Chapter 1

Introduction

The goal of this introduction is to give a quick light presentation of the main ideas developed in this book. Thus we focus on laying out the structure for a basic nonlinear control problem which has the remarkably jargony name, the “two block problem”. This is a nice level of generality for an introduction, because it contains most of the ideas behind the more general four block problem (which is treated in detail in the book), the formulas are simpler, and the linear two block problem corresponds to the paradigm problem of classical control. This paradigm problem is often called the “mixed sensitivity problem”, and Chapter 9 is devoted to its solution. We turn now to our overview of the theory which is the subject of this book. We postpone historical discussion and motivation to §1.11. Linear systems are discussed in §§1.2, 1.3.6.

1.1 The Standard Problem of Nonlinear H_∞ Control

Now we introduce a *special case of the standard problem* of H_∞ control. It entails a description of the plant and controller models, and definitions of the control objectives. This is motivated in §1.8 and actually done carefully in Chapter 9. The standard control problem corresponds to the Figure 1.1, which we now explain.

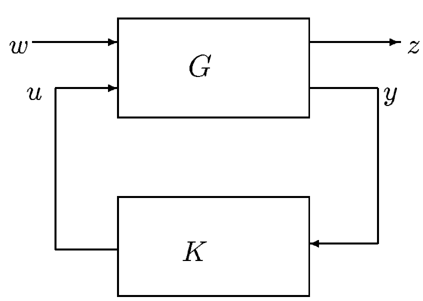

Figure 1.1: The Closed-Loop System (G, K)

1.1.1 The Plant

Let's consider nonlinear plants G with a two block structure

$$G : \begin{cases} \dot{x} = A(x) + B_1(x)w + B_2(x)u \\ z = C_1(x) + D_{12}(x)u \\ y = C_2(x) + w \end{cases} \tag{1.1}$$

Here, $x(t) \in \mathbf{R}^n$ denotes the state of the system, and is not in general directly measurable; instead an output $y(t) \in \mathbf{R}^p$ is observed. The additional output quantity $z(t) \in \mathbf{R}^r$ is a performance measure, depending on the particular problem at hand. The control input is $u(t) \in \mathbf{R}^m$, while $w(t) \in \mathbf{R}^s$ is regarded as an opposing disturbance input. Detailed assumptions concerning the functions appearing in (1.1) will be given in Chapter 2, however we mention here that the origin is an equilibrium and the two block structure requires

$$E_1(x) \stackrel{\triangle}{=} D_{12}(x)' D_{12}(x) > 0. \tag{2.5}$$

(More generally, in the four block case the coefficient of w in the third equation of (1.1) is a matrix D_{21} which will satisfy a condition (2.3) given in Chapter 2.)

1.1.2 The Class of Controllers

The plant G was described by an explicit state space model and is assumed given. However, in the spirit of optimal control theory, we do *not* prescribe a state space model for the controller K, since it is an unknown to be determined from the control objectives. Rather, we simply stipulate some basic input-output properties required of any *admissible controller*, namely that the controller must be a causal function of the output

$$K \ : y(\cdot) \mapsto u(\cdot)$$

and the resulting closed loop system be well-defined in the sense that trajectories and signals exist and are unique. The controller K will be said to be *null-initialized* if $K(0) = 0$, regardless of whether or not a state space realization of K is given.

1.1.3 Control Objectives

The H_∞ control problem is commonly thought of as having two objectives: find a controller K such that the closed loop system (G, K) is

(i) dissipative, and

(ii) stable.

In §1.2 we define what is meant by these terms in the case of linear systems. We now describe their meanings for nonlinear systems; this gives an extension of H_∞ control to nonlinear systems.*

The closed loop system (G, K) is γ-*dissipative* if there exist $\gamma > 0$ and a function $\beta(x_0) \geq 0$ $\beta(0) = 0$, such that

$$\begin{cases} \frac{1}{2}\int_0^T |z(s)|^2\, ds \;\leq\; \gamma^2 \frac{1}{2}\int_0^T |w(s)|^2\, ds + \beta(x_0) \\ \\ \text{for all } \; w \in L_{2,T} \text{ and all } T \geq 0. \end{cases} \tag{1.2}$$

This definition is saying that the nonlinear input-output map $(G, K) : w \mapsto z$ defined by the closed loop system has finite L_2 gain with a bias term due to the initial state x_0 of the plant G.

While dissipation captures the notion of performance of a control system, another issue with H_∞ control is stability of the system. The closed loop system will be called *weakly internally stable* provided that if G is initialized at any x_0, then if $w(\cdot) \in L_2[0, \infty)$, all signals $u(\cdot), y(\cdot), z(\cdot)$ in the loops as well as $x(\cdot)$ converge to 0 as $t \to \infty$. By *internal stability* we mean that the closed loop is weakly internally stable and in addition if the controller has a state space realization, then the controller state will converge to an equilibrium as $t \to \infty$.

Dissipation and stability are closely related; see, e.g. [Wil72], [HM76], [HM77], [vdS96]. Indeed, dissipative systems which enjoy a detectability or observability property also enjoy a stability property. In our context, suppose the system (G, K) is z-*detectable*, that is, $w(\cdot)$ and $z(\cdot) \in L^2[0, \infty)$ imply $x(\cdot) \in L^2[0, \infty)$ and $x(t) \to 0$ as $t \to \infty$. By z-*observable* we mean that if $w(\cdot) = 0$, $(\cdot) = 0$, then $x(\cdot) = 0$. *If (G, K) is γ-dissipative and z-detectable, then (G, K) is weakly internally stable (see Theorem 2.1.3).*

1.1.4 A Classic Example

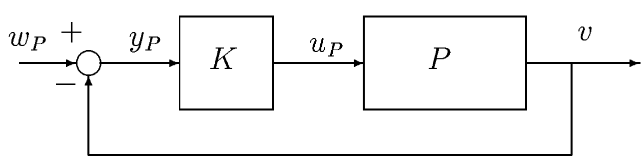

Figure 1.2: Mixed Sensitivity Setup

The book is written for readers with many different interests, so it is worth emphasizing for the reader with a classical control bent that the problem

*The term “nonlinear H_∞ control” has no precise mathematical meaning, but it has come into common use in the control engineering community and refers to nonlinear generalizations of H_∞ control (which has precise meaning for linear systems).

> Given plant P, find a controller K achieving a given H_∞ performance specification (see Figure 1.2).

is of the type we just introduced. In linear H_∞ control the designer selects certain weights and optimizes a worst case frequency domain performance. This is called the *mixed sensitivity* problem of H_∞ control; see Chapter 9. If the weights are chosen correctly for a mixed sensitivity problem then one gets the standard two block problem of H_∞ control which we just presented. Choice of weights is a serious business in practice, and some serious investigation of how this should be done for nonlinear systems is in its infancy. In §1.8 and Chapter 9 we describe some basic considerations in selecting weights.

1.2 The Solution for Linear Systems

The H_∞ problem is well understood when the systems are linear. The plant is linear provided

$$A(x) = Ax, \quad C_1(x) = C_1 x, \ \ C_2(x) = C_2 x,$$
$$B_1(x) = B_1, \ \ B_2(x) = B_2, \quad D_{12}(x) = D_{12},$$

where

$$A, \quad C_1, \, C_2, \quad B_1, \, B_2, \, D_{12},$$

are matrices of appropriate dimension. We recall here the well-known solution to the H_∞ control problem for the two block linear systems, see [DGKF89], [PAJ91], [GL95], etc. (these references also contain the "standard assumptions").

1.2.1 Problem Formulation

The class of admissible controllers K are those with finite dimensional linear state space realizations

$$K \ : \ \begin{cases} \dot{\eta} \ = A_K \eta + B_{1K} y + B_{2K} u \\ \\ u \ = C_K + D_K y \end{cases}$$

Given $\gamma > 0$, the H_∞ control problem for G is to find, if possible, a compensator K such that the resulting closed loop system $(G, K) : w \mapsto z$ satisfies:

(i) *Dissipation:* The required dissipation property is expressed in the frequency domain in terms of the H_∞ norm of the closed loop transfer function $(G, K)(s)$ as follows:

$$\| \ (G, K) \ \|_{H_\infty} = \sup_{\omega \in \mathbf{R}} \sigma_{max}[(G, K)(j\omega)] < \gamma.$$

(ii) *Stability*: We require that the closed loop system

$$(G, K) \ \text{ is } \textit{internally stable}.$$

Some discussion of the classical transfer function (classical loop shaping) pictures vs the state space picture of control is found in §1.8.

1.2.2 Background on Riccati Equations

Recall a few facts about Riccati equations. An algebraic Riccati equation

$$\Sigma\ A + A\Sigma + \Sigma\ R\ \Sigma + Q = 0 \tag{1.3}$$

with real matrix entries A, R, Q and R, Q selfadjoint, meeting suitable positivity and technical conditions (see, e.g., [ZDG96, Chapter 13]), has upper and lower solutions Σ_a, Σ_r so that any other self adjoint solution Σ lies between them

$$\Sigma_a \leq \Sigma \leq \Sigma_r.$$

The bottom solution is called the *stabilizing solution* because it has and is characterized by the property

$$A + R\Sigma_a \tag{1.4}$$

is asymptotically stable. Likewise Σ_r is *antistabilizing* in that

$$-(A + R\Sigma_r) \tag{1.5}$$

is asymptotically stable.

1.2.3 Standard Assumptions

There are a number of "standard assumptions" that are needed for the necessity and sufficiency theorems about H_∞ control. These can be expressed in various ways and here we follow [PAJ91].

The first condition we have already seen, viz. the rank condition (2.5). The D_{12} rank condition ensures that the cost term $|z|^2$ is strictly positive definite in the control u (while the more general 4-block condition (2.3), Chapter 2, relates to the solvability for w given x, y in the output equation $y = C_2 x + D_{12} w$).

Next are two important technical conditions which take the form

$$\text{rank} \begin{bmatrix} A - j\omega I & B_2 \\ \\ C_1 & D_{12} \end{bmatrix} = n + m \ \text{ for all } \omega \geq 0, \tag{1.6}$$

and

$$\text{rank} \begin{bmatrix} A - j\omega I & B_1 \\ \\ C_2 & D_{21} \end{bmatrix} = n + l \ \text{ for all } \omega \geq 0. \tag{1.7}$$

The condition (1.6) can be replaced by a stronger H_2 state feedback assumption (Chapter 5), while (1.7) can be replaced by a stronger H_2 filtering assumption (Chapter 11). These two conditions are commonly used in H_2 control and filtering, and concern the controllability and observability of underlying systems.

1.2.4 Problem Solution

The necessary and sufficient conditions for solvability of the H_∞ problem under the standard assumptions are:

Condition 1: *State feedback Control.* There exists $X_e \geq 0$ solving the control-type Riccati equation

$$\begin{aligned} \mathcal{R}_{state}(X_e) &\triangleq (A - B_2E_1^{-1}D_{12}'C_1)'X_e + X_e(A - B_2E_1^{-1}D_{12}'C_1) \\ &+X_e(\tfrac{1}{\gamma^2}B_1B_1' - B_2E_1^{-1}B_2')X_e + C_1'(I - D_{12}E_1^{-1}D_{12}')C_1 = 0 \end{aligned} \tag{1.8}$$

which is stabilizing, i.e.,

$$A - B_2E_1^{-1}D_{12}'C_1 + (\gamma^{-2}B_1B_1' - B_2E_1^{-1}B_2')X_e \quad \text{is asymptotically stable.} \tag{1.9}$$

Condition 2: *State Estimation*. There exists $Y_e \geq 0$ solving the filter-type Riccati equation

$$\begin{aligned} &\mathcal{R}_{est}(Y_e) \\ &\triangleq (A - B_1C_2)Y_e + Y_e(A - B_1C_2)' + Y_e(\gamma^{-2}C_1'C_1 - C_2'C_2)Y_e = 0 \end{aligned} \tag{1.10}$$

which is stabilizing, i.e.,

$$A - B_1C_2 + Y_e(\gamma^{-2}C_1'C_1 - C_2'C_2) \;\text{ is asymptotically stable.} \tag{1.11}$$

Condition 3: *Coupling.* The matrix X_eY_e has spectral radius strictly less than γ.

THEOREM 1.2.1 *([DGKF89]) ([DGKF89], [PAJ91], [GL95]) The H_∞ control problem for G, meeting certain technical conditions is solvable if and only if the above three conditions are satisfied. If these conditions are met, one controller, called the central controller, is given by*

$$K^* \;:\; \begin{cases} \dot{\hat{x}} &= \big(A - B_1C_2 + Y_e(\gamma^{-2}C_1'C_1 - C_2'C_2)\big)\,\hat{x} \\ & \qquad +(B_1 + Y_eC_2')y + (B_2 + \gamma^{-2}Y_eC_1'D_{12})u \\ u &= -E_1^{-1}[D_{12}'C_1 + B_2'X_e](I - \gamma^{-2}Y_eX_e)^{-1}\hat{x}. \end{cases} \tag{1.12}$$

We sometimes refer K^* as the *"DGKF" central controller*, after its discoverers J. Doyle, K. Glover, P. Kargonekar and B. Francis, [DGKF89].

