12 LQG systems and the Kalman filter

We introduce the LQG model and show how to reduce it to a stochastic controllable dynamical system using the Kalman filter. The LQG system is the system of equations

$$X_{n+1} = AX_n + BU_n + \varepsilon_{n+1},$$
$$Y_{n+1} = CX_n + \eta_{n+1}, \qquad n \geqslant 0.$$

Here A, B and C are given matrices and the random variables $X_0, \binom{\varepsilon_1}{\eta_1}, \binom{\varepsilon_2}{\eta_2}, \dots$ are independent Gaussians, X_0 having mean x and variance Σ_0 and, for $n \geqslant 1$, ε_n and η_n having mean 0 and

$$\mathrm{var}(\varepsilon_n) = N, \quad \mathrm{cov}(\varepsilon_n, \eta_n) = L, \quad \mathrm{var}(\eta_n) = M.$$

The *state* X_n takes values in $\mathbb{R}^d$, the *observation* Y_n takes values in $\mathbb{R}^p$ and the control values U_n are in $\mathbb{R}^m$. We complete the model by specifying a *control*, which is a function $u : (\mathbb{R}^p)^* \to \mathbb{R}^m$, and setting $U_n = u_n(Y_1, \dots, Y_n)$. We emphasise that what is different now is that we no longer observe the state, but have to estimate the state value on the basis of the observations. Set

$$V_n^u(x, \Sigma_0) = \mathbb{E}^u_{(x,\Sigma_0)}\left(\sum_{k=0}^{n-1} c(X_k, U_k) + c(X_n)\right), \quad V_n(x, \Sigma_0) = \inf_u V_n^u(x, \Sigma_0).$$

Lemma 12.1. *Let X and Y be jointly Gaussian, with mean 0 and with*

$$\mathrm{var}(X) = U, \quad \mathrm{cov}(X, Y) = W, \quad \mathrm{var}(Y) = V,$$

with V invertible. Set $\hat{X} = WV^{-1}Y$ and $Z = X - \hat{X}$. Then Z is independent of Y with

$$\mathrm{var}(Z) = U - WV^{-1}W^T.$$

Proof. Note that Y and Z are jointly Gaussian, so zero covariance will imply independence. We compute

$$\mathrm{cov}(Z, Y) = \mathrm{cov}(X, Y) - WV^{-1}\,\mathrm{var}(Y) = W - WV^{-1}V = 0$$

and

$$\mathrm{var}(Z) = \mathrm{cov}(Z, X) = \mathrm{var}(X) - WV^{-1}\,\mathrm{cov}(Y, X) = U - WV^{-1}W^T.$$

□

We now obtain a recursive scheme, called the *Kalman filter*, which determines for $n \geqslant 1$ the mean and variance of the conditional distribution of X_n, given the observations $Y_1, \dots, Y_n$. Suppose inductively that *we can write $X_n = \hat{X}_n + \Delta_n$, where $\hat{X}_n$ is a function of $Y_1, \dots, Y_n$, and where Δ_n is independent of $Y_1, \dots, Y_n$, with distribution $N(0, \Sigma_n)$*. This is true for $n = 0$, with $\hat{X}_0 = x$. We have

$$X_{n+1} = A\hat{X}_n + BU_n + \xi_{n+1}, \qquad \xi_{n+1} = \varepsilon_{n+1} + A\Delta_n,$$
$$Y_{n+1} = C\hat{X}_n + \zeta_{n+1}, \qquad \zeta_{n+1} = \eta_{n+1} + C\Delta_n.$$

