Automatic Control 2

Optimal control and estimation

Prof. Alberto Bemporad

University of Trento

Academic year 2010-2011

Linear quadratic regulation (LQR)

- State-feedback control via pole placement requires one to assign the closed-loop poles
- Any way to place closed-loop poles *automatically* and *optimally* ?
- The main control objectives are
  1. Make the state $x(k)$ "small" (to converge to the origin)
  2. Use "small" input signals $u(k)$ (to minimize actuators' effort)

These are conflicting goals !

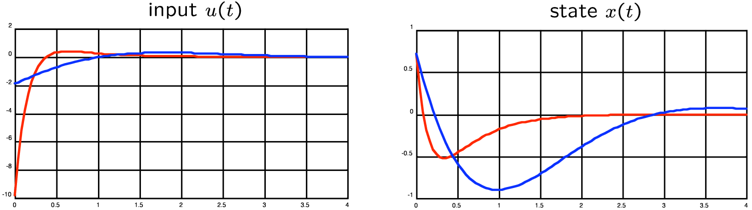

- LQR is a technique to place *automatically* and *optimally* the closed-loop poles

Finite-time optimal control

- Let the model of the open-loop process be the dynamical system

$$x(k+1) = Ax(k) + Bu(k)$$ linear system

with initial condition $x(0)$

- We look for the optimal sequence of inputs

$$U = \{u(0),\ u(1),\ \ldots,\ u(N-1)\}$$

driving the state $x(k)$ towards the origin and that minimizes the performance index

$$J(x(0), U) = x'(N)Q_N x(N) + \sum_{k=0}^{N-1} x'(k)Qx(k) + u'(k)Ru(k)$$ quadratic cost

where $Q = Q' \succeq 0$, $R = R' \succ 0$, $Q_N = Q_N' \succeq 0$[1]

[1] For a matrix $Q \in \mathbb{R}^{n\times n}$, $Q \succ 0$ means that Q is a *positive definite* matrix, i.e., $x'Qx > 0$ for all $x \neq 0$, $x \in \mathbb{R}^n$. $Q_N \succeq 0$ means *positive semidefinite*, $x'Qx \geq 0$, $\forall x \in \mathbb{R}^n$

Finite-time optimal control

- Example: Q diagonal $Q = \mathrm{Diag}(q_1, \dots, q_n)$, single input, $Q_N = 0$

$$J(x(0), U) = \sum_{k=0}^{N-1} \left(\sum_{i=1}^{n} q_i x_i^2(k) \right) + Ru^2(k)$$

- Consider again the general performance index of the posed *linear quadratic (LQ)* problem

$$J(x(0), U) = x'(N) Q_N x(N) + \sum_{k=0}^{N-1} x'(k) Q x(k) + u'(k) R u(k)$$

  - N is called the *time horizon* over which we optimize performance
  - The first term $x'Qx$ penalizes the deviation of x from the desired target $x = 0$
  - The second term $u'Ru$ penalizes actuator authority
  - The third term $x'(N)Q_N x(N)$ penalizes how much the final state $x(N)$ deviates from the target $x = 0$
- Q, R, Q_N are the tuning parameters of the optimal control design (cf. the parameters of the PID controller K_p, T_i, T_d), and are directly related to physical/economic quantities

Minimum-energy controllability

- Consider again the problem of controllability of the state to zero with minimum energy input

$$\begin{array}{ll} \min_U & \left\| \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \end{bmatrix} \right\| \\ \text{s.t.} & x(N) = 0 \end{array}$$