1.2.5 The Linear Solution from a Nonlinear Viewpoint

Of course the solution to the nonlinear control problem which we present in this book when specialized to linear systems solves the linear H_∞ control problem. The solution looks a bit different from the classical one we just saw. The linear solution has been put in coordinates which make degenerate cases appear un-pathological. However, it is not easy to change coordinates in nonlinear solutions, so what we get is forced upon us. Let us see what the linear specialization of the nonlinear solution looks like.

If $Y_e > 0$, and hence invertible, the coupling condition is equivalent to

$$-\gamma^2 Y_e^{-1} + X_e < 0.$$

This foreshadows the nonlinear theory in that it focusses on the *inverse* of Y_e. Moreover, we shall see that one does not actually need the stabilizing properties of X_e and Y_e; positive definite inequalities will do. Indeed if we take the main results of this book given in Chapter 4, §4.10, and specialize them to the linear case we get:

THEOREM 1.2.2 *A solution to the linear H_∞ control problem exists (and there are formulas for producing it) if there exists solutions $X \geq 0$ and $Y > 0$ to the DGKF Riccati equations which satisfy strict coupling*

$$-\gamma^2 Y^{-1} + X < 0. \tag{1.13}$$

Conversely, if a solution to the linear H_∞ control problem exists the stabilizing solutions X_e and Y_e to the Riccati equations are nonnegative definite and if $Y_e > 0$, we have

$$-\gamma^2 Y_e^{-1} + X_e < 0. \tag{1.14}$$

Note that the lower bounding properties

$$X_e \leq X \quad \text{and} \quad Y_e \leq Y$$

of X_e, Y_e imply

$$-\gamma^2 Y_e^{-1} + X_e \leq -\gamma^2 Y^{-1} + X < 0.$$

So the DGKF Theorem 1.2.1 has for simplicity presented the extreme case of the possible solutions.

As we soon see this funny way of writing the X, Y coupling condition is exactly the way it presents itself for general nonlinear systems. Also we have only discussed $Y > 0$. Actually, for the theory to hold Y need not be invertible. This may sound like a fine point but a rank one or two Y contains much less information than a rank seventeen Y, and such economies of information translate into major computational savings in the nonlinear case. Thus in the book we give considerable attention to the "singular cases", that is where Y^{-1} is "not finite".

1.3 The Idea of the Nonlinear Solution

This section is the heart of the introductory outline of the book, and contains a discussion of the main ideas of the solution to the nonlinear H_∞ control problem defined above.

State feedback control. The nature of the information available to the controller has a very significant bearing on the complexity of the problem, and of the resulting controller. Accordingly, we begin with an easier problem in which the controller is allowed to read the state of the plant. This simpler problem is known as the state feedback H_∞ control problem (essentially one with full information), and is well-understood in the literature.

Estimation - the information state. Next we turn to the general output feedback problem. Here the state is not known perfectly and so we must estimate it. This estimation is done with something called the information state, a function on the state space of the plant G which satisfies a PDE. Thus the information state is produced by an infinite dimensional controlled dynamical system. Much of this book is concerned with properties of this dynamical system and how it can be used to solve the H_∞ control problem.

Coupling - information state feedback. Using the information state, the output feedback problem is converted to state feedback problem for a new system. This new system uses the information state as its state variable, and the solution of the new state feedback problem leads to the solution of the output feedback H_∞ control problem. This is a coupling of control and estimation.

This indicates the layout of the remainder of the Introduction.

1.3.1 The State Feedback Control Problem

The state feedback H_∞ has been extensively studied in the literature and is well understood; see [vdS96] and the references contained therein.

1.3.1.1 Problem Statement

A block diagram illustrating the state feedback H_∞ problem is given in Figure 1.3. The state space model for the plant is

$$G : \begin{cases} \dot{x} = A(x) + B_1(x)w + B_2(x)u \\ \\ z = C_1(x) + D_{12}(x)u. \end{cases} \tag{1.15}$$

The controller can read measurements of the plant state x, so that

$$u = K(x).$$

(For simplicity we only consider static state feedback controllers. Alternatively, one could work with full information controllers, where K is a causal function of the disturbance, and this would yield the same optimal controller, under appropriate regularity assumptions.)

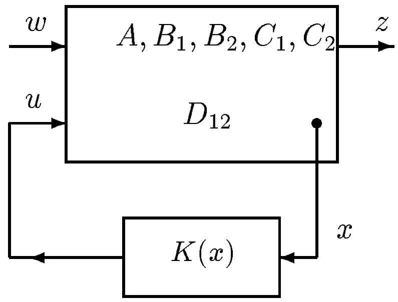

Figure 1.3: The State Feedback Closed-Loop System (G, K).

The state feedback H_∞ control problem is to find a controller $u = K(x)$ which is dissipative in the sense of §1.1 and stable in the sense that the vector field

$$A + B_2 K \text{ is } \textit{asymptotically stable}.$$

1.3.1.2 Problem Solution

The solution is determined by the *state feedback Hamilton-Jacobi-Bellman-Isaacs PDE* (HJBI PDE)

$$\begin{aligned} &\nabla_x V \cdot (A - B_2 E_1^{-1} D_{12}' C_1) + \tfrac{1}{2} \nabla_x V \cdot (\gamma^{-2} B_1 B_1' - B_2 E_1^{-1} B_2') \cdot \nabla_x V' \\ &\quad + \tfrac{1}{2} C_1' (I - D_{12} E_1^{-1} D_{12}') C_1 = 0. \end{aligned} \tag{1.16}$$

If one can find a strictly positive proper smooth solution V ($V(x) > 0$ if $x \neq 0$, $V(0) = 0$) which makes the vector field

$$A - B_2 E_1^{-1} D_{12}' C_1 + (\gamma^{-2} B_1 B_1' - B_2 E_1^{-1} B_2') \nabla_x V' \text{ asymptotically stable}$$

then a solution to the state feedback problem is:

1.3.1.3 The State Feedback Central Controller

$$K^*(x) = \mathbf{u}^*_{state}(x) = -E_1(x)^{-1} [D_{12}(x)' C_1(x) + B_2(x)' \nabla_x V(x)']$$

Using this controller, the closed loop system (G, K_s^*) becomes

$$\begin{cases} \dot{x} = A(x) + B_1(x) w - B_2(x) E_1(x)^{-1} [D_{12}(x)' C_1(x) + B_2(x)' \nabla_x V(x)'] \\ z = C_1(x) - D_{12}(x) E_1(x)^{-1} [D_{12}(x)' C_1(x) + B_2(x)' \nabla_x V(x)'] \end{cases}$$

and integration of the PDE (1.16) yields the dissipation inequality

$$V(x(t)) + \tfrac{1}{2}\int_0^T |z(t)|^2 dt \le \gamma^2 \tfrac{1}{2}\int_0^T |w(t)|^2 dt + V(x(0)).$$

The desired dissipation property (1.2) follows from this on setting $\beta = V$ since $V \geq 0$. Also, stability of the vector field $A + B_2K^*$ follows (see, e.g., [vdS96]).

It is important to note for practical reasons that the designer can solve for V *off line* (i.e. not in real time). This requires the solution of a PDE in n-dimensions.

EXAMPLE 1.3.1 For linear systems (cf. §1.2), the state feedback HJBI PDE has a quadratic solution

$$V(x) = \tfrac{1}{2}x'Xx$$

if it has any solution at all. One can substitute $\nabla_x V(x) = x'X$, into the HJI, which we illustrate when $D_{12} = I = D_{21}$, and get

$$\begin{aligned}
0 \;&=\; \tfrac{1}{2}x'(A - B_2C_1)'\nabla V(x)' + \tfrac{1}{2}\nabla V(x)(A - B_2C_1)x \\
&\quad + \tfrac{1}{2}\nabla V(x)(\gamma^{-2}B_1B_1' - B_2B_2')\nabla V(x)' \\
&=\; \tfrac{1}{2}[\; x'(A - B_2C_1)'Xx + x'X(A - B_2C_1)x \\
&\quad + x'X(\gamma^{-2}B_1B_1' - B_2B_2')Xx \;] \\
\text{i.e. } 0 \;&=\; \tfrac{1}{2}x'\mathcal{R}_{state}(X)x
\end{aligned}$$

which is the DGKF state matrix Riccati equation $\mathcal{R}_{state}(X) = 0$. Take $X_e \geq 0$ to be the stabilizing solution of this Riccati equation to get the optimal state feedback controller

$$K^*(x) = \mathbf{u}^*_{state}(x) = -[C_1 + B_2'X_e]x.$$

∇

REMARK 1.3.2 Actually to solve the state feedback H_∞ problem it is enough to find a function $V(x) \geq 0$, $V(0) = 0$ satisfying the HJBI PDE (1.16) plus a detectability assumption. For example, if the closed loop system (G, K_V^*) is detectable (here $K_V^*(x)$ is determined by the solution $V(x)$: $K_V^*(x) = -E_1(x)^{-1}[D_{12}(x)'C_1(x) + B_2(x)\nabla_x V(x)]$), then one can obtain stability of $A + B_2K_V^*$ from the dissipation inequality analogously to what is done in Theorem 2.1.3. This approach will be used frequently in the sequel. Note, however, that it is in general difficult to check detectability; however, the generic system is detectable (which of course does not imply that a system derived from some generic optimization process is detectable). Another addition to solutions of the HJBI which produce solutions to the state feedback control problem is the class of strict positive V with $V(0) = 0$ solving the strict HJBI inequality

$$\nabla_x V \cdot (A - B_2 E_1^{-1} D_{12}' C_1) + \tfrac{1}{2} \nabla_x V \cdot (\gamma^{-2} B_1 B_1' - B_2 E_1^{-1} B_2') \cdot \nabla_x V'$$
$$+ \tfrac{1}{2} C_1' (I - D_{12} E_1^{-1} D_{12}') C_1 \leq 0.$$

∇

1.3.2 The Information State

We return to the output feedback problem. To solve it, we use an information state. This converts the output feedback problem to a new state feedback problem with a new state, namely the information state (this methodology is an old one from stochastic optimal control). We now give definitions which lead directly to the construction of the controller dynamics.

1.3.2.1 Reversing Arrows

We start by defining the *reverse arrow system*. It is a new system $\tilde{G}$ which, in two block case, is obtained from G by reversing the w and y arrows. While the definition is algebraic, pictures help a lot, see Figures 1.4 (a) and (b).

The *reverse arrow system* is defined by

$$\tilde{G} : \begin{cases} \dot{\xi} = A(\xi)^{\times} + B_1(\xi) y + B_2(\xi) u \\ z = C_1(\xi) + D_{12}(\xi) u \\ w = -C_2(\xi) + y. \end{cases} \tag{1.17}$$

with $A^{\times}$ defined by

$$A^{\times} \stackrel{\triangle}{=} A - B_1 C_2.$$

Note that G and $\tilde{G}$ have the same state space. Clearly this is derived by substituting $w = y - C_2(x)$ into the G dynamics to produce

$$\dot{x} = A(x) + B_1(x)[y - C_2(x)] + B_2(x) u$$

which is the same as the dynamics defined above in (1.17) for $\tilde{G}$.

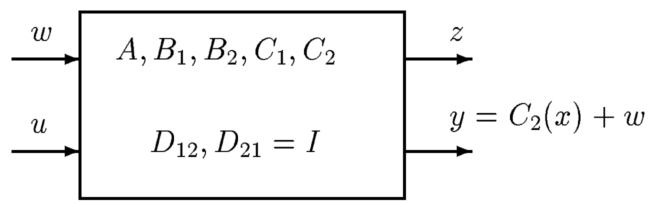

(a) The Original System G.

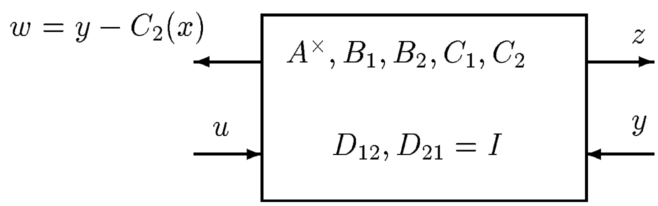

(b) The Reverse Arrow System $\tilde{G}$.

Figure 1.4: Reversing arrows.

