These notes are supposed to run on from those in the previous scan of Chapter 23. I've tried to number the sections alternatively in this pdf to suggest the continuation of the material. The status of these lecture notes is : subject to further revisions!

1. Fréchet derivative [23. 9]

This section complements a discussion in the Lecture centered on section 23.2 concerning the variation of a functional and a definition of strong derivative which briefly is as follows.

The setting needs to be a normed vector space X, that is a vector space X with a norm $||.||_X$. Typical for applications is X taken to be the vector space of continuously differentiable functions $C^2[a,b]$ with norm

$$||x||_X = ||x||_\infty + ||\dot{x}||_\infty.$$

Here $||.||_\infty$ denotes the supremum norm. The current notes need amplification here.

Definition: Version 1.

If $F : X \to Y$ where X and Y are normed vector spaces, we say that a linear transformation $A : X \to Y$ is a strong derivative (or Fréchet derivative) of F at x if for every $\varepsilon > 0$ there is $\delta > 0$ such that it is the case that

$$||F(x+h) - F(x) - Ah||_Y \leq \varepsilon ||h||_X,$$

for all h with $||h||_X \leq \delta$.

This definition identifies the derivative of F with the 'linear part' of $F(x+h) - F(x)$ when this difference is regarded as a function of h.

One can show that A is unique (if it exists).

One should also attempt to re-read the above definition for the context of $X = R^n$ with Euclidean norm.

It is normal to strengthen the definition above to require that in addition A is **continuous** at x. This is to ensure that differentiability of F at x entails continuity of F at x.

What sense of continuity is required here?

One wants to say

$$\lim_{h\to 0} A(x+h) = A(x)$$

i.e. that for every $\varepsilon > 0$ there is $\delta > 0$ such that it is the case that

$$||A(x+h) - A(x)||_Y \leq \varepsilon$$

for all h with $||h||_X \leq \delta$.

However, since A is linear this turns out to be equivalent to

$$\lim_{h\to 0} A(h) = 0,$$

merely because $A(0) = 0$ and

$$A(x+h) - A(x) = A(h).$$

More interestingly though we discover that A is continuous at 0 iff for some M

$$||A(h)||_Y \leq M||h||_X.$$

You can read this as saying that a cannot magnify a vector by more that a bounded amount. In R^n if A is a matrix the eigenvalue of largest modulus identifies the maximal amount of magnification that A can effect.

Indeed taking $\varepsilon = 1$ we have for some δ that

$$||A(h)||_Y \leq 1$$

for all h with $||h||_X \leq \delta$. But now consider $h \neq 0$, then

$$||\frac{\delta h}{||h||}|| = \delta \frac{||h||}{||h||} = \delta.$$

So

$$\left|\left| A\left(\frac{\delta h}{||h||}\right) \right|\right|_Y \leq 1$$

or

$$\frac{\delta}{||h||} \, ||A\,(h)||_Y \leq \varepsilon.$$

hence

$$||A\,(h)||_Y \leq \frac{1}{\delta} ||h||_X$$

so the result follows fro $M = 1/\delta$.

Examples.

All the examples (i)-(iii) listed in Section 23.2 exhibit strong differentiability.

Definition: Version 2.

If $F : X \to Y$ where X and Y are normed vector spaces, we say that a **continuous** linear transformation $A : X \to Y$ is a strong derivative (or Fréchet derivative) of F at x if for every $\varepsilon > 0$ there is $\delta > 0$ such that it is the case that

$$||F(x + h) - F(x) - Ah||_Y \leq \varepsilon ||h||_X,$$

for all h with $||h||_X \leq \delta$.

When the derivative exists it is denoted $DF(x)$.

Easy conclusion: Connection with Weak derivative

$$D_h F(x) = DF(x)h.$$

Proof: Make sense of this....

$$||F(x + sh) - F(x) - Ash||_Y \leq \varepsilon ||sh||_X = |s| \cdot \varepsilon ||h||_X,$$

for all s with $||sh||_X \leq \delta$, i.e. for all small enough s. Hence for such s with s non-zero we have

$$||\frac{F(x + sh) - F(x)}{s} - Ah||_Y \leq \varepsilon ||h||_X.$$

The right-hand side may be arbitrarily small so this shows that

$$\lim_{s \to 0} \frac{F(x + sh) - F(x)}{s} = Ah = DF(x)h.$$

Comment: Blanket Assumptions

The definition of strong differentiability given above is modelled after the one occurring in ordinary calculus, except that the environment has widened from R, (or R^n, if you will) to a normed vector space context. Throughout, the methods we employ will require a background optimization theory that is developed by analogy with ordinary calculus. Strong differentiability is always assumed of the functionals etc.

Natural Notation for function evaluation. We need to clarify a common source of some notational confusion. We think of the points in R^n as column vectors. Nevertheless for F with domain R^n it is natural when evaluating F at the point $x = (x_1, ..., x_n)^T$ to write $F(x_1, ..., x_n)$ rather than the more correct but cumbersome $F((x_1, ..., x_n)^T)$. This natural convention calls for caution in interpreting notation, as the following example indicates

We begin with a cautionary example which will particularly help with the next examples. Both are of use to us later.

Cautionary Example. Let $F(x) = x^T b$. Since $b^T x = x^T b$ is $DF(x) = b$, or is it $DF(x) = b^T$ which?

