

September, 2000

LIDS- TH 2509

Research Supported By:

NSF/KDI Contract ECS-9873451
ARO Contract DAAG55-97-1-0304 and
DAAL03-92-G0115 and John Hopkins
Subcontract 8905-48164

Control Under Communication Constraints

Sekhar Chandra Tatikonda

Control Under Communication Constraints

by

Sekhar Chandra Tatikonda

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2000

© Massachusetts Institute of Technology 2000. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
September 1, 2000

Certified by
Sanjoy K. Mitter
Professor of Electrical Engineering
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Control Under Communication Constraints

by

Sekhar Chandra Tatikonda

Submitted to the Department of Electrical Engineering and Computer Science
on September 1, 2000, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

An important aspect of distributed control is the role of communication between the different components in the system. Traditional information theory is asymptotic, has delays, and does not completely deal with feedback. Since feedback is an essential element of control in the presence of uncertainty and since delays have to be taken into account in control problems, especially for unstable systems, it is natural to look for a unification of information theory and stochastic control.

We present a unified view of control and communication which clarifies many of the conceptual issues underlying the distributed control problem. This view consists of considering a distributed system as an interconnection of different probabilistic systems: be they channels, plants, etc. We discuss the importance of centralized design for this distributed implementation and the conceptual role of dynamic programming.

We provide a very general coding theorem for channels with feedback. We show that the directed mutual information, as introduced by Massey, is the correct notion of capacity for channels with and without feedback. For Markov channels we show that one can solve the capacity optimization problem via dynamic programming.

We formulate the sequential rate distortion problem and provide a coding theorem. For Markov sources we show that one can solve the sequential rate distortion infimization problem via dynamic programming. Finally we show that the successive refinement problem is a special case of the sequential rate distortion problem.

For the general problem of control under communication constraints we examine a distributed system with a plant, a channel encoder, a channel, a channel decoder, and a controller. We give conditions on the capacity of the channel to ensure different control objectives: observability, stability, controllability, and performance. For deterministic systems we introduce the notion of covering number. For the LQG problem we give suitable assumptions on the information pattern to ensure the optimality of the certainty equivalent controller.

Thesis Supervisor: Sanjoy K. Mitter

Title: Professor of Electrical Engineering

Acknowledgments

I would like to thank my advisor Professor Sanjoy Mitter for his encouragement, his advisement, his generous research spirit, and his unwavering confidence in my abilities.

I would like to thank my thesis committee members, Professor Robert Gallager and Professor John Tsitsiklis, for their helpful comments. Professor Vivek Borkar has been a constant source of ideas and inspiration. I thank Anant Sahai for many discussions on the topic of control under communication constraints.

Special thanks go to my “spiritual advisors:” Nicola Elia, Venkatesh Saligrama, and Tushar Shah.

I thank all my friends at LIDS, past and present, who make it such a wonderful place to learn and do research.

Finally this thesis is dedicated to my family. I could not have completed it without the love and support of Mom, Dad, Ramu, and Aruna.

This research was partially supported by NSF/KDI Contract ECS-9873451, ARO Contract DAAG55-97-1-0304, ARO Contract DAAL03-92-G-0115, and ARO/John Hopkins University Subcontract 8905-48164. They are very much appreciated.

Contents

1	Introduction	8
1.1	Control Under Communication Constraints	8
1.1.1	Complexity and Communication	9
1.1.2	Signifier versus Signified	10
1.1.3	Centralized versus Decentralized Design	10
1.2	General Framework	11
1.3	Summary of Thesis	13
1.3.1	Chapter 2	13
1.3.2	Chapter 3	13
1.3.3	Chapter 4	13
1.3.4	Chapter 5	15
1.3.5	Chapter 6	15
2	A General Framework for the Distributed Control Problem	17
2.1	Introduction	17
2.2	General Framework	18
2.2.1	Main Definition	18
2.2.2	The Finite Horizon Control Problem	20
2.2.3	System and Policy Knowledge	20
2.2.4	Partial Orderings on the Signal and System Knowledge	21
2.2.5	Extensions	22
2.3	Dynamic Programming Formulation	26
2.3.1	Difficulties with Decentralized Design	28
2.4	Summary	29
3	Control of Deterministic Systems Under Communication Constraints	30
3.1	Introduction	30
3.2	Problem Setup	32
3.2.1	Information Pattern and System Knowledge	33
3.3	Lower Bounds that are Independent of the Information Pattern	35
3.3.1	Observability	35
3.3.2	Stability	37
3.3.3	Controllability	38
3.4	Encoder Classes, Primitive Quantizers, and the Key Technical Lemma	40

