

Lecture 9

Linear-Quadratic Control

In this lecture we derive an optimal state-feedback controller known as the linear-quadratic regulator.

9.1 An optimization framework

In the remainder of the course we will view the controller design problem as an optimization problem. A general controller optimization framework is illustrated in Figure 9.1 (see also Figure 1.5 in Lecture 1). The “Plant” block is a generalized process description that includes all dynamics and input–output relations that are relevant for the design problem at hand. The vector w contains all exogenous signals, including disturbances and reference signals. The vector z contains all performance outputs (error signals, control signals, etc.) and should in general be minimized. The “Controller” block takes a vector y as input that contains any signals that the controller could use, including measurements and reference signals. Finally the vector u contains the control signals as well as signals that should be included in the performance vector z .




Figure 9.1 A general controller optimization framework.

Assuming a linear MIMO plant, the general objective is to find a linear controller that optimizes the closed-loop transfer matrix $G_{zw}(s)$, subject to various constraints. The most general form of the optimization problem will be studied in Lecture 13. In this and the following two lectures we will focus on linear plants and quadratic performance criteria, which yields so-called linear-quadratic (LQ) controllers. In the current lecture we will assume that the full state vector x is available for feedback. In the next lecture we will design optimal state estimators (Kalman filters), and in the following lecture we combine the optimal state feedback and the Kalman filter into an linear-quadratic Gaussian (LQG) controller.

9.2 The linear-quadratic control problem

The linear-quadratic (LQ) optimal control problem is formulated as

$$\text{minimize} \quad J = \int_0^\infty \begin{pmatrix} x(t) \\ u(t) \end{pmatrix}^T \begin{pmatrix} Q_1 & Q_{12} \\ Q_{12}^T & Q_2 \end{pmatrix} \begin{pmatrix} x(t) \\ u(t) \end{pmatrix} dt, \quad (9.1)$$

$$\text{subject to} \quad \dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0. \quad (9.2)$$

Here, the matrix $Q = \begin{pmatrix} Q_1 & Q_{12} \\ Q_{12}^T & Q_2 \end{pmatrix}$ should be positive definite (“ $Q > 0$ ”) and describes how states and control signals different from zero are penalized. The process (A, B) must be *stabilizable*, meaning that any unstable modes of A must be controllable. The objective is to control the process from the initial state x_0 to the origin (which is the assumed setpoint) while incurring a minimal cost J .

We shall see that this problem formulation yields a static, linear, stabilizing control law that is independent of the initial state. The problem has an analytic solution that can be calculated by hand for small systems. Additional attractive properties are that the method is directly applicable to MIMO systems and that there are guaranteed stability margins. The LQ problem is also the foundation of more advanced, non-linear schemes such as model-predictive control (MPC).

Dynamic programming

The LQ controller can be derived in many different ways. Here we will give a solution based on the technique of *dynamic programming*. Dynamic programming (introduced by Richard Bellman) means optimal sequential decision making, and is most easily explained using an example.

EXAMPLE 1. Determine u_0 and u_1 if the objective is to minimize

$$x_1^2 + x_2^2 + u_0^2 + u_1^2,$$

when

$$\begin{aligned} x_1 &= x_0 + u_0, \\ x_2 &= x_1 + u_1. \end{aligned}$$

This can be viewed as a discrete-time optimal control problem in two time steps. The trick in dynamic programming is to break the problem into smaller parts that can be solved sequentially. In this case we can rewrite the problem as

$$\min_{u_0, u_1} \{x_1^2 + x_2^2 + u_0^2 + u_1^2\} = \min_{u_0} \left\{ x_1^2 + u_0^2 + \underbrace{\min_{u_1} \{x_2^2 + u_1^2\}}_{J_1(x_1)}(x_1) \right\}$$

Here we see that the cost in the last time step can be optimized first, using only knowledge of x_1 . The minimization is performed using completion of the square as

$$J_1(x_1) = \min_{u_1} \{(x_1 + u_1)^2 + u_1^2\} = \min_{u_1} \left\{ 2\left(u_1 + \frac{1}{2}x_1\right)^2 + \frac{1}{2}x_1^2 \right\} = \frac{1}{2}x_1^2$$

where the optimum is attained for $u_1 = -\frac{1}{2}x_1$. Proceeding one step backwards, we can now minimize the total cost using

$$J_0(x_0) = \min_{u_0} \{(x_0 + u_0)^2 + u_0^2 + J_1(x_0 + u_0)\} = \min_{u_0} \left\{ \frac{5}{2}\left(u_0 + \frac{3}{5}x_0\right)^2 + \frac{3}{5}x_0^2 \right\} = \frac{3}{5}x_0^2$$

where the optimum is attained for $u_0 = -\frac{3}{5}x_0$. It is interesting to note that the optimal decision in each step is a linear feedback from the current state. \square

Dynamic programming uses Bellman’s *principle of optimality*, which states that, whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision. Applying this to the LQ control problem, we assume that the first decision is taken at time t in the state $x(t)$. The resulting state is $x(t + \epsilon)$ a short time interval ϵ later, when a new optimal decision should be taken. See Figure 9.2.



