EN530.603 Applied Optimal Control
Lecture 4: Trajectory Optimization Basics

September 13, 2017

Lecturer: Marin Kobilarov and Matt Sheckells

1 Variations of functions

We are interested in solving optimal control problems such as

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

subject to $\dot{x} = f(x, u, t)$ and other constraints. The cost J is called a *functional*, i.e. it is a function of functions since the trajectories $x(\cdot)$ and $u(\cdot)$ are functions of time. It is possible to optimize a functional in a similar way we optimize a regular function. In particular, there is a functional analog to the necessary conditions for a minimum of a function g given by $\nabla g = 0$. This analog is

$$\text{differential of a function} \quad \nabla g = 0 \quad \Leftrightarrow \quad \text{variation of a functional} \quad \delta J = 0$$

Next, define the *change* in a functional, after *varying* $x(t)$ by $\delta x(t)$ at each t, by

$$\Delta J(x(\cdot), \delta x(\cdot)) = J(x(\cdot) + \delta x(\cdot)) - J(x(\cdot)).$$

The *variation* $\delta J(x(\cdot), \delta x(\cdot))$ is a linear function of $\delta x(\cdot)$ defined by the following relationship

$$\Delta J(x(\cdot), \delta x(\cdot)) = \delta J(x(\cdot), \delta x(\cdot)) + o(\|\delta x\|),$$

where the *small-o* notation was used. For a positive integer p and a function $h : \mathbb{R}^n \to \mathbb{R}^m$ we have

$$h(x) = o(\|x\|^p),$$

if $\lim_{k\to\infty} \frac{h(x_k)}{\|x_k\|^p} = 0$ for all sequences x_k such that $x_k \to 0$ and $x_k \neq 0$ for all k.

In other words, roughly speaking, if $h(x)$ denotes the the second and higher order terms in ΔJ, then $h(x)$ has second and higher-order multiples of δx and thus $h(x)/\|\delta x\|$ is at least proportional to δx which tends to zero as $\delta x \to 0$.

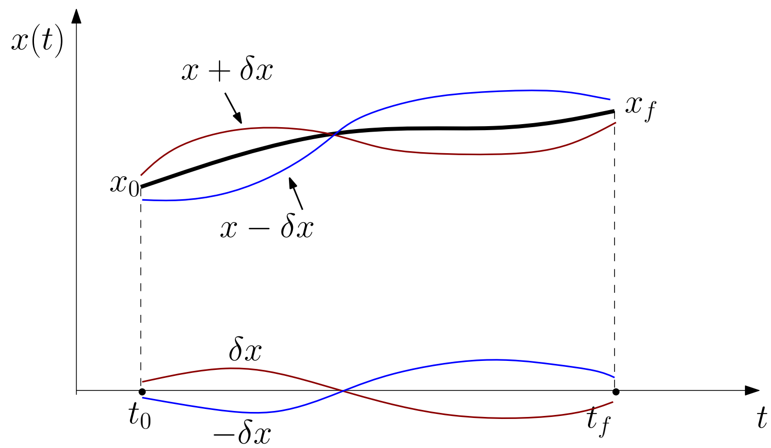

Similarly to standard function optimization, we can use the argument that at an optimum $x^*(\cdot)$ we have

$$\Delta J(x^*(\cdot), \delta x(\cdot)) \geq 0$$

since any change of the cost away from optimum must be positive. Taking $\delta x \to 0$ this implies that

$$\delta J(x^*(\cdot), \delta x(\cdot)) \geq 0.$$

But this must also hold for variations in the direction $-\delta x(\cdot)$, i.e.

$$\delta J(x^*(\cdot), -\delta x(\cdot)) \geq 0 \quad \Rightarrow \quad \delta J(x^*(\cdot), \delta x(\cdot)) \leq 0,$$

by the linearity of δJ. Both conditions can only be true when

$$\delta J(x^*(\cdot), \delta x(\cdot)) = 0,$$

for any $\delta x(\cdot)$.