Note that the *innovations* ξ_{n+1} and ζ_{n+1} are zero-mean Gaussians and are independent of $Y_1, \ldots, Y_n$, with

$$\operatorname{var}(\xi_{n+1}) = \tilde{N} = N + A\Sigma_n A^T, \quad \operatorname{var}(\zeta_{n+1}) = \tilde{M} = M + C\Sigma_n C^T,$$

$$\operatorname{cov}(\xi_{n+1}, \zeta_{n+1}) = \tilde{L} = L + A\Sigma_n C^T.$$

Set

$$H_{n+1} = H(\Sigma_n) = \tilde{L}\tilde{M}^{-1}, \quad \Sigma_{n+1} = \sigma(\Sigma_n) = \tilde{N} - \tilde{L}\tilde{M}^{-1}\tilde{L}^T.$$

By the lemma, $\xi_{n+1} = \hat{\varepsilon}_{n+1} + \Delta_{n+1}$, where

$$\hat{\varepsilon}_{n+1} = H_{n+1}\zeta_{n+1} = H_{n+1}(Y_{n+1} - C\hat{X}_n)$$

and where Δ_{n+1} is independent of ζ_{n+1}, and hence of $Y_1, \ldots, Y_{n+1}$, with distribution $N(0, \Sigma_{n+1})$. Note that

$$\operatorname{var}(\hat{\varepsilon}_{n+1}) = H_{n+1} \operatorname{var}(\zeta_{n+1}) H_{n+1}^T = \tilde{L}\tilde{M}^{-1}\tilde{L}^T = \tilde{N} - \Sigma_{n+1} = N + A\Sigma_n A^T - \Sigma_{n+1}.$$

Now $X_{n+1} = \hat{X}_{n+1} + \Delta_{n+1}$, where

$$\hat{X}_{n+1} = A\hat{X}_n + BU_n + \hat{\varepsilon}_{n+1},$$

which is a function of $Y_1, \ldots, Y_{n+1}$, as required. This establishes the induction.

Note that

$$\begin{aligned}\mathbb{E}(c(X_k, U_k)) &= \mathbb{E}(c(\hat{X}_k + \Delta_k, U_k)) \\ &= \mathbb{E}(\Delta_k^T R \Delta_k) + \mathbb{E}(c(\hat{X}_k, U_k)) = \operatorname{trace}(R\Sigma_k) + \mathbb{E}(c(\hat{X}_k, U_k))\end{aligned}$$

and, similarly,

$$\mathbb{E}(c(X_n)) = \operatorname{trace}(\Pi_0 \Sigma_n) + \mathbb{E}(c(\hat{X}_n)).$$

Hence

$$V_n(x, \Sigma_0) = \hat{V}_n(x, \Sigma_0) + \sum_{k=0}^{n-1} \operatorname{trace}(R\Sigma_k) + \operatorname{trace}(\Pi_0 \Sigma_n),$$

where $\hat{V}_n$ is the infimal cost function of the stochastic controllable dynamical system

$$\hat{X}_{n+1} = A\hat{X}_n + BU_n + \hat{\varepsilon}_{n+1}, \quad \Sigma_{n+1} = \sigma(\Sigma_n),$$

where $(\hat{\varepsilon}_n)_{n \geqslant 1}$ are independent, and $\hat{\varepsilon}_{n+1}$ has distribution $N(0, \hat{N}(\Sigma_n))$, with

$$\hat{N}(\Sigma) = N + A\Sigma A^T - \sigma(\Sigma).$$

This system can be treated by a small variation of the method in the preceding section. In particular, certainty-equivalence holds: the optimal control for the n-horizon problem is $U_k = K(\Pi_{n-1-k})\hat{X}_k$. The product form of this control is remarkable as, given A, B and n, $K(\Pi_{n-1-k})$ depends only on the cost functions, whilst the controllable dynamical system for $\hat{X}_k$ is independent of the costs. This is called the *separation principle.*

Example. We investigate from first principles one of the simplest control problems with noisy observation. We shall follow the same lines as in the general theory and use similar notation. The system has scalar state and observations and is given by

$$X_{n+1} = X_n + U_n, \quad Y_{n+1} = X_{n+1} + \eta_{n+1}, \quad n \geqslant 0,$$

where the random variable X_0 and η_n, $n \geqslant 1$ are independent, with $X_0 \sim N(x, v)$ and $\eta_n \sim N(0, 1)$, for all n, and where $U_n = u_n(Y_1, \dots, Y_n)$. We fix a time-horizon n and aim to choose u to minimize

$$V_n^u(x, v) = \mathbb{E}^u_{(x,v)}\left(\sum_{k=0}^{n-1} U_k^2 + DX_n^2\right).$$

