5.2 Linear Quadratic Problems

In this section we study a particular important case where the difficulties mentioned at the end of the previous section do not apply, namely where the problem of designing an optimal controller for the partial information problem can be broken into two successive parts. First, design an optimal estimator for the state of the system, and then use this estimator directly in a controller that is optimal for the system under perfect information. This phenomenon is called the *separation principle* (or the principle of separation of estimation and control) and significantly simplifies the controller design, since the controller has no influence on the quality of the estimator. The problem considered here is a generalization of the LQR problem studied in chapter 4, where we now assume noisy measurements that are linear in the state of the system. Our discussion follows [Ber07, section 5.2], which shows the nice result that the separation principle follows from using a linear model and a quadratic cost function, *without making any Gaussian assumption on the noise models.* I mention this because very often, the discussion of linear quadratic methods under imperfect information, and even sometimes of LQR, make this Gaussian assumption from the start, but this has the drawback of blurring our understanding of when the separation principle holds (this phenomenon applies more generally than to the linear quadratic case). Hence we introduce the Gaussian assumption on the disturbance and observation noises only at the end to give an example where the estimation problem is tractable, using the Kalman filter.

The dynamics of the system are as in chapter 4

$$x_{k+1} = A_k x_k + B_k u_k + w_k,\ k = 0, 1, \ldots, N-1,$$

assuming that the initial state x_0 is random, with known finite mean $\bar{x}_0$ and covariance matrix Σ_0^-. In addition, the measurements are linear in the state, of the form

$$y_k = C_k x_k + v_k,\ k = 0, \ldots, N-1.$$

Here $x_k \in \mathbb{R}^n, u_k \in \mathbb{R}^m, y^k \in \mathbb{R}^p$ and the matrices A_k, B_k, C_k are of appropriate dimensions. As in chapter 4, the variables w_k are assumed to be independent, and independent of x_0, and to have zero mean and finite covariance matrices W_k. Now we make the same assumptions for the observation noise variables v_k, denote the covariance matrices V_k, and in addition we assume that these variables are independent of the process disturbances w_k and of x_0. The cost is still the quadratic cost

$$E\left\{\sum_{k=0}^{N-1}(x_k^T Q_k x_k + u_k^T R_k u_k) + x_N^T Q_N x_N\right\}, \tag{5.11}$$

with $Q_k \succeq 0, R_k \succ 0$.

The first step of the DP algorithm, which we write here using the full information vector, is:

$$J_N^*(\mathcal{I}_N) = E(x_N^T Q_N x_N | \mathcal{I}_N).$$

Faced with this problem for the first time, we would actually have to compute a few more steps of the DP algorithm to see what's going on and make our induction hypothesis, and after some tedious work, we would settle on trying to prove that for all k, we have

$$J_k^*(\mathcal{I}_k) = E(x_k^T P_k x_k|\mathcal{I}_k) + \sum_{j=k}^{N-1} E[e_j^T \tilde{P}_j e_j|\mathcal{I}_k] + \sum_{j=k}^{N-1} \mathrm{Tr}(P_{j+1}W_j), \tag{5.12}$$

for some matrices $P_j \succeq 0, \tilde{P}_j \succeq 0$, and with the definition

$$e_k := x_k - E[x_k|\mathcal{I}_k].$$

It is good at this point to look again at the formula for the cost in the perfect information case, say (4.4), to see that the only differences here are the second term, which captures some additional cost due to estimation errors, as well as the presence of the conditional expectation in the first term. We will show that indeed the matrices P_k follow the same recursion (Riccati difference equation) as in the perfect information case! In any case, (5.12) clearly holds for $k = N$ with $P_N = Q_N$.