- The minimum-energy control problem can be seen as a particular case of the LQ optimal control problem by setting

$$R = I, \quad Q = 0, \quad Q_N = \infty \cdot I$$

Solution to LQ optimal control problem

- By substituting $x(k) = A^k x(0) + \sum_{i=0}^{k-1} A^i B u(k-1-i)$ in

$$J(x(0),U) = \sum_{k=0}^{N-1} x'(k)Qx(k) + u'(k)Ru(k) + x'(N)Q_N x(N)$$

  we obtain

$$J(x(0),U) = \frac{1}{2}U'HU + x(0)'FU + \frac{1}{2}x(0)'Yx(0)$$

  where $H = H' \succ 0$ is a positive definite matrix

- The optimizer U^* is obtained by zeroing the gradient

$$\begin{aligned} 0 &= \nabla_U J(x(0),U) = HU + F'x(0) \\ \longrightarrow \quad U^* &= \begin{bmatrix} u^*(0) \\ u^*(1) \\ \vdots \\ u^*(N-1) \end{bmatrix} = -H^{-1}F'x(0) \end{aligned}$$

[LQ problem matrix computation]

$$J(x(0),U) = x'(0)Qx(0) + \begin{bmatrix} x(1) \\ x(2) \\ \vdots \\ x(N-1) \\ x(N) \end{bmatrix}' \overbrace{\begin{bmatrix} Q & 0 & 0 & \dots & 0 \\ 0 & Q & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & 0 & Q & 0 \\ 0 & 0 & \dots & 0 & Q_N \end{bmatrix}}^{\bar{Q}} \begin{bmatrix} x(1) \\ x(2) \\ \vdots \\ x(N-1) \\ x(N) \end{bmatrix} +$$

$$\begin{bmatrix} u'(0) & u'(1) & \dots & u'(N-1) \end{bmatrix} \underbrace{\begin{bmatrix} R & 0 & \dots & 0 \\ 0 & R & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & R \end{bmatrix}}_{\bar{R}} \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \end{bmatrix}$$

$$\begin{bmatrix} x(1) \\ x(2) \\ \vdots \\ x(N) \end{bmatrix} = \overbrace{\begin{bmatrix} B & 0 & \dots & 0 \\ AB & B & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A^{N-1}B & A^{N-2}B & \dots & B \end{bmatrix}}^{\bar{S}} \begin{bmatrix} u(0) \\ u(1) \\ \dots \\ u(N-1) \end{bmatrix} + \underbrace{\begin{bmatrix} A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}}_{\bar{N}} x(0)$$

$$\begin{aligned} J(x(0),U) &= x'(0)Qx(0) + (\bar{S}U + \bar{N}x(0))'\bar{Q}(\bar{S}U + \bar{N}x(0)) + U'\bar{R}U \\ &= \frac{1}{2}U' \underbrace{2(\bar{R} + \bar{S}'\bar{Q}\bar{S})}_{H} U + x'(0) \underbrace{2\bar{N}'\bar{Q}\bar{S}}_{F} U + \frac{1}{2}x'(0) \underbrace{2(Q + \bar{N}'\bar{Q}\bar{N})}_{Y} x(0) \end{aligned}$$

Solution to LQ optimal control problem

- The solution

$$U^* = \begin{bmatrix} u^*(0) \\ u^*(1) \\ \vdots \\ u^*(N-1) \end{bmatrix} = -H^{-1}F'x(0)$$

  is an *open-loop* one: $u(k) = f_k(x(0))$, $k = 0, 1, \ldots, N-1$
- Moreover the dimensions of the H and F matrices is proportional to the time horizon N
- We use optimality principles next to find a better solution (computationally more efficient, and more elegant)

Bellman's principle of optimality

Bellman's principle

Given the optimal sequence $U^* = [u^*(0), \dots, u^*(N-1)]$ (and the corresponding optimal trajectory $x^*(k)$), the subsequence $[u^*(k_1), \dots, u^*(N-1)]$ is optimal for the problem on the horizon $[k_1, N]$, starting from the optimal state $x^*(k_1)$

Richard Bellman (1920-1984)

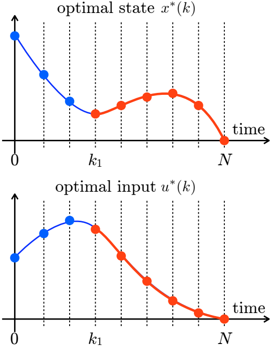

- Given the state $x^*(k_1)$, the optimal input trajectory u^* on the remaining interval $[k_1, N]$ only depends on $x^*(k_1)$
- Then each optimal move $u^*(k)$ of the optimal trajectory on $[0, N]$ only depends on $x^*(k)$
- The optimal control policy can be always expressed in state feedback form $u^*(k) = u^*(x^*(k))$!

