19 LINEAR QUADRATIC REGULATOR

19.1 Introduction

The simple form of loopshaping in scalar systems does not extend directly to multivariable (MIMO) plants, which are characterized by transfer matrices instead of transfer functions. The notion of optimality is closely tied to MIMO control system design. Optimal controllers, i.e., controllers that are the *best* possible, according to some figure of merit, turn out to generate only stabilizing controllers for MIMO plants. In this sense, optimal control solutions provide an automated design procedure – we have only to decide what figure of merit to use. The linear quadratic regulator (LQR) is a well-known design technique that provides practical feedback gains.

19.2   Full-State Feedback

For the derivation of the linear quadratic regulator, we assume the plant to be written in state-space form $\dot{x} = Ax + Bu$, and that *all* of the n states x are available for the controller. The feedback gain is a matrix K, implemented as $u = -K(x - x_{desired})$. The system dynamics are then written as:

$$\dot{x} = (A - BK)x + BKx_{desired}. \tag{210}$$

$x_{desired}$ represents the vector of desired states, and serves as the external input to the closed-loop system. The "A-matrix" of the closed loop system is $(A - BK)$, and the "B-matrix" of the closed-loop system is BK. The closed-loop system has exactly as many outputs as inputs: n. The column dimension of B equals the number of channels available in u, and must match the row dimension of K. Pole-placement is the process of placing the poles of $(A - BK)$ in stable, suitably-damped locations in the complex plane.

19.3   The Maximum Principle

First we develop a general procedure for solving optimal control problems, using the calculus of variations. We begin with the statement of the problem for a fixed end time t_f:

$$\text{choose } u(t) \text{ to minimize } J = \psi(x(t_f)) + \int_{t_o}^{t_f} L(x(t), u(t), t)dt \tag{211}$$

$$\text{subject to } \dot{x} = f(x(t), u(t), t) \tag{212}$$

$$x(t_o) = x_o \tag{213}$$

where $\psi(x(t_f), t_f)$ is the *terminal cost*; the total cost J is a sum of the terminal cost and an integral along the way. We assume that $L(x(t), u(t), t)$ is nonnegative. The first step is to augment the cost using the costate vector $\lambda(t)$.

$$\bar{J} = \psi(x(t_f)) + \int_{t_o}^{t_f} (L + \lambda^T(f - \dot{x}))dt \tag{214}$$

Clearly, $\lambda(t)$ can be anything we choose, since it multiplies $f - \dot{x} = 0$. Along the optimum trajectory, variations in J and hence $\bar{J}$ should vanish. This follows from the fact that J is chosen to be continuous in x, u, and t. We write the variation as

$$\delta\bar{J} = \psi_x \delta x(t_f) + \int_{t_o}^{t_f} \left[L_x \delta x + L_u \delta u + \lambda^T f_x \delta x + \lambda^T f_u \delta u - \lambda^T \delta\dot{x} \right] dt, \tag{215}$$

where subscripts denote partial derivatives. The last term above can be evaluated using integration by parts as

$$-\int_{t_o}^{t_f} \lambda^T \delta\dot{x}dt = -\lambda^T(t_f)\delta x(t_f) + \lambda^T(t_o)\delta x(t_o) + \int_{t_o}^{t_f} \dot{\lambda}^T \delta x dt, \tag{216}$$

and thus

$$\delta \bar{J} = \psi_x(x(t_f))\delta x(t_f) + \int_{t_o}^{t_f}(L_u + \lambda^T f_u)\delta u dt + \tag{217}$$

$$\int_{t_o}^{t_f}(L_x + \lambda^T f_x + \dot{\lambda}^T)\delta x dt - \lambda^T(t_f)\delta x(t_f) + \lambda^T(t_o)\delta x(t_o).$$

Now the last term is zero, since we cannot vary the initial condition of the state by changing something later in time - it is a fixed value. This way of writing $\bar{J}$ makes it clear that there are three components of the variation that must independently be zero (since we can vary any of x, u, or $x(t_f)$):

$$L_u + \lambda^T f_u = 0 \tag{218}$$
$$L_x + \lambda^T f_x + \dot{\lambda}^T = 0 \tag{219}$$
$$\psi_x(x(t_f)) - \lambda^T(t_f) = 0. \tag{220}$$

The second and third requirements are met by explicitly setting

$$\dot{\lambda}^T = -L_x - \lambda^T f_x \tag{221}$$
$$\lambda_T(t_f) = \psi_x(x(t_f)). \tag{222}$$

The evolution of λ is given in *reverse time*, from a final state to the initial. Hence we see the primary difficulty of solving optimal control problems: the state propogates forward in time, while the costate propogates backward. The state and costate are coordinated through the above equations.

