Optimal Control of Systems

Given a system: $\dot{x} = f(x, u, n)$; $x \in \mathbb{R}^n$, $u \in \Omega \subset \mathbb{R}^p$; $x(0) = x_0$

Find the control u(t) for $t \; \varepsilon \; [0,\mathrm{T}]$ such that

$$u = \arg \lim_{u \epsilon \Omega} \underbrace{\int_0^T \overbrace{L(x,u)}^{\textit{Instantaneous (Stage) Cost}} dt + \overbrace{V\big(x(T), u(T)\big)}^{\textit{Terminal Cost}}}_{J(x,u)}$$

- Can include constraints on control u and state x
  - (along trajectory or at final time):
- Final time *T* may or may not be free (we'll first derive fixed *T* case)

Function Optimization

- Necessary condition for optimality is that gradient is zero
  - Characterizes local extrema; need to check sign of second derivative
  - Need convexity to guarantee global optimum

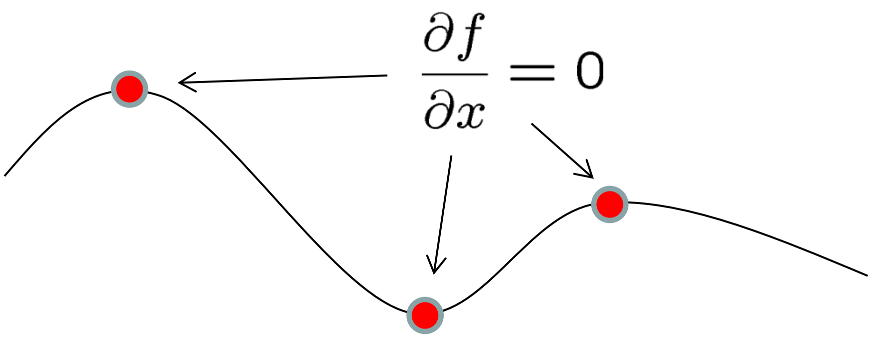

Keep this analogy in mind:
at an extremum

$$\lim_{\Delta x \to 0} \{ f(x^* + \Delta x) - f(x^*) \} = 0$$

$$\lim_{\Delta x \to 0} \left\{ f(x^*) + \frac{\partial f(x^*)}{\partial x} \Delta x + \ ... \ - f(x^*) \right\} = 0$$

Gradient of f(x)
orthogonal to level sets

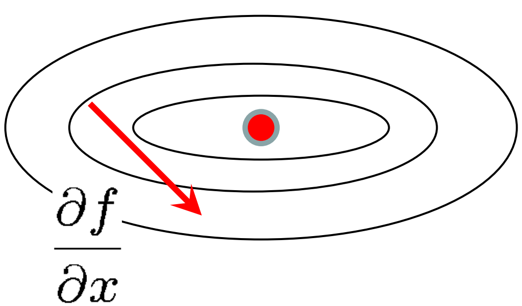

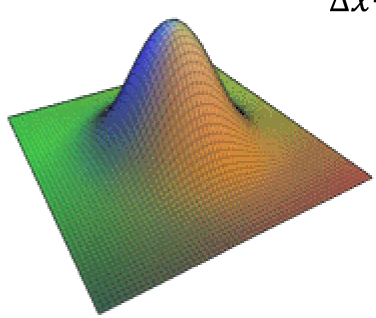

Constrained Function optimization

Given $F : \mathbb{R}^n \to \mathbb{R}$ and $G_i : \mathbb{R}^n \to \mathbb{R}$, $i = 1 \ldots k$, then find $x^* \in \mathbb{R}^n$ such that $G_i(x^*) = 0 \, \forall i$ and $F(x^*) \geq F(x)$ for all x satisfying $G_i(x) = 0 \, \forall i$.

