Stochastic Adaptive Control (02421)

Lecture 4

Tobias K. S. Ritschel

Assistant Professor

Section for Dynamical Systems, DTU Compute

DTU Compute

Department of Applied Mathematics and Computer Science

Lecture Plan

1. ARX models
2. ARX prediction + control
3. ARX estimation
4. ARX model validation + adaptive control
5. ARMAX control
6. ARMAX estimation + adaptive control
7. Systems and control theory
8. Stochastic systems + Kalman filtering
9. SS estimation (recursive) + control
10. SS control
11. SS estimation (batch)
12. SS estimation (recursive)
13. SS nonlinear control

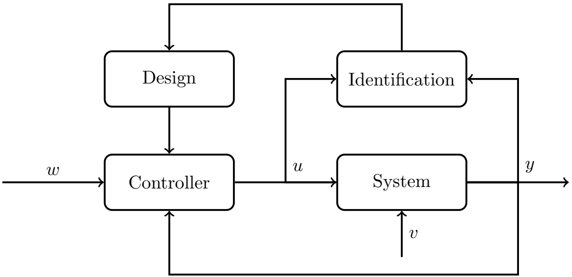

Today's Agenda

- Cautious controllers
- Dual adaptive controllers
- Suboptimal dual adaptive controllers

Example system

Known Systems and Control

ARX model with 1-step delay

$$A(q^{-1})y_t = B(q^{-1})u_{t-1} + e_t, \qquad e_t \sim N(0, \sigma^2) \tag{1}$$

Cost function

$$J = \mathbb{E}\left[\sum_{i=1}^{N}(y_{t+i} - w_{t+i})^2\right] \tag{2}$$

Example: MV_0 controller ($N = 1$)

$$u_{t-1} = \frac{1}{B}w_t - \frac{S}{B}y_{t-1} = \frac{1}{B}w_t - \frac{q(1-A)}{B}y_{t-1} \tag{3}$$

$$y_t = w_t + e_t \tag{4}$$

Known Systems and Control

Alternative form: System

$$y_t = \phi_t^T \theta + e_t = b_0 u_{t-1} + \varphi_t^T \vartheta + e_t \tag{5}$$

$$\phi_t^T = (-y_{t-1}, -y_{t-2}, \ldots, u_{t-1}, u_{t-2}, \ldots), \quad \theta^T = (a_1, a_2, \ldots, b_0, b_1, \ldots) \tag{6}$$

$$\varphi_t^T = (-y_{t-1}, -y_{t-2}, \ldots, 0, u_{t-2}, \ldots), \quad \vartheta^T = (a_1, a_2, \ldots, 0, b_1, \ldots) \tag{7}$$

Alternative form: Controller

$$u_{t-1} = \frac{1}{b_0} w_t - \frac{\varphi_t^T \vartheta}{b_0} \tag{8}$$

$$y_t = w_t + e_t \tag{9}$$

Relation between ϕ and φ

$$\varphi = \phi - \mathrm{diag}(l)\phi, \qquad \vartheta = \theta - \mathrm{diag}(l)\theta \tag{10}$$

$$l^T = (0, 0, \ldots, 1, 0, \ldots) \tag{11}$$

The nonzero entry corresponds to the position of b_0 and u_{t-1}

Adaptive control - Method Overview

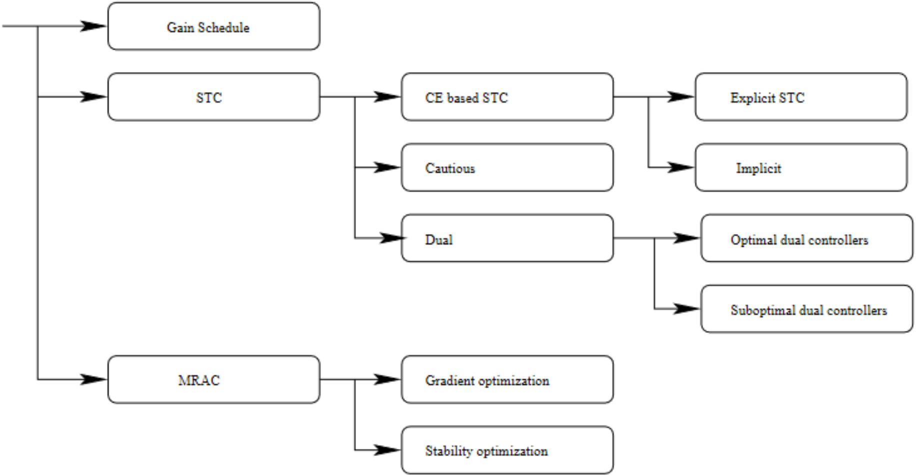

Adaptive Control - CE Control (Explicit)

