Stochastic Adaptive Control (02421)

Lecture 3

Tobias K. S. Ritschel

Section for Dynamical Systems, DTU Compute

DTU Compute
Department of Applied Mathematics and Computer Science

Lecture Plan

1. Systems theory
2. Stochastics
3. **State estimation - Kalman filter 1**
4. State estimation - Kalman filter 2
5. Optimal control 1 - internal models
6. External models
7. Prediction + optimal control 1 - external models
8. Optimal control 2 - external models
9. System identification 1
10. System identification 2
11. System identification 3 + model validation
12. Adaptive control 1
13. Adaptive control 2

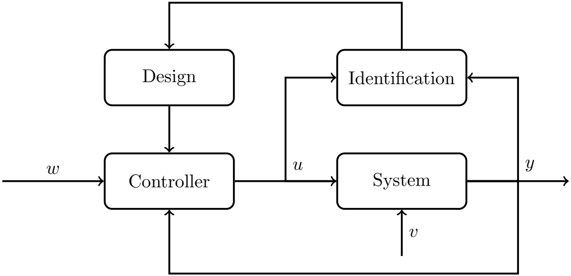

Today's Topics

- Follow-up from last lecture
- Filter theory
- State space estimation
- The Kalman filter

Follow-Up from Last Lecture

Question 2.6: Sketch the 95% confidence interval for the two last signals.
Answer: The uploaded Niveau.m function provides such plots.

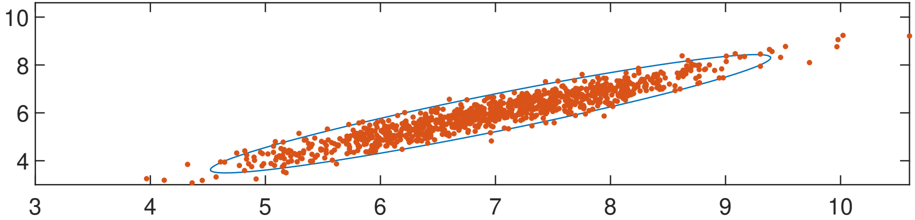

If we consider the variable

$$X \sim N(m, P), \tag{1}$$

the command becomes Niveau(m, P^{-1}, f) where f is a chi2 quantile level with 2 degrees of freedom.

Filters and Estimation

The purpose of state estimation, is to obtain an estimate $\hat{x}_t$ of the signal x_t, based on measurement data $Y_{t_0:t_1}$ from the time period t_0 to t_1. based on the noisy relation:

$$y_t = g(x_t, e_t) \tag{2}$$

Filter: A filter is an approach to estimating signals from known data by filtering out the noise.

We distinguish between filters based on the data period (t_0, t_f) and the time of interest t.

1. Smoothing $(t < t_f)$: Use both past and future data to estimate the states.
2. Filtering: $(t = t_f)$: Estimate the current states based on current and past data.
3. Prediction: $(t > t_f)$: Predict future states based on past data.

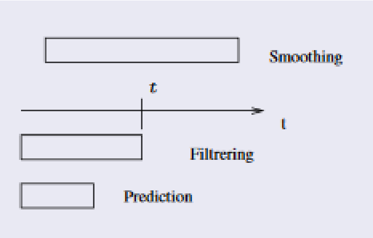

Filter Theory

For our discussion on filter theory, let us consider the discrete system

$$x_{t+1} = Ax_t + Bu_t + v_t, \qquad x_0 \in N(m_0, P_0), \qquad v_t \in N(0, R_1) \tag{3}$$

$$y_t = Cx_t + e_t, \qquad e_t \in N(0, R_2) \tag{4}$$

We will only consider filtering and prediction in this lecture because, in control, we will not have access to future data.

Filter Theory - The Good Estimate

Filter designs consist of 4 core concepts.

1. Characteristics of the signal and noise.
2. Observation model (relation between y, x, e).
3. Criterion (what is a good estimate).
4. Restrictions (what information is available).

