

Symplectic Geometry of Constrained Optimization

Andrey A. Agrachev^{1,2*} and Ivan Yu. Beschastnyi^{2**}

¹PSI RAS

ul. Petra I 4a, Pereslavl-Zalessky, 152020 Russia

²SISSA

via Bonomea 265, Trieste, 34136 Italy

Received September 10, 2017; accepted November 07, 2017

Abstract—In this paper, we discuss geometric structures related to the Lagrange multipliers rule. The practical goal is to explain how to compute or estimate the Morse index of the second variation. Symplectic geometry allows one to effectively do it even for very degenerate problems with complicated constraints. The main geometric and analytic tool is an appropriately rearranged Maslov index.

We try to emphasize the geometric framework and omit analytic routine. Proofs are often replaced with informal explanations, but a well-trained mathematician will easily rewrite them in a conventional way. We believe that Vladimir Arnold would approve of such an attitude.

MSC2010 numbers: 49K15, 65K10

DOI: 10.1134/S1560354717060119

Keywords: optimal control, second variation, Lagrangian Grassmanian, Maslov index

To the memory of Vladimir Arnold

1. FIRST VARIATION IN FINITE DIMENSIONS

Our goal in this note will be to develop a general machinery for optimization problems using the language and results from symplectic geometry.

We start by discussing the Lagrange multiplier rule in the finite-dimensional setting. Let U be a finite-dimensional manifold, $\varphi : U \rightarrow \mathbb{R}$ a smooth function and $\Phi : U \rightarrow M$ a smooth submersion onto a finite-dimensional manifold M . We would like to find critical points of φ restricted to the level sets of Φ . It is well known that this can be done via the Lagrange multiplier rule.

Theorem 1 (Lagrange multiplier rule). *A point $u \in U$ is a critical point of $\varphi|_{\Phi^{-1}(q)}$, $q \in M$, if and only if there exists a covector $\lambda \in T_q^*M$ such that*

$$\begin{aligned} d_u\varphi &= \lambda D_u\Phi, \\ \Phi(u) &= q. \end{aligned} \tag{1.1}$$

This fact has an important geometric implication. Suppose that we have locally a smooth function $q \rightarrow u(q)$ such that each u satisfies the Lagrange multiplier rule. Then the covectors λ corresponding to $u(q)$ are just the values of the differential of the cost function $c(q) = \varphi(u(q))$. Indeed, we can differentiate the constraint equation $\Phi(u(q)) = q$ to get

$$\text{id} = (D_u\Phi) \frac{\partial u}{\partial q} \quad \Rightarrow \quad \lambda = (\lambda D_u\Phi) \frac{\partial u}{\partial q} = (d_u\varphi) \frac{\partial u}{\partial q} = d_q c.$$

So if we choose a branch of $u(q)$, the set of corresponding Lagrange multipliers is a graph of a differential of a smooth function that is a Lagrangian submanifold of the symplectic manifold T^*M . Let us briefly recall the symplectic terminology.

*E-mail: agrachevaa@gmail.com

**E-mail: i.beschastnyi@gmail.com

2. BASIC SYMPLECTIC TOPOLOGY

In this subsection we give some basic definitions from symplectic geometry. For further results and proofs see [6, 8].

A *symplectic space* is a pair (W, σ) of an even-dimensional vector space W and a skew-symmetric nondegenerate bilinear form σ . One can always choose a basis in W such that σ is of the form

$$\sigma(\lambda_1, \lambda_2) = \lambda_1^T J \lambda_2, \quad \lambda_i \in W,$$

where

$$J = \begin{pmatrix} 0 & \text{id} \\ -\text{id} & 0 \end{pmatrix}.$$

Such a basis is called a *Darboux basis*.

A *symplectic map* is a linear map $F : W \rightarrow W$ that preserves the symplectic structure, i. e.,

$$\sigma(F\lambda_1, F\lambda_2) = \sigma(\lambda_1, \lambda_2).$$

In a Darboux basis we can equivalently write

$$F^T J F = J.$$

We define the *skew-orthogonal complement* of a subspace Γ in a symplectic space W as a subspace

$$\Gamma^\perp = \{ \lambda \in W : \sigma(\lambda, \mu) = 0, \forall \mu \in \Gamma \}.$$

One has the following special situations:

- If $\Gamma \subset \Gamma^\perp$, then Γ is called *isotropic*;
- If $\Gamma \supset \Gamma^\perp$, then Γ is called *coisotropic*;
- If $\Gamma = \Gamma^\perp$, then Γ is called *Lagrangian*.

From the definition we can see that Γ is isotropic if and only if the restriction $\sigma|_\Gamma$ vanishes. Since σ is nondegenerate, we have

$$\dim \Gamma + \dim \Gamma^\perp = \dim W.$$

Therefore, a subspace Γ is Lagrangian if and only if Γ is isotropic and has dimension $(\dim W)/2$. Any one-dimensional subspace is isotropic by the skew-symmetry of σ . For the same reasons any codimension one subspace is coisotropic.

In a Darboux basis each vector $\lambda \in W$ has coordinates $(p_1, \dots, p_n, q^1, \dots, q^n) = (p, q)$, where $n = (\dim W)/2$. Then the subspaces defined by the equations $p = 0$ or $q = 0$ are Lagrangian. To construct more examples, we can consider a graph (p, Sp) of a linear map S between those subspaces. Then it is easy to check that (p, Sp) gives a Lagrangian subspace if and only if S is symmetric.

There exists a close relation between symplectic maps and Lagrangian subspaces. Given (W, σ) , we can construct a new symplectic space $(W \times W, (-\sigma) \oplus \sigma)$ of double dimension. It can be used to give an alternative definition of a symplectic map.

Proposition 1. *Let $F : W \rightarrow W$ be a linear map. F is symplectic if and only if the graph of F in $(W \times W, (-\sigma) \oplus \sigma)$ is Lagrangian.*

We can extend all these definitions to the nonlinear setting. A *symplectic manifold* is a pair (W, σ) where W is a smooth manifold and σ is a closed nondegenerate differential two-form. Similarly to the linear case, one can show that locally all symplectic manifolds have the same structure.

Theorem 2 (Darboux). *For any point x of a symplectic manifold (W, σ) one can find a neighborhood U and a local diffeomorphism $\psi : U \rightarrow \mathbb{R}^{2n}$ such that*

$$\sigma = \psi^*(dp_i \wedge dq^i),$$

where (p, q) are coordinates in \mathbb{R}^{2n} .

Note that a tangent space $T_x W$ has naturally a structure of a symplectic space. Therefore, we can say that a submanifold $N \subset M$ is *isotropic/coisotropic/Lagrangian* if the same property is true for each subspace $T_x N \subset T_x W$ for any $x \in N$. Similarly to the linear case, a submanifold N is isotropic if and only if $\sigma|_N = 0$ and Lagrangian if additionally $\dim N = (\dim W)/2$.

A *symplectomorphism* of (W, σ) is a smooth map $f : W \rightarrow W$ that preserves the symplectic structure, i. e.,

$$f^* \sigma = \sigma.$$

Given a smooth function $h : W \rightarrow \mathbb{R}$, a *Hamiltonian vector field* \vec{h} is defined by the identity $dh = \sigma(\cdot, \vec{h})$. The flow generated by the *Hamiltonian system* $\dot{x} = \vec{h}(x)$ preserves the symplectic structure. In Darboux coordinates, the Hamiltonian system has the form

$$\dot{p} = -\frac{\partial h}{\partial q}, \quad \dot{q} = \frac{\partial h}{\partial p}.$$

The nonlinear analogue of Proposition 1 holds as well.

Proposition 2. *A diffeomorphism $f : W \rightarrow W$ of a symplectic manifold (W, σ) is a symplectomorphism if and only if the graph of f in $(W \times W, (-\sigma) \oplus \sigma)$ is a Lagrangian submanifold.*

The most basic and important examples of symplectic manifolds are the cotangent bundles T^*M . To define invariantly the symplectic form, we use the projection map $\pi : T^*M \rightarrow M$. Its differential is a well-defined map $\pi_* : T(T^*M) \rightarrow TM$. We can define the *Liouville one-form* $s \in \Lambda^1(T^*M)$ at $\lambda \in T^*M$ as

$$s_\lambda = \lambda \circ \pi_*.$$

Then the canonical symplectic form on T^*M is simply given by the differential $\sigma = ds$.

In local coordinates T^*M is locally diffeomorphic to $\mathbb{R}^n \times \mathbb{R}^n$ with coordinates (p, q) , where q are coordinates on the base and p are coordinates on the fiber. In these coordinates the Liouville form s is written as $s = p_i dq^i$. Thus, (p, q) are actually Darboux coordinates. We can use this fact to construct many Lagrangian manifolds. Namely:

Proposition 3. *Let $S : M \rightarrow \mathbb{R}$ be a smooth function. Then the graph of the differential $d_q S$ is a Lagrangian submanifold in T^*M .*

The proof is a straightforward computation in the Darboux coordinates and follows from the commutativity of the second derivative of S .

We have seen in the previous subsection that the set of Lagrange multipliers is often a graph of the differential of a smooth function. One can reformulate this by saying that the set of Lagrange multipliers is actually a Lagrangian submanifold. In the next sections we will see that this is a rather general fact, but the resulting “Lagrangian set” can be quite complicated. So we will linearize our problem and extract optimality information from the behavior of what is going to be “tangent spaces” to this “Lagrangian sets”. In this way we obtain a geometric theory of second variation that is applicable to a very large class of optimization problems.

