) = \xi$. The optimal value in $P(\tau,\xi)$, parameterized by (τ,ξ) defines the *value function* $V : [0,+\infty) \times I\!R^n \mapsto \overline{I\!R}$ associated with the Bolza problem. Convexity assumptions on L and g guarantee that $V(\tau,\cdot)$ is convex, for any fixed τ (and in fact proper and lower semicontinuous, under the assumptions we state below).

Mild requirements on the cost functions L and g give the extended problems of Bolza a wide range of modeling possibilities. Control problems with explicit dynamics, control constraints, and various penalty functions can be expressed in the current format, as well as problems with a terminal cost and an initial condition (for which the value function is the "cost to go"). For further details on modeling capabilities see the Introduction in Clarke [5], or Rockafellar [11]. We note though that insisting on convexity essentially limits us to control problems with linear dynamics. Exact assumptions, in force in what follows, are:

(A1) *The functions $g(\cdot)$ and $L(\cdot,\cdot)$ are convex, proper and lsc on, respectively, $I\!R^n$ and $I\!R^n \times I\!R^n$.*

(A2) *The set $F(x) = \{v \mid L(x,v) < \infty\}$ is nonempty for all x, and there is a constant ρ such that* $\operatorname{dist}(0,F(x)) \le \rho(1+|x|)$ *for all x.*

(A3) *There are constants α and β and a coercive, proper, nondecreasing function $\theta(\cdot)$ on $[0,\infty)$ such that $L(x,v) \ge \theta(\max\{0,|v|-\alpha|x|\}) - \beta|x|$ for all x and v.*

[1]Center for Control Engineering & Computation, Department of Electrical & Computer Engineering, University of California, Santa Barbara, CA 93106; *rafal@ece.ucsb.edu*.

The assumptions can be equivalently stated in terms of the *Hamiltonian* $H : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ associated with the given control problem:

$$H(x,y) = \sup_{v\in\mathbb{R}^n} \{\langle y,v\rangle - L(x,v)\}\,. \tag{2}$$

Lagrangian L satisfies *(A1), (A2), (A3)* whenever the Hamiltonian H is finite, concave in x for any fixed y, convex in y for any fixed x, and has some mild growth properties, always in place for example when H – and equivalently, L – is piecewise linear-quadratic, or when the Hamiltonian has the form $H(x,y) = \langle Ax,y\rangle - h_1(x) + h_2(y)$ for some matrix A and convex functions h_1, h_2.

Our interest in the Hamiltonian and the value function comes in part from their relation to the optimal feedback – a map describing optimal controls – optimal velocities in our setting – at any given state of the system. We define the possibly set-valued $\Phi : \operatorname{int}\operatorname{dom}\partial_\xi V \rightrightarrows \mathbb{R}^n$ by

$$\begin{aligned}\Phi(t,x) &= \partial_y H(x, \partial_\xi V(t,x))\\ &= \{\partial_y H(x,\eta)\,|\,\eta\in\partial_\xi V(t,x)\}\,.\end{aligned} \tag{3}$$

Above, $\partial_y H(x,\eta)$ and $\partial_\xi V(t,x)$ denote, respectively, the subdifferentials of convex functions $H(x,\cdot)$ and $V(\tau,\cdot)$. Referring to Φ as the *optimal feedback mapping* is justified by the following result:

Theorem 1.1 (Goebel [8])**.** *Fix any $(\tau,\xi) \in [0,+\infty)\times\mathbb{R}^n$. If $x(\cdot)$ is an optimal arc for $\mathcal{P}(\tau,\xi)$, and $\partial_\xi V(\tau,\xi)\neq\emptyset$, then*

$$\dot{x}(t)\in\Phi(t,x(t)) \quad \text{for almost all } t\in[0,\tau]. \tag{4}$$

On the other hand, if $x(\cdot)$ is such that $x(\tau)=\xi$, (4) holds, and $x(t)\in\operatorname{int}\operatorname{dom}V(t,\cdot)$ for almost all $t\in[0,\tau]$ (as always is the case if either L or g is finite), then $x(\cdot)$ is optimal for $\mathcal{P}(\tau,\xi)$.

The mapping (3) has closed and compact values, however they need not be convex. In Section 3 we give several conditions on the Hamiltonian which guarantee that the optimal feedback mapping, and consequently the differential inclusion (4), is more regular.