Bellman's principle of optimality

- The principle also applies to nonlinear systems and/or non-quadratic cost functions: the optimal control law can be always written in state-feedback form

$$u^*(k) = f_k(x^*(k)), \quad \forall k = 0, \ldots, N-1$$

optimal state trajectories x^*

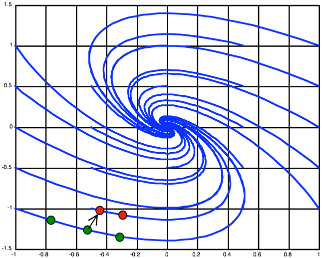

- Compared to the open-loop solution $\{u^*(0), \ldots, u^*(N-1)\} = f(x(0))$ the feedback form $u^*(k) = f_k(x^*(k))$ has the big advantage of being more *robust* with respect to perturbations: at each time k we apply the best move on the remaining period $[k, N]$

Dynamic programming

- At a generic instant k_1 and state $x(k_1) = z$ consider the optimal *cost-to-go*

$$V_{k_1}(z) = \min_{u(k_1),\dots,u(N-1)} \left\{ \sum_{k=k_1}^{N-1} x'(k)Qx(k) + u'(k)Ru(k) + x'(N)Q_N x(N) \right\}$$

- $V_0(z)$ is the optimal cost in the remaining interval $[k_1, N]$ starting at $x(k_1) = z$

Principle of dynamic programming

$$\begin{aligned} V_0(z) &= \min_{U \triangleq \{u(0),\dots,u(N-1)\}} J(z,U) \\ &= \min_{u(0),\dots,u(k_1-1)} \left\{ \sum_{k=0}^{k_1-1} x'(k)Qx(k) + u'(k)Ru(k) + V_{k_1}(x(k_1)) \right\} \end{aligned}$$

- Starting at $x(0)$, the minimum cost over $[0, N]$ equals the minimum cost spent until step k_1 plus the optimal cost-to-go from k_1 to N starting at $x(k_1)$

Riccati iterations

By applying the dynamic programming principle, we can compute the optimal inputs $u^*(k)$ recursively as a function of $x^*(k)$ (*Riccati iterations*):

1. Initialization: $P(N) = Q_N$
2. For $k = N, \ldots, 1$, compute recursively the following matrix

$$P(k-1) = Q - A'P(k)B(R + B'P(k)B)^{-1}B'P(k)A + A'P(k)A$$

3. Define

$$K(k) = -(R + B'P(k+1)B)^{-1}B'P(k+1)A$$

The optimal input is

$$u^*(k) = K(k)x^*(k)$$

Jacopo Francesco Riccati (1676-1754)

The optimal input policy $u^*(k)$ is a (linear time-varying) state feedback !

Linear quadratic regulation

- Dynamical processes operate on a very long time horizon (in principle, for ever). Let's send the optimal control horizon $N \to \infty$

$$V^{\infty}(x(0)) = \min_{u(0),u(1),\ldots} \sum_{k=0}^{\infty} x'(k)Qx(k) + u'(k)Ru(k)$$

Result

Let (A,B) be a stabilizable pair, $R \succ 0$, $Q \succeq 0$. There exists a unique solution P_∞ of the *algebraic Riccati equation (ARE)*

$$P_\infty = A'P_\infty A + Q - A'P_\infty B(B'P_\infty B + R)^{-1}B'P_\infty A$$

such that the optimal cost is $V^{\infty}(x(0)) = x'(0)P_\infty x(0)$ and the optimal control law is the constant linear state feedback $u(k) = K_{\rm LQR}x(k)$ with

$$K_{\rm LQR} = -(R + B'P_\infty B)^{-1}B'P_\infty A$$

MATLAB

P_∞ = dare(A,B,Q,R)

