

EN530.603 Applied Optimal Control
Lecture 6: Linear-Quadratic Regulator Basics
October 1, 2018

Lecturer: Marin Kobilarov

1 Continuous Linear Quadratic Regulator (LQR)

1.1 Finite-Time LQR

Consider a system with dynamics

$$\dot{x} = Ax + Bu$$

which must optimally reach the origin, a task specified by the cost function

$$J = \frac{1}{2}x^T(t_f)P_f x(t_f) + \int_{t_0}^{t_f} \frac{1}{2} [x(t)^T Q(t)x(t) + u(t)^T R(t)u(t)] dt,$$

where P_f and Q are symmetric positive semi-definite matrices and R is a symmetric positive definite matrix. Applying the optimality conditions using the Hamiltonian

$$H = \frac{1}{2} (x^T Q x + u^T R u) + \lambda^T (Ax + Bu)$$

we obtain

$$\dot{\lambda} = -Qx - A^T \lambda, \tag{1}$$

while the control is computed according to

$$Ru + B^T \lambda = 0 \quad \Rightarrow \quad u = -R^{-1} B^T \lambda$$

and transversality conditions become

$$\lambda(t_f) = P_f x(t_f). \tag{2}$$

The optimal state $(x(t), \lambda(t))$ then evolves according to the EL equations

$$\begin{pmatrix} \dot{x} \\ \dot{\lambda} \end{pmatrix} = \underbrace{\begin{pmatrix} A & -BR^{-1}B^T \\ -Q & -A^T \end{pmatrix}}_{=M} \begin{pmatrix} x \\ \lambda \end{pmatrix}.$$

with a final boundary condition (2). When A, B, Q, R are constant this can be solved for given $y(0) = (x(0), \lambda(0))$ as $y(t) = e^{tM} y(0)$ for instance by computing the Laplace transform of $(sI - M)^{-1}$. This becomes difficult in high-dimensions, but thankfully there's an easier way.

Kalman showed that the multipliers $\lambda(t)$ are in fact linear function of the states, i.e.

$$\lambda(t) = P(t)x(t).$$

Hence the control can be written according to

$$u = -R^{-1}B^T Px \equiv Kx.$$

The matrix P can now be computed by noting that

$$\dot{\lambda} = \dot{P}x + P\dot{x}$$

which is equivalent to

$$-Qx - A^T Px = \dot{P}x + PAx - PBR^{-1}B^T Px.$$

A solution exists then if we can find a P which satisfies

$$\dot{P} = -A^T P - PA + PBR^{-1}B^T P - Q, \quad P(t_f) = P_f$$

This is called the *Riccati ODE* and is integrated from t_f to t_0 backwards in time. After $P(t)$ is found the control is updated according to

$$u(t) = -R^{-1}B^T P(t)x(t).$$

Note that it turns out that the optimal control u is in a linear feedback form, i.e. it is a linear function of the state x . This means that we have obtained not only a single optimal control signal from the start state $x(t_0)$ but also an optimal feedback controller from any state $x(t)$ for $t > t_0$. Therefore, we have completely eliminated the need for an additional controller to physically bring the system to the equilibrium state $x = 0$. Furthermore, when the system deviates from the initially computed path e.g. due to disturbances, the same $P(t)$ computed once in the beginning (at time $t = t_0$) can be used from the perturbed state $x(t)$.

Stability. In this context it is instructive to study the stability of the optimal control regarded as a feedback controller $u = Kx$. The closed-loop matrix is

$$A + BK = A - BR^{-1}B^T P$$

and one should be able to verify that the real parts of its eigenvalues are negative, i.e. that

$$\text{Real}[\text{eig}(A + BK)] < 0.$$

Here we have assumed that the system (A, B) is controllable.

Example 1. Linear-quadratic problem. Consider the system $\dot{x}(t) = Ax(t) + Bu(t)$ with initial condition $x(0) = x_0$ and quadratic cost functional

$$J = \frac{1}{2} \int_0^{t_f} x(t)^2 + u(t)^2 dt$$


We have $A = B = Q = R = 1$ and the Riccati equation becomes

$$\dot{P} = -2P + P^2 - 1,$$

with final condition $P(t_f) = 0$. The solution can be obtained analytically and is

$$P(t) = 1 - \sqrt{2} \tanh(\sqrt{2}(t - t_f + (\sqrt{2} \text{atanh}(\sqrt{2}/2))/2))$$

The gain computed with $t_f = 5$ is given below



Note: one standard way to get the solution above is to set $P(t) = -\frac{\dot{b}(t)}{b(t)}$ for some $b(t)$:

$$\dot{P} = -\frac{\ddot{b}}{b} + \frac{\dot{b}^2}{b^2} = -\frac{\ddot{b}}{b} + k^2 \Rightarrow \ddot{b} = -2\dot{b} + b$$

The solution of the second order linear ODE

$$\ddot{b} = -2\dot{b} + b, \quad t < t_f, \quad (3)$$

$$\dot{b}(t_f) = 0, \quad b(t_f) = 1 \quad (4)$$

is then computed from which we find $P(t)$.