Let us denote the constant term $q_k := \sum_{j=k}^{N-1} \mathrm{Tr}(P_{j+1}W_j)$, and assume that the induction hypothesis is true for $k+1$. We want to show that it is true for k. Then, using the induction hypothesis in the DP equation, we have

$$\begin{aligned} J_k^*(\mathcal{I}_k) = \min_{u_k \in \mathsf{U}_k} E\Big[& x_k^T Q_k x_k + u_k^T R_k u_k + E(x_{k+1}^T P_{k+1} x_{k+1}|\mathcal{I}_{k+1}) \\ & + \sum_{j=k+1}^{N-1} E[e_j^T \tilde{P}_j e_j|\mathcal{I}_{k+1}] + q_{k+1} \Big| \mathcal{I}_k \Big]. \end{aligned}$$

First we isolate the terms that clearly are not affected by the minimization over u_k, and also use the tower property of conditional expectations, to get

$$\begin{aligned} J_k^*(\mathcal{I}_k) =& E[x_k^T Q_k x_k|\mathcal{I}_k] + q_{k+1} \\ & + \min_{u_k \in \mathsf{U}_k} \Big\{ u_k^T R_k u_k + E[x_{k+1}^T P_{k+1} x_{k+1}|\mathcal{I}_k] + \sum_{j=k+1}^{N-1} E[e_j^T \tilde{P}_j e_j|\mathcal{I}_k] \Big\}. \end{aligned} \tag{5.13}$$

Now it will be easy to take care of the second term in the minimization term by using the dynamics equation as for LQR, but the last term looks problematic. It is not clear how the choice of u_k influences the future error terms $e_{k+1}, \ldots e_{N-1}$. Fortunately, it turns out that u_k doesn't influence these terms at all! This is due to the linearity of the system and observation equations (and of the fact that the noises are independent of the control inputs), and hence a very particular situation. In contrast to our comments at the end of section 5.1, this is a case where the control inputs cannot influence the quality of the future state estimate errors as expressed by e_k.

Lemma 5.2.1. *For all k, there is a function M_k such that*

$$e_k = M_k(x_0, w_{0:k-1}, v_{0:k}),$$

independently of the policy being used.

Proof. (from [Ber07, Vol. I, p.231]) We consider two systems, one driven by some control inputs, the other with zero control inputs. The two systems are driven by the same noise realizations $\{w_k\}$ and $\{v_k\}$, which is valid under our assumptions that the noise does not depend on the control input values. Then we show that for these two systems the terms e_k are the same. Hence let the two systems be

$$x_{k+1} = A_k x_k + B_k u_k + w_k, \quad y_k = C_k x_k + v_k,$$

and

$$\tilde{x}_{k+1} = A_k \tilde{x}_k + w_k, \quad \tilde{y}_k = C_k \tilde{x}_k + v_k.$$

Assume $x_0 = \tilde{x}_0$, and denote the information vectors for the two systems as

$$\mathcal{I}_k = \{y_{0:k}, u_{0:k-1}\}, \quad \tilde{\mathcal{I}}_k = \{\tilde{y}_{0:k}\}.$$

Now we refer to the classical computations of the response of a discrete-time linear dynamical system. Define the state transition matrix

$$\Phi(k,l) = \begin{cases} A_{k-1}A_{k-2}\dots A_l, & k > l \geq 0 \\ I, & k = l \end{cases}.$$

Then we have

$$x_k = \Phi(k,0)x_0 + \sum_{l=0}^{k-1} \Phi(k,l+1)B_l u_l + \sum_{l=0}^{k-1} \Phi(k,l+1)w_l$$
$$\tilde{x}_k = \Phi(k,0)x_0 + \sum_{l=0}^{k-1} \Phi(k,l+1)w_l$$

and so we immediately get

$$\begin{aligned} x_k - E[x_k|\mathcal{I}_k] &= \Phi(k,0)(x_0 - E[x_0|\mathcal{I}_k]) + \sum_{l=0}^{k-1} \Phi(k,l+1)(w_l - E[w_l|\mathcal{I}_k]) \\ &= \tilde{x}_k - E[\tilde{x}_k|\mathcal{I}_k]. \end{aligned}$$