2 The Euler-Lagrange Equations

Consider the cost function

$$J = \int_{t_0}^{t_f} g(x(t), \dot{x}(t))dt.$$

Its variation becomes

$$\delta J = \int_{t_0}^{t_f} \left[g_x(x, \dot{x})\delta x + g_{\dot{x}}(x, \dot{x})\delta\dot{x}\right] dt,$$

since the variation corresponds to taking first-order terms (i.e. derivatives) of the functions at time t. Integrate by parts (recall $\int u\dot{v} = uv - \int \dot{u}v$) to get

$$\delta J = \int_{t_0}^{t_f} \left[g_x(x, \dot{x})\delta x - \frac{d}{dt} g_{\dot{x}}(x, \dot{x})\delta x\right] dt + g_{\dot{x}}(x(t_f), \dot{x}(t_f))\delta x(t_f)$$

Fixed boundary conditions. If $x(t_0)$ and $x(t_f)$ are given then $\delta x(0) = \delta x(t_f) = 0$ and

$$\delta J = \int_{t_0}^{t_f} \left[g_x(x, \dot{x}) - \frac{d}{dt} g_{\dot{x}}(x, \dot{x})\right] \delta x(t)dt$$

Since $\delta x(t)$ are arbitrary and independent then $\delta J = 0$ only when

$$g_x(x, \dot{x}) - \frac{d}{dt} g_{\dot{x}}(x, \dot{x}) = 0,$$

which are called the *Euler-Lagrange equations (EL).*

Free boundary conditions. When $x(t_f)$ is not fixed, in addition to the EL equation, the following must hold

$$g_{\dot{x}}(x(t_f), \dot{x}(t_f)) = 0.$$

2.1 Example: shortest path curve

Consider a one-dimensional problem with state $x(t)$ and the cost function

$$J = \int_{t_0}^{t_f} \sqrt{1+\dot{x}^2} dt,$$

which in fact corresponds to the length of a curve in the (x, t) plane. The goal is to compute the shortest length curve between two given points (x_0, t_0) and (x_f, t_f).

Applying the EL equations we get

$$\frac{d}{dt}\frac{\dot{x}}{(1+\dot{x}^2)^{1/2}} = 0 \quad \Leftrightarrow \quad \frac{\ddot{x}}{(1+\dot{x})^{3/2}} = 0,$$

which is satisfied when $\ddot{x} = 0$ or when

$$x(t) = c_1 t + c_0,$$

i.e. when $x(t)$ is a straight line. Let $t_0 = 0$ and $t_f = 1$. It is easy to see that $c_0 = x_0$ and $c_1 = x_f - x_0$.

2.2 Particle in 3-D

Let $g(x, \dot{x}, t) = \frac{1}{2}m\|\dot{x}\|^2 - V(x)$, where m denotes the mass and V denotes the potential energy of a particle with position $x \in \mathbb{R}^3$. The function g is actually the Lagrangian of the particle and the EL equations lead to

$$m\ddot{x} = -\nabla_x V,$$

which is simply Newton's law. This is one of the simplest examples that illustrates that Lagrangian mechanics can be considered as a special case of optimal control.

3 Free final-time

When the final time t_f is allowed to vary, the variation of J is expressed as

$$\begin{aligned}\delta J = &\int_{t_0}^{t_f} \left[g_x(x^*, \dot{x}^*, t)\delta x - \frac{d}{dt} g_{\dot{x}}(x^*, \dot{x}^*, t)\delta x \right] dt \\ &+ g_{\dot{x}}(x^*(t_f), \dot{x}^*(t_f), t)\delta x(t_f) + g(x^*(t_f), \dot{x}^*(t_f), t_f)\delta t_f,\end{aligned}$$

where we employ the spatially-optimized trajectory $x^*(t)$ instead of just $x(t)$ to signify that after finding the optimal path without considering δt_f, then variations of t_f infinitesimally contribute to the cost by the term

$$g(x^*(t_f), \dot{x}^*(t_f), t_f)\delta t_f$$

which is simply the cost at the last points multiplied by the time variations: think of it as a first-order approximation to the integral $\int_{t_f}^{t_f+\delta t_f} g(x^*, \dot{x}^*, t)dt$ which is what must be added to the cost when varying time.