- At the optimal solution, gradient of F(x) must be parallel to gradient of G(x):

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

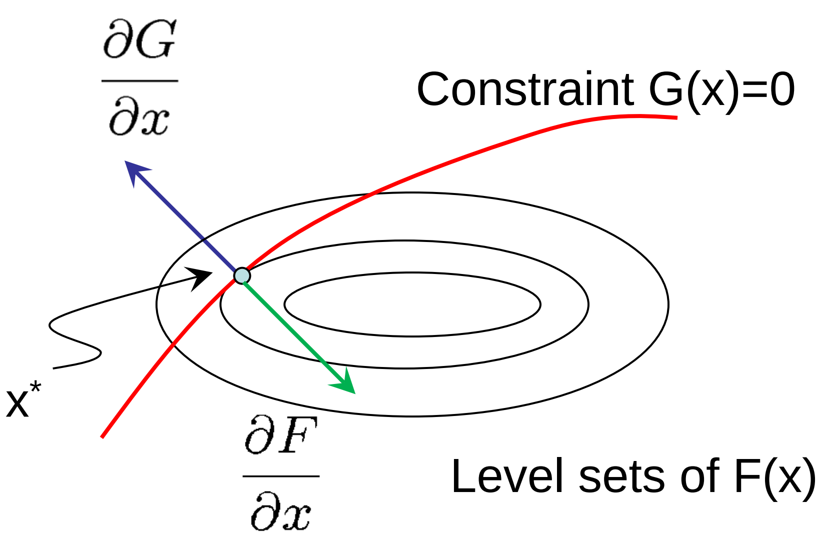

Consider:

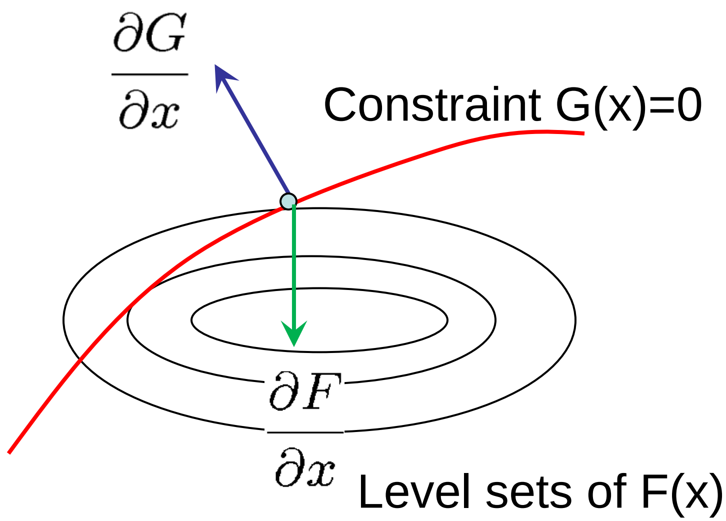

Constrained Function optimization

Given $F : \mathbb{R}^n \rightarrow \mathbb{R}$ and $G_i : \mathbb{R}^n \rightarrow \mathbb{R}$, $i = 1 \ldots k$, then find $x^* \in \mathbb{R}^n$ such that $G_i(x^*) = 0 \, \forall i$ and $F(x^*) \geq F(x)$ for all x satisfying $G_i(x) = 0 \, \forall i$.

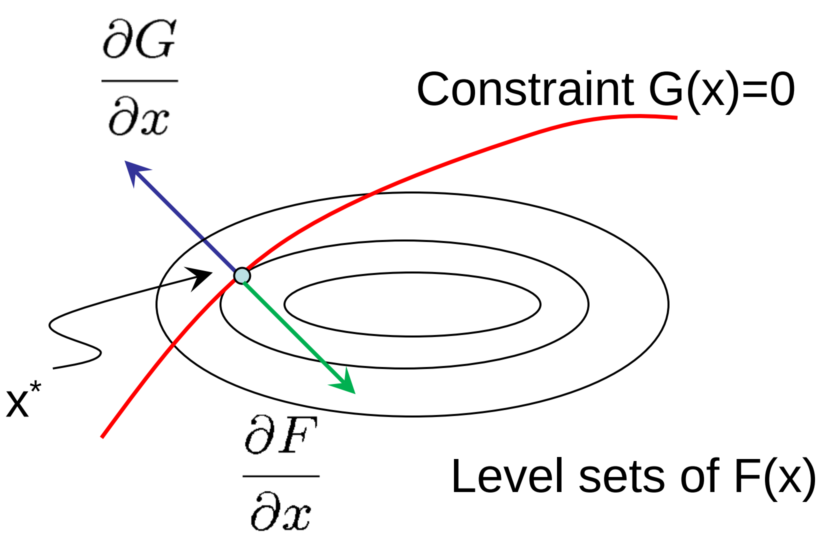

- At the optimal solution, gradient of F(x) must be parallel to gradient of G(x):

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

- More generally, define:

Lagrange multiplier

$$\tilde{F} = F + \lambda^T G$$

- Then a necessary condition is:

$$\frac{\partial \tilde{F}}{\partial \tilde{x}}(\tilde{x}^*) = 0 \qquad \tilde{x} = \begin{bmatrix} x \\ \lambda \end{bmatrix}$$

- The *Lagrange multipliers* λ are the sensitivity of the cost to a change in G

Solution approach

- Add Lagrange multiplier λ(t) for dynamic constraint
  - And additional multipliers for terminal constraints or state constraints
- Form augmented cost functional:

$$\begin{aligned}\tilde{J}(x,u,\lambda) &= J(x,u) + \int_0^T \lambda^T\left(f(x,u) - \dot{x}\right)\mathrm{d}t \\ &= \int_0^T \Big(L(x,u) + \lambda^T(f(x,u) - \dot{x})\Big)\,\mathrm{d}t + V(x(T)) \\ &= \int_0^T \Big(H(x,u,\lambda) - \lambda^T\dot{x}\Big)\,\mathrm{d}t + V(x(T))\end{aligned}$$

- where the ***Hamiltonian*** is: $H \triangleq L + \lambda^T f$
- Necessary condition for optimality: $\delta\tilde{J}$ vanishes for any perturbation (variation) in x, u, or λ about optimum:

"variations"

$$\left.\begin{array}{lll} x(t) = x^*(t) + \underline{\delta x(t)}; & u(t) = u^*(t) + \underline{\delta u(t)}; & \lambda(t) = \lambda^*(t) + \underline{\delta\lambda\,(t)}; \end{array}\right\}$$

Variations must satisfy path end conditions!

