Continuous-time optimal control—indirect approach

Introduction to calculus of variations; LQ-optimal control

Zdeněk Hurák
March 24, 2021

In this lecture we will go on with optimal control design, but this time we will step into the continuous-time domain. Our philosophy will remain the same as in the discrete-time case. First we will aim at finding the first-order necessary conditions of optimality. We will see that these come in the form of differential equations with boundary constraints—two-point boundary value problem (BVP). The solution functions that satisfy these conditions are called *extremals*. We will then study under which conditions a given extremal is actually minimizing the cost function.

Since in the continuous-time setting the optimization is performed over functions, the mathematical formalism of *differential calculus* is no longer appropriate here. Instead, we need to invoke its extension for optimization over functions—the *calculus of variations*. In the first part of the lecture we will ask what will be the appropriate replacement for the derivatives and differentials and we will see that *variational derivatives* and *variations* constitute such replacement. We will then show how these could be applied within the context of continuous-time optimal control. In this lecture we will only consider the optimal control on a fixed time interval (or time horizon).

Finally, we will show a solution to both the fixed- and free-final state LQ-optimal control. The former leads to open-loop control, the latter leads to time-varying linear state feedback, which can be computed by solving *differential Riccati equation*. Similarly as in the discrete-time case, one can aim at finding a constant state-feedback, which would be suboptimal on a finite time interval, but optimal on an infinite one. Such state feedback gains can be either determined from the limiting (steady-state) solution to differential Riccati equation, or as one of the solutions of *continous-time algebraic Riccati equation* (CARE). Identical questions pop up as in the discrete-time case: under which conditions is there a stabilizing solution to the infinite-horizon LQ-optimal control problem and can it be found by solving the CARE? The answers will be pretty much identicle as in the discrete-time case.

This is the problem we are going to solve in this lecture

$$\min_{\mathbf{x}(t),\mathbf{u}(t)} \left[\phi(\mathbf{x}(t_\mathrm{f})) + \int_{t_\mathrm{i}}^{t_\mathrm{f}} L(\mathbf{x}(t),\mathbf{u}(t),t)\ \mathrm{d}t \right], \tag{1}$$

subject to

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x},\mathbf{u},t) \quad \mathbf{x}(t_\mathrm{i}) = \mathrm{r}_\mathrm{i}, \tag{2}$$

where the meaning of the latter constraint is that the state at the initial time is given.

Unlike in the discrete-time case, here we are optimizing not over the n-tuples of real numbers (vectors or finite sequences) but we are optimizing over functions. This

is emphasized by means of Fig.1, which show instances of such $x(t)$ and $u(t)$ functions/trajectories/curves.

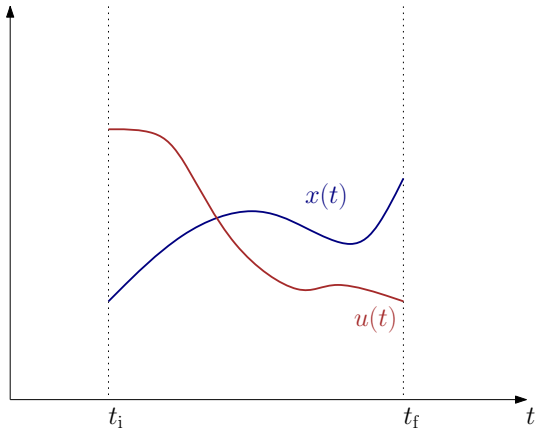

Figure 1: Two instances of state and input trajectories $x(t)$ and $u(t)$, respectively.

To make the assignment of the optimal control design task complete, we have to specify in which space of functions we perform the search for the optimal "values". In this introductory lecture we will consider smooth functions, that is, continuous function for which their first derivative exists. We will discuss possible needs for extensions when they are needed.

Besides building a framework for a general problem, our ultimate goal is to find a solution to the popular LQ problem

$$\min_{\mathbf{x}(t),\mathbf{u}(t)} \left[\mathbf{x}^{\mathrm{T}}(t_\mathbf{f})\mathbf{S_f}\mathbf{x}(t_\mathbf{f}) + \int_0^{t_\mathbf{f}} \left(\mathbf{x}^{\mathrm{T}}(t)\mathbf{Q}\mathbf{x}(t) + \mathbf{u}^{\mathrm{T}}(t)\mathbf{R}\mathbf{u}(t) \right) \mathrm{d}t \right], \tag{3}$$

subject to

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t), \quad \mathbf{x}(0) = \mathbf{r}_0. \tag{4}$$

We will provide such solutions for both the fixed and free final state, including a discussion of the prolongation of the control horizont to infinity.

1 Calculus of variations

Since in our quest for the optimal control we optimize over functions, we can view our optimization as running in an infinite-dimensional vector space. The mathematical discipline of *calculus of variations* provides concepts and tools for such optimization. The general task (in a scalar version) is to find

$$\min_{y(x)\in\mathcal{C}^1[a,b]} J(y(x)), \tag{5}$$

where we relabelled the variables in the following sense: the optimization is performed over $y(x)$, which is a function of the independent scalar variable x. The reason for this choice is that many of the results in calculus of variations were motivated by problems where the independent variable was length or position. This change of notation is

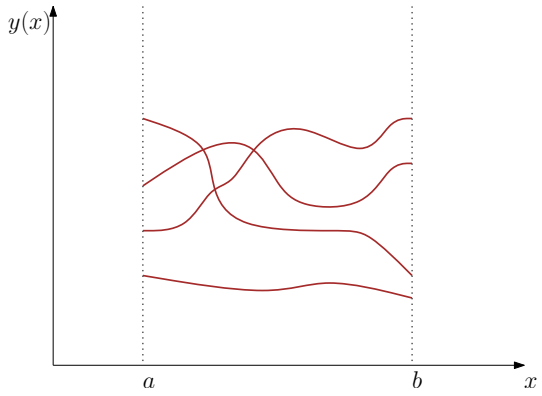

Figure 2: The variables in calculus of variation are functions

shown in Fig.2, which shows a few members of the space in which we search for a minimizer.

The optimization criterion is now a function of a function. An established name for such function is *functional*[1] The optimization is conducted over the set $\mathcal{C}^1[a,b]$ of continuously differentiable (that is, smooth) functions.

It is now important to reinvoke the very definition of *local minimum* that we introduced in the lecture on finite-dimensional optimization. The cost function has a local minimum at a given point if there exists some neighbourhood within which all the other points achieve equal or higher value. With our current notation, J attains a local minimum at y^* if

$$J(y^*) \leq J(y) \tag{6}$$

for all y in some *neighbourhood* of y^*. The neighbourhood is given as a set of all those y for which

$$\|y - y^*\| \leq \epsilon. \tag{7}$$

1.1 Strong vs. weak minimum

The question is, which norm is used in the expression above. In finite-dimensional spaces the choice of a norm did not have an impact on whether a given point was classified as a minimum or not. We could use 2-norm (the popular Euclidean norm), 1-norm (also called Manhattan norm) or ∞-norm (also called max norm). But the situation is dramatically different in infinite dimensional vector spaces; and the spaces of function can be viewed as having an infinite dimension. There are two main norms that are popular in calculus of variations. First, the so-called 0-norm, which is defined as

$$\|y\|_0 = \max_{x\in[a,b]} |y(x)|. \tag{8}$$

The second type of norm that we will use is called 1-norm

$$\|y\|_1 = \max_{x\in[a,b]} |y(x)| + \max_{x\in[a,b]} |y'(x)|. \tag{9}$$

[1]Note that in modern mathematics, the term functional can be applied to any function which assigns a real number to an element of a vectors space, not just a space of functions.

True, this notation is somewhat confusing because the notation $\|.\|_1$ is typically used in a different meaning (the sum-of-absolute-values norm). But that's life too...

When $\|.\|_0$ norm is used to define the neighbourhood, we say that J attains at y^* a *strong* minimum. When $\|.\|_1$ norm is used instead, the minimum is *weak*.

It may take a few seconds to realize that if y^* is a strong minimum, it is automatically weak. The oposite is not true. If y^* is a weak minimum, it is not necessarily strong. In other words, for a fixed ε, the set $N_1 = \{y : \|y - y^*\|_0 \leq \varepsilon\}$ contains all the members of the set $N_2 = \{y : \|y - y^*\|_1 \leq \varepsilon\}$ but the other direction does not hold. The visualization in Fig. 3 may help see this.