Figure 9.2 The principle of optimality states that an optimal trajectory on the time interval $[t, T]$ must be optimal also on each of the subintervals $[t, t + \epsilon]$ and $[t + \epsilon, T]$.

For a linear-quadratic problem, it can be shown that the optimal cost on the time interval $[t, \infty)$ is quadratic in the initial state $x(t)$:

$$\min_{u[t, \infty)} \int_t^\infty \begin{pmatrix} x(\tau) \\ u(\tau) \end{pmatrix}^T Q \begin{pmatrix} x(\tau) \\ u(\tau) \end{pmatrix} d\tau = x^T(t) S x(t). \quad (9.3)$$

Here, S is a symmetric, positive definitive matrix.

We now apply dynamic programming to the LQ problem starting in the state $x(t)$. For an infinitesimal time step of length ϵ , the process dynamics (9.2) can be written as

$$x(t + \epsilon) = x(t) + (Ax(t) + Bu(t))\epsilon \quad \text{as } \epsilon \rightarrow 0 \quad (9.4)$$

where $x(t)$ and $u(t)$ are assumed constant over the short interval ϵ . Invoking the principle of optimality, a new optimal decision should be taken at time $t + \epsilon$. Combining (9.3) and (9.4) we can write the optimal cost as

$$\begin{aligned} x^T(t) S x(t) &= \min_{u[t, \infty)} \int_t^\infty \begin{pmatrix} x(\tau) \\ u(\tau) \end{pmatrix}^T Q \begin{pmatrix} x(\tau) \\ u(\tau) \end{pmatrix} d\tau \\ &= \min_{u[t, \infty)} \left\{ \begin{pmatrix} x(t) \\ u(t) \end{pmatrix}^T Q \begin{pmatrix} x(t) \\ u(t) \end{pmatrix} \epsilon + \int_{t+\epsilon}^\infty \begin{pmatrix} x(\tau) \\ u(\tau) \end{pmatrix}^T Q \begin{pmatrix} x(\tau) \\ u(\tau) \end{pmatrix} d\tau \right\} \\ &= \min_{u(t)} \left\{ x^T(t) Q_1 x(t) + 2x^T(t) Q_{12} u(t) + u^T(t) Q_2 u(t) \right. \\ &\quad \left. + [x(t) + (Ax(t) + Bu(t))\epsilon]^T S [x(t) + (Ax(t) + Bu(t))\epsilon] \right\} \end{aligned}$$

where the last equality follows from the definition of S . Neglecting the ϵ^2 terms gives *Bellman's equation*:

$$0 = \min_{u(t)} \left\{ \left(x^T(t) Q_1 x(t) + 2x^T(t) Q_{12} u(t) + u^T(t) Q_2 u(t) \right) + 2x^T(t) S (Ax(t) + Bu(t)) \right\}$$

As a final step, we find the value of $u(t)$ that minimizes the right-hand side of Bellman's equation by completion of squares:

$$\begin{aligned} 0 &= \min_{u(t)} \left\{ \left(x^T(t) Q_1 x(t) + 2x^T(t) Q_{12} u(t) + u^T(t) Q_2 u(t) \right) + 2x^T(t) S (Ax(t) + Bu(t)) \right\} \\ &= \min_{u(t)} \left\{ x^T(t) [Q_1 + A^T S + SA] x(t) + 2x^T(t) [Q_{12} + SB] u(t) + u^T(t) Q_2 u(t) \right\} \\ &= x^T(t) \left(Q_1 + A^T S + SA - (SB + Q_{12}) Q_2^{-1} (SB + Q_{12})^T \right) x(t) \end{aligned}$$

The optimum is attained for

$$u(t) = -Q_2^{-1} (SB + Q_{12})^T x(t).$$

The quadratic matrix equation

$$0 = Q_1 + A^T S + SA - (SB + Q_{12}) Q_2^{-1} (SB + Q_{12})^T \quad (9.5)$$

is called the (continuous-time) *algebraic Riccati equation* and can be solved by hand for simple problems or by computer software (care in Matlab).