Solution. Of course $DF(x)h = b^T h = h^T b$. So since $DF(x)$ is a matrix and acts on column vectors as arguments the answer has to be

$$DF(x) = b^T.$$

The answer is a row, and in any case the answer must agree for $F : R^n \to R$ with the notation

$$DF(x_1, ..., x_n) = (\frac{\partial F}{\partial x_1}, ..., \frac{\partial F}{\partial x_n}).$$

Moral Whenever computing derivatives of functions involving transposition, first rewrite the relevant term so as to eliminate transposition at the point of differentiation and then differentiate.

Example. Verify the product rule for differentiable functions $u, v : R^n \to R^n$, namely

$$D[u(x)^T v(x)] = v^T Du + u^T Dv.$$

Comment. We leave this as an exercise, but note the appearance of the first term which must contain Du since that is applied to column vectors (as is the left-hand side). If we follow the moral, we know to write $v^T u$ first and then differentiate u.

Example. Show that $DF(x) = 2x^T A$ when $F(x) = x^T Ax$ and A is symmetric.

Solution. Perhaps the simplest way to find and understand the answer is to compute from first principles, as follows.

$$\begin{aligned} F(x+h) - F(x) &= (x+h)^T A(x+h) - x^T Ax \\ &= h^T Ax + x^T Ah + h^T Ah \\ &= x^T A^T h + x^T Ah + h^T Ah \\ &= x^T Ah + x^T Ah + h^T Ah \\ &= 2x^T Ah + h^T Ah. \end{aligned}$$

So we identify the linear part as

$$DF(x) = 2x^T A.$$

We can of course use the product rule with $u(x) = x$ (where $Du(x) = I$) and $v(x) = Ax$ (where $Dv(x) = A$), and from there we have

$$DF(x) = (Ax)^T I + x^T A = 2x^T A.$$

2. Towards Pontryagin's Principle [23.10]

In an important class of optimisation problems one seeks to steer over the time interval $a \le t \le b$ the state $x(t)$ of a dynamical system from a given initial condition to some final state and the system is governed by a first-order differential equation; steering is by means of a control variable $u(t)$ and involves a cost which is to be minimized. We model the cost cumulatively, as the cost is incurred throughout the time interval. To simplify the problem we assume first of all that there are no bounds on the values $u(t)$ may take.

2.1. The unrestricted control problem

The assumption that $u(t)$ may take any value leads to the unrestricted control problem:

$$\min_{u(t)} \int_a^b f_0(x,u)dt$$

subject to

$$\begin{aligned} \dot{x} &= f_1(x,u), \\ x(a) &= x_0, x(b) = x_1. \end{aligned}$$

In applications it is usual to place bounds on u and we will return to this issue in a moment.

Acting on intuition one is tempted to write the Lagrangian in the form:

$$L(x,u,\lambda) = \int_a^b f_0(x,u)dt + \int_a^b [f_1(x,u) - \dot{x}]\lambda(t)dt.$$

on the grounds that this is a limiting sum of terms of the form 'instantaneous Lagrangian $\times$ instantaneous constraint', where at each instant the constraint is $\dot{x} = f_1(x,u)$. But for this limiting sum to exist the constraints would need to change continuously with time. However, there is a problem with the end-point constraints. A small change in u may make the solution to the differential equation with the given the initial condition $x(a) = x_1$ inconsistent with the terminal condition $t = b$. It may therefore be necessary to allow into the Lagrangian a term in addition to the integral so as to take care of a discontinuity in constraints at $t = b$;in other words it may be necessary to use Riemann-Stieltjes integration to included an extra contribution at $t = b$.

Technical point. The formula above is correct provided

(i) there is no constraint placed on $x(b)$; (i.e. $x(b) = x_1$ is omitted);

or,

(ii) in the presence of the constraint $x(b) = x_1$ the Lagrange multiplier need not be continuous at $t = b$.

We consider this particular point briefly in an exercise and we discuss the justifiability of the Lagrangian approach in the next section.

We return to the control problem. Writing

$$L(x,u,\lambda) = \int_a^b f(x,\dot{x},u,\lambda)dt,$$

where

$$f(x,\dot{x},u,\lambda) = [f_0(x,u) + \lambda f_1(x,u)] - \dot{x}\lambda,$$

and for convenience defining the Hamiltonian to consist of the first two terms:

$$H(x, u, \lambda) = f_0(x, u) + \lambda f_1(x, u),$$

the Euler-Lagrange equations read:

$$\begin{aligned} \frac{d}{dt}(f_{\dot{x}})\,0 &= \frac{\partial}{\partial x}(f_x) \text{ i.e. } \dot{\lambda} = -\frac{\partial H}{\partial x}, \\ \frac{d}{dt}(f_{\dot{\lambda}})\,0 &= \frac{\partial}{\partial u}(f_\lambda) \text{ i.e. } \dot{x} = \frac{\partial H}{\partial \lambda}, \\ \frac{d}{dt}(f_{\dot{u}})\,0 &= \frac{\partial}{\partial u}(f_u) \text{ i.e. } 0 = \frac{\partial H}{\partial u}. \end{aligned}$$

The first equation is known as the co-state equation. The second restates the governing differential equation (the state equation). The last equation requires H to be stationary; when u is bounded it is this condition which needs modification.

Example: Minimize

$$\int_0^1 (x^2 + u^2)dt$$

subject to $\dot{x} = -x + u$ and $x(0) = x_0, \quad x(T) = 0.$

Solution. Evidently $\mathcal{H} = x^2 + u^2 + \lambda(u - x)$.

