
Optimization-Based Control

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

DRAFT v1.6c, 15 February 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.

Preface

These notes serve as a supplement to *Feedback Systems* by Åström and Murray and expand on some of the topics introduced there. Our focus is on the use of optimization-based methods for control, including optimal control theory, receding horizon control and Kalman filtering. Each chapter is intended to be a standalone reference for advanced topics that are introduced in *Feedback Systems*.

Contents

Chapter 1. Trajectory Generation and Tracking	1
1.1 Two Degree of Freedom Design	1
1.2 Trajectory Tracking and Gain Scheduling	3
1.3 Trajectory Generation and Differential Flatness	7
1.4 Further Reading	13
Chapter 2. Optimal Control	17
2.1 Review: Optimization	17
2.2 Optimal Control of Systems	20
2.3 Examples	23
2.4 Linear Quadratic Regulators	25
2.5 Choosing LQR weights	28
2.6 Advanced Topics	30
2.7 Further Reading	31
Chapter 3. Receding Horizon Control	38
3.1 Optimization-Based Control	38
3.2 Receding Horizon Control with CLF Terminal Cost	45
3.3 Receding Horizon Control Using Differential Flatness	48
3.4 Implementation on the Caltech Ducted Fan	51
3.5 Further Reading	59
Chapter 4. Stochastic Systems	62
4.1 Review of Random Variables	62
4.2 Introduction to Random Processes	66
4.3 Continuous-Time, Vector-Valued Random Processes	68
4.4 Linear Stochastic Systems	71
4.5 Random Processes in the Frequency Domain	74
4.6 Further Reading	77
Chapter 5. Kalman Filtering	79
5.1 Linear Quadratic Estimators	79
5.2 Extended Kalman Filters	81
5.3 LQG Control	82

CONTENTS	iii
5.4 Further Reading	83
Chapter 6. Sensor Fusion	84
6.1 Kalman Filters in Discrete Time	84
6.2 Predictor-Corrector Form	86
6.3 Sensor Fusion	87
6.4 Information Filters	87
6.5 Further Reading	87
Bibliography	88

Chapter 1

Trajectory Generation and Tracking

This set of notes expands on Section 7.5 of *Feedback Systems* by Åström and Murray (ÅM08), which introduces the use of feedforward compensation in control system design. We begin with a review of the two degree of freedom design approach and then focus on the problem of generating feasible trajectories for a (nonlinear) control system. We make use of the concept of differential flatness as a tool for generating feasible trajectories.

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. Although this material supplements concepts introduced in Chapter 7 of ÅM08, no knowledge of observers is required.

1.1 Two Degree of Freedom Design

A large class of control problems consist of planning and following a trajectory in the presence of noise and uncertainty. Examples include autonomous vehicles maneuvering in city streets, mobile robots performing tasks on factor floors (or other planets), manufacturing systems that regulate the flow of parts and materials through a plant or factory, and supply chain management systems that balance orders and inventories across an enterprise. All of these systems are highly nonlinear and demand accurate performance.

To control such systems, we make use of the notion of *two degree of freedom* controller design. This is a standard technique in linear control theory that separates a controller into a feedforward compensator and a feedback compensator. The feedforward compensator generates the nominal input required to track a given reference trajectory. The feedback compensator corrects for errors between the desired and actual trajectories. This is shown schematically in Figure 1.1.

In a nonlinear setting, two degree of freedom controller design decouples the trajectory generation and asymptotic tracking problems. Given a desired output trajectory, we first construct a state space trajectory x_d and a nominal input u_d that satisfy the equations of motion. The error system can then be written as a time-varying control system in terms of the error, $e = x - x_d$. Under the assumption that that tracking error remains




Figure 1.1: Two degree of freedom controller design for a process P with uncertainty Δ . The controller consists of a trajectory generator and feedback controller. The trajectory generation subsystem computes a feedforward command u_d along with the desired state x_d . The state feedback controller uses the measured (or estimated) state and desired state to compute a corrective input u_{fb} . Uncertainty is represented by the block Δ , representing unmodeled dynamics, as well as disturbances and noise.

small, we can linearize this time-varying system about $e = 0$ and stabilize the $e = 0$ state. (Note: in ÅM08 the notation u_{ff} was used for the desired (feedforward) input. We use u_d here to match the desired state x_d .)

More formally, we assume that our process dynamics can be described by a nonlinear differential equation of the form

$$\begin{aligned} \dot{x} &= f(x, u) & x &\in \mathbb{R}^n, u \in \mathbb{R}^p, \\ y &= h(x, u) & y &\in \mathbb{R}^q, \end{aligned} \quad (1.1)$$

where x is the system state, u is a vector of inputs and f is a smooth function describing the dynamics of the process. The smooth function h describes the output y that we wish to control. We are particularly interested in the class of control problems in which we wish to track a time-varying reference trajectory $r(t)$, called the *trajectory tracking* problem. In particular, we wish to find a control law $u = \alpha(x, r(\cdot))$ such that

$$\lim_{t \rightarrow \infty} (y(t) - r(t)) = 0.$$

We use the notation $r(\cdot)$ to indicate that the control law can depend not only on the reference signal $r(t)$ but also derivatives of the reference signal.

A *feasible trajectory* for the system (1.1) is a pair $(x_d(t), u_d(t))$ that satisfies the differential equation and generates the desired trajectory:

$$\dot{x}_d = f(x_d, u_d) \quad r(t) = h(x_d, u_d).$$

The problem of finding a feasible trajectory for a system is called the *trajectory generation* problem, with x_d representing the desired state for the

(nominal) system and u_d representing the desired input or the *feedforward* control. If we can find a feasible trajectory for the system, we can search for controllers of the form $u = \alpha(x, x_d, u_d)$ that track the desired reference trajectory.

In many applications, it is possible to attach a cost function to trajectories that describe how well they balance trajectory tracking with other factors, such as the magnitude of the inputs required. In such applications, it is natural to ask that we find the optimal controller. We can again use the two degree of freedom paradigm with an optimal control computation for generating the feasible trajectory. This subject is examined in more detail in Chapter 2. In addition, we can take the extra step of updating the generated trajectory based on the current state of the system. This additional feedback path is denoted by a dashed line in Figure 1.1 and allows the use of so-called *receding-horizon-control* techniques: a (optimal) feasible trajectory is computed from the current position to the desired position over a finite time T horizon, used for a short period of time $\delta < T$, and then recomputed based on the new position. Receding horizon control is described in more detail in Chapter 3.

A key advantage of optimization-based approaches is that they allow the potential for customization of the controller based on changes in *mission*, *condition* and *environment*. Because the controller is solving the optimization problem online, updates can be made to the cost function, to change the desired operation of the system; to the model, to reflect changes in parameter values or damage to sensors and actuators; and to the constraints, to reflect new regions of the state space that must be avoided due to external influences. Thus, many of the challenges of designing controllers that are robust to a large set of possible uncertainties become embedded in the online optimization.

1.2 Trajectory Tracking and Gain Scheduling

We begin by considering the problem of tracking a feasible trajectory. Assume that a trajectory generator is able to generate a trajectory (x_d, u_d) that satisfies the dynamics (1.1) and satisfies $r(t) = h(x_d(t), u_d(t))$. To design the controller, we construct the *error system*. Let $e = x - x_d$ and $v = u - u_d$ and compute the dynamics for the error:

$$\begin{aligned} \dot{e} &= \dot{x} - \dot{x}_d = f(x, u) - f(x_d, u_d) \\ &= f(e + x_d, v + u_d) - f(x_d) =: F(e, v, x_d(t), u_d(t)). \end{aligned}$$

In general, this system is time-varying.

For trajectory tracking, we can assume that e is small (if our controller

is doing a good job), and so we can linearize around $e = 0$:

$$\frac{de}{dt} \approx A(t)e + B(t)v, \quad A(t) = \left. \frac{\partial F}{\partial e} \right|_{(x_d(t), u_d(t))}, \quad B(t) = \left. \frac{\partial F}{\partial v} \right|_{(x_d(t), u_d(t))}.$$

It is often the case that $A(t)$ and $B(t)$ depend only on x_d , in which case it is convenient to write $A(t) = A(x_d)$ and $B(t) = B(x_d)$.

We start by reviewing the case where $A(t)$ and $B(t)$ are constant, in which case our error dynamics become

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

This occurs, for example, if the original nonlinear system is linear. We can then search for a control system of the form

$$v = -Ke + k_r r.$$

In the case where r is constant, we can apply the results of Chapter 6 of ÅM08 and solve the problem by finding a gain matrix K that gives the desired close loop dynamics (e.g., by eigenvalue assignment) and choosing k_r to give the desired output value at equilibrium. The equilibrium point is given by

$$x_e = -(A - BK)^{-1} B k_r r \implies y_e = -C(A - BK)^{-1} B k_r r$$

and if we wish the output to be $y = r$ it follows that

$$k_r = -1/(C(A - BK)^{-1}B).$$

It can be shown that this formulation is equivalent to a two degree of freedom design where x_d and u_d are chosen to give the desired reference output (Exercise 1.1).

Returning to the full nonlinear system, assume now that x_d and u_d are either constant or slowly varying (with respect to the performance criterion). This allows us to consider just the (constant) linearized system given by $(A(x_d), B(x_d))$. If we design a state feedback controller $K(x_d)$ for each x_d , then we can regulate the system using the feedback

$$v = K(x_d)e.$$

Substituting back the definitions of e and v , our controller becomes

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

Note that the controller $u = \alpha(x, x_d, u_d)$ depends on (x_d, u_d) , which themselves depend on the desired reference trajectory. This form of controller is called a *gain scheduled* linear controller with *feedforward* u_d .

More generally, the term gain scheduling is used to describe any controller that depends on a set of measured parameters in the system. So, for example, we might write

$$u = -K(x, \mu) \cdot (x - x_d) + u_d,$$




Figure 1.2: Gain scheduling. A general gain scheduling design involves finding a gain K at each desired operating point. This can be thought of as a gain surface, as shown on the left (for the case of a scalar gain). An approximation to this gain can be obtained by computing the gains at a fixed number of operating points and then interpolated between those gains. This gives an approximation of the continuous gain surface, as shown on the right.

where $K(x, \mu)$ depends on the *current* system state (or some portion of it) and an external parameter μ . The dependence on the current state x (as opposed to the desired state x_d) allows us to modify the closed loop dynamics differently depending on our location in the state space. This is particularly useful when the dynamics of the process vary depending on some subset of the states (such as the altitude for an aircraft or the internal temperature for a chemical reaction). The dependence on μ can be used to capture the dependence on the reference trajectory, or they can reflect changes in the environment or performance specifications that are not modeled in the state of the controller.

One limitation of gain scheduling as we have described it is that a separate set of gains must be designed for each operating condition x_d . In practice, gain scheduled controllers are often implemented by designing controllers at a fixed number of operating points and then interpolating the gains between these points, as illustrated in Figure 1.2. Suppose that we have a set of operating points $x_{d,j}$, $j = 1, \dots, N$. Then we can write our controller as

$$u = u_d - K(x)e \quad K(x) = \sum (\alpha_j(x)K_j),$$

where K_j is a set of gains designed around the operating point $x_{d,j}$ and $\alpha_j(x)$ is a weighting factor. For example, we might choose the weights $\alpha_j(x)$ such that we take the gains corresponding to the nearest two operating points and weight them according to the Euclidean distance of the current state from that operating point; if the distance is small then we use a weight very near to 1 and if the distance is far then we use a weight very near to 0.

While the intuition behind gain scheduled controllers is fairly clear, some caution is required in using them. In particular, a gain scheduled controller is not guaranteed to be stable even if $K(x, \mu)$ locally stabilizes the system



Figure 1.3: Vehicle steering using gain scheduling.

around a given equilibrium point. Gain scheduling can be proven to work in the case when the gain varies sufficiently slow (Exercise 1.3).

Example 1.1 Steering control with velocity scheduling

Consider the problem of controlling the motion of a automobile so that it follows a given trajectory on the ground, as shown in Figure 1.3. We use the model derived in ÅM08, choosing the reference point to be the center of the rear wheels. This gives dynamics of the form

$$\begin{aligned} \dot{x} &= \cos \theta v \\ \dot{y} &= \sin \theta v \\ \dot{\theta} &= \frac{v}{l} \tan \phi, \end{aligned} \tag{1.2}$$

where (x, y, θ) is the position and orientation of the vehicle, v is the velocity and ϕ is the steering angle, both considered to be inputs, and l is the wheelbase.

A simple feasible trajectory for the system is to follow a straight line in the x direction at lateral position y_r and fixed velocity v_r . This corresponds to a desired state $x_d = (v_r t, y_r, 0)$ and nominal input $u_d = (v_r, 0)$. Note that (x_d, u_d) is not an equilibrium point for the system, but it does satisfy the equations of motion.

Linearizing the system about the desired trajectory, we obtain

$$\begin{aligned} A_d &= \left. \frac{\partial f}{\partial x} \right|_{(x_d, u_d)} = \begin{bmatrix} 0 & 0 & -\sin \theta \\ 0 & 0 & \cos \theta \\ 0 & 0 & 0 \end{bmatrix} \bigg|_{(x_d, u_d)} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \\ B_d &= \left. \frac{\partial f}{\partial u} \right|_{(x_d, u_d)} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & v_r/l \end{bmatrix}. \end{aligned}$$

We form the error dynamics by setting $e = x - x_d$ and $w = u - u_d$:

$$\dot{e}_x = w_1, \quad \dot{e}_y = e_\theta, \quad \dot{e}_\theta = \frac{v_r}{l} w_2.$$

We see that the first state is decoupled from the second two states and hence we can design a controller by treating these two subsystems separately.

Suppose that we wish to place the closed loop eigenvalues of the longitudinal dynamics (e_x) at λ_1 and place the closed loop eigenvalues of the lateral dynamics (e_y, e_θ) at the roots of the polynomial equation $s^2 + a_1s + a_2 = 0$. This can be accomplished by setting

$$w_1 = -\lambda_1 e_x$$

$$w_2 = \frac{l}{v_r} (a_1 e_y + a_2 e_\theta).$$

Note that the gains depend on the velocity v_r (or equivalently on the nominal input u_d), giving us a gain scheduled controller.

In the original inputs and state coordinates, the controller has the form

$$\begin{bmatrix} v \\ \phi \end{bmatrix} = - \underbrace{\begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \frac{a_1 l}{v_r} & \frac{a_2 l}{v_r} \end{bmatrix}}_{K_d} \underbrace{\begin{bmatrix} x - v_r t \\ y - y_r \\ \theta \end{bmatrix}}_e + \underbrace{\begin{bmatrix} v_r \\ 0 \end{bmatrix}}_{u_d}.$$

The form of the controller shows that at low speeds the gains in the steering angle will be high, meaning that we must turn the wheel harder to achieve the same effect. As the speed increases, the gains become smaller. This matches the usual experience that at high speed a very small amount of actuation is required to control the lateral position of a car. Note that the gains go to infinity when the vehicle is stopped ($v_r = 0$), corresponding to the fact that the system is not reachable at this point. ∇

1.3 Trajectory Generation and Differential Flatness

We now return to the problem of generating a trajectory for a nonlinear system. Consider first the case of finding a trajectory $x_d(t)$ that steers the system from an initial condition x_0 to a final condition x_f . We seek a feasible solution $(x_d(t), u_d(t))$ that satisfies the dynamics of the process:

$$\dot{x}_d = f(x_d, u_d), \quad x_d(0) = x_0, \quad x_d(T) = x_f. \quad (1.3)$$

In addition, we may wish to satisfy additional constraints on the dynamics:

- Input saturation: $|u(t)| < M$;
- State constraints: $g(x) \leq 0$
- Tracking: $h(x) = r(t)$
- Optimization:

$$\min \int_0^T L(x, u) dt + V(x(T), u(T))$$

Formally, this problem corresponds to a two-point boundary value problem and can be quite difficult to solve in general.




Figure 1.4: Simple model for an automobile. We wish to find a trajectory from an initial state to a final state that satisfies the dynamics of the system and constraints on the curvature (imposed by the limited travel of the front wheels).

As an example of the type of problem we would like to study, consider the problem of steering a car from an initial condition to a final condition, as show in Figure 1.4. To solve this problem, we must find a solution to the differential equations (1.2) that satisfies the endpoint conditions. Given the nonlinear nature of the dynamics, it seems unlikely that one could find explicit solutions that satisfy the dynamics except in very special cases (such as driving in a straight line).

A closer inspection of this system shows that it is possible to understand the trajectories of the system by exploiting the particular structure of the dynamics. Suppose that we are given a trajectory for the rear wheels of the system, $x(t)$ and $y(t)$. From equation (1.2), we see that we can use this solution to solve for the angle of the car by writing

$$\frac{\dot{y}}{\dot{x}} = \frac{\sin \theta}{\cos \theta} \quad \Longrightarrow \quad \theta = \tan^{-1}(\dot{y}/\dot{x}).$$

Furthermore, given θ we can solve for the velocity using the equation

$$\dot{x} = v \cos \theta \quad \Longrightarrow \quad v = \dot{x} / \cos \theta,$$

assuming $\cos \theta \neq 0$ (if it is, use $v = \dot{y} / \sin \theta$). And given θ , we can solve for ϕ using the relationship

$$\dot{\theta} = \frac{v}{l} \tan \phi \quad \Longrightarrow \quad \phi = \tan^{-1}\left(\frac{l\dot{\theta}}{v}\right).$$

Hence all of the state variables and the inputs can be determined by the trajectory of the rear wheels and its derivatives. This property of a system is known as *differential-flatness*.

Definition 1.1 (Differential flatness). A nonlinear system (1.1) is *differentially flat* if there exists a function α such that

$$z = \alpha(x, u, \dot{u} \dots, u^{(p)})$$

and we can write the solutions of the nonlinear system as functions of z and an finite number of derivatives

$$\begin{aligned} x &= \beta(z, \dot{z}, \dots, z^{(q)}), \\ u &= \gamma(z, \dot{z}, \dots, z^{(q)}). \end{aligned}$$

For a differentially flat system, all of the feasible trajectories for the system can be written as functions of a flat output $z(\cdot)$ and its derivatives. The number of flat outputs is always equal to the number of system inputs. The kinematic car is differentially flat with the position of the rear wheels as the flat output. Differentially flat systems were originally studied by Fliess et al. [FLMR92].