Linear-quadratic optimal state feedback

Summarizing the result of the previous subsection, we find that the optimal control problem

$$\begin{aligned} &\text{minimize} && \int_0^\infty \left(x^T(t)Q_1x(t) + 2x^T(t)Q_{12}u(t) + u^T(t)Q_2u(t) \right) dt \\ &\text{subject to} && \dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0 \end{aligned}$$

is solved by the unique matrix $S = S^T > 0$ that satisfies the algebraic Riccati equation

$$0 = Q_1 + A^T S + SA - (SB + Q_{12})Q_2^{-1}(SB + Q_{12})^T.$$

The optimal control law is $u = -Lx$ with

$$L = Q_2^{-1}(SB + Q_{12})^T,$$

and the optimal cost is given by

$$J^* = x_0^T S x_0.$$

Note that, just like in the example above, the optimal control law does not depend on the initial state x_0 . Moreover, the feedback gain L is static because we are solving an infinite-horizon problem. (A similar calculation as above can be performed also for finite-horizon problems with time-varying system and cost matrices. We then obtain a Riccati differential equation for $S(t)$ and a time-varying optimal state feedback law, $u(t) = -L(t)x(t)$.)

EXAMPLE 2. Consider optimal control of an integrator process,

$$\dot{x}(t) = u(t), \quad x(0) = x_0.$$

The objective is to minimize the quadratic cost function

$$J = \int_0^\infty \{x(t)^2 + \rho u(t)^2\} dt, \quad \rho > 0$$

With $A = 0$, $B = 1$, $Q_1 = 1$, $Q_2 = \rho$, and $Q_{12} = 0$, the algebraic Riccati equation becomes

$$0 = 1 - S^2/\rho \quad \Rightarrow \quad S = \sqrt{\rho}$$

The optimal feedback gain is

$$L = S/\rho = 1/\sqrt{\rho}$$

yielding the closed-loop system

$$\dot{x} = -x/\sqrt{\rho}.$$

The corresponding optimal cost is

$$J^* = x_0^T S x_0 = x_0^2 \sqrt{\rho}.$$

From the solution we see that we can *tune* the LQ controller by selecting the control signal penalty ρ . A small value of ρ gives a high gain, a fast closed loop, and a large total cost. The opposite holds for large values of ρ . \square

EXAMPLE 3. Consider control of a double integrator process,

$$\begin{aligned} \dot{x} &= \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 1 \end{pmatrix} u, & x(0) &= \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ y &= \begin{pmatrix} 1 & 0 \end{pmatrix} x, \end{aligned}$$


and the cost function

$$J = \int_0^{\infty} (y^2(t) + \rho u^2(t)) dt.$$

The relevant parameters for this LQ problem are

$$A = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \quad B = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad Q_1 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad Q_2 = \rho$$

which enter the Riccati equation (9.5). Solving the problem using Matlab for $\rho = 0.01$, $\rho = 0.1$, $\rho = 1$, $\rho = 10$ gives the closed loop responses shown below (state in full, control signal in dashed):



In this example, the closed loop poles can be calculated as $s = 2^{-1/2} \rho^{-1/4} (-1 \pm i)$, which clearly shows the influence of the control signal penalty ρ on the speed of the closed-loop system. \square

Stochastic interpretation of linear-quadratic control

When we derived the optimal controller above we assumed that the initial state is x_0 and that there are no disturbances. The same solution is valid when the process is disturbed by white noise (essentially setting a new initial state at each time instant). We can hence equivalently formulate a stochastic version of the LQ problem as

$$\begin{aligned} &\text{minimize} && J = \mathbb{E} |z^2| = \mathbb{E} \{ x^T Q_1 x + 2x^T Q_{12} u + u^T Q_2 u \} \\ &\text{subject to} && \dot{x}(t) = Ax(t) + Bu(t) + w(t) \end{aligned}$$

where w is white noise with intensity R . The goal is now to minimize the stationary variance of the performance output vector z . The same Riccati equation and associated solution (S , L) is valid also for this case. The optimal cost is given by

$$J^* = \mathbb{E} w^T S w = \text{tr} S R.$$

9.3 An alternative solution

Dynamic programming provides an elegant way to derive the optimal control law among all possible policies, including non-linear and linear ones, time-varying and constant ones, etc. If

we restrict ourselves to seek out the best possible linear, static feedback law for an infinite horizon LQ problem, a simpler solution is possible. Again consider the optimization problem (9.1), (9.2). Postulating a linear, stabilizing control law,

$$u = -Lx,$$

the cost function can be written as

$$J = \int_0^{\infty} x^T(t) (Q_1 - Q_{12}L - L^T Q_{12}^T + L^T Q_2 L) x(t) dt.$$