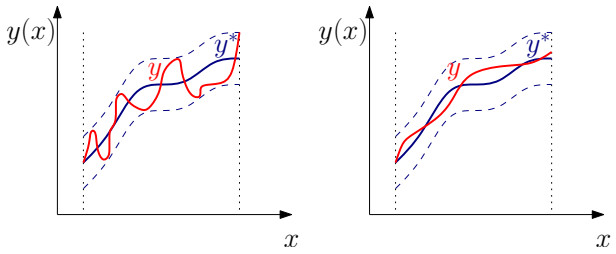

Figure 3: Optimal function y^* and a neighbourhood in the sense of $\|.\|_0$ norm in the left and $\|.\|_1$ norm in the right. The difference is that in the left any smooth function can be considered that fits into the graphical bounds, whereas in the right not only the magnitude (the graphical boundary) matters but the rate of change is bounded as well.

Example 1.1. *Consider the minimization of the functional* $J(y) = \int_1^1 [(y'(x))^2(1 - (y'(x))^2)]dx$ *for which it is requested that* $y(a) = y(b) = 0$. *Clearly* $y(x) = 0$ *is a weak minimum but is not a strong minimum. Just observe that even for tiny perturbations in magnitude, if the derivative is high, that is,* $(y'(x))^2 > 1$, *the functional* J *is negative.*

What is the role of these two norms in our course? The former—the $\|.\|_0$ norm—and the related concept of strong minimum are what we would like to test for, while the latter—the $\|.\|_1$ norm—and the related concept of weak minimum are just mathematically more convenient. It is much easier to show that a point is a weak minimum. Nonetheless, the distinction between these two will only be relevant once we want to find sufficient conditions of optimality. But we are not there yet. First we need to find the (first-order) necessary conditions of optimality.

1.2 Variation, variational derivative and first-order conditions of optimality

Similarly as in the finite-dimensional optimization, we will build the necessary conditions of optimality by studying how the cost function changes if we perturb the independent variable a bit. Let's denote the minimizing function as y^*. The perturbed function will be denoted as y and is given by

$$y(x) = y^*(x) + \delta y(x), \tag{10}$$

where $\delta y(x)$ is *variation of function* and it is a function itself. The variation plays the same role in calculus of variations as the term $\mathrm{d}x$ plays in differential calculus.

Recall that one aproach to deriving the first-order necessary conditions of optimality in the case of vector variables was based on fixing the direction first and then analyzing how the function evolves along this direction. Namely, we considered evolution of the cost function

$$f(\mathbf{x}^* + \alpha\,\mathbf{d}) \tag{11}$$

for given $\mathbf{x}^* \in \mathbb{R}^n$ and $\mathbf{d} \in \mathbb{R}^n$ while varying $\alpha \in \mathbb{R}$. This enables us to convert the vector problem into a scalar one. We can follow this procedure while perturbing function. Namely, we can build the variation of a function by writing it as

$$\delta y(x) = \alpha\eta(x), \tag{12}$$

where $\eta(x)$ is a given (but arbitrary) function in $\mathcal{C}^1$ (playing the role of $\mathbf{d}$ in the finite-dimensional optimization) and $\alpha \in \mathbb{R}$. This way we are about to convert optimization over functions to optimization over real numbers.

Before we proceed, let's elaborate a bit more on the above expression. Let's assume that the function $y(x)$ in the neighbourhood of the minimizing function $y^*(x)$ is actually parameterized by some real parameter α and for $\alpha = 0$ is becomes the minimizing function $y^*(x)$. The Taylor expansion around $\alpha = 0$ is

$$y(x,\alpha) = \underbrace{y(x,0)}_{y^*(x)} + \left.\frac{\partial y(x,\alpha)}{\partial\alpha}\right|_{\alpha=0} \alpha + \mathcal{O}(\alpha^2). \tag{13}$$

The second term on the right is then the variation δy of the function y. We will write it down here for later convenience

$$\delta y(x) = \underbrace{\left.\frac{\partial y(x,\alpha)}{\partial\alpha}\right|_{\alpha=0}}_{\eta(x)} \alpha. \tag{14}$$

This adds an interpretation for our previous choice in (12). Note also that for a fixed x the variation is just a differential with respect to α.

So far, so good. We have now discussed in quite some detail the concept of variation of the function, that is, a concept that will be used to describe the perturbation of the input argument of a cost functional. But now we want to see if another analogy can be found with differential calculus. Note that the first-order condition of optimality of a cost function $f(x)$ of a scalar real argument x can be stated as a condition on the differential of the cost function

$$\mathrm{d}f = 0. \tag{15}$$

But we also know that that the differential is defined as the first-order approximation to the increment in the input argument, that is

$$\mathrm{d}f = \underbrace{f'}_{\frac{\mathrm{d}f}{\mathrm{d}x}}\mathrm{d}x = 0, \tag{16}$$

from which it follows that the first-order condition of optimality can be given as a condition on the derivative

$$f'(x) = 0. \tag{17}$$

In the vector case of $\mathbf{x} \in \mathbb{R}^n$, we rewrite the above condition on a differential as

$$\mathrm{d}f = (\nabla f)^{\mathrm{T}}\, \mathbf{dx} = 0, \tag{18}$$

from which it follows that

$$\nabla f = \mathbf{0}. \tag{19}$$

Having recapitulated these basic facts from differential calculus, we are now curious if we can do similar development within calculus of variations. Namely, we would like to express the variation of the cost functional using the variation of the function, thus mimicking (18). Note that the product in (18) is actually the *inner product*. And inner products are also defined in other vector spaces, not just Euclidean spaces of n-tuples. For continuous functions they are defined using integrals. Therefore we can anticipate the analogy to (18) in the following form

$$\delta J = \int_a^b \frac{\delta J}{\delta y(x)}\, \delta y(x) \mathrm{d}x, \tag{20}$$

where the fraction in the above expression is called *variational derivative*. Please be aware that the whole fraction should be regarded just as one symbol. You should not really treat it as a true ratio (and cancel the denominator term with the other δy term). This is the same type of a trap[2] that you can encounter in differential calculus using Leibniz's notation.

Now, following (14) we may want to express the variation of J as

$$\delta J = \left.\frac{\mathrm{d}J}{\mathrm{d}\alpha}\right|_{\alpha=0} \alpha. \tag{21}$$

Therefore we will compute the derivative of the cost with respect to the real parameter α first (recall that both the y and the η functions are considered as fixed here)

$$\left.\frac{\mathrm{d}}{\mathrm{d}\alpha} J(y(x) + \alpha\eta(x))\right|_{\alpha=0} = \lim_{\alpha\to 0} \frac{J(y(x) + \alpha\eta(x)) - J(y(x))}{\alpha}. \tag{22}$$

Once we have it, we will just multiply the result by α and we have the desired variation of the cost.

It seems that we are now perfectly ready for narrowing down the family of cost functionals. Let's start by considering some concrete examples. We will then extract the common features and characterize some general and yet narrow enough class of cost functionals.

1.3 Some examples of calculus of variations

1.3.1 Minimum distance between two points

Two points are given in the plane. The task is to find the curve that connects these two points and minimizes the total length. Without a loss of generality consider the two ends on the x-axis as in Fig. 4. Although the answer to this problem is trivial, the problem serves a good job of demonstrating the essence of calculus of variations.

The total length of the curve is

$$J(y) = \int_a^b \sqrt{(\mathrm{d}x)^2 + (\mathrm{d}y)^2} = \int_a^b \sqrt{1 + (y'(x))^2}\mathrm{d}x. \tag{23}$$

[2] Search the internet for phrases such as "Leibniz notation derivative fraction" and you will find numerous discussions of this topic.

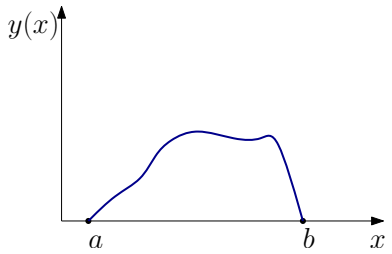

Figure 4: Minimum distance between two points.

1.3.2 Dido's problem

Given a rope of length C, what is the maximum area this rope can circumscribe? Here we have a problem with an equality-type constraint

$$\min_{y\in\mathcal{C}^1} \int_a^b y(x)\mathrm{d}x, \tag{24}$$

subject to

$$\int_a^b \sqrt{1+(y'(x))^2}\mathrm{d}x = C. \tag{25}$$

1.3.3 Brachistochrone problem

The third classical problem mentioned in every textbook on calculus of variations is the problem of *brachistochrone*, where the task is to find a shape of a thin wire with a bead sliding along it (with no friction) in the shortest time, see Fig. 5 (and possibly the video by Vsauce at `https://youtu.be/skvnj67YGmw`).

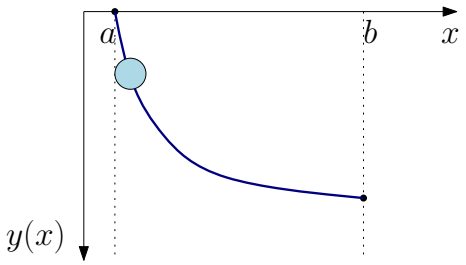

Figure 5: Brachistochrone.