Differentially flat systems are useful in situations where explicit trajectory generation is required. Since the behavior of flat system is determined by the flat outputs, we can plan trajectories in output space, and then map these to appropriate inputs. Suppose we wish to generate a feasible trajectory for the the nonlinear system

$$\dot{x} = f(x, u), \quad x(0) = x_0, x(T) = x_f.$$

If the system is differentially flat then

$$\begin{aligned} x(0) &= \beta(z(0), \dot{z}(0), \dots, z^{(q)}(0)) = x_0 \\ x(T) &= \gamma(z(T), \dot{z}(T), \dots, z^{(q)}(T)) = x_f \end{aligned} \tag{1.4}$$

and any trajectory for z that satisfies these boundary conditions will be a feasible trajectory for the system.

In particular, given initial and final conditions on z and its derivatives that satisfy equation (1.4), any curve $z(\cdot)$ satisfying those conditions will correspond to a feasible trajectory of the system. We can parameterize the flat output trajectory using a set of smooth basis functions $\psi^i(t)$:

$$z(t) = \sum_{i=1}^N \alpha_i \psi_i(t), \quad \alpha_i \in \mathbb{R}.$$

We seek a set of coefficients α_i , $i = 1, \dots, N$ such that $z(t)$ satisfies the boundary conditions (1.4). The derivatives of the flat output can be computed in terms of the derivatives of the basis functions:

$$\begin{aligned} \dot{z}(t) &= \sum_{i=1}^N \alpha_i \dot{\psi}_i(t) \\ &\vdots \\ \dot{z}^{(q)}(t) &= \sum_{i=1}^N \alpha_i \dot{\psi}_i^{(q)}(t). \end{aligned}$$

We can thus write the conditions on the flat outputs and their derivatives as

$$\begin{bmatrix} \psi_1(0) & \psi_2(0) & \dots & \psi_N(0) \\ \dot{\psi}_1(0) & \dot{\psi}_2(0) & \dots & \dot{\psi}_N(0) \\ \vdots & \vdots & & \vdots \\ \psi_1^{(q)}(0) & \psi_2^{(q)}(0) & \dots & \psi_N^{(q)}(0) \\ \psi_1(T) & \psi_2(T) & \dots & \psi_N(T) \\ \dot{\psi}_1(T) & \dot{\psi}_2(T) & \dots & \dot{\psi}_N(T) \\ \vdots & \vdots & & \vdots \\ \psi_1^{(q)}(T) & \psi_2^{(q)}(T) & \dots & \psi_N^{(q)}(T) \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_N \end{bmatrix} = \begin{bmatrix} z(0) \\ \dot{z}(0) \\ \vdots \\ z^{(q)}(0) \\ z(T) \\ \dot{z}(T) \\ \vdots \\ z^{(q)}(T) \end{bmatrix}$$

This equation is a *linear* equation of the form $M\alpha = \bar{z}$. Assuming that M has a sufficient number of columns and that it is full column rank, we can solve for a (possibly non-unique) α that solves the trajectory generation problem.

Example 1.2 Nonholonomic integrator

A simple nonlinear system called a *nonholonomic integrator* [Bro81] is given by the differential equations

$$\begin{aligned} \dot{x}_1 &= u_1 \\ \dot{x}_2 &= u_2 \\ \dot{x}_3 &= x_2 u_1 \end{aligned}$$

This system is differentially flat with flat output $z = (x_1, x_3)$. The relationship between the flat variables and the states is given by

$$\begin{aligned} x_1 &= z_1 \\ x_2 &= \dot{x}_3 / \dot{x}_1 = \dot{z}_2 / \dot{z}_1 \\ x_3 &= z_2. \end{aligned}$$

Using simple polynomials as our basis functions,

$$\begin{aligned} \psi_{1,1}(t) &= 1 & \psi_{1,2}(t) &= t\psi_{1,3}(t) &= t^2\psi_{1,4}(t) &= t^3 \\ \psi_{2,1}(t) &= 1 & \psi_{2,2}(t) &= t\psi_{2,3}(t) &= t^2\psi_{2,4}(t) &= t^3, \end{aligned}$$

the equations for the feasible (flat) trajectory become

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & T & T^2 & T^3 & 0 & 0 & 0 & 0 \\ 0 & 1 & 2T & 3T^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & T & T^2 & T^3 \\ 0 & 0 & 0 & 0 & 0 & 1 & 2T & 3T^2 \end{bmatrix} \begin{bmatrix} \alpha_{11} \\ \alpha_{12} \\ \alpha_{13} \\ \alpha_{14} \\ \alpha_{21} \\ \alpha_{22} \\ \alpha_{23} \\ \alpha_{24} \end{bmatrix} = \begin{bmatrix} x_{1,0} \\ 1 \\ x_{3,0} \\ x_{2,0} \\ x_{1,f} \\ 1 \\ x_{3,f} \\ x_{2,f} \end{bmatrix}.$$

This is a set of 8 linear equations in 8 variables. It can be shown that the matrix M is full rank and the system can be solved numerically. ∇

Note that no ODEs need to be integrated in order to compute the feasible trajectories for a differentially flat system (unlike optimal control methods that we will consider in the next chapter, which involve parameterizing the *input* and then solving the ODEs). This is the defining feature of differentially flat systems. The practical implication is that nominal trajectories and inputs which satisfy the equations of motion for a differentially flat system can be computed in a computationally efficient way (solution of algebraic equations). Since the flat output functions are completely free, the only constraints that must be satisfied are the initial and final conditions on the endpoints, their tangents, and higher order derivatives. Any other constraints on the system, such as bounds on the inputs, can be transformed into the flat output space and (typically) become limits on the curvature or higher order derivative properties of the curve.

If there is a performance index for the system, this index can be transformed and becomes a functional depending on the flat outputs and their derivatives up to some order. By approximating the performance index we can achieve paths for the system that are suboptimal but still feasible. This approach is often much more appealing than the traditional method of approximating the system (for example by its linearization) and then using the exact performance index, which yields optimal paths but for the wrong system.

In light of the techniques that are available for differentially flat systems, the characterization of flat systems becomes particularly important. Unfortunately, general conditions for flatness are not known, but all (dynamic) feedback linearizable systems are differentially flat, as are all driftless systems that can be converted into chained form (see [vNRM94] for details). Another large class of differentially flat systems are those in “pure feedback form”:

$$\begin{aligned}\dot{x}_1 &= f_1(x_1, x_2) \\ \dot{x}_2 &= f_2(x_1, x_2, x_3) \\ &\vdots \\ \dot{x}_n &= f_n(x_1, \dots, x_n, u).\end{aligned}$$

Under certain regularity conditions these systems are differentially flat with output $y = x_1$. These systems have been used for so-called “integrator backstepping” approaches to nonlinear control by Kokotovic et al. [KKM91]. Figure 1.5 shows some additional systems that are differentially flat.

Example 1.3 Planar ducted fan

Consider the dynamics of a planar, vectored thrust flight control system as shown in Figure 1.6. This system consists of a rigid body with body fixed forces and is a simplified model for the Caltech ducted fan [?]. Let (x, y, θ)




Figure 1.5: Examples of flat systems.

denote the position and orientation of the center of mass of the fan. We assume that the forces acting on the fan consist of a force f_1 perpendicular to the axis of the fan acting at a distance r from the center of mass, and a force f_2 parallel to the axis of the fan. Let m be the mass of the fan, J the moment of inertia, and g the gravitational constant. We ignore aerodynamic forces for the purpose of this example.

The dynamics for the system are

$$\begin{aligned} m\ddot{x} &= f_1 \cos \theta - f_2 \sin \theta, \\ m\ddot{y} &= f_1 \sin \theta + f_2 \cos \theta - mg, \\ J\ddot{\theta} &= rf_1. \end{aligned} \tag{1.5}$$

Martin et al. [MDP94] showed that this system is differentially flat and that one set of flat outputs is given by

$$\begin{aligned} z_1 &= x - (J/mr) \sin \theta, \\ z_2 &= y + (J/mr) \cos \theta. \end{aligned} \tag{1.6}$$

Using the system dynamics, it can be shown that

$$\ddot{z}_1 \cos \theta + (\ddot{z}_2 + g) \sin \theta = 0 \tag{1.7}$$

and thus given $z_1(t)$ and $z_2(t)$ we can find $\theta(t)$ except for an ambiguity of π and away from the singularity $\ddot{z}_1 = \ddot{z}_2 + g = 0$. The remaining states and the forces $f_1(t)$ and $f_2(t)$ can then be obtained from the dynamic equations, all in terms of z_1 , z_2 , and their higher order derivatives. ∇




Figure 1.6: Planar ducted fan engine. Thrust is vectored by moving the flaps at the end of the duct.

1.4 Further Reading

The two degree of freedom controller structure introduced in this chapter is described in a bit more detail in ÅM08 [ÅM08] (in the context of output feedback control) and a description of some of the origins of this structure are provided in the “Further Reading” section of Chapter 8. Gain scheduling is a classical technique that is often omitted from introductory control texts, but a good description can be found in the survey article by Rugh [Rug90] and the work of Shamma [Sha90]. Differential flatness was originally developed by Fliess, Levin, Martin and Rouchon [FLMR92]. See [Mur97] for a description of the role of flatness in control of mechanical systems and [vNM98] for more information on flatness applied to flight control systems.

Exercises

1.1 (Feasible trajectory for constant reference) Consider a linear input/output system of the form

$$\dot{A}x + Bu, \quad y = Cx \quad (1.8)$$

in which we wish to track a constant reference r . A feasible (steady state) trajectory for the system is given by solving the equation

$$\begin{bmatrix} 0 \\ r \end{bmatrix} = \begin{bmatrix} A & B \\ C & 0 \end{bmatrix} \begin{bmatrix} x_d \\ u_{\text{ff}} \end{bmatrix}$$

for x_d and u_{ff} .

(a) Show that these equations always has a solution as long as the linear system (1.8) is reachable.

(b) In Section 6.2 of ÅM08 we showed that the reference tracking problem could be solved using a control law of the form $u = -Kx + k_r r$. Show that this is equivalent to a two degree of freedom control design using x_d and u_{ff} and give a formula for k_r in terms of x_d and u_{ff} . Show that this formula matches that given in ÅM08.

1.2 A simplified model of the steering control problem is derived in Åström and Murray, Example 5.12. The model has the form

$$\begin{aligned} \dot{z} &= \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} z + \begin{bmatrix} \gamma \\ 1 \end{bmatrix} u \\ y &= z_1 \end{aligned}$$

where $z \in \mathbb{R}^2$ is the (normalized) state of the system and γ is a parameter related to the speed of the vehicle. Suppose that we wish to track a piecewise constant reference trajectory

$$r = \text{square}(2\pi t/20),$$

where `square` is the square wave function in MATLAB. Suppose further that the speed of the vehicle varies according to the formula

$$\gamma = 2 + 2 \sin(2\pi t/50).$$

Design and implement a gain-scheduled controller for this system by first designing a state space controller that places the closed loop poles of the system at the roots of $s^2 + 2\zeta\omega_0 s + \omega_0^2$, where $\zeta = 0.7$ and $\omega_0 = 1$. You should design controllers for three different parameter values: $\gamma = 0, 2, 4$. Then use linear interpolation to compute the gain for values of γ between these fixed values. Compare the performance of the gain scheduled controller to a simple controller that assumes $\gamma = 2$ for the purpose of the control design (but leaving γ time-varying in your simulation).

Note: a MATLAB file with the vehicle dynamics is available on the course web page. You can use this if you like to get the reference trajectory and parameter variation.

1.3 (Stability of gain scheduled controllers for slowly varying systems (ÅM08)) Consider a nonlinear control system with gain scheduled feedback

$$\dot{e} = f(e, v) \quad v = k(\mu)e,$$

where $\mu(t) \in \mathbb{R}$ is an externally specified parameter (eg, the desired trajectory) and $k(\mu)$ is chosen such that the linearization of the closed loop system around the origin is stable for each fixed μ .

Show that if $|\dot{\mu}|$ is sufficiently small then the equilibrium point is locally asymptotically stable for the full nonlinear, time-varying system. (Hint: find a Lyapunov function of the form $V = x^T P(\mu)x$ based on the linearization of the system dynamics for fixed μ and then show this is a Lyapunov function for the full system.)

1.4 (Flatness of systems in reachable canonical form) Consider a single input system in reachable canonical form [ÅM08, Sec. 6.1]:

$$\frac{dx}{dt} = \begin{bmatrix} -a_1 & -a_2 & -a_3 & \dots & -a_n \\ 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & & & 1 & 0 \end{bmatrix} x + \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} u, \quad (1.9)$$

$$y = [b_1 \ b_2 \ b_3 \ \dots \ b_n] x + du.$$


Suppose that we wish to find an input u that moves the system from x_0 to x_f . This system is differentially flat with flat output given by $z = x_n$ and hence we can parameterize the solutions by a curve of the form

$$x_n(t) = \sum_{k=0}^N \alpha_k t^k, \quad (1.10)$$

where N is a sufficiently large integer.

- Compute the state space trajectory $x(t)$ and input $u(t)$ corresponding to equation (1.10) and satisfying the differential equation (1.9). Your answer should be an equation similar to equation (1.10) for each state x_i and the input u .
- Find an explicit input that steers a double integrator system between any two equilibrium points $x_0 \in \mathbb{R}^2$ and $x_f \in \mathbb{R}^2$.
- Show that all reachable systems are differentially flat and give a formula for the flat output.


1.5 Consider the lateral control problem for an autonomous ground vehicle as described in Example 1.1 and Section 1.3. Using the fact that the system is differentially flat, find an explicit trajectory that solves the following parallel parking maneuver:



Your solution should consist of two segments: a curve from x_0 to x_i with $v > 0$ and a curve from x_i to x_f with $v < 0$. For the trajectory that you

determine, plot the trajectory in the plane (x versus y) and also the inputs v and ϕ as a function of time.

1.6 Consider first the problem of controlling a truck with trailer, as shown in the figure below:



$$\begin{aligned}\dot{x} &= \cos \theta u_1 \\ \dot{y} &= \sin \theta u_1 \\ \dot{\phi} &= u_2 \\ \dot{\theta} &= \frac{1}{l} \tan \phi u_1 \\ \dot{\theta}_1 &= \frac{1}{d} \cos(\theta - \theta_1) \sin(\theta - \theta_1) u_1,\end{aligned}$$

The dynamics are given above, where (x, y, θ) is the position and orientation of the truck, ϕ is the angle of the steering wheels, θ_1 is the angle of the trailer, and l and d are the length of the truck and trailer. We want to generate a trajectory for the truck to move it from a given initial position to the loading dock. We ignore the role of obstacles and concentrate on generation of feasible trajectories.

- Show that the system is differentially flat using the center of the rear wheels of the trailer as the flat output.
- Generate a trajectory for the system that steers the vehicle from an initial condition with the truck and trailer perpendicular to the loading dock into the loading dock.
- Write a simulation of the system stabilizes the desired trajectory and demonstrate your two-degree of freedom control system maneuvering from several different initial conditions into the parking space, with either disturbances or modeling errors included in the simulation.

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.

2.1 Review: Optimization

Consider first the problem of finding the maximum 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 $F(x^*) \geq F(x)$ for all $x \in \mathbb{R}^n$. A necessary condition for x^* to be a maximum is that the gradient of the function be zero at x^* ,

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

Figure 2.1 gives a graphical interpretation of this condition. Note that these are *not* sufficient conditions; the points x_1 and x_2 and x^* in the figure all




Figure 2.1: Optimization of functions. The maximum 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 maximizes $F(x)$ while lying on the surface $G(x) = 0$. (b) We can parameterize the constrained directions by computing the gradient of the constraint G .

satisfy the necessary condition but only one is the (global) maximum.

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 $F(x^*) \geq F(x)$ 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 maximum of the cost function along this surface, as illustrated in Figure 2.2a.

A necessary condition for being at a maximum is that there are no directions tangent to the constraints that also increase the cost. The normal 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 $\frac{\partial F}{\partial x}$ is the usual (gradient) optimality condition while the term $\frac{\partial G}{\partial x}$ 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*. Minimize \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 need to choose x such that $G(x) = 0$. 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 maximum at $x = (a, b)$. Suppose that we add a constraint $G(x) = 0$ given by

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

With this constrain, 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 maximum

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

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

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

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 maximum 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*:

$$\min_u \underbrace{\int_0^T L(x, u) dt}_{\text{integrated cost}} + \underbrace{V(x(T), u(T))}_{\text{final cost}}$$

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

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

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.

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

Infinite Horizon. 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 $T < \infty$.

Linear Quadratic. 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 (assumes $x_d = 0$), $R > 0$ penalizes the input (*must* be positive definite) and $P_1 > 0$ penalizes terminal state.

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 $\psi_i(x(T)) = x_i(T) - x_{i,f}$.

Time Optimal. 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

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

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*:

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

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*

$$\begin{aligned} \dot{x}_i &= \frac{\partial H}{\partial \lambda_i} & -\dot{\lambda}_i &= \frac{\partial H}{\partial x_i} & x(0) &\text{ given, } \psi(x(T)) = 0 \\ & & & & \lambda(T) &= \frac{\partial V}{\partial x}(x(T)) + \nu^T \frac{\partial \psi}{\partial x} \end{aligned}$$

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 (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:

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

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

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

We can now “linearize” the cost function around the optimal solution $x(t) = x^*(t) + \delta x(t)$, $u(t) = u^*(t) + \delta u(t)$. Using Leibnitz’s rule, we have \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

$$\dot{x} = ax + bu, \quad (2.1)$$

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 position.

To solve the problem, we define the various elements used in the maximum principle. Our integrated 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:

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

This is a final value problem for a linear differential equation 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

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

and finally we can write

$$u^*(t) = -\frac{2abc e^{a(2t_f-t_o-t)} x(t_o)}{2a - b^2 c (1 - e^{2a(t_f-t_o)})}$$

$$x^*(t) = x(t_o) e^{a(t-t_o)} + \frac{b^2 c e^{a(t_f-t_o)} x(t_o)}{2a - b^2 c (1 - e^{2a(t_f-t_o)})} \left[e^{a(t_f-t)} - e^{a(t+t_f-2t_o)} \right].$$

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:

$$\dot{x} = \frac{\partial H}{\partial \lambda} = Ax + Bu$$

$$-\dot{\lambda} = \frac{\partial H}{\partial x} = A^T \lambda$$

$$u = \arg \min H = -\text{sgn}(\lambda^T B)$$

The optimal solution always satisfies this equation (necessary condition) with $x(0) = x_0$ and $x(T) = 0$. It follows that the input is always $u = \pm 1 \implies$ “bang-bang”. ▽

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 linear quadratic regulator.

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 two.)

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 problem 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 (Exercise ??). Equation (2.3) is called the .

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 don't 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 $K = \text{lqr}(A, B, Q_x, Q_u)$ 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

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

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.

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

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.



(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.

