3.3.2 Cost functional

Consider functions $L(t,x,u)$ and $K(t,x)$. Here $L : \mathbb{R}\times\mathbb{R}^n\times\mathbb{R}^m \to \mathbb{R}$, $K : \mathbb{R}\times\mathbb{R}^n \to \mathbb{R}$ —sufficiently regular.

(To be consistent with calculus of variations we would have to take $L(t,x,\dot{x})$, not $L(t,x,u)$. However, we can always go from $(t,x,\dot{x})$ to (t,x,u) by plugging in $\dot{x} = f(t,x,u)$, but not necessarily back, so the latter choice of variables is preferable.)

Unlike with f, u, there are no a priori conditions here. All derivatives that will arise will be assumed to exist, so depending on the analysis method we will see at the end what differentiability assumptions on K, L we need.

$$J := \int_{t_0}^{t_f} L(t,x(t),u(t))dt + K(t_f,x_f)$$

Here $J = J(u)$ or, in more detail, $J(t_0,x_0,t_0,u,t_f,x_f)$, where:

- t_0 and $x_0 = x(t_0)$ are the initial time and state;
- t_f and $x_f = x(t_f)$ are the final time and state (x_f depends on u);
- L is the *running cost* (or Lagrangian);
- K is the *terminal cost*.

The above form of J is *Bolza form*. Two special cases:

Lagrange form $J = \int_{t_0}^{t_f} L(t,x,u)dt$ (no terminal cost). This comes from calculus of variations (Lagrange problem)

Mayer form $J = K(t_f,x_f)$ (no running cost)

We can pass back and forth between these:

Mayer $\mapsto$ Lagrange

$$\begin{aligned} K(t_f,x_f) &= \underbrace{K(t_0,x_0)}_{\text{constant, can ignore}} + \int_{t_0}^{t_1} \frac{d}{dt}(K(t,x(t))dt \\ &= K(t_0,x_0) + \int_{t_0}^{t_1} (K_t + K_x\cdot f(t,x,u))dt \end{aligned}$$

Lagrange $\mapsto$ Mayer Introduce additional state x^0:

$$\dot{x}^0 = L(t,x,u), \qquad x^0(t_0) = 0$$

then

$$\int_{t_0}^{t_f} L(t,x,u)dt = x^0(t_f)$$

which is a terminal cost. (The value of $x^0(t_0)$ actually doesn't matter, it's just a constant and we can subtract it.)

$\longrightarrow$ Note that a similar trick can be used to eliminate time-dependence (of f and L): Let $x_{n+1} := t$, so $\dot{x}_{n+1} = 1$, and treat is as additional state (in this case we need f to be $\mathcal{C}^1$ in t for existence and uniqueness of solutions).

3.3.3 Target set

(t_0, x_0) are given, but what about (t_f, x_f)? One or both can be free or fixed, or need to belong to some set. All these possibilities are captured by introducing a *target set* $S \subset [t_0, \infty) \times \mathbb{R}^n$ (or use $\mathbb{R}$ for time) such that $(t_f, x_f) \in S$.

Examples (see [A-F, Sect. 4-13 and 5-12]):

- $S = [t_0, \infty) \times \{x_1\}$, where x_1 is fixed. This gives a *fixed-endpoint, free-time problem.*
- $S = \{t_1\} \times \mathbb{R}^n$, where t_1 is fixed. This gives a *free-endpoint, fixed-time problem.*
- $S = [t_0, \infty) \times S_1$, where S_1 is a surface (manifold) in $\mathbb{R}^n$. *Note:* this is what we get if we start with a free-endpoint, fixed-time problem and consider auxiliary state x_{n+1} as before: $\dot{x}_{n+1} = 1$. Then $S_1 \subset \mathbb{R}^{n+1}$ is $\{x : x_{n+1} = t_1\}$.
- More generally, $S = T \times S_1$ where S_1 is as in the previous item and T is some subset of $[t_0, \infty)$. As a special case of this, we have $S = \{t_1\} \times \{x_1\}$ (*fixed-endpoint, fixed-time problem*, the most restrictive).
- $S = \{(t, g(t)) : t \in [t_0, \infty)\}$ for some $g : \mathbb{R} \to \mathbb{R}^n$. This is *trajectory tracking (moving target).* Or $g(t)$ can be generalized to a set, i.e., g can be set-valued.

Finally, *Problem:*

Find $u(\cdot)$ that minimizes J.