3. FIRST VARIATION FOR CLASSICAL CALCULUS OF VARIATIONS

Let us consider a geometric formulation of the classical problem of calculus of variations, which is an infinite-dimensional optimization problem. We denote by \mathcal{U} the set of Lipschitzian curves $\gamma : [\tau, t] \rightarrow M$, where M is a finite-dimensional manifold. Assume that this set is endowed with a nice topology of a Hilbert manifold. We consider a family of functionals

$$J_\tau^t : \gamma \mapsto \int_\tau^t l(\gamma(s), \dot{\gamma}(s)) ds.$$

We define the evaluation map $F_s : \mathcal{U} \rightarrow M$ to be a map that takes a curve and returns a point on it at a time $s \in [\tau, t]$, i. e., $F_s(\gamma) = \gamma_s$. We look for the critical points of the restriction

$$J_\tau^t|_{\{F_\tau^{-1}(q_\tau), F_t^{-1}(q_t)\}}.$$

So we apply the Lagrange multiplier rule to $\varphi = J_\tau^t$ with $\Phi = (F_\tau, F_t)$ and find that there exists a pair $(-\lambda_\tau, \lambda_t) \in T_{q(\tau)}^*M \times T_{q(t)}^*M$ such that

$$d_\gamma J_\tau^t = \lambda_t D_\gamma F_t - \lambda_\tau D_\gamma F_\tau. \tag{3.1}$$

The evaluation map is obviously a submersion. This fact implies that once we fix γ and one of the covectors $(-\lambda_\tau, \lambda_t)$, the other one will be determined automatically. Indeed, suppose, for example, that $(-\lambda_\tau, \lambda_t)$ and $(-\lambda_\tau, \lambda'_t)$ both satisfy the Lagrange multiplier rule. Then in addition to (3.1) we have

$$d_\gamma J_\tau^t = \lambda'_t D_\gamma F_t - \lambda_\tau D_\gamma F_\tau.$$

We subtract this equation from (3.1) and obtain

$$(\lambda'_t - \lambda_t) D_\gamma F_t = 0,$$

which gives a contradiction with the fact that F_t is a submersion.

As in the finite-dimensional case, under some regularity assumptions the Lagrange multipliers form a Lagrangian submanifold in $T^*M \times T^*M$ endowed with a symplectic form $(-\sigma) \oplus \sigma$. Since each λ_τ determines a unique λ_t , we get that this Lagrangian submanifold can be identified with a graph of some map $A_\tau^t : T^*M \rightarrow T^*M$. Such a graph is a Lagrangian submanifold if and only if A_τ^t is symplectic. Moreover, the identity $J_\tau^t = J_\tau^s + J_s^t$ implies that A_τ^t is actually a symplectic flow, i. e., $A_\tau^t = A_s^t \circ A_\tau^s$, where $s \in [\tau, t]$, $A_\tau^\tau = \text{Id}$.

Since we have a symplectic flow, it should come from a Hamiltonian system. Let us find an expression for the corresponding Hamiltonian. We introduce some local coordinates (p, q) on T^*M . Then our extremal curve γ is given by a map $s \rightarrow q(s)$. We denote $\dot{q} = v$ and write down the Eq. (3.1)

$$\int_\tau^t \left(\frac{\partial l}{\partial q} dq_s + \frac{\partial l}{\partial v} dv_s \right) ds = p_t dq_t - p_\tau dq_\tau.$$

We differentiate this expression w.r.t. time t :

$$\frac{\partial l}{\partial q} dq_t + \frac{\partial l}{\partial v} dv_t = \dot{p}_t dq_t + p_t dv_t.$$

Then we obtain

$$\begin{aligned} \dot{q}_t &= v_t, \\ \dot{p}_t &= \frac{\partial l}{\partial q}, \\ p_t &= \frac{\partial l}{\partial v}. \end{aligned} \tag{3.2}$$

The first two equations can be seen to be a Hamiltonian system

$$\begin{aligned} \dot{q}_t &= \frac{\partial H}{\partial p}, \\ \dot{p}_t &= -\frac{\partial H}{\partial q}, \end{aligned} \tag{3.3}$$

with a Hamiltonian

$$H(p, q, v) = \langle p, v \rangle - l(q, v).$$

The third equation gives the condition

$$\frac{\partial H}{\partial v} = 0.$$

If the second derivative of the Hamiltonian H w.r.t. v is nondegenerate, then by the inverse function theorem we can locally resolve this condition to obtain a function $v = v(p, q)$. Substituting it in $H(p, q, v)$, we get an autonomous Hamiltonian system with a Hamiltonian $H(p, q, v(p, q))$.

4. SECOND VARIATION

Now we are going back to the general setting with $\varphi : U \rightarrow \mathbb{R}$ and constraints $\Phi : U \rightarrow M$. Once we have found a critical point $u \in U$, we would like to study the index of the Hessian, which is a quadratic form

$$\text{Hess}_u \varphi|_{\Phi^{-1}(q)} = \ker D_u \Phi \times \ker D_u \Phi \rightarrow \mathbb{R}.$$

We can write an explicit expression for the Hessian without resolving the constraints in the spirit of the Lagrange multiplier rule. Consider a curve $u(t) \in \Phi^{-1}(q)$ such that $v = \dot{u}(0) \in \ker D_{u(0)} \Phi$. Then using the Lagrange multiplier rule, we obtain

$$\frac{d^2}{dt^2} \Big|_{t=0} \varphi = \frac{d}{dt} \Big|_{t=0} \frac{d\varphi}{du} \dot{u} = \left\langle \frac{d^2\varphi}{du^2} \dot{u}, \dot{u} \right\rangle + \frac{d\varphi}{du} \ddot{u} = \left\langle \frac{d^2\varphi}{du^2} v, v \right\rangle + p \frac{d\Phi}{du} \ddot{u}(0).$$

On the other hand, we can twice differentiate the constraints $\Phi(u(t)) = q$. We get similarly

$$\left\langle \frac{d^2\Phi}{du^2} v, v \right\rangle + \frac{d\Phi}{du} \ddot{u}(0) = 0.$$

If we assume that U is a Hilbert manifold, then the two expressions give a formula for the Hessian in local coordinates that can be written as follows:

$$\text{Hess}_u \varphi(v, v) = \left\langle \left(\frac{d^2\varphi}{du^2} - p \frac{d^2\Phi}{du^2} \right) v, v \right\rangle := \langle Qv, v \rangle$$

such that v satisfies

$$\frac{d\Phi}{du} v = 0.$$

We define

$$L = \left\{ (p, q, u) : \Phi(u) = q, \frac{d\varphi}{du} - p \frac{d\Phi}{du} = 0 \right\},$$

and

$$\mathcal{L} = \pi(L),$$

where $\pi(u, \lambda) = \lambda$. \mathcal{L} is the set of all Lagrangian multipliers. We say that (Φ, φ) is a *Morse pair* (or a *Morse problem*) if Eq. (1.1) is regular, i. e., zero is a regular value for the map

$$(p, q, u) \mapsto \frac{d\varphi}{du} - p \frac{d\Phi}{du}. \tag{4.1}$$

If $\dim U < \infty$, then generically constrained optimization problems are Morse. Not all functions $\varphi|_{\Phi^{-1}(q)}$ are Morse, though, as the name could suggest. The Morse property of a constrained optimization problem implies the following important facts.

Proposition 4. *Let (φ, Φ) be a Morse problem. Then L is a smooth manifold and $\pi|_L$ is a Lagrangian immersion into T^*M .*

The main corollary of this proposition is that \mathcal{L} has a well-defined tangent Lagrangian subspace at each point

$$L_{(u,\lambda)}(\varphi, \Phi) = \left\{ (\delta p, \delta q) : \exists \delta u \in T_u U; \frac{d\Phi}{du} \delta u = \delta q, Q \delta u = \delta p \frac{d\Phi}{du} \right\}, \tag{4.2}$$

and these subspaces will be the main objects of our study.

Before proving the last proposition, we prove a lemma.

Lemma 1. *The point (p, q, u) is regular for the map (4.1) if and only if $\text{im } Q$ is closed and*

$$\ker Q \cap \ker \frac{d\Phi}{du} = 0.$$

Proof. We compute the differential of the map (4.1):

$$Q\delta u - \delta p \frac{d\Phi}{du}.$$

The differential is surjective if and only if its image is closed and has a trivial orthogonal complement or, equivalently, the existence of $w \in T_u U$ such that for any $(\delta u, \delta p)$

$$\langle Q\delta u, w \rangle - \delta p \frac{d\Phi}{du} w = 0 \tag{4.3}$$

implies that $w = 0$. But since $(\delta u, \delta p)$ are arbitrary and Q is symmetric, (4.3) is equivalent to the existence of $w \in T_u U$, such that we have simultaneously

$$Qw = 0, \quad \frac{d\Phi}{du} w = 0.$$

Then by assumption we have $w = 0$ and the result follows. □

Proof (of Proposition 4). The fact that L is a manifold is just a consequence of the implicit function theorem. We prove now that π_M is an immersion. The differential of this map takes the tangent space to L and maps it to the space $L_{(u,\lambda)}(\varphi, \Phi)$ (see (4.2)). It “forgets” δu . So a nontrivial kernel of the differential must lie in the subspace $\delta p = \delta q = 0$. But from the definition of $L_{u,\lambda}(\varphi, \Phi)$ we have that in this case $Q\delta u = 0$ and $D_u\Phi\delta u = 0$, which contradicts the fact that the problem is Morse, as it can be seen from the previous lemma. Thus, the differential is injective.

To prove that this immersion is Lagrangian, it is enough to prove that $L_{(u,\lambda)}(\varphi, \Phi)$ is a Lagrangian subspace. Since Q is symmetric, it is easy to see that this subspace is isotropic. Take $(\delta p_1, \delta q_1)$ and $(\delta p_2, \delta q_2)$ in $L_{(u,\lambda)}(\varphi, \Phi)$. Then we compute

$$\begin{aligned} \sigma((\delta p_1, \delta q_1), (\delta p_2, \delta q_2)) &= \delta p_1 \delta q_2 - \delta p_2 \delta q_1 \\ &= \delta p_1 \frac{d\Phi}{du} \delta u_2 - \delta p_2 \frac{d\Phi}{du} \delta u_1 = \langle Q\delta u_1, \delta u_2 \rangle - \langle Q\delta u_2, \delta u_1 \rangle = 0. \end{aligned}$$

Now it just remains to prove that the dimension of this space is equal to n . We are going to do it only in the finite-dimensional setting, but this is true in general [3]. Note that if we fix $(\delta p, \delta u)$ as in the definition of $L_{(u,\lambda)}(\varphi, \Phi)$, then δq is determined automatically. So it is enough to study the map

$$S : (\delta u, \delta p) \mapsto Q\delta u - \delta p \frac{d\Phi}{du}.$$

Then, clearly, $\dim L_{(u,\lambda)}(\varphi, \Phi) = \dim \ker S$. But we have seen in the proof of the previous lemma that this map is actually surjective. Then $\dim \operatorname{im} S = \dim U$ and we have

$$\dim \ker S = \dim(U \times \mathbb{R}^n) - \dim U = n. \tag{□}$$

In a Morse problem, the Lagrange submanifold \mathcal{L} often contains all necessary information about the Morse index and the nullity of the Hessian. We can already give a geometric characterization of the nullity, while for the geometric characterization of the index we will need more facts from linear symplectic geometry.

