
Optimization-Based Control

Richard M. Murray
Control and Dynamical Systems
California Institute of Technology

DRAFT v2.0a, June 6, 2008
© California Institute of Technology
All rights reserved.

This manuscript is for review purposes only and may not be reproduced, in whole or in part, without written consent from the author.

Chapter 2

Optimal Control

This set of notes expands on Chapter 6 of *Feedback Systems* by Åström and Murray (ÅM08), which introduces the concepts of reachability and state feedback. We also expand on topics in Section 7.5 of ÅM08 in the area of feedforward compensation. Beginning with a review of optimization, we introduce the notion of Lagrange multipliers and provide a summary of the Pontryagin's maximum principle. Using these tools we derive the linear quadratic regulator for linear systems and describe its usage.

Prerequisites. Readers should be familiar with modeling of input/output control systems using differential equations, linearization of a system around an equilibrium point and state space control of linear systems, including reachability and eigenvalue assignment. Some familiarity with optimization of nonlinear functions is also assumed.

2.1 Review: Optimization

Consider first the problem of finding the minimum of a smooth function $F : \mathbb{R}^n \rightarrow \mathbb{R}$. That is, we wish to find a point $x^* \in \mathbb{R}^n$ such that $F(x^*) \leq F(x)$ for all $x \in \mathbb{R}^n$. A necessary condition for x^* to be a minimum is that the gradient of the function be zero at x^* :

$$\frac{\partial F}{\partial x}(x^*) = 0.$$

The function $F(x)$ is often called a cost function and x^* is the optimal value for x . Figure 2.1 gives a graphical interpretation of the necessary condition for a minimum. Note that these are *not* sufficient conditions; the points x_1 and x_2 and x^* in the figure all satisfy the necessary condition but only one is the (global) minimum.

The situation is more complicated if constraints are present. Let $G_i : \mathbb{R}^n \rightarrow \mathbb{R}$, $i = 1, \dots, k$ be a set of smooth functions with $G_i(x) = 0$ representing the constraints. Suppose that we wish to find $x^* \in \mathbb{R}^n$ such that $G_i(x^*) = 0$ and $F(x^*) \leq F(x)$ for all $x \in \{x \in \mathbb{R}^n : G_i(x) = 0, i = 1, \dots, k\}$. This situation can be visualized as constraining the point to a surface (defined by the constraints) and searching for the minimum of the cost function along this surface, as illustrated in Figure 2.2a.

A necessary condition for being at a minimum is that there are no directions tangent to the constraints that also decrease the cost. The normal

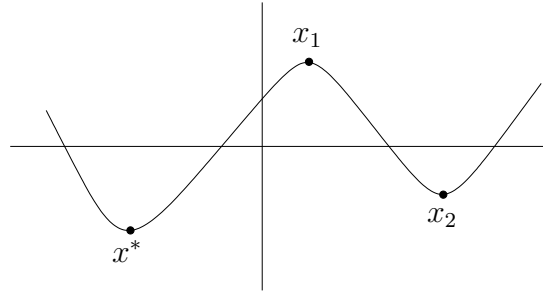


Figure 2.1: Optimization of functions. The minimum of a function occurs at a point where the gradient is zero.

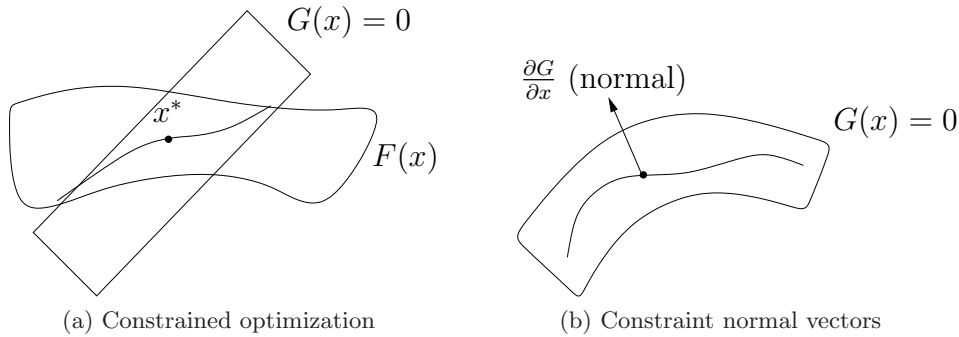


Figure 2.2: Optimization with constraints. (a) We seek a point x^* that minimizes $F(x)$ while lying on the surface $G(x) = 0$. (b) We can parameterize the constrained directions by computing the gradient of the constraint G .

directions to the surface are spanned by $\partial G_i / \partial x$, as shown in Figure 2.2b. A necessary condition is that the gradient of F is spanned by vectors that are normal to the constraints, so that the only directions that increase the cost violate the constraints. We thus require that there exist scalars λ_i , $i = 1, \dots, k$ such that

$$\frac{\partial F}{\partial x}(x^*) + \sum_{i=1}^k \lambda_i \frac{\partial G_i}{\partial x}(x^*) = 0.$$

If we let $G = [G_1 \ G_2 \ \dots \ G_k]^T$, then we can write this condition as

$$\frac{\partial F}{\partial x} + \lambda^T \frac{\partial G}{\partial x} = 0$$

the term $\partial F / \partial x$ is the usual (gradient) optimality condition while the term $\partial G / \partial x$ is used to “cancel” the gradient in the directions normal to the constraint.

An alternative condition can be derived by modifying the cost function to incorporate the constraints. Defining $\tilde{F} = F + \sum \lambda_i G_i$, the necessary

condition becomes

$$\frac{\partial \tilde{F}}{\partial x}(x^*) = 0.$$

The scalars λ_i are called *Lagrange multipliers*. Minimizing \tilde{F} is equivalent to the optimization given by

$$\min_x (F(x) + \lambda^T G(x)).$$

The variables λ can be regarded as free variables, which implies that we need to choose x such that $G(x) = 0$ in order to insure the cost is minimized. Otherwise, we could choose λ to generate a large cost.

Example 2.1 Two free variables with a constraint

Consider the cost function given by

$$F(x) = F_0 + (x_1 - a)^2 + (x_2 - b)^2,$$

which has an unconstrained minimum at $x = (a, b)$. Suppose that we add a constraint $G(x) = 0$ given by

$$G(x) = x_1 - x_2.$$

With this constraint, we seek to optimize F subject to $x_1 = x_2$. Although in this case we could easily do this by simple substitution, we instead carry out the more general procedure using Lagrange multipliers.