Consider first the control problem for $x_k = \mathbb{E}(X_k)$: we seek to minimize $\sum_{k=0}^{n-1} u_k^2 + Dx_n^2$ subject to $x_{k+1} = x_k + u_k$ and $x_0 = x$. The minimum is $Dx^2/(1 + Dn)$, achieved when $u_k = -Dx_k/(1 + D(n-k))$.

Next, we calculate the Kalman filter. We determine recursively for $n \geqslant 0$ a function $\hat{X}_n$ of $Y_1, \dots, Y_n$ such that $X_n = \hat{X}_n + \Delta_n$, with Δ_n independent of $Y_1, \dots, Y_n$. Write $v_n = \mathrm{var}(\Delta_n)$. For $n = 0$ we can take $\hat{X}_0 = x$ and $v_0 = v$. At the nth step, we write

$$\begin{aligned} X_{n+1} &= \hat{X}_n + U_n + \xi_{n+1}, & \xi_{n+1} &= \Delta_n, \\ Y_{n+1} &= \hat{X}_n + U_n + \zeta_{n+1}, & \zeta_{n+1} &= \Delta_n + \eta_{n+1}, \end{aligned}$$

where the innovations ξ_{n+1} and ζ_{n+1} are independent of $Y_1, \dots, Y_n$. We aim to split

$$\xi_{n+1} = H_{n+1}\zeta_{n+1} + \Delta_{n+1},$$

where Δ_{n+1} is independent of ζ_{n+1} and hence of $Y_1, \dots, Y_{n+1}$. On taking variances in this equation, we obtain

$$v_n = H_{n+1}^2(v_n + 1) + v_{n+1}.$$

On the other hand, by taking the covariance with ζ_{n+1}, we have

$$v_n = H_{n+1}(v_n + 1).$$

These equations imply that $H_{n+1} = v_{n+1}$ and determine a recursion $v_{n+1}^{-1} = 1 + v_n^{-1}$, so $v_n^{-1} = n + v_0^{-1}$, and so $v_n = v/(1 + vn)$.

Now $\hat{X}_{n+1} = \hat{X}_n + U_n + \hat{\varepsilon}_{n+1}$, where $\hat{\varepsilon}_{n+1} = H_{n+1}\zeta_{n+1}$, so

$$\mathrm{var}(\hat{\varepsilon}_{n+1}) = s_{n+1} = H_{n+1}^2(1 + v_n) = \left(\frac{v}{1 + (n+1)v}\right)^2\left(1 + \frac{v}{1 + nv}\right).$$

By certainty-equivalence, the optimal control for the n-horizon problem is given by $U_k = -D\hat{X}_k/(1 + D(n-k))$, so

$$\hat{X}_{k+1} = \frac{1 + D(n-k-1)}{1 + D(n-k)}\hat{X}_k + \hat{\varepsilon}_{n+1}.$$

On taking variances, we obtain the recursion

$$\mathrm{var}(\hat{X}_{k+1}) = \left(\frac{1 + D(n-k-1)}{1 + D(n-k)}\right)^2 \mathrm{var}(\hat{X}_k) + s_{n+1}.$$

Finally, the minimal expected cost is

$$\mathbb{E}^u_{(x,v)}\left(\sum_{k=0}^{n-1} U_k^2 + DX_n^2\right) = \sum_{k=0}^{n-1} \frac{D^2}{(1 + D(n-k))^2} \mathrm{var}(\hat{X}_k) + D\,\mathrm{var}(\hat{X}_n) + D\,\mathrm{var}(\Delta_n).$$