Proposition 5. *Hess_u φ has a nontrivial kernel if and only if λ is a critical point of the map π_M|_ℒ, where π_M : T**M* → *M* is the standard projection. The dimension of the kernel of the Hessian is equal to the dimension of the kernel of the differential of π_M|_ℒ.*

Schematically this situation is depicted in Fig. 1 on the right.




Fig. 1. Lagrangian manifolds in the simplest case of $\dim M = 1$.

Proof. Note that λ is a critical point of $\pi_M|_{\mathcal{L}}$ if and only if the tangent space $L_{(u,\lambda)}(\varphi, \Phi)$ contains a vertical direction $(\delta p, 0)$. If this is the case, by definition there exists $\delta u \in \ker D_u \Phi$ such that

$$Q\delta u = \delta p \frac{d\Phi}{du}. \tag{4.4}$$

Then, clearly, for any $v \in \ker D_u \Phi$, we have $\langle Q\delta u, v \rangle = 0$.

On the contrary, if $D_u \Phi \delta u = 0$ and δu belongs to the kernel of the Hessian, then for any $v \in \ker D_u \Phi$ we have $\langle Q\delta u, v \rangle = 0$ and $Q\delta u$ must be a linear combination of the rows of $D_u \Phi$, i. e., there exists $\delta p \in T_\lambda(T_q^*M)$ such that (4.4) holds. \square

If the problem is Morse, then the subspace $L_{(u,\lambda)}(\varphi, \Phi)$ defined in Section 4.2 is Lagrangian. However, our goal is to handle degenerate cases, and a starting point is the following surprising fact, which is valid without any regularity assumption on (φ, Φ) ; in particular, u may be a critical point of Φ .

Proposition 6. *If U is finite-dimensional, then the space $L_{(u,\lambda)}(\varphi, \Phi)$ defined in (4.2) is a Lagrangian subspace of $T_\lambda(T^*M)$.*

We leave the proof of this interesting linear algebra exercise to the reader. The next example shows that finite-dimensionality of U is essential.

Let $M = \mathbb{R}$ and U be a Hilbert space. Then $a = \Phi'_u$ is an element of the dual space U^* , $T^*M = \mathbb{R}^2$, and Lagrangian subspaces are just one-dimensional subspaces of \mathbb{R}^2 . We set $L = L_{(u,\lambda)}(\varphi, \Phi)$. By the definition we have

$$L = \{(\delta p, \delta q) : \exists v \in U, Qv = (\delta p)a, \delta q = \langle a, v \rangle\}.$$

Assume that $Q : U \rightarrow U^*$ is injective and not surjective. If $a \notin \text{im } Q$, then $L = \{(0, 0)\}$ with a unique lift $v = 0$.

Injectivity and symmetricity of Q imply that $\text{im } Q$ is everywhere dense in U , hence $\text{im } Q$ is not closed in the example just described. Now we drop the injectivity assumption but assume that $\text{im } Q = \overline{\text{im } Q}$. Let us show that L is 1-dimensional in this case. Indeed, the self-adjointness of Q implies that $\ker Q \oplus \text{im } Q = U$. Then we have two possible situations:

1. $a \in \text{im } Q$. Then there is a unique preimage of a in $\text{im } Q$, which we denote by $v = Q^{-1}a$. We get

$$L = \text{span} \{(1, \langle a, Q^{-1}a \rangle)\},$$

where Q^{-1} is a pseudo-inverse.

2. $a \notin \text{im } Q$, then we must have $\delta p = 0$ and $v \in \ker Q$. There exists $v \in \ker Q$ such that $\langle a, v \rangle \neq 0$, and we obtain

$$L = \text{span} \{(0, \langle a, v \rangle)\}.$$

If $\dim U < \infty$, then $\text{im } Q$ is automatically closed and L is Lagrangian as we have seen.

5. LAGRANGIAN GRASSMANIAN AND MASLOV INDEX

We are going to give a geometric interpretation of the Morse index of the Hessian in terms of some curves of Lagrangian subspaces. To do this, we need some results about the geometry of the set of all Lagrangian subspaces of a given symplectic space (W, σ) . This set has a structure of a smooth manifold and is called the *Lagrangian Grassmanian* $L(W)$. We give just the basic facts about $L(W)$. For more information see [6].

To construct a chart of this manifold, we fix a Lagrangian subspace $\Lambda_2 \in L(W)$ and consider the set of all Lagrangian subspaces transversal to Λ_2 , which we denote by Λ_2^\pitchfork (the symbol \pitchfork means “transversal”). By applying a Gram-Schmidt like procedure that involves σ , we can find some Darboux coordinates (p, q) on W such that $\Lambda_2 = \{(0, q)\}$ (see [3] for details). Then $\Lambda_0 = \{(p, 0)\}$ belongs to Λ_2^\pitchfork and any other $\Lambda_1 \in \Lambda_2^\pitchfork$ can be defined as a graph of a linear map from Λ_0 to Λ_2 . As we have seen in Section 2, the matrix of this map is symmetric and we obtain the identification of Λ_2^\pitchfork with the space of symmetric $n \times n$ -matrices that gives the desired local coordinates on $L(W)$.




Fig. 2. Defining a Lagrangian plane Λ_1 as a graph of a quadratic form from Λ_0 to Λ_2 .

Since symplectic maps preserve the symplectic form, they also map Lagrangian subspaces to Lagrangian subspaces. This action is transitive, i.e., there are no invariants of the symplectic group acting on $L(W)$. If we consider the action of the symplectic group on pairs of Lagrangian spaces, then the only invariant is the dimension of their intersection [6]. But triples $(\Lambda_0, \Lambda_1, \Lambda_2)$ do have a nontrivial invariant, which is called the *Maslov index* of the triple or the *Kashiwara index*.

To define it, suppose that $\Lambda_0, \Lambda_1, \Lambda_2 \in L(W)$, such that $\Lambda_0 \pitchfork \Lambda_1$ and $\Lambda_0 \pitchfork \Lambda_2$. Since the symplectic group acts transitively on the space of pairs of transversal Lagrangian planes, we can assume without any loss of generality that $\Lambda_0 = \{(p, 0)\}$ and $\Lambda_2 = \{(0, q)\}$. Then we can identify Λ_1 with a graph $\{(p, Sp)\}$, where S is a symmetric matrix. The invariant μ of this triple is then defined as

$$\mu(\Lambda_0, \Lambda_1, \Lambda_2) = \text{sgn } S.$$

This invariant can be also defined intrinsically as the signature of a quadratic form $\tilde{q} : \Lambda_1 \times \Lambda_1 \rightarrow R$ that is defined as follows. Since $\Lambda_0 \pitchfork \Lambda_2$, any $\lambda \in \Lambda_1$ can be decomposed as $\lambda = \lambda_0 + \lambda_2$, where $\lambda_i \in \Lambda_i$. Then we set $\tilde{q}(\lambda) = \sigma(\lambda_0, \lambda_2)$. One can check that those definitions agree.

This invariant has a couple of useful algebraic properties. The simplest one is the antisymmetry:

$$\mu(\Lambda_2, \Lambda_1, \Lambda_0) = -\mu(\Lambda_0, \Lambda_1, \Lambda_2);$$

$$\mu(\Lambda_0, \Lambda_2, \Lambda_1) = -\text{sgn}(S^{-1}) = -\mu(\Lambda_0, \Lambda_1, \Lambda_2).$$

We are going to state and prove another important property, called the chain rule, after we look more carefully at the geometry of $L(W)$. Let us fix some $\Delta \in L(W)$. The *Maslov train* \mathcal{M}_Δ is the set $\mathcal{M}_\Delta = L(W) \setminus \Delta^\pitchfork$ of Lagrangian planes that have a nontrivial intersection with Δ . It is an algebraic hypersurface with singularities, and its intersection with a coordinate chart containing Δ can be identified with the set of degenerate symmetric matrices.

The set of nonsmooth points of the hypersurface \mathcal{M}_Δ consists of Lagrangian subspaces that have an intersection with Δ of dimension two or more. It is easy to check that this singular part has codimension two in \mathcal{M}_Δ . It follows that the intersection number mod 2 of \mathcal{M}_Δ with any continuous curve whose endpoints do not belong to \mathcal{M}_Δ is well defined and is homotopy invariant. For example, when $\dim W = 4$, we have that the intersection of \mathcal{M}_Δ with a coordinate chart is identified with 2×2 symmetric matrices with zero determinant. This is a cone whose points except the origin correspond to Lagrangian planes that have a one-dimensional intersection with Δ . The origin represents the dimension two intersection with Δ , which is equal to Δ itself in this case. Clearly, a general position curve in $L(W)$ does not intersect the origin (see Fig. 3) as well as a general position homotopy of curves, and so the intersection number mod 2 is well defined.




Fig. 3. A curve in a local chart of the Lagrangian Grassmanian $L(\mathbb{R}^4)$ and the Maslov train.

We would like to define the integer-valued intersection number of curves in $L(W)$ with \mathcal{M}_Δ , and for this we need a coorientation of the smooth part of \mathcal{M}_Δ . Let $\Lambda_t \in L(W)$ such that $\Lambda_0 = \Lambda$ and consider any $\lambda_t \in \Lambda_t$ such that $\lambda_0 = \lambda$. Then we define

$$\dot{\underline{\Delta}}(\lambda) = \sigma(\lambda, \dot{\lambda}).$$

Thus, we see that to any tangent vector $\dot{\Lambda}$ we can associate a quadratic form $\dot{\underline{\Delta}}$. It is easy to see that $\dot{\underline{\Delta}}(\lambda)$ is indeed a well-defined quadratic form, i. e., that $\sigma(\lambda, \dot{\lambda})$ depends only on $\dot{\Lambda}$ and λ . Moreover, $\dot{\Lambda} \mapsto \dot{\underline{\Delta}}(\lambda)$, $\dot{\Lambda} \in T_\Lambda L(W)$ is an isomorphism of $T_\Lambda L(W)$ on the space of quadratic forms on Λ .