The augmented cost function is given by

$$\tilde{F}(x) = F_0 + (x_1 - a)^2 + (x_2 - b)^2 + \lambda(x_1 - x_2),$$

where λ is the Lagrange multiplier for the constraint. Taking the derivative of F , we have

$$\frac{\partial F}{\partial x} = [2x_1 - 2a + \lambda \quad 2x_2 - 2b - \lambda].$$

Setting each of these equations equal to zero, we have that at the minimum

$$x_1^* = a - \lambda/2, \quad x_2^* = b + \lambda/2.$$

The remaining equation that we need is the constraint, which requires that $x_1^* = x_2^*$. Using these three equations, we see that $\lambda^* = a - b$ and we have

$$x_1^* = \frac{a - b}{2}, \quad x_2^* = \frac{a - b}{2}.$$

To verify the geometric view described above, note that the gradients of F and G are given by

$$\frac{\partial F}{\partial x} = [2x_1 - 2a \quad 2x_2 - 2b], \quad \frac{\partial G}{\partial x} = [1 \quad -1].$$

At the optimal value of the (constrained) optimization, we have

$$\frac{\partial F}{\partial x} = [b - a \quad a - b], \quad \frac{\partial G}{\partial x} = [1 \quad -1].$$

Although the derivative of F is not zero, it is pointed in a direction that is normal to the constraint, and hence we cannot decrease the cost while staying on the constraint surface. ∇

We have focused on finding the minimum of a function. We can switch back and forth between maximum and minimum by simply negating the cost function:

$$\max_x F(x) = \min_x (-F(x))$$

We see that the conditions that we have derived are independent of the sign of F since they only depend on the gradient being zero in approximate directions. Thus finding x^* that satisfies the conditions corresponds to finding an *extremum* for the function.

Very good software is available for solving optimization problems numerically of this sort. The NPSOL and SNOPT libraries are available in FORTRAN (and C). In MATLAB, the `fmin` function can be used to solve a constrained optimization problem.

2.2 Optimal Control of Systems

Consider now the *optimal control problem*:

$$\min_{u(\cdot)} \int_0^T L(x, u) dt + V(x(T))$$

subject to the constraint

$$\dot{x} = f(x, u) \quad x \in \mathbb{R}^n, u \in \mathbb{R}^m.$$

Abstractly, this is a constrained optimization problem where we seek a *feasible trajectory* $(x(t), u(t))$ that minimizes the cost function

$$J(x, u) = \int_0^T L(x, u) dt + V(x(T)).$$

More formally, this problem is equivalent to the “standard” problem of minimizing a cost function $J(x, u)$ where $(x, u) \in L_2[0, T]$ (the set of square integrable functions) and $h(z) = \dot{x}(t) - f(x(t), u(t)) = 0$ models the dynamics. The term $L(x, u)$ is referred to as the integral cost and $V(x(T))$ is the final (or terminal) cost.

There are many variations and special cases of the optimal control problem. We mention a few here:

Infinite horizon optimal control. If we let $T = \infty$ and set $V = 0$, then we seek to optimize a cost function over all time. This is called the *infinite horizon* optimal control problem, versus the *finite horizon* problem with $T < \infty$. Note that if an infinite horizon problem has a solution with finite cost, then the integral cost term $L(x, u)$ must approach zero as $t \rightarrow \infty$.

Linear quadratic (LQ) optimal control. If the dynamical system is linear and the cost function is quadratic, we obtain the *linear quadratic* optimal control problem:

$$\dot{x} = Ax + Bu \quad J = \int_0^T (x^T Qx + u^T Ru) dt + x^T(T)P_1x(T).$$

In this formulation, $Q \geq 0$ penalizes state error, $R > 0$ penalizes the input and $P_1 > 0$ penalizes terminal state. This problem can be modified to track a desired trajectory (x_d, u_d) by rewriting the cost function in terms of $(x - x_d)$ and $(u - u_d)$.

Terminal constraints. It is often convenient to ask that the final value of the trajectory, denoted x_f , be specified. We can do this by requiring that $x(T) = x_f$ or by using a more general form of constraint:

$$\psi_i(x(T)) = 0, \quad i = 1, \dots, q.$$

The fully constrained case is obtained by setting $q = n$ and defining $\psi_i(x(T)) = x_i(T) - x_{i,f}$. For a control problem with a full set of terminal constraints, $V(x(T))$ can be omitted (since its value is fixed).

Time optimal control. If we constrain the terminal condition to $x(T) = x_f$, let the terminal time T be free (so that we can optimize over it) and choose $L(x, u) = 1$, we can find the *time-optimal* trajectory between an initial and final condition. This problem is usually only well-posed if we additionally constrain the inputs u to be bounded.

A very general set of conditions are available for the optimal control problem that captures most of these special cases in a unifying framework. Consider a nonlinear system

$$\begin{aligned} \dot{x} &= f(x, u), \quad x \in \mathbb{R}^n, \\ x(0) &\text{ given,} \quad u \in \Omega \subset \mathbb{R}^p, \end{aligned}$$

where $f(x, u) = (f_1(x, u), \dots, f_n(x, u)) : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}^n$. We wish to minimize a cost function J with terminal constraints:

$$J = \int_0^T L(x, u) dt + V(x(T)), \quad \psi(x(T)) = 0.$$

The function $\psi : \mathbb{R}^n \rightarrow \mathbb{R}^q$ gives a set of q terminal constraints. Analogous to the case of optimizing a function subject to constraints, we construct the *Hamiltonian*:

$$H = L + \lambda^T f = L + \sum \lambda_i f_i.$$

The variables λ are functions of time and are often referred to as the *costate variables*. A set of necessary conditions for a solution to be optimal was derived by Pontryagin [PBG62].