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 you 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.x

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

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

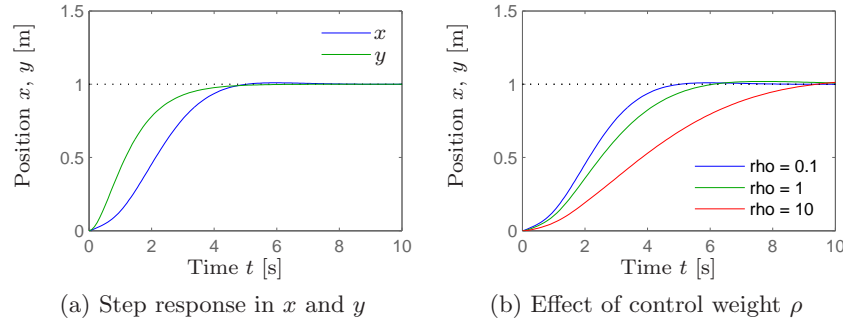


Figure 2.4: Step response for a vectored thrust aircraft. The plot in (a) shows the x and y positions of the aircraft when it is commanded to move 1 m in each direction. In (b) the x motion is shown for control weights $\rho = 1, 10^2, 10^4$. A higher weight of the input term in the cost function causes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




Figure 2.5: Step response for a vector thrust aircraft using physically motivated LQR weights.

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

2.6 Advanced Topics

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.

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 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 $Y \in \text{reals}^{m \times m}$ is a skew symmetric matrix.

(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,

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

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: minimizing $\sqrt{\|\dot{x}\|}$ is the same as minimizing $\dot{x}^T \dot{x}$; this will simplify the algebra a bit.)

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 Consider the optimal control problem for the system

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

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. 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) Derive a set of differential equations for the optimal control $u^*(t)$ and the corresponding state $x^*(t)$ in terms of t_0 , t_f , $x(t_0)$ and t . Be sure to provide any initial or final conditions required for your equations to be solved.
- (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) Increase M by one and show how to choose the free parameter to minimize the cost function.

2.6 Consider the problem of moving a two-wheeled mobile robot (eg, 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 can't 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

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

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 is derived in the lecture notes for week 6 and is shown to be

$$\begin{aligned} u^*(t) &= -\frac{2abc e^{a(2t_f - t_0 - t)} x(t_0)}{2a - b^2c (1 - e^{2a(t_f - t_0)})} \\ x^*(t) &= x(t_0) e^{a(t - t_0)} + \frac{b^2c e^{a(t_f - t_0)} x(t_0)}{2a - b^2c (1 - e^{2a(t_f - t_0)})} \left[e^{a(t_f - t)} - e^{a(t + t_f - 2t_0)} \right]. \end{aligned} \tag{2.7}$$

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 In this problem we will explore the effect of constraints on control of the linear unstable system given by

$$\begin{aligned} \dot{x}_1 &= 0.8x_1 - 0.5x_2 + 0.5u \\ \dot{x}_2 &= x_1 + 0.5u \end{aligned}$$

subject to the constraint that $|u| \leq a$ where a is a positive constant.

(a) Ignore the constraint ($a = \infty$) and design an LQR controller to stabilize the system. Plot the response of the closed system from the initial condition given by $x = (1, 0)$.

(b) Use SIMULINK or `ode45` to simulate the the system for some finite value of a with an initial condition $x(0) = (1, 0)$. Numerically (trial and error) determine the smallest value of a for which the system goes unstable.

(c) Let $a_{\min}(\rho)$ be the smallest value of a for which the system is unstable from $x(0) = (\rho, 0)$. Plot $a_{\min}(\rho)$ for $\rho = 1, 4, 16, 64, 256$.

(d) *Optional:* Given $a > 0$, design and implement a receding horizon control law for this system. Show that this controller has larger region of attraction than the controller designed in part (b). (Hint: solve the finite horizon LQ problem analytically, using the bang-bang example as a guide to handle the input constraint.)

2.11 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:

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

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

- (a) Design an LQR controller that stabilizes the position y to the origin. 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.

Chapter 3

Receding Horizon Control

This set of notes builds on the previous two chapters and explores the use of online optimization as a tool for control of nonlinear control. We begin with an high-level discussion of optimization-based control, refining some of the concepts initially introduced in Chapter 1. We then describe the technique of receding horizon control (RHC), including a proof of stability for a particular form of receding horizon control that makes use of a control Lyapunov function as a terminal cost. We conclude the chapter with a detailed design example, in which we can explore some of the computational tradeoffs in optimization-based control.

Prerequisites. Readers should be familiar with the concepts of trajectory generation and optimal control as described in Chapters 1 and 2. For the proof of stability for the receding horizon controller that we use, familiarity with Lyapunov stability analysis at the level given in ÅM08, Chapter 4 (Dynamic Behavior) is required.

3.1 Optimization-Based Control

Optimization-based control refers to the use of online, optimal trajectory generation as a part of the feedback stabilization of a (typically nonlinear) system. The basic idea is to use a *receding horizon* control technique: a (optimal) feasible trajectory is computed from the current position to the desired position over a finite time T horizon, used for a short period of time $\delta < T$, and then recomputed based on the new position. Development and application of receding horizon control (also called model predictive control, or MPC) originated in process control industries where plants being controlled are sufficiently slow to permit its implementation. An overview of the evolution of commercially available MPC technology is given in [QB97] and a survey of the current state of stability theory of MPC is given in [MRRS00].

Design approach

The basic philosophy that we propose is illustrated in Figure 3.1. We begin with a nonlinear system, including a description of the constraint set. We linearize this system about a representative equilibrium point and perform

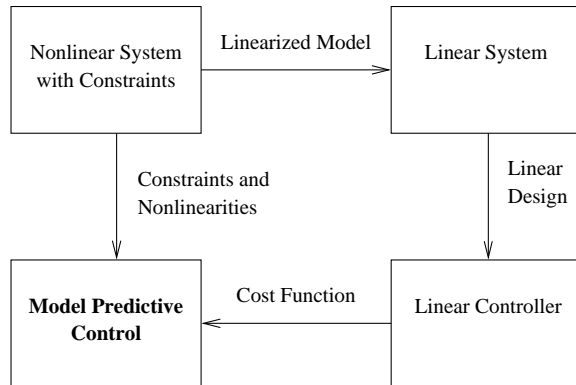


Figure 3.1: Optimization-based control approach.

a linear control design using standard (modern) tools. Such a design gives provably robust performance around the equilibrium point and, more importantly, allows the designer to meet a wide variety of formal and informal performance specifications through experience and the use of sophisticated linear design tools.

This linear control law then serves as a *specification* of the desired control performance for the entire nonlinear system. We convert the control law specification into a receding horizon control formulation, chosen such that for the linearized system, the receding horizon controller gives comparable performance. However, because of its use of optimization tools that can handle nonlinearities and constraints, the receding horizon controller is able to provide the desired performance over a much larger operating envelope than the controller design based just on the linearization. Furthermore, by choosing cost formulations that have certain properties, we can provide proofs of stability for the full nonlinear system and, in some cases, the constrained system.

The advantage of the proposed approach is that it exploits the power of humans in designing sophisticated control laws in the absence of constraints with the power of computers to rapidly compute trajectories that optimize a given cost function in the presence of constraints. New advances in on-line trajectory generation serve as an enabler for this approach and their demonstration on representative flight control experiments shows their viability. This approach can be extended to existing nonlinear paradigms as well, as we describe in more detail below.

A key advantage of optimization-based approaches is that they allow the potential for customization of the controller based on changes in *mission*, *condition*, and *environment*. Because the controller is solving the optimization problem online, updates can be made to the cost function, to change the desired operation of the system; to the model, to reflect changes in parameter values or damage to sensors and actuators; and to the constraints,

to reflect new regions of the state space that must be avoided due to external influences. Thus, many of the challenges of designing controllers that are robust to a large set of possible uncertainties become embedded in the online optimization.

A number of approaches in receding horizon control employ the use of terminal state equality or inequality constraints, often together with a terminal cost, to ensure closed loop stability. In Primbs et al. [PND99], aspects of a stability-guaranteeing, global control Lyapunov function were used, via state and control constraints, to develop a stabilizing receding horizon scheme. Many of the nice characteristics of the CLF controller together with better cost performance were realized. Unfortunately, a global control Lyapunov function is rarely available and often not possible.

Motivated by the difficulties in solving constrained optimal control problems, we have developed an alternative receding horizon control strategy for the stabilization of nonlinear systems [JYH01]. In this approach, closed loop stability is ensured through the use of a terminal cost consisting of a control Lyapunov function that is an incremental upper bound on the optimal cost to go. This terminal cost eliminates the need for terminal constraints in the optimization and gives a dramatic speed-up in computation. Also, questions of existence and regularity of optimal solutions (very important for online optimization) can be dealt with in a rather straightforward manner.

Inverse Optimality

The philosophy presented here relies on the synthesis of an optimal control problem from specifications that are embedded in an externally generated controller design. This controller is typically designed by standard classical control techniques for a nominal plant, absent constraints. In this framework, the controller's performance, stability and robustness specifications are translated into an equivalent optimal control problem and implemented in a receding horizon fashion.

One central question that must be addressed when considering the usefulness of this philosophy is: *Given a control law, how does one find an equivalent optimal control formulation?* The seminal paper by R. E. Kalman [Kal64] lays a solid foundation for this class of problems, known as *inverse optimality*. In this paper, Kalman considers the class of linear time-invariant (LTI) plants with full-state feedback and a single input variable, with an associated cost function that is quadratic in the input and state variables. These assumptions set up the well-known linear quadratic regulator (LQR) problem, by now a staple of optimal control theory.

In Kalman's paper, the mathematical framework behind the LQR problem is laid out, and necessary and sufficient algebraic criteria for optimality are presented in terms of the algebraic Riccati equation, as well as in terms of a condition on the return difference of the feedback loop. In terms of the

LQR problem, the task of synthesizing the optimal control problem comes down to finding the integrated cost weights Q and R given only the dynamical description of the plant represented by matrices A and B and of the feedback controller represented by K . Kalman delivers a particularly elegant frequency characterization of this map [Kal64].

There are two natural extensions of these results: extension to more general dynamical systems and extension to more general optimal control formulations. The contribution of this paper is the simultaneous extension of this approach to systems with constraints along with the extension to the more general receding horizon control framework. A first step in this approach is extension of inverse optimal results to the finite horizon case.

It is important to note that Kalman's results are restricted to the infinite horizon case ($T \rightarrow \infty$) in addition to the assumptions of linearity, time-invariance and quadratic costs. This additional assumption is necessary to derive the results associated with the algebraic Riccati equation ($\dot{P} = 0$). However, we will show that through proper application of terminal costs, the same inverse optimality problem can be soundly addressed in the case of finite horizon length. This problem is addressed by the authors in this paper in the context of Kalman's work; the review of these results will be made mathematically explicit in the next section.

The above results can be generalized to nonlinear systems, in which one takes a nonlinear control system and attempts to find a cost function such that the given controller is the optimal control with respect to that cost.

The history of inverse optimal control for nonlinear systems goes back to the early work of Moylan and Anderson [MA73]. More recently, Sepulchre et al. [SJK97] showed that a nonlinear state feedback obtained by Sontag's formula from a control Lyapunov function (CLF) is inverse optimal. The connections of this inverse optimality result to passivity and robustness properties of the optimal state feedback are discussed in Jankovic *et al.* [JSK99]. The past research on inverse optimality does not consider the constraints on control or state. However, the results on the unconstrained inverse optimality justify the use of a more general nonlinear loss function in the integrated cost of a finite horizon performance index combined with a real-time optimization-based control approach that takes the constraints into account.

Control Lyapunov Functions

For the optimal control problems that we introduce in the next section, we will make use of a terminal cost that is also a control Lyapunov function for the system. Control Lyapunov functions are an extension of standard Lyapunov functions and were originally introduced by Sontag [Son83]. They allow constructive design of nonlinear controllers and the Lyapunov function that proves their stability. A more complete treatment is given in [KKK95].

Consider a nonlinear control system

$$\dot{x} = f(x, u), \quad x \in \mathbb{R}^n, u \in \mathbb{R}^m. \quad (3.1)$$

Definition 3.1 (Control Lyapunov Function). A locally positive function $V : \mathbb{R}^n \rightarrow \mathbb{R}_+$ is called a *control Lyapunov function (CLF)* for a control system (3.1) if

$$\inf_{u \in \mathbb{R}^m} \left(\frac{\partial V}{\partial x} f(x, u) \right) < 0 \quad \text{for all } x \neq 0.$$

In general, it is difficult to find a CLF for a given system. However, for many classes of systems, there are specialized methods that can be used. One of the simplest is to use the Jacobian linearization of the system around the desired equilibrium point and generate a CLF by solving an LQR problem.

It is a well known result that the problem of minimizing the quadratic performance index,

$$J = \int_0^\infty (x^T(t)Qx(t) + u^T Ru(t))dt \quad \text{subject to} \quad \begin{aligned} \dot{x} &= Ax + Bu, \\ x(0) &= x_0, \end{aligned} \quad (3.2)$$

results in finding the positive definite solution of the following Riccati equation:

$$A^T P + PA - PBR^{-1}B^T P + Q = 0 \quad (3.3)$$

The optimal control action is given by

$$u = -R^{-1}B^T Px$$

and $V = x^T Px$ is a CLF for the system.

In the case of the nonlinear system $\dot{x} = f(x, u)$, A and B are taken as

$$A = \left. \frac{\partial f(x, u)}{\partial x} \right|_{(0,0)} \quad B = \left. \frac{\partial f(x, u)}{\partial u} \right|_{(0,0)}$$

where the pairs (A, B) and $(Q^{\frac{1}{2}}, A)$ are assumed to be stabilizable and detectable respectively. Obviously the obtained CLF $V(x) = x^T Px$ will be valid only in a region around the equilibrium $(0, 0)$.

More complicated methods for finding control Lyapunov functions are often required and many techniques have been developed. An overview of some of these methods can be found in [Jad01].

Finite Horizon Optimal Control

We briefly review the problem of optimal control over a finite time horizon as presented in Chapter 2 to establish the notation for the chapter and set some more specific conditions required for receding horizon control. Given an initial state x_0 and a control trajectory $u(\cdot)$ for a nonlinear control system $\dot{x} = f(x, u)$, the state trajectory $x^u(\cdot; x_0)$ is the (absolutely continuous) curve

in \mathbb{R}^n satisfying

$$x^u(t; x_0) = x_0 + \int_0^t f(x^u(\tau; x_0), u(\tau)) d\tau$$

for $t \geq 0$.

The performance of the system will be measured by a given incremental cost $L : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ that is C^2 and fully penalizes both state and control according to

$$L(x, u) \geq c_q(\|x\|^2 + \|u\|^2), \quad x \in \mathbb{R}^n, u \in \mathbb{R}^m$$

for some $c_q > 0$ and $L(0, 0) = 0$. It follows that the quadratic approximation of L at the origin is positive definite,

$$\left. \frac{\partial L}{\partial x} \right|_{(0,0)} \geq c_q I > 0.$$

To ensure that the solutions of the optimization problems of interest are well behaved, we impose some convexity conditions. We require the set $f(x, \mathbb{R}^m) \subset \mathbb{R}^n$ to be convex for each $x \in \mathbb{R}^n$. Letting $p \in \mathbb{R}^n$ represent the co-state, we also require that the pre-Hamiltonian function $u \mapsto p^T f(x, u) + L(x, u) =: K(x, u, p)$ be strictly convex for each $(x, p) \in \mathbb{R}^n \times \mathbb{R}^n$ and that there is a C^2 function $\bar{u}^* : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^m : (x, p) \mapsto \bar{u}^*(x, p)$ providing the global minimum of $K(x, u, p)$. The Hamiltonian $H(x, p) := K(x, \bar{u}^*(x, p), p)$ is then C^2 , ensuring that extremal state, co-state, and control trajectories will all be sufficiently smooth (C^1 or better). Note that these conditions are trivially satisfied for control affine f and quadratic L .

The cost of applying a control $u(\cdot)$ from an initial state x over the infinite time interval $[0, \infty)$ is given by

$$J_\infty(x, u(\cdot)) = \int_0^\infty L(x^u(\tau; x), u(\tau)) d\tau.$$

The optimal cost (from x) is given by

$$J_\infty^*(x) = \inf_{u(\cdot)} J_\infty(x, u(\cdot))$$

where the control functions $u(\cdot)$ belong to some reasonable class of admissible controls (e.g., piecewise continuous or measurable). The function $x \mapsto J_\infty^*(x)$ is often called the *optimal value function* for the infinite horizon optimal control problem.

For the class of f and L considered, we know that $J_\infty^*(\cdot)$ is a positive definite C^2 function on a neighborhood of the origin. This follows from the *geometry* of the corresponding Hamiltonian system (see [HO01] and the references therein). In particular, since $(x, p) = (0, 0)$ is a hyperbolic critical point of the C^1 Hamiltonian vector field $X_H(x, p) := (D_2 H(x, p), -D_1 H(x, p))^T$, the local properties of $J_\infty^*(\cdot)$ are determined by the linear-quadratic approximation to the problem and, moreover, $D^2 J_\infty^*(0) = P > 0$ where P is the

stabilizing solution of the appropriate algebraic Riccati equation.

For practical purposes, we are interested in finite horizon approximations of the infinite horizon optimization problem. In particular, let $V(\cdot)$ be a nonnegative C^2 function with $V(0) = 0$ and define the finite horizon cost (from x using $u(\cdot)$) to be

$$J_T(x, u(\cdot)) = \int_0^T L(x^u(\tau; x), u(\tau)) d\tau + V(x^u(T; x)) \quad (3.4)$$

and denote the optimal cost (from x) as

$$J_T^*(x) = \inf_{u(\cdot)} J_T(x, u(\cdot)).$$

As in the infinite horizon case, one can show, by geometric means, that $J_T^*(\cdot)$ is locally smooth (C^2). Other properties will depend on the choice of V and T .

Let Γ^∞ denote the domain of $J_\infty^*(\cdot)$ (the subset of \mathbb{R}^n on which J_∞^* is finite). It is not too difficult to show that the cost functions $J_\infty^*(\cdot)$ and $J_T^*(\cdot)$, $T \geq 0$, are continuous functions on Γ_∞ [Jad01]. For simplicity, we will allow $J_\infty^*(\cdot)$ to take values in the extended real line so that, for instance, $J_\infty^*(x) = +\infty$ means that there is no control taking x to the origin.