Thus we have

$$\begin{aligned} -\dot{\lambda} &= \frac{\partial H}{\partial x} = 2x - \lambda, \\ 0 &= \frac{\partial H}{\partial u} = 2u + \lambda. \end{aligned}$$

and $\quad \dot{x} = -x + u.$

We thus have three equations in three unknowns. Now

$$\dot{\lambda} = -2(\dot{x} + x),$$

so $\quad \dot{\lambda} = -2(\dot{x} + \ddot{x})$, hence

$$\begin{aligned} 2(\dot{x} + \ddot{x}) &= 2x + 2(\dot{x} + x) \\ \dot{x} + \ddot{x} &= x + \dot{x} + x \end{aligned}$$

i.e. $\quad \ddot{x} = 2x.$

Roots of the auxiliary equation $t^2 - 2 = 0$ are $t = \pm\sqrt{2}$, so

$$x(t) = Ae^{t\sqrt{2}} + Be^{-t\sqrt{2}}.$$

But $x(0) = x_0$ and $x(T) = 0$, so we calculate as follows.

$$\begin{aligned} x_0 &= A + B \\ 0 &= Ae^{T\sqrt{2}} + (x_0 - A)e^{-T\sqrt{2}}, \\ 0 &= A(e^{T\sqrt{2}} - e^{-T\sqrt{2}}) + x_0 e^{-T\sqrt{2}}. \end{aligned}$$

so that

$$A = -x_0 \frac{e^{-T\sqrt{2}}}{e^{T\sqrt{2}} - e^{-T\sqrt{2}}}$$

Hence

$$B = x_0\left(1 + \frac{e^{-T\sqrt{2}}}{e^{T\sqrt{2}} - e^{-T\sqrt{2}}}\right) = x_0\frac{e^{T\sqrt{2}}}{e^{T\sqrt{2}} - e^{-T\sqrt{2}}}$$

so

$$\begin{aligned} x(t) &= -x_0\frac{e^{-T\sqrt{2}}}{e^{T\sqrt{2}} - e^{-T\sqrt{2}}}e^{t\sqrt{2}} + x_0\frac{e^{T\sqrt{2}}}{e^{T\sqrt{2}} - e^{-T\sqrt{2}}}e^{-t\sqrt{2}} \\ &= x_0\frac{e^{(T-t)\sqrt{2}}}{e^{T\sqrt{2}} - e^{-T\sqrt{2}}} - x_0\frac{e^{-(t-T)\sqrt{2}}}{e^{T\sqrt{2}} - e^{-T\sqrt{2}}} \end{aligned}$$

or

$$x(t) = x_0\frac{\sinh(T-t)\sqrt{2}}{\sinh\sqrt{2}T}, \qquad 0 \le t \le T.$$

Remark. This form of the state could have been anticipated from the condition $x(T) = 0$. It would have been more natural to consider the general solution format

$$x(t) = K\cosh(T-t)\sqrt{2} + L\sinh(T-t)\sqrt{2}.$$

Turning now to the form of the control we note that we can recover it from the state equation.

$$\begin{aligned} u(t) &= \dot{x} + x \\ &= x_0\frac{\sinh(T-t)\sqrt{2}}{\sinh\sqrt{2}T} - x_0\sqrt{2}\frac{\cosh(T-t)\sqrt{2}}{\sinh\sqrt{2}T}, \qquad 0 \le t \le T. \end{aligned}$$

2.2. Relative stationarity and Lagrange multipliers

We review the problem of **constrained optimization**. The general framework is this

$$\text{maximize } F(x)$$

(where $F : \mathcal{X} \to \mathbf{R}$ and $\mathcal{X}$ may typically be a space of functions),

$$\text{subject to } G(x) = o,$$

where $G : \mathcal{X} \to \mathcal{Y}$ and o refers to the zero vector of $\mathcal{Y}$.

Before we embark on some theory let us see some examples. We take the opportunity to note that there may be more than one way to formulate a constrained optimization; what is state and what is control may be fluid; what is the space of functions can also vary across formulations. This choice may have several implications.. Whilst the precise formulation may often not be critical to a solution method, the choices may nevertheless on occasion strike at the heart of a correct identification of the optimal trajectory. We shall see one such situation arising from our final example (compare section ?? for implications).

Example 1 The *isoperimetric problem* takes this format when we let:

$$F(x) = \int_{-1}^{1} x(t)dt$$

and

$$G(x) = \int_{-1}^{1} \sqrt{1+\dot{x}^2}dt - l = 0,$$

so that G: Differentiable functions $\rightarrow$ **R**.

Contrast this with the formuation

$$\begin{aligned} G(x,y)(s) &= y(s) - \int_{-1}^{s} \sqrt{1+\dot{x}^2}dt = 0, \\ y(-1) &= 0, y(1) = l \end{aligned}$$

so that G: Differentiable functions$\times$Continuous functions $\rightarrow$ Continuous functions.Here there is one constraint per moment of time. More properly this should be stated as:

$$G(x,y)(s) := (y(-1), y(1) - l, y(s) - \int_{-1}^{s} \sqrt{1+\dot{x}^2}dt) = (0, 0, o(s)),$$

where $o(.)$ denotes the zero function.

Example 2 Minimise $\quad F(x,y) = \int_0^2 \frac{1}{2}\dot{y}^2 dt$ subject to:

$$G(x,y) = y - \dot{x} = 0$$

i.e. $\quad G(x,y)$ is the function $y(t) - \dot{x}(t)$ and so F: Differentiable functions $\times$ Differentiable functions $\rightarrow$ **R** whence also G: Differentiable $\times$ Differentiable $\rightarrow$ Continuous functions.