Now it is sufficient to prove $E[\tilde{x}_k|\mathcal{I}_k] = E[\tilde{x}_k|\tilde{\mathcal{I}}_k]$. Note however that

$$\tilde{y}_k = y_k - \sum_{l=0}^{k-1} \Phi(k,l+1)CB_l u_l,$$

so that the information about $\tilde{x}_k$ contained in $\mathcal{I}_k$ is summarized in $\tilde{\mathcal{I}}_k$. In conclusion

$$e_k = \bar{x}_k - E[\bar{x}_k|\bar{\mathcal{I}}_k] =: M_k(x_0, v_{0:k}, w_{0:k-1}),$$

which is independent of the control policy used. □

Proof. (alternative proof, using the innovation sequence). Define the *innovation sequence*

$$\tilde{y}_0 = y_0 - E[y_0],\ \tilde{y}_k = y_k - E[y_k|y_{0:k-1}],\ k = 1, \dots, N-1.$$

Thus $\tilde{y}_k$ records the deviation in observation k with respect to the value which could have been estimated from the past measurements. Since $y_0 = \tilde{y}_0 + E[y_0] = \tilde{y}_0 + CE[x_0]$ with $E[x_0]$ a known constant, we have $\sigma(y_0) = \sigma(\tilde{y}_0)$, where $\sigma\{X_1, \dots, X_n\}$ denotes the σ-algebra generated by the variables $\{X_1, \dots, X_n\}$, or in other non-mathematical terms, the "information" contained in these variables. Assume by induction that $\sigma(y_{0:k}) = \sigma(\tilde{y}_{0:k})$, which is true for $k = 0$. Then

$$\begin{aligned} y_{k+1} &= \tilde{y}_{k+1} + E[y_{k+1}|y_{0:k}] \\ &= \tilde{y}_{k+1} + E[y_{k+1}|\tilde{y}_{0:k}], \quad \text{(using the induction hypothesis)} \end{aligned}$$

so $y_{k+1} \in \sigma(\tilde{y}_{0:k+1})$, and clearly $\tilde{y}_{k+1} \in \sigma(y_{0:k+1})$ so $\sigma(y_{0:k+1}) = \sigma(\tilde{y}_{0:k+1})$ and the induction step is complete.

Now note that $\tilde{y}_0$ and $e_0 = x_0 - E[x_0|y_0]$ do not depend on the control inputs. Again by induction, assume that $\tilde{y}_{0:k}$ and e_k do not depend on the control policy. We have, noting that u_k must be a function of $y_{0:k}$ to be admissible, hence $E[CBu_k|y_{0:k}] = CBu_k$,

$$\begin{aligned} \tilde{y}_{k+1} &= CAx_k + CBu_k + Cw_k - E[CAx_k + CBu_k + Cw_k|y_{0:k}] \\ &= CAe_k + Cw_k - CE[w_k|y_{0:k}] \\ &= CAe_k + Cw_k - CE[w_k|\tilde{y}_{0:k}], \end{aligned}$$

so $\tilde{y}_{k+1}$ is independent of the control policy. Moreover

$$\begin{aligned} e_{k+1} &= x_{k+1} - E[x_{k+1}|y_{0:k+1}] \\ &= x_{k+1} - E[x_{k+1}|\tilde{y}_{0:k+1}] \\ &= \Phi(k,0)(x_0 - E[x_0|\tilde{y}_{0:k+1}]) + \sum_{l=0}^{k-1} \Phi(k, l+1)(w_l - E[w_l|\tilde{y}_{0:k+1}]), \end{aligned}$$

which is independent of the control input, so we are done. □

Let us go back to the computation (5.13). From lemma 5.2.1, we now have

$$\begin{aligned} J_k^*(\mathcal{I}_k) =& E[x_k^T Q_k x_k|\mathcal{I}_k] + q_{k+1} + \sum_{j=k+1}^{N-1} E[e_j^T \tilde{P}_j e_j|\mathcal{I}_k] \\ &+ \min_{u_k \in \mathsf{U}_k} \Big\{ u_k^T R_k u_k + E[x_{k+1}^T P_{k+1} x_{k+1}|\mathcal{I}_k] \Big\}. \end{aligned}$$