3.4 A variational approach to the optimal control problem

We'll first try to see how far we can get with variational techniques. The problems we'll encounter will motivate us to modify our approach and will lead to the the Maximum Principle . We'll actually also arrive at the correct statement of the Maximum Principle .

References: [A-F, Sect. 5-7, 5-9, 5-10] (5-8 is sufficient conditions); [ECE 515 Class Notes, Sect. 11.2] .

3.4.1 Free-endpoint, fixed-time case

$$\dot{x} = f(t, x, u), \qquad x(t_0) = x_0$$

where $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$ (unconstrained),

$$J(u) = \int_{t_0}^{t_1} L(t, x, u)dt + K(x(t_1)) \tag{3.2}$$

(no direct dependance of terminal cost on t_1), t_1 fixed, $x(t_1)$ free ($S = \{t_1\} \times \mathbb{R}^n$).

Let $u^*(t)$ be an optimal control (assumed to exist), and let $x^*(t)$ be the corresponding optimal trajectory.

Goal: derive necessary conditions.

To simplify matters, assume that u^* is a *global* min:

$$J(u^*) \leq J(u)$$

for all admissible (i.e., piecewise continuous) u.

Begin optional:

If we want to work with local optimality, then it can be shown that the right topology for which the Maximum Principle works is uniform convergence of x and L_1 convergence of u —this is *Pontryagin's topology* [Milyutin and Osmolovskii, Calculus of Variations and Optimal Control, AMS, 1998, p. 23]. On the other hand, the variational approach is with respect to $\|\cdot\|_0$ topology on $\mathcal{U}$. This is close to weak minima: if $\dot{x} = u$ then this is exactly weak minimum, $(\alpha\eta, \alpha\xi)$ for α small.

End optional

Right now, we ignore these distinctions as we work with a global minimum.

Perturbation:

We cannot perturb $x^*(t)$ directly, only through the control. This is new, compared to calculus of variations with constraints where we had perturbations of the curve subject to $\delta C(\eta) = 0$, C—constraint.

So, consider control

$$u = u^* + \alpha\eta$$

It gives the trajectory

$$x = x^* + \alpha\xi$$

where we have $\xi(t_0) = 0$. How are η and ξ related? On one hand,

$$\dot{x} = \dot{x}^* + \alpha\dot{\xi} = f(t, x^*, u^*) + \alpha\dot{\xi}$$

On the other hand, using Taylor with respect to α around $\alpha = 0$,

$$\dot{x} = f(t, x, u) = f(t, x^* + \alpha\xi, u^* + \alpha\eta) = f(t, x^*, u^*) + f_x|_* \xi\alpha + f_u|_* \eta\alpha + o(\alpha)$$

where $|_*$ always means evaluation at the optimal trajectory. Hence

$$\dot{\xi} = f_x|_* \xi + f_u|_* \eta + \frac{o(\alpha)}{\alpha} \tag{3.3}$$

$\longrightarrow$ Note: this is a forced linear *time-varying* system, of the form

$$\dot{\xi} = A(t)\xi + B(t)\eta + \text{h.o.t.} \tag{3.4}$$

(f_x is a matrix!) This is nothing but the *linearization* of the original system around the trajectory $x^*(t)$.

We have J given by (3.2), and we have the constraint $\dot{x}(t) - f(t, x, u) = 0$. Recall Lagrange's idea for treating non-integral constraints: for an arbitrary function $p(t)$, rewrite the cost as

$$J(u) = K(x(t_1)) + \int_{t_0}^{t_1} (L(t, x, u) + p(t)(\dot{x} - f(t, x, u)))\, dt$$

Here $p(t)$ is a Lagrange multiplier function (what we called $\lambda(x)$ in calculus of variations , λ is also often used), it is unspecified for now (can be anything). (The relation with Lagrange multipliers will become clearer in Exercise 10 below.)

Also note that p, f, $\dot{x}$ are n-vectors, so we are dealing with inner products. We will write this more explicitly from now on:

$$\langle p, \dot{x} - f(t, x, u)\rangle$$

Let's introduce the *Hamiltonian:*

$$H(t, x, u, p) := \langle p, f(t, x, u)\rangle - L(t, x, u)$$

$\longrightarrow$ In calculus of variations we had $H = \langle p, y'\rangle - L$ which matches the current definition (since in new notation, $y' \mapsto \dot{x}$).

We have

$$J(u) = K(x(t_1)) + \int_{t_0}^{t_1} (\langle p, \dot{x}\rangle - H(t, x, u, p))\, dt$$

Similarly,

$$J(u^*) = K(x^*(t_1)) + \int_{t_0}^{t_1} (\langle p, \dot{x}^*\rangle - H(t, x^*, u^*, p))\, dt$$

$\longrightarrow$ We know: $J(u) \geq J(u^*)$.

