EE365: Linear Quadratic Regulator

Linear quadratic regulator

- $x_{t+1} = A_t x_t + B_t u_t + w_t$
- $\mathbf{E}\, w_t = 0$, $\mathbf{E}\, w_t w_t^\mathsf{T} = W_t$
- stage cost is (convex quadratic)
$$\frac{1}{2}(x_t^\mathsf{T} Q_t x_t + u_t^\mathsf{T} R_t u_t)$$
with $Q_t \geq 0$, $R_t > 0$
- terminal cost $\frac{1}{2} x_T^\mathsf{T} Q_T x_T$, $Q_T \geq 0$
- variation: terminal constraint $x_T = 0$

Linear quadratic regulator: DP

- value functions are quadratic plus constant (linear terms are zero):

$$v_t(x) = \frac{1}{2}(x^\mathsf{T} P_t x + r_t)$$

- $P_T = Q_T$, $r_T = 0$

- optimal expected tail cost:

$$\begin{aligned} &\mathbf{E}\, v_{t+1}(f_t(x, u, w_t)) \\ &\quad = \tfrac{1}{2}(r_{t+1} + \mathbf{E}(A_t x + B_t u + w_t)^\mathsf{T} P_{t+1}(A_t x + B_t u + w_t)) \\ &\quad = \tfrac{1}{2}(r_{t+1} + (A_t x + B_t u)^\mathsf{T} P_{t+1}(A_t x + B_t u) + \mathbf{Tr}(P_{t+1} W_t)) \end{aligned}$$

  using $\mathbf{E}\, w_t = 0$ and

$$\mathbf{E}\, w_t^\mathsf{T} P_{t+1} w_t = \mathbf{E}\, \mathbf{Tr}(P_{t+1} w_t w_t^\mathsf{T}) = \mathbf{Tr}(P_{t+1} W_t)$$

Linear quadratic regulator: DP

- minimize over u to get optimal policy:

$$\begin{aligned}\mu_t(x) &= \operatorname*{argmin}_u \left(u^\mathsf{T} R_t u + u^\mathsf{T} B_t^\mathsf{T} P_{t+1} B_t u + 2(B_t^\mathsf{T} P_{t+1} A_t x)^\mathsf{T} u\right) \\ &= -\left(R_t + B_t^\mathsf{T} P_{t+1} B_t\right)^{-1} B_t^\mathsf{T} P_{t+1} A_t x \\ &= K_t x\end{aligned}$$

- optimal policy is linear (as opposed to affine)

- using $u = K_t x$ we then have

$$\begin{aligned}v_t(x) = \tfrac{1}{2}(&r_{t+1} + \mathbf{Tr}(P_{t+1} W_t) + x^\mathsf{T}(Q_t + K_t^\mathsf{T} R_t K_t)x + \\ &x^\mathsf{T}(A_t + B_t K_t)^\mathsf{T} P_{t+1}(A_t + B_t K_t)x)\end{aligned}$$

- so coefficients of v_t are

$$\begin{aligned}&P_t = Q_t + K_t^\mathsf{T} R_t K_t + (A_t + B_t K_t)^\mathsf{T} P_{t+1}(A_t + B_t K_t), \\ &r_t = r_{t+1} + \mathbf{Tr}(P_{t+1} W_t)\end{aligned}$$

Linear quadratic regulator: Riccati recursion

- set $P_T = Q_T$
- for $t = T-1, \ldots, 0$

$$K_t = -(R_t + B_t^\mathsf{T} P_{t+1} B_t)^{-1} B_t^\mathsf{T} P_{t+1} A_t$$
$$P_t = Q_t + K_t^\mathsf{T} R_t K_t + (A_t + B_t K_t)^\mathsf{T} P_{t+1} (A_t + B_t K_t)$$

