Lecture 2-1, Optimal Control
CDS 110b, 14 Jan 08

CALIFORNIA INSTITUTE OF TECHNOLOGY
Control and Dynamical Systems

CDS 110b

R. M. Murray

Lecture 1 – Optimal Control

4 January 2006

This lecture provides an overview of optimal control theory. Beginning with a review of optimization, we introduce the notion of Lagrange multipliers and provide a summary of the Pontryagin's maximum principle.

1 Introduction: Optimization-Based Control

Skip intro material (already covered) and summarize optimal control problem instead.

Basic idea: two degree of freedom control design:

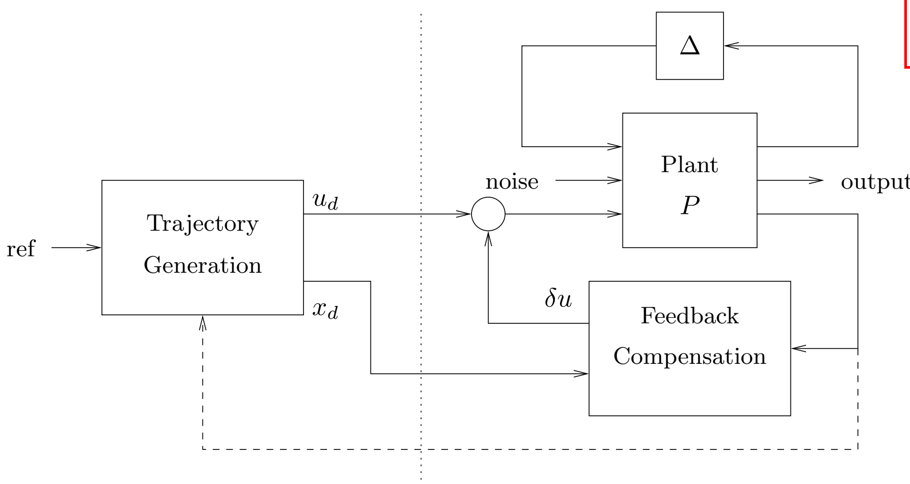

- Compute feasiable trajectory (x_d, u_d) for the system
- Use local controller to track feasible trajectory

Variation (later): update the trajectory based on the current state of the system $\Longrightarrow$ "receding horizon control".

2 Review: Optimization

Optimization of *functions*: given $F : \mathbb{R}^n \to \mathbb{R}$, find $x^* \in \mathbb{R}^n$ such that $F(x^*) \geq F(x)$ for all $x \in \mathbb{R}^n$.

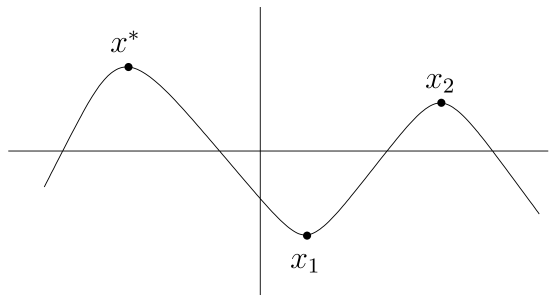

Necessary conditions:

$$\frac{\partial F}{\partial x}(x^*) = 0$$

These are *not* sufficient conditions; the points x_1 and x_2 and x^* all satisfy the necessary condition but only one is the (global) maximum.

Optimization with *constraints*: given $F : \mathbb{R}^n \to \mathbb{R}$ and $G_i : \mathbb{R}^n \to \mathbb{R}$, $i = 1, \ldots, k$, find $x^* \in \mathbb{R}^n$ such that $G_i(x^*) = 0$ (satisfies constraints) and

$$F(x^*) \geq F(x)$$

for all x such that $G_i(x) = 0$.

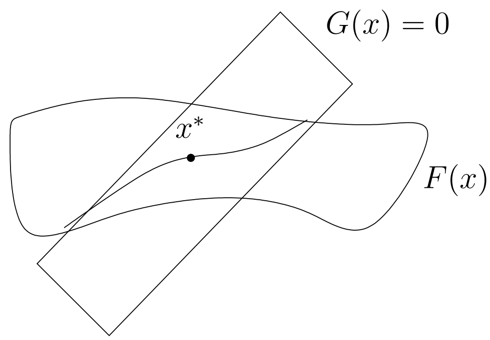

Necessary conditions: $\exists\ \lambda_i$, $i = 1, \ldots, k$ such that

$$\frac{\partial F}{\partial x}(x^*) + \sum_{i=1}^{k} \lambda_i \frac{\partial G_i}{\partial x}(x^*) = 0$$

or define $\widetilde{F} = F + \sum \lambda_i G_i \implies$

$$\frac{\partial \widetilde{F}}{\partial x}(x^*) = 0$$

The scalars λ_i are called *Lagrange multipliers*.