We will assume that f and L are such that the minimum value of the cost functions $J_\infty^*(x)$, $J_T^*(x)$, $T \geq 0$, is attained for each (suitable) x . That is, given x and $T > 0$ (including $T = \infty$ when $x \in \Gamma^\infty$), there is a (C^1 in t) optimal trajectory $(x_T^*(t; x), u_T^*(t; x))$, $t \in [0, T]$, such that $J_T(x, u_T^*(\cdot; x)) = J_T^*(x)$. For instance, if f is such that its trajectories can be bounded on finite intervals as a function of its input size, e.g., there is a continuous function β such that $\|x^u(t; x_0)\| \leq \beta(\|x_0\|, \|u(\cdot)\|_{L_1[0,t]})$, then (together with the conditions above) there will be a minimizing control (cf. [LM67]). Many such conditions may be used to good effect; see [Jad01] for a more complete discussion.

It is easy to see that $J_\infty^*(\cdot)$ is proper on its domain so that the sub-level sets

$$\Gamma_r^\infty := \{x \in \Gamma^\infty : J_\infty^*(x) \leq r^2\}$$

are compact and path connected and moreover $\Gamma^\infty = \bigcup_{r \geq 0} \Gamma_r^\infty$. Note also that Γ^∞ may be a proper subset of \mathbb{R}^n since there may be states that cannot be driven to the origin. We use r^2 (rather than r) here to reflect the fact that our incremental cost is quadratically bounded from below. We refer to sub-level sets of $J_T^*(\cdot)$ and $V(\cdot)$ using

$$\Gamma_r^T := \text{path connected component of } \{x \in \Gamma^\infty : J_T^*(x) \leq r^2\} \text{ containing } 0,$$

and

$$\Omega_r := \text{path connected component of } \{x \in \mathbb{R}^n : V(x) \leq r^2\} \text{ containing } 0.$$

These results provide the technical framework needed for receding horizon control.

3.2 Receding Horizon Control with CLF Terminal Cost

In receding horizon control, a finite horizon optimal control problem is solved, generating an open-loop state and control trajectories. The resulting control trajectory is then applied to the system for a fraction of the horizon length. This process is then repeated, resulting in a sampled data feedback law. Although receding horizon control has been successfully used in the process control industry, its application to fast, stability critical nonlinear systems has been more difficult. This is mainly due to two issues. The first is that the finite horizon optimizations must be solved in a relatively short period of time. Second, it can be demonstrated using linear examples that a naive application of the receding horizon strategy can have disastrous effects, often rendering a system unstable. Various approaches have been proposed to tackle this second problem; see [MRRS00] for a comprehensive review of this literature. The theoretical framework presented here also addresses the stability issue directly, but is motivated by the need to relax the computational demands of existing stabilizing MPC formulations.

Receding horizon control provides a practical strategy for the use of model information through on-line optimization. Every δ seconds, an optimal control problem is solved over a T second horizon, starting from the current state. The first δ seconds of the optimal control $u_T^*(\cdot; x(t))$ is then applied to the system, driving the system from $x(t)$ at current time t to $x_T^*(\delta, x(t))$ at the next sample time $t + \delta$ (assuming no model uncertainty). We denote this receding horizon scheme as $\mathcal{RH}(T, \delta)$.

In defining (unconstrained) finite horizon approximations to the infinite horizon problem, the key design parameters are the terminal cost function $V(\cdot)$ and the horizon length T (and, perhaps also, the increment δ). What choices will result in success?

It is well known (and easily demonstrated with linear examples), that simple truncation of the integral (i.e., $V(x) \equiv 0$) may have disastrous effects if $T > 0$ is too small. Indeed, although the resulting value function may be nicely behaved, the “optimal” receding horizon closed loop system can be unstable.

A more sophisticated approach is to make good use of a suitable terminal cost $V(\cdot)$. Evidently, the best choice for the terminal cost is $V(x) = J_\infty^*(x)$ since then the optimal finite and infinite horizon costs are the same. Of course, if *the* optimal value function were available there would be no need to solve a trajectory optimization problem. What properties of the optimal value function should be retained in the terminal cost? To be effective, the terminal cost should account for the discarded tail by ensuring that the

origin can be reached from the terminal state $x^u(T; x)$ in an efficient manner (as measured by L). One way to do this is to use an appropriate control Lyapunov function which is also an upper bound on the cost-to-go.

The following theorem shows that the use of a particular type of CLF is in fact effective, providing rather strong and specific guarantees.

Theorem 3.1. [JYH01] *Suppose that the terminal cost $V(\cdot)$ is a control Lyapunov function such that*

$$\min_{u \in \mathbb{R}^m} (\dot{V} + L)(x, u) \leq 0 \quad (3.5)$$

for each $x \in \Omega_{r_v}$ for some $r_v > 0$. Then, for every $T > 0$ and $\delta \in (0, T]$, the resulting receding horizon trajectories go to zero exponentially fast. For each $T > 0$, there is an $\bar{r}(T) \geq r_v$ such that $\Gamma_{\bar{r}(T)}^T$ is contained in the region of attraction of $\mathcal{RH}(T, \delta)$. Moreover, given any compact subset Λ of Γ^∞ , there is a T^* such that $\Lambda \subset \Gamma_{\bar{r}(T)}^T$ for all $T \geq T^*$.

Theorem 3.1 shows that for *any* horizon length $T > 0$ and *any* sampling time $\delta \in (0, T]$, the receding horizon scheme is exponentially stabilizing over the set $\Gamma_{r_v}^T$. For a given T , the region of attraction estimate is enlarged by increasing r beyond r_v to $\bar{r}(T)$ according to the requirement that $V(x_T^*(T; x)) \leq r_v^2$ on that set. An important feature of the above result is that, for operations with the set $\Gamma_{\bar{r}(T)}^T$, there is no need to impose stability ensuring constraints which would likely make the online optimizations more difficult and time consuming to solve.

Sketch of proof. Let $x^u(\tau; x)$ represent the state trajectory at time τ starting from initial state x and applying a control trajectory $u(\cdot)$, and let $(x_T^*, u_T^*)(\cdot, x)$ represent the optimal trajectory of the finite horizon, optimal control problem with horizon T . Assume that $x_T^*(T; x) \in \Omega_r$ for some $r > 0$. Then for any $\delta \in [0, T]$ we want to show that the optimal cost $x_T^*(\delta; x)$ satisfies

$$J_T^*(x_T^*(\delta; x)) \leq J_T^*(x) - \int_0^\delta q(L(x_T^*(\tau; x), u_T^*(\tau; x))) d\tau. \quad (3.6)$$

This expression says that solution to the finite-horizon, optimal control problem starting at time $t = \delta$ has cost that is less than the cost of the solution from time $t = 0$, with the initial portion of the cost subtracted off.. In other words, we are closer to our solution by a finite amount at each iteration of the algorithm. It follows using Lyapunov analysis that we must converge to the zero cost solution and hence our trajectory converges to the desired terminal state (given by the minimum of the cost function).

To show equation (3.6) holds, consider a trajectory in which we apply the optimal control for the first T seconds and then apply a closed loop controller using a stabilizing feedback $u = -k(x)$ for another T seconds. (The stabilizing compensator is guaranteed to exist since V is a control

Lyapunov function.) Let $(x_T^*, u_T^*)(t; x)$, $t \in [0, T]$ represent the optimal control and $(x^k, u^k)(t - T; x_T^*(T; x))$, $t \in [T, 2T]$ represent the control with $u = -k(x)$ applied where k satisfies $(\dot{V} + L)(x, -k(x)) \leq 0$. Finally, let $(\tilde{x}(t), \tilde{u}(t))$, $t \in [0, 2T]$ represent the trajectory obtained by concatenating the optimal trajectory (x_T^*, u_T^*) with the CLF trajectory (x^k, u^k) .

We now proceed to show that the inequality (??) holds. The cost of using $\tilde{u}(\cdot)$ for the first T seconds starting from the initial state $x_T^*(\delta; x)$, $\delta \in [0, T]$ is given by

$$\begin{aligned} J_T(x_T^*(\delta; x), \tilde{u}(\cdot)) &= \int_{\delta}^{T+\delta} L(\tilde{x}(\tau), \tilde{u}(\tau)) d\tau + V(\tilde{x}(T + \delta)) \\ &= J_T^*(x) - \int_0^{\delta} L(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau - V(x_T^*(T; x)) \\ &\quad + \int_T^{T+\delta} L(\tilde{x}(\tau), \tilde{u}(\tau)) d\tau + V(\tilde{x}(T + \delta)). \end{aligned}$$

Note that the second line is simply a rewriting of the integral in terms of the optimal cost J_T^* with the necessary additions and subtractions of the additional portions of the cost for the interval $[\delta, T + \delta]$. We can now use the bound

$$L(\tilde{x}(\tau), \tilde{u}(\tau)) \leq \dot{V}(\tilde{x}(\tau), \tilde{u}(\tau)), \quad \tau \in [T, 2T],$$

which follows from the definition of the CLF V and stabilizing controller $k(x)$. This allows us to write

$$\begin{aligned} J_T(x_T^*(\delta; x), \tilde{u}(\cdot)) &\leq J_T^*(x) - \int_0^{\delta} L(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau - V(x_T^*(T; x)) \\ &\quad - \int_T^{T+\delta} \dot{V}(\tilde{x}(\tau), \tilde{u}(\tau)) d\tau + V(\tilde{x}(T + \delta)) \\ &= J_T^*(x) - \int_0^{\delta} L(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau - V(x_T^*(T; x)) \\ &\quad - V(\tilde{x}(\tau)) \Big|_T^{T+\delta} + V(\tilde{x}(T + \delta)) \\ &= J_T^*(x) - \int_0^{\delta} L(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau. \end{aligned}$$

Finally, using the optimality of u_T^* we have that $J_T^*(x_T^*(\delta; x)) \leq J_T(x_T^*(\delta; x), \tilde{u}(\cdot))$ we obtain equation (3.6). \square

An important benefit of receding horizon control is its ability to handle state and control constraints. While the above theorem provides stability guarantees when there are no constraints present, it can be modified to include constraints on states and controls as well. In order to ensure stability when state and control constraints are present, the terminal cost $V(\cdot)$ should be a local CLF satisfying $\min_{u \in \mathcal{U}} \dot{V} + L(x, u) \leq 0$ where \mathcal{U} is the set of

controls where the control constraints are satisfied. Moreover, one should also require that the resulting state trajectory $x^{CLF}(\cdot) \in \mathcal{X}$, where \mathcal{X} is the set of states where the constraints are satisfied. (Both \mathcal{X} and \mathcal{U} are assumed to be compact with origin in their interior). Of course, the set Ω_{r_v} will end up being smaller than before, resulting in a decrease in the size of the guaranteed region of operation (see [MRRS00] for more details).

3.3 Receding Horizon Control Using Differential Flatness

In this section we demonstrate how to use differential flatness to find fast numerical algorithms for solving the optimal control problems required for the receding horizon control results of the previous section. We consider the affine nonlinear control system

$$\dot{x} = f(x) + g(x)u, \quad (3.7)$$

where all vector fields and functions are smooth. For simplicity, we focus on the single input case, $u \in \mathbb{R}$. We wish to find a trajectory of equation (3.7) that minimizes the performance index (3.4), subject to a vector of initial, final, and trajectory constraints

$$\begin{aligned} lb_0 &\leq \psi_0(x(t_0), u(t_0)) \leq ub_0, \\ lb_f &\leq \psi_f(x(t_f), u(t_f)) \leq ub_f, \\ lb_t &\leq S(x, u) \leq ub_t, \end{aligned} \quad (3.8)$$

respectively. For conciseness, we will refer to this optimal control problem as

$$\min_{(x,u)} J(x, u) \quad \text{subject to} \quad \begin{cases} \dot{x} = f(x) + g(x)u, \\ lb \leq c(x, u) \leq ub. \end{cases} \quad (3.9)$$

Numerical Solution Using Collocation

A numerical approach to solving this optimal control problem is to use the direct collocation method outlined in Hargraves and Paris [HP87]. The idea behind this approach is to transform the optimal control problem into a nonlinear programming problem. This is accomplished by discretizing time into a grid of $N - 1$ intervals

$$t_0 = t_1 < t_2 < \dots < t_N = t_f \quad (3.10)$$

and approximating the state x and the control input u as piecewise polynomials \hat{x} and \hat{u} , respectively. Typically a cubic polynomial is chosen for the states and a linear polynomial for the control on each interval. Collocation is then used at the midpoint of each interval to satisfy equation (3.7). Let $\hat{x}(x(t_1), \dots, x(t_N))$ and $\hat{u}(u(t_1), \dots, u(t_N))$ denote the approximations to x and u , respectively, depending on $(x(t_1), \dots, x(t_N)) \in \mathbb{R}^{nN}$ and $(u(t_1), \dots, u(t_N)) \in \mathbb{R}^N$ corresponding to the value of x and u at the grid

points. Then one solves the following finite dimension approximation of the original control problem (3.9):

$$\min_{y \in \mathbb{R}^M} F(y) = J(\hat{x}(y), \hat{u}(y)) \quad \text{subject to} \quad \begin{cases} \dot{\hat{x}} - f(\hat{x}(y)) + g(\hat{x}(y))\hat{u}(y) = 0, \\ lb \leq c(\hat{x}(y), \hat{u}(y)) \leq ub, \\ \forall t = \frac{t_j + t_{j+1}}{2} \quad j = 1, \dots, N-1 \end{cases} \quad (3.11)$$

where $y = (x(t_1), u(t_1), \dots, x(t_N), u(t_N))$, and $M = \dim y = (n+1)N$.

Seywald [Sey94] suggested an improvement to the previous method (see also [Bry99, p. 362]). Following this work, one first solves a subset of system dynamics in equation (3.9) for the the control in terms of combinations of the state and its time derivative. Then one substitutes for the control in the remaining system dynamics and constraints. Next all the time derivatives \dot{x}_i are approximated by the finite difference approximations

$$\dot{\hat{x}}(t_i) = \frac{x(t_{i+1}) - x(t_i)}{t_{i+1} - t_i}$$

to get

$$\left. \begin{array}{l} p(\dot{\hat{x}}(t_i), x(t_i)) = 0 \\ q(\dot{\hat{x}}(t_i), x(t_i)) \leq 0 \end{array} \right\} \quad i = 0, \dots, N-1.$$

The optimal control problem is turned into

$$\min_{y \in \mathbb{R}^M} F(y) \quad \text{subject to} \quad \begin{cases} p(\dot{\hat{x}}(t_i), x(t_i)) = 0 \\ q(\dot{\hat{x}}(t_i), x(t_i)) \leq 0 \end{cases} \quad (3.12)$$

where $y = (x(t_1), \dots, x(t_N))$, and $M = \dim y = nN$. As with the Hargraves and Paris method, this parameterization of the optimal control problem (3.9) can be solved using nonlinear programming.

The dimensionality of this discretized problem is lower than the dimensionality of the Hargraves and Paris method, where both the states and the input are the unknowns. This induces substantial improvement in numerical implementation.

Differential Flatness Based Approach

The results of Seywald give a constrained optimization problem in which we wish to minimize a cost functional subject to $n-1$ equality constraints, corresponding to the system dynamics, at each time instant. In fact, it is usually possible to reduce the dimension of the problem further. Given an output, it is generally possible to parameterize the control and a part of the state in terms of this output and its time derivatives. In contrast to the previous approach, one must use more than one derivative of this output for this purpose.

When the whole state and the input can be parameterized with one




Figure 3.2: Spline representation of a variable.

output, the system is differentially flat, as described in Section 1.3. When the parameterization is only partial, the dimension of the subspace spanned by the output and its derivatives is given by r the *relative degree* of this output [Isi89]. In this case, it is possible to write the system dynamics as

$$\begin{aligned} x &= \alpha(z, \dot{z}, \dots, z^{(q)}) \\ u &= \beta(z, \dot{z}, \dots, z^{(q)}) \\ \Phi(z, \dot{z}, \dots, z^{n-r}) &= 0 \end{aligned} \quad (3.13)$$

where $z \in \mathbb{R}^p$, $p > m$ represents a set of outputs that parameterize the trajectory and $\Phi : \mathbb{R}^n \times \mathbb{R}^m$ represents $n - r$ remaining differential constraints on the output. In the case that the system is flat, $r = n$ and we eliminate these differential constraints.

Unlike the approach of Seywald, it is not realistic to use finite difference approximations as soon as $r > 2$. In this context, it is convenient to represent z using B-splines. B-splines are chosen as basis functions because of their ease of enforcing continuity across knot points and ease of computing their derivatives. A pictorial representation of such an approximation is given in Figure 3.2. Doing so we get

$$z_j = \sum_{i=1}^{p_j} B_{i,k_j}(t) C_i^j, \quad p_j = l_j(k_j - m_j) + m_j$$

where $B_{i,k_j}(t)$ is the B-spline basis function defined in [dB78] for the output z_j with order k_j , C_i^j are the coefficients of the B-spline, l_j is the number of knot intervals, and m_j is number of smoothness conditions at the knots. The set $(z_1, z_2, \dots, z_{n-r})$ is thus represented by $M = \sum_{j \in \{1, r+1, \dots, n\}} p_j$ coefficients.

In general, w collocation points are chosen uniformly over the time interval $[t_o, t_f]$ (though optimal knots placements or Gaussian points may

also be considered). Both dynamics and constraints will be enforced at the collocation points. The problem can be stated as the following nonlinear programming form:

$$\min_{y \in \mathbb{R}^M} F(y) \quad \text{subject to} \quad \begin{cases} \Phi(z(y), \dot{z}(y), \dots, z^{(n-r)}(y)) = 0 \\ lb \leq c(y) \leq ub \end{cases} \quad (3.14)$$

where

$$y = (C_1^1, \dots, C_{p_1}^1, C_1^{r+1}, \dots, C_{p_{r+1}}^{r+1}, \dots, C_1^n, \dots, C_{p_n}^n).$$

The coefficients of the B-spline basis functions can be found using nonlinear programming.

A software package called Nonlinear Trajectory Generation (NTG) has been written to solve optimal control problems in the manner described above (see [MMM00] for details). The sequential quadratic programming package NPSOL by [GMSW] is used as the nonlinear programming solver in NTG. When specifying a problem to NTG, the user is required to state the problem in terms of some choice of outputs and its derivatives. The user is also required to specify the regularity of the variables, the placement of the knot points, the order and regularity of the B-splines, and the collocation points for each output.

3.4 Implementation on the Caltech Ducted Fan

To demonstrate the use of the techniques described in the previous section, we present an implementation of optimization-based control on the Caltech Ducted Fan, a real-time, flight control experiment that mimics the longitudinal dynamics of an aircraft. The experiment is shown in Figure 3.3.