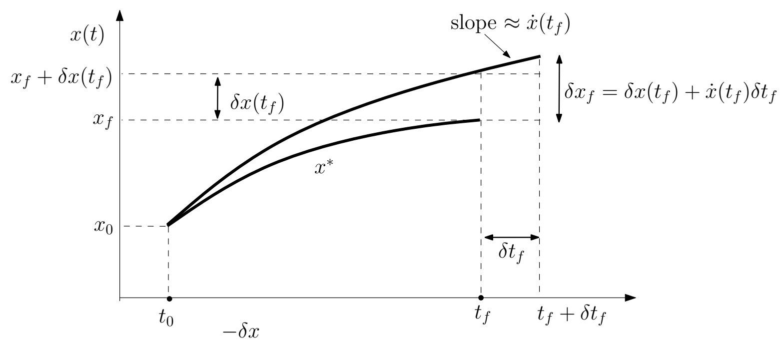

Define the total space-time variation δx_f by

$$\delta x_f = \delta x(t_f) + \dot{x}(t_f)\delta t_f,$$

i.e. this variation combines the effects of varying the final state by keeping time t_f fixed, and then by varying t_f (See figure above).

We have

$$\delta J = \int_{t_0}^{t_f} \left[g_x(x^*, \dot{x}^*, t)\delta x - \frac{d}{dt} g_{\dot{x}}(x^*, \dot{x}^*, t)\delta x \right] dt$$
$$+ \Big(g_{\dot{x}}(x^*, \dot{x}^*, t_f)\delta x_f + \left[g(x^*, \dot{x}^*, t_f) - g_{\dot{x}}(x^*, \dot{x}^*, t_f)\dot{x} \right] \delta t_f \Big)_{t=t_f}.$$

Unrelated t_f and $x(t_f)$. We next consider the case when t_f and $x(t_f)$ are unrelated in which we have

$$g_{\dot{x}}(t_f) = 0, \qquad [g(t_f) - g_{\dot{x}}(t_f)\dot{x}]\,\delta t_f = 0 \quad \Rightarrow \quad g(t_f) = 0$$

Function $\Theta(t_f)$. We next consider the case when the boundary constraint is given by

$$x(t_f) = \Theta(t_f)$$

Variations δt_f and δx_f are related by

$$\delta x_f = \frac{d\Theta}{dt}\delta t_f,$$

where the total variation does not contain a spatial component $\delta x(t_f)$ since the final state is completely determined by time only. The necessary conditions become

$$g_{\dot{x}}(t_f) \left[\frac{d\Theta}{dt} - \dot{x}^* \right]_{t=t_f} + g(t_f) = 0,$$

which are called *trasversality conditions.*

4 Differential Constraints

Consider the optimization of

$$J = \int_{t_0}^{t_f} g(x, \dot{x}, t)dt$$

where $x \in \mathbb{R}^{n+m}$ subject to

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

where $f = (f_1, \dots, f_n)$ are n constraints, and $x(t_f)$ and t_f are fixed. To obtain the necessary conditions, we define the Lagrangian multipliers $\lambda : [t_0, t_f] \to \mathbb{R}^n$ and the augmented cost

$$J_a = \int_{t_0}^{t_f} \left\{ g(x, \dot{x}, t) + \lambda^T f(x, \dot{x}, t) \right\} dt$$

Taking variations

$$J_a = \int_{t_0}^{t_f} \left\{ [g_x + \lambda^T f_x]\delta x(t) + [g_{\dot{x}} + \lambda^T f_{\dot{x}}]\delta \dot{x} + \delta \lambda^T f \right\} dt$$