19.4   Gradient Method Solution for the General Case

Numerical solutions to the general problem are iterative, and the simplest approach is the gradient method. It is outlined as follows:

1. For a given x_o, pick a control history $u(t)$.

2. Propogate $\dot{x} = f(x, u, t)$ forward in time to create a state trajectory.

3. Evaluate $\psi_x(x(t_f))$, and the propogate the costate backward in time from t_f to t_o, using Equation 221.

4. At each time step, choose $\delta u = -K(L_u + \lambda^T f_u)$, where K is a positive scalar or a positive definite matrix in the case of multiple input channels.

5. Let $u = u + \delta u$.

6. Go back to step 2 and repeat loop until solution has converged.

The first three steps are consistent in the sense that x is computed directly from $x(t_o$ and u, and λ is computed directly from x and $x(t_f)$. All of $\delta \bar{J}$ in Equation 217 except the integral with δu is therefore eliminated explicitly. The choice of δu in step 4 then achieves $\delta \bar{J} < 0$, unless $\delta u = 0$, in which case the problem is solved.

The gradient method is quite popular in applications and for complex problems may be the only way to effectively develop a solution. The major difficulties are computational expense, and the requirement of having a reasonable control trajectory to begin.

19.5   LQR Solution

In the case of the Linear Quadratic Regulator (with zero terminal cost), we set $\psi = 0$, and

$$L = \frac{1}{2}x^T Q x + \frac{1}{2}u^T R u, \tag{223}$$

where the requirement that $L \geq 0$ implies that both Q and R are positive definite. In the case of linear plant dynamics also, we have

$$
\begin{aligned}
L_x &= x^T Q & (224)\\
L_u &= u^T R & (225)\\
f_x &= A & (226)\\
f_u &= B, & (227)
\end{aligned}
$$

so that

$$
\begin{aligned}
\dot{x} &= Ax + Bu & (228)\\
x(t_o) &= x_o & (229)\\
\dot{\lambda} = -Qx - A^T\lambda & & (230)\\
\lambda(t_f) &= 0 & (231)\\
Ru + B^T\lambda &= 0. & (232)
\end{aligned}
$$

Since the systems are clearly linear, we try a connection $\lambda = Px$. Inserting this into the $\dot{\lambda}$ equation, and then using the $\dot{x}$ equation, and a substitution for u, we obtain

$$PAx + A^T P x + Q x - PBR^{-1}B^T P x + \dot{P} = 0. \tag{233}$$

This has to hold for all x, so in fact it is a matrix equation, the *matrix Riccati equation*. The steady-state solution is given satisfies

$$PA + A^T P + Q - PBR^{-1}B^T P = 0. \tag{234}$$

19.6   Optimal Full-State Feedback

This equation is the *matrix algebraic Riccati equation* (MARE), whose solution P is needed to compute the optimal feedback gain K. The MARE is easily solved by standard numerical tools in linear algebra.

The equation $Ru + B^T \lambda = 0$ gives the feedback law:

$$u = -R^{-1}B^T Px. \tag{235}$$

19.7   Properties and Use of the LQR

Static Gain. The LQR generates a static gain matrix K, which is not a dynamical system. Hence, the order of the closed-loop system is the same as that of the plant.

Robustness. The LQR achieves infinite gain margin: $k_g = \infty$, implying that the loci of (PC) (scalar case) or $(det(I+PC)-1)$ (MIMO case) approach the origin along the imaginary axis. The LQR also guarantees phase margin $\gamma \geq 60$ degrees. This is in good agreement with the practical guidelines for control system design.