Intuition on Lagrange multiplers

View #1 (geometric):

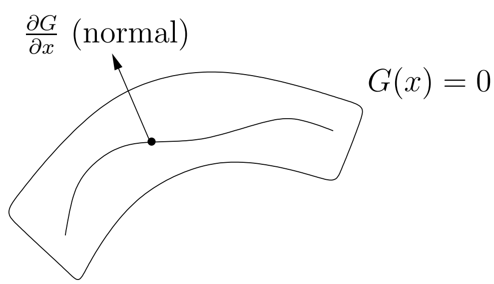

In the expression

$$\frac{\partial F}{\partial x} + \lambda^T \frac{\partial G}{\partial x} = 0$$

the term $\frac{\partial F}{\partial x}$ is the usual (gradient) optimality condition while the term $\frac{\partial G}{\partial x}$ is used to "cancel" the gradient in the directions normal to the constraint.

View #2 (algebraic): In the expression

$$\min_x F(x) + \lambda^T G(x)$$

the variables λ can be regarded as free variables $\implies$ need to choose x such that $G(x) = 0$. Otherwise, we can choose λ to generate a large cost.

Remarks:

1. We will often switch back and forth betweeen max and min:

$$\max_x F(x) = \min_x \big(-F(x)\big)$$

2. Very good software is available for solving optimization problems numerically: NPSOL, SNOPT, fmin (in MATLAB).

Add example:
F(x) = (x1-a)^2 + (x2-b)^2
G(x) = x1 - x2

3 Optimal Control of Systems

Consider now the *optimal control problem*:

$$\min_u \underbrace{\int_0^T L(x,u)\,dt}_{\text{integrated cost}} + \underbrace{V\big(x(T),u(T)\big)}_{\text{final cost}}$$

subject to the constraint

$$\dot{x} = f(x,u) \qquad x \in \mathbb{R}^n,\ u \in \mathbb{R}^m$$

This problem is equivalent to the "standard" problem of minimizing a cost function $J(z)$ where $z \in L_2[0,T]$ and $h(z) = 0$ models the dynamics.

Basic idea: optimize over *curves* $x(t)$, $u(t)$ that satisfy the *constraint* $\dot{x}(t) - f(x(t),u(t)) = 0$.

Remarks

1. A special case is *linear quadratic* optimizal control:

$$\dot{x} = Ax + Bu \qquad J = \tfrac{1}{2}\int_0^T \left(x^T Q x + u^T R u\right)\,dt + x^T(T) P_1 x(T)$$

$Q \geq 0$	penalizes state error (assumes $x_d = 0$)
$R > 0$	penalizes input (*must* be positive definite)
$P_1 > 0$	penalizes terminal state

2. Variation: let $T = \infty$, $V = 0 \implies$ "infinite horizon" problem

3. Use of optimal control does not require linear systems or quadratic cost. *Very* general results are available.

Maximum Principle (Pontryagin, 1960s)

System: nonlinear ODE

$$\begin{aligned} &\dot{x} = f(x,u) \quad x = \mathbb{R}^n \\ &x(0) \text{ given} \quad u \in \Omega \subset \mathbb{R}^p \end{aligned}$$

where $f(x,u) = (f_1(x,u), \ldots f_n(x,u)) : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^n$.

Cost: integral + terminal cost and constraint

$$\begin{aligned} &J = \int_0^T L(x,u)\,dt + V(x(T)) \\ &\psi(x(T)) = 0 \quad \leftarrow \text{terminal constraint} \end{aligned}$$

Hamiltonian:

$$H = L + \lambda^T f = L + \sum \lambda_i f_i$$

Theorem 1. *If (x^*, u^*) is optimal, then $\exists$ $\lambda^*(t)$ and ν^* such that*

$$\dot{x}_i = \frac{\partial H}{\partial \lambda_i} \qquad -\dot{\lambda}_i = \frac{\partial H}{\partial x_i} \qquad \underbrace{\begin{matrix} x(0) \text{ given, } \psi(x(T)) = 0 \\ \lambda(T) = \frac{\partial V}{\partial x}(x(T)) + \frac{\partial \psi^T}{\partial x}\nu \end{matrix}}_{\text{Boundary conditions } (2n \text{ total})}$$

and

$$H(x^*(t), u^*(t), \lambda^*(t)) \leq H(x^*(t), u, \lambda^*(t)) \quad \forall \quad u \in \Omega$$

Remarks

1. Note that even though we are *minimizing* the cost, this is still usually called the maximum principle (artifact of history).
2. This is a *very* general (and elegant) theorem. Generalizes optimization to optimal control.
3. If $u = \text{argmin} H(x,u,\lambda)$ exists, we can use this to *choose* control law u
4. If $\Omega = \mathbb{R}^m$ (unconstrained input) and H is differentiable, then a necessary condition is

$$\frac{\partial H}{\partial u} = 0$$

5. Proof: Lewis and Syrmos [1]