We can now consider the minimization problem, which is very similar to the perfect information case. The term to minimize over u_k is

$$\begin{aligned} & u_k^T R_k u_k + E[(A_k x_k + B_k u_k + w_k)^T P_{k+1} (A_k x_k + B_k u_k + w_k)|\mathcal{I}_k] \\ =& u_k^T (R_k + B_k^T P_{k+1} B_k) u_k + 2 u_k^T B_k^T P_{k+1} A_k E[x_k|\mathcal{I}_k] \\ &+ E[x_k^T A_k^T P_{k+1} A_k x_k|\mathcal{I}_k] + \mathrm{Tr}(P_{k+1} W_k). \end{aligned}$$

The last term of this equation correspond to the term quadratic in w_k, and uses the fact that w_k is independent of $\mathcal{I}_k$. Also because $E[w_k|\mathcal{I}_k] = 0$, the terms linear in w_k vanish. Now we have a quadratic minimization problem, which we can solve by setting the derivative of u_k to zero, or by using directly the Schur complement result from section 4.2. We get

$$\mu_k^*(\mathcal{I}_k) = K_k E[x_k|\mathcal{I}_k], \text{ with } K_k = -(R_k + B_k^T P_{k+1} B_k)^{-1} B_k^T P_{k+1} A_k,$$

exactly like in the perfect information case (provided we show that the matrices P_k are also the same). The optimum value is

$$\begin{aligned} J_k^*(\mathcal{I}_k) =& E[x_k^T(Q_k + A_k^T P_{k+1} A_k) x_k|\mathcal{I}_k] + q_{k+1} + \text{Tr}(P_{k+1} W_k) \\ & - E[x_k|\mathcal{I}_k]^T \tilde{P}_k E[x_k|\mathcal{I}_k] + \sum_{j=k+1}^{N-1} E[e_j^T \tilde{P}_j e_j|\mathcal{I}_k], \end{aligned}$$

where

$$\tilde{P}_k = A_k^T P_{k+1} B_k (R_k + B_k^T P_{k+1} B_k)^{-1} B_k^T P_{k+1} A_k.$$

Finally, note the identity

$$E[x_k^T M x_k|\mathcal{I}_k] = E[x_k|\mathcal{I}_k]^T M E[x_k|\mathcal{I}_k] + E[e_k^T M e_k|\mathcal{I}_k], \tag{5.14}$$

which holds for any $M \succeq 0$. Using it with $M = \tilde{P}_k$ in the fourth term of the cost value above, and replacing the value of q_{k+1} we get exactly (5.12), with

$$P_k = A_k^T P_{k+1} A_k + Q_k - \tilde{P}_k,$$

which as we announced before is the same Riccati equation (4.3) as in the perfect information case (hence, incidentally, positive semidefiniteness of P_k follows directly from the proof in that).

Exercise 12. Prove equation (5.14).

So finally we have

$$J_0^*(y_0) = E(x_0^T P_0 x_0|y_0) + \sum_{j=0}^{N-1} E[e_j^T \tilde{P}_j e_j|y_0] + \sum_{j=0}^{N-1} \text{Tr}(P_{k+1} W_k), \tag{5.15}$$

and the optimal cost $J^* = E[J^*(y_0)]$:

$$J^* = \text{Tr}(P_0 X_0) + \sum_{j=0}^{N-1} \text{Tr}(\tilde{P}_j E[\Sigma_j]) + \sum_{j=0}^{N-1} \text{Tr}(P_{k+1} W_k), \tag{5.16}$$

with $X_0 = E[x_0 x_0^T]$, and $\Sigma_j = E[e_j e_j^T|\mathcal{I}_j]$ the conditional *error covariance matrix* (note that $E[e_j^T \tilde{P}_j e_j|y_0] = \text{Tr}(\tilde{P}_j E[\Sigma_j|y_0])$ by the tower property of conditional expectations).

