3/19: Sampling-based MPC

Announcements

- HW3 due Thurs, 3/26
- Project milestone due today
  - Late days allowed, one from each student.

Last time:

- constraints (SCP)
- MPC fundamentals

Today:

- sampling MPC

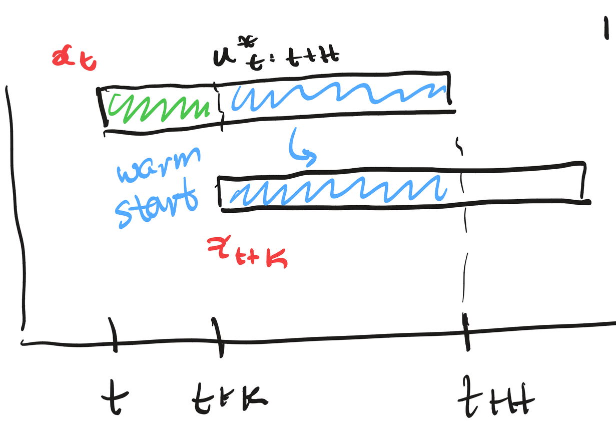

MPC:

1. solve OL optimal control problem from $x_t \to u^*_{t:t+H}$
2. apply first K $u^*_{t:t+K}$
3. replan from x_{t+K}

Two important considerations:

1. Terminal cost x_T
2. Loop rates / timing

Today: sampling-based MPC

Q: what happens when dynamics aren't differentiable? $\cancel{\frac{\partial f}{\partial x}}$ $\cancel{\frac{\partial f}{\partial u}}$

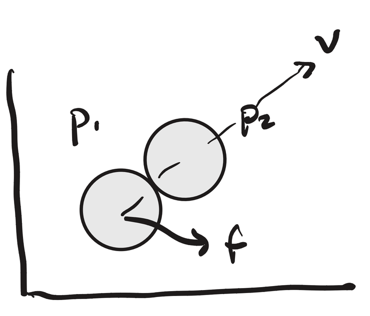

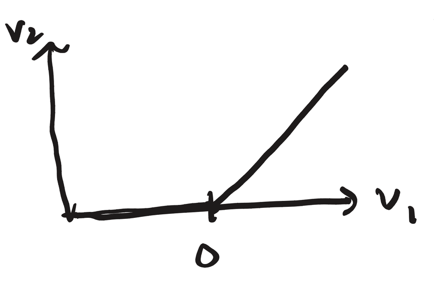

when p_1 not @ p_2

$\partial f / \partial u = 0$

Idea: use zero-order methods.

Problem: these usually don't handle constraints

$$\min_{x_{0:T},\, u_{0:T-1}} \quad c_T(x_T) + \sum_{t=0}^{T-1} c_t(x_t, u_t)$$

s.t. $x_{t+1} = f(x_t, u_t)$ ← necessary

$x_0 = x(0)$

$x_t \in \mathcal{X}$ $\quad u_t \in \mathcal{U}$

Idea: we can reparameterize!

$a \le u_t \le b$

$\tilde{u}_t = a + \sigma(u_t) \cdot (b-a$

implicit

$x_1 = f(x_0, u_0)$

$x_2 = f(x_1, u_1) = \tilde{f}(x_0, u_0, u_1)$

$c_2(x_2, u_2) = \tilde{c}_2(x_0, u_0, u_1, u_2)$

$$\boxed{\min_{u_{0:T-1}} \quad \tilde{c}_T(u_{0:T-1}) + \sum_{t=0}^{T-1} \tilde{c}_t(u_{0:t})}$$

"shooting methods"

$$\frac{\partial \tilde{c}_T}{\partial u_0} = \frac{\partial c_T}{\partial x_T} \cdot \frac{\partial x_T}{\partial x_{T-1}} \cdots \frac{\partial x_1}{\partial u_0}$$

$$\frac{\partial c_T}{\partial x_T} \cdot A_{T-1} \cdot A_{T-2} \cdot \cdots \cdot B_0$$

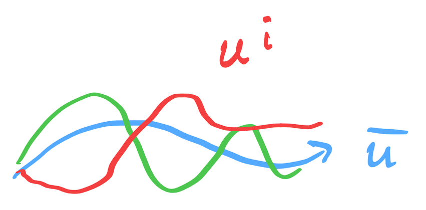

Alg 1 (Sampling MPC)

Input: an initial guess $\bar{u}_{0:T-1}$, x_0

repeat until convergence:

$\quad u^i \sim p(\cdot \mid \bar{u}_{0:T-1})$

$\quad$ for $t = 0, \ldots, t+H$ $\quad$ // in parallel

$\quad\quad x^i_{t+1} \leftarrow f(x^i_t, u^i_t)$

$\quad$ end for

$\quad J^i \leftarrow c_T(x^i) + \sum_{t=0}^{T-1} c_t(x^i_t, u^i_t)$

$\quad \bar{u} \leftarrow \text{update}(u^i, J^i)$

Three ways to implement this:

1. Predictive sampling [MJPC, Howell]

$$u^i = \bar{u} + \varepsilon^i, \quad \varepsilon^i \sim \mathcal{N}(0, \Sigma) \quad (\bar{u} \in \mathbb{R}^{H \cdot m})$$

$$\Sigma = \begin{bmatrix} \sigma^2 & & \\ & \ddots & \\ & & \sigma^2 \end{bmatrix} \quad (\text{detail: } \varepsilon^0 = 0)$$

$$\bar{u} = u^{i*}, \quad i* = \arg\min_i \{J^i\}$$

2. Model-Predictive Path Integral Control (MPPI)

$$u^i = \bar{u} + \varepsilon^i, \quad \varepsilon^i \sim \mathcal{N}(0, \Sigma)$$

$$\bar{u} = \frac{\sum_{i=1}^N w^i u^i}{\sum_{i=1}^N w^i} \qquad w^i = \exp\left(-\frac{1}{\lambda} J^i\right)$$

Temperature: $\lambda > 0$

$\lambda \to 0$, PS

$\lambda \to \infty$, uniform weights

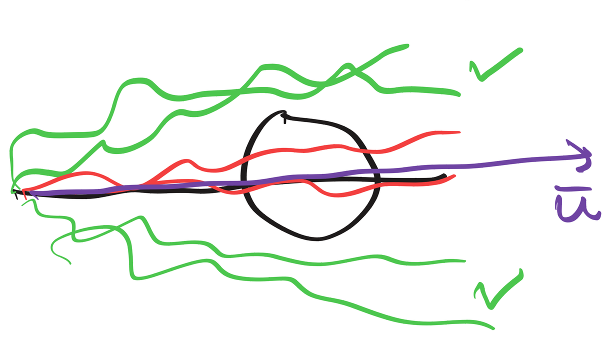

3. Cross-Entropy Method

$\mu_k, \Sigma_k,$ $\quad u^i \sim \mathcal{N}(\mu_k, \Sigma_k)$

$\quad = \mu_k + \Sigma_k^{1/2} \varepsilon^i, \quad \varepsilon^i \sim \mathcal{N}(0, I)$

$$\mathcal{I} = \{ i \mid J^i \in \text{bottom } 10\% \}$$

$$\mu_{k+1} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} u^i$$

$$\Sigma_{k+1} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} (u^i - \mu_{k+1})(u^i - \mu_{k+1})^T$$