Example 2. 2-dim Linear-quadratic problem. Consider the dynamics


$$\dot{x}_1 = x_2, \quad (5)$$

$$\dot{x}_2 = 2x_1 - x_2 + u \quad (6)$$

and cost function

$$J = \frac{1}{2} \int_0^{t_f} [x_1^2 + \frac{1}{2}x_2^2 + \frac{1}{4}u^2] dt$$

We integrate the Riccati ODE numerically using $t_f = 5$ and start state $x(0) = (-4, 4)$. The following matrix $P(t)$, inputs $u(t)$ and state histories $x(t)$ are obtained:



Optimal Cost. Assume the system is at state $x(t)$. We can compute the resulting optimal cost from time t to time t_f as follows:

$$\begin{aligned} J(t) &= \frac{1}{2}x^T(t_f)P_f x(t_f) + \int_t^{t_f} \frac{1}{2} [x^T Q x + u^T R u] dt, \\ &= \frac{1}{2}x^T(t_f)P_f x(t_f) + \int_t^{t_f} \frac{1}{2} [x(t)^T (Q + K^T R K) x] dt, \\ &= \frac{1}{2}x^T(t_f)P_f x(t_f) - \int_t^{t_f} \frac{d}{dt} \left(\frac{1}{2}x^T P x \right) dt = \frac{1}{2}x^T(t)P(t)x(t) \end{aligned}$$

This will be important for several reasons: stability, cost-to-go, etc...

1.2 Infinite-Time LQR

Consider the state equations $\dot{x} = Ax + Bu$ with cost function

$$J = \int_{t_0}^{\infty} \frac{1}{2} [x(t)^T Q(t)x(t) + u(t)^T R(t)u(t)] dt,$$

The Riccati ODE has the same form but at $t = \infty$ reaches the stationary value

$$0 = -A^T P - PA + PBR^{-1}B^T P - Q,$$

which called the *algebraic Riccati equation*.

The equation can be solved using Matlab

$$[K, P] = \text{lqr}(A, B, Q, R)$$

and $u = Kx$ can then be used as the input from state x .

1.3 Trajectory Tracking

Consider the problem of not stabilizing to the origin, i.e. $x \rightarrow 0$ but tracking a given reference trajectory $x_d(t)$, i.e. $x \rightarrow x_d$. This is often useful when x_d was an optimized trajectory for a complex nonlinear system with constraints, which we cannot reoptimize in real time but can easily track.

One approach is to formulate the *error state*

$$e = x - x_d,$$

the control error (assuming we have the control u_d which produced x_d)

$$v = u - u_d,$$

and essentially apply LQR to the dynamics of e subject to “virtual” inputs v . In particular, note that in the general nonlinear case we have

$$\dot{e} = \dot{x} - \dot{x}_d = f(x, u) - f(x_d, u_d) = f(x_d + e, u_d + v) - f(x_d, u_d) \equiv F(e, v, x_d(t), u_d(t)),$$

or in other words we have obtained a new ODE in for e , v , and time-varying parameters. In the linear case the we have

$$\dot{e} = Ae + Bv,$$

and so once the optimal $v = Ke$ is computed using standard LQR the actual control u is recovered by

$$u = K(x - x_d) + u_d.$$

Another approach is to directly obtain necessary conditions. In particular, let the cost be defined as

$$J = \frac{1}{2} \|x(t_f) - x_d(t_f)\|_{P_f}^2 + \frac{1}{2} \int_{t_0}^{t_f} \left\{ \|x(t) - x_d(t)\|_{Q(t)}^2 + \|u(t)\|_{R(t)}^2 \right\} dt$$

The Hamiltonian is

$$H = \frac{1}{2} \|x(t) - x_d(t)\|_{Q(t)}^2 + \frac{1}{2} \|u(t)\|_{R(t)}^2 + \lambda^T (Ax + Bu)$$

and the necessary conditions become

$$\dot{\lambda} = -Q(x - x_d) - A^T \lambda,$$

and

$$u = -R^{-1} B^T \lambda,$$

while the transversality conditions are

$$\lambda(t_f) = P_f(x(t_f) - x_d(t_f))$$

In order to derive a control law, we follow the same reasoning as in the regular case and assume that the multiplier is of the form

$$\lambda = Px + s$$

We will now attempt to derive expressions for P and s that satisfy the necessary conditions. Differentiating we have

$$\dot{\lambda} = \dot{P}x + P\dot{x} + \dot{s},$$

which is equivalent to

$$-Qx + Qx_d - A^T(Px + s) = \dot{P}x + (PA - PBR^{-1}B^T P)x - PBR^{-1}B^T s + \dot{s}$$

or

$$(-A^T P - PA - Q + PBR^{-1}B^T P - \dot{P})x + (\dot{s} + A^T s - PBR^{-1}B^T s - Qx_d) = 0.$$

If we set

$$\dot{P} = -A^T P - PA - Q + PBR^{-1}B^T P, \tag{7}$$

$$\dot{s} = -A^T s + PBR^{-1}B^T s + Qx_d \tag{8}$$

and integrate them backwards starting with

$$P(t_f) = P_f, \tag{9}$$

$$s(t_f) = -P_f x_d \tag{10}$$

then the necessary conditions will be satisfied.

Example 3. Non-zero signal tracking. Consider the dynamics

$$\dot{x}_1 = x_2, \tag{11}$$

$$\dot{x}_2 = 2x_1 - x_2 + u \tag{12}$$

and cost function

$$J = 4(x_1 - 1)^2 + \frac{1}{2} \int_0^{t_f} [2(x_1 - 1)^2 + 0.005u^2] dt$$

which corresponds to driving *only* the x_1 coordinate to 1 while minimizing control effort. This is a simplified version of the more general condition $x(t) \rightarrow x_d(t)$, where $x_{d1} = 1$ is constant. We integrate the Riccati ODE numerically using $t_f = 5$ to obtain the following matrix $P(t)$, inputs $u(t)$ and state histories $x(t)$:

