On the validity of the transversality condition for different concepts of tangent cone to a set

H. J. Sussmann

Department of Mathematics
Rutgers University
Piscataway, NJ 08854, USA
sussmann@math.rutgers.edu

THE 45TH IEEE CDC (CONFERENCE ON DECISION AND CONTROL)
San Diego, CA, December 13-15, 2006

These slides are available at http://www.math.rutgers.edu/~sussmann/slides.html

This work is partially supported by grant NSF-DMS 05-09930. The author is grateful to NSF for the support.

I. A ROUGH CLASIFICATION OF VERSIONS OF THE FDPMP (FINITE-DIMENSIONAL PONTRYAGIN MAXIMUM PRINCIPLE)

Every known version of the FDPMP is of one of the following two types:

- Type T. (The “T” stands for “topological.”)
- Type L. (The “L” stands for “limiting.”)

In the transversality condition:

- Type T versions involve some kind of Boltyanskii tangent cone to the terminal set.
- Type L versions involve the Clarke tangent cone to the terminal set, or the Mordukhovich normal cone

The proofs of Type T versions typically use a topological separation argument, based on the Brouwer fixed point theorem or some variant thereof.

All versions of the finite-dimensional Pontryagin maximum principle with high-order conditions (Knobloch, Krener, Bianchini-Stefani, Agrachev, Sarychev, Gamkrelidze, and many others) appear to be Type T.

The finite dimensionality comes in where the Brouwer fixed-point theorem is used, since that theorem depends esssentially on being in a finite-dimensional space.

The proofs of Type L versions usually produce a sequence $\{\bar{p}_k\}_{k\in\mathbb{N}}$ of "approximate terminal adjoint covectors" (using, for example, the Ekeland variational principle) and then extract a convergent (or weakly convergent) subsequence whose limit $\bar{p}_\infty$ is the terminal value of the adjoint covector.

The finite dimensionality comes in when one tries to establish that $\bar{p}_\infty \neq 0$. The $\bar{p}_k$ can be normalized so that $\|\bar{p}_k\| = 1$, and the existence of a weak*-convergent subsequence (if, say, we are working on a Hilbert space) follows from the weak*-compactness of the closed unit ball, but in infinite dimensions one cannot prove in general that $\bar{p}_\infty \neq 0$, since the unit sphere is not weak*-compact.

NATURAL QUESTIONS:

- Is it possible to unify all these versions, and their proofs, into a single general theorem?

- If so, would that theorem be Type T, Type L, or of some new type, involving techniques that somehow combine or go beyond those of the two basic types?

- In particular, is there a FDPMP with high-order conditions that would apply to a dynamical law that in some portion of the reference trajectory is only Lipschitz, and with a transversality condition involving a Clarke or Mordukhovich cone?

It turns out that the key issue is whether the following property is true:

The Transversal Intersection Property (TIP)

If two subsets S_1, S_2 of $\mathbb{R}^n$ have tangent cones C_1, C_2 at a point $p \in \mathbb{R}^n$, and the cones C_1, C_2 are strongly transversal, then $S_1 \cap S_2$ contains a sequence of points p_j converging to p and $\neq p$.

II. CONES

A. Definition

A *cone* in a real linear space X is a subset C of X which is nonempty, and closed under multiplication by nonnegative scalars. (In particular, if C is a cone then necessarily $0 \in C$.)

B. Definition

The *polar* of a cone C in a real linear normed space X is the set $C^{\perp}$ of all $w \in X^{\dagger}$ such that $\langle w, c \rangle \leq 0$ for all $c \in C$. Clearly, $C^{\perp}$ is always a closed convex cone. If X is finite-dimensional (so $X \sim X^{\dagger\dagger}$ canonically), then $C^{\perp\perp}$ is the smallest closed convex cone containing C, from which it follows in particular that $C^{\perp\perp} = C$ if and only if C is closed and convex.

REMARK: $X^{\dagger}$ is the dual of X.

C. Definition

Assume that $S \subseteq \mathbb{R}^n$ and $p \in S$. The *Bouligand tangent cone* to S at p is the set of all vectors $v \in \mathbb{R}^n$ such that there exist

(i) a sequence $\{p_j\}_{j\in\mathbb{N}}$ of points of S converging to p,

(ii) a sequence $\{h_j\}_{j\in\mathbb{N}}$ of positive real numbers converging to 0,

such that

$$v = \lim_{j\to\infty} \frac{p_j - p}{h_j}.$$

D. Notation

We use $T_p^B S$ to denote the Bouligand tangent cone to S at p. (It is then clear that $T_p^B S$ is always a closed cone.)