3.4.1	Encoder Classes	40
3.4.2	Equi-memory	42
3.4.3	Primitive Quantizer	43
3.4.4	Key Technical Lemma	45
3.4.5	Comments	47
3.5	Encoder Class One	49
3.5.1	Observability	49
3.5.2	Stability	54
3.5.3	Controllability	55
3.6	Encoder Class Two	57
3.6.1	Observability	57
3.6.2	Stability	58
3.6.3	Controllability	62
3.6.4	Lyapunov synthesis	62
3.6.5	Minimum communication	64
3.7	Systems with Multiple Sensors	66
3.7.1	Problem setup	66
3.7.2	Relationship to Slepian-Wolf Coding	69
3.8	Sampling of Continuous Time Systems	70
3.8.1	Observability	70
3.8.2	Stability	72
3.9	Performance	74
3.9.1	Problem setup	74
3.9.2	LQ	75
3.9.3	Covering numbers	76
3.9.4	Discussion	77
3.10	Summary	78
4	Channel Coding With Feedback	79
4.1	Introduction	79
4.2	Directed Mutual Information	81
4.3	Channel Capacity	84
4.3.1	Channels	84
4.3.2	Message Set, Encoder, and Decoder	86
4.3.3	Interconnection of Code-Functions to the Channel	86
4.3.4	Directed Mutual Information: The Intuition	88
4.3.5	Operational Channel Capacity	89
4.3.6	Channel Capacity	90
4.4	Coding Theorem for Finite Alphabet Channels	93
4.4.1	Main Technical Lemmas	93
4.4.2	Feedback Channel Coding Theorem	98
4.4.3	Error Exponents	102
4.5	Markov Channels and the Dynamic Programming Formulation	106
4.5.1	Setup	106
4.5.2	Coding Theorem	110

4.5.3	Stochastic Control Formulation	111
4.5.4	Extensions and Limitations of our Formulation	117
4.6	Coding Theorem for Gaussian Channels	118
4.6.1	Setup	118
4.6.2	Main Theorem	120
4.6.3	Error Exponents	127
4.7	Examples	129
4.8	Channel Realizations and the Directed Data Processing Inequality	131
4.9	Summary	135
5	Sequential Rate Distortion	136
5.1	Introduction	136
5.2	Rate Distortion	138
5.2.1	Review	138
5.2.2	Issue of Delay	139
5.2.3	Induced Channel and Approximate Factorization	140
5.2.4	Examples	143
5.2.5	High Rate Approximation	145
5.2.6	Summary	147
5.3	Sequential Rate Distortion	148
5.3.1	Setup	148
5.3.2	Converse Theorem	151
5.4	Coding Theorem for the Noiseless Digital Channel	153
5.4.1	Finite Alphabet Source Coding Theorem	154
5.4.2	Stochastic Control Formulation for Finite Alphabet Source Coding	160
5.4.3	Gaussian Source Coding Theorem	167
5.4.4	Stochastic Control Formulation for Gaussian Source Coding	170
5.4.5	Summary and Extensions	174
5.5	Successive Refinement	175
5.5.1	Setup and Results	175
5.5.2	Examples	176
5.5.3	Rate of Convergence	182
5.5.4	The Relationship between Successive Refinement and Feedback Channel Capacity	185
5.6	Solution Characterization and High Rate Approximation	188
5.6.1	Sequential Rate Distortion for Gauss-Markov Sources	188
5.6.2	Sequential Rate Distortion for Differing Information Patterns	191
5.6.3	High Rate Asymptotics	193
5.7	Summary	196
6	Control of Stochastic Systems Under Communication Constraints	197
6.1	Introduction	197
6.2	Stochastic Control Problem	199
6.3	Linear Quadratic Gaussian Problem	203
6.3.1	Tracking	204

6.3.2	Quadratic Performance	205
6.4	Summary	208
7	Summary	209
A	Background and Supporting Results	211
A.1	Chapter 2	211
A.2	Chapter 3	213
A.3	Chapter 4	215
A.4	Chapter 5	221
	References	224

Chapter 1

Introduction

The Internet, wireless networks, and the like are making it possible to have many remote plants linked together via communication channels. Thus a substantive theory of distributed control is increasingly becoming important in today's control systems.

An important aspect of distributed control is the role of communication between the different components in the system. In this setting traditional information theory, which codifies the fundamental limitations to reliable communication over noisy channels, is not directly applicable. The reason is because traditional information theory is asymptotic, has delays, and does not completely deal with feedback. Since feedback is an essential element of control in the presence of uncertainty and since delays have to be taken into account in control problems, especially for unstable systems, it is natural to look for a unification of information theory and stochastic control when components of control systems are interconnected through communication channels.

In this thesis we present a unified view of control and communication which clarifies many of the conceptual issues underlying the distributed control problem. We discuss the interaction between information and control. Our main goal is to understand the fundamental limits of control performance in distributed systems when there are communication constraints.