Example 3. Minimize

$$F(u) = \frac{1}{2}\int_0^1 (x^2 + 2xu + u^2)dt$$

where x is defined by the state equation

$$\dot{x}(t) = f(x(t), u(t)), x(0) = x_0.$$

The relation between x and u may be regarded as a constraint. The constraint may be reformulated as

$$G(x,u)(t) = x(t) - x_0 - \int_0^t f(x(s), u(s))ds = o(t).$$

The latter formulation has the advantage of not referring to $\dot{x}$.so the domain of G here refers only to continuous functions. How useful this is will become apparent at the end of this section.

Example. Minimize

$$J = \frac{1}{2}\int_0^{3\pi/2} (x_2^2 - x_1^2)dt$$

s.t.

$$\begin{aligned} \dot{x}_1 &= x_2, \\ \dot{x}_2 &= u, \end{aligned}$$

with $|u| \leq 1$ and $x_1(0) = 0$ and $x_2(0) = 1$.

Without loss of generality we may put $u = \dot{v}$. Thus in effect $v = x_2$. which converts our view of x_2 from being a state component to being a control variable. Writing $x = x_1$ we obtain the formulation:

$$\tilde{J} = \frac{1}{2}\int_0^{3\pi/2} (v^2 - x^2)dt$$

so that

$$\dot{x} = v$$

with $x(0) = 0$.In this formulation only x is the state and v (previously the state x_2) is the control variable. Not only has some simplification occurred but also the Hamiltonian

$$\frac{1}{2}(v^2 - x^2) + \lambda v$$

has become quadratic rather than linear in the control. Evidently we have the constraint $|\dot{v}| \leq 1$.Here we have $v = -\lambda$ and

$$\dot{\lambda} = -H_x = x.$$

sos

$$\ddot{\lambda} = \dot{x} = v = -\lambda.$$

With hindsight write $\lambda = -A\cos(t + \varepsilon)$. Hence $v = -\lambda = A\cos(t + \varepsilon)$ and we therefore have $|A| \leq 1$.But $1 = \dot{x}(0) = v(0) = A\cos\varepsilon$, so that $A = \pm 1$ and respectively $\cos\varepsilon = \pm 1$. Since $\sin\varepsilon = 0$ we conclude that

$$x = A\sin(t + \varepsilon) = \sin t.$$

Theorem (Deep, though plausible) *If x is optimal for the problem of maximising F subject to $G = 0$, then*

$$D_hG(x) = 0 \quad \Rightarrow D_hF(x) = 0.$$

(The theorem assumes strong differentiability of F and G and some extra regularity properties on G akin to being of 'full rank' - we will see why.)

This is a condition of **relative stationarity**. Here is a plausiblity argument. Consider any increment h, then so long as $x + sh$ satisfies the constraint $G(x + sh) = 0$ we are safe in asserting that $\psi(s) = F(x + sh)$ has a local maximum at $s = 0$. This is a far cry from saying straight off that $D_hF(x) = 0$. However, if $D_hG(x) = 0$ then we can argue that:

$$G(x + sh) = G(x) + DG(x)sh + e(sh) = 0 + 0 + e(sh),$$

where by differentiability the error term $e(sh)$ is $o(|s|)$. (Indeed, $e(k)$ has to be $o(||k||)$, so for fixed h, we readily see that $e(sh)$ is $o(|s|)$).

Thus for practical purposes, for all small enough s the right-hand side is zero, and then $x + sh$ satisfies (well, nearly so) the constraint equation, and so $F(x + sh)$ decreases as we move away from $s = 0$ in either direction. This isn't a bad argument, but it requires quite a lot of machinery to make water-tight, in particular it requires knowing that a correction to $x + sh$ when $D_hG(x) = 0$ can be made so that the corrected term does indeed satisfy the constraint equation.At this point one requires something like local invertibility of the transformation $DG(x)$. See for instance Luenberger.

Deduction We really need to write our relative stationarity condition in the equivalent format that:

$$DG(x)h = 0 \Rightarrow DF(x)h = 0$$

and this says that $DF(x) \perp h$ whenever $h \in \mathcal{N}(DG(x))$, i.e. $DF(x) \in \mathcal{N}(DG(x))^{\perp}$. At t his point one is reminded of the duality result in Euclidean spaces, namely that $\mathcal{R}(A^{\top}) = \mathcal{N}(A)^{\perp}$ where A is a matrix (see section??). One hopes for a generalization. Indeed one is available, subject to technical qualifications. It does however call for the notion of transpose to be extended beyond the Euclidean spaces R^n to the wider context of vector spaces.

Let us see how this is done.

To begin with one must recognize a geometric insight into the connection between transposition from column vector x to row vector and to draw on the basic inspiration of the geometric duality in R^2 between points and lines, that transposition of the word point for the word line enables theorems about lines and points to be translated into valid theorems about points and lines. This translation needs sensitive editing of the English and of the Mathematics. To see this at work at its simplest observe how the theorem that two distinct points determine a unique line translates into an assertion about two lines determining a point: two lines with distinct normal direction vectors (i.e. non-parallel lines) determine a unique point as their intersection. This connection is made possible by reading an equation like

$$ax + by = 0$$

as asserting either that the line identified by the row/normal (a, b) contains the column/point $(x, y)^T$, or that the column/point $(x, y)^T$ lies on the line identified by a row/normal (a, b).

The tradition of translating between dual statements by transposition back and forth between rows and columns albeit helpful in identifying the dual objects of line and point obscures the essence of duality as residing in two interpretions of an equation involving real numbers that arise when combining the two vectors (a, b) and (x, y) via the evaluation of an inner product.