Theorem 2.1 (Maximum Principle). *If (x^*, u^*) is optimal, then there exists $\lambda^*(t) \in \mathbb{R}^n$ and $\nu^* \in \mathbb{R}^q$ such that*

$$\dot{x}_i = \frac{\partial H}{\partial \lambda_i} \quad -\dot{\lambda}_i = \frac{\partial H}{\partial x_i} \quad x(0) \text{ given, } \psi(x(T)) = 0$$

$$\lambda(T) = \frac{\partial V}{\partial x}(x(T)) + \nu^T \frac{\partial \psi}{\partial x}$$

and

$$H(x^*(t), u^*(t), \lambda^*(t)) \leq H(x^*(t), u, \lambda^*(t)) \quad \text{for all } u \in \Omega$$

The form of the optimal solution is given by the solution of a differential equation with boundary conditions. If $u = \operatorname{argmin} H(x, u, \lambda)$ exists, we can use this to choose the control law u and solve for the resulting feasible trajectory that minimizes the cost. The boundary conditions are given by the n initial states $x(0)$, the q terminal constraints on the state $\psi(x(T)) = 0$ and the $n - q$ final values for the Lagrange multipliers

$$\lambda(T) = \frac{\partial V}{\partial x}(x(T)) + \nu^T \frac{\partial \psi}{\partial x}.$$

In this last equation, ν is a free variable and so there are n equations in $n + q$ free variables, leaving $n - q$ constraints on $\lambda(T)$. In total, we thus have $2n$ boundary values.

The maximum principle is a very general (and elegant) theorem. It allows the dynamics to be nonlinear and the input to be constrained to lie in a set Ω , allowing the possibility of bounded inputs. If $\Omega = \mathbb{R}^m$ (unconstrained input) and H is differentiable, then a necessary condition for the optimal input is

$$\frac{\partial H}{\partial u} = 0.$$

We note that even though we are *minimizing* the cost, this is still usually called the maximum principle (an artifact of history).

Sketch of proof. We follow the proof given by Lewis and Syrmos [LS95], omitting some of the details required for a fully rigorous proof. We use the method of Lagrange multipliers, augmenting our cost function by the dynamical constraints and the terminal constraints:

$$\begin{aligned} \tilde{J}(x(\cdot), u(\cdot), \lambda(\cdot), \nu) &= J(x, u) + \int_0^T \lambda^T(t)(\dot{x}(t) - f(x, u)) dt + \nu^T \psi(x(T)) \\ &= \int_0^T (L(x, u) + \lambda^T(t)(\dot{x}(t) - f(x, u))) dt \\ &\quad + V(x(T)) + \nu^T \psi(x(T)). \end{aligned}$$

Note that λ is a function of time, with each $\lambda(t)$ corresponding to the instantaneous constraint imposed by the dynamics. The integral over the interval $[0, T]$ plays the role of the sum of the finite constraints in the regular optimization.

Making use of the definition of the Hamiltonian, the augmented cost becomes

$$\tilde{J}(x(\cdot), u(\cdot), \lambda(\cdot), \nu) = \int_0^T (H(x, u) - \lambda^T(t)\dot{x}) dt + V(x(T)) + \nu^T \psi(x(T)).$$

We can now “linearize” the cost function around the optimal solution $x(t) = x^*(t) + \delta x(t)$, $u(t) = u^*(t) + \delta u(t)$, $\lambda(t) = \lambda^*(t) + \delta \lambda(t)$ and $\nu = \nu^* + \delta \nu$. Taking T as fixed for simplicity (see [LS95] for the more general case), the incremental cost can be written as

$$\begin{aligned} \delta \tilde{J} &= \tilde{J}(x^* + \delta x, u^* + \delta u, \lambda^* + \delta \lambda, \nu^* + \delta \nu) - \tilde{J}(x^*, u^*, \lambda^*, \nu^*) \\ &\approx \int_0^T \left(\frac{\partial H}{\partial x} \delta x + \frac{\partial H}{\partial u} \delta u - \lambda^T \delta \dot{x} + \left(\frac{\partial H}{\partial \lambda} - \dot{x}^T \right) \delta \lambda \right) dt \\ &\quad + \frac{\partial V}{\partial x} \delta x(T) + \nu^T \frac{\partial \psi}{\partial x} \delta x(T) + \delta \nu^T \psi(x(T), u(T)), \end{aligned}$$

where we have omitted the time argument inside the integral and all derivatives are evaluated along the optimal solution.

We can eliminate the dependence on $\delta \dot{x}$ using integration by parts:

$$- \int_0^T \lambda^T \delta \dot{x} dt = -\lambda^T(T) \delta x(T) + \lambda^T(0) \delta x(0) + \int_0^T -\lambda^T \delta x dt.$$

Since we are requiring $x(0) = x_0$, the first term vanishes and substituting this into $\delta \tilde{J}$ yields

$$\begin{aligned} \delta \tilde{J} &\approx \int_0^T \left[\left(\frac{\partial H}{\partial x} - \lambda^T \right) \delta x + \frac{\partial H}{\partial u} \delta u + \left(\frac{\partial H}{\partial \lambda} - \dot{x}^T \right) \delta \lambda \right] dt \\ &\quad + \left(\frac{\partial V}{\partial x} + \nu^T \frac{\partial \psi}{\partial x} - \lambda^T(T) \right) \delta x(T) + \delta \nu^T \psi(x(T), u(T)). \end{aligned}$$

To be optimal, we require $\delta \tilde{J} = 0$ for all δx , δu , $\delta \lambda$ and $\delta \nu$, and we obtain the (local) conditions in the theorem. \square

2.3 Examples

To illustrate the use of the maximum principle, we consider a number of analytical examples. Additional examples are given in the exercises.

Example 2.2 Scalar linear system

Consider the optimal control problem for the system

$$\dot{x} = ax + bu, \tag{2.1}$$

where $x \in \mathbb{R}$ is a scalar state, $u \in \mathbb{R}$ is the input, the initial state $x(t_0)$ is given, and $a, b \in \mathbb{R}$ are positive constants. We wish to find a trajectory

$(x(t), u(t))$ that minimizes the cost function

$$J = \frac{1}{2} \int_{t_0}^{t_f} u^2(t) dt + \frac{1}{2} cx^2(t_f),$$

where the terminal time t_f is given and $c > 0$ is a constant. This cost function balances the final value of the state with the input required to get to that state.