Characteristics: The nature of the state and noises, or the dynamics of (3)

Observation: The relation of y, state x, and the noise or the outputs (4).

The criterion: We define a good estimate as one with minimum expected squared deviation from the truth, $||x - \hat{x}||^2$.

Restrictions: What set of data Y is available, i.e., do we want to filter, predict, or smooth.

The Filter Problem: A Good Filter?

1. Characteristics of the signal and noise.
2. Observation model (relation between y, x, e).
3. Criterion (what is a good estimate).
4. Restrictions (what information is available).

The first two items are given by the system. Therefore, let us consider item 3 and 4

Criterium:

$$J = E\{||x - \hat{x}||^2\} \tag{5}$$

Restrictions:

$$\hat{x} = func(Y) \tag{6}$$

The Filter Problem: A Good Estimator

The law of total expectation,

$$E\{g(x)\} = E_Y\{E\{g(x)|Y\}\} \tag{7}$$

allows us to analyze the nature of a good estimate:

$$\begin{aligned} J &= E\{(x-\hat{x})^T(x-\hat{x})\} && (8) \\ &= E_Y\{E\{(x-\hat{x})^T(x-\hat{x})|Y\}\} = E_Y(J_{in}) && (9) \end{aligned}$$

For a good estimate, the inner term should be constant with respect to changes in estimate

$$\begin{aligned} J_{in} &= E\{x^Tx - \hat{x}^Tx - x^T\hat{x} + \hat{x}^T\hat{x}|Y\} && (10) \\ &= E\{x^Tx|Y\} - \hat{x}^TE\{x|Y\} - E\{x|Y\}^T\hat{x} + \hat{x}^T\hat{x} && (11) \\ \frac{dJ_{in}}{d\hat{x}} &= 2\hat{x} - 2E\{x|Y\} = 0 && (12) \\ \hat{x} &= E(x|Y) && (13) \end{aligned}$$

Filter Theory - Projection Theorem

With the restriction for a good estimate found, we can consider how the information can be used.

But let us first consider the stochastic normal-distributed vector:

$$Z = \begin{bmatrix} X \\ Y \end{bmatrix} \in N\left(\begin{bmatrix} m_x \\ m_y \end{bmatrix}, \begin{bmatrix} P_x & P_{xy} \\ P_{xy}^T & P_y \end{bmatrix} \right) \tag{14}$$

Then, according to the **projection theorem**, the conditional distribution $X|Y \in N(\hat{x}, P_1)$ is given by

$$\hat{x} = m_x + P_{xy} P_y^{-1} (y - m_y) \tag{15}$$

$$P_1 = P_x - P_{xy} P_y^{-1} P_{xy}^T \tag{16}$$

$$X - \hat{x} \perp Y \tag{17}$$

As we will see, this is a useful relation in filter theory.

Filter Theory - Projection Theorem

$$Z = \begin{bmatrix} X \\ Y \end{bmatrix} \in N\left(\begin{bmatrix} m_x \\ m_y \end{bmatrix}, \begin{bmatrix} P_x & P_{xy} \\ P_{xy}^T & P_y \end{bmatrix} \right), \tag{18}$$

$$\hat{x} = m_x + P_{xy} P_y^{-1} (y - m_y) \tag{19}$$

$$P_1 = P_x - P_{xy} P_y^{-1} P_{xy}^T \tag{20}$$

Assume that X and Y are scalar.

1. What happens if we measure exactly the value we expected?
2. What happens if the measurement is an outlier?
3. What if X and Y are uncorrelated?
4. Can P_1 become negative?
5. What happens as P_x or P_y approach zero?

Think about it for yourself for one minute and then discuss with the person next to you for two minutes.