From the previous discussion we know that $\dim(\Lambda_t \cap \Delta) = 1$ if Λ_t is a smooth point of \mathcal{M}_Δ . Let $\lambda \in \Lambda_t \cap \Delta$, $\lambda \neq 0$; the intersection is transversal at the point Λ_t if and only if $\dot{\underline{\Delta}}(\lambda) \neq 0$. We say that the sign of the intersection is positive if $\dot{\underline{\Delta}}(\lambda) > 0$ and negative otherwise. The intersection number of a continuous curve $t \mapsto \Lambda_t$ with \mathcal{M}_Δ is called the Maslov index of the curve with respect to \mathcal{M}_Δ . Note that the Maslov index of a closed curve (i. e., a curve without endpoint) does not depend on the choice of Δ . Indeed, the train \mathcal{M}_Δ can be transformed to any other train by a continuous one-parametric family of symplectic transformations and the Maslov index is a homotopy invariant.

The Maslov index of a curve and the Maslov index of a triple are closely related. Let $\gamma : [0, 1] \rightarrow L(W)$ such that the whole curve does not leave the chart Δ^\natural . Then we have

$$2\gamma \circ \mathcal{M}_\Lambda = \mu(\Lambda, \gamma(1), \Delta) - \mu(\Lambda, \gamma(0), \Delta), \tag{5.1}$$

which easily follows from definitions.

Now we can state the last property.

Lemma 2 (The chain rule). *Let $\Lambda_i \in L(W)$, $i = 0, 1, 2, 3$. Then*

$$\mu(\Lambda_0, \Lambda_1, \Lambda_2) + \mu(\Lambda_1, \Lambda_2, \Lambda_3) + \mu(\Lambda_2, \Lambda_3, \Lambda_0) + \mu(\Lambda_3, \Lambda_0, \Lambda_1) = 0.$$

Proof. Connect Λ_0, Λ_2 with two curves: one that is completely in Λ_3^\natural and another that is completely in Λ_1^\natural . Schematically this situation is depicted in Fig. 4. This gives a closed curve γ in $L(W)$ and we can compute its intersection with \mathcal{M}_{Λ_1} and \mathcal{M}_{Λ_3} . The chain rule follows from the identities $\gamma \circ \mathcal{M}_{\Lambda_1} = \gamma \circ \mathcal{M}_{\Lambda_3}$ and (5.1). □




Fig. 4. An illustration to the proof. Λ_3^{\natural} correspond to the region with vertical lines and Λ_1^{\natural} to the region with the horizontal ones.

The formula (5.1) allows us to compute the Maslov index of a continuous curve, without putting it in general position and really computing the intersection. We just need to split the whole curve into small pieces such that each of them lies in a single coordinate chart, and then compute the index of the corresponding triples. This motivates the following definition. A curve $\gamma(t) \in L(W)$ is called *simple* if there exists $\Delta \in L(W)$ such that $\gamma(t) \in \Delta^{\natural}$.

6. MORSE INDEX

Now we return to the study of the second variation of a finite-dimensional Morse problem. We have seen that \mathcal{L} is an immersed Lagrangian submanifold. This allows us to define the Maslov cocycle in the following way. We fix a curve $\gamma : [t_0, t_1] \rightarrow \mathcal{L}$ and consider the corresponding curve $L_t : t \mapsto T_{\gamma(t)}\mathcal{L}$. Along $\gamma(t)$ we have vertical subspaces $\Pi = T_{\gamma(t)}(T_{q(t)}^*M)$ which lie in different symplectic spaces. Vector bundles over the segment $[t_0, t_1]$ are trivial, so by using a homotopy argument, we can assume that the symplectic space and the vertical subspace Π are fixed. We define the *Maslov cocycle* as

$$\mu(\gamma(t)) = L_t \circ \mathcal{M}_{\Pi}.$$

Then we have the following theorem.

Theorem 3. *Let $\gamma(t)$ be a curve connecting q_{t_0} with q_{t_1} . Then*

$$-2\mu(\gamma) = \delta(\text{sgn Hess}_u \varphi|_{\Phi^{-1}(q)}) := \text{sgn Hess}_{u_{t_1}} \varphi|_{\Phi^{-1}(q_{t_1})} - \text{sgn Hess}_{u_{t_0}} \varphi|_{\Phi^{-1}(q_{t_0})}.$$

Proof (A sketch). We denote $\overline{Q}(v) = \langle Qv, v \rangle$. Then

$$\text{Hess}_u \varphi|_{\Phi^{-1}(q)} = \overline{Q}|_{\ker \Phi'_u}, \quad \Phi(u) = q$$

and the difference of the signatures will be the difference of the corresponding signatures of $\overline{Q}|_{\ker \Phi'_u}$. Now we do not restrict Q to the kernel of Φ'_u and use the fact that it depends on the choice of coordinates. Indeed, the vertical subspace is fixed, but we have a freedom in choosing the horizontal space.

Exercise: Given a Lagrangian subspace $\Lambda \subset T_{\lambda}(T^*M)$ that is transversal to the fiber T_q^*M , there exist local coordinates in which $\Lambda = \{(0, x) : x \in \mathbb{R}^n\}$. Operator Q is nondegenerate iff the horizontal subspace $\{(0, x) : x \in \mathbb{R}^n\}$ is transversal to $L_{(u,\lambda)}(\varphi, \Phi)$.

To prove the theorem, we may divide the curve into small pieces and check the identity separately for each piece. In other words, we can assume that the curve is contained in the given coordinate chart and, according to the exercise, that Q_t is not degenerate for all $t \in [t_0, t_1]$.

Then we can apply the following linear algebra lemma.

Lemma 3. *Let E be a possibly infinite-dimensional Hilbert space, Q a quadratic form on E that is positive definite on a finite codimension subspace and V a closed subspace of E . Then if we denote by V_Q^\perp the orthogonal complement of V in E w.r.t. Q , the following formula is valid:*

$$\text{ind}^- Q = \text{ind}^- Q|_V + \text{ind}^- Q_{V_Q^\perp} + \dim(V \cap V_Q^\perp) - \dim(V \cap \ker Q).$$

Moreover, if $Q|_V$ is nondegenerate, then $\dim(V \cap V_Q^\perp) = \dim(V \cap \ker Q) = 0$.

Thus, by construction we get

$$\text{sgn } \overline{Q}_{t_i} = \text{sgn } \overline{Q}_{t_i}|_{\ker \Phi'_u} + \text{sgn } \overline{Q}_{t_i}|_{(\ker \Phi'_u)_{Q_{t_i}}^\perp}.$$

Since Q_t is nondegenerate and continuously depends on t , we have $\text{sgn } Q_{t_1} = \text{sgn } Q_{t_0}$. Then the first summand is just the Hessian and one can show that

$$\text{sgn } \overline{Q}_{t_i}|_{(\ker \Phi'_u)_{Q_{t_i}}^\perp} = \text{sgn } S_{t_i} = \mu(\text{Ver}, L_t, \text{Hor}),$$

where $\text{Ver} = \{(\xi, 0) : \xi \in \mathbb{R}^n\}$, $\text{Hor} = \{(0, x) : x \in \mathbb{R}^n\}$.

The statement of the theorem now follows from (5.1). □

What about the infinite-dimensional Morse problem? The signature of the Hessian is not defined in this case, but the difference of the signatures can be replaced with the difference of the Morse indices if $\text{Hess}_u \varphi|_{\Phi^{-1}(q)}$ are positive definite on a finite codimension subspace and $\text{Hess}_{u_{t_i}} \varphi|_{\Phi^{-1}(q_{t_i})}$, $i = 0, 1$, are nondegenerate.

7. GENERAL CASE

Consider now a general, not necessarily Morse, constrained optimization problem (φ, Φ) and a couple (u, λ) that satisfies the Lagrange multiplier rule $d_u \varphi = \lambda D_u \Phi$. In local coordinates:

$$\lambda = (p, q); \quad \varphi'_u = p\Phi'_u, \quad \Phi(u) = q.$$

We would like to consider the subspace $L_{(u,\lambda)}(\varphi, \Phi)$ defined in 4, but in general it is just an isotropic subspace. Nevertheless, if $V \in \mathcal{U}$ is finite-dimensional, then the space L corresponding to $(\varphi, \Phi)|_V$ is Lagrangian and we denote it by $L_{(u,\lambda)}(\varphi, \Phi)|_V$.

The set of all finite-dimensional subspaces has a partial ordering given by inclusion. Moreover, it is a directed set, therefore, we can take a generalized limit over the sequence of nested subspaces. The existence of this limit is guaranteed by the following

Theorem 4 ([2]). *The limit*

$$\mathfrak{L}_{(u,\lambda)}(\varphi, \Phi) = \lim_{V \nearrow U} L_{(u,\lambda)}(\varphi, \Phi)|_V$$

exists if and only if $\text{ind}^- Q|_{\Phi'_u} < \infty$.

If the limit exists, we call it the *L-derivative* and denote it by the gothic symbol \mathfrak{L} to distinguish it from the isotropic subspace that we would have got otherwise. The *L-derivative* constructed over some finite-dimensional subspace of the source space will be called a *L-prederivative*. In what follows, we also omit for brevity (u, λ) in the notations. The following property allows to find efficient ways to compute $\mathfrak{L}(\varphi, \Phi)$.

Theorem 5. *Suppose that U is a topological vector space such that φ, Φ are continuous on U and $U_0 \subset U$ is a dense subspace. Then*

$$\mathfrak{L}_{(u,\lambda)}(\varphi, \Phi)|_{U_0} = \mathfrak{L}_{(u,\lambda)}(\varphi, \Phi)|_U.$$

One can use this theorem in two different directions. Given a topology on U_0 , one can look for a weaker topology on U_0 such that φ and Φ are continuous in that topology. Then we extend U_0 to U by completion. This trick was previously used in [2].

Another way is to take a smaller subspace. For example, if U is separable, then we can take a dense countable subset e_1, e_2, \dots and compute the limit as

$$\lim_{n \rightarrow \infty} L(\varphi, \Phi)|_{\text{span}\{e_1, \dots, e_n\}} = \mathfrak{L}(\varphi, \Phi).$$

This allows us to see how the L -derivative changes as we add variations. Under some additional assumptions we can also compute the change in the Maslov index after adding additional subspaces (see the Appendix).

