CDS 131 Homework 9: System Level Synthesis

Winter 2025

Due 3/14 at 11:59 PM

Instructions

This homework is divided into three parts:

1. *Optional Exercises: the exercises are entirely optional but are recommended to be completed before looking at the problems. They consist of easier, more computational questions to help you get a feel for the material.*
2. *Required Problems: the problems are the required component of the homework, and might require more work than the exercises to complete.*
3. *Optional Problems: the optional problems are some additional, recommended problems - some of these might go a little beyond the standard course material.*

All you need to turn in is the solutions to the required problems - the others are recommended but not required.

Preliminaries

We will consider linear, time-invariant systems with state equation

$$x_{t+1} = Ax_t + Bu_t + w_t \quad \forall\, t \in [0, \ldots, T], \tag{1}$$

over an interval $[0, T]$. The system level parameterization (SLP) for the closed-loop maps (CLMs) θ_x, θ_u is

$$\begin{aligned} \theta_x(t+1, \tau+1) &= A\theta_x(t, \tau) + B\theta_u(t, \tau)\,. \\ \theta_x(t, 0) &= I \quad \forall\, t \geq 0\,. \end{aligned} \tag{2}$$

The SLP is alternatively written in block matrix form as,

$$[I - \mathbf{A}, -\mathbf{B}] \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} = I, \tag{3}$$

where θ_x and θ_u are block lower triangular matrices containing $\theta_x(t, \tau)$ and $\theta_u(t, \tau)$ and $\mathbf{A}, \mathbf{B}$ are block subdiagonal matrices containing A and B on their block subdiagonals and zeros elsewhere.

1 Optional Exercises

1.1 No Disturbances

In the setting of no disturbances, that is, $w_t = 0$ for all t, what is the system level parameterization?

1.2 Feasibility and Controllability

Consider the no-disturbance setting, in which the control responds only to the initial condition. If we constrain the final state response matrix, $\theta_x(T)$, to be equal to zero, this means that the state can be driven from any state to zero in time T. Prove that if A, B are a controllable pair in time T, then the constraint (2) with the added constraint $\theta_x(T) = 0$ is feasible.

2 Required Problems

2.1 Disturbance Estimation

Consider dynamics of the form (1). Let $x = [x_0; \dots; x_T]$ and $w = [w_0; \dots; w_T]$. One implementation of SLS involves writing the control input in terms of the expected state, $\hat{x}$, and the state estimation error, $\hat{w}$. Let θ_x, θ_u satisfy the SLP (2) and

$$\begin{aligned} u_t &= \sum_{\tau=1}^{T} \theta_u(t,\tau)\hat{w}_{t-\tau} , \\ \hat{x}_{t+1} &= \sum_{\tau=1}^{T-1} \theta_x(t+1,\tau+1)\hat{w}_{t-\tau} , \\ \hat{w}_t &= x_{t+1} - \hat{x}_{t+1} . \end{aligned}$$

Note that $\theta_x(t,\tau) = 0$ for $\tau > t$, with the same holding for θ_u.

(a) Assume $x_0 = 0$. What is $\hat{w}_0$ in terms of the disturbance?

(b) Still assuming $x_0 = 0$, what is $\hat{w}_1$? Use a recursion argument to show that $\hat{w}_t = w_t$ for all $t \geq 0$.

(c) Now consider the setting where $x_0 \neq 0$. Shift the indices for the formulas of u_t, $\hat{x}_{t+1}$, and $\hat{w}_t$ so that

$$\begin{aligned} \hat{w}_t &= w_{t-1} \quad \forall\, t > 0 \\ \hat{w}_0 &= x_0 . \end{aligned}$$

2.2 Alternating Direction Method of Multipliers: Convergence

We will prove convergence for a quadratic cost function and linear constraint, corresponding to the LQR cost with the system level parameterization (SLP). This problem will be implemented in the next question. The optimization problem is

$$\begin{aligned} \min_{\theta_x,\theta_u} & \left\| \begin{bmatrix} \mathbf{Q}^{1/2} & \\ & \mathbf{R}^{1/2} \end{bmatrix} \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} \right\|_F^2 \\ \text{subject to } & [I - \mathbf{A}, -\mathbf{B}] \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} = I . \end{aligned}$$

Both the cost function and the constraint are column-separable, so this problem could be solved in one step. However, in the output feedback setting, the constraints are a combination of row and column separable constraints, so we would like to implement ADMM in a way that will generalize to output feedback.

The ADMM algorithm for this problem can be written in terms of a rho-separable problem, a column-separable problem, and an update step:

$$\begin{aligned} \phi^{k+1} &= \operatorname*{argmin}_{\phi} \left(h^{(r)}(\phi) + \frac{\rho}{2} \left\| \phi - \psi^k + \Lambda^k \right\|_F^2 \right) \\ \psi^{k+1} &= \operatorname*{argmin}_{\psi} \left(h^{(c)}(\psi) + \frac{\rho}{2} \left\| \psi - \phi^{k+1} - \Lambda^k \right\|_F^2 \right) \\ \Lambda^{k+1} &= \Lambda^k + \phi^{k+1} - \psi^{k+1} , \end{aligned}$$

where $h^{(r)}(\phi)$ is the row-separable component of the cost function with $h^{(r)}(\phi) = +\infty$ if ϕ does not satisfy the row-separable constraints, and $h^{(c)}(\psi)$ is the column-separable component of the cost function with $h^{(c)}(\psi) = +\infty$ if ψ does not satisfy the column-separable constraints. The original optimization problem is solved if $\phi = \psi$.