Remark. Here we see that we have a further sufficient statistic for the linear quadratic control problem, in terms of $\hat{x}_k = E[x_k|\mathcal{I}_k]$ and $\Sigma_k = E[e_k e_k^T|\mathcal{I}_k]$.

Remark. The derivation above is often found with the assumption from the beginning that the noises are Gaussian, in which case it turns out that the conditional error covariance matrices $\Sigma_k = E[e_k e_k^T|\mathcal{I}_k]$ are in fact constant (i.e., independent of the observations $\mathcal{I}_k$) and the derivation becomes somewhat simpler. We haven't made that assumption yet, but we will when we introduce the Kalman filter. The point is to show that pretty much everything works in the proof without the Gaussian assumption, using only linearity of the dynamics and the quadratic cost. If you can design for your particular noise model a nice recursive filter to obtain $E[x_k|\mathcal{I}_k]$, you can just use it directly in the optimal control law.

The Separation Principle

As announced in the introduction, the separation principle holds for the linear quadratic problem with imperfect information. The optimal control problem indeed decomposes into two successive parts

1. An estimator, the conditional mean $E[x_k|\mathcal{I}_k]$. Note that this estimator is the Minimum Mean Square Estimator (MMSE), i.e., the estimate $\hat{x}_k$ of x_k given $\mathcal{I}_k$ which minimizes $E[\|x_k - \hat{x}_k\|^2|\mathcal{I}_k]$ is precisely $\hat{x}_k = E[x_k|\mathcal{I}_k]$ (it's good to check this fact if you have forgotten).

2. A control law which is the optimal policy for the full information case with x_k simply replaced by its conditional mean estimate $E[x_k|\mathcal{I}_k]$.

Recursive Estimation Using the Kalman Filter

The optimal controller requires the computation of the same gain matrices K_k as in the perfect information case, which in turn requires the backward computation of the P_k matrices using the Riccati difference equation. However, in the partial information case, it also requires the computation of the MMSE:

$$\hat{x}_k = E[x_k|\mathcal{I}_k],$$

which is not easy in general. We discuss this estimation problem in this section, following the earlier discussion of the Bayes filter. Let us repeat for convenience the hypothesis on our system:

$$x_{k+1} = A_k x_k + B_k u_k + w_k, \tag{5.17}$$

$$y_k = C_k x_k + v_k, \tag{5.18}$$

and the noise processes $\{w_k\}$ and $\{v_k\}$ are white, zero-mean, uncorrelated, with known covariance matrices W_k and V_k:

$$E[w_k] = E[v_k] = 0,$$
$$E[w_k w_l^T] = W_k \delta_{j-l},\ E[v_k v_l^T] = V_k \delta_{j-l},\ E[w_k v_l^T] = 0,$$

where δ_j is the Kronecker delta ($\delta_j = 1$ if $j = 0$, $\delta_j = 0$ otherwise). Recall that we denoted the error term $e_k = x_k - \hat{x}_k$. Since we will follow the Bayes filter and its propagation and measurement update steps, we also need the following quantities

$$\begin{aligned}
&\mathcal{I}_k^- = \{\mathcal{I}_{k-1}, u_{k-1}\} = \{y_{0:k-1}, u_{0:k-1}\},\ \mathcal{I}_0^- = \emptyset \\
&\hat{x}_k^- = E[x_k|\mathcal{I}_k^-],\ \hat{x}_0^- = \bar{x}_0, \\
&e_k^- = x_k - \hat{x}_k^-,\ \Sigma_k^- = E[e_k^-(e_k^-)^T|\mathcal{I}_k^-], \\
&\Sigma_0^- = E[(x_0 - \bar{x}_0)(x_0 - \bar{x}_0)^T] = X_0 - \bar{x}_0\bar{x}_0^T, \\
&e_k = x_k - \hat{x}_k,\ \Sigma_k = E[e_k e_k^T|\mathcal{I}_k].
\end{aligned}$$

Here $\hat{x}_k^-$ is the conditional expectation of x_k before incorporating the last measurement y_k. We let $\Sigma_0^- = E[(x_0 - \bar{x}_0)(x_0 - \bar{x}_0)^T]$ since it is the error covariance before incorporating the first measurement y_0. Note that the notation Σ_k for the conditional error covariance was already introduced earlier to express the optimal cost function.