Output Variables. In many cases, it is not the states x which are to be minimized, but the output variables y. In this case, we set the weighting matrix $Q = C^T Q' C$, since $y = Cx$, and the auxiliary matrix Q' weights the plant output.

Behavior of Closed-Loop Poles: Expensive Control. When $R >> C^T Q' C$, the cost function is dominated by the control effort u, and so the controller minimizes the control action itself. In the case of a completely stable plant, the gain will indeed go to zero, so that the closed-loop poles approach the open-loop plant poles in a manner consistent with the scalar root locus.

The optimal control *must* always stabilize the closed-loop system, however, so there should be some account made for unstable plant poles. The expensive control solution puts stable closed-loop poles at the *mirror* images of the unstable plant poles.

Behavior of Closed-Loop Poles: Cheap Control. When $R << C^T Q' C$, the cost function is dominated by the output errors y, and there is no penalty for using large u. There are two groups of closed-loop poles. First, poles are placed at stable plant zeros, and at the mirror images of the unstable plant zeros. This part is akin to the high-gain limiting case of the root locus. The remaining poles assume a Butterworth pattern, whose radius increases to infinity as R becomes smaller and smaller.

The Butterworth pattern refers to an arc in the stable left-half plane, as shown in the figure. The angular separation of n closed-loop poles on the arc is constant, and equal to $180°/n$. An angle $90°/n$ separates the most lightly-damped poles from the imaginary axis.
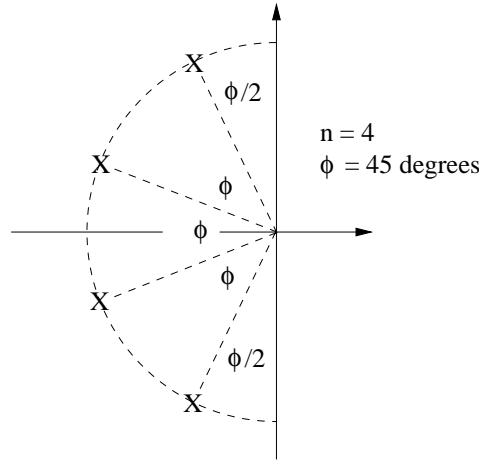
19.8   Proof of the Gain and Phase Margins

The above stated properties of the LQR control are not difficult to prove. We need an intermediate tool (which is very useful in its own right) to prove it: the Liapunov function. Consider a scalar, continuous, positive definite function of the state $V(x)$. If this function is always decreasing, except at the origin, then the origin is asymptotically stable. There is an intuitive feel to the Liapunov function; it is like the energy of a physical system. If it can be shown that energy is leaving such a system, then it will "settle." It suffices to find one Liapunov function to show that a system is stable.

Now, in the case of the LQR control, pick
$V(x) = \frac{1}{2}x^T P x$.
It can be shown that P is positive definite; suppose that instead of the design gain B, we have actual gain B', giving (with constant P based on the design system)

$$
\begin{aligned}
\dot{V} &= 2x^T P \dot{x} & (236) \\
&= 2x^T P(Ax - B'R^{-1}B^T Kx) & (237) \\
&= x^T(-Q + PBR^{-1}B^T P)x - 2x^T KB'R^{-1}B^T Px. & (238)
\end{aligned}
$$

where we used the Riccati equation in a substitution. Since we require $\dot{V} < 0$ for $x \neq 0$, we must have

$$
Q + PBG - 2P(B - B')G > 0, \tag{239}
$$

where $G = R^{-1}B^T P$. Let $B' = BN$, where N is a diagonal matrix of random, unknown gains on the control channels. The above equation becomes

$$
Q - PB(I - 2N)R^{-1}B^T P > 0. \tag{240}
$$

This is satisfied if, for all channels, $N_{i,i} > 1/2$. Thus, we see that the LQR provides for one-half gain reduction, and infinite gain amplification, in all channels.

The phase margin is seen (roughly) by allowing N to be a diagonal matrix of transfer functions $N_{i,i} = e^{j\phi_i}$. In this case, we require that the *real part* of $N_{i,i}$ is greater than one-half. This is true for all $|\phi_i| < 60°$, and hence sixty degrees of phase margin is provided in all channels.