The cost function is simply the total time, that is

$$J = \int_{t_\mathrm{i}}^{t_\mathrm{f}} \mathrm{d}t = t_\mathrm{f} - t_\mathrm{i}. \tag{26}$$

Note that it does not quite fit into the framework that we currently use because time enters here as the independent variable. But there is an easy fix to this. We will express time as a ratio of the distance and velocity. In particular,

$$J = \int_a^b \frac{\mathrm{d}s}{v}. \tag{27}$$

We are already well familiar with the numerator but the velocity in the denominator needs to be determined too. We will use a physical argument here: when the bead is in the initial position, the velocity is zero and the height (as measured along the y axis) is zero. Therefore the total energy given as a sum of kinetic and potential energies $\mathcal{T} + \mathcal{V}$ is zero. But since we assume no friction, the total energy remains constant along the whole trajectory, that is,

$$\frac{1}{2}mv^2 - mgy = 0, \tag{28}$$

from which we can write

$$v(x) = \sqrt{2gy(x)}. \tag{29}$$

We can finally write the expression for the total time as

$$J = \int_a^b \frac{\mathrm{d}s}{v} = \int_a^b \frac{\sqrt{1 + (y'(x))^2}}{\sqrt{2gy(x)}}\mathrm{d}x. \tag{30}$$

1.4 Basic problem of calculus of variations with fixed ends

The only motivation for including those few simple examples was to justify the following general problem. We will call this *the basic problem of calculus of variations with fixed ends*. We will keep considering $\mathcal{C}^1$ functions of x defined on an interval $[a, b]$ with the values at the beginning and end of the interval fixed

$$y(a) = y_a, \qquad y(a) = y_b \tag{31}$$

see Fig. 6 and the task is to find $y^* \in \mathcal{C}^1$ minimizing the functional of the following type

$$J(y) = \int_a^b L(x, y, y')\mathrm{d}x. \tag{32}$$

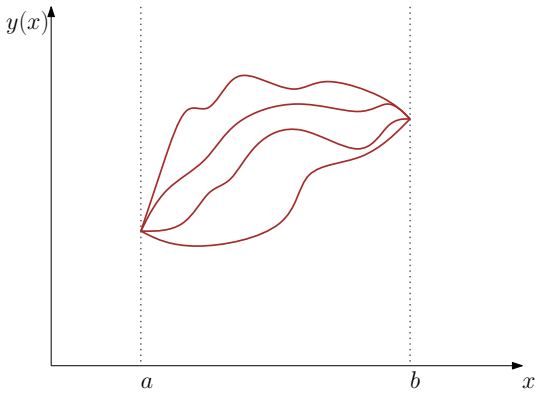

Figure 6: Basic problem of calculus of variations with both ends fixed.

It is possible to extend this basic problem formulation into something more complicated, for example by relaxing the ends, but this will only be done later. First let us solve the basic problem and see if we can use it in the optimal control framework.

In order to state the first-order necessary condition of optimality, we need to find the variation of the cost functional. But we already know that we can form it from

the partial derivative of the cost functional with respect to some real parameter as in (22), that is,

$$\begin{aligned}\frac{\mathrm{d}J(y^*(x)+\alpha\eta(x))}{\mathrm{d}\alpha} &= \frac{\mathrm{d}}{\mathrm{d}\alpha}\int_a^b [L(x,y^*+\alpha\eta,(y^*)'+\alpha\eta')]\mathrm{d}x,\\ &= \int_a^b \frac{\mathrm{d}}{\mathrm{d}\alpha}[L(x,y^*+\alpha\eta,(y^*)'+\alpha\eta')]\mathrm{d}x,\\ &= \int_a^b \left[\frac{\partial L(x,y,y')}{\partial y}\eta(x)+\frac{\partial L(x,y,y')}{\partial y'}\eta'(x)\right]\mathrm{d}x\end{aligned} \tag{33}$$

Now, setting this equal to zero, we do not learn much because the arbitrary η and also its derivative appear in the conditions. It will be much better if we can modify this into something like $\int_a^b [(\quad)\eta(x)]\,\mathrm{d}x$.

Why do we care? It is the content of the *fundamental lemma of calculus of variations* that if the following condition is to hold for any $\eta \in \mathcal{C}^1$ vanishing at a and b

$$\int_a^b \xi(x)\eta(x)\mathrm{d}x = 0, \tag{34}$$

then necessarily $\xi(x) = 0$ identically on the whole interval $[a, b]$. The proof is given elsewhere.

Hence we are motivated to bring the formula for the variation into the format where the derivative of η is missing. This will be accomplished by applying *per partes integration* to the term $\int_a^b \frac{\partial L(x,y,y')}{\partial y'}\eta'(x)\mathrm{d}x$

$$\int_a^b \frac{\partial L(x,y,y')}{\partial y'}\eta'(x)\mathrm{d}x = \left[\frac{\partial L(x,y,y')}{\partial y'}\eta(x)\right]_a^b - \int_a^b \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial L(x,y,y')}{\partial y'}\eta(x)\mathrm{d}x. \tag{35}$$

Substituting back to our expression for the variation, we get

$$\frac{\mathrm{d}J}{\mathrm{d}\alpha} = \left[\frac{\partial L(x,y,y')}{\partial y'}\eta(x)\right]_a^b + \int_a^b \left(\frac{\partial L(x,y,y')}{\partial y} - \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial L(x,y,y')}{\partial y'}\right)\eta(x)\mathrm{d}x. \tag{36}$$

The first term on the right is zero because we assumed at the very beginning that the function $y(x)$ is fixed at both ends, hence the variation $\delta y(x)$ is zero at both ends, hence $\eta(a) = \eta(b) = 0$. As a result, we have the following equation

$$\boxed{\frac{\partial L(x,y,y')}{\partial y} - \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial L(x,y,y')}{\partial y'} = 0} \tag{37}$$

or

$$\boxed{\frac{\partial L(x,y,y')}{\partial y} = \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial L(x,y,y')}{\partial y'}.} \tag{38}$$

This is the famous Euler-Lagrange equation. My biased opinion is that it can be rated as a result that deserves its position in top ten of results in applied mathematics. Smooth function which satisfy it are called *extremals*; but note that all that we know about them is they are just candidate functions for a minimizer, the Euler-Lagrange equation provides just necessary conditions of optimality.

In order to finish bringing this result into the variational format, we can now invoke (21) and we can write

$$\delta J = \int_a^b \underbrace{\left[\frac{\partial L(x,y,y')}{\partial y} - \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial L(x,y,y')}{\partial y'}\right]}_{\frac{\delta J}{\delta y(x)}} \underbrace{\delta y(x)}_{\alpha\eta(x)} \mathrm{d}x, \tag{39}$$

from which we can see that the left hand side of the Euler-Lagrange equation gives us the variational derivative that we were looking for.

You may now wonder why on earth did we actually bother to introduce the new concept of a variation (of a function and of a functional)? We were able to derive the Euler-Lagrange equation just using a partial derivative with respect to α. Good point. In fact, the major motivation was to develop a framework that would resemble that of differential calculus as closely as possible. Knowing now the resulting format of the first-order necessary conditions of optimality, let's now try to rederive it in the fully variational style. That is, we want to find the variation δJ of the functional $J(y(x))$ that is given by the integral:

$$\delta J = \delta \int_a^b L(x,y,y')\mathrm{d}x. \tag{40}$$

The variation now constitutes an operation pretty much mimicking the differential when it comes to dealing with composite functions, products of two function and other situations. Namely, for constant lower and upper bounds in the integral we can move the variation operation into the integral

$$\delta J = \int_a^b \delta L(x,y,y')\mathrm{d}x \tag{41}$$

and then proceeding following those standard rules (shared with operation of differentiation) we get

$$\delta J = \int_a^b \left[\frac{\partial L(x,y,y')}{\partial x}\delta x + \frac{\partial L(x,y,y')}{\partial y}\delta y + \frac{\partial L(x,y,y')}{\partial y'}\delta y'\right]\mathrm{d}x. \tag{42}$$

Now, x is an independent variable, hence it does not vary and $\delta x = 0$. Furthermore, the operations of variation and derivative with respect to x commute, therefore $\delta y'$, which is a shorthand notation for $\delta\left(\frac{\mathrm{d}}{\mathrm{d}x}y(x)\right)$ can be rewritten as

$$\delta y'(x) = \frac{\mathrm{d}}{\mathrm{d}x}(\delta y(x)) \tag{43}$$

and we can write the variation of the cost function as

$$\delta J = \int_a^b \left[\frac{\partial L(x,y,y')}{\partial y}\delta y + \frac{\partial L(x,y,y')}{\partial y'}(\delta y(x))'\right]\mathrm{d}x. \tag{44}$$

Identically as in our previous development, we can get rid of the derivative of the variation using integration *per partes*, which gives

$$\boxed{\delta J = \left[\frac{\partial L(x,y,y')}{\partial y'}\delta y(x)\right]_a^b + \int_a^b \left(\frac{\partial L(x,y,y')}{\partial y} - \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial L(x,y,y')}{\partial y'}\right)\delta y(x)\mathrm{d}x,} \tag{45}$$

which under the assumption of fixed both ends, that is, $\delta y(a) = \delta y(b) = 0$, gives the Euler-Lagrange equation. Elegant procedure, isn't it?