Feedback results applicable to other nonsmooth control problems are known, see for example Berkovitz [1], Cannarsa and Frankowska [3] and Clarke, Ledyaev, Stern and Wolenski [6]. While these can handle nonlinear dynamics, they have limited applicability to problems with unbounded control sets or non-Lipschitz costs, this is outlined in further detail by Goebel [8], for a discussion of related issues in the Hamilton-Jacobi theory see Rockafellar and Wolenski [14] and Galbraith [7]. We add that the well-developed notion of viscosity solutions has a limited application in presence of convexity.

2 Hamiltonian Dynamical System

By a *Hamiltonian trajectory* on an interval $[a,b]$ we understand a pair of absolutely continuous arcs $(x(\cdot),y(\cdot)) : [a,b]\mapsto\mathbb{R}^n$ such that

$$\begin{cases} -\dot{y}(t)\in\tilde{\partial}_x H\left(x(t),y(t)\right), \\ \dot{x}(t)\in\partial_y H\left(x(t),y(t)\right), \end{cases} \quad \begin{matrix}\text{for almost}\\ \text{all } t\in[a,b].\end{matrix} \tag{5}$$

Above, $\tilde{\partial}_x H(x,y)$ is the subdifferential (in the concave sense) of the concave function $H(\cdot,y)$ and equals $-\partial_x(-H)(x,y)$, where the latter subdifferential should be understood in the convex sense. Both subdifferentials in (5) reduce to gradients if the Hamiltonian is differentiable. Key to our analysis of the regularity of the value function is the following description of its subgradients.

Theorem 2.1 (Rockafellar and Wolenski [14])**.** *A point (x_t,y_t) is in the graph of $\partial_\xi V(t,\cdot)$ if and only if for some $(x_0,y_0)\in\operatorname{gph}\partial g$, there is a Hamiltonian trajectory $(x(\cdot),y(\cdot))$ on $[0,t]$ with $(x(0),y(0))=(x_0,y_0)$, $(x(t),y(t))=(x_t,y_t)$.*

As an illustration of the above theorem, we show that the value function need not be piecewise linear-quadratic, even if the terminal cost function g is quadratic and the Hamiltonian is smooth, with piecewise linear gradient.

Example 2.2 Consider a one-dimensional problem of Bolza with the cost functions

$$L(x,v)=\frac{1}{2}v^2+\begin{cases}0 & x<0,\\ \frac{1}{2}x^2 & x\geq 0,\end{cases}\qquad g(x)=\frac{1}{2}(x+3)^2.$$

The corresponding Hamiltonian and its gradient are:

$$H(x,y)=\begin{cases}\frac{1}{2}y^2 & x<0,\\ -\frac{1}{2}x^2+\frac{1}{2}y^2 & x\geq 0,\end{cases}$$

$$\nabla H(x,y)=\begin{cases}(0,y) & x<0,\\ (-x,y) & x\geq 0.\end{cases}$$

A Hamiltonian trajectory $(x(\cdot),y(\cdot))$ must satisfy $\dot{x}(t)=y(t)=\mathit{const}$ when $x(t)<0$, and $x(t)=\alpha e^t+\beta e^{-t}$, $y(t)=\alpha e^t-\beta e^{-t}$ for suitably chosen α, β when $x(t)>0$. We parameterize the segment between $(-2,1)$ and $(-1,2)$ by $(x_s(0),y_s(0))=(s-2,s+1)$, $s\in[0,1]$, and calculate the Hamiltonian trajectories originating at points $(x_s(0),y_s(0))$. We get $(x_s(t),y_s(t))=((s+1)t+s-2,s+1)$ for $0\leq t\leq\frac{2-s}{s+1}$, and $(x_s(t),y_s(t))=(s+1)\Big(\sinh(t-\frac{2-s}{s+1}),\cosh(t-\frac{2-s}{s+1})\Big)$ for $t\geq\frac{2-s}{s+1}$. It is easy to check that for any $t>1$, the set $\{(x_s(t),y_s(t)),s\in[0,1]\}$ is not a straight line segment, nor is it a union of segments. Now note that the segment between $(-2,1)$ and $(-1,2)$ is in $\operatorname{gph}\nabla g$, while $\{(x_s(t),y_s(t)),s\in[0,1]\}\subset\operatorname{gph}\partial_\xi V(t,\cdot)$. Thus, $V(t,\cdot)$ is not piecewise linear-quadratic.