This introduction is divided into three parts. Section 1.1 reviews the issues involved in distributed control over communication channels. Section 1.2 discusses our general framework for dealing with these problems. And section 1.3 gives a summary of the chapters that follow. Within each chapter summary is a review of the relevant existing literature.

1.1 Control Under Communication Constraints

What is a distributed system? Qualitatively it consists of a set of plants and controllers linked by communication channels. See figure 1-1. There the boxes can be plants or controllers. The arrows represent different communication channels. A line from a plant to a controller represents a sensor signal. A line from a controller to a plant represents an

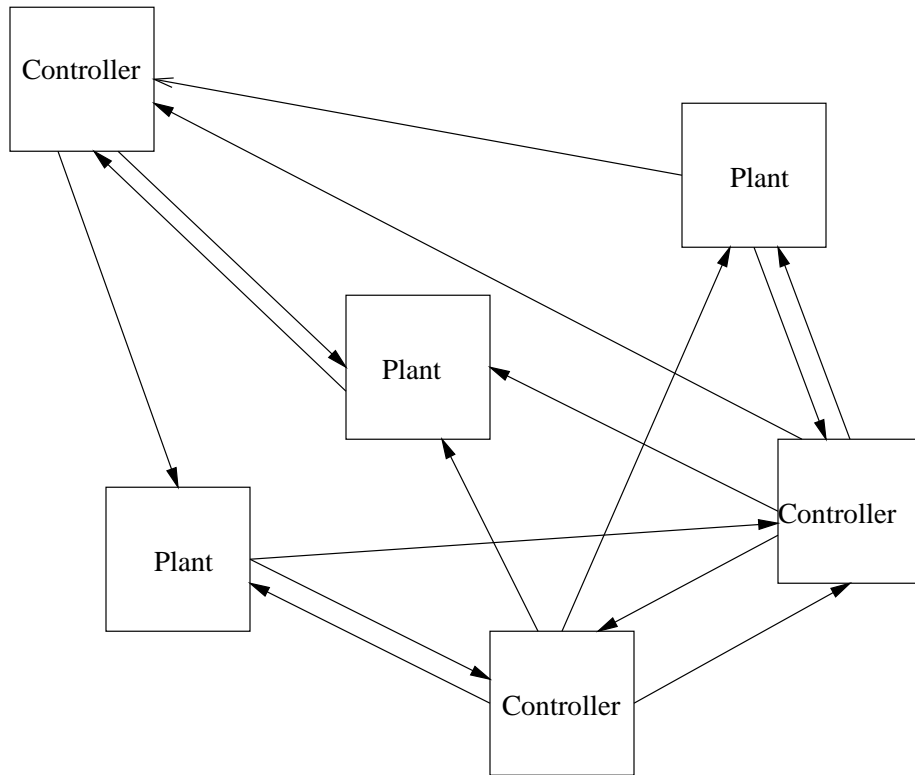


Figure 1-1: Distributed System

actuation signal. And a line from a controller to another controller represents a coordination signal.

1.1.1 Complexity and Communication

If all the communication links in figure 1-1 are of infinite bandwidth and zero delay then there is no reason to treat the problem as a distributed problem. We can easily construct a new centralized controller with links to all the plants.

There are two main reasons for studying distributed systems. One is the issue of complexity in the design of the controllers. We will not treat this issue in this thesis. The second is the issue of communication. If the channels are finite bandwidth then thought has to be put into what signals we want to send across them. Thus, for us, distributed control arises because we have information bottlenecks due to the finite bandwidth communication links connecting the different components of the system.

These communication links can be noisy, have delays, and drop signals. Furthermore they may have memory. Thus these communication channels can be considered to be plants

themselves. The channel encoders and channel decoders can be considered to be controllers. By viewing channels as plants and encoders and decoders as controllers we are able to unify the different components of the distributed system.

1.1.2 Signifier versus Signified

Given a signal source and a channel we have to decide what part of that signal we want to transmit across the channel. Furthermore we need to decide how we want to represent that signal over the channel. This dichotomy is essentially captured in Saussure's distinction between the "signifier" and the "signified." [Sau] We need to determine first what part of the source signal we want to transmit across the channel. This is the "signified" part. Then we need to decide how we are going to represent that signal. This is the "signifier" part.

Information theory has addressed this problem in its two major founding results. Shannon's channel coding theorem tells us how many signifiers we can transmit across a channel reliably. And Shannon's rate distortion theorem tells us how well we can approximate the signified part of the source signal. [Sha1] The issue can be restated as follows. Data without knowledge of what is being signified is not useful information. Similarly connectivity without knowledge of the signifiers does not lead to useful coordination.

1.1.3 Centralized versus Decentralized Design

As stated before the systems we are interested in studying are distributed. There are two options though for the design of the controllers: centralized design and decentralized design.