1. What is $h^{(c)}$ and $h^{(r)}$ for the LQR problem?

Let ϕ, ψ be the primal variables and Λ be the dual variable. Define the residual as $r^k = \phi^k - \psi^k$, and assume that some optimal $\phi^* = \psi^*$ exists with a corresponding Λ^*. Define the Lyapunov function $V^k = \rho \left\|\Lambda^k - \Lambda^*\right\|_F^2 + \rho \left\|\psi^k - \psi^*\right\|_F^2$.

2. It can be shown that $V^{k+1} \leq V^k - \rho \left\|r^{k+1}\right\|_F^2 - \rho \left\|\psi^{k+1} - \psi^k\right\|_F^2$. Use this to show that as $k \to \infty$, it holds that $r^k \to 0$ and $\left\|\psi^{k+1} - \psi^k\right\|_F^2 \to 0$.

3. [Optional] Sketch a proof of the convergence of $h^{(c)}(\psi^k) + h^{(r)}(\phi^k)$ to the optimal cost as $k \to \infty$. Hint: the calculation details for the inequalities for showing ADMM convergence can be found in Section 3 and Appendix A of these lecture notes. For details on ADMM for SLS, see Section 5.5 in Anderson 2019.

2.3 Alternating Direction Method of Multipliers: Coding

See Python notebook. Please attach a link to your solution code (e.g. on Google Drive, Github, etc.). *Do your best to complete the required coding portions—if you don't have prior experience with CVX or if you find yourself spending an overly large amount of time, you may submit your best guess for a solution.*

2.4 Alternating Direction Method of Multipliers: Constraints

See Python notebook. Please attach a link to your solution code (e.g. on Google Drive, Github, etc.). *Do your best to complete the required coding portions—if you don't have prior experience with CVX or if you find yourself spending an overly large amount of time, you may submit your best guess for a solution.*

2.5 Dynamic Programming

In this exercise, we will solve the for the closed-loop maps in closed-form using dynamic programming. For ease of notation, consider the no-disturbance setting, $w_t = 0$ for all $t \geq 0$. From exercise 1.1, we have the system level parametrization for this setting. Let the cost function be given by

$$C(\theta_x, \theta_u) = \sum_{\tau=0}^{T} \left\| \begin{bmatrix} Q_\tau^{1/2} \theta_x(\tau) \\ R_\tau^{1/2} \theta_u(\tau) \end{bmatrix} \right\|_F^2 .$$

(a) What is the state x_t and control u_t in terms of the closed-loop maps and initial conditions?

(b) Consider the cost-to-go at the last time step. In this step, we optimize over $\theta_u(T-1)$, which determines $\theta_x(T)$. Defining

$$V_{T-1}(\theta_u(T-1); \theta_x(T-1)) = \left\| \begin{bmatrix} Q_T^{1/2} \theta_x(T) \\ R_{T-1}^{1/2} \theta_u(T-1) \end{bmatrix} \right\|_F^2 = \left\| \begin{bmatrix} Q_T^{1/2}(A\theta_x(T-1) + B\theta_u(T-1)) \\ R_{T-1}^{1/2} \theta_u(T-1) \end{bmatrix} \right\|_F^2,$$

why is it unnecessary to account for $\theta_u(T)$? Solve the following optimization problem in closed-form, that is, write $\theta_u(T-1)$ as a function of $\theta_x(T-1)$:

$$\theta_u(T-1)(\theta_x(T-1)) = \underset{\tilde{\theta}_u(T-1)}{\operatorname{argmin}} V_{T-1}(\tilde{\theta}_u(T-1), \theta_x(T-1))$$

(c) The previous computation allows us to write V_{T-1} as a function of $\theta_x(T-1)$ only. Write the cost-to-go at time t, $V_t(\cdot)$, in terms of the current control closed loop map, $\theta_u(t)$ and state map $\theta_x(t)$. You do not need to solve explicitly.

(d) Explain in words how to determine all the closed-loop maps using the initial condition $\theta_x(0) = I$.

3 Optional Problem

3.1 Output Feedback

Consider the state and observer dynamics, where $w_t^{(x)}, w_t^{(y)}$ are the disturbances and observation noise, respectively,

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + w_t^{(x)} \\ y_t &= Cx_t + w_t^{(y)} \qquad\qquad \forall\, t \in [0, \ldots, T] \,. \end{aligned}$$

We now parameterize the state and control as the response to *both* the disturbances and observation noise,

$$\begin{bmatrix} x \\ u \end{bmatrix} = \begin{bmatrix} \theta_{xx} & \theta_{xy} \\ \theta_{ux} & \theta_{uy} \end{bmatrix} \begin{bmatrix} w^{(x)} \\ w^{(y)} \end{bmatrix} .$$

In this output feedback setting, the controller has access to only y_t, and not x_t. The controller is given by $K = \theta_{uy} - \theta_{ux}\theta_{xx}^{-1}\theta_{xy}$, and the SLP is

$$\begin{bmatrix} I - \mathbf{A} & -\mathbf{B} \end{bmatrix} \begin{bmatrix} \theta_{xx} & \theta_{xy} \\ \theta_{ux} & \theta_{uy} \end{bmatrix} = \begin{bmatrix} I & 0 \end{bmatrix} , \qquad \begin{bmatrix} \theta_{xx} & \theta_{xy} \\ \theta_{ux} & \theta_{uy} \end{bmatrix} \begin{bmatrix} I - \mathbf{A} \\ \mathbf{C} \end{bmatrix} = \begin{bmatrix} I \\ 0 \end{bmatrix} .$$

Prove that the trajectories (x, u) resulting from output feedback controllers is equivalent to those generated by $\theta_{xx}, \theta_{ux}, \theta_{uy}, \theta_{xy}$ which satisfy the SLP. This is the same idea as the theorem shown in class, here extended to output feedback.