1.3.2.2 Definition

Given time $t \geq 0$, past measurement $y \in L_2[0,t]$ and past control signal $u \in L_2[0,t]$ introduce a function $p_t(x) = p(x,t)$ on the states x of the plant G by

$$\begin{aligned} p_t(x) &= p_0(\xi(0)) + \tfrac{1}{2}\int_0^t [|C_1(\xi(s)) + D_{12}(\xi(s))u(s)|^2 - \gamma^2 |y(s) - C_2(\xi(s))|^2] ds \\ &= \text{initial state energy } + \int_0^t [\text{output}^2 - \text{input}^2]\, ds \\ &\qquad \text{given } u, y \text{ over } [0,t] \text{ and given } \xi(t) = x. \end{aligned} \tag{1.18}$$

Here ξ follows the state trajectory from 0 to t of the reverse arrow system (1.17) with *final* state $\xi(t) = x$. This is the tricky part. Given x to define $p_t(x)$ we must run the $\tilde{G}$ system backwards for t time units, using the given u, y. We see how much energy $E_{[0,t]}$ was consumed by the system and what state $\xi(0)$ the trajectory hits, with energy E_0. Then

$$p_t(x) = E_0 + E_{[0,t]},$$

the sum of two cost terms. This function is called the *information state* for the H_∞ control problem and plays the role of a "sufficient statistic", [JBE94], [JB95]. In [BB89], this function is called the "cost to come". For $p_t = p(\cdot, t)$ to be defined everywhere we must assume that for each $u, y \in L_2[0,t]$ the differential equation (1.1) has trajectories whose end points at time t sweep out the whole state space provided the endpoints at time 0 do. (See Theorem 3.1.8.) Figure 1.5 illustrates some of the common shapes of information states; generally they are bounded above, point downwards, and may take the value $-\infty$.

If the information state $p_t(x)$ is smooth then it satisfies the *information state PDE*

$$\begin{aligned} \frac{\partial p_t(x)}{\partial t} &= F(p_t, u(t), y(t)) \stackrel{\triangle}{=} -\nabla_x p_t(x) \cdot [A^\times(x) + B_2(x)u(t) + B_1(x)y(t)] \\ &\quad + \tfrac{1}{2}|C_1(x) + D_{12}(x)u(t)|^2 - \gamma^2 \tfrac{1}{2}|y(t) - C_2(x)|^2 \end{aligned} \tag{1.19}$$

which is readily obtained by differentiating (1.18). Often we write this differential equation even when p is not smooth, but one should interpret it as the integral equation (1.18), or perhaps in the viscosity sense (see Appendix B). We can think of the PDE (1.19) as describing an *infinite dimensional dynamical system*, which can be written in shorthand form

$$\dot{p} = F(p, u, y).$$

This system has a "state" p belonging to an infinite dimensional function space, and it is driven by input signals u and y. The solution of the H_∞ problem depends on properties of this system. Some references include [JB95], [BB89], [DBB93].

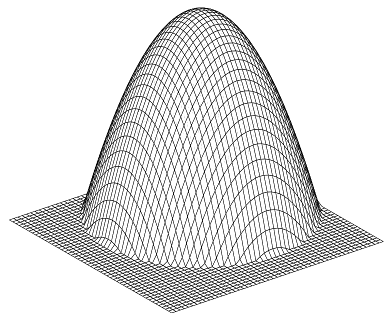

(a) Nonsingular (everywhere finite).

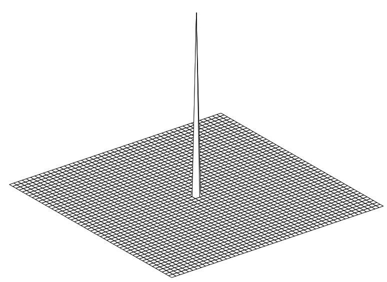

(b) Purely Singular (equal to $-\infty$ everywhere except at $x = 0$).

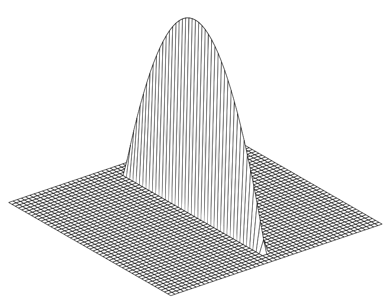

(c) Mixed Singular (finite on a subset of $\mathbf{R}^n$ and equal to $-\infty$ elsewhere).

Figure 1.5: Common Information States. (The information state is a function $p(x)$ defined on the plant state space, the horizontal plane in the figure, with coordinates $x = (x_1, \ldots, x_n)$.)

1.3.3 The Central Controller

We now give a high level formula for the structure and dynamics of a controller which (as the book unfolds) turns out to be a good candidate for the solution of the H_∞ problem.

State Space is a space $\mathcal{X}_e$ of functions $p = p(x)$ on the state space $\mathbf{R}^n$ of the plant G.

Dynamics are the PDE

$$\dot{p} = F(p, u, y).$$

Output u is function

$$u = \mathbf{u}(p)$$

defined (on a subset of) the information state space $\mathcal{X}_e$.

We call this the *information state controller*, illustrated in Figure 1.6.

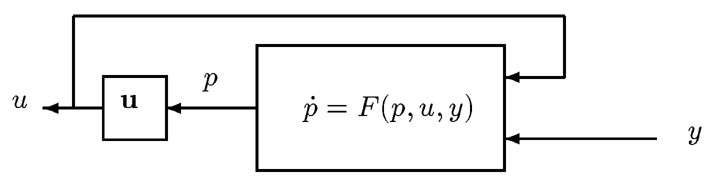

Figure 1.6: Information State Controller

In §1.3.5 we show how $\mathbf{u}$ is constructed. Ultimately we shall focus on particular information state controller called the central controller; it is obtained by optimization (yielding $\mathbf{u}^*(p)$) and suitable initialization.

An important point is that for this controller to be implementable *one must solve the information state PDE online*. This is a PDE in n-dimensions.

1.3.4 Equilibrium Information States

Our definition of the controller dynamics is not complete because in order to define its dynamics we must specify an initial information state p_0. As we shall see careful choice of this initial state p_0 makes a big difference in the implementability of the controller and strongly affects the dynamical behavior. Thus we devote substantial effort and several subsections to the question: *which initial state p_0 do we use?*

An obvious requirement of p_0 stemming from the null initializing property $K(0) = 0$ is

$\mathbf{u}(p_t) = 0$ for all t when p_t solves $\dot{p} = F(p, 0, 0)$ initialized at p_0.

However a stronger highly desirable condition is

$$0 = F(p_0, 0, 0) \ \text{ and } \ \mathbf{u}(p_0) = 0.$$

That is, p_0 is said to be an *equilibrium solution* p_e to the information state PDE. This is the correct initialization of the central controller: $p_0 = p_e$. (Below we discuss convergence of p_t to p_e—stability.)

As we shall see the equilibria for two block information states have a surprising form. It is surprising enough that we had better retreat to an example before describing it; this is done in §1.3.6. In the meantime, we consider the problem of choosing the controller output function $\mathbf{u}(p)$.

1.3.5 Finding u^* and Validating the Controller

We give now some details on the construction of the function $\mathbf{u}(p)$ which is a key component of the information state and central controllers. This is chosen optimally as follows (so that we will take $\mathbf{u} = \mathbf{u}^*$): solve an infinite dimensional state feedback control problem. The HJBI PDE for this problem is

$$\inf_{u \in \mathbf{R}^m} \sup_{y \in \mathbf{R}^p} \nabla_p W(p)[F(p, u, y)] = 0 \ \text{ in } \mathrm{dom} W. \tag{1.20}$$

Here, $\nabla_p W(p)$ is interpreted as a Frechet derivative (more general interpretations are discussed in Chapters 4 and 10). One attempts to solve this PDE for a smooth function $W(p)$ defined on a domain $\mathrm{dom} W$, a subset of the state space, and satisfying auxiliary conditions such as $W(p) \geq \sup_x \{p(x)\}$, and $W(p_0) = 0$ for some $p_0 \in \mathrm{dom} W$. The function $W(p)$ is called the *value function* for the H_∞ control problem, and can be regarded as an analog of the state feedback value function $V(x)$ (see §1.3) for the information state system. The information state feedback function $\mathbf{u}^*(p)$ is obtained by

$$\begin{aligned} \mathbf{u}^*(p) &= \operatorname*{argmin}_{u \in \mathbf{R}^m} \{ \sup_{y \in \mathbf{R}^p} \nabla_p W(p)[F(p, u, y)] \} \\ &= (\nabla_p W(p)[E_1])^{-1} \nabla_p W(p)[-D_{12}' C_1 + B_2' \nabla_x p']. \end{aligned}$$

Necessary and sufficient conditions for the solvability of the H_∞ control problem can be expressed in terms of the function $W(p)$ and the PDE (1.20). The following "metatheorem" states the main idea without the clutter of technical details:

RESULT 1.3.3 *If there exists some controller which solves the H_∞ problem, then there exists a function $W(p)$ solving the PDE (1.20) (in some sense) as well as auxiliary technical conditions. If the function $W(p)$ is smooth, then the central controller $K_{p_e}^*$ obtained from $W(p)$ solves the H_∞ problem. Key to this is the "coupling condition"ensuring that the controller is well-defined for all time and along trajectories of the closed-loop system,*

$$p_t(x(t)) \le \sup_x \{p_t(x)\} \le W(p_t) \le W(p_e) < +\infty, \qquad (1.21)$$

where $u(t) = \mathbf{u}^*(p_t)$. Conversely, if one can solve the PDE (1.20) for a smooth function $W(p)$ satisfying some auxiliary technical conditions, then the central controller $K_{p_e}^*$ obtained from $W(p)$ solves the H_∞ problem.

The major objective of this book is to present intuition and theory for results of this type.

1.3.5.1 Construction of the Central Controller

Now we summarize the procedure for building the central controller:

(i) Obtain a function $W(p)$ and $\mathbf{u}^*(p)$ solving the PDE (1.20) and the coupling (1.21).

(ii) Compute p_e and check $\mathbf{u}^*(p_e) = 0$.

(iii) Use $\mathbf{u}^*$ as the output term of the central controller.

(iv) The information state PDE (1.19) initialized at $p_0 = p_e$ gives the dynamics of the controller $K_{p_e}^*$.

1.3.5.2 Validating the Controller

We review the context in which we sit. Let $w(\cdot) \in L_2$ and $x_0 \in \mathbf{R}^n$ be given. These determine signals $y(\cdot)$, $z(\cdot)$, and $u(\cdot)$ and trajectories $x(\cdot)$, $p_.$ from the dynamics of the closed loop $(G, K_{p_e}^*)$ with $p_0 = p_e$, $u(\cdot) = \mathbf{u}^*(p_.)$. The idea behind confirming dissipativity of the closed loop system is:

(i) Integrate the PDE (1.20) along the trajectory p_t:

$$W(p_t) \le W(p_0) = W(p_e) < +\infty.$$

Then use the property

$$p_t(x(t)) \le \sup_x \{p_t(x)\} \le W(p_t) \le W(p_e) = 0 \tag{1.22}$$

and the definition of the information state to obtain

$$p_e(\xi(0)) + \tfrac{1}{2}\int_0^t [|C_1(\xi(s)) + D_{12}(\xi(s))u(s)|^2 - \gamma^2|y(s) - C_2(\xi(s))|^2]\, ds$$

$$= p_t(x) \le W(p_e)$$

where $\xi(\cdot)$ is the solution of (1.17) with $\xi(t) = x$. Now if $w(\cdot)$ is input to the plant G with initial state x_0 we obtain signals $u(\cdot) = K^*_{p_e}(y(\cdot))$ and state $x(\cdot)$ in closed-loop, and so if we set $x = x(t)$ we have $\xi(\cdot) = x(\cdot)$ and so

$$\tfrac{1}{2}\int_0^t [|C_1(x(s)) + D_{12}(x(s))u(s)|^2 - \gamma^2|w(s)|^2]ds \le W(p_e) - p_e(x(0))$$

$$= -p_e(x_0)$$

which is the dissipation inequality (1.2) with $\beta = -p_e$.

(ii) If p_e is nonsingular and if $(G, K^*_{p_e})$ is detectable, then $K^*_{p_e}$ solves the H_∞ control problem. If p_e is singular, then with extra work and stronger conditions it is possible to prove that $K^*_{p_e}$ solves the H_∞ control problem. See Chapter 4.

(iii) The stability results discussed in §1.5 below for the information state system are used to deduce the asymptotic behavior of the information state in closed loop (Chapter 4).