In this thesis we focus on the centralized design. (Though in parts of chapter two we treat some aspects of decentralized design.) The idea is that although the system, when running, is distributed, the design of the policies at each controller can be done with common knowledge of the other controller's policies. In the decentralized design case some of the controllers do not know the other controller's policies. Thus there is an inherent game-theoretic aspect to the decentralized design that makes it much more difficult to analyze. The design problem becomes more tractable if we assume that there is coordination between the design of the controllers. Furthermore many distributed systems are centrally designed even though the real-time functioning of the system is distributed.

1.2 General Framework

We now present an overview of our general framework. Plants and channels can be described by stochastic kernels. They can be thought of as a partial specification of a joint probability measure on the variables of interest. Optimal control has to do with optimizing, over the allowed controllers, some performance objective. To compute this performance we require a complete probability measure on the variables of interest. The maps defined by the controllers allows us to “complete” this joint measure from the partial specification given by the plants and the channels. In fact controllers are nothing more than stochastic kernels themselves. We are most familiar with control policies that are functions. But this is just a stochastic kernel that maps inputs to Dirac measures.

The view taken throughout this thesis is that we interconnect controller stochastic kernels with the plant and channel stochastic kernels to form a joint measure.

An important issue that has caused a lot of confusion in the literature in both information theory and control theory is the role of causality in control systems when feedback is present. We treat this by explicitly putting a time ordering on the random variables of interest. Causality is then defined by this ordering. An example will help elucidate this point.

Consider a traditional control problem with time horizon T and state transition maps: $\{P(dX_{t+1}|x_t, u_t)\}_{t=1}^T$. The random variables of interest are the state variables X_1, \dots, X_T and the control variables U_1, \dots, U_T . The natural time ordering is

$$X_1, U_1, \dots, X_T, U_T.$$

Any joint measure $P(dX^T, dU^T)$ can be factored according to this *causal* ordering as follows

$$P(dX^T, dU^T) = \bigotimes_{t=1}^T P(dU_t|x^t, u^{t-1}) \otimes P(dX_t|x^{t-1}, u^{t-1}).$$

We already know $P(dX_t|x^{t-1}, u^{t-1}) = P(dX_t|x_{t-1}, u_{t-1})$ for $t = 1, \dots, T$. The controller is described by the sequence of stochastic kernels: $\{P(dU_t|x^t, u^{t-1})\}_{t=1}^T$. We design these controller kernels. Note that we can incorporate differing dependencies of the control on the past data. For example we may want U_t to only depend on the current state X_t . These sorts of restrictions are called the information pattern of the decision variable U_t .

A model of a system is defined to be the set of all joint measures on the variables of interest that satisfy:

- (1) A time ordering on the random variables of interest.
- (2) A specification of the stochastic kernels representing the plants and channels in the system.
- (3) A specification of the information patterns for the different decision variables.

We ‘can ‘complete’ this joint measure by specifying a sequence of stochastic kernels representing controllers.

We show that within this framework and the assumption of a centralized designer we can formulate a broad class of distributed control problems as dynamic programming problems. The next section summarizes the key results in the chapters to follow.

1.3 Summary of Thesis

Figure 1-2 lists the main results and logical flow of the chapters in this thesis.

1.3.1 Chapter 2

In this chapter we present the aforementioned general framework for the distributed control problem. The field of distributed control is quite large and we could never do it justice by summarizing it here. The work in this chapter, though, is heavily influenced by the work of Witsenhausen. He wrote an important paper in 1971 where he defines the notion of information pattern and discusses conditions for the separation of estimation and control. [Wit] Later he showed general conditions for the policy independence of conditional expectations. [Wit2] We expand on his idea of information patterns by discussing system and policy knowledge. In the resulting dynamic program we show that for nontraditional information patterns we no longer have the policy independence of conditional expectations. We make appropriate assumptions on centralized design to deal with this problem.