Filter Theory - Proof of Projection Theorem

Recall that the joint and conditional pdf of the normal distributions are

$$f_Z(z) = f_{X,Y}(x,y) = \frac{1}{\sqrt{Det(P_z)}\sqrt{(2\pi)^{n_x+n_y}}} e^{-\frac{1}{2}(z-m_z)^T P_z^{-1}(z-m_z)} \tag{21}$$

$$f_{X|Y}(x|y) = \frac{f_{X,Y}(x,y)}{f_Y(y)} \tag{22}$$

$$= \sqrt{\frac{Det(P_y)}{Det(P_z)(2\pi)^{n_x}}} e^{-\frac{1}{2}(z-m_z)^T P_z^{-1}(z-m_z)+\frac{1}{2}(y-m_y)^T P_y^{-1}(y-m_y)}$$

$$= \kappa e^{-\frac{1}{2}\alpha} \tag{23}$$

Using block inversion on the variance, P_z, we get

$$D = P_x - P_{xy} P_y^{-1} P_{xy}^T \tag{24}$$

$$P_z^{-1} = \begin{bmatrix} D^{-1} & -D^{-1}P_{xy}P_y^{-1} \\ -P_y^{-1}P_{xy}^T D^{-1} & P_y^{-1}(I + P_{xy}^T D^{-1} P_{xy} P_y^{-1}) \end{bmatrix} \tag{25}$$

$$\Rightarrow Det(P_z) = Det(P_y)Det(D) \tag{26}$$

Filter Theory - Proof of Projection Theorem

From the inversion, we obtain

$$\alpha = [x - (m_x + P_{xy}P_y^{-1}(y - m_y))]^T D^{-1}[x - (m_x + P_{xy}P_y^{-1}(y - m_y)] \tag{27}$$

$$\kappa = \frac{1}{\sqrt{Det(D)(2\pi)^{n_x}}} \tag{28}$$

Therefore, it is proven that the conditional distribution has the desired mean and variance:

$$E(X|Y) = m_{x|y} = m_x + P_{xy}P_y^{-1}(y - m_y) \tag{29}$$

$$Var(X|Y) = P_{x|y} = D = P_x - P_{xy}P_y^{-1}P_{xy}^T \tag{30}$$

For the claim of independence, we simply check the covariance:

$$CoV\{X - m_{x|y}, Y\} = CoV(X, Y) - P_{xy}P_y^{-1}CoV(Y, Y) \tag{31}$$

$$= P_{xy} - P_{xy}P_y^{-1}P_y = 0 \tag{32}$$

As X and Y are Gaussian, they are independent.

State Estimation

Now, let us see how the information can be used for the purpose of estimation.

We consider the system:

$$x_{t+1} = Ax_t + Bu_t + v_t \quad v_t \in N(0, R_1) \quad x_{t_0} \in N(\hat{x}_0, P_0) \tag{33}$$

$$y_t = Cx_t + e_t \quad e_t \in N(0, R_2) \quad e_t, v_t \textit{ white } \perp x_s \quad \forall s \leq t \tag{34}$$

First, we focus on the relation

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} | Y_{t-1} \in N\left(\begin{bmatrix} \times \\ \times \end{bmatrix}, \begin{bmatrix} \times & \times \\ \times & \times \end{bmatrix} \right), \quad Y_t = \begin{bmatrix} Y_{t-1} \\ y_t \end{bmatrix} \tag{35}$$

State Estimation 2: A Recursion

Let us define the conditional state distributions as

$$x_t|Y_{t-1} \in N(\hat{x}_{t|t-1}, P_{t|t-1}) \tag{36}$$

$$x_t|y_t, Y_{t-1} = x_t|Y_t \in N(\hat{x}_{t|t}, P_{t|t}) \tag{37}$$

Standard computation considering $y_t|Y_{t-1}$ provides:

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} |Y_{t-1} \in N\left(\begin{bmatrix} \hat{x}_{t|t-1} \\ C\hat{x}_{t|t-1} \end{bmatrix}, \begin{bmatrix} P_{t|t-1} & \times \\ \times & CP_{t|t-1}C^T + R_2 \end{bmatrix} \right) \tag{38}$$

State Estimation 3: Covariance

Consider the conditional covariance

$$CoV(y_t, x_t | Y_{t-1}) = CP_{t|t-1} \tag{39}$$

Then, the conditional distribution becomes

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} | Y_{t-1} \in N\left(\begin{bmatrix} \hat{x}_{t|t-1} \\ C\hat{x}_{t|t-1} \end{bmatrix}, \begin{bmatrix} P_{t|t-1} & P_{t|t-1}C^T \\ CP_{t|t-1} & CP_{t|t-1}C^T + R_2 \end{bmatrix} \right) \tag{40}$$

State Estimation 4: The Current Estimate

Now, for the estimation, remember that

$$x_t|y_t, Y_{t-1} = x_t|Y_t \in N(\hat{x}_{t|t}, P_{t|t}) \tag{41}$$

Then, using the projection theorem and our conditional distribution, we can obtain a state estimate:

$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + P_{t|t-1}C^T(CP_{t|t-1}C^T + R_2)^{-1}(y_t - C\hat{x}_{t|t-1}) \tag{42}$$

$$P_{t|t} = P_{t|t-1} - P_{t|t-1}C^T(CP_{t|t-1}C^T + R_2)^{-1}CP_{t|t-1} \tag{43}$$

State Estimation 5: Prediction Estimate

Finally, we obtain a prediction estimates of the state* by using the system in (33)

$$\hat{x}_{t+1|t} = A\hat{x}_{t|t} + Bu_t \tag{44}$$

$$P_{t+1|t} = AP_{t|t}A^T + R_1 \tag{45}$$

The errors of the different estimates are defined as follows.

Estimation error of the data update:

$$\tilde{x}_{t|t} = x_t - \hat{x}_{t|t} \in N(0, P_{t|t}) \tag{46}$$

Estimation error of the time update:

$$\tilde{x}_{t|t-1} = x_t - \hat{x}_{t|t-1} \in N(0, P_{t|t-1}) \tag{47}$$

Estimation error of the measurement prediction:

$$\epsilon_{t|t-1} = y_t - C\hat{x}_{t|t-1} \tag{48}$$

*The system matrices can even be time-variant (but deterministic).

Filter Theory - Kalman Filter

The estimator discussed so far is the Kalman filter named after Rudolph E. Kalman.

Kalman Filter: The optimal linear estimate of the conditional state $x_t|Y_s$, relying on data up to time t_s.

$$Y_s = [y_{t_0}, y_{t_1}, ..., y_{t_s}]^T \tag{49}$$

Data update (inference)

$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + \kappa_t(y_t - C\hat{x}_{t|t-1}) \tag{50}$$

$$\kappa_t = P_{t|t-1}C^T(CP_{t|t-1}C^T + R_2)^{-1} \tag{51}$$

$$P_{t|t} = P_{t|t-1} - \kappa_t C P_{t|t-1} \tag{52}$$

Time update (prediction)

$$\hat{x}_{t+1|t} = A\hat{x}_{t|t} + Bu_t, \qquad \hat{x}_{0|0} = \hat{x}_0 \tag{53}$$

$$P_{t+1|t} = AP_{t|t}A^T + R_1, \qquad P_{0|0} = P_0 \tag{54}$$

Filter Theory - Kalman Filter

Data update (inference)

$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + \kappa_t(y_t - C\hat{x}_{t|t-1}) \quad (55)$$

$$\kappa_t = P_{t|t-1}C^T(CP_{t|t-1}C^T + R_2)^{-1} \quad (56)$$

$$P_{t|t} = P_{t|t-1} - \kappa_t C P_{t|t-1} \quad (57)$$

Time update (prediction)

$$\hat{x}_{t+1|t} = A\hat{x}_{t|t} + Bu_t, \qquad \hat{x}_{0|0} = \hat{x}_0 \quad (58)$$

$$P_{t+1|t} = AP_{t|t}A^T + R_1, \qquad P_{0|0} = P_0 \quad (59)$$

How do the measurements, y_t, affect the covariances and the Kalman gain? Is it intuitive that it is that way, and can we use it to our advantage?

Think about it for yourself for one minute and then discuss with the person next to you for one minute.

Continuous-time Kalman Filter

The continuous-time Kalman filter is also known as Kalman-Bucy filter.

Consider the system:

$$\frac{d}{dt}x = Ax + Bu + v \tag{60}$$

$$y = Cx + e \tag{61}$$

For this system, the Kalman filter is

$$\frac{d}{dt}\hat{x} = A\hat{x} + Bu + K(y - C\hat{x}) \tag{62}$$

$$\frac{d}{dt}P = AP + PA^T + R_1 - KR_2K^T \tag{63}$$

$$K = PC^T R_2^{-1} \tag{64}$$

Example: Pseudocode - Kalman Filter/Simulation Implementation

Initial values: $x_{0|-1}$, $P_{0|-1}$, x_0.

---

for $t = 0, \ldots, N$

Measurement from true system:

$y_t = \texttt{Measurement}(x_t, e_t)$

Data update:

$[\hat{x}_{t|t}, P_{t|t}, \kappa_t] = \texttt{DataUpdate}(y_t, \hat{x}_{t|t-1}, P_{t|t-1}; C, R_2)$

Apply Control:

$u_t = \texttt{Actuator}(\hat{x}_{t|t})$

$x_{t+1} = \texttt{Simulator}(x_t, u_t, v_t)$

Time update:

$[\hat{x}_{t+1|t}, P_{t+1|t}] = \texttt{TimeUpdate}(\hat{x}_{t|t}, P_{t|t}, u_t; A, B, R_1)$

end

The Two Kalman Filters

When designing Kalman filters, it is important to acknowledge the type of available measurement data Y_s:

$$\hat{x}_{t|s} = E\{x_t|Y_s\} \tag{65}$$

The data determines the estimation type: Prediction or filtering.

This give us two different Kalman filters for systems on the form:

$$x_{t+1} = Ax_t + Bu_t + v_t \qquad x_{t_0} \in N(\hat{x}_0, P_0) \qquad v_t \in N(0, R_1) \tag{66}$$

$$y_t = Cx_t + e_t \qquad e_t \in N(0, R_2) \tag{67}$$

The Two Kalman Filters

The first type of Kalman filter is known as the ordinary Kalman filter:

$$\underbrace{\begin{bmatrix} \hat{x}_{t|t} \\ P_{t|t} \end{bmatrix} \rightarrow \begin{bmatrix} \hat{x}_{t+1|t} \\ P_{t+1|t} \end{bmatrix} \rightarrow \begin{bmatrix} \hat{x}_{t+1|t+1} \\ P_{t+1|t+1} \end{bmatrix}}_{\text{Ordinary Kalman Filter}} \tag{68}$$

The second type of Kalman filter is known as the predictive Kalman filter:

$$\underbrace{\begin{bmatrix} \hat{x}_{t|t-1} \\ P_{t|t-1} \end{bmatrix} \rightarrow \begin{bmatrix} \hat{x}_{t|t} \\ P_{t|t} \end{bmatrix} \rightarrow \begin{bmatrix} \hat{x}_{t+1|t} \\ P_{t+1|t} \end{bmatrix}}_{\text{Predictive Kalman Filter}} \tag{69}$$

The Ordinary Kalman Filter

The ordinary Kalman filter, provides an filter estimate $\hat{x}_{t|t}$, based on current data.

This provides a real-time estimate of the system.