8. MONOTONICITY

To discuss the Morse theorem in the general setting, we need a very useful notion of monotonicity. Assume that the curve γ_t is contained in a chart $\Delta^{\mathfrak{h}}$, then one can associate a one-parametric family of quadratic forms S_t , where $\gamma_t = \{(p, S_t p) : p \in \mathbb{R}^n\}$, $\Delta = \{(0, q) : q \in \mathbb{R}^n\}$. We say that γ_t is *increasing* if S_t is increasing, i. e., $S_t - S_\tau$ is positive definite when $t > \tau$. It is important that the property of a smooth curve to be increasing does not depend on the choice of a coordinate chart. Indeed, the quadratic form $p \mapsto \langle \dot{S}_t p, p \rangle$ is equivalent by a linear change of variables to the form $\dot{\gamma}_t$ defined in Section 5, and the definition of $\dot{\gamma}_t$ is intrinsic. It does not use local coordinates.

Moreover, if γ_t , $t_0 \leq t \leq t_1$ is simple and increasing, then the Maslov index $\gamma \circ \mathcal{M}_\Lambda$ depends only on $\gamma_{t_0}, \gamma_{t_1}, \Lambda$ and can be explicitly expressed via the Maslov index of this triple. More precisely, assume that $\gamma_{t_0}, \gamma_{t_1}, \Lambda$ are mutually transversal and let \tilde{q} be a quadratic form on λ defined by the formula: $\tilde{q}(\lambda) = \sigma(\lambda_1, \lambda_0)$, $\lambda \in \Lambda$, where $\lambda_0 \in \gamma_{t_0}, \lambda_1 \in \gamma_{t_1}, \lambda = \lambda_0 + \lambda_1$. Actually, if we define

$$\text{Ind}_\Lambda(\gamma_{t_0}, \gamma_{t_1}) = \text{ind}^- \tilde{q},$$

then one can show [1] that

$$\gamma \circ \mathcal{M}_\Lambda = \text{Ind}_\Lambda(\gamma_{t_0}, \gamma_{t_1}).$$

A corollary of this fact is the following triangle inequality:

Proposition 7. *Let $\Pi, \Lambda_i, i = 0, 1, 2$ be Lagrangian subspaces in $L(\Sigma)$. Then*

$$\text{Ind}_\Pi(\Lambda_0, \Lambda_2) \leq \text{Ind}_\Pi(\Lambda_0, \Lambda_1) + \text{Ind}_\Pi(\Lambda_1, \Lambda_2).$$

Proof. We consecutively connect Λ_0 with Λ_1 , Λ_2 with Λ_3 , and Λ_3 with Λ_0 by simple monotone curves such that we get a closed curve γ . Then we have

$$\gamma \circ \mathcal{M}_\Pi = \text{Ind}_\Pi(\Lambda_0, \Lambda_1) + \text{Ind}_\Pi(\Lambda_1, \Lambda_2) + \text{Ind}_\Pi(\Lambda_2, \Lambda_0).$$

From the definition of Ind , one has

$$\text{Ind}_\Pi(\Lambda_2, \Lambda_0) = n - \text{Ind}_\Pi(\Lambda_0, \Lambda_2).$$

So it is enough to show that $\gamma \circ \mathcal{M}_\Pi \geq n$. This again follows from the fact that the intersection index of a closed curve does not depend on the choice of Π . Recall that the group of symplectic transformations acts transitively on the set of pairs of transversal Lagrangian planes. Hence we can find $\Delta \in L(\Sigma)$ such that Λ_0 and Λ_1 belong to the coordinate chart $\Delta^{\mathfrak{h}}$ and, moreover, Λ_0 is represented by a negative definite symmetric matrix in this chart, while Λ_1 is represented by a positive definite symmetric matrix. Then

$$\gamma \circ \mathcal{M}_\Pi = \gamma \circ \mathcal{M}_\Delta = \text{Ind}_\Delta(\Lambda_0, \Lambda_1) + \text{Ind}_\Delta(\Lambda_1, \Lambda_2) + \text{Ind}_\Delta(\Lambda_2, \Lambda_0) \geq n$$

since by definition $\text{Ind}_\Delta(\Lambda_0, \Lambda_1) = n$ and $\text{Ind}_\Delta(\Lambda_i, \Lambda_j) \geq 0$. □

So if we take a curve $\gamma(t) \in L(\Sigma)$ and its subdivision at times $0 = t_0 < t_1 < \dots < t_N = 1$, we can consider the sum

$$\sum_{i=0}^{N-1} \text{Ind}_{\Pi} (\gamma(t_i), \gamma(t_{i+1}))$$

which grows as the partition gets finer and finer. For monotone increasing curves this sum will stabilize and be equal to a finite number. This motivates the next definition. A (perhaps discontinuous) curve $\gamma : [0, 1] \rightarrow L(\Sigma)$ is called *monotone increasing* if

$$\sup_D \sum_{t_i \in D} \text{Ind}_{\Pi} (\gamma(t_i), \gamma(t_{i+1})) < \infty,$$

where the supremum is taken over all possible finite partitions D of the interval $[0, 1]$.

Monotone curves have properties similar to monotone functions. For example, they have only jump discontinuities and are almost everywhere differentiable.

It is instructive to see how monotonicity works on curves in $L(\mathbb{R}^2)$. The Lagrangian Grassmanian $L(\mathbb{R}^2)$ topologically is just an oriented circle and a curve is monotone increasing if it runs in the counterclockwise direction. A coordinate chart is the circle with a removed point. On the left picture of Fig. 5 the blue curve is monotone and its Maslov index is equal to zero. On the right the black curve is not monotone increasing. Indeed, if we take any two points, the Maslov index of a triple $(\Lambda_{t_i}, \Pi, \Lambda_{t_{i+1}})$ will be equal to one (the red curve).




Fig. 5. Monotone increasing blue curve and monotone decreasing black curve with the same end-points.

9. L-DERIVATIVE FOR OPTIMAL CONTROL PROBLEMS WITH FREE CONTROLS

In this subsection we consider an optimal control problem with fixed end-points¹⁾.

$$\dot{q} = f(q, u(t)), \quad q \in M, \quad u \in \mathbb{R}^k, \tag{9.1}$$

$$q(0) = q_0, \quad q(t_1) = q_1,$$

$$J_0^{t_1} = \int_0^{t_1} l(q(s), u(s)) ds.$$

We may assume that time t_1 is fixed. Otherwise, we can make a time scaling $s = \alpha\tau$, where $\alpha > 0$ is a constant that will be used as an additional control function. If we denote $\hat{u}(\tau) = u(\alpha\tau)$ and $\tau_1 = t_1/\alpha$, then we get an equivalent optimal control problem

$$\dot{q} = \alpha f(q, \hat{u}(\tau)), \quad (\alpha, \hat{u}(\tau)) \in \mathbb{R} \times L_k^2[0, \tau_1],$$

$$q(0) = q_0, \quad q(\tau_1) = q_1,$$

¹⁾We do not study free endpoint problems here: see the recent paper [7] and references therein for 2nd order optimality conditions in the free endpoint case.

$$\hat{J}_0^{\tau_1} = \alpha \int_0^{\tau_1} l(q(\tau), \hat{u}(\tau)) d\tau$$

with fixed time τ_1 . We see that by varying α we vary the final time t_1 , so we can assume from the beginning that t_1 is fixed.

A curve $q(t)$ that satisfies (9.1) for some locally bounded measurable function $u(t)$ is called an *admissible trajectory*. As in the case of the classical calculus of variations, we can define the evaluation map F_t that takes an admissible curve and maps it to the corresponding point at time t . We easily recover the Lagrange multiplier rule (1.1)

$$\nu d_\gamma J_0^t = \lambda_t D_\gamma F_t - \lambda_0 D_\gamma F_0, \tag{9.2}$$

where ν can be normalized in such a way that it takes value 1 or 0. If $\nu = 1$, we call the corresponding extremal *normal*, otherwise it is called *abnormal*. In the case of calculus of variations one has only normal extremals because (F_t, F_0) is a submersion.

We can derive the Hamiltonian system also in this case. Locally we assume that $\lambda_t = (p_t, q_t)$. Then

$$\nu \int_0^t \left(\frac{\partial l}{\partial q} dq_s + \frac{\partial l}{\partial u} du_s \right) ds = p_t dq_t - p_0 dq_0.$$

We differentiate this expression w.r.t. time t :

$$\nu \frac{\partial l}{\partial q} dq_t + \nu \frac{\partial l}{\partial u} du_t = \dot{p}_t dq_t + p_t df,$$

where $df = \frac{\partial f}{\partial q} dq_t + \frac{\partial f}{\partial u} du_t$. Now we collect terms and obtain

$$\begin{aligned} \frac{\partial}{\partial u} (\langle p_t, f(q_t, u) \rangle - \nu l(q_t, u)) &= 0, \\ \dot{p}_t &= \frac{\partial}{\partial q_t} (\nu l(q_t, u) - \langle p_t, f(q_t, u) \rangle), \\ \dot{q} &= f(q, u). \end{aligned}$$

Thus, if we set

$$H(p, q, u) = \langle p, f(q, u) \rangle - \nu l(q, u),$$

we see that the equations above are equivalent to a Hamiltonian system, where the Hamiltonian satisfies

$$\frac{\partial H}{\partial u} = 0.$$

We can rewrite these equations in the coordinate-free form:

$$\dot{\lambda}_t = \vec{H}(\lambda_t, u), \quad \frac{\partial H}{\partial u} = 0, \tag{9.3}$$

where \vec{H} is the Hamiltonian vector field with the Hamiltonian H .