1.3.5.3 Storage Functions

Associated with a dissipative system are functions $e(x,p)$ on its state space called *storage functions*. Of course we are interested in the closed loop system $(G, K^*_{p_e})$ and a storage function e for it is defined to be non-negative and satisfy the "dissipation inequality":

$$e(x(t), p(t)) + \tfrac{1}{2}\int_0^t |z(s)|^2\, ds \le e(x_0, p_e) + \gamma^2 \tfrac{1}{2}\int_0^t |w(s)|^2\, ds \tag{1.23}$$

for all $t \ge 0$ and all $w \in L_2[0,t]$. It is fairly remarkable that there is a storage function $e(x,p)$ for the closed loop system $(G, K^*_{p_0})$ which has a very simple and explicit formula:

$$e(x,p) = -p(x) + W(p).$$

It is interesting to note that the content of (1.23) is the same as that of (1.22) as can be verified by adding minus the information state equation which p_t satisfies to (1.21). Also compare (1.23) with the dissipation inequality (1.2) of §1.1 (note $\beta(x) = e(x, p_e) = -p_e(x)$). This storage function gives a handy tool for validating that $K^*_{p_0}$ is γ-dissipative provided $p_0(x_0)$ is finite.

1.3.6 Example: Linear Systems

1.3.6.1 $W(p)$ for linear systems

The information state for linear systems are quadratic and will be described immediately below. For now we discuss the form of W. One has

$$W(p) = \max_x \{p(x) + \tfrac{1}{2}x'X_e x\}.$$

Thus

(i) The integrated form of (1.20)

$$\begin{aligned} \max_x\{p_t(x)\} &\le W(p_t) = \max_x\{p_t(x) + \tfrac{1}{2}x'X_e x\} \\ &\le W(p_0) = \max_x\{p_0(x) + \tfrac{1}{2}x'X_e x\} \end{aligned}$$

is equivalent to $W(p)$ being finite, and X_e being positive semidefinite.

(ii) The equilibrium information state is $p_e = -\gamma^2\frac{1}{2}x'Y_e^{-1}x$ where Y_e solves the DGKF Y equation. $W(p_e)$ finite is equivalent to the matrix $-\gamma^2 Y_e^{-1} + X_e$ being negative semidefinite. If γ is suboptimal, then this is negative definite since, small perturbations of this will be negative.

Thus we have that the DGKF conditions 1, 2 and 3 (except for the strictness) of §1.2.4 are implied by the existence of (finite) W and the existence of p_e solving (1.20) and (1.21). The converse is true and can be checked with a little effort.

1.3.6.2 The Information State

For linear systems, one can check that if $Y(t)$ is invertible, then solutions to the information state equation have the form

$$p_t(x) = -\gamma^2\tfrac{1}{2}[x - \hat{x}(t)]'Y(t)^{-1}[x - \hat{x}(t)] + \phi(t)$$

whenever p_0 has this form, where

$$\begin{aligned} \dot{\hat{x}} &= [A^\times + Y(\gamma^{-2}C_1'C_1 - C_2'C_2)]\hat{x} + \\ &\quad + [B_1 + YC_2']y + [B_2 + \gamma^{-2}YC_1'D_{12}]u \\ \dot{Y} &= A^\times Y + YA^{\times\prime} + Y(\gamma^{-2}C_1'C_1 - C_2'C_2)Y. \end{aligned}$$

Now we compare this to the dynamics of the DGKF central controller (1.12) to the linear H_∞ problem. The $\hat{x}$ equation is exactly (1.12) if we take $Y(t)$ equal to Y_e, the stabilizing solution to the DGKF Y equation (1.10). The above Riccati differential equation for $Y(t)$ can be initialized in many ways which lead to a solution to the H_∞ control problem. However, the equilibrium solution has the great advantage that $dY_e/dt = 0$ so we have no Y differential equation to solve in real time (since $Y(t) = Y_e$ for all $t \geq 0$).

Now comes a crucial pair of exercises. They are so crucial that the reader should think for a minute and not race to the answers:

> *Exercise 1.* Suppose $A^\times$ is a stable matrix. What is Y_e?

Answer: $Y_e = 0$. The reason is that the DGKF Y equation is homogeneous in Y, so $Y_e = 0$ certainly satisfies it. But is it stabilizing? Well yes since $A^\times$ is stable even without perturbing it.

> *Exercise 2.* $A^\times$ has no pure imaginary eigenvalues. What can we say about Y_e?

Answer: Y_e is 0 on the stable eigenspace of $A^\times$.

In the first exercise the DGKF Y equation disappears when $Y_0 = Y_e$ since the stabilizing solution Y_e is zero, so the controller formulas will only involve the DGKF X equation. In the second exercise Y_e is usually low rank, so maybe *the controller will have a low dimension* (in some sense) if we initialize $Y_0 = Y_e$. For the nonlinear case this suggests a big simplification since Y determines the state estimator (the online part of the computation).

We return to the equilibrium information state p_e, which in the linear case is formally of the form

$$p_e(x) = -\gamma^2 \tfrac{1}{2} x' Y_e^{-1} x,$$

and immediately worry because Y_e is typically not invertible. Indeed if $Y_e = 0$ we suspect that $p_e(x) = -\infty$. While this is close to correct it is not quite and so we now embark on definitions and a discussion of singular functions. Later we give precise formulas for singular information states and resulting controllers.

1.4 Singular Functions

When Y_e is not of full rank, the function $p_e(x) = -\gamma^2 \frac{1}{2} x' Y_e^{-1} x$ interpreted as a singular function. In the first exercise $A^\times$ is stable and corresponds to $Y_e = 0$, so we define then

$$p_e = \delta_0,$$

where

$$\delta_0(x) = \begin{cases} 0 & \text{if } x = 0, \\ -\infty & \text{if } x \neq 0. \end{cases}$$

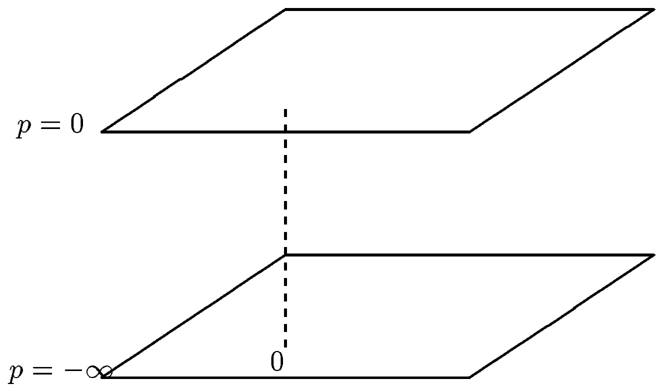

Figure 1.7: The singular function δ_0.

In the second exercise

$$p_e = \delta_{M_{as}} + \breve{p}_e.$$

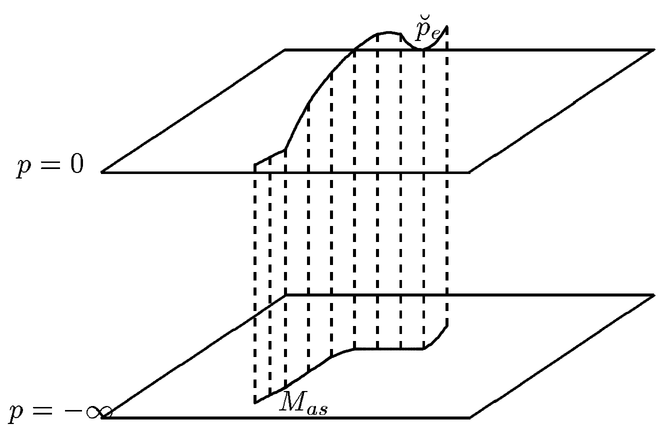

Figure 1.8: The singular function $p_e = \delta_{M_{as}} + \breve{p}_e$.

Here

$$\delta_{M_{as}}(x) = \begin{cases} 0 & \text{if } x \in M_{as}, \\ -\infty & \text{if } x \notin M_{as} \end{cases}$$

and M_{as} is the antistable subspace of $A^{\times}$ and $\breve{p}_e$ is a quadratic form on M_{as} (the analogous notation $p = \delta_M + \breve{p}$ will be used frequently where $M \subset \mathbf{R}^n$ and $\breve{p}$ is a function defined on M).

We emphasize again that in mixed sensitivity control applications M_{as} *is usually low dimensional!* Thus $p_0 = p_e$ is supported on a very thin set.

1.4.1 Singular Equilibrium Information States

For the nonlinear two block problem we shall assume that

$$A^{\times} \text{ is a hyperbolic vector field}$$

with global stable M_s and antistable M_{as} submanifolds. As we shall see the equilibrium information state p_e is given by

$$p_e = \delta_{M_{as}} + \breve{p}_e$$

where M_{as} is the *antistable submanifold* of $A^{\times}$ and $\breve{p}_e$ is a smooth function on M_{as}. There are two important special cases:

(i) when $A^{\times}$ is stable, $p_e = \delta_0$ is a purely singular function, and

(ii) when $A^{\times}$ is antistable, p_e is a finite, smooth, nonsingular function.

In any case once M_{as} is computed $\breve{p}_e$ can be determined by computing for each $x \in M_{as}$ the integral

$$\breve{p}_e(x) = \tfrac{1}{2}\int_{-\infty}^{0} [|C_1(\xi(s))|^2 - \gamma^2|C_2(\xi(s))|^2]ds,$$

where $\xi(\cdot)$ is the solution in backward time to

$$\dot{\xi} = A^{\times}(\xi), \quad \xi(0) = x \in M_{as}.$$

See §3.2 for a derivation.

1.4.2 The Central Controller Dynamics

The *central controller* is obtained by initializing the optimal information state controller ($\mathbf{u} = \mathbf{u}^*$) at equilibrium, $p_0 = p_e$, and is denoted $K^*_{p_e}$. If p_e is singular, the resulting information state dynamics is still quite concrete to write down, manipulate, and compute numerically.

The formula for the dynamics.

Suppose $p_0 = p_e = \delta_{M_{as}} + \breve{p}_e$. Then for any $u, y \in L_2,\ p_t$ is of the form

$$\delta_{M_t} + \breve{p}(\cdot, t)$$

where

$$M_t \stackrel{\triangle}{=} \{\xi(t)\ :\ \dot{\xi} = A^{\times}(\xi) + B_1(\xi)y + B_2(\xi)u,\ 0 \le s \le t, \xi(0) \in M_{as}\},$$

and also $\breve{p}(\cdot,t)$ is the function on M_{as} given by

$$\breve{p}(x,t) \;=\; \breve{p}_e(\xi(0)) + \frac{1}{2}\int_0^t [|C_1(\xi(s)) + D_{12}(\xi(s))u(s)|^2 - \gamma^2|y(s) - C_2(\xi(s))|^2]ds$$

$$= \text{ initial state energy} + \text{the energy it takes to get from } \xi(0) \text{ to } \xi(t) = x$$

where $\xi(s)$, $0 \le s \le t$, is given by the reversed system dynamics (1.17) with $\xi(t) = x \in M_t$.

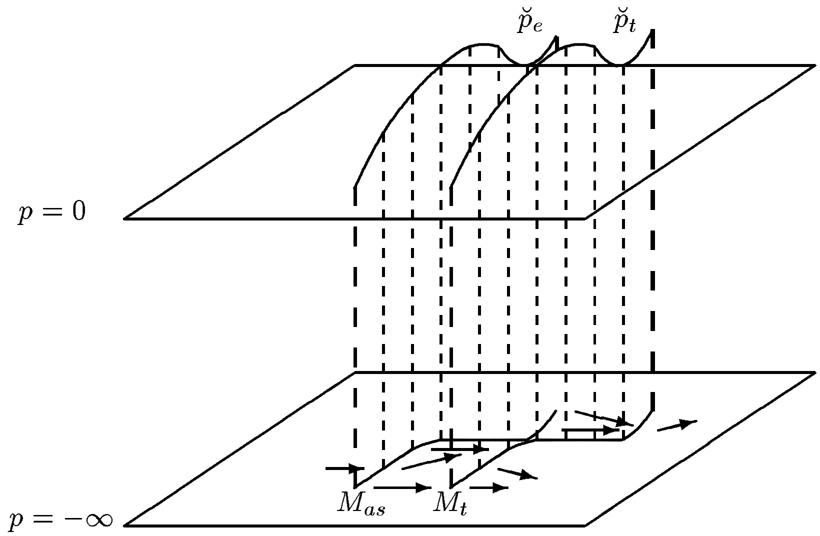

Figure 1.9: Flow of Singular Information States

This evolution of functions $\breve{p}_t$ on the $M \equiv \mathbf{R}^k$ together with the evolution (2.6) constitute a "reduced dimensional" picture of the compensator dynamics (to be discussed further in Chapter 10).

1.4.2.1 Computational Requirements

The amount of computation required for a singular function indeed is less than for a nonsingular function. This is because one need only solve a $k << n$-dimensional PDE (in real time). Intuitively, the singularity of information states reflects a degree of knowledge concerning the state trajectory, and this means that less computational effort is required.