Description of the Caltech Ducted Fan Experiment

The Caltech ducted fan is an experimental testbed designed for research and development of nonlinear flight guidance and control techniques for Uninhabited Combat Aerial Vehicles (UCAVs). The fan is a scaled model of the longitudinal axis of a flight vehicle and flight test results validate that the dynamics replicate qualities of actual flight vehicles [MM99].

The ducted fan has three degrees of freedom: the boom holding the ducted fan is allowed to operate on a cylinder, 2 m high and 4.7 m in diameter, permitting horizontal and vertical displacements. Also, the wing/fan assembly at the end of the boom is allowed to rotate about its center of mass. Optical encoders mounted on the ducted fan, gearing wheel, and the base of the stand measure the three degrees of freedom. The fan is controlled by commanding a current to the electric motor for fan thrust and by commanding RC servos to control the thrust vectoring mechanism.




Figure 3.3: Caltech ducted fan.

The sensors are read and the commands sent by a dSPACE multi-processor system, comprised of a D/A card, a digital I/O card, two Texas Instruments C40 signal processors, two Compaq Alpha processors, and a ISA bus to interface with a PC. The dSPACE system provides a real-time interface to the 4 processors and I/O card to the hardware. The NTG software resides on both of the Alpha processors, each capable of running real-time optimization.

The ducted fan is modeled in terms of the position and orientation of the fan, and their velocities. Letting x represent the horizontal translation, z the vertical translation and θ the rotation about the boom axis, the equations of motion are given by

$$\begin{aligned}
 m\ddot{x} + F_{X_a} - F_{X_b} \cos \theta - F_{Z_b} \sin \theta &= 0 \\
 m\ddot{z} + F_{Z_a} + F_{X_b} \sin \theta - F_{Z_b} \cos \theta &= mg_{\text{eff}} \\
 J\ddot{\theta} - M_a + \frac{1}{r_s} I_p \Omega \dot{x} \cos \theta - F_{Z_b} r_f &= 0,
 \end{aligned} \tag{3.15}$$

where $F_{X_a} = D \cos \gamma + L \sin \gamma$ and $F_{Z_a} = -D \sin \gamma + L \cos \gamma$ are the aerodynamic forces and F_{X_b} and F_{Z_b} are thrust vectoring body forces in terms of

the lift (L), drag (D), and flight path angle (γ). I_p and Ω are the moment of inertia and angular velocity of the ducted fan propeller, respectively. J is the moment of ducted fan and r_f is the distance from center of mass along the X_b axis to the effective application point of the thrust vectoring force. The angle of attack α can be derived from the pitch angle θ and the flight path angle γ by

$$\alpha = \theta - \gamma.$$

The flight path angle can be derived from the spatial velocities by

$$\gamma = \arctan \frac{-\dot{z}}{\dot{x}}.$$

The lift (L), drag (D), and moment (M) are given by

$$L = qSC_L(\alpha) \quad D = qSC_D(\alpha) \quad M = \bar{c}SC_M(\alpha),$$

respectively. The dynamic pressure is given by $q = \frac{1}{2}\rho V^2$. The norm of the velocity is denoted by V , S the surface area of the wings, and ρ is the atmospheric density. The coefficients of lift ($C_L(\alpha)$), drag ($C_D(\alpha)$) and the moment coefficient ($C_M(\alpha)$) are determined from a combination of wind tunnel and flight testing and are described in more detail in [MM99], along with the values of the other parameters.

Real-Time Trajectory Generation

In this section we demonstrate the trajectory generation results by using NTG to generate minimum time trajectories in real time. An LQR-based regulator is used to stabilize the system, allow us to focus on the trajectory generation properties. We focus in this section on aggressive, forward flight trajectories. The next section extends the controller to use a receding horizon controller, but on a simpler class of trajectories.

Stabilization Around Reference Trajectory

The results in this section rely on the traditional two degree of freedom design paradigm described in the Introduction. In this approach, a local control law (inner loop) is used to stabilize the system around the trajectory computed based on a nominal model. This compensates for uncertainties in the model, which are predominantly due to aerodynamics and friction. Elements such as the ducted fan flying through its own wake, ground effects, and thrust not modeled as a function of velocity and angle of attack contribute to the aerodynamic uncertainty. The friction in the vertical direction is also not considered in the model. The prismatic joint has an unbalanced load creating an effective moment on the bearings. The vertical frictional force of the ducted fan stand varies with the vertical acceleration of the ducted fan as well as the forward velocity. Actuation models are not

used when generating the reference trajectory, resulting in another source of uncertainty.

Since only the position of the fan is measured, we must estimate the velocities. We use an extended Kalman filter with the optimal gain matrix is gain scheduled on the (estimated) forward velocity. The Kalman filter outperformed other methods that computed the derivative using only the position data and a filter.

The stabilizing LQR controllers were gain scheduled on pitch angle, θ , and the forward velocity, \dot{x} . The pitch angle was allowed to vary from $-\pi/2$ to $\pi/2$ and the velocity ranged from 0 to 6 m/s. The weights were chosen differently for the hover-to-hover and forward flight modes. For the forward flight mode, a smaller weight was placed on the horizontal (x) position of the fan compared to the hover-to-hover mode. Furthermore, the z weight was scheduled as a function of forward velocity in the forward flight mode. There was no scheduling on the weights for hover-to-hover. The elements of the gain matrices for each of the controller and observer are linearly interpolated over 51 operating points.

Nonlinear Trajectory Generation Parameters

We solve a minimum time optimal control problem to generate a feasible trajectory for the system. The system is modeled using the nonlinear equations described above and computed the open loop forces and state trajectories for the nominal system. This system is not known to be differentially flat (due to the aerodynamic forces) and hence we cannot completely eliminate the differential constraints.

We choose three outputs, $z_1 = x$, $z_2 = z$, and $z_3 = \theta$, which results in a system with one remaining differential constraint. Each output is parameterized with four (intervals), sixth order, C^4 (multiplicity), piecewise polynomials over the time interval scaled by the minimum time. A fourth output, $z_4 = T$, is used to represent the time horizon to be minimized and is parameterized by a scalar. By choosing the outputs to be parameterized in this way, we are in effect controlling the frequency content of inputs. Since we are not including the actuators in the model, it would be undesirable to have inputs with a bandwidth higher than the actuators. There are a total of 37 variables in this optimization problem. The trajectory constraints are enforced at 21 equidistant breakpoints over the scaled time interval.

There are many considerations in the choice of the parameterization of the outputs. Clearly there is a trade between the parameters (variables, initial values of the variables, and breakpoints) and measures of performance (convergence, run-time, and conservative constraints). Extensive simulations were run to determine the right combination of parameters to meet the performance goals of our system.




Figure 3.4: Forward flight test case: (a) θ and \dot{x} desired and actual, (b) desired F_{X_b} and F_{Z_b} with bounds.




Figure 3.5: Forward flight test case: altitude and x position (actual (solid) and desired (dashed)). Airfoil represents actual pitch angle (θ) of the ducted fan.

Forward Flight

To obtain the forward flight test data, the operator commanded a desired forward velocity and vertical position with the joysticks. We set the trajectory update time, δ to 2 seconds. By rapidly changing the joysticks, NTG produces high angle of attack maneuvers. Figure 3.4(a) depicts the reference trajectories and the actual θ and \dot{x} over 60 sec. Figure 3.4(b) shows the commanded forces for the same time interval. The sequence of maneuvers corresponds to the ducted fan transitioning from near hover to forward flight, then following a command from a large forward velocity to a large negative velocity, and finally returning to hover.

Figure 3.5 is an illustration of the ducted fan altitude and x position for these maneuvers. The air-foil in the figure depicts the pitch angle (θ). It is apparent from this figure that the stabilizing controller is not tracking well in the z direction. This is due to the fact that unmodeled frictional

effects are significant in the vertical direction. This could be corrected with an integrator in the stabilizing controller.

An analysis of the run times was performed for 30 trajectories; the average computation time was less than one second. Each of the 30 trajectories converged to an optimal solution and was approximately between 4 and 12 seconds in length. A random initial guess was used for the first NTG trajectory computation. Subsequent NTG computations used the previous solution as an initial guess. Much improvement can be made in determining a “good” initial guess. Improvement in the initial guess will improve not only convergence but also computation times.

Receding Horizon Control

The results of the previous section demonstrate the ability to compute optimal trajectories in real time, although the computation time was not sufficiently fast for closing the loop around the optimization. In this section, we make use of a shorter update time δ , a fixed horizon time T with a quadratic integral cost, and a CLF terminal cost to implement the receding horizon controller described in Section 3.2. We also limit the operation of the system to near hover, so that we can use the local linearization to find the terminal CLF.

We have implemented the receding horizon controller on the ducted fan experiment where the control objective is to stabilize the hover equilibrium point. The quadratic cost is given by

$$\begin{aligned} L(x, u) &= \frac{1}{2} \hat{x}^T Q \hat{x} + \frac{1}{2} \hat{u}^T R \hat{u} \\ V(x) &= \gamma \hat{x}^T P \hat{x} \end{aligned} \tag{3.16}$$

where

$$\begin{aligned} \hat{x} &= x - x_{eq} = (x, z, \theta - \pi/2, \dot{x}, \dot{z}, \dot{\theta}) \\ \hat{u} &= u - u_{eq} = (F_{X_b} - mg, F_{Z_b}) \\ Q &= \text{diag}\{4, 15, 4, 1, 3, 0.3\} \\ R &= \text{diag}\{0.5, 0.5\}, \end{aligned}$$

$\gamma = 0.075$ and P is the unique stable solution to the algebraic Riccati equation corresponding to the linearized dynamics of equation (1.5) at hover and the weights Q and R . Note that if $\gamma = 1/2$, then $V(\cdot)$ is the CLF for the system corresponding to the LQR problem. Instead V is a relaxed (in magnitude) CLF, which achieved better performance in the experiment. In either case, V is valid as a CLF only in a neighborhood around hover since it is based on the linearized dynamics. We do not try to compute off-line a region of attraction for this CLF. Experimental tests omitting the terminal cost and/or the input constraints leads to instability. The results in this section show the success of this choice for V for stabilization. An inner-loop PD controller on $\theta, \dot{\theta}$ is implemented to stabilize to the receding horizon




Figure 3.6: Receding horizon input trajectories

states θ_T^* , $\dot{\theta}_T^*$. The θ dynamics are the fastest for this system and although most receding horizon controllers were found to be nominally stable without this inner-loop controller, small disturbances could lead to instability.

The optimal control problem is set-up in NTG code by parameterizing the three position states (x, z, θ) , each with 8 B-spline coefficients. Over the receding horizon time intervals, 11 and 16 breakpoints were used with horizon lengths of 1, 1.5, 2, 3, 4 and 6 seconds. Breakpoints specify the locations in time where the differential equations and any constraints must be satisfied, up to some tolerance. The value of $F_{X_b}^{\max}$ for the input constraints is made conservative to avoid prolonged input saturation on the real hardware. The logic for this is that if the inputs are saturated on the real hardware, no actuation is left for the inner-loop θ controller and the system can go unstable. The value used in the optimization is $F_{X_b}^{\max} = 9$ N.

Computation time is non-negligible and must be considered when implementing the optimal trajectories. The computation time varies with each optimization as the current state of the ducted fan changes. The following notational definitions will facilitate the description of how the timing is set-up:

i	Integer counter of MPC computations
t_i	Value of current time when MPC computation i started
$\delta_c(i)$	Computation time for computation i
$u_T^*(i)(t)$	Optimal output trajectory corresponding to computation i , with time interval $t \in [t_i, t_i + T]$

A natural choice for updating the optimal trajectories for stabilization is to do so as fast as possible. This is achieved here by constantly resolving the optimization. When computation i is done, computation $i + 1$ is immediately started, so $t_{i+1} = t_i + \delta_c(i)$. Figure 3.6 gives a graphical picture of the timing set-up as the optimal input trajectories $u_T^*(\cdot)$ are updated. As shown in the figure, any computation i for $u_T^*(i)(\cdot)$ occurs for $t \in [t_i, t_{i+1}]$ and the

resulting trajectory is applied for $t \in [t_{i+1}, t_{i+2}]$. At $t = t_{i+1}$ computation $i + 1$ is started for trajectory $u_T^*(i + 1)(\cdot)$, which is applied as soon as it is available ($t = t_{i+2}$). For the experimental runs detailed in the results, $\delta_c(i)$ is typically in the range of $[0.05, 0.25]$ seconds, meaning 4 to 20 optimal control computations per second. Each optimization i requires the current measured state of the ducted fan and the value of the previous optimal input trajectories $u_T^*(i - 1)$ at time $t = t_i$. This corresponds to, respectively, 6 initial conditions for state vector x and 2 initial constraints on the input vector u . Figure 3.6 shows that the optimal trajectories are advanced by their computation time prior to application to the system. A dashed line corresponds to the initial portion of an optimal trajectory and is not applied since it is not available until that computation is complete. The figure also reveals the possible discontinuity between successive applied optimal input trajectories, with a larger discontinuity more likely for longer computation times. The initial input constraint is an effort to reduce such discontinuities, although some discontinuity is unavoidable by this method. Also note that the same discontinuity is present for the 6 open-loop optimal state trajectories generated, again with a likelihood for greater discontinuity for longer computation times. In this description, initialization is not an issue because we assume the receding horizon computations are already running prior to any test runs. This is true of the experimental runs detailed in the results.

The experimental results show the response of the fan with each controller to a 6 meter horizontal offset, which is effectively engaging a step-response to a change in the initial condition for x . The following details the effects of different receding horizon control parameterizations, namely as the horizon changes, and the responses with the different controllers to the induced offset.

The first comparison is between different receding horizon controllers, where time horizon is varied to be 1.5, 2.0, 3.0, 4.0 or 6.0 seconds. Each controller uses 16 breakpoints. Figure 3.7(a) shows a comparison of the average computation time as time proceeds. For each second after the offset was initiated, the data corresponds to the average run time over the previous second of computation. Note that these computation times are substantially smaller than those reported for real-time trajectory generation, due to the use of the CLF terminal cost versus the terminal constraints in the minimum-time, real-time trajectory generation experiments.

There is a clear trend toward shorter average computation times as the time horizon is made longer. There is also an initial transient increase in average computation time that is greater for shorter horizon times. In fact, the 6 second horizon controller exhibits a relatively constant average computation time. One explanation for this trend is that, for this particular test, a 6 second horizon is closer to what the system can actually do. After 1.5 seconds, the fan is still far from the desired hover position and the terminal cost CLF is large, likely far from its region of attraction. Figure 3.7(b)

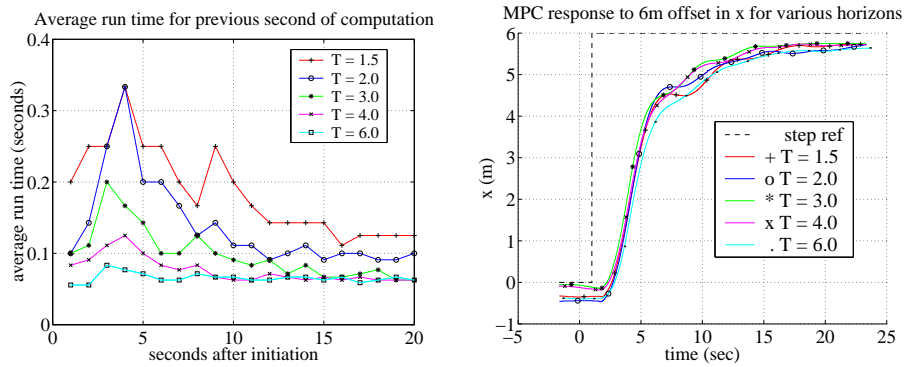


Figure 3.7: Receding horizon control: (a) moving one second average of computation time for MPC implementation with varying horizon time, (b) response of MPC controllers to 6 meter offset in x for different horizon lengths.

shows the measured x response for these different controllers, exhibiting a rise time of 8–9 seconds independent of the controller. So a horizon time closer to the rise time results in a more feasible optimization in this case.

3.5 Further Reading

Exercises

3.1 Consider a nonlinear control system

$$\dot{x} = f(x, u)$$

with linearization

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

Show that if the linearized system is reachable, then there exists a (local) control Lyapunov function for the nonlinear system. (Hint: start by proving the result for a stable system.)

3.2 Consider the optimal control problem given in Example 2.2:

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

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 control for a finite time $T > 0$ 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?

Using the solution given in equation (2.7), implement the finite-time optimal controller in a receding horizon fashion with an update time of $\delta = 0.5$. Using the parameter values in problem 1(b), Compare the responses of the receding horizon controllers to the LQR controller you designed for problem 1, from the same initial condition. What do you observe as c and t_f increase?

(Hint: you can write a MATLAB script to do this by performing the following steps:

- (i) set $t_0 = 0$
- (ii) using the closed form solution for x^* from problem 1, plot $x(t)$, $t \in [t_0, t_f]$ and save $x_\delta = x(t_0 + \delta)$
- (iii) set $x(t_0) = x_\delta$ and repeat step (ii) until x is small.)

3.3 In this problem we will explore the effect of constraints on control of the linear unstable system given by

$$\begin{aligned} \dot{x}_1 &= 0.8x_1 - 0.5x_2 + 0.5u \\ \dot{x}_2 &= x_1 + 0.5u \end{aligned}$$

subject to the constraint that $|u| \leq a$ where a is a positive constant.

(a) Ignore the constraint ($a = \infty$) and design an LQR controller to stabilize the system. Plot the response of the closed system from the initial condition given by $x = (1, 0)$.

(b) Use SIMULINK or `ode45` to simulate the the system for some finite value of a with an initial condition $x(0) = (1, 0)$. Numerically (trial and error) determine the smallest value of a for which the system goes unstable.

(c) Let $a_{\min}(\rho)$ be the smallest value of a for which the system is unstable from $x(0) = (\rho, 0)$. Plot $a_{\min}(\rho)$ for $\rho = 1, 4, 16, 64, 256$.

(d) *Optional:* Given $a > 0$, design and implement a receding horizon control law for this system. Show that this controller has larger region of attraction than the controller designed in part (b). (Hint: solve the finite horizon LQ problem analytically, using the bang-bang example as a guide to handle the input constraint.)

Chapter 4

Stochastic Systems

In this chapter we present a focused review of stochastic systems, oriented toward the material that is required in Chapters 5 and 6. After a brief review of random variables, we define discrete-time and continuous-time random processes, including the expectation, (co-)variance and correlation functions for a random process. These definitions are used to describe linear stochastic systems (in continuous time) and the stochastic response of a linear system to a random process (e.g., noise). We initially derive the relevant quantities in the state space, followed by a presentation of the equivalent frequency domain concepts.

Prerequisites. Readers should be familiar with basic concepts in probability, including random variables and standard distributions. We do not assume any prior familiarity with random processes.