Now we want to compute $\hat{x}_k$ for all k, and in addition let us consider the computation of the error covariance matrices Σ_k, which we would need anyway if we wanted to evaluate the costs (5.15) or (5.16). Suppose that we knew $\hat{x}_k$ and Σ_k. Following the Bayes filter, the next step would be to compute $\hat{x}_{k+1}^-$ and Σ_{k+1}^- (propagation or time-update step) resulting from the input u_k and disturbance w_k. At this point, we do not yet take into account the measurement y_{k+1}. Computing $\hat{x}_{k+1}^-$ is simply a matter of taking expectations in the system equation (5.17) and using linearity

$$\hat{x}_{k+1}^- = A_k \hat{x}_k + B_k u_k.$$

For the covariance update, it is not too hard to see that we have

$$\Sigma_{k+1}^- = A_k \Sigma_k A_k^T + W_k.$$

The next step, corresponding to the measurement update step in the Bayes filter, asks us to compute $\hat{x}_{k+1}$ and Σ_{k+1}, based on the knowledge of $\hat{x}_{k+1}^-, \Sigma_{k+1}^-$ and the new measurement y_{k+1}. This step is more complicated. Consider first the following situation. We have a pair of vectors $X \in \mathbb{R}^n$, $Y \in \mathbb{R}^p$ which are *jointly Gaussian*, with mean and covariance

$$E \begin{bmatrix} X \\ Y \end{bmatrix} = \begin{bmatrix} \bar{x} \\ \bar{y} \end{bmatrix},\ \mathrm{Cov}\left(\begin{bmatrix} X \\ Y \end{bmatrix}\right) = \begin{bmatrix} C_{xx} & C_{xy} \\ C_{xy}^T & C_{yy} \end{bmatrix}.$$

Then the distribution of X given Y is again Gaussian, and we have

$$E[X|Y] = \bar{x} + C_{xy} C_{yy}^{-1}(Y - \bar{y}) \tag{5.19}$$

$$\begin{aligned}
\mathrm{Cov}(X|Y) &:= E[(X - E(X|Y))(X - E(X|Y))^T|Y] \\
&= C_{xx} - C_{xy} C_{yy}^{-1} C_{xy}^T.
\end{aligned} \tag{5.20}$$

Remark. Note the following facts:

- The conditional mean (5.19) is an affine function of Y.
- The formula for the conditional covariance (5.20) does not actually depend on Y! This is a very special property of the Gaussian distribution.
- Note the Schur complement in (5.20). A pseudo-inverse can replace C_{yy}^{-1} if this matrix is not invertible.

Now let us assume that the disturbances $\{w_k\}, \{v_k\}$ and the initial condition x_0 have a Gaussian distribution, in addition to the previous assumptions on their mean and covariance. Then it turns out that the random vectors X_{k+1}, Y_{k+1} are jointly Gaussian. Their mean and covariance, conditioned on $\mathcal{I}_{k+1}^-$, are

$$\begin{bmatrix} \hat{x}_{k+1}^- \\ C_{k+1}\hat{x}_{k+1}^- \end{bmatrix}, \ \begin{bmatrix} \Sigma_{k+1}^- & \Sigma_{k+1}^- C_{k+1}^T \\ C_{k+1}\Sigma_{k+1}^- & C_{k+1}\Sigma_{k+1}^- C_{k+1}^T + V_{k+1} \end{bmatrix}.$$