To solve the problem, we define the various elements used in the maximum principle. Our integral and terminal costs are given by

$$L = \frac{1}{2}u^2(t) \quad V = \frac{1}{2}cx^2(t_f).$$

We write the Hamiltonian of this system and derive the following expressions for the costate λ :

$$\begin{aligned} H &= L + \lambda f = \frac{1}{2}u^2 + \lambda(ax + bu) \\ \dot{\lambda} &= -\frac{\partial H}{\partial x} = -a\lambda, \quad \lambda(t_f) = \frac{\partial V}{\partial x} = cx(t_f). \end{aligned}$$

This is a final value problem for a linear differential equation in λ and the solution can be shown to be

$$\lambda(t) = cx(t_f)e^{a(t_f-t)}.$$

The optimal control is given by

$$\frac{\partial H}{\partial u} = u + b\lambda = 0 \quad \Rightarrow \quad u^*(t) = -b\lambda(t) = -bcx(t_f)e^{a(t_f-t)}.$$

Substituting this control into the dynamics given by equation (2.1) yields a first-order ODE in x :

$$\dot{x} = ax - b^2cx(t_f)e^{a(t_f-t)}.$$

This can be solved explicitly as

$$x^*(t) = x(t_0)e^{a(t-t_0)} + \frac{b^2c}{2a}x^*(t_f) \left[e^{a(t_f-t)} - e^{a(t+t_f-2t_0)} \right].$$

Setting $t = t_f$ and solving for $x(t_f)$ gives

$$x^*(t_f) = \frac{2a e^{a(t_f-t_0)}x(t_0)}{2a - b^2c(1 - e^{2a(t_f-t_0)})}$$

and finally we can write

$$\begin{aligned} u^*(t) &= -\frac{2abc e^{a(2t_f-t_0-t)}x(t_0)}{2a - b^2c(1 - e^{2a(t_f-t_0)})} \\ x^*(t) &= x(t_0)e^{a(t-t_0)} + \frac{b^2c e^{a(t_f-t_0)}x(t_0)}{2a - b^2c(1 - e^{2a(t_f-t_0)})} \left[e^{a(t_f-t)} - e^{a(t+t_f-2t_0)} \right]. \end{aligned}$$

We can use the form of this expression to explore how our cost function affects the optimal trajectory. For example, we can ask what happens to the terminal state $x^*(t_f)$ and $c \rightarrow \infty$. Setting $t = t_f$ in equation (2.2) and taking the limit we find that

$$\lim_{c \rightarrow \infty} x^*(t_f) = 0.$$

▽

Example 2.3 Bang-bang control

The time-optimal control program for a linear system has a particularly simple solution. Consider a linear system with bounded input

$$\dot{x} = Ax + Bu, \quad |u| \leq 1,$$

and suppose we wish to minimize the time required to move from an initial state x_0 to a final state x_f . Without loss of generality we can take $x_f = 0$. We choose the cost functions and terminal constraints to satisfy

$$J = \int_0^T 1 dt, \quad \psi(x(T)) = x(T).$$

To find the optimal control, we form the Hamiltonian

$$H = 1 + \lambda^T(Ax + Bu) = 1 + (\lambda^T A)x + (\lambda^T B)u.$$

Now apply the conditions in the maximum principle:

$$\begin{aligned} \dot{x} &= \frac{\partial H}{\partial \lambda} = Ax + Bu \\ -\dot{\lambda} &= \frac{\partial H}{\partial x} = A^T \lambda \\ u &= \arg \min H = -\text{sgn}(\lambda^T B) \end{aligned}$$

The optimal solution always satisfies this equation (since the maximum principle gives a necessary condition) with $x(0) = x_0$ and $x(T) = 0$. It follows that the input is always either +1 or -1, depending on $\lambda^T B$. This type of control is called “bang-bang” control since the input is always on one of its limits. If $\lambda^T(t)B = 0$ then the control is not well defined, but if this is only true for a specific time instant (e.g., $\lambda^T(t)B$ crosses zero) then the analysis still holds. ▽

2.4 Linear Quadratic Regulators

In addition to its use for computing optimal, feasible trajectories for a system, we can also use optimal control theory to design a feedback law $u = \alpha(x)$ that stabilizes a given equilibrium point. Roughly speaking, we do this by continuously resolving the optimal control problem from our current state $x(t)$ and applying the resulting input $u(t)$. Of course, this approach is

impractical unless we can solve explicitly for the optimal control and somehow rewrite the optimal control as a function of the current state in a simple way. In this section we explore exactly this approach for the linear quadratic optimal control problem.

We begin with the the finite horizon, linear quadratic regulator (LQR) problem, given by

$$\begin{aligned}\dot{x} &= Ax + Bu, & x \in \mathbb{R}^n, u \in \mathbb{R}^n, x_0 \text{ given,} \\ \tilde{J} &= \frac{1}{2} \int_0^T (x^T Q_x x + u^T Q_u u) dt + \frac{1}{2} x^T(T) P_1 x(T),\end{aligned}$$

where $Q_x \geq 0$, $Q_u > 0$, $P_1 \geq 0$ are symmetric, positive (semi-) definite matrices. Note the factor of $\frac{1}{2}$ is usually left out, but we included it here to simplify the derivation. (The optimal control will be unchanged if we multiply the entire cost function by 2.)

To find the optimal control, we apply the maximum principle. We begin by computing the Hamiltonian H :

$$H = x^T Q_x x + u^T Q_u u + \lambda^T (Ax + Bu).$$

Applying the results of Theorem 2.1, we obtain the necessary conditions

$$\begin{aligned}\dot{x} &= \left(\frac{\partial H}{\partial \lambda} \right)^T = Ax + Bu & x(0) = x_0 \\ -\dot{\lambda} &= \left(\frac{\partial H}{\partial x} \right)^T = Q_x x + A^T \lambda & \lambda(T) = P_1 x(T) \\ 0 &= \frac{\partial H}{\partial u} = Q_u u + \lambda^T B.\end{aligned} \quad (2.2)$$

The last condition can be solved to obtain the optimal controller

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

which can be substituted into the dynamic equation (2.2) To solve for the optimal control we must solve a *two point boundary value problem* using the initial condition $x(0)$ and the final condition $\lambda(T)$. Unfortunately, it is very hard to solve such problems in general.