Theorem 2.1 suggests that the Hamiltonian dynamical system propagates all the information on ∂g, unless some Hamiltonian trajectories escape to infinity in finite time. A stronger conclusion can be then made, two different initial costs will not yield equal value functions at any time $\tau > 0$:

Theorem 2.3 (Goebel and Rockafellar [10]). *Let V_1, V_2 be value functions corresponding to initial costs g_1, g_2 and a Lagrangian L. Assume that for the Hamiltonian dynamical system associated with L, no trajectories escape to infinity in finite time. Then the following are equivalent, and either implies that $V_1 = V_2$.*

(a) For some $\tau > 0$, $V_1(\tau, \xi) = V_2(\tau, \xi)$ for all $\xi \in I\!R^n$.

(b) $g_1(x) = g_2(x)$ for all $x \in I\!R^n$.

Note that the result above does not require the uniqueness of Hamiltonian trajectories originating at any given point. We add that the assumption about no finite escape time holds for example for Hamiltonians depending only on x or only on y, for any piecewise linear-quadratic Hamiltonian, and for $H(x, y) = \langle Ax, y\rangle - h_1(x) + h_2(y)$.

The concavity and convexity of the Hamiltonian, and consequently the monotone structure of $\tilde{\partial}_x H$, $\partial_y H$, displayed in the fact that the mapping $(x, y) \mapsto (-\tilde{\partial}_x H(x, y), \partial_y H(x, y))$ is maximal monotone, has various implications on the nature of Hamiltonian trajectories. In particular, if $(x_i(\cdot), y_i(\cdot))$, $i = 1, 2$ are Hamiltonian trajectories on $[a, b]$, then the function

$$f(t) = \langle x_1(t) - x_2(t), y_1(t) - y_2(t)\rangle$$

is nondecreasing on $[a, b]$. This was noted by Rockafellar [12], where the study of generalized Hamiltonian systems for saddle Hamiltonians was initiated. Further assumptions of strict concavity or convexity of the Hamiltonian lead to f strictly increasing.

3 Regularity of the Value Function

We now employ the properties of the Hamiltonian differential inclusion to investigate the single-valuedness and continuity of the subdifferential $\partial_\xi V(\tau, \cdot)$. For simplicity of presentation, we assume:

- Value functions under consideration are finite.

This is always the case if either L or g is finite. An important consequence of single-valuedness of $\partial_\xi V(\tau, \cdot)$ is displayed in the equivalence of the following statements:

(a) For every $\tau > 0$, the subdifferential mapping $\partial_\xi V(\tau, \cdot)$ is single-valued.

(b) V is differentiable on $(0, \infty) \times I\!R^n$,

Under either of the above conditions, ∇V is continuous on $(0, \infty) \times I\!R^n$. These facts follow from in part from the Hamilton-Jacobi equation appropriate for the current setting, as described by Rockafellar and Wolenski [14].

For any finite convex function f, single-valuedness of ∂f is equivalent to the following: for any $y_1 \in \partial f(x_1)$, $y_2 \in \partial f(x_2)$, $\langle x_1 - x_2, y_1 - y_2\rangle > 0$ unless $y_1 = y_2$. The description of $\partial_\xi V(\tau, \cdot)$ of Theorem 2.1, combined with the behavior of $\langle x_1(1) - x_2(t), y_(1) - y_(2)\rangle$ for Hamiltonian trajectories $(x_i(\cdot), y_i(\cdot))$, as described at the end of Section 2, leads to the conclusion: if $H(x, \cdot)$ is strictly convex for every x, then $\partial_\xi V(\tau, \cdot)$ is single-valued for all $\tau > 0$. Thus, strict convexity of the Hamiltonian in y guarantees differentiability of the value function, on $(0, +\infty) \times I\!R^n$, independently of the regularity of g. However, the strict convexity of $H(x, \cdot)$ is equivalent to smoothness of $L(x, \cdot)$, a condition rarely satisfied, especially in presence of hard control constraints.

In the method of characteristics for the Hamilton-Jacobi equation, a "shock" is said to occur when two Hamiltonian trajectories (characteristic curves) originating on $\mathrm{gph}\, \partial g$ have different endpoints but with the same x-coordinate. The proof of the result below shows that in our convex setting, a shock can only happen in presence of nonuniqueness of Hamiltonian trajectories. This is in great contrast to the nonconvex case.

Theorem 3.1 (Goebel [8]). *Assume that the Hamiltonian H satisfies the following "uniqueness condition", for every $\tau > 0$:*

> *If $(x(\cdot), y_i(\cdot))$, $i = 1, 2$, are Hamiltonian trajectories on $[0, \tau]$ with the same initial condition $(x(0), y_i(0)) = (x_0, y_0)$, then $y_1(\cdot) = y_2(\cdot)$.*

If the initial cost function $g(\cdot)$ is differentiable, then so is $V(\tau, \cdot)$, for every $\tau > 0$.

Proof: Fix $\tau > 0$ and pick two points (x', y_1') and (x', y_2') in the graph of $\partial_\xi V(\tau, \cdot)$. Theorem 2.1 implies that there exist points $(x_i, y_i) \in \mathrm{gph}\, \partial g$ and Hamiltonian trajectories on $[0, \tau]$, $(x_i(\cdot), y_i(\cdot))$, from (x_i, y_i) to (x', y_i'), for $i = 1, 2$. We get

$$\begin{aligned}&\langle x_1 - x_2, y_1 - y_2\rangle = \langle x_1(0) - x_2(0), y_1(0) - y_2(0)\rangle\\ &\leq \langle x_1(\tau) - x_2(\tau), y_1(\tau) - y_2(\tau)\rangle = \langle x' - x', y_1' - y_2'\rangle,\end{aligned}$$

with the first quantity being nonnegative and the last one equal to 0. So $\langle x_1 - x_2, y_1 - y_2\rangle = 0$. Differentiability of g, and so single-valuedness of ∂g, is equivalent to $\langle x_1 - x_2, y_1 - y_2 \rangle > 0$ unless $y_1 = y_2$, and thus the last equality must hold. The above estimation shows also that $\langle x_1(t) - x_2(t), y_1(t) - y_2(t) \rangle$ is constant on $[0, \tau]$. Under this condition, the following pair of arcs is also a Hamiltonian trajectory, for any $\alpha, \beta \in [0, 1]$: $x'(\cdot) = (1-\alpha)x_1(\cdot) + \alpha x_2(\cdot)$, $y' = (1-\beta)y_1(\cdot) + \beta y_2(\cdot)$ (Rockafellar [12]). By taking $\alpha = 0$, $\beta = 1$ we see that $(x_1(\cdot), y_2(\cdot))$ is a Hamiltonian trajectory. Notice that $(x_1(0), y_2(0)) = (x_1, y_2) = (x_1, y_1) = (x_1(0), y_1(0))$. The uniqueness assumption implies that $y_1(\cdot) = y_2(\cdot)$, and in particular that $y_1' = y_2'$. This shows the single-valuedness of $\partial_\xi V(\tau, \cdot)$. ∎

Example 3.2 Consider a control problem

$$\text{minimize } g(x(0)) + \int_0^\tau l(x(t), u(t))\, dt$$
$$\text{subject to } \dot{x}(t) = u(t),\ u(t) \in U(x(t)),$$

with the cost and state-dependent control constraint set given by

$$\begin{aligned} l(x,u) &= \sum_{i=1}^n \begin{cases} 0 & \text{for} \quad u_i \geq 0, \\ \frac{1}{2}u_i^2 & \text{for} \quad u_i < 0, \end{cases} \\ U(x) &= \{u \mid u_i \leq x_i, i = 1, 2, \ldots, n\}. \end{aligned}$$

Translating this control problem to a Bolza setting yields the Lagrangian

$$L(x, v) = \frac{1}{2}|v_-|^2 + \delta_{\boldsymbol{R}_+^n}(x - v).$$

Here, the negative part v_- of v is described as $(v_-)_i = 0$ if $v_i \geq 0$ and $(v_-)_i = v_i$ if $v_i < 0$, while the indicator function $\delta_{\boldsymbol{R}_+^n}(z)$ equals 0 if $z \in \mathbb{R}_+^n$ and $+\infty$ otherwise.

As the Lagrangian is separable:

$$L(x, v) = \sum_{i=1}^n \left(\frac{1}{2}|(v_i)_-|^2 + \delta_{\boldsymbol{R}_+}(x_i - v_i) \right),$$

and for any convex function $f(z_1, z_2) = f_1(z_1) + f_2(z_2)$ one has $f^*(w_1, w_2) = f_1^*(w_1) + f_2^*(w_2)$, we only consider the Hamiltonian for $n = 1$, and obtain

$$H(x, y) = \begin{cases} \begin{cases} xy, & \text{for} \quad y \geq 0 \\ \frac{1}{2}y^2, & \text{for} \quad y < 0 \end{cases}, & \text{if} \quad x \geq 0, \\ \begin{cases} xy - \frac{1}{2}x^2, & \text{for} \quad y \geq x \\ \frac{1}{2}y^2, & \text{for} \quad y < x \end{cases}, & \text{if} \quad x < 0. \end{cases}$$

In other words, the Hamiltonian is given by $xy - \frac{1}{2}x^2$ in the region A, xy in B, and $\frac{1}{2}y^2$ in C.

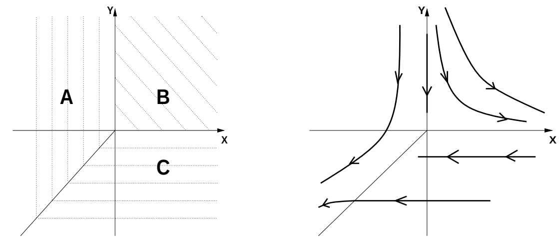

In the neighborhood of any point (x_0, y_0) with $x_0 \leq 0$ or $y_0 \neq 0$, the Hamiltonian is differentiable, and ∇H is Lipschitz continuous. As in such cases $\tilde{\partial}_x H = \nabla_x H$ and $\partial_y H = \nabla_y H$, trajectories emanating from such (x_0, y_0) are unique for small time intervals. Global analysis of such trajectories (approximate sketch is above on the right) shows that they remain in the described region where ∇H is Lipschitz continuous. Therefore, they are unique.

For points (x, y) such that $x \geq 0$, $y = 0$, we have $\tilde{\partial}_x H(x, y) = \{0\}$, $\partial_y H(x, y) = [0, x]$. Any trajectory originating at $(x_0, 0)$ with $x_0 \geq 0$ must satisfy $\dot{x}(t) \in [0, x(t)]$, $\dot{y}(t) = 0$, and so $y(t) = 0$, while $x(t)$ is nondecreasing. Thus, the "uniqueness condition" holds. Consequently, if the initial cost $g(\cdot)$ is differentiable, then so is the value function.

In the example above, while the Hamiltonian is not smooth, it is differentiable in x for any fixed y. Lipschitz behavior of $\nabla_x H$ is sufficient for the "uniqueness condition" to hold, that is, the latter is implied by the following: the Hamiltonian H is differentiable in the x-variable, and for every (x_0, y_0), there exists a neighborhood $X_0 \times Y_0$ of (x_0, y_0) and a constant K such that the following is true

$$\begin{aligned} &\text{for all } x \in X_0,\ y_1, y_2 \in Y_0, \\ &\quad |\nabla_x H(x, y_1) - \nabla_x H(x, y_2)| \leq K|y_1 - y_2|. \end{aligned}$$

Note that differentiability of V, even in presence of the condition above, does not guarantee that the optimal feedback map (3) is single-valued (it must, though, have convex values). Also notice that the lack of Lipschitz behavior of $\nabla_x H$ is not necessary for uniqueness of Hamiltonian trajectories even in both components: a simple analysis shows that Hamiltonian trajectories for $H(x, y) = -|x| + |y|$ are unique (this Hamiltonian corresponds to the Lagrangian $L(x, v) = |x|$ with the velocity v constrained to $[-1, 1]$).

Stronger regularity properties of the value function, possibly the Lipschitz continuity of its gradient, should be expected if the Hamiltonian is differentiable, and its gradients are Lipschitz continuous. A setting where smoothness of a function automatically entails global Lipschitz continuity of its gradient is that of piecewise linear-quadratic functions. Hamiltonians fitting this format naturally occur in various extensions of the classical linear-quadratic regulator problem. In the latter, the Lagrangian is a quadratic function

$$L(x,v) = \frac{1}{2}\langle x, Px\rangle + \frac{1}{2}\langle v, Qv\rangle,$$

with the matrices P, Q symmetric and positive semidefinite, and Q in fact positive definite. The Hamiltonian corresponding to the L above is also quadratic: $H(x,y) = -\frac{1}{2}\langle x, Px\rangle + \frac{1}{2}\langle y, Q^{-1}y\rangle$. Addition of simple constraints of the form $v(t) \in U$ for some polyhedral set U, not necessarily bounded nor conical, leads to a piecewise linear-quadratic and differentiable Hamiltonian

$$H(x,y) = -\frac{1}{2}\langle x, Px\rangle + \sup_{v\in U}\left\{\langle y, v\rangle - \frac{1}{2}\langle v, Qv\rangle\right\}.$$

Similarly, a linear-quadratic regulator with an extra state-dependent constraint $v(t) \leq Cx(t) + d$ for any matrix C, any vector d, leads to a differentiable and piecewise linear-quadratic H. On the other hand, constraints essentially guarantee (here, and in general convex problems) that the Hamiltonian will not be C^2. This poses an obstacle to Riccati-like descriptions of the gradient of the value function, as given for example by Byrnes [2].