This can perhaps be regarded as culmination of our attempts to develop calculus of variations as an analogy to differential calculus. Let's now move on to learning brand new things.

To make the notation a bit more compact, we will often write the partial derivative of $L(x, y, y')$ with respect to $y(x)$ as L_y. Similarly, the partial of the same function with respect to $y'(x)$ as $L_{y'}$. Using this notation we can immediately show that Euler-Lagrange equation is actually a second-order ordinary differential equation

$$\boxed{L_y - L_{y'x} - L_{y'y}y' - L_{y'y'}y'' = 0.} \tag{46}$$

In order to see how we got this, first recall that $L()$ is a function of x, $y(x)$ and $y'(x)$. Let's write it explicitly as $L(x, y(x), y'(x))$. Then, in the Euler-Lagrange equation we need to find the **total derivative** of $L_{y'}$ with respect to x, that is, we need $\frac{\mathrm{d}}{\mathrm{d}x}L_{y'}$ and we invoke the chain rule for this. Remember, that the function $L_{y'}$ is generally (!) a function of three arguments too: x, y and y'. Therefore, applying the chain rule we will get three terms:

$$\frac{\mathrm{d}}{\mathrm{d}x}L_{y'}(x, y, y') = \underbrace{\frac{\partial L_{y'}}{\partial x}}_{L_{y'x}} + \underbrace{\frac{\partial L_{y'}}{\partial y}}_{L_{y'y}}\underbrace{\frac{\mathrm{d}y(x)}{\mathrm{d}x}}_{y'} + \underbrace{\frac{\partial L_{y'}}{\partial y'}}_{L_{y'y'}}\underbrace{\frac{\mathrm{d}y'(x)}{\mathrm{d}x}}_{y''}.$$

That is it. Combine with the other term L_y from the E.-L. equation and we are done.

We know that in order to specify a solution of a second-order ODE completely, we need to provide two values. Sometimes we specify them at the beginning of the interval, in which case we would give the value of the function and its derivative. This is the well-known *initial value problem* (IVP). On some other occasions we specify the values at two different points on the interval. And this is our case here. In particular, here we have $y(a) = y_a$ and $y(b) = y_b$. which turns the problem into so-called *boundary value problem* (BVP). Both analysis and (numerical) methods for solution for such boundary value problems are a way more difficult that for initial value problems. But there are dedicated solvers (see bvp4c and bvp5c in Matlab, for example).

Nonetheless, before jumping into calling some numerical solvers, let's get some insight for two special cases by analyzing the situations carefully. First, assume that L does not depend on y—a "no y case". Then

$$0 = \frac{\mathrm{d}}{\mathrm{d}x}L_{y'}, \tag{47}$$

and, as a consequence, $L_{y'}$ is constant, independent of x.

The second special case is the "no x case". Then

$$L_y - L_{y'y}y' - L_{y'y'}y'' = 0. \tag{48}$$

By multiplying both sides by y', the equation turns into

$$\frac{\mathrm{d}}{\mathrm{d}x}(L_{y'}y' - L) = 0. \tag{49}$$

(Certainly this needs at least one line on a paper to complete.) As a consequence, $L_{y'}y' - L$ is constant along the optimal curve.

The two new functions that are preserved along the optimal solution are so special that they deserve their own symbols and names:

$$\boxed{p(x) := L_{y'}} \tag{50}$$

and you will see that the choice of the letter "p" is intentional becase this variable will be seen to play the role of *momentum* as you know it from physics, and

$$\boxed{H(x, y, y', p) := py' - L} \tag{51}$$

and you will see that the choice of the letter "H" is intentional becase this variable will be seen to play the role of *Hamiltonian* as you know it from physics.

You may be familiar with these two objects from the courses on physics and/or theoretical mechanics[3].

Let us now see how y and p develop as functions of x and we will use Hamiltonian for that purpose. First, it is immediate from the definition of Hamiltonian that

$$y' = H_p. \tag{52}$$

Similarly, the derivative of momentum is

$$p' = \frac{\mathrm{d}}{\mathrm{d}x} L_{y'} = L_y = -H_y, \tag{53}$$

where the second equality comes from E.-L. equation and the third from the definition of H. We will now view the two differential equations as one vector equation

$$\boxed{\begin{bmatrix} y' \\ p' \end{bmatrix} = \begin{bmatrix} H_p \\ -H_y \end{bmatrix}} \tag{54}$$

This version of first-order necessary conditions is no less famous in physics and theoretical mechanics—*Hamilton's canonical equations*—and some of our results on optimal control will come in this format.

Example 1.2. *Let us now see how the analysis of the special cases can be practically useful. We will only have a look at the minimum distance problem. The Lagrangian is*

$$L(y') = \sqrt{1 + (y')^2}, \tag{55}$$

which is clearly independent of y (and of x as well). Therefore

$$L_{y'} = \frac{1}{2} \frac{2y'}{\sqrt{1 + (y')^2}} \tag{56}$$

must be constant. The only way is to have $y'(x)$ constant, that is, the graph of the function $y(x)$ must be a line. Introducing the boundary conditions, it is now obvious that the solution is a line connecting the two points.

[3] Unfortunately, we have to give a warning here, that while the above definitions are well accepted in the physics-related fields of science, most control theory books adopt a slightly different notation, which may be confusing. We will have more on that in a while (or see directly 3.4.4 in Liberzon's book)

One important general property will be revealed if we differentiate H with respect to y'

$$H_{y'}(x, y, y', p) = p - L_{y'}. \tag{57}$$

If p is chosen as $L_{y'}$, the derivative of Hamiltonian vanishes. In other words, when H is evaluated on the *extremal* curve, it has a stationary point with respect to the third variable. To support this mental step of regarding the third (input) argument as independent from the rest, we write the Hamiltonian for the extremal with the third variable relaxed as $H(x, y, z, p)$. The above result says that

$$\left.\frac{\partial H(x, y, z, p)}{\partial z}\right|_{z=y'} = 0 \tag{58}$$

when evaluated on the optimal trajectory. In fact, as we will see shortly, Hamiltonian is not only stationary along the optimal trajectory but it achieves the maximum value. It will turn out that it is actually this property—the Hamiltonian being maximum along the optimal trajectory with respect to y'—is the core property of the problem. The fact that the derivative is zero is just a consequence in the special case when such derivative exists.

1.5 Sufficient conditions of optimality (minimum)

What remains to be done before we come to applying the EL equation to control problems is to discuss how we can distinguish the minimum from the maximum. What if the solution to EL equation actually is maximizing the cost functional? Or what if it is just a saddle "point"? The mathematics needed to answer these questions is quite delicate, hence we will only sketch the direction of reasoning and for complete proofs refer to the literature.

Knowing that the first variation vanishes for an extremal, higher order terms need to be investigated, starting with the term in the Taylor's expansion corresponding to the squared variations. Similarly as in the finite-dimensional optimization we first argue that for small enough α the second order term dominates all the higher order terms and then we study under which conditions is the second order term nonnegative (for the second-order necessary condition) or positive (for the second-order sufficient condition).

The answer for the necessity part, which relies heavily on the fact that we have decided to work with the $\|.\|_1$ norm, is that

$$L_{y'y'} \geq 0 \tag{59}$$

needs to be satisfied. This is called the *Legendre necessary condition*. We have certainly skipped a lot of nontrivial work that needs to be done to show this result. Check the literature for details if you are interested.

The sufficiency part is even more complicated. It turns out that merely sharpening the necessary Legendre condition into

$$L_{y'y'} > 0 \tag{60}$$

is not enough to guarantee the minimality. The additionial constraint is quite involved even to be merely stated (it is called Jacobi condition and has something to do with

absence of conjugate points on the interval of control[4]). Instead we will state that the optimality is guaranteed if the inverval of x is not too long.

Crucial observation now is that

$$H_{y'y'} = -L_{y'y'}. \tag{61}$$

Hence, if a given function y minimizes J, then $L_{y'y'} \geq 0$ and

$$H_{y'y'} \leq 0, \tag{62}$$

which reads that Hamiltonian achieves maximum when evaluated on the optimal curve. This is a key property and constitutes a preview of the celebrated Pontryagin's principle of minimum. It can be restated as

$$\boxed{H(t, y, y', p) \geq H(t, y, z, p)} \tag{63}$$

for all $z \in \mathcal{C}^1$ on the interval $[a, b]$ and close to $y'(t)$ (in the sense of 1-norm).