Let us analyze $J(u) - J(u^*)$, by *Taylor:*

$$K(x(t_1)) - K(x^*(t_1)) \approx \langle K_x(x^*(t_1)), x(t_1) - x^*(t_1)\rangle = \langle K_x(x^*(t_1)), \xi(t_1)\rangle\, \alpha$$

$\longrightarrow$ Here and below, $\approx$ stands for *equality up to* $o(\alpha)$.

$$\begin{aligned} H(t, x, u, p) - H(t, x^*, u^*, p) &\approx \langle\, H_x|_* \,, x - x^*\rangle + \langle\, H_u|_* \,, u - u^*\rangle \\ &= \langle\, H_x|_* \,, \xi\rangle\, \alpha + \langle\, H_u|_* \,, \eta\rangle\, \alpha \end{aligned}$$

Hence,

$$J(u) - J(u^*) \approx \alpha\, \langle K_x(x^*(t_1)), \xi(t_1)\rangle + \alpha \int_{t_0}^{t_1} \Big(\langle p, \dot{\xi}\rangle - \langle\, H_x|_* \,, \xi\rangle - \langle\, H_u|_* \,, \eta\rangle\Big)\, dt$$

$\longrightarrow$ We have just computed the first variation $\delta J|_{u^*}(\eta)$! It's all the first-order terms in α.

As in calculus of variations , let us simplify the term containing $\dot{\xi}$, integrating it by parts:

$$\int \langle p, \dot{\xi}\rangle dt = \langle p, \xi\rangle|_{t_0}^{t_1} - \int \langle \dot{p}, \xi\rangle dt$$

Since $x = x^* + \alpha\xi$ and $x(t_0)$ is fixed, we must have $\xi(t_0) = 0$. So,

$$\int \langle p, \dot{\xi}\rangle dt = \langle p(t_1), \xi(t_1)\rangle - \int \langle \dot{p}, \xi\rangle dt$$

and we see that

$$\delta J|_{u^*}(\eta) = \langle K_x(x^*(t_1)) + p(t_1), \xi(t_1)\rangle - \int_{t_0}^{t_1} \Big(\langle\, H_x|_* + \dot{p}, \xi\rangle + \langle\, H_u|_* \,, \eta\rangle\Big) dt$$

The familiar first-order necessary condition for optimality (Section 1.3.2) says that we must have $\delta J|_{u^*}(\eta) = 0$. But this makes the choice of $p(\cdot)$ clear:

$$\dot{p} = -H_x|_* = -(f_x)^T\Big|_* p + L_x|_*$$

(recalling that $H = f^T p - L$), and

$$p(t_1) = -K_x(x^*(t_1))$$

is the boundary condition. (If $K = 0$, i.e., no terminal cost, then $p(t_1) = 0$.)

$\longrightarrow$ Note: we have a *final* (not initial) condition for p.

Note that the ODE for $p(t)$ is also linear (affine, or forced):

$$\dot{p} = -A^T(t)p + L_x|_*$$

where A is the same as in the ODE (3.4) for the state perturbation ξ.

Two linear systems with matrices A and $-A^T$ are called *adjoint*. For this reason, p is called the *adjoint vector. Call it* p^* *from now on.*

$\longrightarrow$ The significance of the concept of adjoint will become clear in the proof of the Maximum Principle ($\langle p, x\rangle =$ const).

p is also called the *costate*, because it *acts* on $x, \dot{x}$ by $p^T\dot{x}$ (dual, *co*vector).

It is also a (distributed) Lagrange multiplier.

Note that x^* satisfies $\dot{x} = f = H_p|_*$ (recall again that $H = \langle p, f\rangle - L$), so x^* and p^* satisfy *Hamilton's canonical equations*:

$$\begin{aligned} \dot{x}^* &= H_p|_* \\ \dot{p}^* &= -H_x|_* \end{aligned} \tag{3.5}$$

(From now on, $|_*$ means evaluation at (x^*, p^*)—only then are the equations valid.)

If we had a fixed endpoint problem, then the term $\langle p(t_1), \xi(t_1)\rangle$ would disappear and the boundary condition on $p(t_1)$ would not be enforced; we would use the endpoint condition on x instead. *More on this case later* (problems arise).