Now fix $q_0 \in M$; then $F_0^{-1}(q_0)$ is just the space of control functions $u(\cdot)$. Let $E_t : u(\cdot) \mapsto q(t)$ be the *endpoint map*, $E_t = F_t|_{F_0^{-1}(q_0)}$. It is easy to see that the relation (9.2) is equivalent to the relation

$$\nu d_u J_0^t = \lambda_t D_u E_t.$$

Let $\tilde{u}(\cdot)$ satisfy this relation, i. e., let there exist λ_t , $0 \leq t \leq t_1$ such that (9.3) is satisfied with $u = \tilde{u}(t)$. We are going to compute the L -derivative $\mathfrak{L}_{(\tilde{u}, \lambda_t)}(\nu J_0^t, E_t)$. Let $P_t : M \rightarrow M$ be the flow generated by the differential equation $\dot{q} = f(q, \tilde{u}(t))$, $P_0 = Id$. We set $G_t = P_t^{-1} \circ E_t$. The map G_t is obtained from E_t by a “time-dependent change of variables” in M . The intrinsic nature of the L -derivative now implies that

$$\mathfrak{L}_{(\tilde{u}, \lambda_t)}(\nu J_0^t, G_t) = P_t^* (\mathfrak{L}_{(\tilde{u}, \lambda_0)}(\nu J_0^t, E_t))$$

and we may focus on the computation of the L -derivatives $\mathfrak{L}_{(\tilde{u}, \lambda_0)}(\nu J_0^t, G_t) \subset T_{\lambda_0}(T^*M)$, $0 \leq t \leq t_1$, which is more convenient, since all the L -derivatives lie in the same symplectic space $T_{\lambda_0}(T^*M)$ and this way we don't need a connection or a homotopy argument to compute the Maslov index.

To find G_t we apply a time-dependent change of variables

$$y = P_t^{-1}(q).$$

In this way we get an equivalent control system

$$\dot{y} = (P_t^{-1})_*(f(u, \cdot) - f(\tilde{u}(t), \cdot))(y) = g(t, u, y),$$

where in the center we have the pull-back of vector fields. Then the definitions give us

$$y_0 = q_0, \quad g(\tilde{u}(t), y) = 0.$$

Similarly, in the functional we define $\psi(t, u, y) = l(P_t(y), u)$.

To write down an explicit expression for $\mathfrak{L}_{(\tilde{u}, \lambda_0)}(\nu J_0^t, G_t)$, we must characterize first all the critical points and Lagrange multipliers in these new coordinates. We apply the Lagrange multiplier rule exactly as above and find that if $(y(t), u(t))$ is a critical point, then there exists a curve of covectors $\mu_t \in T_{y(t)}^*M$ that satisfies a Hamiltonian system

$$\dot{\mu}_t = \vec{h}(t, u, \mu_t), \tag{9.4}$$

where the Hamiltonian

$$h(t, u, \mu_t) = \langle \mu_t, g(t, u, y) \rangle - \nu \psi(t, u, y)$$

satisfies also

$$\frac{\partial h(t, u, \mu_t)}{\partial u} = 0. \tag{9.5}$$

Note that, for the referenced critical point $(y_0, \tilde{u}(t))$, the corresponding curve of covectors is simply $\mu_t = \lambda_0$ and $h(t, \tilde{u}(t), \lambda_0) = 0$.

Recall that the L -derivative is obtained by linearizing the relation for the Lagrange multiplier rule. But the Lagrange multiplier rule is equivalent to this weak version of the Pontryagin maximum principle, so it is enough to linearize (9.4) and (9.5) at $\mu_t = \lambda_0$ and $u(t) = \tilde{u}(t)$. Since $\vec{h}(t, \tilde{u}, \lambda_0) = 0$, the linearization of (9.4) gives

$$\dot{\eta}_t = \left. \frac{\partial \vec{h}(t, u, \lambda_0)}{\partial u} \right|_{u=\tilde{u}} v_\tau \iff \eta_t = \eta_0 + \int_0^t X_\tau v_\tau d\tau, \quad X_\tau := \left. \frac{\partial \vec{h}(t, u, \lambda_0)}{\partial u} \right|_{u=\tilde{u}}.$$

Similarly, we can linearize Eq. (9.5) and use the definition of a Hamiltonian vector field to derive:

$$\left. \frac{\partial^2 h(t, u, \lambda_0)}{\partial u^2} \right|_{u=\tilde{u}} (v(t), \cdot) + \left\langle d_{\mu_t} \left. \frac{\partial h(t, u, \lambda_0)}{\partial u} \right|_{u=\tilde{u}}, \eta_t \right\rangle = 0 \iff b_t(v(t), \cdot) + \sigma(\eta_t, X_t \cdot) = 0,$$

where b_t is the second derivative of $h(t, u, \lambda_0)$ w.r.t. u at $u = \tilde{u}$. Combining these two expressions, we find that a L -prederivative is defined as

$$L(\nu J_0^t, E_t)|_V = (P_0^t)_* \left\{ \eta_t = \eta_0 + \int_0^t X_\tau v_\tau d\tau : \eta_0 \in T_{\lambda_0}(T_{q_0}^*M), v_\tau \in V : \int_0^t \sigma(\eta_\tau, X_\tau w_\tau) + b_\tau(v_\tau, w_\tau) d\tau = 0, \forall w_\tau \in V \right\}.$$

From this description and the definition of the L -derivative as a limit of L -prederivatives we can deduce an interesting property that can be successfully used to construct approximations for the former. This property says that if we know a L -derivative at time $s_1 \in (0, t_1)$, which we denote by \mathfrak{L}_1 , then the L -derivative at time $s_1 < s_2 \leq t_1$ can be computed using the vectors from \mathfrak{L}_1 and variations with support in $[s_1, s_2]$. A precise statement is the following

Lemma 4. Take $0 < s_1 < s_2$ and suppose that $\text{ind}^- Q|_{\ker D_{\tilde{u}}E_t}$ is finite along an extremal curve defined on $[0, s_2]$. Let \mathfrak{L}_1 and \mathfrak{L}_2 be the two L -derivatives for times s_1 and s_2 , respectively. We denote by V_2 some finite-dimensional subspace of $L^2_k[s_1, s_2]$ and consider the following equation:

$$\int_{s_1}^{s_2} \left[\sigma \left(\lambda + \int_{s_1}^\tau X_\theta v_2(\theta) d\theta, X_\tau w(\tau) \right) + b_\tau(v_2(\tau), w(\tau)) \right] d\tau = 0, \quad \forall w(\tau) \in V_2, \tag{9.6}$$

where $v_2(\tau) \in V_2$ and $\lambda \in \mathfrak{L}_1$. Then \mathfrak{L}_2 is a generalized limit of the Lagrangian subspaces $L^2_1[V_2]$ defined as

$$L^2_1[V_2] = \left\{ \lambda + \int_{s_1}^{s_2} X_\tau v(\tau) d\tau : \lambda \in \mathfrak{L}_1, v(\tau) \in V_2 \text{ satisfies (9.6) for any } w(\tau) \in V_2 \right\}.$$

This lemma implies that the curve of L -derivatives has some sort of flow property, i.e., the L -derivative at the current instant of time can be recovered from the L -derivatives at previous instants. This observation is the key point in our algorithm for computation of the L -derivative with arbitrary good precision.

The algorithm can be summarized in the following steps:

1. Take a partition $0 = s_0 < s_1 < \dots < s_N = t$ of the interval $[0, t]$. The finer the partition is, the better will be the approximation of the L -derivative at time t ;
2. Compute inductively $L_{(u,\lambda)}(\nu J^t_0, G_t)|_{V_i}$, $V_i = \mathbb{R}^k \chi_{[s_0, s_1]} \oplus \dots \oplus \mathbb{R}^k \chi_{[s_i, s_{i+1}]}$ starting from $L_{(u,\lambda)}(\nu J^t_0, G_t)|_{V_0} = \Pi$, where $\chi_{[s_i, s_{i+1}]}$ are characteristic functions.

When $\max |s_{i+1} - s_i| \rightarrow 0$, we get in the limit the real L -derivative, since piecewise constant functions are dense in L^2 . In this way we reduce the problem to solving iteratively systems of k linear equations. In Theorem 8 of the Appendix an explicit solution to this system is given.

This algorithm allows one not only to approximate the L -derivative, but also to compute the index of the Hessian restricted to the subspace of piece-wise constant variations.

Theorem 6. Let $D = \{0 = s_0 < s_1 < \dots < s_N = t_1\}$ be a partition of the interval $[0, t_1]$ and let V_D be the space of piece-wise constant functions with jumps at times s_i . We denote by $V_i \subset V_D$ the subspace of functions that are zero for $t > s_i$, $V_i^0 = V_i \cap \ker d_{\tilde{u}}G_{t_1}$ and $\Lambda_i = L(\nu J^{t_1}_0, G_{t_1})|_{V_i}$. Then the following formula holds:

$$\text{ind}^- Q|_{V_D^0} = \sum_{i=0}^N \text{Ind}_\Pi(\Lambda_i, \Lambda_{i+1}) + \dim \left(\bigcap_{i=0}^N \Lambda_i \right) - n, \tag{9.7}$$

where $\Lambda_0 = \Lambda_{N+1} = \Pi$.

Moreover, one can prove the following result, which is the basis of the whole theory.

Theorem 7. Suppose that (\tilde{q}, \tilde{u}) is an extremal of the problem (9.1) such that the index of the corresponding Hessian is finite and $\mathfrak{L}_t = \mathfrak{L}_{(\tilde{u}, \lambda_0)}(\nu J^t_0, G_t)$, $\mathfrak{L}_t \in L(T_{\lambda_0}(T^*M))$, $t \in [0, t_1]$ is the family of L -derivatives associated to it; then $t \mapsto \mathfrak{L}_t$ is a monotone curve and

$$\text{ind}^- \text{Hess}_{\tilde{u}} J^{t_1}_0|_{E^{-1}(q_0)} \geq \sup_D \sum_{s_i \in D} \text{Ind}_\Pi(\mathfrak{L}_{s_i}, \mathfrak{L}_{s_{i+1}}) + \text{Ind}_\Pi(\mathfrak{L}_{t_1}, \mathfrak{L}_0) + \dim \left(\bigcap_{t=0}^{t_1} \mathfrak{L}_t \right) - n,$$

where the supremum is taken over all possible finite partitions D of the interval $[0, t_1]$.

10. *L*-DERIVATIVE WITH CONSTRAINTS ON THE CONTROLS

In our construction of the *L*-derivative we have heavily used the fact that all variations are two-sided, but in optimal control theory this is often not the case. The control parameters may take values in some closed set *U*. Then on the boundary ∂U we can only vary along smooth directions of ∂U . To cover also this kind of situations, we are going to use the change of variables introduced in the previous section.