Given the linear nature of the dynamics, we attempt to find a solution by setting $\lambda(t) = P(t)x(t)$ where $P(t) \in \mathbb{R}^{n \times n}$. Substituting this into the necessary condition, we obtain

$$\begin{aligned}\dot{\lambda} &= \dot{P}x + P\dot{x} = \dot{P}x + P(Ax - BQ_u^{-1}B^T P)x, \\ -\dot{P}x - PAx + PBQ_u^{-1}BPx &= Q_x x + A^T Px.\end{aligned}$$

This equation is satisfied if we can find $P(t)$ such that

$$-\dot{P} = PA + A^T P - PBQ_u^{-1}B^T P + Q_x, \quad P(T) = P_1. \quad (2.3)$$

This is a *matrix differential equation* that defines the elements of $P(t)$ from

a final value $P(T)$. Solving it is conceptually no different than solving the initial value problem for vector-valued ordinary differential equations, except that we must solve for the individual elements of the matrix $P(t)$ backwards in time. Equation (2.3) is called the *Riccati ODE*.

An important property of the solution to the optimal control problem when written in this form is that $P(t)$ can be solved without knowing either $x(t)$ or $u(t)$. This allows the two point boundary value problem to be separated into first solving a final-value problem and then solving a time-varying initial value problem. More specifically, given $P(t)$ satisfying equation (2.3), we can apply the optimal input

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

and then solve the original dynamics of the system forward in time from the initial condition $x(0) = x_0$. Note that this is a (time-varying) *feedback* control that describes how to move from *any* state to the origin.

An important variation of this problem is the case when we choose $T = \infty$ and eliminate the terminal cost (set $P_1 = 0$). This gives us the cost function

$$J = \int_0^\infty (x^T Q_x x + u^T Q_u u) dt. \quad (2.4)$$

Since we do not have a terminal cost, there is no constraint on the final value of λ or, equivalently, $P(t)$. We can thus seek to find a constant P satisfying equation (2.3). In other words, we seek to find P such that

$$PA + A^T P - PBQ_u^{-1}B^T P + Q_x = 0. \quad (2.5)$$

This equation is called the *algebraic Riccati equation*. Given a solution, we can choose our input as

$$u = -Q_u^{-1}B^T P x.$$

This represents a *constant gain* $K = Q_u^{-1}B^T P$ where P is the solution of the algebraic Riccati equation.

The implications of this result are interesting and important. First, we notice that if $Q_x > 0$ and the control law corresponds to a finite minimum of the cost, then we must have that $\lim_{t \rightarrow \infty} x(t) = 0$, otherwise the cost will be unbounded. This means that the optimal control for moving from any state x to the origin can be achieved by applying a feedback $u = -Kx$ for K chosen as described as above and letting the system evolve in closed loop. More amazingly, the gain matrix K can be written in terms of the solution to a (matrix) quadratic equation (2.5). This quadratic equation can be solved numerically: in MATLAB the command `K = lqr(A, B, Qx, Qu)` provides the optimal feedback compensator.

In deriving the optimal quadratic regulator, we have glossed over a number of important details. It is clear from the form of the solution that we must have $Q_u > 0$ since its inverse appears in the solution. We would typically also have $Q_x > 0$ so that the integral cost is only zero when $x = 0$,

but in some instances we might only care about certain states, which would imply that $Q_x \geq 0$. For this case, if we let $Q_x = H^T H$ (always possible), our cost function becomes

$$J = \int_0^\infty x^T H^T H x + u^T Q_u u dt = \int_0^\infty \|Hx\|^2 + u^T Q_u u dt.$$

A technical condition for the optimal solution to exist is that the pair (A, H) be *observable*. This makes sense intuitively by considering $y = Hx$ as an output. If y is not observable then there may be non-zero initial conditions that produce no output and so the cost would be zero. This would lead to an ill-conditioned problem and hence we will require that $Q_x \geq 0$ satisfy an appropriate observability condition.

We summarize the main results as a theorem.

Theorem 2.2. *Consider a linear control system with quadratic cost:*

$$\dot{x} = Ax + Bu \quad J = \int_0^\infty x^T Q_x x + u^T Q_u u dt.$$

Assume that the system defined by (A, B) is reachable, $Q_x = Q_x^T \geq 0$ and $Q_u = Q_u^T > 0$. Further assume that $Q_u = H^T H$ and that the linear system with dynamics matrix A and output matrix H is observable. Then the optimal controller satisfies

$$u = -Q_u^{-1} B^T P x, \quad PA + A^T P - PBQ_u^{-1} B^T P = -Q_x,$$

and the minimum cost from initial condition $x(0)$ is given by $J^ = x^T(0)Px(0)$.*

The basic form of the solution follows from the necessary conditions, with the theorem asserting that a constant solution exists for $T = \infty$ when the additional conditions are satisfied. The full proof can be found in standard texts on optimal control, such as Lewis and Syrmos [LS95] or [AF06].

Example 2.4 Optimal control of a double integrator

Consider a double integrator system

$$\frac{dx}{dt} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} x + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u.$$

with quadratic cost given by

$$Q_x = \begin{bmatrix} q^2 & 0 \\ 0 & 0 \end{bmatrix}, \quad Q_u = 1.$$

The optimal control is given by the solution of matrix Riccati equation (2.5). Let P be a symmetric positive definite matrix of the form

$$P = \begin{bmatrix} a & b \\ b & c \end{bmatrix}.$$

Then the Riccati equation becomes

$$\begin{bmatrix} -b^2 + q^2 & a - bc \\ a - bc & 2b - c^2 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix},$$

which has solution

$$P = \begin{bmatrix} \sqrt{2q} & q \\ q & \sqrt{2q^3} \end{bmatrix}.$$

The controller is given by

$$K = R^{-1}B^T P = [1/q \quad \sqrt{2/q}].$$

The feedback law minimizing the given cost function is then $u = -Kx$.