Integrating by parts we get

$$J_a = \int_{t_0}^{t_f} \left\{ [g_x + \lambda^T f_x - \frac{d}{dt}[g_{\dot{x}} + \lambda^T f_{\dot{x}}]\delta x + \delta \lambda^T f \right\} dt.$$

If we define the augmented cost g_a by

$$g_a = g + \lambda^T f$$

the Euler-Lagrange equations become

$$\frac{\partial g_a}{\partial x} - \frac{d}{dt}\frac{\partial g_a}{\partial \dot{x}} = 0$$

which along with the constraint

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

constitute the necessary conditions.

5 General Boundary Constraints

Let $x \in \mathbb{R}^n$ and consider the optimization of

$$J(x(\cdot), t_f) = \varphi(x(t_f), t_f) + \int_{t_0}^{t_f} g(x, \dot{x}, t)dt$$

subject to the free final time t_f and general boundary conditions

$$\psi(x(t_f), t_f) = 0,$$

where ψ is a vector of m functions. To obtain the necessary conditions define the augmented cost

$$J_a = \varphi(x(t_f), t_f) + \nu^T \psi(x(t_f), t_f) + \int_{t_0}^{t_f} g(x, \dot{x}, t)dt,$$

where $\nu \in \mathbb{R}^m$. Let

$$w(x(t_f), \nu, t_f) = \varphi(x(t_f), t_f) + \nu^T \psi(x(t_f), t_f)$$

Taking variations

$$\delta J_a = w_x(t_f)\delta x_f + w_{t_f}(t_f)\delta t_f + \delta\nu^T\psi(t_f) + \int_{t_0}^{t_f} [g_x \cdot \delta x(t) + g_{\dot{x}} \cdot \delta\dot{x}(t)]dt + g(t_f)\delta t_f,$$

where the last term is due to cost accrued from final time variations. Using integration by parts as well as the total variation definition $\delta x_f = \delta x(t_f) + \dot{x}(t_f)\delta t_f$ we have

$$\begin{aligned}\int_{t_0}^{t_f} g_{\dot{x}} \cdot \delta\dot{x}(t)dt &= g_{\dot{x}}(t_f)\delta x(t_f) - \int_{t_0}^{t_f} \frac{d}{dt} g_{\dot{x}} \cdot \delta x(t) \\ &= g_{\dot{x}}(t_f)(\delta x_f - \dot{x}\delta t_f) - \int_{t_0}^{t_f} \frac{d}{dt} g_{\dot{x}} \cdot \delta x(t)\end{aligned}$$

which results in

$$\delta J_a = [w_x(t_f) + g_{\dot{x}}]\delta x_{t_f} + [w_{t_f}(t_f) + g(t_f) - g_{\dot{x}}(t_f)\dot{x}(t_f)]\delta t_f + \delta\nu^T\psi(t_f) + \int_{t_0}^{t_f} [g_x - \frac{d}{dt} g_{\dot{x}}]\delta x(t)]dt,$$

The necessary conditions require that $\delta J_a = 0$ for arbitrary $\delta x(t), \delta\nu$ which is only possible if the following necessary conditions hold:

$$\nabla_x w(x(t_f), \nu, t_f) + \nabla_{\dot{x}} g(x(t_f), \dot{x}(t_f), t_f) = 0, \tag{1}$$

$$\frac{\partial}{\partial t_f} w(x(t_f), \nu, t_f) + g(x(t_f), \dot{x}(t_f), t_f) - \nabla_{\dot{x}} g(x(t_f), \dot{x}(t_f), t_f)^T \dot{x}(t_f) = 0, \tag{2}$$

$$\psi(x(t_f), t_f) = 0 \tag{3}$$

$$\nabla_x g(x(t), \dot{x}(t), t) - \frac{d}{dt} \nabla_{\dot{x}} g(x(t), \dot{x}(t), t) = 0, \qquad t \in (t_0, t_f). \tag{4}$$