MATLAB

[$-K_\infty$,P_∞] = dlqr(A,B,Q,R)

Linear quadratic regulation

- Go back to Riccati iterations: starting from $P(\infty) = P_\infty$ and going backwards we get $P(j) = P_\infty$, $\forall j \geq 0$
- Accordingly, we get

$$K(j) = -(R + B'P_\infty B)^{-1}B'P_\infty A \triangleq K_{\text{LQR}}, \quad \forall j = 0, 1, \ldots$$

- The LQR control law is linear and time-invariant

MATLAB

```
» [-K∞,P∞,E] = lqr(sysd,Q,R)
```

E= closed-loop poles
= eigenvalues of $(A + BK_{\text{LQR}})$

- Closed-loop stability is ensured if (A, B) is stabilizable, $R \succ 0$, $Q \succeq 0$, and $(A, Q^{\frac{1}{2}})$ is detectable, where $Q^{\frac{1}{2}}$ is the *Cholesky factor*[2] of Q
- LQR is an automatic and optimal way of placing poles !
- A similar result holds for continuous-time linear systems (**MATLAB :** `lqr`)

[2] Given a matrix $Q = Q' \succeq 0$, its Cholesky factor is an upper-triangular matrix C such that $C'C = Q$ (**MATLAB :** `chol`)

LQR with output weighting

- We often want to regulate only $y(k) = Cx(k)$ to zero, so define

$$V^\infty(x(0)) = \min_{u(0),u(1),\ldots} \sum_{k=0}^{\infty} y'(k)Q_y y(k) + u'(k)Ru(k)$$

- The problem is again an LQR problem with equivalent state weight

$Q = C'Q_yC$

MATLAB

```
» [−K∞,P∞,E] = dlqry(sysd,Qy,R)
```

Corollary

Let (A,B) stabilizable, (A,C) detectable, $R > 0$, $Q_y > 0$. Under the LQR control law $u(k) = K_{\mathrm{LQR}}x(k)$ the closed-loop system is asymptotically stable

$$\lim_{t\to\infty} x(t) = 0, \quad \lim_{t\to\infty} u(t) = 0$$

- Intuitively: the minimum cost $x'(0)P_\infty x(0)$ is finite $\Rightarrow y(k) \to 0$ and $u(k) \to 0$. $y(k) \to 0$ implies that the observable part of the state $\to 0$. As $u(k) \to 0$, the unobservable states remain undriven and goes to zero spontaneously (=detectability condition)

LQR example

- Two-dimensional single input single output (SISO) dynamical system (double integrator)

$$\begin{aligned} x(k+1) &= \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x(k) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(k) \\ y(k) &= \begin{bmatrix} 1 & 0 \end{bmatrix} x(k) \end{aligned}$$

- LQR (infinite horizon) controller defined on the performance index

$$V^{\infty}(x(0)) = \min_{u(0),u(1),\ldots} \sum_{k=0}^{\infty} \frac{1}{\rho} y^2(k) + u^2(k), \quad \rho > 0$$

- Weights: $Q_y = \frac{1}{\rho}$ (or $Q = \left[\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right] \cdot \frac{1}{\rho} \cdot \left[\begin{smallmatrix} 1 & 0 \end{smallmatrix}\right] = \left[\begin{smallmatrix} \frac{1}{\rho} & 0 \\ 0 & 0 \end{smallmatrix}\right]$), $R = 1$
- Note that only the ratio $Q_y/R = \frac{1}{\rho}$ matters, as scaling the cost function does not change the optimal control law