To better understand the structure of the optimal solution, we examine the eigenstructure of the closed loop system. The closed-loop dynamics matrix is given by

$$A_{cl} = A - BK = \begin{bmatrix} 0 & 1 \\ -1/q & -\sqrt{2/q} \end{bmatrix}.$$

The characteristic polynomial of this matrix is

$$\lambda^2 + \sqrt{\frac{2}{q}}\lambda + \frac{1}{q}.$$

Comparing this to $\lambda^2 + 2\zeta\omega_0\lambda + \omega_0^2$, we see that

$$\omega_0 = \sqrt{\frac{1}{q}}, \quad \zeta = \frac{1}{\sqrt{2}}.$$

Thus the optimal controller gives a closed loop system with damping ratio $\zeta = 0.707$, giving a good tradeoff between rise time and overshoot (see ÅM08). ∇

2.5 Choosing LQR weights

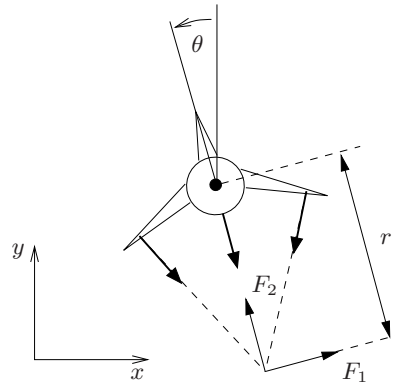
One of the key questions in LQR design is how to choose the weights Q_x and Q_u . To choose specific values for the cost function weights Q_x and Q_u , we must use our knowledge of the system we are trying to control. A particularly simple choice is to use diagonal weights

$$Q_x = \begin{bmatrix} q_1 & & 0 \\ & \ddots & \\ 0 & & q_n \end{bmatrix}, \quad Q_u = \begin{bmatrix} \rho_1 & & 0 \\ & \ddots & \\ 0 & & \rho_n \end{bmatrix}.$$

For this choice of Q_x and Q_u , the individual diagonal elements describe how much each state and input (squared) should contribute to the overall cost. Hence, we can take states that should remain small and attach higher weight



(a) Harrier “jump jet”



(b) Simplified model

Figure 2.3: Vectored thrust aircraft. The Harrier AV-8B military aircraft (a) redirects its engine thrust downward so that it can “hover” above the ground. Some air from the engine is diverted to the wing tips to be used for maneuvering. As shown in (b), the net thrust on the aircraft can be decomposed into a horizontal force F_1 and a vertical force F_2 acting at a distance r from the center of mass.

values to them. Similarly, we can penalize an input versus the states and other inputs through choice of the corresponding input weight ρ_j .

Choosing the individual weights for the (diagonal) elements of the Q_x and Q_u matrix can be done by deciding on a weighting of the errors from the individual terms. Bryson and Ho [BH75] have suggested the following method for choosing the matrices Q_x and Q_u in equation (2.4): (1) choose q_i and ρ_j as the inverse of the square of the maximum value for the corresponding x_i or u_j ; (2) modify the elements to obtain a compromise among response time, damping and control effort. This second step can be performed by trial and error.

It is also possible to choose the weights such that only a given subset of variable are considered in the cost function. Let $z = Hx$ be the output we want to keep small and verify that (A, H) is observable. Then we can use a cost function of the form

$$Q_x = H^T H \quad Q_u = \rho I.$$

The constant ρ allows us to trade off $\|z\|^2$ versus $\rho\|u\|^2$.

We illustrate the various choices through an example application.

Example 2.5 Thrust vectored aircraft

Consider the thrust vectored aircraft example introduced in ÅM08, Example 2.9. The system is shown in Figure 2.3, reproduced from ÅM08. The linear quadratic regulator problem was illustrated in Example 6.8, where the weights were chosen as $Q_x = I$ and $Q_u = \rho I$. Figure 2.4 reproduces the step response for this case.

A more physically motivated weighted can be computing by specifying the comparable errors in each of the states and adjusting the weights ac-

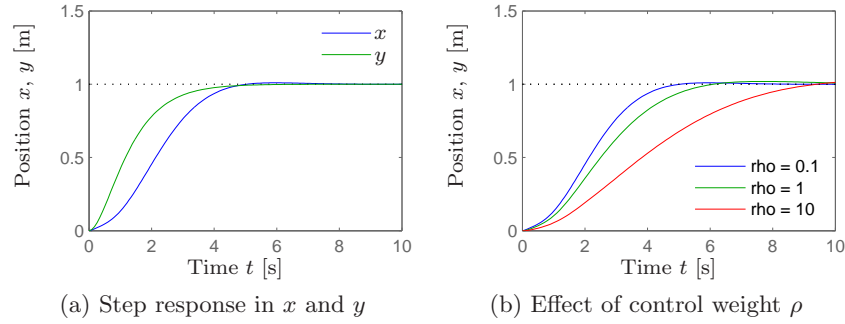


Figure 2.4: Step response for a vectored thrust aircraft. The plot in (a) shows the x and y positions of the aircraft when it is commanded to move 1 m in each direction. In (b) the x motion is shown for control weights $\rho = 1, 10^2, 10^4$. A higher weight of the input term in the cost function causes a more sluggish response.

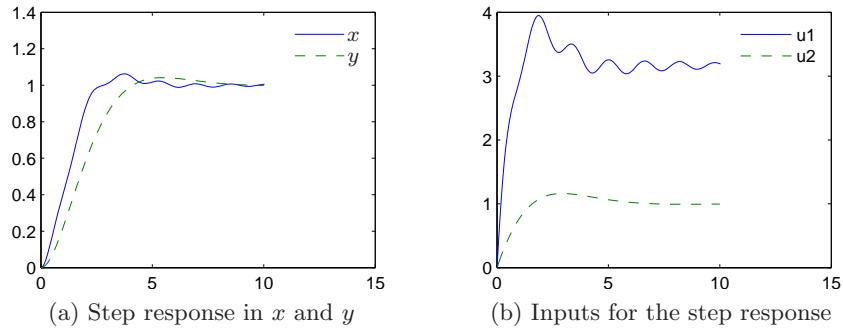


Figure 2.5: Step response for a vector thrust aircraft using physically motivated LQR weights (a). The rise time for x is much faster than in Figure 2.4a, but there is a small oscillation and the inputs required are quite large (b).

cordingly. Suppose, for example that we consider a 1 cm error in x , a 10 cm error in y and a 5° error in θ to be equivalently bad. In addition, we wish to penalize the forces in the sideways direction since these results in a loss in efficiency. This can be accounted for in the LQR weights by choosing

$$Q_x = \begin{bmatrix} 100 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2\pi/9 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad Q_u = 0.1 \times \begin{bmatrix} 1 & 0 \\ 0 & 10 \end{bmatrix}.$$

The results of this choice of weights are shown in Figure 2.5. ▽

2.6 Advanced Topics

In this section we briefly touch on some related topics in optimal control, with reference to more detailed treatments where appropriate.