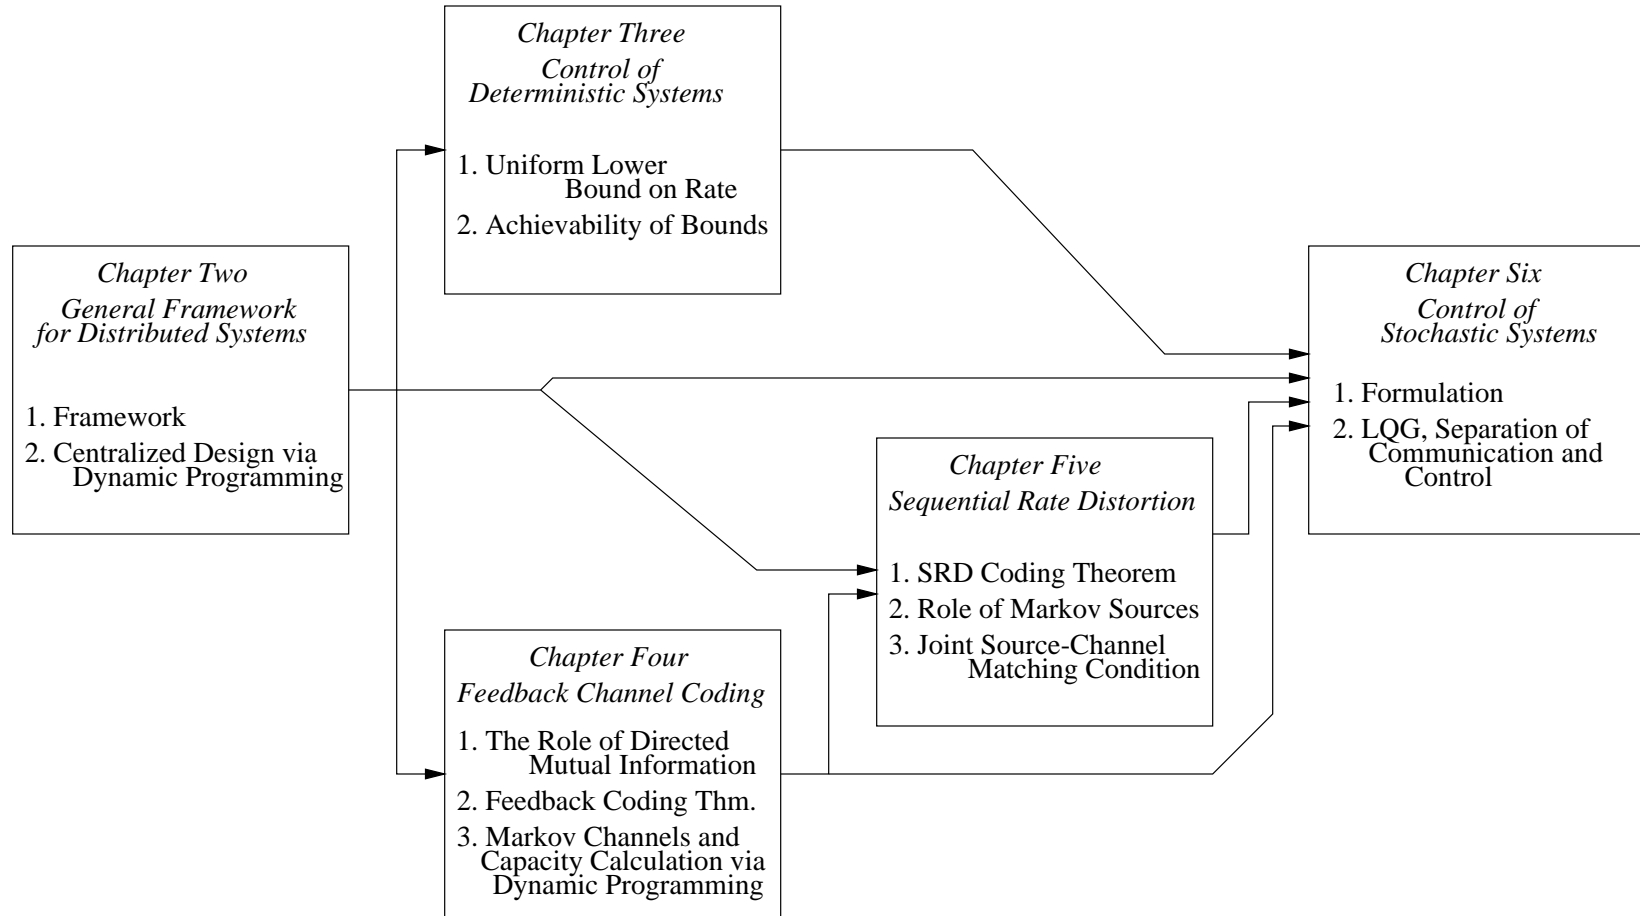
1.3.2 Chapter 3

In this chapter we examine the control of deterministic plants with a noiseless digital channel of finite rate connecting a sensor, measuring system variables at the plant, to the controller. We are interested in computing the minimum rate needed to achieve different control objectives: observability, stability, controllability, and performance. For these different control objectives we provide a lower bound on the required channel rate. This bound is independent of the information patterns of the encoder, decoder, and controller. We give conditions on the information pattern so that this rate is achievable. Finally we introduce the notion of covering number. The covering number counts the minimal number of control trajectories needed to achieve some given performance objective. This in turn is used to compute the required channel rate.