LQR Example

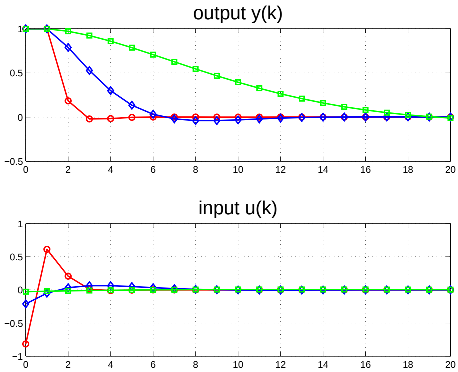

$\rho = 0.1$ (red line)

$K = [-0.8166 \quad -1.7499]$

$\rho = 10$ (blue line)

$K = [-0.2114 \quad -0.7645]$

$\rho = 1000$ (green line)

$K = [-0.0279 \quad -0.2505]$

Initial state: $x(0) = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$

$$V^\infty(x(0)) = \min_{u(0),u(1),\ldots} \sum_{k=0}^{\infty} \frac{1}{\rho} y^2(k) + u^2(k)$$

Kalman filtering – Introduction

- **Problem:** assign observer poles in an optimal way, that is to minimize the state estimation error $\tilde{x} = x - \hat{x}$
- Information comes in two ways: from sensors measurements (*a posteriori*) and from the model of the system (*a priori*)
- We need to mix the two information sources optimally, given a probabilistic description of their reliability (sensor precision, model accuracy)

Rudolf E. Kalman*

(born 1930)

The *Kalman filter* solves this problem, and is now the most used state observer in most engineering fields (and beyond)

*R.E. Kalman receiving the Medal of Science from the President of the USA on October 7, 2009

Modeling assumptions

- The process is modeled as the *stochastic linear system*

$$\begin{aligned} x(k+1) &= Ax(k)+Bu(k)+\xi(k) \\ y(k) &= Cx(k)+\zeta(k) \\ x(0) &= x_0 \end{aligned}$$

- $\xi(k) \in \mathbb{R}^n =$ *process noise*. We assume $E[\xi(k)] = 0$ (zero mean), $E[\xi(k)\xi'(j)] = 0\ \forall k \neq j$ (white noise), and $E[\xi(k)\xi'(k)] = Q \succeq 0$ (covariance matrix)
- $\zeta(k) \in \mathbb{R}^p =$ *measurement noise*, $E[\zeta(k)] = 0$, $E[\zeta(k)\zeta'(j)] = 0\ \forall k \neq j$, $E[\zeta(k)\zeta'(k)] = R \succ 0$
- $x_0 \in \mathbb{R}^n$ is a random vector, $E[x_0] = \bar{x}_0$, $E[(x_0-\bar{x}_0)(x_0-\bar{x}_0)'] = Var[x_0] = P_0$, $P_0 \succeq 0$
- Vectors $\xi(k)$, $\zeta(k)$, x_0 are uncorrelated: $E[\xi(k)\zeta'(j)] = 0$, $E[\xi(k)x_0'] = 0$, $E[\zeta(k)x_0'] = 0$, $\forall k,j \in \mathbb{Z}$
- Probability distributions: we often assume *normal* (=*Gaussian*) distributions $\xi(k) \sim \mathcal{N}(0,Q)$, $\zeta(k) \sim \mathcal{N}(0,R)$, $x_0 \sim \mathcal{N}(\bar{x}_0,P_0)$

Kalman predictor

- Introduce some quantities:

$$\hat{x}(k+1|k) \rightarrow \textit{state prediction}$$
$$\tilde{x}(k+1|k) = x(k+1) - \hat{x}(k+1|k) \rightarrow \textit{prediction error}$$
$$P(k+1|k) = E\left[\tilde{x}(k+1|k)\tilde{x}(k+1|k)'\right] \rightarrow \textit{prediction error covariance}$$

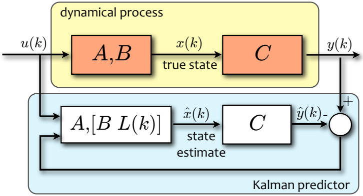

- The observer dynamics is

$$\hat{x}(k+1|k) = A\hat{x}(k|k-1) + Bu(k) + L(k)(y(k) - C\hat{x}(k|k-1))$$

  starting from an initial estimate $\hat{x}(0|-1) = \hat{x}_0$.
- The Kalman predictor minimizes the covariance matrix $P(k+1|k)$ of the prediction error $\tilde{x}(k+1|k)$