General quadratic cost functions.

$$\dot{x} = Ax + Bu \quad J = \int_0^\infty \overbrace{\left(x^T Q_x x + u^T Q_u u + x^T S u \right)}^{L(x,u)} dt,$$

where the S term is almost always left out.

Singular extremals. The necessary conditions in the maximum principle enforce the constraints through the of the Lagrange multipliers $\lambda(t)$. In some instances, we can get an extremal curve that has one or more of the λ 's identically equal to zero. This corresponds to a situation in which the constraint is satisfied strictly through the minimization of the cost function and does not need to be explicitly enforced. We illustrate this case through an example.

Example 2.6 Nonholonomic integrator

Consider the minimum time optimization problem for the nonholonomic integrator introduced in Example 1.2 with input constraints $|u_i| \leq 1$. The Hamiltonian for the system is given by

$$H = 1 + \lambda_1 u_1 + \lambda_2 u_2 + \lambda_3 x_2 u_1$$

and the resulting equations for the Lagrange multipliers are

$$\dot{\lambda}_1 = 0, \quad \dot{\lambda}_2 = \lambda_3 x_2, \quad \dot{\lambda}_3 = 0. \quad (2.6)$$

It follows from these equations that λ_1 and λ_3 are constant. To find the input u corresponding to the extremal curves, we see from the Hamiltonian that

$$u_1 = -\text{sgn}(\lambda_1 + \lambda_3 x_2 u_1), \quad u_2 = -\text{sgn} \lambda_2.$$

These equations are well-defined as long as the arguments of $\text{sgn}()$ are non-zero and we get switching of the inputs when the arguments pass through 0.

An example of an abnormal extremal is the optimal trajectory between $x_0 = (0, 0, 0)$ to $x_f = (\rho, 0, 0)$ where $\rho > 0$. The minimum time trajectory is clearly given by moving on a straight line with $u_1 = 1$ and $u_2 = 0$. This extremal satisfies the necessary conditions but with $\lambda_2 \equiv 0$, so that the “constraint” that $\dot{x}_2 = u_2$ is not strictly enforced through the Lagrange multipliers. ∇

Abnormal extremals.

$$\dot{x}_1 = u_1, \quad \dot{x}_2 = u_2, \quad \dot{x}_3 = x_2 u_1.$$

Minimum length curve:

$$J = \int_0^T u_1^2 + u_2^2 dt$$

Show that x_0 to x_f gives multiple solutions.

2.7 Further Reading

There are a number of excellent books on optimal control. One of the first (and best) is the book by Pontryagin et al. [PBG62]. During the 1960s and 1970s a number of additional books were written that provided many examples and served as standard textbooks in optimal control classes. Athans and Falb [AF06] and Bryson and Ho [BH75] are two such texts. A very elegant treatment of optimal control from the point of view of optimization over general linear spaces is given by Luenberger [Lue97]. Finally, a modern engineering textbook that contains a very compact and concise derivation of the key results in optimal control is the book by Lewis and Syrmos [LS95].

Exercises

2.1 (a) Let G_1, G_2, \dots, G_k be a set of row vectors on a \mathbb{R}^n . Let F be another row vector on \mathbb{R}^n such that for every $x \in \mathbb{R}^n$ satisfying $G_i x = 0$, $i = 1, \dots, k$, we have $Fx = 0$. Show that there are constants $\lambda_1, \lambda_2, \dots, \lambda_k$ such that

$$F = \sum_{i=1}^k \lambda_i G_i.$$

(b) Let $x^* \in \mathbb{R}^n$ be an the extremal point (maximum or minimum) of a function f subject to the constraints $g_i(x) = 0$, $i = 1, \dots, k$. Assuming that the gradients $\partial g_i(x^*)/\partial x$ are linearly independent, show that there are k scalars λ_i , $i = 1, \dots, k$ such that the function

$$\tilde{f}(x) = f(x) + \sum_{i=1}^k \lambda_i g_i(x)$$

has an extremal point at x^* .

2.2 Consider the following control system

$$\begin{aligned} \dot{q} &= u \\ \dot{Y} &= qu^T - uq^T \end{aligned}$$

where $u \in \mathbb{R}^m$ and $Y \in \mathbb{R}^{m \times m}$ is a skew symmetric matrix, $Y^T = -Y$.

(a) For the fixed end point problem, derive the form of the optimal controller minimizing the following integral

$$\frac{1}{2} \int_0^1 u^T u dt.$$

(b) For the boundary conditions $q(0) = q(1) = 0$, $Y(0) = 0$ and

$$Y(1) = \begin{bmatrix} 0 & -y_3 & y_2 \\ y_3 & 0 & -y_1 \\ -y_2 & y_1 & 0 \end{bmatrix}$$

for some $y \in \mathbb{R}^3$, give an explicit formula for the optimal inputs u .

(c) (Optional) Find the input u to steer the system from $(0, 0)$ to $(0, \tilde{Y}) \in \mathbb{R}^m \times \mathbb{R}^{m \times m}$ where $\tilde{Y}^T = -\tilde{Y}$.

(Hint: if you get stuck, there is a paper by Brockett on this problem.)

2.3 In this problem, you will use the maximum principle to show that the shortest path between two points is a straight line. We model the problem by constructing a control system

$$\dot{x} = u,$$

where $x \in \mathbb{R}^2$ is the position in the plane and $u \in \mathbb{R}^2$ is the velocity vector along the curve. Suppose we wish to find a curve of minimal length connecting $x(0) = x_0$ and $x(1) = x_f$. To minimize the length, we minimize the integral of the velocity along the curve,

$$J = \int_0^1 \|\dot{x}\| dt = \int_0^1 \sqrt{\dot{x}^T \dot{x}} dt,$$

subject to the initial and final state constraints. Use the maximum principle to show that the minimal length path is indeed a straight line at maximum velocity. (Hint: try minimizing using the integral cost $\dot{x}^T \dot{x}$ first and then show this also optimizes the optimal control problem with integral cost $\|\dot{x}\|$.)

2.4 Consider the optimal control problem for the system

$$\dot{x} = -ax + bu,$$

where $x \in \mathbb{R}$ is a scalar state, $u \in \mathbb{R}$ is the input, the initial state $x(t_0)$ is given, and $a, b \in \mathbb{R}$ are positive constants. (Note that this system is not quite the same as the one in Example 2.2.) The cost function is given by

$$J = \frac{1}{2} \int_{t_0}^{t_f} u^2(t) dt + \frac{1}{2} cx^2(t_f),$$

where the terminal time t_f is given and c is a constant.