Time update (prediction):

$$\hat{x}_{t+1|t} = A\hat{x}_{t|t} + Bu_t, \quad \hat{x}_{0|0} = \hat{x}_0 \tag{70}$$

$$P_{t+1|t} = AP_{t|t}A^T + R_1, \quad P_{0|0} = P_0 \tag{71}$$

Data update (inference):

$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + \kappa_t(y_t - C\hat{x}_{t|t-1}) \tag{72}$$

$$\kappa_t = P_{t|t-1}C^T(CP_{t|t-1}C^T + R_2)^{-1} \tag{73}$$

$$P_{t|t} = P_{t|t-1} - \kappa_t CP_{t|t-1} \tag{74}$$

Ordinary Kalman Filter - Closed Form

An ordinary Kalman filter can also be described in a one step notation or closed form:

$$\hat{x}_{t|t} = (I - \kappa_t C)(A\hat{x}_{t-1|t-1} + Bu_{t-1}) + \kappa_t y_t \tag{75}$$

$$P_{t|t} = AP_{t-1|t-1}A^T + R_1 - \kappa_t C(AP_{t-1|t-1}A^T + R_1) \tag{76}$$

$$\kappa_t = (AP_{t-1|t-1}A^T + R_1)C^T(C(AP_{t-1|t-1}A^T + R_1)C^T + R_2)^{-1} \tag{77}$$

In this form the time-update is substituted into the data-update.

Estimation error:

$$\tilde{x}_{t|t} = x_t - \hat{x}_{t|t} \in N(0, P_{t|t}) \tag{78}$$

The Predictive Kalman Filter

The predictive Kalman filter, provides an estimate $\hat{x}_{t|t-1}$ based on past data.

Depending on the circumstances, this can be beneficial.

Data update (inference):

$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + \kappa_t(y_t - C\hat{x}_{t|t-1}), \quad \hat{x}_{0|-1} = \hat{x}_0 \tag{79}$$

$$\kappa_t = P_{t|t-1}C^T(CP_{t|t-1}C^T + R_2)^{-1} \tag{80}$$

$$P_{t|t} = P_{t|t-1} - \kappa_t C P_{t|t-1}, \quad P_{0|-1} = P_0 \tag{81}$$

Time update (prediction):

$$\hat{x}_{t+1|t} = A\hat{x}_{t|t} + Bu_t \tag{82}$$

$$P_{t+1|t} = AP_{t|t}A^T + R_1 \tag{83}$$

Benificial circumstances includes fast systems, providing more time for computation/measurement.

Predictive Kalman Filter - Closed form

For the predictive Kalman filter the closed form is given as:

$$\hat{x}_{t+1|t} = (A - K_t C)\hat{x}_{t|t-1} + Bu_t + K_t y_t \tag{84}$$

$$P_{t+1|t} = AP_{t|t-1}A^T + R_1 - K_t C P_{t|t-1} A^T \tag{85}$$

$$K_t = A\kappa_t = AP_{t|t-1}C^T (CP_{t|t-1}C^T + R_2)^{-1} \tag{86}$$

The new gain K_t is called the predictive (Kalman) gain, as opposed to the Kalman gain κ_t

In this form, the data-update is substituted into the time-update

Estimation error:

$$\tilde{x}_{t|t-1} = x_t - \hat{x}_{t|t-1} \in N(0, P_{t|t-1}) \tag{87}$$

The ordinary case can also be given as an output equation, using the prediction as a process equation.