With the above definition of p, we are left with

$$\delta J|_{u^*}(\eta) = -\int_{t_0}^{t_1} \langle H_u|_*, \eta\rangle \, dt = 0 \qquad \forall \eta$$

We know from calculus of variations ("DuBois-Reymond") that this means $H_u|_* = 0$, or

$$H_u(t, x^*(t), u^*(t), p^*(t)) \equiv 0 \tag{3.6}$$

(already saw this in calculus of variations). This means that $H(u)$ has an *extremum* at u^*.

Exercise 10 (due Mar 7) The purpose of this exercise is to recover the earlier conditions from calculus of variations, namely, the Euler-Lagrange equation and the Lagrange multiplier condition (for the case of multiple degrees of freedom and several constraints) from the preliminary necessary conditions for optimal controls derived so far, namely, (3.5) and (3.6).

a) The standard (unconstrained) problem from calculus of variations can be rewritten in the new optimal control language if we consider the control system

$$\dot{x}_i = u_i, \qquad i = 1, \dots n$$

together with the cost $J(u) = \int_{t_0}^{t_1} L(t, x, \dot{x})dt$. Use (3.5) and (3.6) for this system to prove that L must satisfy the Euler-Lagrange equation, $\frac{d}{dt}L_{\dot{x}_i} = L_{x_i}$, along the optimal trajectory.

b) Now suppose that we have $m < n$ non-integral constraints. We can model this by the control system

$$\begin{aligned} \dot{x}_i &= f_i(t, x_1, \dots, x_n, u_1, \dots, u_{n-m}), \qquad & i = 1, \dots, m \\ \dot{x}_{m+i} &= u_i, & i = 1, \dots, n - m \end{aligned}$$

and the same cost as in a). Use (3.5) and (3.6) for this system to prove that there exist functions $\lambda_i(t)$, $i = 1, \dots, m$ such that the augmented Lagrangian $L + \sum_{i=1}^{m} \lambda_i(t)(\dot{x}_i - f_i)$ satisfies the Euler-Lagrange equation along the optimal trajectory.

Hint: part b) is a bit tricky. The Lagrange multipliers λ_i are related to the components of the adjoint vector p^*, but they are not the same. Also note that the Lagrangian has $\dot{x}_i$ as arguments (since $L_{\dot{x}_i}$ appear in Euler-Lagrange equation), whereas the Hamiltonian that you will construct should not (it is a function of t, x, u, p), so you will need to use the differential equations to put H in the right form. A consequence of this, in part b), is that x_i will appear inside H in two different places.

In order to see whether the extremum of H as a function of u along the optimal trajectory is a minimum or a maximum, let us bring in the second variation. In $\delta^2 J\big|_{u^*}$, going back to the earlier expression for $J(u) - J(u^*)$ but this time expanding up to second-order terms in α, we'll get second partial derivatives of K outside the integral and second partial derivatives of H inside the integral. Let us look at the latter (since we care about H).

$x = x^* + \alpha\xi$, $u = u^* + \alpha\eta$.

$$H(t, x, u, p) - H(t, x^*, u^*, p) = \langle H_x|_*, \xi\rangle \alpha + \langle H_u|_*, \eta\rangle \alpha + \begin{pmatrix}\xi & \eta\end{pmatrix} \begin{pmatrix} H_{xx} & H_{xu} \\ H_{xu} & H_{uu} \end{pmatrix}\Bigg|_* \begin{pmatrix}\xi \\ \eta\end{pmatrix} \cdot \frac{\alpha^2}{2} + o(\alpha^2)$$

Which term in the above Hessian matrix is more important?

As in calculus of variations (Section 2.6), we can see that η is more important: we may have a large η producing small ξ (but not vice versa). (In calculus of variations η' was more important than η.)

Since H appears in J with a $-$ sign, the second-order necessary condition is

$$\eta^T H_{uu}|_* \eta \le 0 \qquad \forall \eta$$

i.e., the Hessian of H as a function of u along the optimal trajectory should be negative semidefinite.

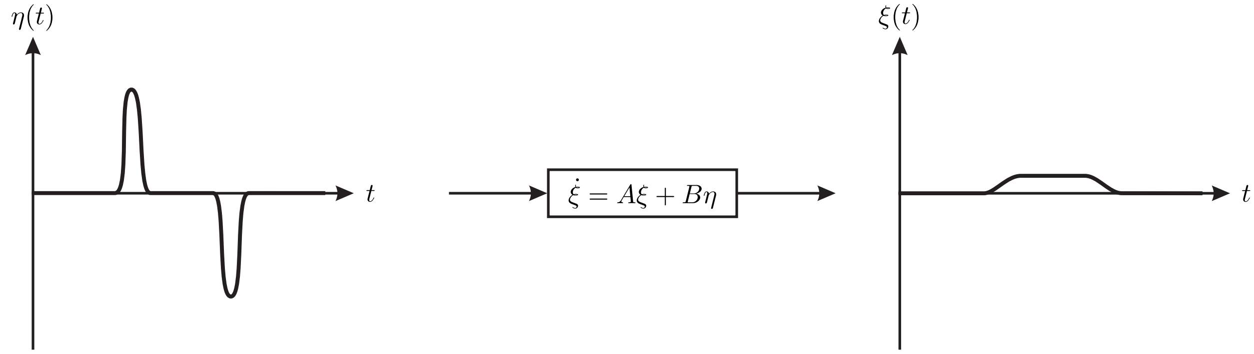

Figure 3.4: Large η produces small ξ

Therefore, the extremum of H at u^* is actually a *maximum*. We arrive at the following *necessary condition for optimality*: If $u^*(t)$ is an optimal control and $x^*(t)$ is the corresponding optimal trajectory, then there exists an adjoint vector $p^*(t)$ (we use superscript $*$ to emphasize that it is associated with the optimal trajectory) such that:

1) $x^*(t), p^*(t)$ satisfy the canonical equations (3.5),where $|_*$ means evaluation at x^*, p^*, u^*, with boundary conditions $x^*(t_0) = x_0$, $p^*(t_1) = -K_x(x^*(t_1))$.

2) $H(t, x^*(t), u, p^*(t))$, viewed as a function of u for each t (here we treat u as a variable independent of x and p, ignoring system dynamics), has a (possibly local) maximum at $u = u^*(t)$, for each $t \in [t_0, t_1]$:
$$H(t, x^*(t), u^*(t), p^*(t)) \geq H(t, x^*(t), u, p^*(t)) \qquad \forall\, u \text{ near } u^*(t),\ \forall\, t$$

Comments:

- Sign difference: we could define a different Hamiltonian and a different costate:
$$\hat{p} := -p, \hat{H} := \langle \hat{p}, f \rangle + L = -H$$
So $\hat{H}$ would have a *minimum* at u^*, and $\hat{p}$ would still satisfy
$$\dot{\hat{p}} = -\dot{p} = H_x = (f_x)^T p - L_x = -(f_x)^T \hat{p} - L_x = -\hat{H}_x$$
This alternative convention is the one followed in [A-F], and it might seem a bit more natural: $\min_u J \mapsto \min_u H$. But the convention we follow is consistent with the definition of Hamiltonian in calculus of variations and its mechanical interpretation as total energy of the system. Note that if the original problem is maximization, then it can also be reduced to minimization by $L \mapsto -L$.

  $\longrightarrow$ So, whether we get *minimum principle* or *maximum principle* has nothing to do with max or min in the problem, and is determined just by the sign convention.

- Note that our calculus of variations definition of the momentum ($p = L_{\dot{x}}$) has disappeared and is replaced by the existential statement. This is more suitable for the Maximum Principle but in the present variational setting it can be recovered: if $\dot{x} = u$ (unconstrained case in calculus of variations) then $H = pu - L$ and we have
$$0 = H_u = p - L_u$$
so $p = L_u$.

- Using canonical equations, we have (note that in general $H = H(t, x, u, p)$, i.e., explicitly depends on t):

$$\left.\frac{dH}{dt}\right|_* = H_t|_* + \langle H_x|_*, \dot{x}^* \rangle + \langle H_p|_*, \dot{p}^* \rangle + \langle H_u|_*, \dot{u}^* \rangle = H_t|_*$$

  because the first two terms cancel each other and the third term is zero. (Assuming that u is differentiable.)

  $\longrightarrow$ In particular, if the problem is time-invariant (no t in f and L) then H is constant along the optimal trajectory.

- Making suitable assumptions so that $\delta^2 J\big|_{u^*} > 0$ and dominates $o(\alpha^2)$, we can get *sufficient* conditions for optimality [A-F, Sect. 5.8]. They apply, e.g., to LQR (same section).

The following problem is on the last point, and also illustrates the necessary conditions.

Exercise 11 (due Mar 7) Consider

$$\dot{x} = Ax + Bu$$

(LTI or LTV, doesn't matter),

$$J = \int_{t_0}^{t_1} (x^T Q x + u^T R u) dt, \qquad Q \geq 0, \ R > 0$$

(or $Q(t) \geq 0$, $R(t) > 0 \ \forall t$), free endpoint, fixed time as above.

a) Write down the above necessary conditions for this case, and find a formula for the control satisfying them.

b) By analyzing the second variation, show that this control is indeed optimal (in what sense?) .

3.4.2 Critique of the variational approach and preview of the Maximum Principle

(We'll completely trash it now.)

Summary of the above approach:

- We took a perturbation

$$u = u^* + \alpha\eta \tag{3.7}$$

  where α is a small number.

- We showed (using the first variation) that

$$\int_{t_0}^{t_1} \langle H_u|_*, \eta \rangle \, dt = 0 \tag{3.8}$$

- We deduced that $H_u|_* \equiv 0$, hence H has an extremum at u^*.

- We then showed (using the second variation) that $H_{uu}|_* \leq 0$, hence the extremum is a *maximum*.

The good thing about this approach is that it leads, quite quickly, to the correct form of necessary conditions (canonical equations plus Hamiltonian maximization).

However, it has several severe limitations which we now discuss.

- Considering perturbations of the form (3.7) is OK for the unconstrained case $u \in \mathbb{R}^m$. However, in control applications we usually have constraints on the control range: $u \in \mathcal{U} \subset \mathbb{R}^m$ —bounded set. $\mathcal{U}$ may even be finite (think of car gears)!

  We'll see in the Maximum Principle that the statement that $H(u)$ has a maximum at u^* (with all other arguments evaluated along the optimal trajectory) is still correct!

  $\longrightarrow$ In fact, u^* is a *global* maximum of $H(u)$.

  But the variational approach doesn't let us show this, because $H_u|_* = 0$ need not hold for u^* on the boundary of $\mathcal{U}$.

- ([A-F, Sect. 5-9]) We treated the free-endpoint case, but in practice we want to hit some target set S. Take the simplest case: $S = \{t_1\} \times \{x_1\}$ (fixed endpoint). Then the control perturbation η is no longer arbitrary. It gives a state perturbation ξ, and we must have $\xi(t_1) = 0$. We know that ξ and η are related by

  $$\dot{\xi} = A\xi + B\eta$$

  so we need

  $$0 = \xi(t_1) = \int_{t_0}^{t_1} \Phi(t_1, t)B(t)\eta(t)dt$$

  where $\Phi(\cdot, \cdot)$ is the transition matrix for $A(\cdot)$. The equation (3.8) holds *only for these perturbations*, and not for all η. So we cannot conclude that $H_u|_* = 0$. (Lemma 5-1 in [A-F, p. 248] doesn't help.)

- Writing conditions like $H_u|_* = 0$, we are assuming that H is differentiable with respect to u. But recall that $H = \langle p, f \rangle - L$. So first, we need L_u to exist. This rules out very reasonable cost functionals like

  $$\int_{t_0}^{t_1} |u(t)|dt$$

  Second, we need f_u to exist.

  Was this part of the standing assumptions needed for existence and uniqueness of solutions?

  No! So, the variational approach requires *extra* regularity assumptions on the system.

  The second variation analysis (needed to conclude that u^* is a maximum) involves H_{uu} —even worse!

  $\longrightarrow$ We want to establish the maximization claim directly, not by working with derivatives.

- $u = u^* + \alpha\eta$, $x = x^* + \alpha\xi$ —this means that we are allowing only small deviations in both x and u. If we had $\dot{x} = u$, this would correspond exactly to the notion of weak minimum in calculus of variations . $\dot{x} = f(t, x, u)$, $\|x - x^*\|_0$ and $\|u - u^*\|_0$ small is an appropriate generalization of that.

  However, we want optimality with respect to control perturbations that may be large, as long as the corresponding trajectories are close (as discussed some time ago); see Figure 3.5.

  Important for bang-bang controls ($\mathcal{U}$ bounded).

  We want to incorporate such perturbations—this would correspond to the notion of strong minima from calculus of variations (actually, we have minima with respect to Pontryagin's topology, see