1.6 Constrained problems in calculus of variations

Once again following our own path in finite-dimensional optimization, now that we have covered the first-order necessary and second-order sufficient conditions of optimality, we need to cover problems with constraints. Here we will only investigate equality type-constraints. The inequality-type constraints are beyond the reach of methods of calculus of variations. But as we will see, an offspring of calculus of variations—Pontryagin's principle—will handle these easily.

The constraints that we are going to encounter in optimal control come in the form of differential equations, and these constitute pointwise constraints

$$F(x, y, y') = 0. \tag{64}$$

For every value of the independent variable (x in the standard setting of calculus of variations) we have one constraint. Therefore we have a continuum of constraints. As a consequence, we will need an infinite number of Lagrange multipliers as well—in other words, the Lagrange multiplier will be a function of x too. The *augmented cost function* is

$$J^{\text{aug}}(y) = \int_a^b L(x, y, y')\mathrm{d}x + \int_a^b \lambda(x) \cdot F(x, y, y')\mathrm{d}x, \tag{65}$$

where the symbol " $\cdot$ " is there to emphasize that in the case when both y and λ are vector functions (hence F is a vector), the second integrand is obtained as an inner product. Rewriting the above expression for the augmented criterion of optimality as

$$J^{\text{aug}}(y) = \int_a^b \left[L(x, y, y') + \lambda(x) \cdot F(x, y, y')\right] \mathrm{d}x \tag{66}$$

suggests that we can introduce an augmented Lagrangian

$$L^{\text{aug}}(x, y, y', \lambda) = L(x, y, y') + \lambda(x) \cdot F(x, y, y') \tag{67}$$

[4]The only motivation for stating these terms here without actually providing any explanation is just to provide you with keywords and search phrases just in case you want to learn more elsewhere.

and continue as we did in the unconstrained case. For completeness, let's state here that in the case of vector functions, the augmented Lagrangian is given as

$$L^{\text{aug}}(x, y, y', \lambda) = L(x, y, y') + \lambda(x)^{\text{T}} \cdot F(x, y, y') \tag{68}$$

A word of warning is needed here, though. Similarly as in the unconstrained case, it can happen that the constraints will be *degenerate*, in which case the EL equation fails to be a necessary condition of optimality. We will not discuss this delicate issue here and rather direct the interested student to the Liberzon's book.

2 Optimal control problem on a finite and fixed time interval with fixed final state

Now we finally seem to be ready for solving our optimal control problems stated at the beginning of the lecture. Equipped with the solution to the fixed-ends basic problem of calculus of variation, we start with the finite-horizon fixed-final state version. We will extend the result for a free final state in due course. The problem to be solved is

$$\min_{\mathbf{x}(t),\mathbf{u}(t)} \left[\int_{t_{\text{i}}}^{t_{\text{f}}} L(\mathbf{x}, \mathbf{u}, t) \mathrm{d}t \right], \tag{69}$$

subject to

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}, \mathbf{u}, t), \qquad \mathbf{x}(t_{\text{i}}) = \mathbf{r}_{\text{i}}. \tag{70}$$

Note that there is no term penalizing the state at the final time since it is requested that the system is brought to some prespecified state

$$\mathbf{x}(t_{\text{f}}) = \mathbf{r}(t_{\text{f}}). \tag{71}$$

The augmented cost function and the augmented Lagrangian are

$$J^{\mathbf{aug}}(t, \mathbf{x}, \dot{\mathbf{x}}, \mathbf{u}, \boldsymbol{\lambda}) = \int_{t_{\text{i}}}^{t_{\text{f}}} \left[\underbrace{L(\mathbf{x}, \mathbf{u}, t) + \boldsymbol{\lambda}^{\text{T}} \left(\dot{\mathbf{x}} - \mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{t}) \right)}_{L^{\text{aug}}} \right] \mathrm{d}t. \tag{72}$$

Note that compared to the original unconstrained calculus of variations setting, here we made a notational shift from x to t as the independent variable, y to (x, u, λ) as the dependent variables and the only derivative appearing in the augmented Lagrangian is $\dot{x}$.

Applying EL equation to this augmented Lagrangian we obtain three equations (or sets of equations), one for each dependent variable.

$$L^{\text{aug}}_{\mathbf{x}} = \frac{\mathrm{d}}{\mathrm{d}t} L^{\text{aug}}_{\dot{\mathbf{x}}}, \tag{73}$$

$$L^{\text{aug}}_{\mathbf{u}} = 0, \tag{74}$$

$$L^{\text{aug}}_{\lambda} = 0. \tag{75}$$

These can be expanded in terms of the unconstrained Lagrangian (and I will assume scalar functions first for notational simplicity)

$$\frac{\partial L}{\partial x} - \lambda \frac{\partial f}{\partial x} = \dot{\lambda}, \tag{76}$$

$$\frac{\partial L}{\partial u} - \lambda \frac{\partial f}{\partial u} = 0, \tag{77}$$

$$\dot{x}(t) - f(x, u, t) = 0. \tag{78}$$

In the vector case (when $\mathbf{x}$, hence $\mathbf{f}()$, and/or $\mathbf{u}$ are vectors)

$$\nabla_{\mathbf{x}} L - \sum_{i=1}^{n} \lambda_i \nabla_{\mathbf{x}} \mathbf{f}_i(\mathbf{x}, \mathbf{u}) = \dot{\boldsymbol{\lambda}}, \tag{79}$$

$$\nabla_{\mathbf{u}} L - \sum_{i=1}^{n} \lambda_i \nabla_{\mathbf{u}} \mathbf{f}_i(\mathbf{x}, \mathbf{u}) = \mathbf{0}, \tag{80}$$

$$\dot{\mathbf{x}}(t) - \mathbf{f}(\mathbf{x}, \mathbf{u}, t) = \mathbf{0}. \tag{81}$$

We can also write the same result in the compact vector form. Recall that we agreed in this course to regard gradients as column vectors and that $\nabla \mathbf{f}$ for a vector function $\mathbf{f}$ is a matrix whose columns are gradients $\nabla_{\mathbf{x}} f_i$ of the individual elements of the vector function. We can then write the first order conditions compactly as

$$\nabla_{\mathbf{x}} L - \nabla_{\mathbf{x}} \mathbf{f}\, \boldsymbol{\lambda} = \dot{\boldsymbol{\lambda}}, \tag{82}$$

$$\nabla_{\mathbf{u}} L - \nabla_{\mathbf{u}} \mathbf{f}\, \boldsymbol{\lambda} = \mathbf{0}, \tag{83}$$

$$\dot{\mathbf{x}}(t) - \mathbf{f}(\mathbf{x}, \mathbf{u}, t) = \mathbf{0}. \tag{84}$$

These three (sets of) equations give the necessary conditions of optimality that we were looking for. We can immediately recognize the last one—the original state equation describing how the state vector $\mathbf{x}$ evolves in time. The other two equations are new, though. The first one is called *costate equation* because the variable $\boldsymbol{\lambda}$, originally introduced as a Lagrange multiplier, now evolves also according to a first-order differential equation. The middle equation is called an *equation of stationarity*. With the exception of some singular cases, it can often be used to determine how the control vector $\mathbf{u}$ depends on the state $\mathbf{x}$ and the costate $\boldsymbol{\lambda}$, in which case u is eliminated from the two differential equations and we end up with differential equations just in $\mathbf{x}$ and $\boldsymbol{\lambda}$.

Note that for differential equations, we always need a sufficient number of boundary conditions to determine the solution uniquely. In particular, for a state vector of dimension n, the costate is also of dimension n, hence we need in total $2n$ boundary conditions. In our current setup these are given by the n specified values of the state vector at the beginning and n values at the end. This class of problems is called *two-point boundary value problem (BVP)* and generally it can only be solved using some numerical algoritms. Such algorithms are also implemented in Matlab (bvp4c and bvp5c functions).