We have already that a real number arises from the combination of $DF(x)$ from $\mathcal{X}^*$ and h from $\mathcal{X}$ to form $DF(x)h$. Duality comes into its own if we interpret the real number

$$x^*(x),$$

arising from the evaluation of the linear transformation x^* in $\mathcal{X}^*$ at a point x of $\mathcal{X}$ as though it was the inner product

$$< x^*, x > .$$

This does mean that lines/normals are represented by members of $\mathcal{X}^*$ when points come from $\mathcal{X}$. This implies that if $S \subset \mathcal{X}$ we may directly define

$$S^{\perp} = \{x^* : x^*(s) = 0 \text{ for all } s \text{ in } S\}$$

as the set of normals orthogonal to the set S. Note that $S^{\perp} \subset \mathcal{X}^*$ so $S^{\perp\perp} \subset \mathcal{X}^{**}$ but $S^{\perp\perp} = S^{\perp}$ for S a subspace only if $\mathcal{X}^* = \mathcal{X}^{**}$; but it is not in general the case that a space is reflexive, that is $\mathcal{X}^* = \mathcal{X}^{**}$.

Continuing in this vein, if $A : \mathcal{X} \to \mathcal{Y}$ is a continuous linear transformation between normed vector spaces, its adjoint A^* may be defined as a continuous linear transformation from the dual space $\mathcal{Y}^*$ to $\mathcal{X}^*$ by the identity:

$$A^*y^*(x) = y^*(Ax).$$

This is in line with thinking of A^T as a transformation between rows, one need only restate the equation $z = Ay$ as

$$z^T = y^T A^T.$$

It is routine to check that the proof in section ?? of $\mathcal{R}(A)^\perp = \mathcal{N}(A^T)$ readily translates into $\mathcal{R}(A)^\perp = \mathcal{N}(A^*)$. One can just as easily show that $\mathcal{R}(A^*) \subseteq \mathcal{N}(A)^\perp$. However, the converse assertion that $\mathcal{N}(A)^\perp \subseteq \mathcal{R}(A^*)$ requires a degree of mathematical sophistication that takes us beyond this textbook's remit.

Returning to the main theme of tthis section, we note that the relative stationarity condition is equivalent to $DF(x) \in \mathcal{N}(DG(x))^\perp = \mathcal{R}(DG(x)^*)$, so $DF(x) = DG(x)^*\mu^*$ for some μ^*. Since $DG(x) : \mathcal{X} \to \mathcal{Y}$ we have $DG(x)^\top : \mathcal{Y}^* \to \mathcal{X}^*$ and

$$\begin{aligned} DF(x) &= DG(x)^*\mu^*(h) \text{ for all } h \\ &= \mu^*(DG(x)h) \\ &= \mu^* DG(x)h. \end{aligned}$$

The last line uses the usual notation for functions applied in composition (as with matrices $AB\mathbf{z} \equiv A(B\mathbf{z})$). Thus

$$[DF(x) - \mu^* DG(x)](h) = 0 \quad \text{for all } h.$$

Indeed since μ^* is continuous we have by Question 2 Exercises 2.7 that $D\mu^*G(x)h = \mu^*(DG(x))h$. Thus if $L(x) = F(x) + \lambda^*G(x)$ with $\lambda^* = -\mu^*$, we have $DL(x) = 0$, giving us the:

Lagrange Multiplier Theorem.

If x is the optimal curve, then for some λ^* the Lagrangian

$$F(x) + \lambda^*G(x)$$

is stationary in x.

The Lagrange multiplier evidently lies in $\mathcal{Y}^*$ when $G : \mathcal{X} \to \mathcal{Y}$ describes the constraint.

The theorem relegates the exercise of constrained optimization to the identification of the dual space $\mathcal{X}^*$. For our purposes it suffices to know one result: any continuous linear functional x^* with domain the space $\mathcal{X} = \mathcal{C}[a, b]$ may be represented by an integrator in the Riemann Stieltjes integral format

$$x^*(x) = \int_a^b x(t)d\lambda(t),$$

where the integrator $\lambda(t)$ is a function of bounded variation (see Section ??).

3. Restricted control Problem: Pontryagin's Principle [23.11]

Theorem (Pontryagin's Principle) For the bounded control problem:

$$\min_{u(t)} \int_a^b f_0(x, u)dt,$$

subject to

$$\begin{aligned} \dot{x} &= f_1(x, u), \\ x(a) &= x_0, x(b) = x_1, \\ A &\leq u \leq B, \end{aligned}$$

let $u(t)$ be the optimal trajectory, let $x(t)$ be the corresponding trajectory satisfying the state equation and let $\lambda(t)$ satisfy the co-state equation corresponding to these two trajectories

$$\dot{\lambda} = -\frac{\partial}{\partial x}\mathcal{H}(x(t), u(t), \lambda(t)), \text{ with } \lambda(b) = 0,$$

then, for each t, if v solves

$$\mathcal{H}(x(t), v, \lambda(t)) = \min_{A \le w \le B} \mathcal{H}(x(t), w, \lambda(t)),$$

then $u(t) = v$, i.e. the optimal control minimizes the Hamiltonian on the optimal trajectory.

Remark 1. Notice that if only $x(a)$ is specified then the solution of the differential equations governing the vector (x, λ) is a two-point boundary problem: boundary specification of the x component is at $t = a$ and of the λ component at $t = b$.This technicxal difficulty can be resolved quite easily for H a quadratic in (x, u) in the case when the problem is (strictly) non-singular, i.e. when $H_{uu} > 0$ along the optimal path.See section ??.