- (a) Solve explicitly for the optimal control $u^*(t)$ and the corresponding state $x^*(t)$ in terms of $t_0, t_f, x(t_0)$ and t and describe what happens to the terminal state $x^*(t_f)$ as $c \rightarrow \infty$.
- (b) Show that the system is differentially flat with appropriate choice of output(s) and compute the state and input as a function of the flat output(s).
- (c) Using the polynomial basis $\{t^k, k = 0, \dots, M - 1\}$ with an appropriate choice of M , solve for the (non-optimal) trajectory between $x(t_0)$ and $x(t_f)$. Your answer should specify the explicit input $u_d(t)$ and state $x_d(t)$ in terms of $t_0, t_f, x(t_0), x(t_f)$ and t .
- (d) Let $a = 1$ and $c = 1$. Use your solution to the optimal control problem and the flatness-based trajectory generation to find a trajectory between $x(0) = 0$ and $x(1) = 1$. Plot the state and input trajectories for each solution and compare the costs of the two approaches.
- (e) (Optional) Suppose that we choose more than the minimal number of basis functions for the differentially flat output. Show how to use the additional degrees of freedom to minimize the cost of the flat trajectory and demonstrate that you can obtain a cost that is closer to the optimal.

2.5 Repeat Exercise 2.4 using the system

$$\dot{x} = -ax^3 + bu.$$

For part (a) you need only write the conditions for the optimal cost.

2.6 Consider the problem of moving a two-wheeled mobile robot (e.g., a Segway) from one position and orientation to another. The dynamics for the system is given by the nonlinear differential equation

$$\begin{aligned}\dot{x} &= \cos \theta v \\ \dot{y} &= \sin \theta v \\ \dot{\theta} &= \omega\end{aligned}$$

where (x, y) is the position of the rear wheels, θ is the angle of the robot with respect to the x axis, v is the forward velocity of the robot and ω is spinning rate. We wish to choose an input (v, ω) that minimizes the time that it takes to move between two configurations (x_0, y_0, θ_0) and (x_f, y_f, θ_f) , subject to input constraints $|v| \leq L$ and $|\omega| \leq M$.

Use the maximum principle to show that any optimal trajectory consists of segments in which the robot is traveling at maximum velocity in either the forward or reverse direction, and going either straight, hard left ($\omega = -M$) or hard right ($\omega = +M$).

Note: one of the cases is a bit tricky and cannot be completely proven with the tools we have learned so far. However, you should be able to show the other cases and verify that the tricky case is possible.

2.7 Consider a linear system with input u and output y and suppose we wish to minimize the quadratic cost function

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

Show that if the corresponding linear system is observable, then the closed loop system obtained by using the optimal feedback $u = -Kx$ is guaranteed to be stable.

2.8 Consider the control system transfer function

$$H(s) = \frac{s + b}{s(s + a)}, \quad a, b > 0$$

with state space representation

$$\begin{aligned} \dot{x} &= \begin{bmatrix} 0 & 1 \\ 0 & -a \end{bmatrix} x + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u, \\ y &= \begin{bmatrix} b & 1 \end{bmatrix} x \end{aligned}$$

and performance criterion

$$V = \int_0^{\infty} (x_1^2 + u^2) dt.$$

(a) Let

$$P = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix},$$

with $p_{12} = p_{21}$ and $P > 0$ (positive definite). Write the steady state Riccati equation as a system of four explicit equations in terms of the elements of P and the constants a and b .

(b) Find the gains for the optimal controller assuming the full state is available for feedback.

(c) Find the closed loop natural frequency and damping ratio.

2.9 Consider the optimal control problem for the system

$$\dot{x} = ax + bu \quad J = \frac{1}{2} \int_{t_0}^{t_f} u^2(t) dt + \frac{1}{2} cx^2(t_f),$$

where $x \in \mathbb{R}$ is a scalar state, $u \in \mathbb{R}$ is the input, the initial state $x(t_0)$ is given, and $a, b \in \mathbb{R}$ are positive constants. We take the terminal time t_f as given and let $c > 0$ be a constant that balances the final value of the state with the input required to get to that position. The optimal trajectory is derived in Example 2.2.

Now consider the infinite horizon cost

$$J = \frac{1}{2} \int_{t_0}^{\infty} u^2(t) dt$$

with $x(t)$ at $t = \infty$ constrained to be zero.

(a) Solve for $u^*(t) = -bPx^*(t)$ where P is the positive solution corresponding to the algebraic Riccati equation. Note that this gives an explicit feedback law ($u = -bPx$).

(b) Plot the state solution of the finite time optimal controller for the following parameter values

$$\begin{aligned} a &= 2, & b &= 0.5, & x(t_0) &= 4, \\ c &= 0.1, 10, & t_f &= 0.5, 1, 10. \end{aligned}$$

(This should give you a total of 6 curves.) Compare these to the infinite time optimal control solution. Which finite time solution is closest to the infinite time solution? Why?

2.10 Consider the lateral control problem for an autonomous ground vehicle from Example 1.1. We assume that we are given a reference trajectory $r = (x_d, y_d)$ corresponding to the desired trajectory of the vehicle. For simplicity, we will assume that we wish to follow a straight line in the x direction at a constant velocity $v_d > 0$ and hence we focus on the y and θ dynamics:

$$\begin{aligned} \dot{y} &= \sin \theta v_d, \\ \dot{\theta} &= \frac{1}{l} \tan \phi v_d. \end{aligned}$$

We let $v_d = 10$ m/s and $l = 2$ m.

(a) Design an LQR controller that stabilizes the position y to $y_d = 0$. Plot the step and frequency response for your controller and determine the overshoot, rise time, bandwidth and phase margin for your design. (Hint: for the frequency domain specifications, break the loop just before the process dynamics and use the resulting SISO loop transfer function.)

(b) Suppose now that $y_d(t)$ is not identically zero, but is instead given by $y_d(t) = r(t)$. Modify your control law so that you track $r(t)$ and demonstrate the performance of your controller on a “slalom course” given by a sinusoidal trajectory with magnitude 1 meter and frequency 1 Hz.