4.1 Review of Random Variables

A (*continuous*) random variable X is a variable that can take on any value according to a probability distribution P :

$P(x_l \leq X \leq x_u)$ = probability that x takes on a value in the range x_l, x_u .

More generally, we write $P(A)$ as the probability that an event A will occur (e.g., $A = \{x_l \leq X \leq x_u\}$). It follows from the definition that if X is a random variable in the range $[L, U]$ then $P(L \leq X \leq U) = 1$. Similarly, if $Y \in [L, U]$ then $P(L \leq X \leq Y) = 1 - P(Y \leq X \leq U)$.

We characterize a random variable in terms of the *probability density function* (pdf), $p(x)$:

$$P(x_l \leq X \leq x_u) = \int_{x_l}^{x_u} p(x) dx \quad (4.1)$$

This can be taken as the definition of the pdf. We will sometimes write $p_X(x)$ when we wish to make explicit that the pdf is associated with the random variable X . Note that we use capital letters to refer to a random variable and lower case letters to refer to a specific value.

Some standard probability distributions include a *uniform distribution*,

$$p(x) = \frac{1}{U - L}, \quad (4.2)$$




Figure 4.1: Probability density function (pdf) for a Gaussian distribution.

and a *Gaussian distribution* (also called a *normal distribution*),

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}. \quad (4.3)$$

In the normal distribution, the parameter μ is called the *mean* of the distribution and σ is called the *standard deviation* of the distribution. Figure 4.1 gives a graphical representation of a Gaussian pdf. There many other distributions that arise in applications, but for the purpose of these notes we focus on uniform distributions and Gaussian distributions.

If two random variables are related, we can talk about their *joint probability*: $P_{X,Y}(A, B)$ is the probability that both event A occurs for X and B occurs for Y . This is sometimes written as $P(A \cap B)$. For continuous random variables, these can be characterized in terms of a *joint probability density function*

$$P(x_l \leq X \leq x_u, y_l \leq Y \leq y_u) = \int_{y_l}^{y_u} \int_{x_l}^{x_u} p(x, y) dx dy. \quad (4.4)$$

The joint pdf thus describes the relationship between X and Y . We say that X and Y are *independent* if $p(x, y) = p(x)p(y)$, which implies that $P_{X,Y}(A, B) = P_X(A)P_Y(B)$ for events A associated with X and B associated with Y .

The *conditional probability* for an event A given that an event B has occurred, written as $P(A|B)$, is given by

$$P(A|B) = \frac{P(A \cap B)}{P(B)}. \quad (4.5)$$

If the events A and B are independent, then $P(A|B) = P(A)$. Note that the individual, joint and conditional probability distributions are all different, so we should really write $P_{X,Y}(A \cap B)$, $P_{X|Y}(A|B)$ and $P_Y(B)$.

If X is dependent on Y then Y is also dependent on X . *Bayes' rule* relates the conditional and individual probabilities:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \quad (4.6)$$

The analog of the probability density function for conditional probability

is the *conditional probability density function* $p(x|y)$

$$p(x|y) = \begin{cases} \frac{p(x, y)}{p(y)} & 0 < p(y) < \infty \\ 0 & \text{otherwise.} \end{cases} \quad (4.7)$$

It follows that

$$p(x, y) = p(x|y)p(y) \quad (4.8)$$

and

$$P(x_l \leq X \leq x_u | y) = \int_{x_l}^{x_u} p(x|y) dx = \frac{\int_{x_l}^{x_u} p(x, y) dx}{p(y)}. \quad (4.9)$$

If X and Y are independent then $p(x|y) = p(x)$ and $p(y|x) = p(y)$. Note that $p(x, y)$ and $p(x|y)$ are different density functions, though they are related through equation (4.8). If X and Y are related with conditional probability distribution $p(x|y)$ then

$$p(x) = \int_{-\infty}^{\infty} p(x, y) dy = \int_{-\infty}^{\infty} p(x|y)p(y) dy.$$

Example 4.1 Conditional probability for sum

Consider three random variables X , Y and Z related by the expression

$$Z = X + Y.$$

In other words, the value of the random variable Z is given by choosing values from two random variables X and Y and adding them. We assume that X and Y are independent Gaussian random variables with mean μ_1 and μ_2 and standard deviation $\sigma = 1$ (the same for both variables).

Clearly the random variable Z is not independent of X (or Y) since if we know the values of X then it provides information about the likely value of Z . To see this, we compute the joint probability between Z and X . Let

$$A = \{x_l \leq x \leq x_u\}, \quad B = \{z_l \leq z \leq z_u\}.$$

The joint probability of both events A and B occurring is given by

$$\begin{aligned} P(A \cap B) &= P(x_l \leq x \leq x_u, z_l \leq x + y \leq z_u) \\ &= P(x_l \leq x \leq x_u, z_l - x \leq y \leq z_u - x). \end{aligned}$$

We can compute this probability by using the probability density functions for X and Y :

$$\begin{aligned} P(A \cap B) &= \int_{x_l}^{x_u} \left(\int_{z_l - x}^{z_u - x} p_Y(y) dy \right) p_X(x) dx \\ &= \int_{x_l}^{x_u} \int_{z_l}^{z_u} p_Y(z - x) p_X(x) dz dx =: \int_{z_l}^{z_u} \int_{x_l}^{x_u} p_{Z, X}(z, x) dx dz. \end{aligned}$$

Using Gaussians for X and Y we have

$$\begin{aligned} p_{Z,X}(z, x) &= \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(z-x-\mu_Y)^2} \cdot \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(x-\mu_X)^2} \\ &= \frac{1}{2\pi} e^{-\frac{1}{2}((z-x-\mu_Y)^2 + (x-\mu_X)^2)}. \end{aligned}$$

▽

Given a random variable X , we can define various standard measures of the distribution. The *expectation* of a random variable is defined as

$$E\{X\} = \int_{-\infty}^{\infty} x p(x) dx,$$

and the *mean square* of a random variable is

$$E\{X^2\} = \int_{-\infty}^{\infty} x^2 p(x) dx.$$

If we let μ represent the expectation (or mean) of X then we define the *variance* of X as

$$E\{(X - \mu)\} = \int_{-\infty}^{\infty} (x - \mu)^2 p(x) dx.$$

We will often write the variance as σ^2 . As the notation indicates, if we have a Gaussian random variable with mean μ and variance σ^2 , then the expectation and variance as computed above return precisely those quantities.

The following properties follow from the definitions.

Proposition 4.1.

1. $E\{\alpha X + \beta Y\} = \alpha E\{X\} + \beta E\{Y\}$

2. If X and Y are Gaussian random processes/variables with

$$p(x) = \frac{1}{\sqrt{2\pi\sigma_x^2}} e^{-\frac{1}{2}\left(\frac{x-\mu_x}{\sigma_x}\right)^2} \quad p(y) = \frac{1}{\sqrt{2\pi\sigma_y^2}} e^{-\frac{1}{2}\left(\frac{y-\mu_y}{\sigma_y}\right)^2}$$

then $X + Y$ is a Gaussian random process/variable with

$$p(x + y) = \frac{1}{\sqrt{2\pi\sigma_z^2}} e^{-\frac{1}{2}\left(\frac{x+y-\mu_z}{\sigma_z}\right)^2}$$

where

$$\mu_z = \mu_x + \mu_y \quad \sigma_z^2 = \sigma_x^2 + \sigma_y^2$$

3. If X is a Gaussian random variable with means μ and variance σ^2 , then αX is Gaussian with mean $\alpha\mu$ and variance $\alpha^2\sigma^2$.

Proof. The first item follows directly from the definition of expectation. The second item is left as an exercise. The third statement is proved using the definitions:

$$\begin{aligned}
 P(x_l \leq \alpha X \leq x_u) &= P\left(\frac{x_l}{\alpha} \leq X \leq \frac{x_u}{\alpha}\right) \\
 &= \int_{\frac{x_l}{\alpha}}^{\frac{x_u}{\alpha}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} dx \\
 &= \int_{x_l}^{x_u} \frac{1}{\alpha\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\left(\frac{y/\alpha-\mu}{\sigma}\right)^2} dy \\
 &= \int_{x_l}^{x_u} \frac{1}{\sqrt{2\pi\alpha^2\sigma^2}} e^{-\frac{1}{2}\left(\frac{y-\alpha\mu}{\alpha\sigma}\right)^2} dy = \int_{x_l}^{x_u} p(y) dy
 \end{aligned}$$

□

4.2 Introduction to Random Processes

In this section we generalize the concept of a random variable to that of a random process. We focus on the scalar state, discrete-time case and try to build up some intuition for the basic concepts that can later be extended to the continuous-time, multi-state case.

A *discrete-time random process* is a stochastic system characterized by the *evolution* of a sequence of random variables $X[k]$. As an example, consider a discrete-time linear system with dynamics

$$x[k+1] = Ax[k] + Bu[k] + Fv[k]. \quad (4.10)$$

As in ÅM08, $x \in \mathbb{R}^n$ represents the state of the system, $u \in \mathbb{R}^p$ is the vector of inputs and $y \in \mathbb{R}^q$ is the vector of outputs. The signal v represents the process disturbances. The (possibly vector-valued) signal v represents disturbances to the process dynamics and w represents noise in the measurements. For simplicity, we will take $u = 0$, $n = 1$ (single state) and $F = 1$.

We wish to describe the evolution of the dynamics when the disturbances and noise are not given as deterministic signals, but rather are chosen from some probability distribution. Thus we will let $V[k]$ be random variables where the values at each instant k are chosen from the probability distribution $P_{V,k}$. As the notation indicates, the distributions might depend on the time instant k , although the most common case is to have a *stationary* distribution in which the distributions are independent of k .

In addition to stationarity, we will often also assume that distribution of values of V at time k is independent of the values of V at time l if $k \neq l$. In other words, $V[k]$ and $V[l]$ are two separate random variables that are independent of each other. We say that the corresponding random process is *uncorrelated* (defined more formally below). As a consequence of our

independence assumption, we have that

$$E\{V[k]V[l]\} = \delta(k-l)E\{V^2[k]\} = \begin{cases} E\{V^2[k]\} & k=l \\ 0 & k \neq l. \end{cases}$$

In the case that $V[k]$ is a Gaussian with mean zero and standard deviation σ , then $E\{V[k]V[l]\} = \delta(k-l)\sigma^2$.

We next wish to describe the evolution of the state x in equation (4.10) in the case when V is a random variable. In order to do this, we describe the state x as a sequence of random variables $X[k]$, $k = 1, \dots, N$. Looking back at equation (4.10), we see that even if $V[k]$ is an uncorrelated sequence of random variables, then the states $X[k]$ are not uncorrelated since

$$X[k+1] = AX[k] + FV[k],$$

and hence the probability distribution for X at time $k+1$ depends on the value of X at time k (as well as the value of V at time k), similar to the situation in Example 4.1.

Since each $X[k]$ is a random variable, we can define the mean and variance as $\mu[k]$ and $\sigma^2[k]$ using the previous definitions at each time k :

$$\begin{aligned} \mu[k] &= E\{X[k]\} = \int_{-\infty}^{\infty} x p(x, k) dx, \\ \sigma^2[k] &= E\{(X[k] - \mu[k])^2\} = \int_{-\infty}^{\infty} (x - \mu[k])^2 p(x, k) dx. \end{aligned}$$

To capture the relationship between the current state and the future state, we define the *correlation function* for a random process as

$$\rho(k_1, k_2) = E\{X[k_1]X[k_2]\} = \int_{-\infty}^{\infty} x_1 x_2 p(x_1, x_2; k_1, k_2) dx_1 dx_2$$

The function $p(x_1, x_2; k_1, k_2)$ is the joint density function, which depends on the times k_1 and k_2 . A process is *stationary* if $p(x, k + \kappa) = p(x, \kappa)$ for all k , $p(x_1, x_2; k_1 + \kappa, k_2 + \kappa) = p(x_1, x_2; k_1, k_2)$, etc. In this case we can write $p(x_1, x_2; \kappa)$ for the joint probability distribution. We will almost always restrict to this case. Similarly, we will write $p(k_1, k_2)$ as $p(\kappa) = p(\kappa, k + \kappa)$.

We can compute the the correlation function by explicitly computing the joint pdf (see Example 4.1) or by directly computing the expectation. Suppose that we take a random process of the form (4.10) with $x[0] = 0$ and V having zero mean and standard deviation σ . The correlation function is

given by

$$\begin{aligned} E\{x[k_1]x[k_2]\} &= E\left\{\left(\sum_{i=0}^{k_1-1} A^{k_1-i} BV[i]\right)\left(\sum_{j=0}^{k_2-1} A^{k_2-j} BV[j]\right)\right\} \\ &= E\left\{\sum_{i=0}^{k_1-1} \sum_{j=0}^{k_2-1} A^{k_1-i} BV[i]V[j]BA^{k_2-j}\right\}. \end{aligned}$$

We can now use the linearity of the expectation operator to pull this inside the summations:

$$\begin{aligned} E\{x[k_1]x[k_2]\} &= \sum_{i=0}^{k_1-1} \sum_{j=0}^{k_2-1} A^{k_1-i} BE\{V[i]V[j]\}BA^{k_2-j} \\ &= \sum_{i=0}^{k_1-1} \sum_{j=0}^{k_2-1} A^{k_1-i} B\sigma^2\delta(i-j)BA^{k_2-j} \\ &= \sum_{i=0}^{k_1-1} A^{k_1-i} B\sigma^2BA^{k_2-i}. \end{aligned}$$

Note that the correlation function depends on k_1 and k_2 .

We can see the dependence of the correlation function on the time more clearly by letting $d = k_2 - k_1$ and writing

$$\begin{aligned} \rho(k, k+d) &= E\{x[k]x[k_d]\} = \sum_{i=0}^{k_1-1} A^{k-i} B\sigma^2BA^{d+k-i} \\ &= \sum_{j=1}^k A^j B\sigma^2BA^{j+d} = \left(\sum_{j=1}^k A^j B\sigma^2BA^j\right)A^d. \end{aligned}$$

In particular, if the discrete time system is stable then $|A| < 1$ and the correlation function decays as we take points that are further departed in time (d large). Furthermore, if we let $k \rightarrow \infty$ (i.e., look at the steady state solution) then the correlation function only depends d (assuming the sum converges) and hence the steady state random process is stationary.

4.3 Continuous-Time, Vector-Valued Random Processes

A *continuous-time random process* is a stochastic system characterized by the evolution of a random variable $X(t)$, $t \in [0, T]$. As in the case of the a discrete-time random process, we are interested in understanding how the (random) state of the system is related at separate times. The process is defined in terms of the ‘‘correlation’’ of $X(t_1)$ with $X(t_2)$.

We call $X(t) \in \mathbb{R}^n$ the *state* of the random process. For the case $n > 1$,

we have a vector of random processes:

$$X(t) = \begin{bmatrix} X_1(t) \\ \vdots \\ X_n(t) \end{bmatrix}$$

We can characterize the state in terms of a (vector-valued) time-varying pdf,

$$P(x_l \leq X_i(t) \leq x_u) = \int_{x_l}^{x_u} p_{X_i}(x; t) dx.$$

Note that the state of a random process is not enough to determine the next state (otherwise it would be a deterministic process). We typically omit indexing of the individual states unless the meaning is not clear from context.

We can characterize the dynamics of a random process by its statistical characteristics, written in terms of *joint probability density functions*:

$$P(x_{1l} \leq X(t_1) \leq x_{1u}, x_{2l} \leq X(t_2) \leq x_{2u}) = \int_{x_{2l}}^{x_{2u}} \int_{x_{1l}}^{x_{1u}} p(x_1, x_2; t_1, t_2) dx_1 dx_2$$

The function $p(x_1, x_2; t_1, t_2)$ is called a *joint probability density function* and depends both on the individual states that are being compared and the time instants over which they are compared. In practice, pdf's are not available for most random processes, so this formulation is mainly useful for analytical derivations. Typically we will *assume* a certain pdf (or class of pdfs) as a model and then do our calculations across this class. One of the most common classes of random variables are Gaussian distributions and, as we shall see, one can often compute closed for solutions in this case.

In general, the distributions used to describe a random process depend on the specific time or times that we evaluate the random variables. However, in some cases the relationship only depends on the difference in time and not the absolute times (similar to the notion of time invariance in deterministic systems, as described in ÅM08). A process is *stationary* if $p(x, t + \tau) = p(x, t)$ for all τ , $p(x_1, x_2; t_1 + \tau, t_2 + \tau) = p(x_1, x_2; t_1, t_2)$, etc. In this case we can write $p(x_1, x_2; \tau)$ for the joint probability distribution and we write $R(\tau) := R(t, t + \tau)$. Stationary distributions roughly correspond to the steady-state properties of a random process and we will often restrict our attention to this case.

In practice we don't usually specify random processes via the joint probability distribution $p(x_i, x_j; t_1, t_2)$ but instead describe them in terms of their mean, covariance and correlation function. The previous definitions for mean, variance and correlation can be extended to the continuous time,




Figure 4.2: Correlation function for a first-order Markov process.

vector-valued case by indexing the individual state:

$$E\{X(t)\} = \begin{bmatrix} E\{X_1(t)\} \\ \vdots \\ E\{X_n(t)\} \end{bmatrix}$$

$$E\{X(t)X^T(t)\} = \begin{bmatrix} E\{X_1(t)X_1(s)\} & \dots & E\{X_1(t)X_n(s)\} \\ & \ddots & \vdots \\ & & E\{X_n(t)X_n(s)\} \end{bmatrix} =: R(t, s)$$

As in the discrete time case, the random variables and their statistical properties are all indexed by the time t (and s). The matrix $R(t, s)$ is called the *correlation matrix* for $X(t) \in \mathbb{R}^n$. If $t = s$ then $R(t, t)$ describes how the elements of x are correlated at time t (with each other) and is called the *covariance matrix*. Note that the elements on the diagonal of $R(t, t)$ are the variances of the corresponding scalar variables.

Example 4.2 First-order Markov process

Consider a first order *Markov process* defined by a Gaussian pdf with $\mu = 0$,

$$p(x, t) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\frac{x^2}{\sigma^2}},$$

and a correlation function given by

$$\rho(t_1, t_2) = \frac{Q}{2\omega_0} e^{-\omega_0|t_2-t_1|}$$

The correlation function is illustrated in Figure 4.2. This is a stationary process. ∇

The terminology and notation for covariance and correlation varies between disciplines. In some communities (e.g., statistics), the term “cross-covariance” is used to refer to the covariance between two random vectors X and Y , to distinguish this from the covariance of the elements of X with each other. The term “cross-correlation” is sometimes also used. **MATLAB** has a number of functions to implement covariance and correlation, which mostly match the terminology here:

- `cov(X)` - this returns the variance of the vector X that represents samples of a given random variable.