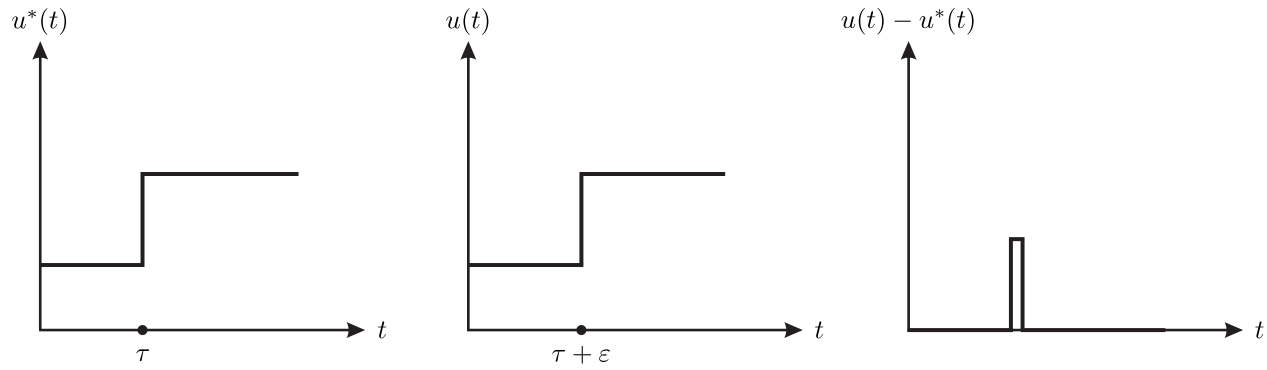

Figure 3.5: Control perturbation

before). A richer perturbation family is crucial for obtaining a better result, i.e., sharper necessary conditions that we want (MP).

See a nice quote on p. 154 of "Stochastic Controls" by Zhou and Yong (cutting the "umbilical cord").

So, while the basic statement of the result (canonical equations plus Hamiltonian maximization) in the Maximum Principle will be very similar to the one obtained using the variational approach, we want to incorporate:

- Constraints on the control set;
- Constraints on the final state;
- Less restrictive notion of "closeness" of controls (this one is key for the others);
- Fewer differentiability assumptions

and for these reasons, the proof will be quite different and much more involved.

the Maximum Principle was proved many years later than all basic calculus of variations results, and it is a very nontrivial extension.

$\longrightarrow$ The proof of the Maximum Principle is also much more *geometric*.

Chapter 4

The Maximum Principle

4.1 The statement of the Maximum Principle

For the Maximum Principle , A-F is a primary reference. Supplementary references: the original book by Pontryagin et al.; Sussmann's lecture notes (more rigorous and modern treatment).

$\longrightarrow$ All standing assumptions on the general control problem still hold, we'll just make them more specific.

We will state and prove the Maximum Principle for two special problems, and later discuss how all other cases can be deduced from this.

4.1.1 Special Problem 1

System: $\dot{x} = f(x,u)$, $x(t_0) = x_0$ (given), no t dependence, f is $\mathcal{C}^1$ in x (no differentiability with respect to u!), f and f_x are $\mathcal{C}^0$ in u (these are conditions for existence and uniqueness we had earlier). $x \in \mathbb{R}^n$, $u \in \mathcal{U}$—some subset of $\mathbb{R}^m$. Can be the whole $\mathbb{R}^m$. Usually closed (Sussmann assumes this, but A-F and Pontryagin don't).

$L = L(x,u)$, no t dependence. Assume that L satisfies the same regularity assumptions as f.

$K = 0$ —no terminal cost.

Target set: $S = \mathbb{R} \times \{x_1\}$—fixed-endpoint, free-time problem.

$J(u) = \int_{t_0}^{t_1} L(x,u)dt$, where t_1 is the *first time* at which $x(t) = x_1$.

Theorem 4 (Maximum principle for Special Problem 1 [Theorem 5-5P in [A-F, p. 305]) *Let $u^*(t)$ be an optimal control and let $x^*(t)$ be the corresponding trajectory. Then there exists a function $p^*(t)$ and a constant $p_0^* \leq 0$, satisfying $(p_0^*, p^*(t)) \neq (0,0)\ \forall\, t$, such that:*

1) $x^*(t)$ and $p^*(t)$ satisfy the canonical equations

$$\dot{x}^* = H_p(x^*, u^*, p^*, p_0^*)$$
$$\dot{p}^* = -H_x(x^*, u^*, p^*, p_0^*)$$

with boundary conditions $x^*(t_0) = x_0$, $x^*(t_1) = x_1$, where the Hamiltonian is defined as

$$H(x,u,p,p_0) := \langle p, f(x,u)\rangle + p_0 L(x,u) \qquad \forall\, u \in \mathcal{U},\ \forall\, t \in [t_0, t_1]$$

2) $H(x^*(t), u^*(t), p^*(t), p_0^*) \geq H(x^*(t), u, p^*(t), p_0^*)\ \forall\, u \in \mathcal{U},\ \forall\, t \in [t_0, t_1]$, *i.e.*, u^* *is a global maximum of* $H(x^*(t), \cdot, p^*(t))\ \forall\, t$.