Now we can use the formulas (5.19), (5.20) to see that X_{k+1} conditioned on $\mathcal{I}_{k+1} = \{\mathcal{I}_{k+1}^-, Y_{k+1}\}$ is again a Gaussian random variable with mean and covariance

$$\begin{aligned} \hat{x}_{k+1} &= E[X_{k+1}|\mathcal{I}_{k+1}] \\ &= \hat{x}_{k+1}^- + \Sigma_{k+1}^- C_{k+1}^T (C_{k+1}\Sigma_{k+1}^- C_{k+1}^T + V_{k+1})^{-1}(Y_{k+1} - C_{k+1}\hat{x}_{k+1}^-) \\ \Sigma_{k+1} &= \Sigma_{k+1}^- - \Sigma_{k+1}^- C_{k+1}^T (C_{k+1}\Sigma_{k+1}^- C_{k+1}^T + V_{k+1})^{-1} C_{k+1}\Sigma_{k+1}^-. \end{aligned}$$

Here we assume for simplicity that the covariance matrices $\{V_k\}$ are positive definite, in order to guarantee the existence of the inverses. In summary, we have obtained a recursive algorithm to compute $\hat{x}_k$ as new measurements arrive, under the additional assumption that the disturbances have a Gaussian distribution. The algorithm can be summarized as follows:

$$\begin{aligned} &\hat{x}_0^- = \bar{x}_0, \ \Sigma_0^- = E[(X_0 - \bar{x}_0)(X_0 - \bar{x}_0)^T] \\ &\hat{x}_k = \hat{x}_k^- + \Sigma_k^- C_k^T (C_k \Sigma_k^- C_k^T + V_k)^{-1}(y_k - C_k \hat{x}_k^-), \ k = 0, \dots, N \\ &\hat{x}_{k+1}^- = A_k \hat{x}_k + B_k u_k, \ k = 0, \dots, N-1 \\ &\Sigma_k = \Sigma_k^- - \Sigma_k^- C_k^T (C_k \Sigma_k^- C_k^T + V_k)^{-1} C_k \Sigma_k^-, \ k = 0, \dots, N \\ &\Sigma_{k+1}^- = A_k \Sigma_k A_k^T + W_k, \ k = 0, \dots, N-1. \end{aligned} \tag{5.21}$$

So starting with $\bar{x}_0, \Sigma_0^-$, we can perform the time-update (or propagation) step to get $\hat{x}_0, \Sigma_0$, then the measurement-update step to obtain $\hat{x}_1^-, \Sigma_1^-$, and so on. This recursive algorithm to compute $\hat{x}_k, \Sigma_k$ is called the *Kalman filter*. Under the assumptions of Gaussian disturbances, the distribution of X_k given $\mathcal{I}_k$ is also Gaussian for all k and so this distribution is completely determined by its mean $\hat{x}_k$ and variance Σ_k which are computed by the Kalman filter. The Kalman filter algorithm produces $\hat{x}_{k+1}$ at time $k+1$ based on $\hat{x}_k$, using only the most recent observation y_{k+1} and control u_k. Here are a few additional facts that you can note at this point regarding the Kalman filter.

1. The Kalman filter update for $\hat{x}_k$ is linear in the new observation y_k.

2. The conditional error covariance matrices $\Sigma_k = E[e_k e_k^T | \mathcal{I}_k]$ turn out not to depend on the actual values of the observations $y_{0:k}$ (and controls $u_{0:k}$, we saw that in a lemma earlier). Only the conditional mean estimates $\hat{x}_k$ do. This is a special feature of the Kalman filter, which depends on the linearity and Gaussian assumption, and does not happen in general if the model is nonlinear. Hence we have also $\Sigma_k = E[e_k e_k^T]$, a constant independent of $\mathcal{I}_k$. No one set of measurements or controls helps any more than any other to eliminate some uncertainty about x_k. Note also that
$$\hat{x}_{k+1} = A_k \hat{x}_k + B_k u_k + L_k(y_{k+1} - C_k(A_k \hat{x}_k + B_k u_k))$$
where $L_k = \Sigma_k^- C_k^T (C_k \Sigma_k^- C_k^T + V_k)^{-1}$ are the *gain matrices* of the filter. These gains are also independent of the values of the measurements. An immediate consequence is that they *can be computed offline before the filter is actually run.* This is a critical property if the filter is used in real-time applications.