Kalman predictor

Theorem

The observer gain $L(k)$ minimizing the trace of $P(k|k-1)$ is

$$L(k) = AP(k|k-1)C' \left[CP(k|k-1)C' + R\right]^{-1}$$

where $P(k|k-1)$ solves the iterative equations

$$P(k+1|k) = AP(k|k-1)A' + Q - AP(k|k-1)C' \left[CP(k|k-1)C' + R\right]^{-1} CP(k|k-1)A'$$

with initial condition $P(0|-1) = P_0$

- The Kalman predictor is a *time-varying* observer, as the gain $L(k)$ depends on the time index k
- In many cases we are interested in a *time-invariant* observer gain L, to avoid computing $P(k+1|k)$ on line
- Note the similarities with Riccati iterations in LQR

Stationary Kalman predictor

Theorem

Let (A, C) observable, and (A, B_q) stabilizable, where $Q = B_q B_q'$ (B_q=Cholesky factor of Q). Then

- the stationary optimal predictor is

$$\begin{aligned}\hat{x}(k+1|k) &= A\hat{x}(k|k-1) + Bu(k) + L(y(k) - C\hat{x}(k|k-1)) \\ &= (A - LC)\hat{x}(k|k-1) + Bu(k) + Ly(k)\end{aligned}$$

  with

$$L = AP_\infty C' \left[CP_\infty C' + R\right]^{-1}$$

  where P_∞ is the only positive-definite solution of the algebraic Riccati equation

$$P_\infty = AP_\infty A' + Q - AP_\infty C' \left[CP_\infty C' + R\right]^{-1} CP_\infty A'$$

- the observer is asymptotically stable, i.e. all the eigenvalues of $(A - LC)$ are inside the unit circle

Example of Kalman predictor: time-varying gain

- Noisy measurements $y(k)$ of an exponentially decaying signal $x(k)$

$$\begin{array}{rclll} x(k+1) & = & ax(k)+\xi(k) & E[\xi(k)]=0 & E[\xi^2(k)]=Q \\ y(k) & = & x(k)+\zeta(k) & E[\zeta(k)]=0 & E[\zeta^2(k)]=R \\ & & |a|<1 & E[x(0)]=0 & E[x^2(0)]=P_0 \end{array}$$

- The equations of the Kalman predictor are

$$\hat{x}(k+1|k) = a\hat{x}(k|k-1) + L(k)\left(y(k) - \hat{x}(k|k-1)\right), \quad \hat{x}(0|-1) = 0$$

$$L(k) = \frac{aP(k|k-1)}{P(k|k-1)+R}$$

$$P(k+1|k) = a^2P(k|k-1) + Q - \frac{a^2P^2(k|k-1)}{P(k|k-1)+R}, \quad P(0|-1) = P_0$$

- Let $a = 0.98$, $Q = 10^{-5}$, $R = 5\cdot 10^{-3}$, $P_0 = 1$, $\hat{x}(0) = 0.4$, $x(0) = 1$