variation

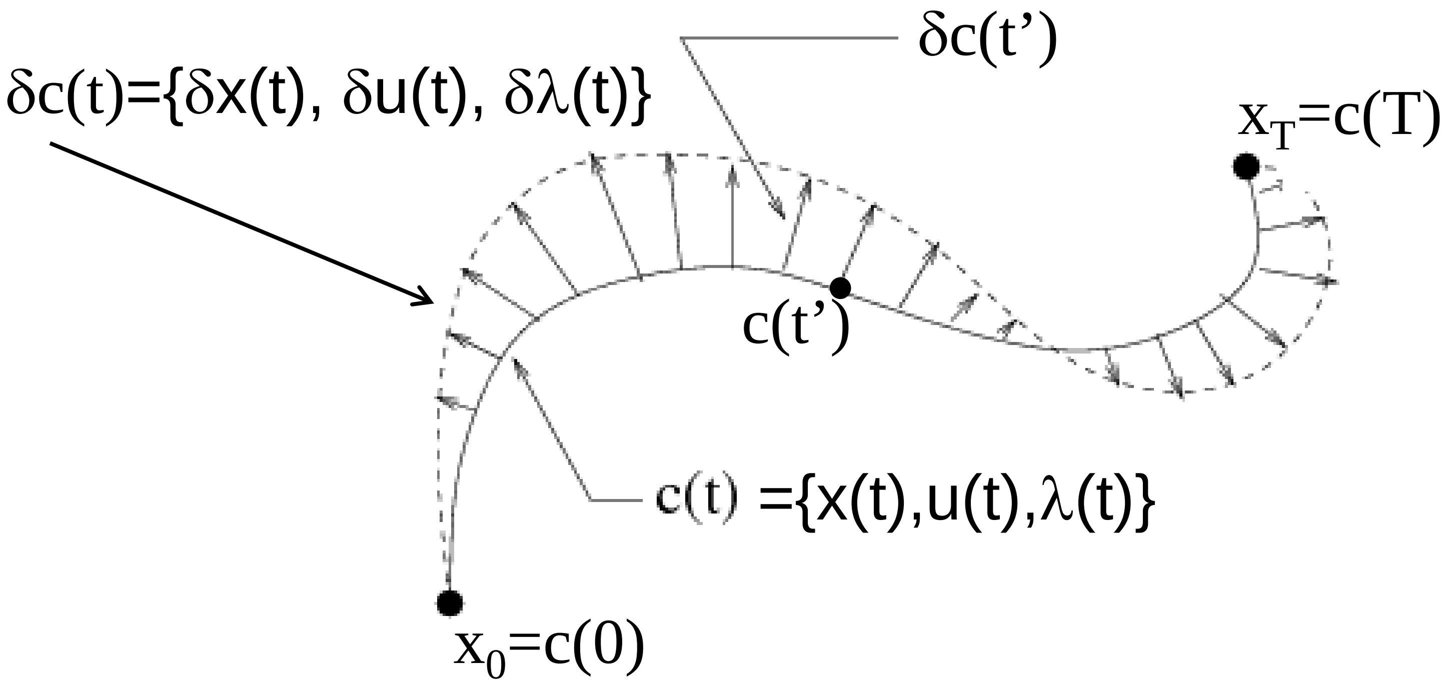

Derivation…

$$\delta\tilde{J} = \tilde{J}\,(x^* + \delta x, u^* + \delta u, \lambda^* + \delta\lambda) \;-\; \tilde{J}(x^*, u^*, \lambda^*)$$

$$\delta\tilde{J} \triangleq \tilde{J} - \tilde{J}^*$$

$$\simeq \int_0^T \left(\frac{\partial H}{\partial x}\delta x + \frac{\partial H}{\partial u}\delta u - \lambda^T\delta\dot{x} + \left(\frac{\partial H}{\partial \lambda} - \dot{x}^T\right)\delta\lambda\right)\mathrm{d}t + \frac{\partial V}{\partial x}\delta x(T)$$

- Note that (integration by parts):

$$\int_0^T \lambda^T\delta\dot{x} = -\int_0^T \dot{\lambda}^T\delta x + \lambda^T(T)\delta x(T) - \lambda^T(0)\delta x(0)$$

- So:

$$\delta\tilde{J} = \int_0^T \left[\left(\frac{\partial H}{\partial x} + \dot{\lambda}^T\right)\delta x + \frac{\partial H}{\partial u}\delta u + \left(\frac{\partial H}{\partial \lambda} - \dot{x}^T\right)\delta\lambda\right]\mathrm{d}t$$
$$+ \left(\frac{\partial V}{\partial x} - \lambda^T(T)\right)\delta x(T) + \lambda^T(0)\delta x(0)$$

We want this to be ***stationary*** for all variations

Pontryagin's Maximum Principle

- Optimal (x^*,u^*) satisfy:

$$\dot{x} = \left(\frac{\partial H}{\partial \lambda}\right)^T \qquad x(0) = x_0$$

$$-\dot{\lambda} = \left(\frac{\partial H}{\partial x}\right)^T \qquad \lambda(T) = \left(\left.\frac{\partial V}{\partial x}\right|_{x=x(T)}\right)^T$$

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

Optimal control is solution to O.D.E.

Unbounded controls

- If $\Omega = \mathbb{R}^m$ and H differentiable then $\partial H / \partial u = 0$
- Can be more general and include terminal constraints
- Follows directly from:

$$\delta \tilde{J} = \int_0^T \left[\underbrace{\left(\frac{\partial H}{\partial x} + \dot{\lambda}^T\right)}_{0} \delta x + \frac{\partial H}{\partial u} \delta u + \underbrace{\left(\frac{\partial H}{\partial \lambda} - \dot{x}^T\right)}_{0} \delta \lambda \right] \mathrm{d}t + \underbrace{\left(\frac{\partial V}{\partial x} - \lambda^T(T)\right)}_{0} \delta x(T) + \underbrace{\lambda^T(0) \delta x(0)}_{0}$$

Interpretation of λ

$$\dot{x} = \left(\frac{\partial H}{\partial \lambda}\right)^T \qquad x(0) = x_0 \qquad \leftarrow \quad \dot{x} = f(x,u)$$

$$-\dot{\lambda} = \left(\frac{\partial H}{\partial x}\right)^T \qquad \lambda(T) = \left(\left.\frac{\partial V}{\partial x}\right|_{x=x(T)}\right)^T$$

- Two-point boundary value problem: λ is solved backwards in time
- λ is the "co-state" (or "adjoint" variable)
- Recall that $H = L(x,u) + \lambda^T f(x,u)$
- If L=0, λ(t) is the sensitivity of the cost to a perturbation in state x(t)
  - In the integral as $\lambda(t)\delta\dot{x}$
  - Recall $\delta J = \ldots + \lambda(0)\delta x(0)$

$$x(\tau^+) = x(\tau^-) + \epsilon$$
$$\Rightarrow \delta\dot{x} = \epsilon\delta_D(t-\tau)$$
$$\Rightarrow \delta\tilde{J} = \int \cdots \lambda^T \delta\dot{x} \cdots = \lambda^T(\tau)\epsilon$$