3) $H(x^*(t), u^*(t), p^*(t), p_0^*) \equiv 0,\ t \in [t_0, t_1]$.

Comments:

- u^* is *optimal*—by this we mean that for any other control u taking values in $\mathcal{U}$ which transfers x from x_0 at $t = t_0$ to x_1 at some (unspecified) time $t = t_1$,

$$J(u^*) \leq J(u)$$

  (We could relax the assumption of *global* minimum to a *local* minimum, but need to know what topology to use, so we won't do that. The proof will make the topology more clear.)

- p_0 is the *abnormal multiplier*. We saw its analog in calculus of variations . It handles pathological cases. If $p_0 \neq 0$, then we can always normalize it to $p_0 = -1$ and forget it, which is what one does when applying the Maximum Principle (almost always). For this reason, we will also suppress the argument p_0 of H from now on.

- We saw in the variational approach (just a little while ago) that $H = const$, but $H \equiv 0$ may seem surprising. This is a feature of the *free-time problem* (earlier we were looking at the fixed-time case). More on this later.

How restrictive is Special Problem 1?

- Time-independence—no problem: if time-dependent, can always set $x_{n+1} = t$ and eliminate t (but we need f to be C^1 with respect to t);

- No terminal cost—no problem: if we have terminal cost, can do

$$K(t_f, x_f) = K(t_0, x_0) + \int_{t_0}^{t_1} (K_t + K_x f(t, x, u))\, dt$$

  and define $\hat{L} := L + K_t + K_x f$.

- $S = \mathbb{R} \times \{x_1\}$—this is not very general, need more general target set.

4.1.2 Special Problem 2

Same as Special Problem 1, *except* $S = \mathbb{R} \times S_1$, where S_1 is a k-dimensional surface in $\mathbb{R}^n$, $1 \leq k \leq n$. This is defined as follows:

$$S_1 = \{x \in \mathbb{R}^n : h_1(x) = h_2(x) = \cdots = h_{n-k}(x) = 0\}$$

Here, h_i are smooth functions from $\mathbb{R}^n$ to $\mathbb{R}$. We also assume that each $x \in S_1$ is a regular point, in the sense discussed before: the gradients $(h_i)_x$ are linearly independent at x.

The special case $k = n$ means $S_1 = \mathbb{R}^n$.

Special Problem 1 can be thought of as a special case with $k = 0$.

Does the case $S_1 = \mathbb{R}^n$ (i.e., both x_f and t_f are free) make sense? When do we stop? Why even start moving?

Answer: yes, because L may be negative. (Also if we have K, which translates to the same thing: moving decreases the cost.)

$\longrightarrow$ When we say "cost", we may think implicitly that $L \geq 0$, but it is not necessarily true.

The difference between the next theorem and the previous one is only in the boundary conditions for the canonical system.

Theorem 5 (Maximum principle for Special Problem 2 [Theorem 5-6P in [A-F, p. 306]) *Same as the Maximum Principle for Special Problem 1, except:*

The final boundary condition in the canonical equations is

$$x^*(t_1) \in S_1$$

and there is one more necessary condition:

4) ("transversality condition")

The vector $p^*(t_1)$ is orthogonal to the tangent space to S_1 at $x^*(t_1)$:

$$\langle p^*(t_1), d \rangle = 0 \qquad \forall\, d \in T_{x^*(t_1)} S_1$$

We know this means that $p^*(t_1)$ is a linear combination of $(h_i)_x(x^*(t_1))$, since

$$T_{x^*(t_1)} S_1 = \{d \in \mathbb{R}^n : \langle (h_i)_x(x^*(t_1)), d \rangle = 0,\ i = 1, 2, \ldots, n-k\}$$

If $k = n$, then $S_1 = \mathbb{R}^n$. In this case, the transversality condition says (or can be extended to say)

$$p^*(t_1) = 0$$

(p^* must be orthogonal to all $d \in \mathbb{R}^n$, since there are no h_i).

Note that we still have n boundary conditions at $t = t_1$: k degrees of freedom for $x^*(t_1) \in S_1$ correspond to $n - k$ equations, and $n - k$ degrees of freedom for $p^*(t_1) \perp S_1$ correspond to k equations/constraints.

$\longrightarrow$ The freer the state, the less free the costate. Each additional degree of freedom for $x^*(t_1)$ kills one degree of freedom for $p^*(t_1)$.