E. Definition

Assume that $S \subseteq \mathbb{R}^n$ and $p \in S$. A *Boltyanskii approximating cone* to S at p is a convex cone C in $\mathbb{R}^n$ having the property that there exist

(i) a nonnegative integer m,
(ii) a closed convex cone D in $\mathbb{R}^m$,
(iii) a neighborhood U of 0 in $\mathbb{R}^m$,
(iv) a continuous map $F : U \cap D \mapsto S$,
(v) a linear map $L : \mathbb{R}^m \mapsto \mathbb{R}^n$,

such that

$$F(x) = p + Lx + o(\|x\|) \quad \text{as} \quad x \to 0\,, \;\; x \in D\,,$$

and $LD = C$.

F. Definition

Assume that $S \subseteq \mathbb{R}^n$, S is closed, and $p \in S$. The *Clarke tangent cone* to S at p is the set of all vectors $v \in \mathbb{R}^n$ such that, whenever $\{p_j\}_{j \in \mathbb{N}}$ is a sequence of points of S converging to p, it follows that there exist Bouligand tangent vectors $v_j \in T^B_{p_j} S$ such that $\lim_{j \to \infty} v_j = v$.

G. Notation

We use $T^C_p S$ to denote the Clarke tangent cone to S at p. Then $T^C_p S$ is a closed convex cone.

III. TRANSVERSALITY

A. Definition

Two convex cones C_1, C_2 in $\mathbb{R}^n$ are *transversal* if

$$C_1 - C_2 = \mathbb{R}^n ,$$

i.e., if for every $x \in \mathbb{R}^n$ there exist $c_1 \in C_1$, $c_2 \in C_2$, such that $x = c_1 - c_2$.

B. Remark

This is a very natural generalization to cones of the ordinary notion of transversality of linear subspaces. For subspaces S_1, S_2, it is customary to require that $S_1 + S_2 = \mathbb{R}^n$, but it would make no difference if we required $S_1 - S_2 = \mathbb{R}^n$ instead.

C. Intuition

The basic idea of transversality is that, if two objects O_1, O_2 have first-order approximations A_1, A_2 near a point p, and A_1 and A_2 are transversal, then $O_1 \cap O_2$ looks, near p, like $A_1 \cap A_2$.

IV. NON-TRANSVERSALITY = LINEAR SEPARATION

Suppose C_1, C_2 are convex cones in $\mathbb{R}^n$. Then the following conditions are equivalent:

- C_1 and C_2 are not transversal,
- $C_1^\perp \cap (-C_2)^\perp \neq \{0\}$,
- there exists a nonzero linear functional $\bar{p} : \mathbb{R}^n \mapsto \mathbb{R}$ such that

$$\langle \bar{p}, c_1 \rangle \leq 0 \quad \text{for all} \quad c_1 \in C_1 \,,$$

and

$$\langle \bar{p}, c_2 \rangle \geq 0 \quad \text{for all} \quad c_2 \in C_2 \,.$$

V. STRONG TRANSVERSALITY

A. Definition

Two convex cones C_1, C_2 in $\mathbb{R}^n$ are *strongly transversal* if they are transversal and in addition $\boxed{C_1 \cap C_2 \neq \{0\}}$.

B. Intuition:

If two sets S_1, S_2 have first-order approximations C_1, C_2 near a point p, and the cones C_1, C_2 are strongly transversal, it should follow that $S_1 \cap S_2$ contains points p_j converging to p and $\neq p$.

Reason:

Near p, $S_1 \cap S_2$ should look like $C_1 \cap C_2$, because C_1 and C_2 are transversal.

Since $C_1 \cap C_2$ contains a full half-line through 0, $S_1 \cap S_2$ should also contains a nontrivial curve through p.

C. An important caveat:

The above intuition is, of course, not a proof, and when one does things carefully, it turns out that, for very reasonable notions of "first-order approximation," all one can prove is that $S_1 \cap S_2$ must contain a nontrivial connected set through p, but this set could fail to be path-connected. And for other reasonable notions one can prove even less. (For example, that $S_1 \cap S_2$ contains a sequence of points $p_j \neq p$ that converges to p.)

The following lemma says that transversality and strong transversality are almost equivalent.

More precisely, the only gap between the two conditions occurs when the cones C_1 and C_2 are linear subspaces such that $C_1 \oplus C_2 = \mathbb{R}^n$, in which case C_1 and C_2 are transversal but not strongly transversal.

D. Lemma

If C_1, C_2 are convex cones in $\mathbb{R}^n$, then C_1 and C_2 are transversal if and only if either

(i) C_1 and C_2 are strongly transversal,

or

(ii) C_1 and C_2 are linear subspaces and $C_1 \oplus C_2 = \mathbb{R}^n$.

PROOF.

It suffices to assume that C_1 and C_2 are transversal but not strongly transversal and show that (ii) holds. (Recall that (ii) says: "C_1 and C_2 are linear subspaces and $C_1 \oplus C_2 = \mathbb{R}^n$.")

Let us prove that C_1 is a linear subspace. Pick $c \in C_1$. Using the transversality of C_1 and C_2 write

$$-c = c_1 - c_2, \quad c_1 \in C_1\,, \quad c_2 \in C_2\,.$$

Then $c_1 + c = c_2$. But $c_1 + c \in C_1$ and $c_2 \in C_2$. So $c_1 + c \in C_1 \cap C_2$, and then $c_1 + c = 0$, since C_1 and C_2 are not strongly transversal. Therefore $-c = c_1$, so $-c \in C_1$. This shows that $c \in C_1 \Rightarrow -c \in C_1$. So C_1 is a linear subspace. A similar argument shows that C_2 is a linear subspace. Then the transversality of C_1 and C_2 implies that $C_1 + C_2 = \mathbb{R}^n$, and the fact that they are not strongly transversal implies that $C_1 \cap C_2 = \{0\}$. Hence $C_1 \oplus C_2 = \mathbb{R}^n$. *END OF PROOF.*

VI. Set separation

Two subsets S_1, S_2 of a Hausdorff topological space space T are locally separated at a point $p \in T$ if there exists a neighborhood U of p in T such that

$$S_1 \cap S_2 \cap U \subseteq \{p\}.$$

VII. The Transversal Intersection Property

If two subsets S_1, S_2 of $\mathbb{R}^n$ have tangent cones C_1, C_2 at a point p, and the cones C_1, C_2 are strongly transversal, then S_1 and S_2 are not locally separated at p.

The statement that "S_1 and S_2 are not locally separated at p" means the following:

$S_1 \cap S_2$ contains a sequence of points p_j converging to p and $\neq p$.

A. Remark. This is exactly the "intuition" discussed earlier.

B. Question. For what notions of "tangent cone to a set at a point" is the TIP (Transversal Intersection Property) true?

VIII. How the TIP is applied to prove versions of the FDPMP

To apply the TIP to prove a version of the FDPMP for optimal control, one carries out the following steps:

St 1. Reduce the optimal control problem to a separation problem in which, for a dynamics $\boxed{\dot{x} = f(x, u, t)}$, and an interval $[a, b]$, it is required that the reachable set $\mathcal{R}(f, [a, b], \bar{x}_{in})$ be locally separated from some other given set S. (This reduction is well known. It amounts to "augmenting the system by adding the cost as a new dynamical variable".)

St 2. Construct a "tangent cone" C_1 to $\mathcal{R}(f, [a, b], \bar{x})$ at the terminal point $\bar{x}_{term}$ of the reference trajectory.

NOTE: $\mathcal{R}(f, [a, b], \bar{x})$ is the set of all points reachable from the initial point $\bar{x}_{in}$ over the interval $[a, b]$ for the dynamics f.

St 3. Compare C_1, the tangent cone to $\mathcal{R}(f,[a,b],\bar{x})$ at $\bar{x}_{term}$, to C_2, the tangent cone to S at $\bar{x}_{term}$.

St 4. Use the TIP to conclude that C_1 and C_2 cannot be strongly transversal, because $\mathcal{R}(f,[a,b],\bar{x}_{in})$ and S are locally separated at $\bar{x}_{term}$.

St 5. If we can go from "not strongly transversal" to "not transversal," then the non-transversality is exactly the existence of a nontrivial covector linearly separating C_1 and C_2, and this yields the desired "adjoint vector" of the Maximum Principle.