To get some more insight and also to develop a practical design tool, let us consider the LQ version of this general problem. That is, the augmented Lagrangian is

$$L^{\text{aug}} = \frac{1}{2} \left(\mathbf{x}^{\mathrm{T}} \mathbf{Q} \mathbf{x} + \mathbf{u}^{\mathrm{T}} \mathbf{R} \mathbf{u} \right) + \boldsymbol{\lambda}^{\mathrm{T}} \cdot (\dot{\mathbf{x}} - \mathbf{A}\mathbf{x} - \mathbf{B}\mathbf{u}). \tag{85}$$

The three necessary conditions of optimality are

$$\mathbf{Q}\mathbf{x} - \mathbf{A}^{\mathrm{T}} \boldsymbol{\lambda} = \dot{\boldsymbol{\lambda}}, \tag{86}$$

$$\mathbf{R}\mathbf{u} - \mathbf{B}^{\mathrm{T}} \boldsymbol{\lambda} = 0, \tag{87}$$

$$\dot{\mathbf{x}} - \mathbf{A}\mathbf{x} - \mathbf{B}\mathbf{u} = 0. \tag{88}$$

Provided $\mathbf{R} > 0$, we can express $\mathbf{u}$ from the second equation (aka the equation of stationarity) and substitute into the third one

$$\boxed{\mathbf{u} = \mathbf{R}^{-1} \mathbf{B}^{\mathrm{T}} \boldsymbol{\lambda}.} \tag{89}$$

This leaves us with just two differential equations

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\boldsymbol{\lambda}, \tag{90}$$

$$\dot{\boldsymbol{\lambda}} = \mathbf{Q}\mathbf{x} - \mathbf{A}^{\mathrm{T}}\boldsymbol{\lambda}. \tag{91}$$

This is a set of $2n$ differential equations of first order (we assume that $\mathbf{x} \in \mathbb{R}^n$). We have $2n$ boundary conditions to fully specify the solution

$$\mathbf{x}(t_{\mathrm{i}}) = \mathbf{r}_{\mathrm{i}}, \tag{92}$$

$$\mathbf{x}(t_{\mathrm{f}}) = \mathbf{r}_{\mathrm{f}}. \tag{93}$$

Let's now rewrite the equations as one large vector equation

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\boldsymbol{\lambda}} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{A} & \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}} \\ \mathbf{Q} & -\mathbf{A}^{\mathrm{T}} \end{bmatrix}}_{\tilde{\mathbf{A}}} \underbrace{\begin{bmatrix} \mathbf{x} \\ \boldsymbol{\lambda}. \end{bmatrix}}_{\mathbf{w}} \tag{94}$$

Similarly as in the discrete-time case, we can find the solution numerically by relating the the state and costate at both ends of the time interval. In the continuous-time setting, it is the exponential of the matrix in the above equation that will be used for that purpose. Let's do it now. By labelling the block matrix in the above as $\tilde{\mathbf{A}}$ and the vector composed of the state and costate vectors as $\mathbf{w}$, we can write

$$\mathbf{w}(t_{\mathrm{f}}) = \underbrace{e^{\tilde{\mathbf{A}}(t_{\mathrm{f}})}}_{\boldsymbol{\Phi}(t_{\mathrm{f}})}\mathbf{w}(0), \tag{95}$$

where $\boldsymbol{\Phi}()$ is a *state-transition matrix* (well, this is how we would call it if it was related to x only but here we use it for relating not only the state but also the costate vector). Labelling the blocks as in

$$\boldsymbol{\Phi}(t) = \begin{bmatrix} \boldsymbol{\Phi}_{11}(t) & \boldsymbol{\Phi}_{12}(t) \\ \boldsymbol{\Phi}_{21}(t) & \boldsymbol{\Phi}_{22}(t) \end{bmatrix}, \tag{96}$$

we can write from the equation for the state

$$\boldsymbol{\lambda}(0) = \boldsymbol{\Phi}_{12}^{-1}\left(\mathbf{x}(t_{\mathrm{f}}) - \boldsymbol{\Phi}_{11}\mathbf{x}(0)\right). \tag{97}$$

Now, the initial value problem in (94) can be solved for both the state and the costate. Finally, the control signal can be computed using (89). Some code is available on the course website.

However feasible the above procedure is, we can get even more insight into the problem after setting $\mathbf{Q} = \mathbf{0}$, which will essentially decouple the second equation from the first. Then we can follow the same procedure as we did in the discrete-time setting. The control design problem can be interpreted as driving the system from a given initial state to some given final state while minimizing the "energy". The procedure for solving the problem

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\boldsymbol{\lambda}, \tag{98}$$

$$\dot{\boldsymbol{\lambda}} = -\mathbf{A}^{\mathrm{T}}\boldsymbol{\lambda} \tag{99}$$

then proceeds by expressing from the costate equation the solution for $\boldsymbol{\lambda}$ as a function of costate at the final time

$$\boldsymbol{\lambda}(t) = e^{-\mathbf{A}^{\mathrm{T}}(t-t_{\mathrm{f}})}\boldsymbol{\lambda}(t_{\mathrm{f}}). \tag{100}$$

We substitute into the state equation

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}} e^{-\mathbf{A}^{\mathrm{T}}(t-t_{\mathrm{f}})}\boldsymbol{\lambda}(t_{\mathrm{f}}) \tag{101}$$

and solve for $\mathbf{x}(t)$

$$\mathbf{x}(t) = e^{\mathbf{A}(t-0)}x_0 + \int_0^t \left[e^{\mathbf{A}^{\mathrm{T}}(t-\tau)}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}} e^{-\mathbf{A}^{\mathrm{T}}(\tau-t_{\mathrm{f}})}\boldsymbol{\lambda}(t_{\mathrm{f}})\right] \mathrm{d}\tau. \tag{102}$$

Evaluating this at $t = t_{\mathrm{f}}$, and moving the costate at the final time outside of the integral

$$\mathbf{x}_f = e^{\mathbf{A}(t_{\mathrm{f}}-0)}x_0 + \underbrace{\int_0^{t_{\mathrm{f}}} \left[e^{\mathbf{A}^{\mathrm{T}}(t_{\mathrm{f}}-\tau)}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}} e^{\mathbf{A}^{\mathrm{T}}(t_{\mathrm{f}}-\tau)}\right] \mathrm{d}\tau}_{\mathbf{G}_R(0,t_{\mathrm{f}})} \boldsymbol{\lambda}(t_{\mathrm{f}}). \tag{103}$$

The as of yet unknown $\lambda(t_{\mathrm{f}})$ can finally be extracted

$$\boldsymbol{\lambda}(t_{\mathrm{f}}) = \mathbf{G}_R(0,t_1)^{-1}\left(\mathbf{x}_f - e^{\mathbf{A}(t_{\mathrm{f}}-0)}x_0\right), \tag{104}$$

where $\mathbf{G}_R(0,t_1)$ is the weighted controllability/reachability gramian. Remember that its inverse only exists if the system is controllable/reachable. The last step constitutes in bringing this value back into the formula for the optimal control

$$\begin{aligned}\mathbf{u}(t) &= \mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\boldsymbol{\lambda}(t), &(105)\\ &= \mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}} e^{-\mathbf{A}^{\mathrm{T}}(t-t_{\mathrm{f}})}\boldsymbol{\lambda}(t_{\mathrm{f}})\mathbf{G}_R(0,t_1)^{-1}\left(\mathbf{x}_f - e^{\mathbf{A}(t_{\mathrm{f}}-0)}x_0\right). &(106)\end{aligned}$$

The conclusion is identical as in the discrete-time setting: the minimum-energy control for a fixed final state assignment can be obtained in the form of a precomputed signal. The necessary condition for the existence of optimal control is controllability of the system.

Honestly, I can hardly view the contribution of the above procedure (exploiting the structure of the problem when $\mathbf{Q} = \mathbf{0}$) as anything else than just giving us some insight. For linear systems we already know from our previous treatment of discrete-time systems (and we will see it in a while for continuous-time systems too) that much more useful feedback controller can be designed if we relax the final state constraint.

3 Optimal control on a finite and fixed time interval with free final state

Apparently, the open-loop nature of the optimal control for the fixed-final-state scenario is not quite satisfactory in most engineering applications. Similarly as in the discrete-time situation we may suspect that by relaxing the final state we may obtain a more useful control scheme. Relaxing the final state does not mean that we resign at the task of controlling the system behavior at the end of the control interval. It is only that now we will have to use the terminal and the running costs to enforce fast enough response of the system. The optimal control criterion is then

$$J = \phi(\mathbf{x}(t_{\mathrm{f}})) + \int_{t_{\mathrm{i}}}^{t_{\mathrm{f}}} L(\mathbf{x},\mathbf{u},t)\mathrm{d}t. \tag{107}$$

First we need to go back to the basic problem of calculus of variations and see how the solution to the basic problem changes if we set one of the ends free. As we are going to see in a minute, the optimal solution will still have to satisfy the Euler-Lagrange equation, the boundary condition will change, however.

Switching to the notation in the calculus of variations temporarily, the problem is described in the Fig. 7.