We consider the optimal control problem (9.1), but now we assume that $u \in U \subset \mathbb{R}^k$, where *U* is a union of a locally finite number of smooth submanifolds U_i without boundaries. In particular, any semianalytic set is available. A typical situation is when the constraints are given by a number of smooth inequalities

$$p_i(u) \leq 0$$

that satisfy

$$p_i(u) = 0 \implies d_u p_i \neq 0.$$

For example, *U* can be a ball or a polytope. In the latter case U_i consists of the interior of polytope and faces of different dimensions.

Recall that in the last subsection we used a time scaling to reduce a free time problem to a fixed one. It is actually very useful to use general time reparameterizations as possible variations, even in the fixed time case. Assume that $t(\tau)$ is an increasing absolutely continuous function such that $t(0) = 0$ and if t_1 is fixed also $\tau(t_1) = \tau_1$. Actually, instead of the last condition, one can simply take the time variable as a new variable satisfying

$$\dot{t} = 1.$$

Then we can consider an optimal control problem

$$\dot{q} = f(q, u(\tau^{-1}(t))), \quad \int_0^{t_1} l(q(s), u(\tau^{-1}(s))) ds \rightarrow \min, \tag{10.1}$$

which is essentially the optimal control problem (9.1) written in a slightly different way. Since $t(\tau)$ is absolutely continuous, it is of the form

$$t(\tau) = \int_0^\tau \alpha(s) ds.$$

We rewrite (10.1) in the new time τ to get

$$\frac{dq}{d\tau} = \alpha(\tau) f(q, u(\tau)), \quad \int_0^{\tau_1} \alpha(\tau) l(q(\tau), u(\tau)) d\tau \rightarrow \min.$$

Variations with respect to $\alpha(t)$ are called time variations. Since $\alpha > 0$, time variations are always two-sided and thus one can include them to study the index of the Hessian via *L*-derivatives. They have been already used to derive necessary and sufficient optimality conditions in the bang-bang case, where no two-sided variations are available if we just vary *u* (see [4, 5]).

Time variations do not give any new contribution to the Morse index of the Hessian if the extremal control $\tilde{u}(t)$ is C^2 . Indeed, assume, for example, that λ_t is an abnormal extremal corresponding to $\tilde{u}(t)$. Let us denote for simplicity $\beta = 1/\alpha$, i. e.,

$$\tau^{-1}(t) = \int_0^t \beta(s) ds,$$

so that we don't have to include differentials of inverse functions in the expressions.

We consider the end-point map $E_t(\tilde{u}(\beta(s)))$ of (10.1) and calculate the Hessian with respect to β at a point $\beta(s) \equiv 1$. We obtain

$$\lambda_t d_\beta^2 E_t(\gamma_1, \gamma_2) = \lambda_t d_{\tilde{u}}^2 E_t \left(\frac{d\tilde{u}}{dt} \int_0^t \gamma_1(s) ds, \frac{d\tilde{u}}{dt} \int_0^t \gamma_2(s) ds \right) + \lambda_t d_{\tilde{u}} E_t \left(\frac{d^2 \tilde{u}}{dt^2} \int_0^t \gamma_1(s) ds \int_0^t \gamma_2(s) ds \right),$$

but the the second term is zero since \tilde{u} is extremal and therefore $\lambda_t d_{\tilde{u}} E_t = 0$. Thus, we see that all the time variations in the Hessian could have been realized by variations of u .

If $\tilde{u}(t)$ has less regularity, then the time variations become nontrivial. For example, in the bang-bang case \tilde{u} is piece-wise constant and the effect of the time variations concentrates at the points of discontinuity of \tilde{u} . This allows one to reduce an infinite-dimensional optimization problem to a finite-dimensional one. This finite-dimensional space of variations corresponds simply to variations of the switching times.

If we include the time variations, we will have enough two-sided variations to cover all the known cases. It only remains to construct the L -derivative over the space of all available two-sided variations. Note that after adding the time variations, this space is not empty. It can be a very difficult computation, but using our algorithm, we can always construct an approximation and obtain a bound on the Morse index.

To apply our algorithm, we must define “constant” variations. The set U is a union of smooth submanifolds $U_i \in \mathbb{R}^{k+1}$ without boundaries. Since each U_i is embedded in \mathbb{R}^{k+1} by assumption, we can take the orthogonal projections $\pi_u^i : \mathbb{R}^{k+1} \rightarrow T_u U_i$ whenever $u \in U_i$. Then we can define a projection of a general variation $v_t \in L_{k+1}^2[0, t]$ to the subspace of two-sided variations as

$$\pi_\tau v_\tau = \sum_{i=0}^n \chi_{U_i}(\tau) \pi_{\tilde{u}(\tau)}^i v_\tau,$$

where $\chi_{U_i}(\tau)$ are the characteristic functions. “Constant” variations for the constrained problem are just projections $\pi_\tau v$ of the constant sections $v \in \mathbb{R}^{k+1}$. Equivalently one can consider directly constant variations in \mathbb{R}^{k+1} and simply replace X_τ in the definition of the L -derivative by $X_\tau \pi_\tau$. Then the algorithm from the previous subsection is applicable without any further modifications.

An important remark is that the orthogonal projections depend on the metric that we choose on \mathbb{R}^{k+1} . This choice would indeed give us different L -prederivatives, since the “constant” variations would be different, but in the limit the L -derivative will be the same, because at the end we just approximate the same space in two different ways.

APPENDIX. INCREMENT OF THE INDEX

Recall that in Theorem 6 we have stated that by adding piece-wise constant variations we can track how the Maslov index of the corresponding Jacobi curve changes. But there is no use in this theorem if we are not able to construct explicitly the corresponding L -prederivatives from our algorithm. To do this, we can use the following theorem.

Theorem 8. *Suppose that we know $\mathfrak{L}(\nu J_0^t, G_t)|_V$, where V is some space of variations defined on $[0, t]$. We identify $\mathfrak{L}(\nu J_0^t, G_t)|_V$ with \mathbb{R}^n and the space of control parameters with \mathbb{R}^k , and put an arbitrary Euclidean metric on both of them. Let E be the space of all $v \in \mathbb{R}^k$ for which*

$$\sigma \left(\eta, \frac{1}{\varepsilon} \int_t^{t+\varepsilon} X_\tau d\tau \cdot v \right) = 0, \quad \forall \eta \in L(\nu J_0^t, G_t)|_V$$

and let $L = \mathfrak{L}(\nu J_0^t, G_t)|_V \cap \mathfrak{L}(\nu J_0^t, G_t)|_{\tilde{V}}$, where $\tilde{V} = V \oplus \mathbb{R}^k \chi_{[t, t+\varepsilon]}$. We define the two bilinear maps $A : \mathfrak{L}(\nu J_0^t, G_t)|_V \times E^\perp \rightarrow \mathbb{R}$, $Q : E^\perp \times E^\perp \rightarrow \mathbb{R}$:

$$A : (\eta, w) \mapsto \sigma \left(\eta, \frac{1}{\varepsilon} \int_t^{t+\varepsilon} X_\tau d\tau \cdot w \right),$$

$$Q : (v, w) \mapsto \frac{1}{\varepsilon} \int_t^{t+\varepsilon} \sigma \left(\int_t^\tau X_\theta d\theta \cdot v, X_\tau w \right) + b_\tau(v, w) d\tau,$$

and we use the same symbols for the corresponding matrices.

Then the new L -prederivative $\mathfrak{L}_{(u, \lambda)}(\nu J_0^t, G_t)|_{\tilde{V}}$ is a span of vectors from the subspace L and vectors

$$\eta_i + \frac{1}{\varepsilon} \int_t^{t+\varepsilon} X_\tau d\tau \cdot v_i,$$

where v_i is an arbitrary basis of E^\perp and η_i are defined as

$$\eta_i = -A^+ Q e_i$$

with A^+ being the Penrose – Moore pseudoinverse.

Although we use some additional structures in the formulation like a Euclidean metric, it will only give a different basis for $\mathcal{L}(\nu J_0^t, G_t)|_{\tilde{V}}$, but the L -derivative itself will be the same.

In general, if we would like to compute the difference between indices of the Hessian $\text{Hess}_{\tilde{u}} \varphi|_{\Phi^{-1}(q_0)}$ restricted to two finite-dimensional subspaces $U_1 \subset U$, the Maslov index will only give us a lower bound:

$$\text{ind}^- Q|_{U^0} - \text{ind}^- Q|_{U_1^0} \geq \text{Ind}_\Pi (L(\varphi, \Phi)|_{U_1}, L(\varphi, \Phi)|_U),$$

where as before $U_1^0 = U_1 \cap \ker d_{\tilde{u}} \Phi$ and the same for U^0 . This formula was proved in [2].

One can ask when this inequality becomes an equality. It seems that there is no general if and only if condition, but one can find some nice situations where the equality holds, as in the piece-wise constant case. Another condition that is quite general is stated in the following theorem.

Theorem 9. *Assume that the index of the Hessian at a point (u, λ) is finite and that we can find a splitting $U_1 \oplus U_2$ of a possibly infinite-dimensional U such that*

1. U_1 and U_2 are orthogonal with respect to Q ;
2. $Q|_{U_2^0} > 0$, where $U_2^0 = U_2 \cap \ker \Phi'_u$;
3. $\dim \mathcal{L}_{(u, \lambda)}(\varphi, \Phi)|_{U_1} \cap \Pi = 0$.

Then

$$\text{ind}^- Q|_{U^0} - \text{ind}^- Q|_{U_1^0} = \text{Ind}_\Pi (\mathcal{L}(\varphi, \Phi)|_{U_1}, \mathcal{L}(\varphi, \Phi)|_U).$$

We are going to apply Lemma 3 twice: first, to the subspace U_1^0 in U^0 and, second, to U_2^0 in $(U_1^0)^\perp$. Assume for now that $\dim U < \infty$. First we clarify what are all the subspaces presented in the formula. We have

$$(U_1^0)^\perp = \{v = v_1 + v_2 \in U : f'_u(v_1 + v_2) = 0, \langle Qv_1, U_1^0 \rangle = 0\}.$$

Here we have used our orthogonality assumptions. We claim that $(U_1^0)^\perp$ is actually equal to the subspace

$$\left\{ v_1 + v_2 \in U : f'_u(v_1 + v_2) = 0, \exists \xi_1 \in T_{f(u)}^* M, \langle Qv_1 + \xi_1 f'_u, U_1 \rangle = 0 \right\}.$$

It is clear that the second space is a subspace of $(U_1^0)^\perp$. We want to prove the converse statement. Assume that $v \in (U_1^0)^\perp$. Firstly, we put any Euclidean metric on $T_{f(u)} M$ and use it to define an isomorphism between $T_{f(u)}^* M$ and $T_{f(u)} M$. Secondly, we choose a subspace $E \subset U_1$ complementary to U_1^0 and a basis e_i of E such that $f'_u(e_i)$ form an orthogonal subset. Then the covector ξ_1 that we need is simply given by

$$\xi_1 = - \sum_{i=1}^{\dim E} \frac{f'_u(e_i)}{|f'_u(e_i)|^2} \langle Qv_1, e_i \rangle.$$

Thus, the claim has been proved.