- called Riccati recursion; gives optimal policies, which are linear functions
- *surprise:* optimal policy does not depend on the disturbance distribution (provided it is zero mean)
- $J^\star = \frac{1}{2}(\mathbf{Tr}(P_0 X_0) + \sum_{t=0}^{T-1} \mathbf{Tr}(P_{t+1} W_t))$, where $X_0 = \mathbf{E}(x_0 x_0^\mathsf{T})$

Linear quadratic regulator: Example

- $n = 5$ states, $m = 2$ inputs, horizon $T = 31$
- A, B chosen randomly; A scaled so $\max_i |\lambda_i(A)| = 1$
- $Q_t = I$, $R_t = I$, $t = 0, \ldots, T-1$, $Q_T = 5I$
- $x_0 \sim \mathcal{N}(0, X_0)$, $X_0 = I$
- $w_t \sim \mathcal{N}(0, W)$, $W = 0.1I$

Linear quadratic regulator: Example

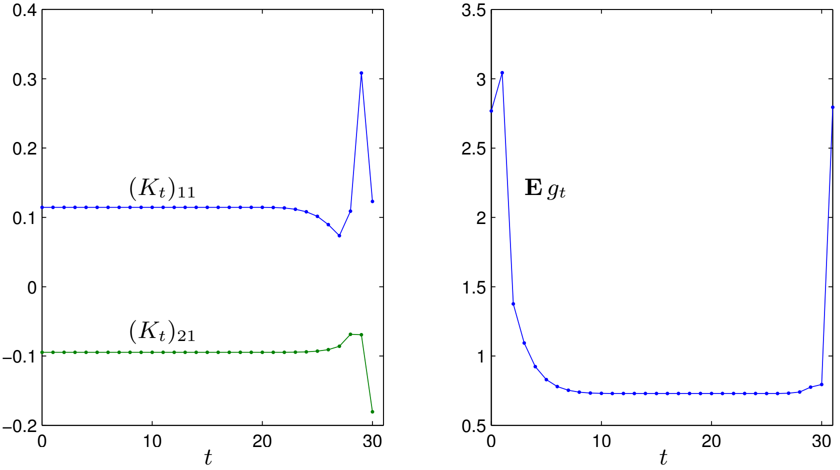

Linear quadratic regulator: Sample trajectory

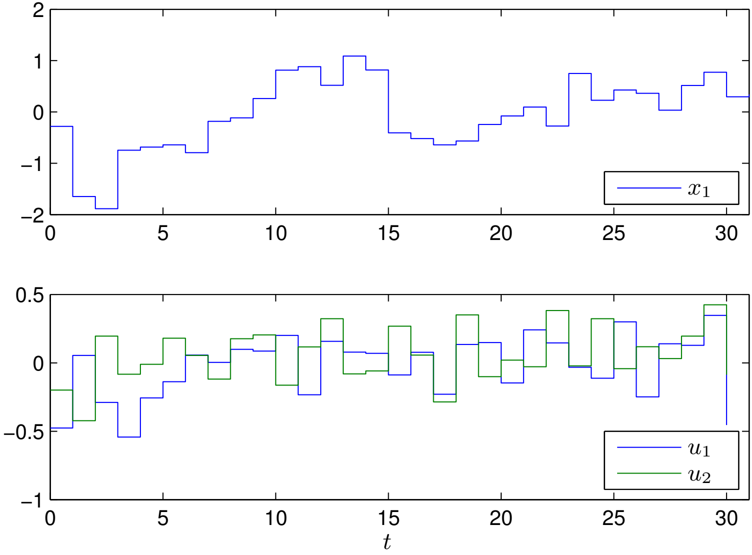

Linear quadratic regulator: Cost comparison

compare cost for

- optimal policy, $J^\star$
- prescient policy, J^{pre}: $w_0 \dots, w_T$ known in advance
- open loop policy, J^{ol}: choose $u_0, \dots, u_T$ with knowledge of x_0 only
- no control (1-step greedy), J^{nc}: $u_0, \dots, u_T = 0$

Linear quadratic regulator: Cost comparison

total stage cost histograms, $N = 5000$ Monte Carlo simulations