- `cov(X, Y)` - I'm not sure what this means yet.
- `xcorr(X, Y)` - the “cross-correlation” between two random sequences. If these sequences came from a random process, this is basically the correlation function.
- `xcov(X, Y)` - this returns the “cross-covariance”, which MATLAB defines as the “mean-removed cross-correlation”.

The MATLAB help pages give the exact formulas used for each, so the main point here is to be careful to make sure you know what you really want.

4.4 Linear Stochastic Systems

We now consider the problem of how to compute the response of a linear system to a random variable. We assume we have a linear system described either in state space or as a transfer function:

$$\begin{aligned} \dot{X} &= AX + BU & H(s) &= C(sI - A)^{-1}B. \\ Y &= CX \end{aligned}$$

Given an input U which is itself a random process with mean $\mu(t)$, variance $\sigma^2(t)$ and correlation $\rho(t)$, what is the description of the random process Y ?

Let U be an uncorrelated, Gaussian random process, with zero mean and covariance Q :

$$\rho(\tau) = Q\delta(\tau).$$

We can write the output of the system in terms of the convolution integral

$$Y(t) = \int_0^t h(t - \tau)U(\tau) d\tau,$$

where $h(t - \tau)$ is the impulse response for the system

$$h(t - \tau) = Ce^{A(t-\tau)}B + D\delta(t - \tau).$$

We now compute the statistics of the output, starting with the mean:

$$\begin{aligned} E\{Y\} &= E\left\{\int_0^t h(t - \eta)U(\eta) d\eta\right\} \\ &= \int_0^t h(t - \eta)E\{U(\eta)\} d\eta = 0. \end{aligned}$$

Note here that we have relied on the linearity of the convolution integral to pull the expectation inside the integral.

We can compute the covariance of the output by computing the correla-

tion $\rho(\tau)$ and setting $\sigma^2 = \rho(0)$. The correlation function for y is

$$\begin{aligned}\rho_Y(t, s) &= E\{Y(t)Y(s)\} = E\left\{\int_0^t h(t-\eta)U(\eta) d\eta \cdot \int_0^s h(s-\xi)U(\xi) d\xi\right\} \\ &= E\left\{\int_0^t \int_0^s h(t-\eta)U(\eta)U(\xi)h(s-\xi) d\eta d\xi\right\}\end{aligned}$$

Once again linearity allows us to exchange expectation and integration

$$\begin{aligned}\rho_y(t, s) &= \int_0^t \int_0^s h(t-\eta)E\{U(\eta)U(\xi)\}h(s-\xi) d\eta d\xi \\ &= \int_0^t \int_0^s h(t-\eta)Q\delta(\eta-\xi)h(s-\xi) d\eta d\xi \\ &= \int_0^t h(t-\eta)Qh(s-\eta) d\eta\end{aligned}$$

Now let $\tau = s - t$ and write

$$\begin{aligned}\rho_y(\tau) &= \rho_y(t, t + \tau) = \int_0^t h(t-\eta)Qh(t+\tau-\eta) d\eta \\ &= \int_0^t h(\xi)Qh(\xi + \tau) d\xi \quad (\text{setting } \xi = t - \eta)\end{aligned}$$

Finally, we let $t \rightarrow \infty$ (steady state)

$$\lim_{t \rightarrow \infty} \rho_y(t, t + \tau) = \bar{\rho}_y(\tau) = \int_0^\infty h(\xi)Qh(\xi + \tau) d\xi \quad (4.11)$$

If this integral exists, then we can compute the second order statistics for the output Y .

We can provide a more explicit formula for the correlation function ρ in terms of the matrices A , F and C by expanding equation (4.11). We will consider the general case where $V \in \mathbb{R}^m$ and $Y \in \mathbb{R}^p$ and use the correlation matrix $R(t, s)$ instead of the correlation function $\rho(t, s)$. Define the *state transition matrix* $\Phi(t, t_0) = e^{A(t-t_0)}$ so that the solution of system (??) is given by

$$x(t) = \Phi(t, t_0)x(t_0) + \int_{t_0}^t \Phi(t, \lambda)Fv(\lambda)d\lambda$$

We will also make use of a special type of random process referred to as “white noise”. A *white noise process* $X(t)$ satisfies $E\{X(t)\} = 0$ and $R(t, s) = W\delta(s - t)$, where $\delta(\tau)$ is the impulse function. White noise is an idealized process, similar to the impulse function or Heaviside (step) function in deterministic systems. In particular, we note that $\rho(0) = E\{X^2(t)\} = \infty$, so we never see this signal in practice. However, like the step function, it is very useful for characterizing the response of a linear system, as described in the following proposition. A more formal definition of white noise, which gives better insight into the terminology, is given in the next section.

Proposition 4.2 (Full stochastic response). *Let $E\{X(t_0)X^T(t_0)\} = P(t_0)$ and V be white noise with $E\{V(\lambda)V^T(\xi)\} = R_V\delta(\lambda - \xi)$. Then the correlation matrix for X is given by*

$$R_X(t, s) = P(t)\Phi^T(s, t)$$

where $P(t)$ satisfies the linear matrix differential equation

$$\dot{P}(t) = AP + PA^T + FR_VF, \quad P(0) = P_0.$$

Proof. Using the definition of the correlation matrix, we have

$$\begin{aligned} E\{X(t)X^T(s)\} &= E\left\{\Phi(t, 0)X(0)X^T(0)\Phi^T(t, 0) + \text{cross terms}\right. \\ &\quad \left.+ \int_0^t \Phi(t, \xi)FV(\xi)d\xi \int_0^s V^T(\lambda)F^T\Phi(s, \lambda)d\lambda\right\} \\ &= \Phi(t, 0)E\{X(0)X^T(0)\}\Phi(s, 0) \\ &\quad + \int_0^t \int_0^s \Phi(t, \xi)FE\{V(\xi)V^T(\lambda)\}F^T\Phi(s, \lambda)d\xi d\lambda \\ &= \Phi(t, 0)P(0)\Phi^T(s, 0) \\ &\quad + \int_0^t \Phi(t, \lambda)FR_V(\lambda)F^T\Phi(s, \lambda)d\lambda \end{aligned}$$

Now use the fact that $\Phi(s, 0) = \Phi(s, t)\Phi(t, 0)$ (and similar relations) to obtain

$$R_X(t, s) = P(t)\Phi^T(s, t)$$

where

$$P(t) = \Phi(t, 0)P(0)\Phi^T(t, 0) + \int_0^t \Phi(t, \lambda)FR_VF^T(\lambda)\Phi^T(t, \lambda)d\lambda$$

Finally, differentiate to obtain

$$\dot{P}(t) = AP + PA^T + FR_VF, \quad P(0) = P_0$$

(see Friedland for details). \square

The correlation matrix for the output Y can be computed using the fact that $Y = CX$ and hence $R_Y = C^T R_X X C$. We will often be interested in the steady-state properties of the output, which are given by the following proposition.

Proposition 4.3 (Steady-state stochastic response). *For a time-invariant linear system driven by white noise, the correlation matrices for the state and output are given by*

$$R_X(\tau) = R_X(t, t + \tau) = Pe^{A^T\tau}, \quad R_Y(\tau) = CR_X(\tau)C^T$$

where P satisfies the algebraic equation

$$AP + PA^T + FR_VF^T = 0 \quad P > 0. \quad (4.12)$$

Equation (4.12) is called the *Lyapunov equation* and can be solved in MATLAB using the function `lyap`.

Example 4.3 First-order system

Consider a scalar linear process

$$\dot{X} = -aX + V, \quad Y = cX,$$

where V is a white, Gaussian random process with variance σ^2 . Using the results of Proposition 4.2, the correlation function for X is given by

$$R_X(t, t + \tau) = p(t)e^{-a\tau}$$

where $p(t) > 0$ satisfies

$$p(t) = -2ap + \sigma^2.$$

We can solve explicitly for $p(t)$ since it is a (non-homogeneous) linear differential equation:

$$p(t) = e^{-2at}p(0) + (1 - e^{-2at})\frac{\sigma^2}{2a}.$$

Finally, making use of the fact that $Y = cX$ we have

$$\rho(t, t + \tau) = c^2(e^{-2at}p(0) + (1 - e^{-2at})\frac{\sigma^2}{2a})e^{-a\tau}.$$

In steady state, the correlation function for the output becomes

$$\rho(\tau) = \frac{c^2\sigma^2}{2a}e^{-a\tau}.$$

Note correlation function has the same form as the first-order Markov process in Example 4.2 (with $Q = c^2\sigma^2$). ∇


4.5 Random Processes in the Frequency Domain

As in the case of deterministic linear systems, we can analyze a stochastic linear system either in the state space or the frequency domain. The frequency domain approach provides a very rich set of tools for modeling and analysis of interconnected systems, relying on the frequency response and transfer functions to represent the flow of signals around the system.

Given a random process $X(t)$, we can look at the frequency content of the properties of the response. In particular, if we let $\rho(\tau)$ be the correlation function for a random process, then we define the *power spectral density function* as the Fourier transform of ρ :

$$S(\omega) = \int_{-\infty}^{\infty} \rho(\tau)e^{-j\omega\tau} d\tau, \quad \rho(\tau) = \frac{1}{2\pi} \int_{-\infty}^{\infty} S(\omega)e^{j\omega\tau} d\omega.$$

The power spectral density provides an indication of how quickly the values of a random process can change through the frequency content: if there



is high frequency content in the power spectral density, the values of the random variable can change quickly in time.

Example 4.4 First-order Markov process

To illustrate the use of these measures, consider a first-order Markov process as defined in Example 4.2.

$$\begin{aligned} \rho(\tau) &= \frac{Q}{2\omega_0} e^{-\omega_0|\tau|} & p(x) &= \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}} \\ S(\omega) &= \int_{-\infty}^{\infty} \frac{Q}{2\omega_0} e^{-\omega|\tau|} e^{-j\omega\tau} d\tau \\ &= \int_{-\infty}^0 \frac{Q}{2\omega_0} e^{(\omega-j\omega)\tau} d\tau + \int_0^{\infty} \frac{Q}{2\omega_0} e^{(-\omega-j\omega)\tau} d\tau \\ &= \frac{Q}{\omega^2 + \omega_0^2} \end{aligned}$$

Be careful: $S(\omega)$ is *not* a transfer function. $S(\omega)$ is always real. ▽

Using the power spectral density, we can more formally define “white noise”: a *white noise process* is a zero-mean, random process with power spectral density $S(\omega) = W = \text{constant}$ for all ω . If $X(t) \in \mathbb{R}^n$ (a random vector), then $W \in \mathbb{R}^{n \times n}$. We see that a random process is white if all frequencies are equally represented in its power spectral density. This spectral property is the reason for the terminology “white”. The following proposition verifies that this formal definition agrees with our previous (time domain) definition.

Proposition 4.4. *For a white noise process,*

$$\rho(\tau) = \frac{1}{2\pi} \int_{-\infty}^{\infty} S(\omega) e^{j\omega\tau} d\omega = W\delta(\tau),$$

where $\delta(\tau)$ is the unit impulse function.

Proof. If $\tau \neq 0$ then

$$\rho(\tau) = \frac{1}{2\pi} \int_{-\infty}^{\infty} W(\cos(\omega\tau) + j \sin(\omega\tau)) d\omega = 0$$

If $\tau = 0$ then $\rho(\tau) = \infty$. Can show that

$$\int_{-\epsilon}^{\epsilon} \int_{-\infty}^{\infty} (\dots) d\omega d\tau = W$$

□

We now compute the spectral density function corresponding to Y

$$\begin{aligned} S_y(\omega) &= \int_{-\infty}^{\infty} \left[\int_0^{\infty} h(\xi) Q h(\xi + \tau) d\xi \right] e^{-j\omega\tau} d\tau \\ &= \int_0^{\infty} h(\xi) Q \left[\int_{-\infty}^{\infty} h(\xi + \tau) e^{-j\omega\tau} d\tau \right] d\xi \\ &= \int_0^{\infty} h(\xi) Q \left[\int_0^{\infty} h(\lambda) e^{-j\omega(\lambda - \xi)} d\lambda \right] d\xi \\ &= \int_0^{\infty} h(\xi) e^{j\omega\xi} d\xi \cdot Q H(j\omega) = H(-j\omega) Q_u H(j\omega) \end{aligned}$$


This is then the response of a linear system to white noise.

Composition:

$$S_y(\omega) = H_2(-j\omega) H_1(-j\omega) Q_u H_1(j\omega) H_2(j\omega)$$


Note that in the frequency domain, we get *multiplication* of frequency responses (same as composition transfer functions)

Example: first order response



We can also compute the correlation function for a random process Y

$$\bar{\rho}_Y(\tau) = \int_0^{\infty} H(\xi) Q H(\xi + \tau) d\xi = \frac{Q}{2\omega_0} e^{-\omega|\tau|}$$



- Shows that $Y(t)$ is correlated to $Y(t + \tau)$ even though input is white (uncorrelated)
- Correlation drops off as τ increases

Spectral factorization

We often want to find a Q and H such that we match the statistics of measured noise. Eg, given $S(\omega)$, find $Q > 0$ and $H(s)$ such that $S(\omega) =$

$$H(-j\omega)QH(j\omega).$$

4.6 Further Reading

Exercises

4.1 A random variable Y is the sum of two independent normally (Gaussian) distributed random variables having means m_1 , m_2 and variances σ_1^2 , σ_2^2 respectively. Show that the probability density function for Y is

$$p(y) = \frac{1}{2\pi\sigma_1\sigma_2} \int_{-\infty}^{\infty} \exp \left\{ -\frac{(y-x-m_1)^2}{2\sigma_1^2} - \frac{(x-m_2)^2}{2\sigma_2^2} \right\} dx$$

and confirm that this is normal (Gaussian) with mean $m_1 + m_2$ and variance $\sigma_1^2 + \sigma_2^2$. (Hint: Use the fact that $p(z|y) = p_x(x) = p_x(z-y)$.)

4.2 (ÅM08, Exercise 7.13) Consider the motion of a particle that is undergoing a random walk in one dimension (i.e., along a line). We model the position of the particle as

$$x[k+1] = x[k] + u[k],$$

where x is the position of the particle and u is a white noise processes with $E\{u[i]\} = 0$ and $E\{u[i]u[j]\} = R_u\delta(i-j)$. We assume that we can measure x subject to additive, zero-mean, Gaussian white noise with covariance 1. Show that the expected value of the particle as a function of k is given by

$$E\{x[k]\} = E\{x[0]\} + \sum_{i=0}^{k-1} E\{u[i]\} = E\{x[0]\} =: \mu_x$$

and the covariance $E\{(x[k] - \mu_x)^2\}$ is given by

$$E\{(x[k] - \mu_x)^2\} = \sum_{i=0}^{k-1} E\{u^2[i]\} = kR_u$$

4.3 Consider a second order system with dynamics

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} -a & 0 \\ 0 & -b \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} v, \quad y = [1 \quad 1] \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

that is forced by Gaussian white noise with zero mean and variance σ^2 . Assume $a, b > 0$.

- Compute the correlation function $\rho(\tau)$ for the output of the system. Your answer should be an explicit formula in terms of a , b and σ .
- Assuming that the input transients have died out, compute the mean and variance of the output.

4.4 Find a constant matrix A and vectors F and C such that for

$$\dot{x} = Ax + Fw, \quad y = Cx$$

the power spectrum of y is given by

$$S(\omega) = \frac{1 + \omega^2}{(1 - 7\omega^2)^2 + 1}$$

Describe the sense in which your answer is unique.

Chapter 5

Kalman Filtering

5.1 Linear Quadratic Estimators

Consider a stochastic system

$$\begin{aligned}\dot{X} &= AX + BU + FV & E\{V(s)V^T(t)\} &= R_V(t)\delta(t-s) \\ Y &= CX + W & E\{W(s)W^T(t)\} &= R_W(t)\delta(t-s)\end{aligned}$$

Assume that the disturbance V and noise W are zero-mean and Gaussian (but not necessarily stationary):

$$p(v) = \frac{1}{\sqrt[2]{2\pi} \sqrt{\det R_V}} e^{-\frac{1}{2}v^T R_V^{-1}v}$$

$$p(w) = \dots \quad (\text{using } R_W)$$

- multi-variable Gaussian with covariance matrix R_V
- in scalar case, $R_V = \sigma^2$

Problem statement:. Find the estimate $\hat{X}(t)$ that minimizes the mean square error $E\{(X(t) - \hat{X}(t))(X(t) - \hat{X}(t))^T\}$ given $\{Y(\tau) : 0 \leq \tau \leq t\}$.

Proposition $\hat{X}(t) = E\{X(t)|Y(\tau), \tau \leq t\}$

- Optimal estimate is just the expectation of the random process X given the *constraint* of the observed output.
- This is the way Kalman originally formulated the problem.
- Can think of this as a *least squares* problem: given all previous $Y(t)$, find the estimate \hat{X} that satisfies the dynamics and minimizes the square error with the measured data.

Proof See text. Basic idea: show that the conditional mean minimizes the mean square error.

Theorem 5.1 (Kalman-Bucy, 1961). *The optimal estimator has the form of a linear observer*

$$\dot{\hat{X}} = A\hat{X} + BU + L(Y - C\hat{X})$$

where $L(t) = P(t)C^T R_W^{-1}$ and $P(t) = E\{(X(t) - \hat{X}(t))(X(t) - \hat{X}(t))^T\}$ and

satisfies

$$\begin{aligned}\dot{P} &= AP + PA^T - PC^T R_W^{-1}(t)CP + FR_V(t)F^T \\ P(0) &= E\{X(0)X^T(0)\}\end{aligned}$$

Proof. (sketch) The error dynamics are given by

$$\begin{aligned}\dot{E} &= (A - LC)E + \xi & \xi &= FV - LW \\ R_\xi &= FR_V F^T + LR_W L^T\end{aligned}$$

The covariance matrix $P_E = P$ for this process satisfies (from last lecture):

$$\dot{P} = (A - LC)P + P(A - LC)^T + FR_V F^T + LR_W L^T.$$

We need to find L such that $P(t)$ is as small as possible. Can show that the L that achieves this is given by

$$R_W L^T = CP \quad \implies \quad L = PC^T R_W^{-1}$$

(See Friedland, Section 9.4). □

Remarks and properties

1. The Kalman filter has the form of a *recursive* filter: given $P(t) = E\{E(t)E^T(t)\}$ at time t , can compute how the estimate and covariance *change*. Don't need to keep track of old values of the output.
2. The Kalman filter gives the estimate $\hat{X}(t)$ and the covariance $P_E(t) \implies$ you can see how well the error is converging.
3. If the noise is stationary (R_V, R_W constant) and if \dot{P} is stable, then the observer gain is constant:

$$L = PC^T R_W^{-1} \quad AP + PA^T - PC^T R_W^{-1} CP + FR_V F^T \quad (\text{algebraic Riccati equation})$$

This is the problem solved by the `lqe` command in MATLAB.