Certainty equivalent self-tuner: Use the parameter estimate as the true parameters

$$y_t = \phi_t^T \theta + \epsilon_t, \quad \theta \to \hat{\theta} \tag{12}$$

Update parameter estimate

$$\epsilon_t = y_t - \phi_t^T \hat{\theta}_{t-1}, \tag{13}$$

$$s_t = 1 + \phi_t^T P_{t-1} \phi_t, \tag{14}$$

$$K_t = \frac{P_{t-1}\phi_t}{s_t}, \tag{15}$$

$$\hat{\theta}_t = \hat{\theta}_{t-1} + K_t \epsilon_t, \tag{16}$$

$$P_t = P_{t-1} - K_t s_t K_t^T \tag{17}$$

Redesign the control law

$$u_t = \frac{1}{\hat{b}_0} w_t - \frac{\varphi_t^T \hat{\vartheta}}{\hat{b}_0} \tag{18}$$

Cautious controllers

Adaptive Control - Cautious Control

Cautious adaptive control: Take estimation uncertainty into account

Conditional cost function

$$J = \mathbb{E}[(y_{t+1} - w_{t+1})^2 | Y_t] \tag{19}$$

$$= (\mathbb{E}[y_{t+1} - w_{t+1} | Y_t])^2 + \mathrm{Var}(y_{t+1} - w_{t+1} | Y_t) \tag{20}$$

Uncertainty of parameter estimate

$$\hat{\theta}_t \sim N(\theta, P_t) \tag{21}$$

$$\hat{b}_{0,t} = l^T \hat{\theta}_t \tag{22}$$

$$p_{b,t} = l^T P_t l \tag{23}$$

Control law

$$u_t = \frac{\hat{b}_{0,t}^2}{\hat{b}_{0,t}^2 + p_{b,t}} \left(\frac{w_{t+1} - \varphi_t^T \hat{\theta}}{\hat{b}_{0,t}} - \frac{\varphi_t^T P_t l}{\hat{b}_{0,t}^2} \right) \tag{24}$$

Adaptive Control - Cautious Control

If $P_t \to 0$, cautious control is equivalent to certainty equivalent control

$$u_t = \frac{\hat{b}_{0,t}^2}{\hat{b}_{0,t}^2 + p_{b,t}} \left(\frac{w_{t+1} - \varphi_t^T \hat{\theta}}{\hat{b}_{0,t}} - \frac{\varphi_t^T P_t l}{\hat{b}_{0,t}^2} \right) \quad \to \quad u_t = \frac{w_{t+1} - \varphi_t^T \hat{\theta}}{\hat{b}_{0,t}} \tag{25}$$

If $\hat{\theta}_t \to \theta$, certainty equivalent control is equivalent to the known control

1. $P_t \to 0$: Cautious = CE $\neq$ known
2. $\hat{\theta}_t \to \theta$: Cautious $\neq$ CE = known
3. $P_t \to 0$, $\hat{\theta}_t \to \theta$: Cautious = CE = known

Adaptive Control - Cautious Control - Usage

Cautious controller

$$u_t = \frac{\hat{b}_{0,t}^2}{\hat{b}_{0,t}^2 + p_{b,t}} \left(\frac{w_{t+1} - \varphi_t^T \hat{\theta}}{\hat{b}_{0,t}} - \frac{\varphi_t^T P_t l}{\hat{b}_{0,t}^2} \right) \tag{26}$$

- *Turn-off phenomenon:* Control is dampened due to high uncertainty of b_0
- Consequence: Less information about b_0 for the next estimate, i.e., the uncertainty increases
- Turn-off usually occurs if b_0 or the control signal is small
- Conclusion: The cautious controller is useful for systems with constant or almost constant parameters, but unsuitable for general time-varying systems

Dual adaptive controllers

Adaptive Control - Optimal Dual Control

Dual control: Conditional expectation of the cost

$$J = \min_{U_t} \mathbb{E}\left[\sum_{i=1}^{N}(y_{t+i} - w_{t+i})^2\right] = \mathbb{E}_{Y_t}\left[\min_{U_t} \mathbb{E}\left[\sum_{i=1}^{N}(y_{t+i} - w_{t+i})^2 \middle| Y_t\right]\right] \tag{27}$$

If the parameter uncertainty is Gaussian, the conditional expectation is Gaussian (even if y_t is not)

$$\xi_t = [\varphi_{t-1}, \quad \hat{\theta}_t, \quad P_t] \tag{28}$$

contains the necessary information

If not Gaussian, it becomes computationally challenging to compute the hyper space and storage requirements increase

Adaptive Control - Optimal Dual Control

Bellman equation

$$V(\xi_t, t) = \min_{u_{t-1}} \mathbb{E}[(y_t - w_t)^2 + V(\xi_{t+1}, t)|Y_{t-1}] \tag{29}$$