To be more specific, suppose one wishes to approximately compute the compensator state $(M_t, \breve{p}_t)$ by numerically solving the ODE

$$\dot{x} = A^{\times}(x) + B_1(x)y + B_2(x)u$$

which propagates M_t. Then $\breve{p}_t$ is computed by evaluating the integral recursively, i.e., one only needs to update the integral at each time step. One begins such a numerical computation by laying out a grid on M_0. For example, if $M_0 = M_{as}$ is k dimensional and one chooses N equally spaced grid points in each dimension, then if M_0 is k dimensional choosing N^k grid points would be natural. One initializes the ODE at each grid point x_g and solves the ODE numerically as the values of $u(t)$ and $y(t)$ become known. Since the ODE is a $n \times n$ system at any time its solution is an n-dimensional vector; and we get one of these per grid point. Thus the memory and operation requirements scale like

$$n\mathrm{O}(N^k).$$

This is a striking improvement over the $\mathrm{O}(N^n)$ required to solve the PDE (3.9) for a smooth solution in $\mathbf{R}^n$.

1.5 Attractors for the Information State

A very important issue is whether or not the equilibrium p_e is an attractor for the information state dynamics (1.19). By this we mean roughly that there exist a set $\mathcal{D}_{attr}(p_e)$ of initial states p_0 and a nonempty set of signals u, y in L_2 for which the resulting information state trajectory p_t converges to p_e (in an appropriate sense). We call an equilibrium p_e a *control attractor* if for all u, y in L_2

$$p_t \to p_\infty = p_e + c \quad \text{as} \;\; t \to \infty$$

for some constant $c \in \mathbf{R}$ (depending on u, y, p_0), for $p_0 \in \mathcal{D}_{attr}(p_e)$.

For the nonlinear two block problem with $A^\times$ hyperbolic we shall prove roughly:

> *Suppose $A^\times = A - B_1C_2$ is hyperbolic, that the initial function p_0 satisfies $F(p_0, 0, 0) < 0$, and that certain technical conditions hold. If u, y are supported on $[0, T]$, then the solution p_t to*
>
> $$\dot{p} = F(p, u, y)$$
>
> *with smooth initial condition p_0 has the stability property*
>
> $$p_t \to p_e + c \stackrel{\triangle}{=} \delta_{M_{as}} + \breve{p}_e + c \;\; \text{as}\; t \to \infty$$
>
> *where M_{as} is the antistable manifold of $A^\times$, and $\breve{p}_e$ is a function on M_{as} and c is a real number depending on u, y and p_0.*

It is important to note that the special case $A^\times$ antistable leads to nonsingular equilibria p_e, whereas when $A^\times$ is stable $p_e = \delta_0$ is purely singular. Moreover, stronger results are proven (Chapter 11):

- With restrictions on p_0 it is possible to prove convergence for arbitrary u, y in L_2.

- Initial states p_0 in $\mathcal{D}_{attr}(p_e)$ are characterized.

Various notions of convergence are required, but these are mainly confined to the more theoretical sections of the book. In particular, we make use of the max-plus notion of weak convergence (analogous to weak convergence of probability measures), and hypoconvergence (from optimization); see Appendix C.

1.6 Solving the PDE and Obtaining $\mathbf{u}^*$

The value function $W(p)$ and the infinite dimensional PDE (1.20) offer a high level framework for solving nonlinear H_∞ control problems. The function PDE (1.20) is to be solved *offline* for $W(p)$, and hence $\mathbf{u}^*(p)$ can be constructed offline. As mentioned, the only online part of the central controller $K^*_{p_e}$ is the information state dynamics (1.19). We now discuss three situations in which $W(p)$ vastly simplifies.

The significance of these results is that the infinite dimensional PDE (1.20) can be solved in terms of a PDE on a finite dimensional space (i.e. one on $\mathbf{R}^n$). Solving such PDE is a traditional pursuit of mathematics and engineering and it bears directly on the (offline) construction of the central controller. Solving these PDE give formulas for $\mathbf{u}^*$ in terms of the optimal state feedback $\mathbf{u}^*_s$ control law applied to carefully selected states. Of course, the PDE for the information state.

1.6.1 Certainty Equivalence

Under the certainty equivalence assumption, Whittle [Whi81], Basar-Bernhard [BB89], it is possible to use the function

$$\hat{W}(p) = \sup_x \{p(x) + V(x)\}$$

as a value function for the H_∞ control problem. Here, $V(x)$ is the state feedback value function of §1.3 which determines the state feedback controller $u^*_{state}(x)$. The *certainty equivalence assumption* requires that the *minimum stress estimate*

$$\bar{x}(t) = \underset{x\in\mathbf{R}^n}{\operatorname{argmax}} \{p_t(x) + V(x)\}$$

is unique. If this assumption holds, then the certainty equivalence controller

$$u^{ce}(t) = \mathbf{u}^*_{state}(\bar{x}(t))$$

coincides with the central controller described above. Further, the function $\hat{W}(p)$ solves the PDE (1.20).

The certainty equivalence controller has dynamics

$$
\begin{aligned}
\dot{\bar{x}} &= A - B_2E_1^{-1}D_{21}'C_1 + (\gamma^{-2}B_1B_1' - B_2E_1^{-1}B_2')\nabla_x V' \\
&+(B_1D_{21}' - \gamma^2(\nabla_x^2 r_t)^{-1}\nabla_x y_s^{*\prime})(y - y_s^*) \qquad (1.24)\\
&+(B_2 + (\nabla_x^2 r_t)^{-1}\nabla_x u_s^{*\prime}E_1)(u - u_s^*),
\end{aligned}
$$

where the RHS is evaluated at $x = \bar{x}(t)$, and $r = r(x,t)$ solves a PDE obtained by combining the PDEs for $p_t(x)$ and $V(x)$ (see Chapter 7). This yields the concrete formula for the u^*:

$$\mathbf{u}^*(p) = \mathbf{u}^*_{state}(\bar{x})$$

since $\bar{x} = \bar{x}(p)$.

REMARK 1.6.1 A generalization of certainty equivalence to cases of multiple maxima has been considered in [HV95]. ∇

1.6.2 Bilinear Systems

There are some classes of systems for which the information state is finite dimensional. Two such classes are those consisting of bilinear and linear systems. The plant is *bilinear* provided

$$
\begin{aligned}
&A(x) = Ax, \quad C_1(x) = C_1x, \;\; C_2(x) = C_2x, \\
&B_1(x) = B_1, \;\; B_2(x) = B_2 + B_3x, \quad D_{12}(x) = D_{12},
\end{aligned}
$$

where

$$A, \quad C_1,\; C_2, \quad B_1,\; B_2,\; B_3,\; D_{12},$$

are matrices of appropriate dimension, and we assume for simplicity that u is one dimensional ($m = 1$).

If $p_0(x) = -\gamma^2\frac{1}{2}x'Y_0^{-1}x$ with $Y_0 > 0$ then the information state is given explicitly by

$$p_t(x) = -\gamma^2\tfrac{1}{2}(x - \hat{x}(t))'Y(t)^{-1}(x - \hat{x}(t)) + \phi(t), \qquad (1.25)$$

where

$$
\begin{aligned}
\dot{\hat{x}} &= \big(A(u) - B_1C_2 + Y(\gamma^{-2}C_1'C_1 - C_2'C_2)\big)\,\hat{x} \\
&\quad +(B_1 + YC_2')y + (B_2 + \gamma^{-2}YC_1'D_{12})u, \qquad (1.26)\\
\dot{Y} &= (A(u) - B_1C_2)Y + Y(A(u) - B_1C_2)' \\
&\quad +Y(\gamma^{-2}C_1'C_1 - C_2'C_2)Y \qquad (1.27)\\
\dot{\phi} &= \tfrac{1}{2}|C_1\hat{x} + D_{12}u|^2 - \gamma^2\tfrac{1}{2}|y - C_2\hat{x}|^2. \qquad (1.28)
\end{aligned}
$$

Here, we have written $A(u) = A + B_2 u$. Thus the information state p_t projects to a finite dimensional triple $(\hat{x}, Y, \phi)$. Consequently, *the online computation of the information state is drastically simplified and feasible.*

If the value function is defined for the triple $(\hat{x}, Y, \phi)$, call it $\hat{W}(\hat{x}, Y, \phi)$, then corresponding PDE for $\hat{W}(\hat{x}, Y, \phi)$ is defined on a finite dimensional space $\mathbf{R}^n \times \mathbf{R}^{n^2} \times \mathbf{R}$ and has the form

$$\inf_{u \in \mathbf{R}^m} \sup_{y \in \mathbf{R}^p} \nabla_{(\hat{x}, Y, \phi)} \hat{W}(\hat{x}, Y, \phi)[\hat{F}(\hat{x}, Y, \phi, u, y)] = 0 \tag{1.29}$$

where $\hat{F}(\hat{x}, Y, \phi, u, y)$ denotes the dynamics defined by (1.26), (1.27), (1.28). Evaluating the infimum in the RHS of the PDE (1.29) yields the central controller function

$$\begin{aligned}
\mathbf{u}^*(\hat{x}, Y, \phi) &= \operatorname*{argmin}_u \{\sup_y \nabla_{(\hat{x}, Y, \phi)} \hat{W}(\hat{x}, Y, \phi)[\hat{F}(\hat{x}, Y, \phi, u, y)]\} \\
&= -E_1^{-1}[D_{12}' C_1 \hat{x} + (B_2' + \gamma^{-2} D_{12}' C_1 Y + \hat{x}' B_3') \nabla_{\hat{x}} \hat{W}(\hat{x}, Y, \phi)' \\
&\quad + \nabla_Y \hat{W}(\hat{x}, Y, \phi) \cdot (B_3 Y + Y B_3')]
\end{aligned}$$

The PDE (1.29) for a function $\hat{W}(\hat{x}, Y, \phi)$ can seldom be solved explicitly and so approximations and numerical methods must be used. However, *it is important to note that this is already feasible in applications where the state space is very low dimensional*. In general, it is not possible to solve this explicitly.

REMARK 1.6.2 For linear systems ($B_3 = 0$), the value function is given explicitly by

$$\hat{W}(\hat{x}, Y, \phi) = \tfrac{1}{2} \hat{x}' X_e (I - \gamma^{-2} Y X_e)^{-1} \hat{x} + \phi,$$

where $X_e \geq 0$ is a solution of the X Riccati equation (1.8). ∇

1.7 Factorization

While engineers have a deep love for feedback diagrams like Figure 1.1 this is not familiar to the average mathematician. Most mathematicians are, however, quite fond of factoring. They will try to factor numbers or mappings or most objects you put in front of them. Fortunately, the H_∞ control problem for the plant G in (1.1) is equivalent under various hypotheses to a type of factorization problem for the reversed arrow system $\widetilde{G}$ in (1.1) or more accurately because of possible degeneracies to what we call a decomposition of $\widetilde{G}$. To be more specific we start with a given system $\Sigma = \widetilde{G}$ and seek another system Σ^0 so that the composition $\Sigma^I = \Sigma \circ \Sigma^0$ is dissipative with respect to a certain signed bilinear form and so that Σ^0 satisfies a fairly weak partial left invertibility type of assumption. If Σ^0 is invertible, this is equivalent to Σ having the factorization $\Sigma = \Sigma^I \circ (\Sigma^0)^{-1}$.

Notice that if Σ is a system whose input space is $\mathbf{R}^{p+m}$, then the output space of Σ^0 is constrained to be $\mathbf{R}^{m+p}$, but its input space can be of any dimension. Traditionally, investigators found factors whose input space is $\mathbf{R}^{m+p}$, which if Σ^0 is linear means that its transfer function has values which are square matrices. The square case does not correspond precisely to the H_∞ control problem in Figure 1.1, but it can be used to parameterize many solutions to the problem; thus having a good square factoring is more than is needed to solve the control problem. The bulk of Chapter 8 treats square factorization. Actually *equivalent to the control problem* is having a good factor Σ^0 whose input space is $\mathbf{R}^p$. This is described in §8.10.

Factoring of various types as a subject independent of control is presented in the first and last parts of Chapter 8. The middle part of the chapter treats the connection between factoring and control. A mathematician with little interest in control could skip directly to the factoring chapter after reading the introduction. Much of it is self contained with only a few (key) proofs requiring machinery from the first of the book.

1.8 A Classical Perspective on H_∞ Control

Most people who learn H_∞ control these days for linear systems see state space problems and state space theory. In fact the subject began as a purely input-output frequency domain theory; H_∞ engineering began with amplifier design and later came into control and gained prominence there, see §1.11. In this section we sketch some of these ideas. We start with H_∞ control and then mention a few ideas and connections with broadband impedance matching, an ingredient of classical amplifier design.