Example of Kalman predictor: time-varying gain

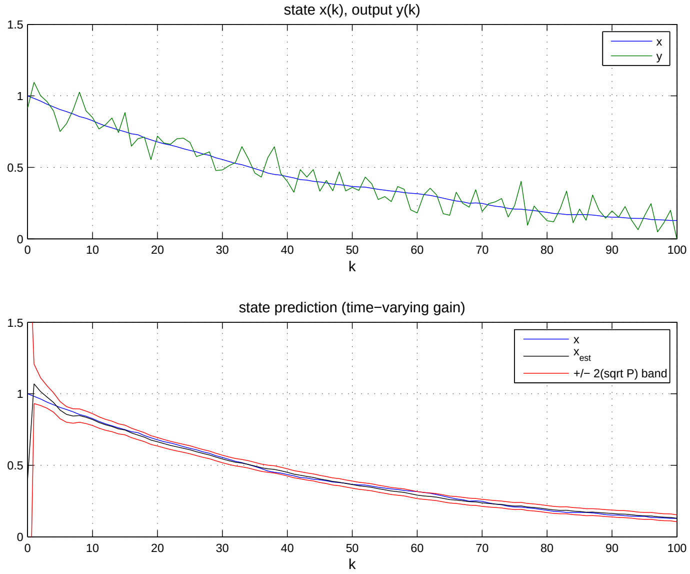

Example of Kalman predictor: stationary gain

- Let's compute the stationary Kalman predictor (the pair $(a, 1)$ is observable, the pair (a, B_q), with $B_q = \sqrt{10^{-5}}$ is stabilizable)
- Solve the algebraic Riccati equation for $a = 0.98$, $Q = 10^{-5}$, $R = 5 \cdot 10^{-3}$, $P_0 = 1$. We get two solutions:

$$P_\infty = -3.366 \cdot 10^{-4}, \quad P_\infty = 1.486 \cdot 10^{-4}$$

  the positive solution is the correct one
- Compute $L = \dfrac{aP_\infty}{P_\infty + R} = 2.83 \cdot 10^{-2}$
- The stationary Kalman predictor is

$$\hat{x}(k+1|k) = 0.98\,\hat{x}(k) + 2.83 \cdot 10^{-2}(y(k) - C\hat{x}(k|k-1))$$

- Check asymptotic stability:

$$\tilde{x}(k+1|k) = (a-L)\tilde{x}(k|k-1) = 0.952\,\tilde{x}(k|k-1) \Rightarrow \lim_{k\to +\infty} \tilde{x}(k|k-1) = 0$$

Kalman predictor: example

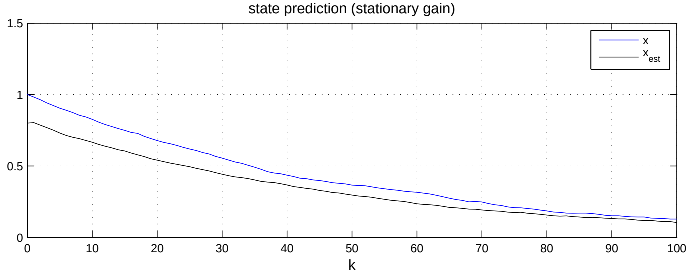

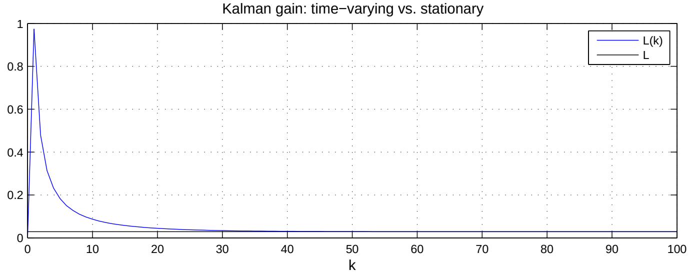

Kalman filter

- The Kalman filter provides an optimal estimate of $x(k)$ given the measurements *up to time* k
- The Kalman filter proceeds in two steps:
  1. *measurement update* based on the most recent $y(k)$

$$
\begin{array}{rcll}
M(k) & = & P(k|k-1)C'[CP(k|k-1)C'+R]^{-1} & P(0|-1)=P_0 \\
\hat{x}(k|k) & = & \hat{x}(k|k-1)+M(k)\left(y(k)-C\hat{x}(k|k-1)\right) & \hat{x}(0|-1)=\hat{x}_0 \\
P(k|k) & = & (I-M(k)C)P(k|k-1) &
\end{array}
$$

  2. *time update* based on the model of the system

$$
\begin{aligned}
\hat{x}(k+1|k) &= A\hat{x}(k|k)+Bu(k) \\
P(k+1|k) &= AP(k|k)A'+Q
\end{aligned}
$$