The original impetus for the work in this chapter came from two papers written by Wong and Brockett. [WB1], [WB2] They introduced the “systems with finite communication bandwidth constraints” problem. They give sufficient conditions, in the form of explicit schemes, for state estimation and stability. We extend their preliminary results in many directions. Nair and Evans’ work has evolved in parallel with the work in this chapter. [NE1], [NE2] They examine state estimation for a more general class of processes over a bit-rate constrained channel. They also provide sufficient conditions in the form of explicit schemes. Elia and Mitter have treated the problem of stability where the quantizer has fixed levels. [EM] See also Liberzon and Brockett. [LB]

1.3.3 Chapter 4

In this chapter we examine the problem of channel coding for channels with different kinds of memory and feedback. Shannon was the first to consider the feedback channel. [Sha2]



Main Results and Logical Flow of the Chapters

Figure 1-2

Dobrushin and Wolfowitz extended Shannon's results. [Dob1] [Wol] We prove a very general coding theorem for finite alphabet channels with different forms of feedback. This feedback coding theorem is a generalization of the non-feedback coding theorem presented by Verdu and Han. [VH] We prove the coding theorem by using Dobrushin's idea of defining an interconnection between a source and a channel. [Dob2] See also Gallager's book. [Gal]

We show that the directed mutual information, as introduced by Massey, is the correct notion of capacity for channels with and without feedback. [Mas] We also extend the work of Kramer on determining the properties of directed mutual information for the single-user channel with feedback. [Kra] Kramer treats the memoryless multiple access channel and the memoryless two-way channel but does not treat the single user case with channel memory and feedback.

For Markov channels we show that one can solve the capacity optimization problem via dynamic programming. We treat channels with ISI and differing side information at the transmitter and receiver. This dynamic programming formulation allows us to capture many existing coding theorems for different channels within one framework. Furthermore, by using this framework, we provide coding theorems for new channels. Extensions to the Gaussian channel are provided. For related work see [CS], [CP], [GV1], [GV2], [SK], [Sha3], and [Vis]. We discuss our work in relation to these works in the chapter.

Finally we discuss channel realizations and provide a causal generalization of the data processing inequality.

1.3.4 Chapter 5

In this chapter we formulate the sequential rate distortion problem. The sequential rate distortion problem is a generalization of the traditional rate distortion problem to processes over time with the added restriction that the reconstructions be computed causally and without delay. We define the sequential rate distortion function to be the infimum over all causal channels of the directed mutual information under a constraint on the distortion. This optimization problem was first formulated by Gorbunov and Pinsker. [GP] We, though, formulated the problem independently of them. They provide many structural results. We give an operational meaning to the sequential rate distortion function. Specifically we provide a coding theorem for the digital noiseless channel. For Markov sources we show that the underlying optimization problem can be solved via dynamic programming. We show that the successive refinement problem is a special case of the sequential rate distortion problem. [EC], [Rim]

We treat the joint-source channel coding problem. [VVS] We show that if the source and channel are *matched* then one can achieve the sequential rate distortion bound.

1.3.5 Chapter 6

In this chapter we examine the stochastic control problem under a communication constraint. There is one communication channel connecting the sensor to the controller. We

first formulate the problem using the framework of chapter two. For the LQG problem we then provide sufficient conditions on the information pattern of the encoder and decoder to ensure the optimality of the certainty equivalent controller. This result generalizes the separation result proved by Borkar and Mitter. [BM] This separation property allows us to design the controller and the encoder and decoder separately. The optimal cost separates into two pieces: a full observation cost and a sequential rate distortion cost. Bansal and Basar showed that one could lower bound the performance by using information theoretic quantities. [BB] We provide a more general lower bound using the directed data processing inequality proved in chapter three. The idea of using the directed data processing inequality here comes from a more general lower bounding technique discussed by Mitter. [Mit] We show that this bound can be achieved if the channel is matched to the source.

Chapter 2

A General Framework for the Distributed Control Problem

2.1 Introduction

In this chapter we provide a general framework for modeling distributed control problems. The framework we present allows us to treat plants and communication channels on an equal footing. We can also treat controllers, channel encoders, and channel decoders on an equal footing. As a result we are able to present a conceptual view of the design of a large class of distributed control problems where sensors, controller, and actuators are interconnected through communication channels.

In general a model of a system consists of a set of stochastic kernels, an information pattern, and causality constraints codified by a time-ordering on the variables of interest. These stochastic kernels can represent plant transitions or channel transmissions. These different kernels provide a partial specification of a joint measure over all the random variables of interest. The job of a designer then is to “complete” this joint measure by providing appropriate interconnections between the given plants and channels. Note that this completion must respect the information pattern and causality constraints imposed by the model. In practice these interconnections will be provided through the design of a controller, encoder, or decoder.

