15-884/484 – Control 2: Dynamical Systems

J. Zico Kolter

October 31, 2013

Dynamical Systems

- Many extensions or special cases possible for the general dynamical system

$$x_{t+1} = f(x_t, u_t)$$

1. Linear systems
2. Partially observable systems
3. Differential algebraic equations
4. Stochastic systems, MDPs
5. Many others

Linear Dynamical Systems

- Already seen a brief mention of (discrete time) *linear* systems

$$x_{t+1} = Ax_t + Bu_t$$

- A simple extension: affine systems

$$x_{t+1} = Ax_t + Bu_t + a_t$$

- Continuous-time analogue

$$\dot{x} = Ax + Bu$$

- Often, one of the few cases where we *can* solve optimal control problems exactly

- More complex example: multi-room heating

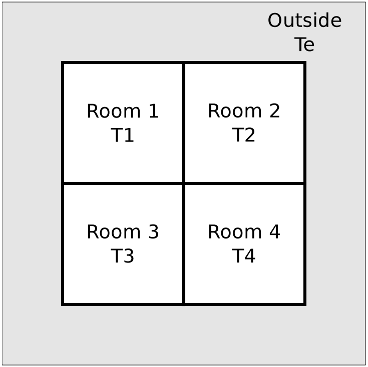

$$\dot{T}_1 = 2k(T_e - T_1) + k(T_2 - T_1) + k(T_3 - T_1) + du_1$$

- Let $x_i = T_i$, assume heaters/coolers in rooms one and two; dynamics are given by

$$x_{t+1} = x_t + \Delta t \begin{bmatrix} -4k & k & k & 0 \\ k & -4k & 0 & k \\ k & 0 & -4k & k \\ 0 & k & k & -4k \end{bmatrix} x_t$$

$$+ \ \Delta t \begin{bmatrix} 2kT_e \\ 2kT_e \\ 2kT_e \\ 2kT_e \end{bmatrix} + \Delta t \begin{bmatrix} d & 0 \\ 0 & d \\ 0 & 0 \\ 0 & 0 \end{bmatrix} u_t$$

$$x_t \in \mathbb{R}^4, \ \ u_t \in \mathbb{R}^2$$

- Interesting question: with just these two heaters/coolers, can we control the temperature any way desired?

Questions for linear systems

- **Controllability:** Given linear system

$$x_{t+1} = Ax_t + Bu_t$$

is it possible to reach some state $x^\star$ from some initial state x_1 in k time steps?

$$\begin{aligned} x_{k+1} &= (A(Ax_1 + Bu_1) + Bu_2 + \ldots) \\ &= A^k x_1 + \begin{bmatrix} A^{k-1}B & A^{k-2}B & \cdots & B \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_k \end{bmatrix} \end{aligned}$$

- Can reach any state after k steps if

$$\begin{bmatrix} A^{k-1}B & A^{k-2}B & \cdots & B \end{bmatrix}$$

is full rank (can find n independent columns)

- **Stability:** Given some controller $u_t = Kx_t$, does the *autonomous system*

$$x_{t+1} = Ax_t + BKx_t \equiv \tilde{A}x_t$$

go to zero as or does it diverge to ∞?

- Can understand this by looking at the *eigenvalues* of $\tilde{A}$

$$\tilde{A} = S\Lambda S^{-1} \text{ (Λ diagonal)}$$

- This implies that

$$\tilde{A}^k = S\Lambda S^{-1}S\Lambda S^{-1}\ldots = S\Lambda^n S^{-1}$$

- Since Λ is diagonal, Λ^k is just each entry raised to kth power
  - So, if all eigenvalues have $|\lambda_i| < 1$, $x_k \to 0$

