Lecture 3
Infinite horizon linear quadratic regulator

- infinite horizon LQR problem
- dynamic programming solution
- receding horizon LQR control
- closed-loop system

Infinite horizon LQR problem

discrete-time system $x_{t+1} = Ax_t + Bu_t$, $x_0 = x^{\text{init}}$

problem: choose $u_0, u_1, \ldots$ to minimize

$$J = \sum_{\tau=0}^{\infty} \left(x_\tau^T Q x_\tau + u_\tau^T R u_\tau \right)$$

with given constant state and input weight matrices

$$Q = Q^T \geq 0, \qquad R = R^T > 0$$

. . . an *infinite dimensional problem*

problem: it's possible that $J = \infty$ for all input sequences $u_0, \ldots$

$$x_{t+1} = 2x_t + 0u_t, \qquad x^{\rm init} = 1$$

let's assume (A, B) is controllable

then for any $x^{\rm init}$ there's an input sequence

$$u_0, \ldots, u_{n-1},\ 0,\ 0, \ldots$$

that steers x to zero at $t = n$, and keeps it there

for this u, $J < \infty$

and therefore, $\min_u J < \infty$ for any $x^{\rm init}$

Dynamic programming solution

define **value function** $V : \mathbf{R}^n \to \mathbf{R}$

$$V(z) = \min_{u_0,\ldots} \sum_{\tau=0}^{\infty} \left(x_\tau^T Q x_\tau + u_\tau^T R u_\tau \right)$$

subject to $x_0 = z$, $x_{\tau+1} = Ax_\tau + Bu_\tau$

- $V(z)$ is the minimum LQR cost-to-go, starting from state z
- doesn't depend on time-to-go, which is always ∞; infinite horizon problem is *shift invariant*

Hamilton-Jacobi equation

fact: V is quadratic, *i.e.*, $V(z) = z^T P z$, where $P = P^T \geq 0$
(can be argued directly from first principles)

HJ equation:

$$V(z) = \min_w \left(z^T Q z + w^T R w + V(Az + Bw) \right)$$

or

$$z^T P z = \min_w \left(z^T Q z + w^T R w + (Az + Bw)^T P (Az + Bw) \right)$$

minimizing w is $w^* = -(R + B^T P B)^{-1} B^T P A z$

so HJ equation is

$$\begin{aligned} z^T P z &= z^T Q z + w^{*T} R w^* + (Az + Bw^*)^T P (Az + Bw^*) \\ &= z^T \left(Q + A^T P A - A^T P B (R + B^T P B)^{-1} B^T P A \right) z \end{aligned}$$

this must hold for all z, so we conclude that P satisfies the ARE

$$P = Q + A^T PA - A^T PB(R + B^T PB)^{-1} B^T PA$$

and the optimal input is constant state feedback $u_t = Kx_t$,

$$K = -(R + B^T PB)^{-1} B^T PA$$

compared to finite-horizon LQR problem,

- value function and optimal state feedback gains are time-invariant
- we don't have a recursion to compute P; we only have the ARE

fact: the ARE has only one positive semidefinite solution P

$i.e.$, ARE plus $P = P^T \geq 0$ uniquely characterizes value function

consequence: the Riccati recursion

$$P_{k+1} = Q + A^T P_k A - A^T P_k B(R + B^T P_k B)^{-1} B^T P_k A, \qquad P_1 = Q$$

converges to the unique PSD solution of the ARE
(when (A, B) controllable)

(later we'll see direct methods to solve ARE)

thus, infinite-horizon LQR optimal control is same as steady-state finite horizon optimal control

Receding-horizon LQR control

consider cost function

$$J_t(u_t, \ldots, u_{t+T-1}) = \sum_{\tau=t}^{\tau=t+T} \left(x_\tau^T Q x_\tau + u_\tau^T R u_\tau \right)$$

- T is called *horizon*
- same as infinite horizon LQR cost, truncated after T steps into future

if $(u_t^*, \ldots, u_{t+T-1}^*)$ minimizes J_t, u_t^* is called (T-step ahead) *optimal receding horizon control*

in words:

- at time t, find input sequence that minimizes T-step-ahead LQR cost, starting at current time
- then use only the first input

example: 1-step ahead receding horizon control

find $u_t,\ u_{t+1}$ that minimize

$$J_t = x_t^T Q x_t + x_{t+1}^T Q x_{t+1} + u_t^T R u_t + u_{t+1}^T R u_{t+1}$$

first term doesn't matter; optimal choice for u_{t+1} is 0; optimal u_t minimizes

$$x_{t+1}^T Q x_{t+1} + u_t^T R u_t = (A x_t + B u_t)^T Q (A x_t + B u_t) + u_t^T R u_t$$

thus, 1-step ahead receding horizon optimal input is

$$u_t = -(R + B^T Q B)^{-1} B^T Q A x_t$$

. . . a constant state feedback

in general, optimal T-step ahead LQR control is

$$u_t = K_T x_t, \qquad K_T = -(R + B^T P_T B)^{-1} B^T P_T A$$

where

$$P_1 = Q, \qquad P_{i+1} = Q + A^T P_i A - A^T P_i B (R + B^T P_i B)^{-1} B^T P_i A$$

$i.e.$: same as the optimal finite horizon LQR control, $T - 1$ steps before the horizon N

- a constant state feedback
- state feedback gain converges to infinite horizon optimal as horizon becomes long (assuming controllability)

Closed-loop system

suppose K is LQR-optimal state feedback gain

$$x_{t+1} = Ax_t + Bu_t = (A + BK)x_t$$

is called *closed-loop system*

($x_{t+1} = Ax_t$ is called *open-loop system*)

is closed-loop system stable? consider

$$x_{t+1} = 2x_t + u_t, \qquad Q = 0, \qquad R = 1$$

optimal control is $u_t = 0x_t$, *i.e.*, closed-loop system is unstable

fact: if (Q, A) observable and (A, B) controllable, then closed-loop system is stable