Remark 2. How is the theorem to be re-stated if one is to solve the control problem with maximization of the objective? The answer is remarkably simple: the theorem now state that the optimal control maximizes the Hamiltonian on the optimal trajectory. The reason is this. Obviously one replaces the given objective function of the control problem namely f_0 by its negative $-f_0$ (to turn a minimum into a maximum). Now replace λ by $-\lambda$ then the Hamiltonan becomes

$$-f_0 - \lambda f_1,$$

and it follows that the optimal control must maximize $f_0 + \lambda f_1$. Note that the co-state equation remains unchanged since

$$-\dot{\lambda} = -\frac{\partial}{\partial x}\left(-f_0 - \lambda f_1\right),$$

is the same as

$$\dot{\lambda} = -\frac{\partial}{\partial x}\left(f_0 + \lambda f_1\right).$$

Remark 3. The blanket assumption noted in Section 23.9 here asserts that the functionals $I(x, u) = \int_a^b f_0(x, u)dt$ and $J(x, u) = \int_a^t f_1(x(s), u(s))ds$ are strongly differentiable.

3.1. An optimal Investment Problem

This problem has $H_{uu} = 0$ and part of the technical difficulty arises from the two-point boundary nature of the Maximum Principle.

Example (After Luenberger) A farmer produces wheat at a (flow) rate $x(t)$. He may store wheat or sell it and re-invest the proceeds thereby increasing his rate $x(t)$. Let $u(t)$ denote the fraction of the rate $x(t)$ which is re-invested. Suppose that

$$\begin{aligned} \dot{x} &= u(t)x(t) \\ x(0) &> 0 \end{aligned}$$

and that the farmer seeks to maximize the amount held in store at time $t = T$ viz.

$$\int_0^T (1 - u(t))x(t)dt.$$

What is the optimal control $u(t)$?

Remark. We implicitly have $0 \leq u(t) \leq 1$ for all t. Note that the (flow) rate is interpreted to mean that in the time interval $[t, t+\delta t]$ the amount of wheat produced is $x(t)\delta t$. Assuming unit price he has $u(t)x(t)\delta t$ available to re-invest. This is supposed to improve his rate in the next time interval by an amount δx. Assuming a law of proportionality between improvement and investment (with the constant of proportionality also set equal to 1) gives

$$\delta x = u(t)x(t)\delta t.$$

Hence $dx/dt = u(t)x(t)$. Similarly the stock $= \sum(1-u)x\delta t$.

Solution. The Hamiltonian is

$$\mathcal{H} = (1-u)x + \lambda x u$$

and the co-state equation reads

$$\dot{\lambda} = -\frac{\partial \mathcal{H}}{\partial x} = (1-u) - \lambda u \text{ with } \lambda(T) = 0.$$

Along the optimal trajectory

$$\mathcal{H} = x \ + \ u\{\lambda x - x\} = x + ux(\lambda - 1)$$

and the Hamiltonian is maximized as follows. We expect that $x(t) > 0$, so assuming this is true we have for each time t

$$u = \begin{cases} 1 & \text{if } \lambda > 1, \\ 0 & \text{if } \lambda < 1. \end{cases}$$

We note that we can solve

$$\dot{x} - ux = 0.$$

The integrating factor is $exp(-\int u(t)dt)$, so

$$\frac{d}{dt}(xe^{-\int u}) = 0.$$

Hence $x\exp(-\int u) = \text{const} = \ A$, or $x(t) = Ae^{\int_0^t u(\tau)d\tau}$. To find A, put $t = 0$. Thus

$$x_0 = Ae^0 = A.$$

Hence $x(t) = x_0 \exp\{\int_0^t u(\tau)d\tau\}$,and this remains positive as expected.

We can also set about solving the co-state equation

$$\dot{\lambda} + \lambda u = -1 + u, \text{ with } \lambda(T) = 0$$

near $t = T$. For since $\lambda(T) = 0$, we evidently have $u(T) = 0$. Moreover λ is continuous at $t = T$, since $x(T)$ is unspecified. So near $t = T$, by continuity, $\lambda(t) < 1$ hence $u(t) = 0$ near $t = T$. Thus near $t = T$, $\dot{\lambda} = -1$, i.e. $\lambda(t) = -t \ + \ $ const. But $\lambda(T) = 0$ so we have $0 = -T + \text{const.}$ i.e. $\lambda(t) = T - t$.

Now the interval in which $\lambda < 1$ may be computed, as

$$T - t = \lambda(t) < 1 \quad \Leftrightarrow \quad T - 1 < t.$$

Thus on $(T-1, T)$ we have $u = 0$ and $\lambda(t) = T - t$.

It is possible to continue to work backwards in time by reference to the intervals over which $u = 0$ and $u = 1$. If $u = 0$ to the left of $T - 1$ then $\dot{\lambda} = -1$ and so $\lambda(t) > 1$ for $t < T - 1$, contradicting $u = 0$. Thus, in fact, to the left of $T - 1$ we must have $u = 1$. Whereupon we have

$$\dot{\lambda} + \lambda = 0, with\ \lambda(T-1) = 1.$$

This may now be solved using integrating factors. As

$$\frac{d}{dt}(\lambda(t)e^t) = 0,$$

we have

$$\begin{aligned} \lambda(t) \quad &= Be^{-t}, \\ \lambda(T-1) \quad &= 1 = e^{-(T-1)}B, \end{aligned}$$

Thus $\lambda(t) = e^{(T-1-t)}$.

Running backwards in time we see that λ increases and so stays above 1 and hence $u = 1$ remains true.