- Lots of use for theory of linear systems, even for systems that aren't linear

- Consider non-linear system

$$x_{t+1} = f(x_t, u_t)$$

and *equilibrium point*

$$x^\star = f(x^\star, u^\star)$$

- Then the controllability/stability of the system around this point is the same as that of the linear system

$$A = D_x f(x^\star, u^\star), \;\; B = D_u f(x^\star, u^\star)$$

Partially observable systems

- In many situations, it is not possible to directly observe all the state variables of a system
- Can formulate this as a partially observable system with output $y_t \in \mathbb{R}^p$

$$x_{t+1} = f(x_t, u_t)$$
$$y_t = h(x_t)$$

- Typically the case that $p < n$ (or h is not invertible), so we cannot directly obtain x_t given y_t

- A fundamental question for partially observable systems: given a *sequence* of measurements, can we determine the state of the system? (*state estimation*, *filtering*, many other names)

- Let x_1 be the initial (unknown) state of the system

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} h(x_1) \\ h(f(x_1, u_1)) \\ h(f(f(x_1, u_1), u_2)) \\ \vdots \end{bmatrix}$$

- For k observations, n unknowns (x_1) and kp knowns

- When $kp \geq n$, we might be able to recover state (for example using Newton's method)
  - Lots of subtleties involved for general non-linear systems

Partially observable linear systems

- As before, a very useful special case

$$x_{t+1} = Ax_t + Bu_t$$
$$y_t = Cx_t$$

- Example: room heating with four rooms and only two thermostats

$$y_t = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$$

- **Observability:** When can we estimate state in linear system?

- For simplicity, assume $u_t = 0$ (easy to extend to the case of known inputs also)

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} C \\ CA \\ CA^2 \\ \vdots \end{bmatrix} x_1$$

- We can identify the state of the system after observing k outputs so long as

$$\begin{bmatrix} C \\ CA \\ \vdots \\ CA^k \end{bmatrix}$$

is full rank (i.e., we can find n independent rows)

Differential Algebraic Equations

- Sometimes, we want to include both dynamics and equality constraints in a dynamical system

- Typically formulated in continuous time as a combination of both time derivatives and algebraic equations

$$\dot{x} = f(x, z, u)$$
$$0 = g(x, z, u)$$
$$x \in \mathbb{R}^n \text{ (state variables)}$$
$$z \in \mathbb{R}^p \text{ (state algebraic variables)}$$
$$f : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^m \to \mathbb{R}^n \text{ (state dynamics equations)}$$
$$g : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^m \to \mathbb{R}^p \text{ (state algebraic equations)}$$

- Solution idea: states x, controls u, and algebraic equations g determine the algebraic variables y

- Suppose we could directly invert g

$$y = \hat{g}^{-1}(x, u)$$

  then dynamics reduce to a pure differential equation

$$\dot{x} = f(x, \hat{g}^{-1}(x, u), u)$$

- In practice, often can't invert g, but we can use Newton's method to find y such that $g(x, y, u) = 0$, then plug this into f.

Example: generator and power flow dynamics

- Two relevant states for a generator

$$\theta = \text{ voltage angle },\ \ \omega = \text{ frequency}$$

[recall that generator will spin according to $v(t) = cos(\omega t + \theta)$]

- Dynamics of a generator are given by

$$\dot{\theta} = \omega - \omega^{\text{ref}}$$
$$\dot{\omega} = \frac{1}{2H}(p^{\text{mech}} - p^{\text{elec}})$$

where $H \in \mathbb{R}_+$ is generator *moment of inertia* (physical constant).