The design of the optimal controller for the distributed system can be done in many ways. We distinguish between centralized design and decentralized design. Centralized design occurs when the design of a given controller is done with full knowledge of the other controllers. In decentralized design we do not assume this knowledge. We will show that centrally designed systems can be solved via dynamic programming.

In section 2.2 we provide our general framework. In section 2.3 we show that centrally designed control problem can be formulated as a dynamic programming problem. We conclude in section 2.4.

2.2 General Framework

We now give our general framework. We assume, through some sort of modeling process, that we are given partial knowledge of the system variables. This partial knowledge comes in the form of a set of stochastic kernels on the variables of interest. From this partial knowledge we want to complete a joint measure on these variables. There are many ways to complete this measure. We define a *model* of a system to be the set of all joint measures consistent with the given stochastic kernels. Thus a model is a subset of the set of all measures on the variables of interest. (See section A.1 of the appendix for a summary of stochastic kernels, Markov chains, and factoring joint measures.) We now give the formal definition of a model.

Given a sequence of variables x_1, \dots, x_T we use the notation $x^T \triangleq (x_1, \dots, x_T)$.

2.2.1 Main Definition

Let (Ω, \mathcal{F}) be a measure space. For $t = 1, \dots, T$ let Z_t be a measurable function from (Ω, \mathcal{F}) to the Polish space $(\mathcal{Z}_t, \mathcal{B}(\mathcal{Z}_t))$. (Note that T may equal infinity.) These variables, $\{Z_t\}$, represent different objects in the model. An object, for example, may be the state of a system at a given time, the control at a given time, the channel input at a given time, etcetera. We call them “variables” as opposed to “random variables” because we have not defined a measure on them. Specifically the Z_t ’s will be random variables only after we define a measure on (Ω, \mathcal{F}) . We now define a *model*:

Definition 2.2.1 *A model, \mathcal{M} , is a subset of the set of all measures on the variables Z_1, \dots, Z_T . This subset, $\mathcal{M} \subset \mathcal{P}(Z_1, \dots, Z_T)$, is defined by a three-tuple: causal ordering, system specification, and information pattern.*

(1) *Causal Ordering: This is a specification of a time-ordering on the variables Z_1, \dots, Z_T .*

Note that any measure $P(dZ^T)$ can be factored with respect to this ordering:

$$P(dZ^T) = \bigotimes_{t=1}^T P(dZ_t \mid z^{t-1}).$$

We need to specify these factors. There are two kinds of factors $P(dZ_t \mid z^{t-1})$. Those that are specified by the system and those that are specified by the designer. (For example state versus control.) Separate the set $\{1, \dots, T\}$ into these two disjoint sets:

(a) *The system set: $I = \{i_1, \dots, i_K\}$ with $1 \leq i_1 < \dots < i_K \leq T$*

(b) *The decision set: $J = \{j_1, \dots, j_L\} = \{1, \dots, T\} \setminus I$ with $1 \leq j_1 < \dots < j_L \leq T$.*

(2) *System specification: This is a set of stochastic kernels of the form $\{Q(dZ_{i_k} \mid z^{i_k-1})\}_{k=1}^K$. For all $i_k \in I$ the measure $P \in \mathcal{M}$ must satisfy*

$$P(dZ_{i_k} \mid Z^{i_k-1} = z^{i_k-1}) = Q(dZ_{i_k} \mid z^{i_k-1}) \quad P(dZ^{i_k-1}) - a.s.$$

(Where $P(dZ_{i_k} \mid Z^{i_k-1})$ is the conditional probability under P .)

(3) *Information pattern:* The information pattern defines what each decision maker $j_l \in J$ can base his decision on. Specifically for each $j_l \in J$ the information pattern of decision maker j_l is a subset $J_{j_l} \subseteq \{1, \dots, j_l - 1\}$. The decision kernel $Q(dZ_{j_l} | z^{j_l-1})$ can only be a function of $\{z_t : t \in J_{j_l}\}$. In other words,

$$Q(dZ_{j_l} | z^{j_l-1}) = Q(dZ_{j_l} | \tilde{z}^{j_l-1}) \quad \forall z^{j_l-1}, \tilde{z}^{j_l-1} \text{ such that } z_t = \tilde{z}_t \quad \forall t \in J_{j_l}.$$

For all $j_l \in J$ the measure $P \in \mathcal{M}$ must satisfy

$$P(dZ_{j_l} | Z^{j_l-1} = z^{j_l-1}) = Q(dZ_{j_l} | z^{j_l-1}) \quad P(dZ^{j_l-1}) - a.s.$$

The stochastic kernels chosen in part (3) allow us to interconnect the stochastic kernels given in part (2). This interconnection specifies a joint measure. The stochastic kernels chosen in part (3) are called *policies*. Note that an element in \mathcal{M} is uniquely specified, almost surely, once we specify the policies.