The Predictive Kalman Filter - Structure

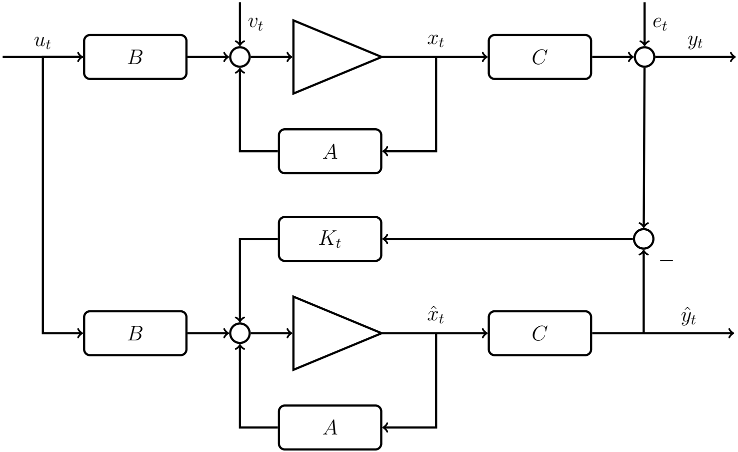

Stationary Kalman Filter

Stationarity in Kalman filters corresponds to constant Kalman gains, κ_t, K_t. A Kalman filter reaches stationarity when its variance converges to a constant P_∞. We denote by κ_∞ and K_∞ the corresponding stationary Kalman gains. $(\kappa_\infty, P^o_\infty)$ is the stationary ordinary Kalman filter, and (K_∞, P^p_∞) is the stationary predictive Kalman filter.
From the closed-form, we obtain

Predictive:

$$P^p_\infty = AP^p_\infty A^T + R_1 - AP^p_\infty C^T (CP^p_\infty C^T + R_2)^{-1} CP^p_\infty A^T \tag{88}$$

Ordinary:

$$P^o_\infty = AP^o_\infty A^T + R_1 \tag{89}$$
$$- (AP^o_\infty A^T + R_1)C^T (C(AP^o_\infty A^T + R_1)C^T + R_2)^{-1} C(AP^o_\infty A^T + R_1)$$

The stationary variances are related by

$$P^p_\infty = AP^o_\infty A^T + R_1 \tag{90}$$
$$(P^o_\infty)^{-1} = (P^p_\infty)^{-1} + C^T R_2^{-1} C \tag{91}$$

Stationary Kalman Filter

The stationary variances are related by

$$P_\infty^p = AP_\infty^o A^T + R_1 \tag{92}$$

$$(P_\infty^o)^{-1} = (P_\infty^p)^{-1} + C^T R_2^{-1} C \tag{93}$$

Assume that x and y are scalar. Which is bigger, P_∞^p or P_∞^o? How is this affected by A, C, R_1, and R_2?

Think about it for yourself for one minute and
then discuss with the person next to you for one minute.

Stationary Kalman Filter - The Riccati Equation

The form of the predictive variance is known as the discrete Riccati equation – named after Jacopo Riccati:

$$X_{t+1} = AX_tA^T + R_1 - AX_tC^T(CX_tC^T + R_2)^{-1}CX_tA^T \qquad (94)$$

Our stationary form in (88), has the form of a discrete algebraic Riccati equation (DARE):

$$X = AXA^T + R_1 - AXC^T(CXC^T + R_2)^{-1}CXA^T \qquad (95)$$

- If (A, C) is observable, a positive semi-definite solution X exists for each X_0.
- If (A, C) is observable, (A, R) is reachable $(RR^T = R_1)$, $R_1 \succeq 0$, and $R_2 \succ 0$, the solution is unique and independent of X_0 and $A - KC$ is asymptotically stable (its eigenvalues are strictly within the unit circle).

Similarly, for the continuous Kalman filter, the equation for the variance,

$$\dot{P} = AP + PA^T + R_1 - PC^TR_2^{-1}CP, \qquad (96)$$

is a continuous Riccati equation, and the stationary variant $(\dot{P} = 0)$ is a continuous algebraic Riccati equation (CARE / ARE)

Theory - Standard Kalman Assumptions

The Kalman filter is based on the assumption that a system has the form:

$$x_{t+1} = Ax_t + Bu_t + v_t \tag{97}$$

$$y_t = Cx_t + e_t \tag{98}$$

and has the following noise properties

1. $x_0 \sim N(\hat{x}_0, P_0)$
2. $v_t \sim N(0, P_v)$, white
3. $e_t \sim N(0, P_e)$, white
4. $CoV(v_t, e_t) = 0$
5. $v_t, e_t \perp x_s, \quad s \leq t$

Questions

Questions?