As these simple examples suggest, a nonsmooth and infinite-valued Lagrangian may lead to a regular Hamiltonian. Further details, and in particular precise results on differentiability of the Hamiltonian for control problems in the extended piecewise linear-quadratic format, as proposed by Rockafellar [13], can be found in Goebel [9].

When the gradient of the Hamiltonian is globally Lipschitz, so is the mapping sending a point (x_0, y_0) to the endpoint of a Hamiltonian trajectory on $[0,\tau]$ originating at (x_0, y_0). In general, small Lipschitz perturbations of sets which are a graph of a Lipschitz map do not damage the latter property. Theorem 2.1 then suggests that $\nabla_\xi V(\tau,\cdot)$ will remain Lipschitz for small values of τ, provided that the initial cost g has a Lipschitz gradient. This in fact may hold, in some local sense, in absence of convexity (see Caroff and Frankowska [4], and note that the results described therein on differentiability of V – for a nonconvex setting – require regularity of the Lagrangian). In our setting, the structure of the Hamiltonian dynamical system, in particular the "monotonicity preserving" properties discussed at the end of Section 2, yield a stronger result.

Theorem 3.3 (Goebel [9])**.** *Suppose that both the Hamiltonian H and the initial cost g are differentiable and have globally Lipschitz gradients. For every fixed $\tau > 0$, $\nabla_\xi V(\tau,\cdot)$ is globally Lipschitz on $\mathbb{R}^n$, and in fact the Lipschitz constant can be chosen uniformly over τ on every bounded interval* $[0,T]$.

Consequently, under the assumptions of Theorem 3.3, the value function V is differentiable on $(0,+\infty)\times\mathbb{R}^n$ and ∇V is locally Lipschitz. The optimal feedback mapping (3) is locally Lipschitz continuous, and consequently the inclusion (4) with the endpoint condition $x(\tau) = \xi$ has unique solutions.

References

[1] L. Berkovitz. Optimal feedback controls. *SIAM J. Control Optim.*, 27(5):991–1006, 1989.

[2] C. Byrnes. On the Riccati partial differential equation for nonlinear Bolza and Lagrange problems. *J. Math. Systems Estim. Control*, 8(1):1–54, 1998.

[3] P. Cannarsa and H. Frankowska. Some characterizations of optimal trajectories in optimal control theory. *SIAM J. Control Optim.*, 29(6):1322–1347, 1991.

[4] N. Caroff and H. Frankowska. Conjugate points and shocks in nonlinear optimal control. *Trans. Amer. Math. Soc.* 348(8):3133–3153, 1996.

[5] F.H. Clarke. *Optimization and Nonsmooth Analysis.* Wiley, 1983.

[6] F.H. Clarke, Yu.S. Ledyaev, R.J. Stern, and P.R. Wolenski. *Nonsmooth Analysis and Control Theory.* Springer-Verlag, 1998.

[7] G.N. Galbraith. Extended Hamilton-Jacobi characterization of value functions in optimal control. *SIAM J. Control Optim.*, 39(1):281–305, 2001.

[8] R. Goebel. Regularity of the optimal feedback and the value function in convex problems of optimal control. submitted.

[9] R. Goebel. Convex optimal control problems with smooth Hamiltonians. submitted.

[10] R. Goebel and R.T. Rockafellar. Generalized conjugacy in Hamilton-Jacobi theory for fully convex Lagrangians. *J. Convex Anal.*, 9(1), 2002.

[11] R.T. Rockafellar. Conjugate convex functions in optimal control and the calculus of variations. *J. Math. Anal. Appl.*, 32:174–222, 1970.

[12] R.T. Rockafellar. Generalized Hamiltonian equations for convex problems of Lagrange. *Pacific J. Math.*, 33(2):411–427, 1970.

[13] R.T. Rockafellar. Linear-quadratic programming and optimal control. *SIAM J. Control Optim.*, 25(3):781–814, 1987.

[14] R.T. Rockafellar and P.R. Wolenski. Convexity in Hamilton-Jacobi theory, 1: Dynamics and duality. *SIAM J. Control Optim.*, 39(5):1323–1350, 2000.