Terminal Constraints

Assume ***q*** terminal constraints of the form: $\psi(x(T))=0$

- Then

$$\lambda(T) = \left(\frac{\partial V}{\partial x}\right)\left(x(T)\right) + \left(\frac{\partial \Psi}{\partial x}\right)(x(T))\, v$$

- Where ν is a set of undetermined Lagrange Multipliers

- Under some conditions, ν is free, and therefore $\lambda(T)$ is free as well

When the final time T is free (i.e., it is not predetermined), then the cost function *J* must be stationary with respect to perturbations T: $T^* + \delta T$. In this case:

$$H(T) = 0$$

General Remarks

- PMP equations in general ***very hard to solve***
  - Linear system with quadratic costs
    - $\tilde{J} = \int_0^T (x^T Q x + u^T R u)\, dt + x^T(T) S x(T)$
    - Closed Form solution exists: u(t) = -K(t) x(t)
    - We will study this case in detail
- For some problems, PMP conditions can lend insight into the properties of the solution:
  - e.g., *bang-bang* control (we'll see this later)
  - Used to understand limits of performance, and characteristics…
    - **Adjoint Equation:** Weather forecasting: what measurements are most useful to make?
    - What perturbations most likely to lead to an El Niño event?
  - Real-time implementation of full non-linear limited to relatively simple systems (e.g. chemical plants, Mars entry/descent/landing (EDL),…)

Example: Bang-Bang Control

Consider time optimal control of linear system

- $\dot{x} = Ax + Bu$
- $x(0) = x_0$; $x(T) = x_F$; $\varphi(x(T)) = x_F - x(T)$
- $|u| \le 1$; $J = \int_0^T 1\, dt$ ← "minimum time control"

Apply PMP:

- $H = L + \lambda^T f = 1 + \lambda^T (Ax + Bu) = 1 + (\lambda^T A)x + (\lambda^T B)u$
- $\dot{x} = \left(\frac{\partial H}{\partial x}\right)^T = Ax + Bu$
- $u = \mathrm{argmin}(H) = -\,\mathrm{sgn}(\lambda^T B)$
- I.e., control u is +1 or -1 in value

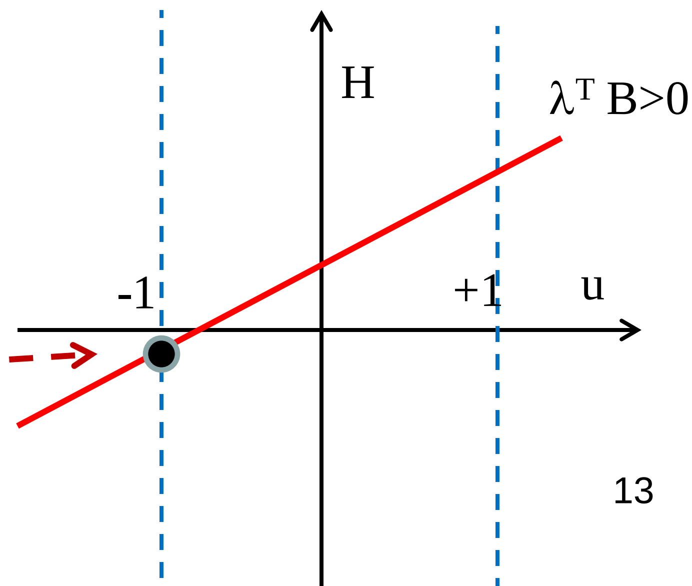

Since H is linear w.r.t. u,
minimization occurs at boundary

Linear system, Quadratic cost

$$\dot{x} = Ax + Bu \qquad J = \frac{1}{2}\int_0^T \left(x^T Q x + u^T R u\right) \mathrm{d}t$$

- Apply PMP:

$$\dot{\lambda} = -A^T\lambda + Qx$$
$$Ru = -B^T\lambda$$

B^T selects the part of the state that is influenced by u, so $B^T\lambda$ is sensitivity of aug. state cost to u

- Guess that $\lambda(t)=P(t)x(t)$:

$$-\dot{P} = PA + A^T P + Q - PBR^{-1}B^T P \qquad P(T) = 0$$
$$u = -R^{-1}B^T P x$$

- $x^T P x$ has an interpretation as the "cost to go"
- Often see the infinite-horizon solution where $dP/dt=0$