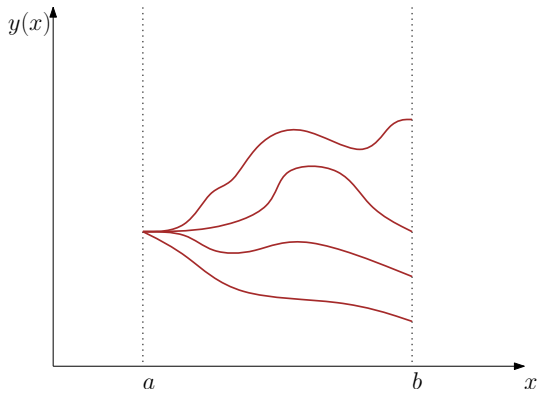

Figure 7: Basic problem of calculus of variations with a free end.

Let us recall here the already derived expression (45) for the first variation

$$\delta J = \int_a^b \left(\frac{\partial L(x,y,y')}{\partial y} - \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial L(x,y,y')}{\partial y'} \right) \delta(x)\mathrm{d}x + \left[\frac{\partial L(x,y,y')}{\partial y'}\delta(x) \right]_a^b. \tag{108}$$

This time, however, the second term on the right is not zero. In particular, $\delta y(a) = 0$ but $\delta y(b) \neq 0$. Therefore the first variation is

$$\delta J = \frac{\partial L(x,y,y')}{\partial y'}\delta y(x)\bigg|_b + \int_a^b \left(\frac{\partial L(x,y,y')}{\partial y} - \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial L(x,y,y')}{\partial y'} \right) \delta y(x)\mathrm{d}x. \tag{109}$$

The sum must be equal to zero. Note that even though we relaxed the condition on one end of the curve, the family of perturbations still contains the functions that vanish at the end point, that is $\delta y(b) = 0$. Therefore when searching for the necessary conditions, it appears that the extremals must still satisfy the EL equation. If the EL equation is satisfied for perturbations vanishing at the end, the first term on the right hand above (the integral) must also vanish for the perturbations not vanishing at the end. As a consequence, the second term on the right must be equal to zero as well. To conclude, setting one of the ends free, the necessary conditions of optimality are still given by the EL equation but the boundary condition $y(b) = y_b$ is replaced by

$$\frac{\partial L(x,y,y')}{\partial y'}\bigg|_b = 0. \tag{110}$$

This result can be immediately applied to the free-final-state optimal control problem. The only deficiency is that the criterion of optimality

$$J = \int_{t_\mathrm{i}}^{t_\mathrm{f}} L(\mathbf{x},\mathbf{u},t)\mathrm{d}t \tag{111}$$

does not include the term penalizing the final state. We will correct this in a moment. For the time being, note that the solution to the current problem is identical to the solution to the fixed final state problem with the final state condition $\mathbf{x}(t_\mathrm{f})$ replaced by the condition

$$\boldsymbol{\lambda}(t_\mathrm{f}) = \mathbf{0}. \tag{112}$$

Now let's include the term penalizing the final state. This is actually quite easy: what we need to do is to bring that term under the integral sign

$$\phi(\mathbf{x}(t_\mathrm{f})) = \int_{t_\mathrm{i}}^{t_\mathrm{f}} \frac{\mathrm{d}\phi}{\mathrm{d}t}\mathrm{d}t + \phi(\mathbf{x}(t_\mathrm{i})), \tag{113}$$

$$= \int_{t_\mathrm{i}}^{t_\mathrm{f}} \left[\frac{\partial\phi}{\partial t} + (\nabla_\mathbf{x}\phi)^\mathrm{T}\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t}\right]\mathrm{d}t + \phi(\mathbf{x}(t_\mathrm{i})). \tag{114}$$

Note that the last term on the right (the one corresponding to $\mathbf{x}(t_\mathrm{i})$) is constant and excluding it from the optimization has no impact on the optimal solution.

Restricting our attention to time-invariant cases in favor of simplicity (assuming $\phi_t = 0$), the augmented Lagrangian is modified to

$$J^\mathrm{aug}(t,\mathbf{x},\dot{\mathbf{x}},\mathbf{u},\boldsymbol{\lambda}) = \int_{t_\mathrm{i}}^{t_\mathrm{f}} \left[\underbrace{L(\mathbf{x},\mathbf{u},t) + (\nabla_\mathbf{x}\phi)^\mathrm{T}\cdot\dot{\mathbf{x}} + \boldsymbol{\lambda}^\mathrm{T}(t)\cdot(\dot{\mathbf{x}}(t) - \mathbf{f}(\mathbf{x},\mathbf{u},\mathbf{t}))}_{L^\mathrm{aug}}\right]\mathrm{d}t. \tag{115}$$

Substituting to (110), the new boundary condition corresponding to the final time is

$$\boxed{(\nabla_\mathbf{x}\phi)(t_\mathrm{f}) + \boldsymbol{\lambda}(t_\mathrm{f}) = 0.} \tag{116}$$

Specializing the result to the LQ case with the final state penalization

$$\phi(\mathbf{x}(t_\mathrm{f})) = \frac{1}{2}\mathbf{x}^\mathrm{T}(t_\mathrm{f})\mathbf{S}_\mathrm{f}\mathbf{x}(t_\mathrm{f}), \tag{117}$$

we get the new boundary condition

$$\boxed{\mathbf{S}_\mathrm{f}\mathbf{x}(t_\mathrm{f}) + \boldsymbol{\lambda}(t_\mathrm{f}) = 0.} \tag{118}$$

This looks already familiar, right? We found an identical relationship between the state and the costate at the final time in the discrete-time setting[5]. The difference is in the sign, we will comment on this in a while. Suffice to say for now that this has no impact on the solution.

Restate here the full necessary conditions for the LQ problem. The state and the costate equations are

$$\boxed{\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^\mathrm{T}\boldsymbol{\lambda},} \tag{119}$$

$$\boxed{\dot{\boldsymbol{\lambda}} = \mathbf{Q}\mathbf{x} - \mathbf{A}^\mathrm{T}\boldsymbol{\lambda}.} \tag{120}$$

The stationarity equation is

$$\boxed{\mathbf{u} = \mathbf{R}^{-1}\mathbf{B}^\mathrm{T}\boldsymbol{\lambda}.} \tag{121}$$

[5]Note that there was actually a difference in sign. This is due to our arbitrary choice while defining the augmented Lagrangian. See the discussion of this at the end of this text.

The two sets of boundary equations are (118) and $\mathbf{x}(t_\mathrm{i}) = 0$.

Similarly as in the previous scenario with fixed final state, here we can also proceed by numerically solving the linear boundary value problem. For completeness I will describe it here, but note that in a few moments we will learn something more about this problem.

Briefly, from the state equation we have

$$\mathbf{x}(t_\mathrm{f}) = \boldsymbol{\Phi}_{11}\mathbf{x}(0) + \boldsymbol{\Phi}_{12}\boldsymbol{\lambda}(0), \tag{122}$$

which after multiplication by $\mathbf{S}_\mathrm{f}$ gives

$$\mathbf{S}_\mathrm{f}\mathbf{x}(t_\mathrm{f}) = \mathbf{S}_\mathrm{f}\boldsymbol{\Phi}_{11}\mathbf{x}(0) + \mathbf{S}_\mathrm{f}\boldsymbol{\Phi}_{12}\boldsymbol{\lambda}(0). \tag{123}$$

The boundary condition in the free final state case is

$$\mathbf{S}_\mathrm{f}\mathbf{x}(t_\mathrm{f}) = -\boldsymbol{\lambda}(t_\mathrm{f}), \tag{124}$$

which immediately invites us to substitute to the right side the solution of the costate equation

$$-\boldsymbol{\Phi}_{21}\mathbf{x}(0) - \boldsymbol{\Phi}_{12}\boldsymbol{\lambda}(0) = \mathbf{S}_\mathrm{f}\boldsymbol{\Phi}_{11}\mathbf{x}(0) + \mathbf{S}_\mathrm{f}\boldsymbol{\Phi}_{12}\boldsymbol{\lambda}(0), \tag{125}$$

from which we can compute the initial value of the costate

$$\boldsymbol{\lambda}(0) = -(\mathbf{S}_\mathrm{f}\boldsymbol{\Phi}_{12} + \boldsymbol{\Phi}_{22})^{-1}(\mathbf{S}_\mathrm{f}\boldsymbol{\Phi}_{11} + \boldsymbol{\Phi}_{12})\mathbf{x}(0). \tag{126}$$

Having computed the initial value of the costate, we can easily solve for the states and costates throughout the time interval. The optimal control then follows from the stationarity equation.