4. The Kalman filter extracts the maximum possible information about output data

$$R = Y - C\hat{X} = \text{residual or } \textit{innovations} \text{ process}$$

Can show that for the Kalman filter, the correlation matrix is

$$R_R(t, s) = W(t)\delta(t - s) \quad \implies \quad \text{white noise}$$

So the output error has *no* remaining dynamic information content (see Friedland section 11.5 for complete calculation)

5.2 Extended Kalman Filters

Consider a *nonlinear* system

$$\begin{aligned}\dot{X} &= f(X, U, V) & X \in \mathbb{R}^n, u \in \mathbb{R}^m \\ Y &= CX + W & V, W \text{ Gaussian white noise processes with} \\ & & \text{covariance matrices } R_V \text{ and } R_W.\end{aligned}$$

Nonlinear observer:

$$\dot{\hat{X}} = f(\hat{X}, U, 0) + L(Y - C\hat{X})$$

Error dynamics: $E = X - \hat{X}$

$$\begin{aligned}\dot{E} &= f(X, U, V) - f(\hat{X}, U, 0) - LC(X - \hat{X}) \\ &= F(E, \hat{X}, U, V) - LCe \quad F(E, \hat{X}, U, V) = f(E + \hat{X}, U, V) - f(\hat{X}, U, 0)\end{aligned}$$

Now linearize around *current* estimate \hat{X}

$$\begin{aligned}\hat{E} &= \frac{\partial F}{\partial E} E + \underbrace{F(0, \hat{X}, U, 0)}_{=0} + \underbrace{\frac{\partial F}{\partial V} V}_{\text{noise}} - \underbrace{LCe}_{\text{observer gain}} + \text{h.o.t} \\ &= \tilde{A}E + \tilde{F}V - LCE\end{aligned}$$

where

$$\left. \begin{aligned}\tilde{A} &= \left. \frac{\partial F}{\partial e} \right|_{(0, \hat{X}, U, 0)} = \left. \frac{\partial f}{\partial X} \right|_{(\hat{X}, U, 0)} \\ \tilde{F} &= \left. \frac{\partial F}{\partial V} \right|_{(0, \hat{X}, U, 0)} = \left. \frac{\partial f}{\partial V} \right|_{(\hat{X}, U, 0)}\end{aligned} \right\} \text{Depend on current estimate } \hat{X}$$

Idea: design observer for the linearized system around *current* estimate

$$\begin{aligned}\dot{\hat{X}} &= f(\hat{X}, U, 0) + L(Y - C\hat{X}) & L &= PC^T R_W^{-1} \\ \dot{P} &= (\tilde{A} - LC)P + P(\tilde{A} - LC)^T + \tilde{F}R_V\tilde{F}^T + LR_WL^T & P(t_0) &= E\{X(t_0)X^T(t_0)\}\end{aligned}$$

This is called the (Schmidt) *extended Kalman filter* (EKF)

Remarks:

1. Can't prove very much about EKF due to nonlinear terms
2. In applications, works *very* well. One of the most used forms of the Kalman filter
3. Unscented Kalman filters
4. Information form of the Kalman filter (next week?)

Application: parameter ID

Consider a linear system with unknown parameters ξ

$$\begin{aligned}\dot{X} &= A(\xi)X + B(\xi)U + FV & \xi \in \mathbb{R}^p \\ Y &= C(\xi)X + W\end{aligned}$$

Parameter ID problem: given $U(t)$ and $Y(t)$, estimate the value of the parameters ξ .

One approach: treat ξ as unknown *state*

$$\left. \begin{aligned}\dot{X} &= A(\xi)X + B(\xi)U + FV \\ \dot{\xi} &= 0\end{aligned}\right\} \rightarrow \frac{d}{dt} \begin{bmatrix} X \\ \xi \end{bmatrix} = \underbrace{\begin{bmatrix} A(\xi) & 0 \\ 0 & 0 \end{bmatrix}}_{f\left(\begin{bmatrix} X \\ \xi \end{bmatrix}, U, V\right)} \begin{bmatrix} X \\ \xi \end{bmatrix} + \underbrace{\begin{bmatrix} B(\xi) \\ 0 \end{bmatrix}}_{h\left(\begin{bmatrix} X \\ \xi \end{bmatrix}, W\right)} U + \begin{bmatrix} F \\ 0 \end{bmatrix} V$$

Now use EKF to *estimate* X and $\xi \implies$ determine unknown parameters $\xi \in \mathbb{R}^p$.

Remark: need various observability conditions on augmented system in order for this to work.

5.3 LQG Control

Return to the original “ H_2 ” control problem

$$\text{Figure} \quad \begin{aligned} \dot{X} &= AX + BU + FV \\ Y &= CX + W \end{aligned} \quad \begin{array}{l} V, W \text{ Gaussian white} \\ \text{noise with covariance} \\ R_V, R_W \end{array}$$

Stochastic control problem: find $C(s)$ to minimize

$$J = E \left\{ \int_0^\infty [(Y - r)^T R_V (Y - r)^T + U^T R_W U] dt \right\}$$

Assume for simplicity that the reference input $r = 0$ (otherwise, translate the state accordingly).

Theorem 5.2 (Separation principle). *The optimal controller has the form*

$$\begin{aligned}\dot{\hat{X}} &= A\hat{X} + BU + L(Y - C\hat{X}) \\ U &= K(\hat{X} - X_d)\end{aligned}$$

where L is the optimal observer gain ignoring the controller and K is the optimal controller gain ignoring the noise.

This is called the *separation principle* (for H_2 control).

5.4 Further Reading

Exercises

Chapter 6

Sensor Fusion

6.1 Kalman Filters in Discrete Time

One of the principal uses of observers in practice is to estimate the state of a system in the presence of *noisy* measurements. We have not yet treated noise in our analysis and a full treatment of stochastic dynamical systems is beyond the scope of this text. In this section, we present a brief introduction to the use of stochastic systems analysis for constructing observers. We work primarily in discrete time to avoid some of the complications associated with continuous time random processes and to keep the mathematical prerequisites to a minimum. This section assumes basic knowledge of random variables and stochastic processes.

Consider a discrete time, linear system with dynamics

$$\begin{aligned}x[k+1] &= Ax[k] + Bu[k] + Fv[k] \\ y[k] &= Cx[k] + w[k],\end{aligned}\tag{6.1}$$

where $v[k]$ and $w[k]$ are Gaussian, white noise processes satisfying

$$\begin{aligned}E\{v[k]\} &= 0 & E\{w[k]\} &= 0 \\ E\{v[k]v[j]^T\} &= \begin{cases} 0 & k \neq j \\ R_v & k = j \end{cases} & E\{w[k]w[j]^T\} &= \begin{cases} 0 & k \neq j \\ R_w & k = j \end{cases} \\ E\{v[k]w[j]^T\} &= 0.\end{aligned}\tag{6.2}$$

We assume that the initial condition is also modeled as a Gaussian random variable with

$$E\{x_0\} = x_0 \quad E\{x_0x_0^T\} = P_0.\tag{6.3}$$

We wish to find an estimate $\hat{x}[k]$ that minimizes the mean square error $E\{(x[k] - \hat{x}[k])(x[k] - \hat{x}[k])^T\}$ given the measurements $\{y(\delta) : 0 \leq \tau \leq t\}$. We consider an observer in the same basic form as derived previously:

$$\hat{x}[k+1] = A\hat{x}[k] + Bu[k] + L[k](y[k] - C\hat{x}[k]).\tag{6.4}$$

The following theorem summarizes the main result.

Theorem 6.1. *Consider a random process $x[k]$ with dynamics (6.1) and noise processes and initial conditions described by equations (6.2) and (6.3). The observer gain L that minimizes the mean square error is given by*

$$L[k] = AP[k]C^T(R_w + CP[k]C^T)^{-1},$$

where

$$\begin{aligned} P[k+1] &= (A - LC)P[k](A - LC)^T + R_v + LR_wL^T \\ P_0 &= E\{X(0)X^T(0)\}. \end{aligned} \quad (6.5)$$

Before we prove this result, we reflect on its form and function. First, note that the Kalman filter has the form of a *recursive* filter: given $P[k] = E\{E[k]E[k]^T\}$ at time k , can compute how the estimate and covariance *change*. Thus we do not need to keep track of old values of the output. Furthermore, the Kalman filter gives the estimate $\hat{x}[k]$ and the covariance $P[k]$, so we can see how reliable the estimate is. It can also be shown that the Kalman filter extracts the maximum possible information about output data. If we form the residual between the measured output and the estimated output,

$$e[k] = y[k] - C\hat{x}[k],$$

we can show that for the Kalman filter the correlation matrix is

$$R_e(j, k) = W\delta_{jk}.$$

In other words, the error is a white noise process, so there is no remaining dynamic information content in the error.

In the special case when the noise is stationary (R_v, R_w constant) and if $P[k]$ converges, then the observer gain is constant:

$$L = APC^T(R_w + CPC^T)^{-1},$$

where P satisfies

$$P = APA^T + R_v - APC^T(R_w + CPC^T)^{-1}CPA^T.$$

We see that the optimal gain depends on both the process noise and the measurement noise, but in a nontrivial way. Like the use of LQR to choose state feedback gains, the Kalman filter permits a systematic derivation of the observer gains given a description of the noise processes. The solution for the constant gain case is solved by the `dlqe` command in MATLAB.

Proof (of theorem). We wish to minimize the mean square of the error, $E\{(x[k] - \hat{x}[k])(x[k] - \hat{x}[k])^T\}$. We will define this quantity as $P[k]$ and then show that it satisfies the recursion given in equation (6.5). By definition,

$$\begin{aligned} P[k+1] &= E\{x[k+1]x[k+1]^T\} \\ &= (A - LC)P[k](A - LC)^T + R_v + LR_wL^T \\ &= AP[k]A^T - AP[k]C^TL^T - LCA^T + L(R_w + CP[k]C^T)L^T. \end{aligned}$$

Letting $R_\epsilon = (R_w + CP[k]C^T)$, we have

$$\begin{aligned} P[k+1] &= AP[k]A^T - AP[k]C^T L^T - LCA^T + LR_\epsilon L^T \\ &= AP[k]A^T + (L - AP[k]C^T R_\epsilon^{-1})R_\epsilon(L - AP[k]C^T R_\epsilon^{-1})^T \\ &\quad - AP[k]C^T R_\epsilon^{-1}CP[k]^T A^T + R_w. \end{aligned}$$

In order to minimize this expression, we choose $L = AP[k]C^T R_\epsilon^{-1}$ and the theorem is proven. \square

6.2 Predictor-Corrector Form

The Kalman filter can be written in a two step form by separating the correction step (where we make use of new measurements of the output) and the prediction step (where we compute the expected state and covariance at the next time instant).

We make use of the notation $\hat{x}[k|j]$ to represent the estimated state at time instant k given the information up to time j (where typically $j = k-1$). Using this notation, the filter can be solved using the following algorithm:

Step 0: Initialization.

$$\begin{aligned} k &= 0 \\ \hat{x}[0|-1] &= E\{x[0]\} \\ P[0|-1] &= E\{x^T[0]x[0]\} \end{aligned}$$

Step 1: Correction.

$$\begin{aligned} \hat{x}[k|k] &= \hat{x}[k|k-1] + L[k](y[k] - C\hat{x}[k|k-1]) \\ P[k|k] &= P[k|k-1] - P[k|k-1]C^T (CP[k|k-1]C^T + R_w[k])^{-1}CP[k|k-1] \end{aligned}$$

Step 2: Prediction.

$$\begin{aligned} \hat{x}[k+1|k] &= A\hat{x}[k|k] + Bu[k] \\ P[k+1|k] &= AP[k|k]A^T + FR_v[k]F^T \end{aligned}$$

Step 3: Iterate. Set k to $k+1$ and repeat steps 1 and 2.

Note that the correction step reduces the covariance by an amount related to the relative accuracy of the measurement, while the prediction step increases the covariance by an amount related to the process disturbance.

This form of the discrete-time Kalman filter is convenient because we can reason about the estimate in the case when we do not obtain a measurement on every iteration of the algorithm. In this case, we simply update the prediction step (increasing the covariance) until we receive new sensor data, at which point we call the correction step (decreasing the covariance).

6.3 Sensor Fusion

6.4 Information Filters

6.5 Further Reading

Exercises

Bibliography

- [AF06] M. Athans and P. L. Falb. *Optimal Control: An Introduction to the Theory and Its Applications*. Dover, 2006. Originally published in 1963.
- [ÅM08] K. J. Åström and R. M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, 2008. Available at <http://www.cds.caltech.edu/~murray/amwiki>.
- [BH75] A. E. Bryson, Jr. and Y.-C. Ho. *Applied Optimal Control: Optimization, Estimation, and Control*. Wiley, New York, 1975.
- [Bro81] R. W. Brockett. Control theory and singular Riemannian geometry. In *New Directions in Applied Mathematics*, pages 11–27. Springer-Verlag, New York, 1981.
- [Bry99] A. E. Bryson. *Dynamic optimization*. Addison Wesley, 1999.
- [dB78] C. de Boor. *A Practical Guide to Splines*. Springer-Verlag, 1978.
- [FLMR92] M. Fliess, J. Levine, P. Martin, and P. Rouchon. On differentially flat nonlinear systems. *Comptes Rendus des Séances de l'Académie des Sciences*, 315:619–624, 1992. Serie I.
- [GMSW] P. E. Gill, W. Murray, M. A. Saunders, and M. Wright. *User's Guide for NPSOL 5.0: A Fortran Package for Nonlinear Programming*. Systems Optimization Laboratory, Stanford University, Stanford, CA 94305.
- [HO01] J. Hauser and H. Otinga. On the geometry of optimal control: The inverted pendulum example. In *American Control Conference*, 2001.
- [HP87] C. Hargraves and S. Paris. Direct trajectory optimization using nonlinear programming and collocation. *AIAA J. Guidance and Control*, 10:338–342, 1987.
- [Isi89] A. Isidori. *Nonlinear Control Systems*. Springer-Verlag, 2nd edition, 1989.
- [Jad01] A. Jadbabaie. *Nonlinear Receding Horizon Control: A Control Lyapunov Function Approach*. PhD thesis, California Institute of Technology, Control and Dynamical Systems, 2001.
- [JSK99] M. Jankovic, R. Sepulchre, and P. V. Kokotović. CLF based designs with robustness to dynamic input uncertainties. *Systems Control Letters*, 37:45–54, 1999.
- [JYH01] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 46, 2001.
- [Kal64] R. E. Kalman. When is a linear control system optimal? *J. Basic Engrg. Trans. ASME Ser. D*, 86:51–60, 1964.
- [KKK95] M. Krstić, I. Kanellakopoulos, and P. Kokotović. *Nonlinear and Adaptive Control Design*. Wiley, 1995.

- [KKM91] I. Kanellakopoulos, P. V. Kokotovic, and A. S. Morse. Systematic design of adaptive controllers for feedback linearizable systems. *IEEE Transactions on Automatic Control*, 36(11):1241–1253, 1991.
- [LM67] E. B. Lee and L. Markus. *Foundations of Optimal Control Theory*. Robert E. Krieger Publishing Company, 1967.
- [LS95] F. L. Lewis and V. L. Syrmos. *Optimal Control*. Wiley, second edition, 1995.
- [Lue97] David G. Luenberger. *Optimization by Vector Space Methods*. Wiley, New York, 1997.
- [MA73] P. J. Moylan and B. D. O. Anderson. Nonlinear regulator theory and an inverse optimal control problem. *IEEE Trans. on Automatic Control*, 18(5):460–454, 1973.
- [MDP94] P. Martin, S. Devasia, and B. Paden. A different look at output tracking—Control of a VTOL aircraft. *Automatica*, 32(1):101–107, 1994.
- [MM99] M. B. Milam and R. M. Murray. A testbed for nonlinear flight control techniques: The Caltech ducted fan. In *Proc. IEEE International Conference on Control and Applications*, 1999.
- [MMM00] M. B. Milam, K. Mushambi, and R. M. Murray. A computational approach to real-time trajectory generation for constrained mechanical systems. In *Proc. IEEE Control and Decision Conference*, 2000.
- [MRRS00] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.
- [Mur97] R. M. Murray. Nonlinear control of mechanical systems: A Lagrangian perspective. *Annual Reviews in Control*, 21:31–45, 1997.
- [PBGM62] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko. *The Mathematical Theory of Optimal Processes*. Wiley-Interscience, 1962. (translated from Russian).
- [PND99] J. A. Primbs, V. Nevistić, and J. C. Doyle. Nonlinear optimal control: A control Lyapunov function and receding horizon perspective. *Asian Journal of Control*, 1(1):1–11, 1999.
- [QB97] S. J. Qin and T. A. Badgwell. An overview of industrial model predictive control technology. In J.C. Kantor, C.E. Garcia, and B. Carnahan, editors, *Fifth International Conference on Chemical Process Control*, pages 232–256, 1997.
- [Rug90] W. J. Rugh. Analytical framework for gain scheduling. In *Proc. American Control Conference*, pages 1688–1694, 1990.
- [Sey94] H. Seywald. Trajectory optimization based on differential inclusion. *J. Guidance, Control and Dynamics*, 17(3):480–487, 1994.
- [Sha90] J. S. Shamma. Analysis of gain scheduled control for nonlinear plants. *IEEE Transactions on Automatic Control*, 35(12):898–907, 1990.
- [SJK97] R. Sepulchre, M. Jankovic, and P. V. Kokotović. *Constructive Nonlinear Control*. Springer, London, 1997.
- [Son83] E. D. Sontag. A Lyapunov-like characterization of asymptotic controllability. *SIAM Journal of Control and Optimization*, 21:462–471, 1983.
- [vNM98] M. J. van Nieuwstadt and R. M. Murray. Rapid hover to forward flight transitions for a thrust vectored aircraft. *Journal of Guidance, Control, and Dynamics*, 21(1):93–100, 1998.

- [vNRM94] M. van Nieuwstadt, M. Rathinam, and R. M. Murray. Differential flatness and absolute equivalence. Technical Report CIT/CDS 94-006, California Institute of Technology, 1994. Submitted *SIAM J. Control and Optimization*. Available electronically from <http://cds.caltech.edu/cds/reports>.