- Generator states and power must obey power flow

$$p_i^{\text{elec}} = \hat{v}_i \sum_{k=1}^{N} \hat{v}_k (G_{ik} \cos(\theta_i - \theta_k) + B_{ik} \sin(\theta_i - \theta_k))$$

$$q_i^{\text{elec}} = \hat{v}_i \sum_{k=1}^{N} \hat{v}_k (G_{ik} \sin(\theta_i - \theta_k) - B_{ik} \cos(\theta_i - \theta_k))$$

(where we use N now to denote the number of buses)

- Putting it all together, dynamics are DAE

$$\dot{\theta}_i = \omega_i - \omega^{\text{ref}}$$

$$\dot{\omega}_i = \frac{1}{2H}(p_i^{\text{mech}} - p_i^{\text{elec}})$$

$$p_i^{\text{elec}} = \hat{v}_i \sum_{k=1}^{N} \hat{v}_k (G_{ik}\cos(\theta_i - \theta_k) + B_{ik}\sin(\theta_i - \theta_k))$$

$$q_i^{\text{elec}} = \hat{v}_i \sum_{k=1}^{N} \hat{v}_k (G_{ik}\sin(\theta_i - \theta_k) - B_{ik}\cos(\theta_i - \theta_k))$$

where

$$\text{state variables} = \theta_i, \omega_i \;\; i \in \text{GEN}$$

$$\text{algebraic variables} = \begin{cases} p_i^{\text{elec}}, q^{\text{elec}_i} & i \in \text{GEN} \\ \hat{v}_i, \theta_i & i \in \text{LOAD} \end{cases}$$

$$\text{control inputs} = p_i^{\text{mech}}, \hat{v}_i, \;\; i \in \text{GEN}$$

Stochastic systems

- Often the state of a system does not evolve deterministically (or, we have only imperfect predictions of how it will evolve)

- These settings can be captured by *stochastic* dynamical systems

$$x_{t+1} = f(x_t, u_t) + \epsilon_t$$

  where $\epsilon_t \in \mathbb{R}^n$ is a *zero-mean* random variable

- A completely general model (assuming ϵ_t can depend on state and control)

- Issues of controllability and observability become even more difficult in general stochastic systems

- Basic problem: even if noise is zero mean, after putting it through the dynamics, it may not be zero mean

$$\mathbf{E}[\epsilon_t] = 0 \quad \not\Rightarrow \quad \mathbf{E}[f(x_t + \epsilon_t, u_t)] = f(x_t, u_t)$$

- In the general case, we need to maintain a general distribution $p(x_t)$ at each time

- An important special case (again): *linear Gaussian* systems (linear or affine dynamics with multi-variate Gaussian noise)

$$x_{t+1} = Ax_t + Bu_t + \epsilon_t, \;\; \epsilon_t \sim \mathcal{N}(0, \Sigma)$$

- Key fact: by linearity of expecattion, we do have

$$\begin{aligned} \mathbf{E}[x_{t+1}] &= \mathbf{E}[Ax_t + Bu_t + \epsilon_t] \\ &= Ax_t + Bu_t + \mathbf{E}[\epsilon_t] \\ &= Ax_t + Bu_t \end{aligned}$$

- In this case, optimal control strategy can be to just ignore stochasticity, proceed as if system were deterministic

- Linear Gaussian case also applies to partially observable setting

$$x_{t+1} = Ax_t + Bu_t + \epsilon_t \;\; \epsilon_t \sim \mathcal{N}(0, \Sigma)$$
$$y_t = Cx_t + q_t \;\; q_t \sim N(0, Q)$$

- In this case, we can never estimate x_t exactly from past measurements, but we can get the best possible estimate (in a least squares sense)
  - Known as the *Kalman filter*, just (weighted) least squares to estimate current state from past measurements

Markov Decision Processes

- A general description of non-linear stochastic systems, but where we can typically only get exact solutions for *discrete* states and controls

- **State**: $x_t \in \{1, 2, \ldots, n\}$

- **Control**: $u_t \in \{1, 2, \ldots, m\}$

- **Transition probabilities**: $P^u \in \mathbb{R}^{n \times n}$, $u = 1, \ldots, m$

$$p(x_{t+1} = j | x_t = i, u_t = k) = P_{ij}^k$$

- **Cost**: $C : \{1, \ldots, n\} \times \{1, \ldots, m\} \to \mathbb{R}$