St 6. How do we go from "not strongly transversal" to "not transversal"? In optimal control this is easy, because the cone C_2 is, typically, the product of a tangent cone to the set of admissible terminal states times a half-line, so it is never a linear subspace.

Naturally, for all this to work one needs the notion of "tangent cone" used in the above steps to be such that the TIP is true.

THEOREM: The TIP is true if "tangent cone" is taken to mean "Boltyanskii approximating cone." (The proof of this is Type T.)

THEOREM: The TIP is true if "tangent cone" is interpreted to mean "Clarke tangent cone." (The proof of this is Type L.)

The first TIP result leads to a number of versions of the FDPMP with a Boltyanskii or Boltyanskii-like tangent cones in the transversality condition. In these versions, high-order conditions can easily be included. (Classical work by Pontryagin et al., work by Knobloch, Krener, Agrachev, Sarychev, Gamkrelidze, Bianchini, Stefani, HJS, and lots of others.) These results are all proved using the TIP for Boltyanskii cones or for some generalization of them, such as the "approximating multicones" used by HJS.

The second TIP result leads to a number of versions of the FDPMP with a Clarke or Mordukhovich normal cone in the transversality condition. (Work by Clarke, Vinter, Rockafellar, Ioffe, Mordukhovich, Loewen, da Pinho, Franskowska, and lots of others.) In these versions, it does not seem that high-order conditions can be incorporated. Most of these results are not proved by explicitly using the TIP for Clarke cones or for some generalization thereof, but work is now in progress by HJS which, it is hoped, will show that they can be proved that way.

It may seem natural to expect that a more general TIP might be true, containing both results. I conjectured (and even briefly believed I had proved) about 10 years ago that such a result was true.

The problem was solved in January, 2006, by Alberto Bressan, who proved the following:

IX. Bressan's Theorem

There exist two closed subsets S_1, S_2 of $\mathbb{R}^4$, and two closed convex cones C_1, C_2 in $\mathbb{R}^4$, such that

- C_1 is a Boltyanskii approximating cone to S_1 at 0;
- C_2 is the Clarke tangent cone to S_2 at 0;
- C_1, C_2 are strongly transversal;
- $S_1 \cap S_2 = \{0\}$.

Using Bressan's example, one can construct an example of a Lagrange optimal control problem in $\mathbb{R}^8$ with a terminal state constraint, and an optimal trajectory-control pair (ξ_*, η_*), defined on an interval $[a_*, b_*]$, such that

- the dynamics and Lagrangian satisfy conditions that lend themselves to Type T arguments,
- the terminal set S has a Clarke tangent cone C at the terminal point of $\xi_*(b)$,
- there does not exist a nontrivial multiplier $(\pi(\cdot), \pi_0)$ (consisting of an adjoint covector $\pi(\cdot)$ and "abnormal multiplier" π_0) that satisfies the adjoint equation, the Hamiltonian maximization condition, and the transversality condition $-\pi(b_*) \in C^{\perp}$.

The actual construction is done in complete detail in the paper, and it's sort of technical.

Remark: In this particular example, the usual nonsmooth "adjoint differential inclusion" is actually a true "adjoint differential equation."

A lot remains to be done. For example,

- find a good counterexample as above, with a very smooth optimal control problem, for which one can get lots of high-order necessary conditions for optimality involving high-order variations in the direction of Lie brackets, but for which the terminal condition on the state involves a set with a Clarke tangent cone.

- carry out the program of proving all Type L versions of the FDPMP using the "Type L" TIP. A first step in that direction was my paper in the Sevilla CDC, where I introduced a concept of "approximating multicones" (called "Mordukhovich-Warga approximating multicones") adapted to Type L arguments, and prove the TIP.

- find a good example of failure of the FDPMP for which the dynamics are appropriate to Type L arguments, but the terminal condition on the state involves a set with a Boltyanskii approximating cone. (Conjecture: this will probably happen for some problem which is governed by a differential inclusion $\dot{x} \in F(x,t)$, and whose adjoint equation is the "intrinsic adjoint equation" involving a partial convexification of the Mordukhovich normal cone to the graph of F.)

Argument for the conjecture: I have tried and tried to derive the intrinsic equation in the Type T setting and wasn't able to. This suggest to me that perhaps the intrinsic equation can only be derived with Type L methods, in which case it is reasonable to expect that it will not "go well" with a Boltyanskii tangent cone to the terminal set.