1.8.1 Control

One is given a system (plant) P and wishes to find a controller K so that the closed loop transfer function

$$T(j\omega) \stackrel{\triangle}{=} P(j\omega)K(j\omega)(1 + P(j\omega)K(j\omega))^{-1}$$

of the system in Figure 1.2 has a certain "shape". The desired shape corresponds to the specs layed out in the control problem. A typical situation is illustrated by the Bode plot in Figure 1.10. It contains two plots which contain equivalent information but in different coordinates.

You see in the top picture of Figure 1.10 that the absolute value of $L \stackrel{\triangle}{=} PK$ must be bigger than the heavy line at low frequency and below the other heavy line at high frequency. At midrange frequencies there is a bit of flexibility so precise constraints are typically not drawn in. Algebraically the low and high frequency constraints are written as

$$|PK| \geq \gamma_\ell >> 1 \quad \text{low } \omega$$

$$|PK| \leq \gamma_h << 1 \quad \text{high } \omega$$

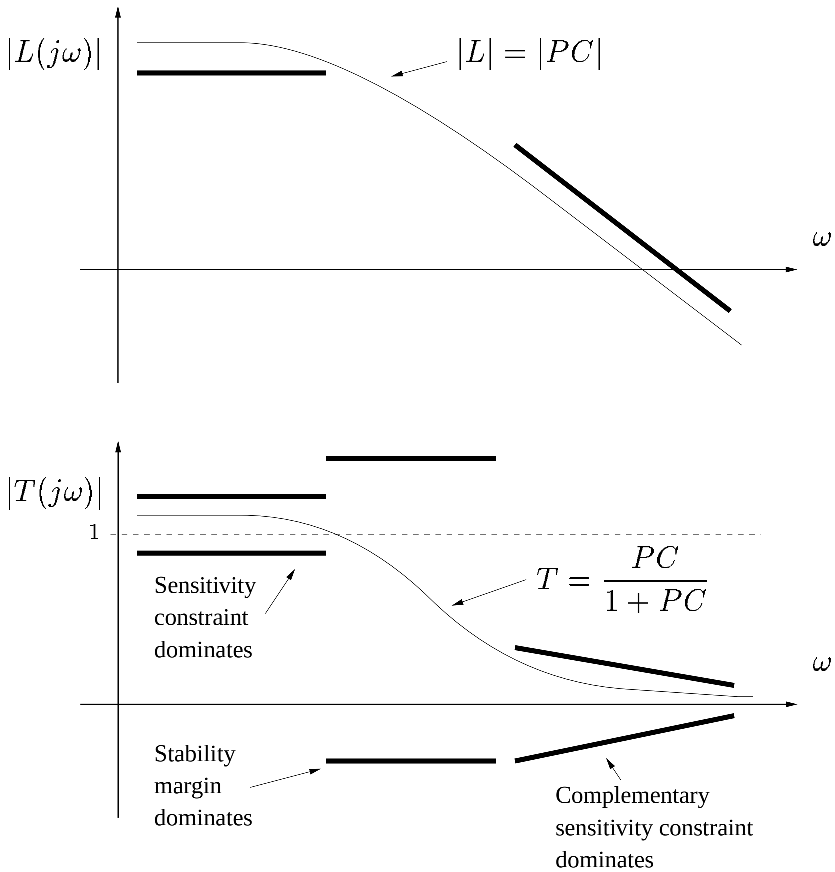

Figure 1.10: Bode diagram (transfer function magnitude vs. angular frequency).

where γ_ℓ and γ_h are given. The bottom figure contains the same information as the top figure but in terms of T, which we now see using simple algebra. At low frequency

$$|1 + PK| \geq \gamma_\ell - 1$$
$$|(1 + PK)^{-1}| \leq (\gamma_\ell - 1)^{-1}$$

so

$$|T - 1| = |PK(1 + PK)^{-1} - 1| = |(1 + PK)^{-1}| \leq (\gamma_\ell - 1)^{-1}$$

is small if γ_ℓ is large. At high frequency

$$|T| \leq \gamma_h(1 + \gamma_h)^{-1}$$

is small if γ_h is near 0. This high frequency constraint is often called the *rolloff* or bandwidth constraint.

We rephrase the constraint on T in the form

$$|T - \kappa|^2 W_\ell + |T|^2 W_h \leq R \tag{1.30}$$

where W_ℓ, W_h are positive weight functions and κ is a function which is 1 at low frequencies, zero at high frequencies and interpolates smoothly in between. Note that (1.30) contains a constraint on frequencies at midrange and the Bode plot above does not. Actually (1.30) constitutes a well posed problem while the Bode plot constraints do not. Adding midrange (like stability margin) constraints to the Bode plot gives a well posed problem. Note $W_h(j\omega) \to \infty$ as $\omega \to \infty$ to force the envelope containing T to pinch to zero at ∞.

We would like to show how the problem of finding a stable closed loop system meeting the constraint (1.30) translates to a familiar state space H_∞ problem. Actually there is a subtle issue on what we mean by the closed loop system being stable. Certainly we want T to have no poles in the closed right half plane (RHP), but we need in addition that small perturbations of P and K also have this property. This is one version of *internal stability* . We will not belabor this view point because that would be time consuming and because internal stability corresponds directly *to stability of the state space equations for the closed loop system* as was previously defined.

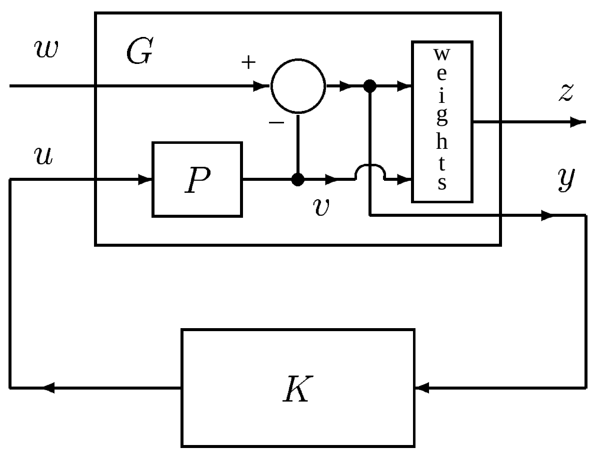

Figure 1.11: Mixed sensitivity embeds in the standard problem.

The next step in conversion of the H_∞ problem to state space form is embedding our H_∞ control problem in *the standard problem* described in §1.8. Figure 1.11 indicates how this is done. The transfer function $G(s)$ incorporates all information in the weights W_1, W_2 and plant P. One can read off the precise formula for G from Figure 1.11 and there is no reason to record it here, since we explicitly give the state space version of the formula in Chapter 9. Thus we have shown that our classical H_∞ problem is equivalent to finding K which makes the closed loop system in Figure 1.11 γ-dissipative, or equivalently internally stable with $w \to z$ transfer function having sup norm less than γ. Now Figure 1.11 has the form of Figure 1.1 and we see that the classical H_∞ control problem has the form of the standard problem of H_∞ control in §1.1

1.8.2 Broadband Impedance Matching

A basic problem in classical circuit theory is: *given an amplifying device, connect it with a passive circuit that produces a total (closed loop) amplifier which maximizes the worst gain over all frequencies.* An easier problem which often bears heavily on the amplifier problem is the broadband impedance matching problem: *Transfer as much power as possible from a given source to a passive (dissipative) load.*

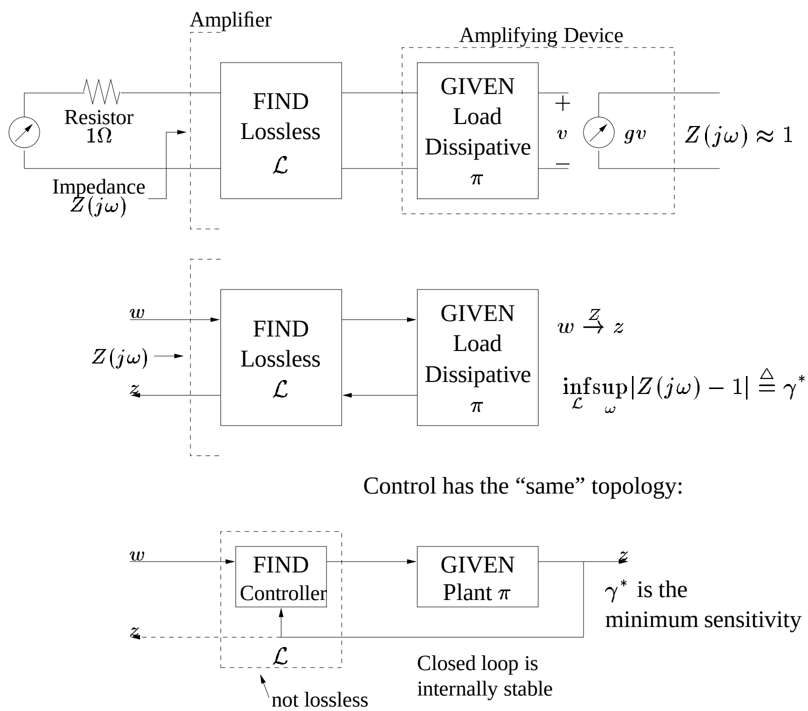

Figure 1.12: Amplifier

This problem is illustrated by Figure 1.12. The top picture shows an amplifier gain maximization problem. The middle picture illustrates the impedance matching problem associated with the amplifier problem. The last picture draws the middle picture in a way which looks much like Figure 1.2 and the classical control problem we discussed in §1.8.

One thing to mention is that the key tradeoff in impedance matching goes under the name *gain-bandwidth limitations*. They have been studied (under this name for decades) and the basic "rule of thumb" is the Bode–Fano integral constraints (an analog of the Freudenberg–Looze constraints of control theory). Gain–Bandwidth limitations are quite literal analogs of performance–rolloff constraints in control.

1.9 Nonlinear "Loop Shaping"

As mentioned in §1.8, in classical linear control the main objectives (in order of importance) are to make the controlled system

(i) stable,

(ii) have prescribed rolloff, and

(iii) achieve high performance at low frequencies.

A metaphor for their implications is that if we design an airplane that fails to be stable it will crash immediately, if rolloff is poor then it will probably crash eventually, and if performance is mediocre the plane will waste something, maybe fuel or time.

Controller design classically often consisted of choosing a candidate controller, and then checking the closed loop transfer function to see if it met given performance and rolloff specs; hence the term *loop shaping*. H_∞ control originated with the goal of making loop shaping more systematic. The H_∞ formalism involves weight selection, which is reasonably intuitive. Once sensible weights are picked solutions to the H_∞ problem often are not so far from desired that a few natural iterations gives a solution.

Of course there are serious tradeoffs between stability, rolloff, and performance constraints. While frequency and hence rolloff have no meaning for nonlinear systems, it is hard to believe that when systems are nonlinear that these tradeoffs dissappear. They must be important in some form.

What is nonlinear loop shaping? This is the subject of much current research and discussion, although the word loop shaping is not used. Indeed the issue is enough in flux that we do not presume to say anything definitive here. Our goal in this section is just to introduce a few issues. The main issues actually emerge in the state feedback problem, say for a system of the form

$$\begin{aligned} \dot{x} &= A(x) + B_2(x)u, \\ z &= C_1(x) + D_{12}(x)u, \end{aligned}$$

so we focus on state feedback in this presentation rather than on the more complicated measurement feedback problem (the next section considers this briefly).

Much attention goes to stabilizing a system and stability might be viewed as a type of performance constraint. This can be facilitated by solving a control Lyapunov inequality,

$$(A(x) + B_2(x)u) \cdot \nabla_x V(x) + \Omega(x) \leq 0 \tag{1.31}$$

where $\Omega \geq 0$. Given V and Ω this can be done explicitly. For example, for single input systems ($\dim u = 1$),

$$u(x) = \frac{-1}{B_2(x) \cdot \nabla_x V(x)} (A(x) \cdot \nabla_x V(x) + \Omega(x))$$

for $x \neq 0$.

More challenging is to decipher the analog or function in nonlinear control of rolloff constraints. Mathematically, rolloff constraints for a linear system look something like

$$\sup_{s=j\omega} \left| \frac{1}{(1+s)^n} T(s) \right| \leq k_\infty$$

or

$$\int_{\substack{j\omega \\ axis}} \left| \frac{1}{(1+s)^n} T(s) u(s) \right|^2 \leq k_2$$

In time domain terms these inequalities punish the size of

$$\frac{d^n}{dt^n} [\widehat{Tu}](t),$$

that is one has rate bounds on the output of the system. The output of the system is $C_1(x(t)) + D_{12}u(t)$ and a rate bound is implied by a rate bound on $u(t)$ and $x(t)$ separately. Thus it suffices to impose constraints of the form

$$|u| \leq k_1, \;\; |\dot{u}| \leq k_2, \;\; \text{and} \;\; x \in \mathcal{R},$$

where $\mathcal{R}$ is a carefully chosen region in state and input space.