From the orthogonality assumption it follows that $U_2^0 \in (U_1^0)^\perp$. The orthogonal complement of U_2^0 in $(U_1^0)^\perp$ is equal to $(U_1^0 + U_2^0)^\perp$, which is equal to

$$(U_1^0 + U_2^0)^\perp = \left\{ v_1 + v_2 \in U : f'_u(v_1 + v_2) = 0, \exists \xi_1 \in T_{f(u)}^* M, \langle Qv_1 + \xi_1 f'_u, U_1 \rangle = 0, \langle Qv_2, U_2^0 \rangle = 0 \right\}.$$

Similarly to the above, this is equivalent to

$$(U_1^0 + U_2^0)^\perp = \left\{ v_1 + v_2 \in U : f'_u(v_1 + v_2) = 0, \exists \xi_1, \xi_2 \in T_{f(u)}^*M, \langle Qv_1 + \xi_1 f'_u, U_1 \rangle = 0, \right. \\ \left. \langle Qv_2 - \xi_2 f'_u, U_2 \rangle = 0 \right\}.$$

We can now compute the quadratic form Q restricted to $(U_1^0 + U_2^0)^\perp$. Again we use the orthogonality assumption and the equivalent definition of $(U_1^0 + U_2^0)^\perp$ above. Assume that $v = v_1 + v_2 \in (U_1^0 + U_2^0)^\perp$ and $\xi = \xi_1 + \xi_2$. Then

$$\langle Q(v_1 + v_2), v_1 + v_2 \rangle = \langle Qv_1, v_1 \rangle + \langle Qv_2, v_2 \rangle = -\xi_1 f'_u v_1 + \xi_2 f'_u v_2 = -\xi f'_u v_1.$$

Now we would like to write down the expression for the matrix S from the definition of the Maslov index. First we write down the definition of the two L -derivatives:

$$L(\varphi, \Phi)|_{U_1} = \{(\eta_1, f'_u v_1^1) : \langle Qv_1^1 + \eta_1 f, U_1 \rangle = 0\};$$

$$L(\varphi, \Phi)|_U = \{(\eta_2, f'_u(v_1^2 + v_2)) : \langle Qv_1^2 + \eta_2 f, U_1 \rangle = 0, \langle Qv_2 + \eta_2 f, U_2 \rangle = 0\}.$$

The quadratic form from the Maslov index is defined on $(L(\varphi, \Phi)|_{U_1} + L(\varphi, \Phi)|_U) \cap \Pi$. We write $v_1^1 + v_1^2 = v_1$, $\xi_1 = \eta_1 + \eta_2$, $\xi_2 = -\eta_2$ and suppose that $f'_u(v_1) + f'_u(v_2) = 0$. Then for the quadratic form \tilde{q} we have

$$\tilde{q} = \sigma((\eta_1, f'_u v_1^1), (\eta_2, f'_u(v_1^2 + v_2))) = \sigma((\eta_1, f'_u v_1^1), (0, f'_u v_2)) = \eta_1 f'_u v_2 = -\eta_1 f'_u v_1 = -\xi f'_u v_1.$$

In the second equality we have used the fact that $(\eta_1, f'_u v_1^1)$ and $(\eta_2, f'_u v_1^2)$ belong to $L(\varphi, \Phi)|_{U_1}$ by definition.

We see that this gives the same expression as for $Q|_{(U_2^0)^\perp}$. But, moreover, both quadratic forms are actually defined on the same space. Indeed, we have

$$(L(\varphi, \Phi)|_{U_1} + L(\varphi, \Phi)|_U) \cap \Pi = \{(\xi_1 + \xi_2, 0) : \exists v_i \in V_i, f'_u(v_1 + v_2) = 0, (\xi_1, f'_u v_1) \in L(\varphi, \Phi)|_{U_1}, \\ (-\xi_2, f'_u v_2) \in L(\varphi, \Phi)|_{U_2}\} = (L(\varphi, \Phi)|_{U_1} + L(\varphi, \Phi)|_{U_2}) \cap \Pi.$$

But if we add to $(\xi_1 + \xi_2) \in (L(\varphi, \Phi)|_{U_1} + L(\varphi, \Phi)|_U) \cap \Pi$ the corresponding v_i and to $v_i \in U_i \cap (U_1^0 + U_2^0)^\perp$ the corresponding $\xi_1 + \xi_2$, we obtain the same space.

Now we compute the other terms from the formula in Lemma 3. We have

$$U_1^0 \cap (U_1^0)^\perp = \{v_1 \in U_1^0 : Q(v_1, U_1^0) = 0\}.$$

Similarly to the discussion at the beginning of the proof, we can show that

$$U_1^0 \cap (U_1^0)^\perp = \{v_1 \in U_1^0 : \exists \xi \in T_{f(u)}^*M, Q(v_1 + \xi f'_u, U_1) = 0\}.$$

We do now the same for $\ker Q|_{U^0} \cap U_1^0$:

$$\ker Q|_{U^0} \cap U_1^0 = \{v_1 \in U_1^0 : \langle Qv_1, U^0 \rangle = 0\} = \{v_1 \in U_1^0 : \langle Qv_1 + \xi f'_u, U \rangle = 0\}.$$

To understand the dimensions, we look carefully at the equation

$$\langle Qv_1 + \xi f'_u, U_1 \rangle = 0.$$

If there are two solutions (ξ, v_1) and (ξ, v'_1) of this equation, then by linearity $(0, v_1 - v'_1)$ is a solution as well and thus all solutions are uniquely defined by different ξ modulo $\ker Q|_{U_1} \cap U_1^0$. These ξ lie in $L(\varphi, \Phi)|_{U_1} \cap \Pi$, as can be seen from the definitions. Therefore,

$$\dim(U_1^0 \cap (U_1^0)^\perp) = \dim(L(\varphi, \Phi)|_{U_1} \cap \Pi) + \dim(\ker Q|_{U_1} \cap U_1^0).$$

Now we do the same for

$$0 = Q(v_1 + \xi f'_u, V) = Q(v_1 + \xi f'_u, U_1) + \xi f'_u U_2.$$

Again ξ are defined uniquely modulo $\ker Q|_{U_1 \cap U_1^0}$, but now they lie in $L(\varphi, \Phi)|_U \cap \Pi$. Therefore,

$$\dim(\ker Q|_{U^0 \cap U_1^0}) = \dim(L(\varphi, \Phi)|_U \cap \Pi) + \dim(\ker Q|_{U_1 \cap U_1^0}).$$

Since Q is positive on U_2^0 , we have $(U_2^0)^\perp \cap U_2^0 = \{0\}$ and so we can collect all the formulas using the fact that $(L(\varphi, \Phi)|_U \cap \Pi) \subset (L(\varphi, \Phi)|_{U_1 \cap \Pi})$:

$$\begin{aligned} \operatorname{ind}^- Q|_{U^0} - \operatorname{ind}^- Q|_{U_1^0} &= \operatorname{Ind}_\Pi(L(\varphi, \Phi)|_{U_1}, L(\varphi, \Phi)|_U) + \frac{1}{2} \dim(L(\varphi, \Phi)|_{U_1 \cap \Pi}) - \\ &\quad - \frac{1}{2} (\dim L(\varphi, \Phi)|_U \cap \Pi). \end{aligned}$$

Under the third assumption the formula is valid also in the infinite-dimensional case. We know that the L -prederivatives will converge and that the quadratic form from the Maslov-type index is continuous. The only possibly discontinuous terms are the dimensions of various intersections, but they are zero now for L -prederivatives close to the L -derivatives.

ACKNOWLEDGMENTS

The work of A. A. Agrachev was supported by the Russian Science Foundation under grant No. 17-11-01387.

REFERENCES

1. Agrachev, A. A., Quadratic Mappings in Geometric Control Theory, *J. Soviet Math.*, 1990, vol. 51, no. 6, pp. 2667–2734; see also: *Itogi Nauki Tekh. Ser. Probl. Geom.*, vol. 20, Moscow: VINITI, 1988, pp. 111–205.
2. Agrachev, A. A., Feedback-Invariant Optimal Control Theory and Differential Geometry: 2. Jacobi Curves for Singular Extremals, *J. Dynam. Control Systems*, 1998, vol. 4, no. 4, pp. 583–604.
3. Agrachev, A. A., Barilari, D., and Boscain, U., Introduction to Riemannian and Sub-Riemannian Geometry: Preprint SISSA 09/2012/M, <https://webusers.imj-prg.fr/~daveide.barilari/Notes.php> (2016).
4. Agrachev, A. A. and Gamkrelidze, R. V., Symplectic Geometry for Optimal Control, in *Nonlinear Controllability and Optimal Control*, H. J. Sussmann (Ed.), Monogr. Textbooks Pure Appl. Math., vol. 133, New York: Dekker, 1990, pp. 263–277.
5. Agrachev, A., Stefani, G., and Zezza, P., Strong Optimality of a Bang-Bang Trajectory, *SIAM J. Control Optim.*, 2002, vol. 41, no. 4, pp. 981–1014.
6. de Gosson, M., *Symplectic Geometry and Quantum Mechanics*, Oper. Theory Adv. Appl., vol. 166, Basel: Birkhäuser, 2006.
7. Frankowska, H. and Hoehener, D., Pointwise Second-Order Necessary Optimality Conditions and Second-Order Sensitivity Relations in Optimal Control, *J. Differential Equations*, 2017, vol. 262, no. 12, pp. 5735–5772.
8. McDuff, D. and Salamon, D., *Introduction to Symplectic Topology*, 2nd ed., Oxford Math. Monogr., New York: Oxford Univ. Press, 1999.