Nonetheless, having computed the solution numerically, some important opportunity escaped our attention. In order to discover it, we need to dig a bit deeper. First, we recall that the boundary condition at the end of the interval gives us a linear relation between the state and the costate. We now assume that this linear relation also holds throughout the interval (the familiar *sweep method*)

$$\mathbf{S}(t)\mathbf{x}(t) = -\boldsymbol{\lambda}(t). \tag{127}$$

Differentiate both sides to obtain

$$\dot{\mathbf{S}}\mathbf{x} + \mathbf{S}\dot{\mathbf{x}} = -\dot{\boldsymbol{\lambda}} \tag{128}$$

Substituting the state equation for $\dot{\mathbf{x}}$ on the left and for $\dot{\boldsymbol{\lambda}}$ on the right we get

$$\dot{\mathbf{S}}\mathbf{x} + \mathbf{S}(\mathbf{A}\mathbf{x} - \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^\mathrm{T}\mathbf{S}\mathbf{x}) = -\mathbf{Q}\mathbf{x} - \mathbf{A}^\mathrm{T}\mathbf{S}\mathbf{x}, \tag{129}$$

which, since $\mathbf{x}$ can be arbitrary, translates to the condition on $\mathbf{S}$

$$\boxed{-\dot{\mathbf{S}} = \mathbf{S}\mathbf{A} + \mathbf{A}^\mathrm{T}\mathbf{S} + \mathbf{Q} - \mathbf{S}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^\mathrm{T}\mathbf{S}.} \tag{130}$$

This is another classical result called *Riccati differential equation*. Initiated at the final time t_f, the differential equation is solved backwards to obtain a function (generally a matrix function) $\mathbf{S}(t)$, which is then substituted into the stationarity equation to obtain the optimal control $\mathbf{u}(t)$

$$\boxed{\mathbf{u}(t) = -\mathbf{R}^{-1}\mathbf{B}^\mathrm{T}\mathbf{S}(t)\mathbf{x}(t).} \tag{131}$$

The story is now completely identical to the discrete-time case—the solution to the Riccati equation evolves in time, but it turns out that for a stabilizable system given by the matrices $(\mathbf{A}, \mathbf{B})$, it converges to some bounded limit. This limit can be either found by implementing the limit searching solver, or it can be retrieved from the *algebraic Riccati equation*

$$\boxed{\mathbf{0} = \mathbf{S}\mathbf{A} + \mathbf{A}^{\mathrm{T}}\mathbf{S} + \mathbf{Q} - \mathbf{S}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\mathbf{S}.} \tag{132}$$

Numerical solution can be found using specialized solvers such as CARE in Matlab. The question remains to see how does the solution to ARE relate to the limiting solution $\mathbf{S}_\infty = \lim_{t\to-\infty} \mathbf{S}(t)$. After all, the ARE is a quadratic equation, which even in the scalar case can have two real solutions. We can afford to be rather short in the remaining analysis since the similarity to the discrete-time case is truly very strong. Thus we can conclude that a unique stabilizing solution of the ARE exists if and only if the artificial system $(\mathbf{A}, \sqrt{\mathbf{Q}})$ is detectable (or observable if we require positive definiteness of $\mathbf{S}(t)$).

4 On notation

4.1 Hamiltonian

Note that there was some ambiguity when forming the augmented Lagrangian. Somehow arbitrarily we have decided in this lecture to define the augmented Lagrangian as

$$L^{\mathrm{aug}} = L(\mathbf{x}, \mathbf{u}, t) + \boldsymbol{\lambda}^{\mathrm{T}}(t) \cdot (\dot{\mathbf{x}}(t) - \mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{t})) \tag{133}$$

but we could have easily formulated it as

$$\hat{L}^{\mathrm{aug}} = L(\mathbf{x}, \mathbf{u}, t) + \hat{\boldsymbol{\lambda}}^{\mathrm{T}}(t) \cdot (\mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{t}) - \dot{\mathbf{x}}(t)) \,. \tag{134}$$

Both are clearly correct and perfectly equivalent. Indeed, although the intermediate steps differ, the final results (Riccati equation, state feedback gain) are identical. What is the motivation for introduction of two differing notations? The former (and the one used in this lecture) enables us writing the augmented Lagrangian using Hamiltonian as

$$L^{\mathrm{aug}} = \boldsymbol{\lambda}^{\mathrm{T}}(t) \cdot \dot{\mathbf{x}}(t) - H(t, \mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) \tag{135}$$

where

$$H(t, \mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) = \boldsymbol{\lambda}^{\mathrm{T}}(t) \cdot \mathbf{f}(\mathbf{x}, \mathbf{u}, t) - L(t, \mathbf{x}, \mathbf{u}), \tag{136}$$

whereas the latter notation supports expressing the augmented Lagrangian as

$$\hat{L}^{\mathrm{aug}} = \hat{H}(t, \mathbf{x}, \mathbf{u}, \hat{\boldsymbol{\lambda}}) - \hat{\boldsymbol{\lambda}}^{\mathrm{T}}(t) \cdot \dot{\mathbf{x}}(t), \tag{137}$$

where

$$\hat{H}(t, \mathbf{x}, \mathbf{u}, \hat{\boldsymbol{\lambda}}) = L(t, \mathbf{x}, \mathbf{u}) + \hat{\boldsymbol{\lambda}}^{\mathrm{T}}(t) \cdot \mathbf{f}(\mathbf{x}, \mathbf{u}, t), \qquad \text{and} \qquad \hat{\boldsymbol{\lambda}} = -\boldsymbol{\lambda}, \tag{138}$$

which is the notation that we decided to use (again, completely arbitrarily, while paying no attention to being consistent) in our lecture on discrete-time systems.

Whether one or the other, the canonical equations are identical. It is only that the second-order sufficiency conditions show maximization of the Hamiltonian in one case and minimization in the other. This can be concluded by observing that

$$H(t, \mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) = -\hat{H}(t, \mathbf{x}, \mathbf{u}, \hat{\boldsymbol{\lambda}}). \tag{139}$$

See more on this issue in 3.4.4 of [4].

4.2 Variation

Upon consulting numerous textbooks and monographs, it appears that the authors are far from an agreement regarding a single definition of variation (within the context of calculus of variations). Essentially, two main definitions appear.

The one that we followed in this lecture defines the variation as an extension of the concept of a differential. That is, a variation δJ of a (cost) functional is a first-order approximation to the increment ΔJ in the (cost) functional J. This we discussed in quite some detail in the text.

The other one defines variation as the derivative of the (cost) functional with respect to the real (perturbation) parameter. In our text, it is the $\frac{\mathrm{d}J}{\mathrm{d}\alpha}$ (for fixed $y(x)$ and $\eta(x)$) that would be called a variation and labelled δJ. The increment in the (cost) functional would be then be approximated by $\delta J\ \alpha$.

Both definitions are often encountered in the literature (see the section on literature below), but we prefer the former because the definitions of variation of a functional δJ and variation of a function $\delta y(x)$ are consistent. Both serve as differentials, that is first-order approximations to increments.

5 Summary

We introduced the basic problem of calculus of variations on a fixed interval and we showed how to solve it both in the situation when both ends are fixed and in the situation when one end is free. Optimal solutions to both problems must satisfy Euler-Lagrange equation, they differ in the boundary condition. We also discussed a bit the sufficient conditions of optimality although we did not go for full derivation. These fundamental results are then applied to the general optimal control problem on a finite horizon, both with fixed and free final state. Then these results are specialized to LTI systems with the popular quadratic cost, leading to the popular LQ optimal control. As the key computational object, the continuous-time algebraic Riccati equation (CARE) is introduced here. The discussion of existence of a uniqueness of a stabilizing solution is very brief here since it mimicks the discussion of discrete-time case completely.

6 Further reading

Calculus of variations is a well-established discipline in mathematical analysis and many dozens books and hundreds if not thousands of online texts are available. This lecture was prepared by using a large number of resources, hence it can hardly be mapped to a single chapter in a single book. But perhaps the chapters 4 and 5 in the affordable reprint [3] were the most influential. The whole book is also very readable. Taking also its low price into consideration, the book is highly recommended.

Furthermore, the slender monograph [2] is a true classic, albeit a bit terse. Another classic for calculus of variations based optimal control is [1], equally terse but very comprehensive.

Some inspiration was also found in the relatively recently published [4], in particular in chapters 2 (intro to calculus of variations), 3 (application of calculus of variations to general problem of optimal control), a partially 6 (application to LQ optimal control). The treatment in the book is fairly rigorous and yet accessible to an engineering student. A draft of the book is available on the author's academic web page. Note however, that the author decides to define the variation of a function in the latter way, that is, for Liberzon the variation of a functional is just the derivative with respect to α.

References

[1] Arthur E. Bryson, Jr. and Yu-Chi Ho. *Applied Optimal Control: Optimization, Estimation and Control*. CRC Press, revised edition edition, January 1975.

[2] I. M. Gelfand and S. V. Fomin. *Calculus of Variations*. Dover Publications, October 2000.

[3] Donald E. Kirk. *Optimal Control Theory: An Introduction*. Dover Publications, April 2004.

[4] Daniel Liberzon. *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton University Press, December 2011.