In our discussion we shall focus on bounding $\dot{u}$, since this is an actuator rate bound and these are very common; so set $C_1 = 0$, $D_{12} = 1$, and $n = 1$. We begin by considering a rate saturation constraint $|\dot{u}| \leq 1$ and use the standard trick of making u a state and adding an input v to get

$$\dot{u} = v$$

with v meeting the saturation constraint $|v| \leq 1$. Incorporate this into the control Lyapunov inequality (1.31), to get

$$(A(x) + B_2(x)u) \cdot \nabla_x V(x,u) + v \nabla_u V(x,u) + \Omega(x) \leq 0.$$

Next we use the saturation constraint on v and the inequality to get that it is sufficient that V satisfy the first order PDI

$$(A(x) + B_2(x)u) \cdot \nabla_x V(x,u) - |\nabla_u V(x,u)| + \Omega(x) \leq 0, \tag{1.32}$$

with

$$v^*(x,u) = -\text{sign}\,\{\nabla_u V(x,u)\}.$$

The corresponding controller satisfies is

$$u(t) = u(0) + \int_0^t v^*(x(s), u(s)) ds.$$

We emphasise that we are just giving sufficient conditions for solution.

In fact many more formulas much more thoughtfully crafted than this can be written out in a large variety of circumstances. This is an effort pioneered by E. Sontag and collaborators, see [LS91], [LS95] for cases like those we have just treated. Other very direct approaches to finding control Lyapunov functions for systems with special structures are described in [KKK95].

Now we describe another approach to imposing constraints on u. Rather than directly impose a hard rate bound $|\dot{u}| \leq 1$ we just punish "large" rates in some way, for example, we combine a cost of the form

$$\gamma^2 \tfrac{1}{2} \int_0^\infty |\dot{u}|^2 \, dt \tag{1.33}$$

with the control Lyapunov inequality (1.31). This gives

$$(A + Bu) \cdot \nabla_x V + v \nabla_u V + \gamma^2 \tfrac{1}{2} v^2 + \Omega \leq 0$$

with $\dot{u} = v$. Just as before we optimize over v, but with this approach there is no constraint on v, so the maximizing v is

$$v^*(x) = -\gamma^{-2} \nabla_u V(x, u)$$

and the PDI becomes

$$(A + Bu) \cdot \nabla_x V - \frac{1}{2\gamma^2} (\nabla_u V)^2 + \Omega \leq 0 \tag{1.34}$$

which is similar to a nonlinear Riccati inequality.

Similarly, we could be more cautious and treat disturbances w entering the system

$$\begin{aligned} \dot{x} &= A(x) + B_1(x) w + B_2(x) u, \\ z &= C_1(x) + D_{12}(x) u, \end{aligned}$$

thereby getting H_∞ type inequalities

$$(A + B_1 w + B_2 u) \cdot \nabla_x V + v \nabla_u V + \gamma_1^2 \tfrac{1}{2} v^2 - \gamma_2^2 \tfrac{1}{2} w^2 + \Omega \leq 0$$

maximize over w and minimize over v to get

$$v^*(x) = -\gamma_1^{-2} \nabla_u V(x, u) \quad \text{and} \quad w^*(x) = \gamma_2^{-2} B_1 \cdot \nabla_x V(x, u)$$

and the PDI

$$\begin{aligned} &(A + B_2 u) \cdot \nabla_x V(x, u) - \tfrac{1}{2\gamma_1^2} (\nabla_u V(x, u))^2 \\ &\quad + \tfrac{1}{2\gamma_2^2} \nabla_x V(x, u) B_1 B_1' \nabla_x V(x, u) + \Omega(x) \leq 0. \end{aligned}$$

We summarize all of this by saying that solution of the stabilization problem together with a rate saturation constraint amounts to solving a particular first order PDI.

The bottom line is that to handle constraints on u it is very likely that some first order PDI or comparably difficulty problem must be addressed. What we have done here is done with the intention of provoking thought and is hardly conclusive.

While this book treats HJBI inequalities much of what is done applies to large classes of first order PDI. The most extreme example is Bellman inequalities, since they are just the special case where $B_1 = 0$. The next section expands on this theme.

1.10 Other Performance Functions

A wide range of problems can be cast into a form that involves the use of optimization techniques, such as optimal control, game theory, and in particular, the dynamic programming method. In this book we emphasize measurement feedback problems, solved using the information state framework. This framework applies to a range of stochastic (or H_2) problems, and as we discuss in detail in this book, deterministic minimax problems.

The integrand $\frac{1}{2}|z|^2 - \gamma^2 \frac{1}{2}|w|^2$ in the cost functional used in this book (see (3.1)) has special meaning due to the L_2 dissipation inequalities and connection to the H_∞ norm (a frequency domain concept) in the case of linear systems. Any integrand $L(x, u, w)$ could be substituted in principle for $\frac{1}{2}|z|^2 - \gamma^2 \frac{1}{2}|w|^2$ and the corresponding solution could be derived using similar methods. In particular, a suitable information state can be defined:

$$p_t(x) \ = p_0(\xi(0)) + \tfrac{1}{2}\int_0^t L(\xi(s), u(s), y(s) - C_2(\xi(s)))\, ds$$

$$\text{given } u, y \text{ over } [0, t] \text{ and given } \xi(t) = x,$$

where as discussed above the trajectory $\xi(\cdot)$ is a solution of the reversed-arrow dynamics (1.17).

Further, measurement feedback versions of stabilization and loop shaping can also be developed. To illustrate, consider a robust version of the hard-constrained rate saturation example discussed above, where $|\dot{u}| \leq 1$, with plant model

$$\begin{aligned} \dot{x} \ &= \ A(x) + B_1(x)w + B_2(x)u \\ y \ &= \ C(x) + w \end{aligned}$$

and Lyapunov integrand $\Omega(x)$. The minimax cost associated with this problem for an output feedback controller $K : y(\cdot) \mapsto v(\cdot)$ is

$$J_{p_0}(K) = \sup_{T \geq 0,\, w(\cdot)} \{p_0(x_0) + \int_0^T [\Omega(x(s)) - \gamma^2 \tfrac{1}{2}|w(s)|^2] ds\}$$

with information state (a function of (x, u)) defined for $v(\cdot), y(\cdot)$

$$p_t(x, u) = p_0(\xi(0), u(0)) + \int_0^t [\Omega(\xi(s)) - \gamma^2 \tfrac{1}{2}|y(s) - C(\xi(s))|^2] ds,$$

where

$$\dot{\xi} = A(\xi) + B_1(\xi)w + B_2(\xi)u$$

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

for $0 \leq s \leq t$ and $\xi(t) = x$, $u(t) = u$. The PDE for the information state is

$$\dot{p} = -\nabla_x p(x,u) \cdot (A(x) + B_1(x)w + B_2(x)u) - \nabla_u p(x,u) \cdot v$$
$$+\Omega(x) - \gamma^2 |y - C(x)|^2.$$

The dynamic programming PDE is

$$\inf_{|v| \leq 1} \sup_y \{\nabla_p W(p)[-\nabla_x p \cdot (A + B_1(y - C) + B_2 \cdot) - \nabla_u p \cdot v$$
$$+\Omega - \gamma^2 |y - C|^2] = 0.$$

If we optimize over $|v| \leq 1$ we get

$$\sup_y \{\nabla_p W(p)[-\nabla_x p \cdot (A + B_1(y - C) + B_2 \cdot) - \nabla_u p \cdot v$$
$$+\Omega - \gamma^2 |y - C|^2] - |\nabla_p W(p)[-\nabla_u p]| = 0,$$

with optimizer

$$v^*(p) = -\mathrm{sign}\,\{\nabla_p W(p)[-\nabla_u p]\}.$$

This gives the controller

$$u(t) = u(0) + \int_0^t v^*(p_s) ds.$$

The integral-constrained control rate example (1.33) can be handled in the same way.

1.11 History

The objective of this section is to give a history of developments preceding this book. Initially the account follows that given in [HM98a]. We have attempted to mention the main developments, and we appologize in advance if we have missed some references.

1.11.1 Linear Frequency Domain Engineering

In commonplace language H_∞ engineering amounts to achieving prescribed worst case frequency domain specs. Optimizing worst case error in the frequency domain along its present lines started not with control but with passive circuits. One issue was to design amplifiers with maximum gain over a given frequency band. Another was the design of circuits with minimum broadband power loss. Indeed H_∞ control is a subset of a broader subject, H_∞ *engineering*, which focusses on worst case design in the frequency domain. In paradigm engineering problems this produces what a mathematician calls an "interpolation problem" for analytic functions. These can be solved by Nevanlinna-Pick techniques. The techniques of Nevanlinna-Pick interpolation had their first serious introduction into engineering in a SISO (single-input single output) circuits paper of Youla and Saito [YS67] in the middle 1960's. Further development waited until the mid-seventies, when Helton [Hel76], [Hel78], [Hel81] applied interpolation and more general techniques from operator theory to amplifier problems. Here the methods of commutant lifting [And63], [SNF70], [Sar67] and of Admajan-Arov-Krein [AAK68], [AAK72], [AAK78] were used to solve MIMO optimization problems.

In the late 1970's G. Zames [Zam79] began to marshall arguments indicating that H_∞ rather than H_2 was the physically proper setting for control. Zames suggested on several occasions that these methods were the appropriate ones for codifying classical control. These efforts yielded a mathematical problem which Helton identified as an interpolation problem solvable by existing means (see [ZF81]). In 1981 Zames and Francis [ZF83] used this to solve the resulting single input single output SISO problem. In 1982 Chang-Pearson [CJ84] and Francis-Helton-Zames [FHZ84] solved it for many-input, many-output MIMO system.

The pioneering work of Zames and Francis treated only sensitivity optimization. In 1983 three independent efforts emphasized bandwidth constrains, formulated the problem as a precise mathematics problem and indicated effective numerical methods for their solution: Doyle [Doy83], Helton [Hel83], and Kwakernaak [Kwa83]. All of these papers described quantitative methods which were soon implemented on computers. It was these papers which actually laid out precisely the tradeoff in control between performance at low frequency and roll off at higher frequency and how one solves the resulting mathematics problem. This is in perfect analogy with amplifier design where one wants large gain over as wide a band as possible, producing the famous gain bandwidth trade off.

Another independent development was Tannenbaum's [Tan80] very clever use of Nevanlinna-Pick interpolation in a control problem in 1980. Also early on the H_∞ stage was Kwakernaak's polynomial theory, [Kwa86]. Another major development that dove tailed closely with the invention of H_∞ control was a tractable theory of plant uncertainty. A good historical treatment appears in [DFT92]. Another application of these techniques is to robust stabilization of systems H. Kimura *et al* [Kim84]. An early book on H_∞ control was [Fra84].

1.11.2 Linear State Space Theory

To describe the origins of state space H_∞ engineering we must back up a bit. Once the power of the commutant lifting-AAK techniques were demonstrated on engineering problems, P. de Wilde played a valuable role by introducing them to signal processing applications, see [DVK78], and to others in engineering. The state space solutions of H_∞ optimization problems originated not in H_∞ control, but in the area of model reduction. The AAK work with a shift of language is in a paper on model reduction (though not in state space coordinates) by Bettayab-Safanov-Silverman [BSS80], which gives a state space viewpoint for SISO systems. Subsequently Glover [Glo84] gave the MIMO state space theory of AAK type model reduction. Since the H_∞ control problem was already known to be solvable by AAK, this quickly gave state space solutions to the H_∞ control problem. These state space solutions were described first in 1984 by Doyle in a report [Doy84] which though never published was extremely influential. Earlier in his thesis (unpublished) he had given state space H_∞ solutions based on converting the geometric (now it would be called behavioral by engineers) version of commutant lifting-AAK due to Ball and Helton to state space.

There was a vast effort on state space H_∞ control by many engineers and mathematicians. We mention now only a few major developments. In the beginning there were only crude bounds on the dimension of the state space of the controller and numerical recipes for the controller relied on substantial cancellation which of course is bad. It was discovered by [LH88] that the dimension of an H_∞ optimal controller equals that of the plant G. Next came the famous paper of [DGKF89] which gave an elegant cancellation free formula for the controller (as discussed in §1.2). The formulas in this paper have become standard. Other closely related results also appeared around this time or a little later, see [Tad90], [Sto92]. An excellent presentation is given in [GL95].