Figure 3. *Graph of* $\lambda(t)$.

Actually, a more elegant way to proceed is so solve

$$\dot{\lambda} + u\lambda = u - 1$$

formally, using integrating factors. We have I.F. $= exp\{\int_0^t u(\tau)d\tau\}$ and so

$$\frac{d}{dt}\left[\lambda(t)e^{\int_0^t u(\tau)d\tau}\right] = (u-1)e^{\int_0^t u(\tau)d\tau}.$$

But $exp\{\int_0^t u(\tau)d\tau\}$ is non-decreasing (as $u \geq 0$) and since

$$u - 1 \leq 0$$

it must be that λ is non-increasing. Suppose τ is the largest time such that $\lambda(\tau) = 1$. Similarly let σ be the least time such that $\lambda(\sigma) = 1$. Then on (τ, T) we have $u = 0$ as $\lambda < 1$ and on $(0, \sigma)$ we have $u = 1$ since $\lambda > 1$. Then as before on (τ, T) we have

$$\lambda(t) = T - t.$$

Finally $1 = \lambda(\tau) = T - \tau$ so $\tau = T - 1$. For $t < \sigma$ evidently we have, as already computed,

$$\lambda(t) = e^{(\sigma - t)}.$$

Note that both (1) and (2) yield $\dot\lambda = -1$ at $t = \tau$ and $t = \sigma$ respectively. Since λ is non-decreasing we will have on $[\sigma, \tau]$ that $\lambda(t) = 1$. Suppose that $\sigma < \tau$ then in (σ, τ) we have $\dot\lambda = 0$ and the co-state equation

$$\dot\lambda + u\lambda = u - 1$$

now reads

$$u = u - 1,$$

a contradiction. Hence after all $\sigma = \tau$. This completes the analysis of the λ curve.

Remark. As the Hamiltonian is linear in u the optimal control is at one or other endpoint of the control interval just so long as the coefficient of u is non-zero. In the current problem we were able to show that this coefficient is zero for only a single moment (at most) and therefore the contribution to the objective over the time when the coefficient is zero is itself evidently zero. That is to say, the fact that optimization of the Hamiltonian offers no information as to the value of u when its coeffcent is zero turns out to be irrelevant to the objective. The example in the next section shares this property. We give later an example in which the coefficient of u is zero on a non-degenerate interval; the optimal control is said to have a singular arc.

4. Time optimality [23.12]

The Pontryagin Principle can be applied when $f_0 = 1$ so that we minimize

$$\int_0^T dt = T,$$

the time taken to steer the dynamical system from a starting point to givena terminal point. We show this on an example which typifies what can happen in the general multi-variable linear dynamical system of the form

$$\dot x = Ax + ub,$$

with A an arbitrary matrix. The reader is invited to consider similar examples in the exercises.

Example. Steer the dynamical system

$$x'' + 3x' + 2x = u.$$

from an arbitrary initial state to the origin when $|u| \le 1$.

We can rewrite the state equation in first-order form provided we take $x_1 = x$ and $x_2 = x'$. Now the first-order formulation with $x_1 = x, x_2 = x'$ is

$$\begin{aligned} \dot x_1 &= x' = f_1(x_1, x_2, u) = x_2, \\ \dot x_2 &= x'' = f_2(x_1, x_2, u) = -(2x_1 + 3x_2) + u, \end{aligned}$$

with $|u| \le 1$. In this case we therefore have

$$\begin{bmatrix} \dot x_1 \\ \dot x_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -2 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + u \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \tag{4.1}$$

Thus the Hamiltonian is

$$H = 1 + \lambda_1 x_2 + \lambda_2(-2x_1 - 3x_2 + u)$$

and this is minimised if

$$\begin{aligned} u &= +1, \text{ for } \lambda_2 < 0, \\ &= -1, \text{ for } \lambda_2 > 0. \end{aligned}$$

To determine the behaviour of the co-state variable we write down the co-state equations:

$$-\dot{\lambda}_1 = \frac{\partial H}{\partial x_1}, -\dot{\lambda}_2 = \frac{\partial H}{\partial x_2}$$

which in this case are

$$-\dot{\lambda}_1 = -2\lambda_2, -\dot{\lambda}_2 = \lambda_1 - 3\lambda_2.$$

Notice that

$$\begin{bmatrix} \dot{\lambda}_1 \\ \dot{\lambda}_2 \end{bmatrix} = \begin{bmatrix} 0 & 2 \\ -1 & 3 \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \tag{4.2}$$

and that the costate coefficient matrix is the transpose of the state coefficient matrix. Now

$$\lambda_2'' - 3\lambda_2' + 2\lambda_2 = 0.$$

The auxiliary equation for this is

$$m^2 - 3m + 2 = 0$$

or

$$(m-1)(m-2) = 0$$

with roots $1, 2$ which are of like sign (and are the eigenvalues of the costate matrix of (4.2). Now

$$\begin{aligned} \lambda_2 &= Ae^t + Be^{2t} \\ &= Be^{2t}\{1 + \frac{B}{A}e^{-t}\}, \end{aligned}$$

so λ_2 changes sign at most once. The coefficient of u in the Hamiltonian is thus zero for only one instant of time at most. Compare the concluding Remark of section ??.

In summary the optimal control is one of the following:
i) $u = 1$ throughout the solution;
ii) $u = -1$ throughout the solution;
iii) $u = 1$ initially followed by $u = -1$;
iv) $u = -1$ initially followed by $u = 1$.

In the last two cases we say that the control *switches* once; since the control takes two extreme values it is said to be of 'bang-bang' type (as in 'full steam ahead' followed by 'full steam reverse').