3. Even though the error covariance and gain matrices can be precomputed, we still need to store them in memory, which could be problematic for problems with a long horizon. Just as for the control gain matrices K_k however, we have a *steady state* version of the Kalman filter, which is obtained by letting $k \to +\infty$ in the equations (5.21) (whereas for the controller, the recurrence was going backwards and we let $k \to -\infty$ in section 4.3). Note that the matrices Σ_k^- also satisfy a Riccati difference equation
$$\Sigma_{k+1}^- = A_k \Sigma_k^- A_k^T + W_k - A_k \Sigma_k^- C_k^T (C_k \Sigma_k^- C_k^T + V_k)^{-1} C_k \Sigma_k^- A_k^T,$$
compare with (4.3). For a time-homogeneous problem, as $k \to \infty$ and under appropriate assumptions (e.g. (A, C) observable and $(A, W_k^{1/2})$ controllable), we have that Σ_k^- converges to a constant matrix Σ solution of the *algebraic Riccati equation* (ARE)
$$\bar{\Sigma} = A\bar{\Sigma}A^T + W - A\bar{\Sigma}C^T(C\bar{\Sigma}C^T + V)^{-1}C\bar{\Sigma}A^T,$$
and the gain matrices L_k of the filter also converge to the constant
$$L = \bar{\Sigma}C^T(C\bar{\Sigma}C^T + V)^{-1}.$$
The error covariance matrices converge to the steady state matrix Σ with
$$\Sigma = \bar{\Sigma} - \bar{\Sigma}C^T(C\bar{\Sigma}C^T + V)^{-1}C\bar{\Sigma}.$$
The steady-state filter is often used, even for finite (but sufficiently long) horizon problems, because it has obviously much lower memory requirements than the time-varying optimal filter.

With the Gaussian assumption on the disturbances, the problem considered in this section is often called the Linear-Quadratic-Gaussian problem (LQG). We have seen that the matrices Σ_k are then constants independent of $\mathcal{I}_k$. The optimal cost (5.16) can then be written

$$\begin{aligned} J^* &= \mathrm{Tr}(P_0 X_0) + \sum_{j=0}^{N-1} \mathrm{Tr}(\tilde{P}_j \Sigma_j) + \sum_{j=0}^{N-1} \mathrm{Tr}(P_{k+1} W_k) \\ &= J_{lqr} + J_{est} \end{aligned}$$

where J_{lqr} is the term corresponding to the cost (4.4) of the LQR problem, and J_{est} is an additional cost for having to estimate the state

$$J_{est} = \sum_{j=0}^{N-1} \mathrm{Tr}(\tilde{P}_j \Sigma_j).$$

The corresponding infinite horizon average cost is

$$J^*_{avge} = \lim_{N \to \infty} E \left\{ \sum_{k=0}^{N-1} x_k^T Q_k x_k + u_k^T R_k u_k \right\} = \mathrm{Tr}(\tilde{P}\Sigma) + \mathrm{Tr}(PW).$$

Here $P, \tilde{P}$ and Σ are the steady state versions of the matrices P_k, $\tilde{P}_k$, and Σ_k. The term $\mathrm{Tr}(PW)$ is what we would get in the LQR problem (perfect information). The additional term $\mathrm{Tr}(\tilde{P}\Sigma)$ is due to the average estimation error. This average cost does not depend on X_0 or of any transient modification of the optimal control law. An optimal policy for this infinite horizon problem consists in using the steady state Kalman filter and using the corresponding estimate in the steady-state LQR controller.