- Same as Kalman predictor, as $L(k)=AM(k)$
- Stationary version: $M=P_\infty C'(CP_\infty C'+R)^{-1}$

MATLAB

» [KEST,L,P_∞,M,Z]=kalman(sys,Q,R)

$$Z=E[(x(k)-x(k|k))(x(k)-x(k|k))']$$

Tuning Kalman filters

- It is usually hard to quantify exactly the correct values of Q and R for a given process
- The diagonal terms of R are related to how noisy are output sensors
- Q is harder to relate to physical noise, it mainly relates to how rough is the (A, B) model
- After all, Q and R are the tuning knobs of the observer (similar to LQR)
- The "larger" is R with respect to Q the "slower" is the observer to converge (L, M will be small)
- On the contrary, the "smaller" is R than Q, the more precise are considered the measurments, and the "faster" observer will be to converge

LQG control

- *Linear Quadratic Gaussian (LQG)* control combines an LQR control law and a stationary Kalman predictor/filter
- Consider the stochastic dynamical system

$$\begin{aligned} x(k+1) &= Ax(k) + Bu(k) + \xi(k),\ w \sim \mathcal{N}(0, Q_{KF}) \\ y(k) &= Cx(k) + \zeta(k),\ v \sim \mathcal{N}(0, R_{KF}) \end{aligned}$$

with initial condition $x(0) = x_0$, $x_0 \sim \mathcal{N}(\bar{x}_0, P_0)$, $P, Q_{KF} \succeq 0$, $R_{KF} \succ 0$, and ζ and ξ are independent and white noise terms.

- The objective is to minimize the cost function

$$J(x(0), U) = \lim_{T \to \infty} \frac{1}{T} E\left[\sum_{k=0}^{T} x'(k) Q_{LQ} x(k) + u'(k) R_{LQ} u(k)\right]$$

when the state x is not measurable

LQG control

If we assume that all the assumptions for LQR control and Kalman predictor/filter hold, i.e.

- the pair (A, B) is reachable and the pair (A, C_q) with C_q such that $Q_{LQ} = C_q C_q'$ is observable (here Q is the weight matrix of the LQ controller)
- the pair (A, B_q), with B_q s.t. $Q_{KF} = B_q B_q'$, is stabilizable, and the pair (A, C) is observable (here Q is the covariance matrix of the Kalman predictor/filter)

Then, apply the following procedure:

1. Determine the optimal stationary Kalman predictor/filter, neglecting the fact that the control variable u is generated through a closed-loop control scheme, and find the optimal gain L_{KF}
2. Determine the optimal LQR strategy assuming the state accessible, and find the optimal gain K_{LQR}

LQG control

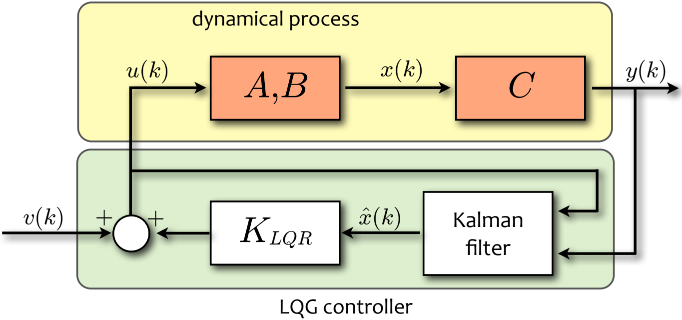

Analogously to the case of output feedback control using a Luenberger observer, it is possible to show that the extended state $[x' \ \tilde{x}']'$ has eigenvalues equal to the eigenvalues of $(A + BK_{LQR})$ plus those of $(A - L_{KF}C)$ ($2n$ in total)

English-Italian Vocabulary

	
positive (semi)definite matrix	*matrice (semi)definita positiva*
dynamic programming	*programmazione dinamica*

Translation is obvious otherwise.