To understand the implications of this observation we must study the general trajectories of constant control $u = u^* = \pm 1$ and in particular those which reach the origin. Now

$$\begin{aligned} x_1' &= x_2, \\ x_2' &= -(2x_1 + 3x_2) + u^*. \end{aligned}$$

so for this control there is a *rest* point, also known as a singular point, defined by the conditions $x_1' \equiv 0, x_2' \equiv 0$, which is given by

$$x_2 = 0, x_1 = u^*/2.$$

We change origin to the rest point taking new co-ordinates $y_1 = x_1 - u^*/2, y_2 = x_2$; then we have

$$\begin{aligned} y_1' &= y_2, \\ y_2' &= -(2y_1 + 3y_2), \end{aligned}$$

so that

$$y_1'' + 3y_1' + 2y_1 = 0.$$

The auxiliary equation here is

$$m^2 + 3m + 2 = 0$$

or

$$(m+2)(m+1) = 0$$

with roots which are negatives of those found in the auxiliary equation for the co-state variable's. They are $-2, -1$ and are of course eigenvalues of

$$\begin{bmatrix} 0 & 1 \\ -2 & 3 \end{bmatrix}.$$

Here the eigenvalues are of like sign and

$$y_1 = Ke^{-t} + Le^{-2t}$$

so that

$$y_2 = y_1' = -Ke^{-t} - 2Le^{-2t}.$$

It follows that we obtain two linear trajectories in the (y_1, y_2)-space, one when $L = 0$ and one when $K = 0$. We have:

$$\frac{y_2}{y_1} = \frac{y_1'}{y_1} = \frac{-Ke^{-t} - 2Le^{-2t}}{Ke^{-t} + Le^{-2t}} = -1, \text{ for } L = 0$$

this one having state approaching the singular point as $t \to +\infty$ since $|K|e^{-t} \to 0$ (and the rest-state is said to be 'attracting'). Similarly

$$\frac{y_2}{y_1} = -2, \text{ for } K = 0$$

is the other linear trajectory which also corresponds to the state approaching the singular point (attracting state). As for the general non-linear trajectory, noting the effect of the change of variable

$$\begin{aligned} Y_1 &= -y_1 - y_2 = Le^{-2t}, \\ Y_2 &= 2y_1 + y_2 = Ke^{-t}, \end{aligned}$$

allows us to write

$$Y_2^2 = (Ke^{-t})^2 = (K^2e^{-2t}) = Y_1 \cdot (K^2/L)$$

which is thus a parabola (relative to either co-ordinate system). The general trajectory approaches its the singular point (i.e. for $t \to \infty$) with gradient

$$\frac{y_2}{y_1} = \frac{-Ke^{-t} - 2Le^{-2t}}{Ke^{-t} + Le^{-2t}} = \frac{-K - 2Le^{-t}}{K + Le^{-t}} \to -1, \text{ for } t \to \infty, \text{unless } K = 0, \text{when the gradient is } -2$$

The equation of the linear trajectory to which all other trajectories (except the second linear trajectory) are tangential is thus

$$x_1 - (\pm\frac{1}{2}) = -x_2.$$

Switching curve

Choosing for each singular point $(\pm\frac{1}{2}, 0)$ a trajectory passing through the origin in the direction towards the singular point we obtain two curves meeting at the origin which give the locus of initial states from which the system may be steered with a constant control of $u = 1$ or $u = -1$.

The two curves just indicated give what is called the *switching curve*. For a general initial starting state in the plane we therefore need to use two opposite control values. When using initially $u = +1$ the system follows a general trajectory until it interesects the switching curve, and then follows the $u = -1$ section of the switching curve to reach the origin. Similarly, When using initially $u = -1$ the system follows a general trajectory until it interesects the switching curve, and then follows the $u = +1$ section of the switching curve to reach the origin.

5. Example of a singular arc [23.13]

This first example is taken from Bell&Jacobson (p. 58) and is an example to sharpen your awareness of complications that can arise when the optimal control might not take an extreme value over a non-degenerate interval of time.

Here again we have $H_{uu} = 0$ and again we face a technical difficulty arising from the two-point boundary nature of the Maximum Principle. In the current case our task is a much simpler one.

Solve the problem

$$\min \ J(x) = \frac{1}{2}\int_0^2 x^2(t)dt,$$

subject to

$$\dot{x} = u, \text{ with } |u| \leq 1,$$

and

$$x(0) = 1.$$

Solution. The Hamiltonian is

$$H = f_0 + \lambda f_1 = \frac{1}{2}x^2 + \lambda u,$$

so that H is minimized by setting

$$u = \begin{cases} +1 & \text{if } \lambda < 0, \\ ? & \text{if } \lambda = 0, \\ -1 & \text{if } \lambda > 0. \end{cases}$$

To settle the middle case in the situation where $\lambda = 0$ on a non-degenerate time interval, note that the co-state equation asserts that

$$-\dot{\lambda} = \frac{\partial H}{\partial x} = x,$$

so in the interior of the time interval when $\lambda = 0$ we must have $x = 0$. That via the state equation in turn tells us that $u = 0$. Moreover the contribution over the singular arc to the objective is zero.

We note that outside of the singular arc we have $u = u^*$ with $u^* = \pm 1$ so that

$$x = u^* t + const.$$

Since we are given $x(0) = 1$ the optimal curve here comprise the arc

$$x = 1 - t \text{ for } 0 \leq t \leq 1,$$

and the singular arc $x = 0$ for $1 \leq t \leq 2$.