1.11.3 Factorization

It might be mentioned that factorization (the subject of Chapter 8) was known from early on to yield all controllers producing a certain performance, as well has other problems, c.f. [HBJP87]. These methods were developed by Ball-Helton and H. Kimura and coworkers in many papers during the 80's and 90's(see [BHV91], [Kim97] and the references therein). This lead to an elegant proof of the original DGKF formulae as well as the first discrete time DGKF formulas by Ball-Ran [BR87]. A J-spectral factorization approach was presented in [GGLD90], [Gre92].

1.11.4 Game Theory

It was observed in [Pet87], [DGKF89] (and elsewhere) that there are close connections between H_∞ control and differential games. Basically, the two quite distinct problems can be solved using the *same* Riccati equations. These connections were pursued in depth by a number of researchers, see e.g. [LAKG92], [BB89] (updated in 1995). The game theory view of H_∞ control is as a minimax game, where the disturbance or uncertainty is modelled by a malicious opponent, and the aim of the controller is to minimize the worst performance under these circumstances. This time domain formulation is very important for nonlinear systems.

1.11.5 Nonlinear H_∞ Control and Dissipative Systems

The efforts to extend H_∞ control to nonlinear systems begin in the middle 80's by mathematicians versed in linear commutant lifting and AAK techniques. Ball-Foias-Helton-Tannenbaum formulated the nonlinear problem and showed that power series (Volterra) expansions lead to reasonable approximate solutions, [BFHT87b], [BFHT87a]. This effort has continued to the produce impressive results, [FGT95], [FGT96], [FGT98]. Ball-Helton pursued several different approaches. One was what would today be described in terms of behaviors or games in extensive form [BH88c]. Another was in state space form [BH92a], [BH92b], [BH88a], [BH88b]. This reduced the solution of the measurement feedback discrete time nonlinear problem for a "strongly" stable plant G to solution of an HJBI equation.

For continuous time state feedback basic work was done by van der Schaft [vdS91], [vdS92], [vdS96]. He reduced the solution of the state feedback problem for a nonlinear plant G to solution of an HJBI equation. This work was influenced by Willems' theory of dissipative systems, [Wil72], [HM76], [HM77], etc. Indeed, van der Schaft emphasizes L_2-gain terminology and the Bounded Real Lemma, [AV73]. This is a powerful and natural formulation. Indeed, it is the L_2-gain inequality (which we refer to as the dissipation inequality in this book) which makes sense for nonlinear systems, whereas the frequency domain concept of H_∞ norm does not apply to nonlinear systems.

1.11.6 Filtering and Measurement Feedback Control

Classical control problems, as discussed earlier, are formulated in the frequency domain and are naturally measurement feedback problems. This is reflected in the Ball-Helton papers of the 80's. Optimal control with measurement feedback is difficult, and this explains in part the length of time it took to obtain a nice state space solution to the linear H_∞ control problem (most of a a decade). The issue is how to represent and use the information contained in the measurements.

Much of optimal control theory (including games) is concerned with state feedback problems. This is natural, since the state of a system is a summary of its status, and together with the current input values can be used to determine future behavior. Engineers are interested in feedback controllers, and solutions to state feedback optimal control problems lead to state feedback solutions (via, say, dynamic programming). However, given that the original problem of interest is a measurement feedback one, there is the difficulty of what to do with the lack of full state information. A common, but often suboptimal approach is to design a state estimator (or observer), and plug the state estimate into the optimal state feedback controller. This is called certainty equivalence. The solution of the Linear Quadratic Gaussian (LQG) problem is an optimal certainty equivalence controller, [Won68]. First, an optimal state feedback controller is designed, and then coupled with the output of the optimal state estimator, i.e. the Kalman-Bucy filter, [Kal60] [KB60]. The certainty equivalence approach is not optimal for the deterministic Linear Quadratic Regulator (LQR) problem. Deterministic LQR designs may employ a Luenberger observer, [Lue66].

The linear LQG problem is a stochastic optimal control problem. What is happening in Kalman's solution is that the optimal state estimate, the conditional mean, becomes the state of a new system, and the optimal controller for this new system turns out to coincide with the optimal state feedback controller for the original system. Actually, the optimal LQG controller feeds back the conditional probability distribution, which being a Gaussian distribution, is completely determined by the conditional mean and covariance (finite parameters). For nonlinear optimal stochastic control problems analogous to LQG, the optimal controller is a function of the conditional distribution. Thus the conditional distribution serves as an "information state" for these optimal control problems. The measurement feedback optimal control problem is transformed into a new state feedback optimal control problem, with the information state serving as the state variable. The evolution of the conditional distribution is described by a stochastic partial differential equation, called the Kushner-Stratonovich equation [Kus64], [Str68], or in unnormalized form, the Duncan-Mortensen-Zakai equation, [Dun67], [Mor66], [Zak69]. These are the stochastic PDEs of nonlinear filtering, and are the nonlinear counterparts to the Kalman filter equations. Thus nonlinear filtering is infinite dimensional, and measurement feedback optimal stochastic control involves the optimal state feedback control of an infinite dimensional system.

The information state approach has been well known since at least the 60's, both in the West and East. A nice explanation of these ideas is given in [KV86]. Of the many publications devoted to this problem, we mention only [Str65], [Nis76], [Ell82], [FP82], [Fle82], [Hij90], [EAM95]. It is still a difficult mathematical problem, and presents challenging implementation issues.

For nonlinear problems analogous to the deterministic LQR problem, there is no information state solution, and one typically uses a suboptimal certainty equivalence design as discussed above. A key difficulty here is the design of the state estimator or observer. This is a major problem in nonlinear control, [KET75], [HK77], [KR85], etc.

In contrast, it is relatively straightforward to write down a nonlinear filter, although one is faced with computational difficulties for implementation. In 1968, R.E. Moretensen derived a deterministic approach to nonlinear filtering, called minimum energy estimation, [Mor68]. This is essentially a least squares approach, and leads to a filter which is a first order nonlinear PDE. An interesting study of this filter was conducted in 1980 by O. Hijab, [Hij80]. These deterministic filters are related to the stochastic filters via small noise limits. These limits are examples of the type which occur in the theory of large deviations. J.S. Baras was intrigued by these filters and their connections, and in [BK82] proposed using these methods as the basis of a design procedure for nonlinear observers, [BBJ88], [JB88], [Jam91].

1.11.7 H_∞ Control, Dynamic Games, and Risk-Sensitive Control

In 1973, D.H. Jacobson, [Jac73], introduces a new type stochastic optimal control problem with an exponential cost function, which today is often called the risk-sensitive problem. He solved a Linear Exponential Quadratic Gaussian (LEQG) problem with full state feedback, and observed that his solution is the same as the solution for a related dynamic game (same Riccati equation). It took until 1981 for the corresponding linear measurement feedback problem to be solved, by Whittle [Whi81]. The structure of the controller is again of the certainty equivalence type, although the Kalman filter estimate is not used. Instead, the Kalman filter is modified with terms coming from the control objective. Whittle's solution was very interesting, since the conditional distribution is not used as the information state. Later, connections with H_∞ control, were discovered, [GD88], [DGKF89]. Thus H_∞ control, dynamic games, and risk-sensitive control are all related.

In the late 80's and early 90's Basar-Bernhard and coworkers developed the certainty equivalence principle for deterministic minimax games and H_∞ control. The key reference here is the 1989 monograph [BB89] (revised in 1995), as well as the papers [Ber91], [DBB93], [BR95]. The book [BB89] contains an excellent account of the minimax game approach and certainty equivalence mainly in the linear context, with some nonlinear results in the second edition. The certainty equivalence solution is very closely related to the solution of Whittle, and is the basis of an important approach to measurement feedback nonlinear H_∞ control.

In the early 90's a number of researchers began exploring the connections between H_∞ control, dynamic games, and risk-sensitive control in the nonlinear context, beginning with Whittle [Whi90a], [Whi90b], [Whi91]. The connections made use of small noise limits. This work inspired Fleming-McEneaney, leading to the papers [FM92], [FM95], and also to papers studying viscosity solutions of the H_∞ PDEs and PDIs [BH96], [Jam93], [McE95b], [McE95a], [Sor96]. Independently, J.S. Baras suggested investigating the risk-sensitive problem using small noise methods, in conjunction with earlier work on nonlinear filters. This led to the papers [Jam92], [JBE94], [JB95], [JB96], [BJ97]. The paper [JBE94] solved the nonlinear measurement feedback (discrete time) stochastic risk-sensitive problem, solved a nonlinear measurement feedback deterministic minimax game, and established connections between them via small noise limits. An information state was used for both problems, and in the risk-sensitive case, the information state was not the conditional probability distribution. The information state definition was inspired by the paper [BvS85], which used a method which generalizes to nonlinear systems. In the minimax case, the information state coincides with Basar-Bernhard's cost-to-come, and is related to the risk-sensitive information state in a manner analogous to the link between Mortensen's minimum energy estimator and stochastic nonlinear filters discussed above. See also the publications [KS89], [Ber96]. A large number of papers have since been written concerning various aspects of risk-sensitive control, filtering, games, and their connections: [PMR96], [CE95], [CH95], [FHH97], [FHH98], [Nag96], [RS91], [Run91], etc.

1.11.8 Nonlinear Measurement Feedback H_∞ Control

While stable plant problems had been known to convert to HJBI inequalities since the late 80's, the unstable measurement feedback problem remained intractable. A substantial number of papers have been written, including: Isidori-Astolfi-Kang [IA92a], [IA92b], [Isi94], [IK95], Ball-Helton-Walker [BHW93], Didinsky-Basar-Bernhard [DBB93], Krener [Kre94], Lin-Byrnes [LB95], Lu-Doyle [LD94], Maas [Maa96], Nguang [Ngu96]. These results illuminated various aspects of the measurement feedback problem, and indeed the results all specialized to the well known DGKF solution when applied to linear systems. The results were generally of a sufficient nature, so that if certain PDEs or PDI could be solved, then a solution to the nonlinear H_∞ control problem would be produced. However, in general these results are far from being necessary: H_∞ controllers could exist but not be of the form given in these papers. This is because nonlinear filtering, and hence optimal measurement feedback control, is intrinsically infinite dimensional.

Information state controllers for nonlinear H_∞ control were obtained by a number of authors in the early 90's. van der Schaft [vdS96] identified some of the key measurement feedback equations, including the coupling condition, and obtained information state controller formulas assuming certainty equivalence. Didinsky-Basar-Bernhard [DBB93] obtained information state controllers assuming certainty equivalence and generalized certainty equivalence. The first general solution to the nonlinear H_∞ problem was given in [JB95] (see also [JBE94]). The information state was employed to give an intrinsically infinite dimensional solution, complete with a clean set of basic necessity and sufficiency theorems. A number of related papers have appeared since then, e.g. [Teo94], [TYJB94], [JY95], [Yul96]. In 1994 Helton-James realized that the information state framework could be used for J inner outer factorization, and preliminary results and formulas were published in [HJ94]. This initiated a detailed investigation and development of the information state solution, leading to the papers [HJ95], [HJ96b], [HJ96a], and ultimately, to this book.

1.11.9 Prehistory

Now we lurch back to sketch the origins of the HJBI equations which play such a big role in this book. This is an extensive subject which is well described in many places, so we give little account of the history and just list some references. Thus we urge the curious to read [Bel57], [Isa65], [You69], [FR75], [FS93], [BO95].

1.12 Comments Concerning PDEs and Smoothness

In this book we make extensive use of optimal control methods and nonlinear PDEs (Hamilton-Jacobi type). In general, solutions to such PDEs are *not globally smooth*, and in Appendix B we discuss these equations and their solutions, in particular, the concept of *viscosity solution*.

We have attempted to minimize technical issues arising because of lack of smoothness, to keep the focus of the book on control-theoretic ideas. In many places we use PDEs on finite dimensional spaces (such as the PDE giving the dynamics of the information state), and use integrated (i.e. dynamic programming) representations which are meaningful without smoothness. In some results we assume smoothness to help keep statements clear (and readily connected to the familiar linear case), and to simply proofs. However, readers should be aware that such results remain valid without the smoothness assumptions, with appropriate interpretations and proofs.

PDEs on infinite dimensional spaces play a major role in this book. There are many unresolved purely mathematical issues concerning these PDEs. We have not attempted to describe in detail issues concerning the concept of solution for such equations (this is still an open question). Instead, we have stated a number of results which have no need of smoothness (these make use of the integrated dynamic programming equation). However, when one uses the dynamic programming PDE to obtain an optimal feedback controller (such as our construction of the central controller) some form of smoothness is required, so we formalize what we need and assume this in order to develop the control-theoretic ideas. We have tried to make clear where smoothness is or is not assumed.

We remark that the results in this book have discrete time analogs (see [JB95]), and differentiability is irrelevant in discrete time. Thus discrete time controllers can be obtained directly from discrete time analogs of the dynamic programming PDE without the need for the value function to be differentiable.