The closed-loop system is given by

$$\dot{x} = (A - BL)x$$

with the solution

$$x(t) = e^{(A-BL)t} x_0.$$

Introducing $S = e^{(A-BL)^T t} (Q_1 - Q_{12}L - L^T Q_{12}^T + L^T Q_2 L) e^{(A-BL)t}$, a Lyapunov equation for the cost $J = x(0)^T S x(0)$ is

$$(A - BL)^T S + S(A - BL) + Q_1 - Q_{12}L - L^T Q_{12}^T + L^T Q_2 L = 0.$$

We minimize S by completing the square,

$$A^T S + SA + Q_1 + (L^T Q_2 - SB - Q_{12}) Q_2^{-1} (L^T Q_2 - SB - Q_{12})^T - (SB + Q_{12}) Q_2^{-1} (SB + Q_{12})^T = 0$$

with minimum attained for

$$L = Q_2^{-1} (B^T S + Q_{12}^T)$$

The minimum S is then given by the algebraic Riccati equation

$$A^T S + SA + Q_1 - (SB + Q_{12}) Q_2^{-1} (SB + Q_{12})^T = 0$$

which is the same as (9.5).

9.4 Stability and robustness of linear-quadratic control

The linear-quadratic optimal controller has remarkable robustness properties, as we will see below. We start by showing that the closed-loop system is guaranteed to be stable.

THEOREM 9.1—STABILITY OF THE CLOSED-LOOP SYSTEM

Assume that

$$Q = \begin{pmatrix} Q_1 & Q_{12} \\ Q_{12}^T & Q_2 \end{pmatrix} > 0$$

and that there exists a solution $S > 0$ to the algebraic Riccati equation (9.5). Then the optimal controller $u(t) = -Lx(t)$ gives a stable closed-loop system $\dot{x}(t) = (A - BL)x(t)$.

PROOF: Taking the time derivative of $x^T S x$ gives

$$\begin{aligned} \frac{d}{dt} x^T S x &= 2x^T S \dot{x} = 2x^T S (Ax + Bu) = \text{[Bellman's equation]} \\ &= -\left(x^T Q_1 x + 2x^T Q_{12} u + u^T Q_2 u\right) < 0 \text{ for } x(t) \neq 0 \end{aligned}$$

Hence $x^T(t) S x(t)$ is decreasing and tends to zero as $t \rightarrow \infty$. □

The linear-quadratic controller also has remarkable robustness properties, as stated by the following theorem:

THEOREM 9.2—ROBUSTNESS OF THE LQ CONTROLLER

For a system with scalar control signal u and the LQ design parameters

$$Q_1 > 0, \quad Q_{12} = 0, \quad Q_2 = \rho > 0$$

the distance from the resulting optimal loop gain $L(i\omega I - A)^{-1}B$ to -1 is never smaller than 1. This implies that the phase margin is at least 60° and that the gain margin is infinite.

PROOF: Using the Riccati equation

$$0 = Q_1 + A^T S + SA - L^T Q_2 L, \quad L = Q_2^{-1}(SB + Q_{12})^T$$

it can be verified that

$$[I + L(i\omega - A)^{-1}B]^* Q_2 [I + L(i\omega - A)^{-1}B] = \begin{bmatrix} (i\omega - A)^{-1}B \\ I \end{bmatrix}^* \begin{bmatrix} Q_1 & Q_{12} \\ Q_{12}^* & Q_2 \end{bmatrix} \begin{bmatrix} (i\omega - A)^{-1}B \\ I \end{bmatrix}$$

In particular, with $Q_1 > 0, Q_{12} = 0, Q_2 = \rho > 0$ it holds that


$$[1 + L(i\omega - A)^{-1}B]^* \rho [1 + L(i\omega - A)^{-1}B] = B^T [(i\omega - A)^{-1}]^* Q_1 (i\omega - A)^{-1}B + \rho \geq \rho$$

Dividing by ρ gives

$$|1 + L(i\omega - A)^{-1}B|^2 \geq 1$$

□

An example of the loop gain resulting from a linear-quadratic controller design is shown below:



It is seen that the Nyquist curve always stays outside a unit circle centered at the critical point -1 .

It should be noted that the robustness result is only valid for the pure state feedback case, i.e., when the full state vector is available for feedback.