The last (N'th) step is identical to the cautious controller

$$V(\xi_N, N) = \min_{u_{N-1}} \mathbb{E}[(y_N - w_N)^2|Y_{N-1}] \tag{30}$$

$$= (\varphi_{N-1}^T \theta_N - w_N)^2 + \sigma^2 + \varphi_{N-1}^T P_N \varphi_{N-1} - \frac{\hat{b}_{0,N} w_N - \varphi_{N-1}^T(\hat{b}_{0,N}\hat{\theta}_N + P_N l)}{\hat{b}_{0,N}^2 - p_{b,N}} \tag{31}$$

substituting into $V(\xi_{N-1}, N-1)$, the second last control can be computed, and so on

This is similar to the LQR – however, it does not have an analytical solution and must be solved numerically

Adaptive Control - Optimal Dual control

Fundamental paradox of adaptive control

1. Control objective: Small signals (control action)
2. Estimation: Large signals (probing action)

For the optimal N-step dual control problem, the solution is a compromise between these goals

1. Improved long-term estimation accuracy; sacrificing short-term loss
2. Probing adds active learning to the method

Cautious control ($N = 1$): The probing effects diminishes and any learning is "accidental"

Issue with dual control: Curse of dimensionality – the computational cost increases drastically with increasing hyperspace dimension and horizon

Sub-optimal dual adaptive controllers

Adaptive Control - Sub-Optimal Dual control

As optimal dual control is impractical, sub-optimal dual controllers exist. They are based on the cautious controller and fix the issue with turn-off

Various approaches

1. Constrain the uncertainty
2. Extend the loss function
3. Serial expansion of the loss function
4. Add perturbation signals to the control

Sub-Optimal Dual Control - Constrained One-step Controller

Constrained one-step controller (minimum distance to zero control)

$$u_t = \begin{cases} u_c & \text{if } |u_c| \geq |u_l| \\ u_l \operatorname{sign}(u_c) & \text{if } |u_c| < |u_l| \end{cases} \tag{32}$$

u_c is the cautious controller input and u_l is a lower limit determined by us

The constraints do not prevent turn-off, but add extra perturbation when it happens

Sub-Optimal Dual Control - Constrained Uncertainty

Alternatively, constrain the uncertainty

$$\mathrm{Tr}(P_{t+1}^{-1}) \geq M \tag{33}$$

or constrain only p_b

$$p_{b,t+1} \leq \begin{cases} \gamma \hat{b}_{0,t+1}^2 & \text{if } p_{b,t} \leq \hat{b}_{0,t}^2 \\ \alpha p_{b,t} & \text{otherwise} \end{cases} \tag{34}$$

Sub-Optimal Dual control - extended loss function

Add uncertainty to the cost function

$$J = \mathbb{E}[(y_{t+1} - w_{t+1})^2 + \rho f(P_{t+1})] \tag{35}$$

f can be formulated in many ways

1. $f(P_{t+1}) = p_{b,t+1}$
2. $f(P_{t+1}) = R_2 \frac{p_{b,t+1}}{p_{b,t}}$
3. $f(P_{t+1}) = -\frac{\det(P_t)}{\det(P_{t+1})}$
4. $f(P_{t+1}) = -\epsilon_{t+1}^2$

This might lead to multiple local minima, and numerical optimization is required. Alternatively, use a second order serial expansion (e.g., a Taylor expansion)

Sub-Optimal Dual control - Extended Loss Function - Example

Third version of f

$$J = \mathbb{E}\left[(y_{t+1} - w_{t+1})^2 - \rho \frac{\det(P_t)}{\det(P_{t+1})} \middle| Y_t \right] \tag{36}$$

Ratio between determinants

$$\frac{\det(P_t)}{\det(P_{t+1})} = 1 + \phi_{t+1}^T P_t \phi_{t+1} \tag{37}$$

Analytical control law

$$u_t = \frac{\hat{b}_0(w_{t+1} - \varphi_{t+1}^T \hat{\vartheta}_t) + \rho (P_t l)^T \varphi_{t+1}}{\hat{b}_0^2 - \rho p_{b,t}} \tag{38}$$

Depending on ρ, we get specific controllers

1. $\rho = 0$: the CE controller
2. $\rho = -1$: the cautious controller
3. $\rho > 0$: an active learning controller

Sub-Optimal Dual control - Probing

Add probing/perturbation signal

$$u_t = u_t^c + u_t^x \tag{39}$$

Possible probing/perturbation signals include

1. PRBS
2. DOX: Design of excitation signal

They can be applied both at certain points in time (low uncertainty) or constantly

Questions?