Our definition of a model closely parallels Witsenhausen's model. [Wit] In our model, though, we emphasize the two different kinds of stochastic kernels that make up the joint measure: those specified by the system and those that are policies.

Lemma 2.2.1 *Under any measure $P \in \mathcal{M}$ and $\forall j_l \in J$ the following forms a Markov chain:*

$$Z_{j_l} - \{Z_t : t \in J_{j_l}\} - \{Z_t : t \in \{1, \dots, j_l - 1\} \setminus J_{j_l}\}$$

Proof: Note that $P(dZ^{j_l}) = P(dZ_{j_l} | z^{j_l-1}) \otimes P(dZ^{j_l-1})$. By the information pattern constraint we know $P(dZ_{j_l} | Z^{j_l-1} = z^{j_l-1}) = Q(dZ_{j_l} | z^{j_l-1}) \quad P(dZ^{j_l-1}) - a.s.$ Thus $P(dZ^{j_l}) = P(dZ_{j_l} | \{z_t : t \in J_{j_l}\}) \otimes P(dZ^{j_l-1})$. By definition A.1.2 this is a Markov chain. \square

The following example will elucidate the previous definition. It is an abstract description of the models used in partially observed Markov decision problems.

Example 2.2.1 *The variables of interest are $X_1, \dots, X_T, Y_1, \dots, Y_T$ and U_1, \dots, U_T . The X 's are state variables, the Y 's are observation variables, and the U 's are control variables. The causal ordering is*

$$X_1, Y_1, U_1, \dots, X_T, Y_T, U_T.$$

The system specification is given by

$$\{Q(dX_t | x^{t-1}, y^{t-1}, u^{t-1})\}_{t=1}^T \text{ and } \{Q(dY_t | x^t, y^{t-1}, u^{t-1})\}_{t=1}^T.$$

Where $X_t - (X_{t-1}, U_{t-1}) - (X^{t-2}, Y^{t-1}, U^{t-2})$ and $Y_t - X_t - (X^{t-1}, Y^{t-1}, U^{t-1})$ are Markov chains. These are descriptions of the plant and the observation mechanism.

The information pattern specifies $U_t - (Y^t, U^{t-1}) - X^t$ to be a Markov chain. That is the control U_t is allowed to observe only the past observations and controls and not the past

states. A policy, then, is a sequence of stochastic kernels $\{Q(dU_t|x^t, y^t, u^{t-1})\}_{t=1}^T$ such that

$$Q(dU_t|x^t, y^t, u^{t-1}) = Q(dU_t|\tilde{x}^t, y^t, u^{t-1}) \quad \forall x^t, \tilde{x}^t.$$

2.2.2 The Finite Horizon Control Problem

So far we have defined a model. The model, though, only tells us what policies are acceptable. It does not tell us how to choose a policy. To do that we need an objective. Here we define the control objective in terms of minimizing an objective function.

Definition 2.2.2 A cost is an integrable function $f : \prod_{t=1}^T \mathcal{Z}_t \rightarrow \mathbb{R}^+$. The control problem consists of computing the cost

$$\inf_{P \in \mathcal{M}} E_P (f(Z_1, \dots, Z_T)) \tag{2.1}$$

and finding the infimizing P .

Example 2.2.1 continued For the partially observed Markov decision problem the cost is usually of the form

$$f(X^T, Y^T, U^T) = \sum_{t=1}^T c(X_t, U_t)$$

for some “running cost” c .

We are left now with the question of how to compute the optimal solution to the control problem. The solution depends on whether we allow centralized or decentralized design. Roughly speaking centralized design occurs if decision maker j_l can decide $Q(dZ_{j_l} | z^{j_l-1})$ based on the policies of all the other decision makers and complete knowledge of the system specifications. In decentralized design decision maker j_l does not have access to all of the other decision maker’s policies or all of the system specifications.

In the next subsection we discuss system and policy knowledge. Then we define centralized and decentralized design.

2.2.3 System and Policy Knowledge

One component of our model is the information pattern of each decision maker. This information pattern specifies who knows what and when. But that “what” refers to knowledge of the actual signals in the system and not to the knowledge of the different stochastic kernels that make up the joint measure. In computing the solution to (2.1) we need to determine whether the design of the system, i.e. the specification of the decision kernels, can be done in a centralized manner or a decentralized manner. To get a handle on this we first define what we mean by system and policy knowledge:

Definition 2.2.3 *The system knowledge, \mathcal{K}_j , of decision maker j is a subset of*

$$\mathcal{K}_j \subseteq \{Q(dZ_{i_k} \mid z^{i_k-1}), k = 1, \dots, K\}.$$

The policy knowledge, \mathcal{L}_j , of decision maker j is a subset of the policies

$$\mathcal{L}_j \subseteq \{Q(dZ_{j_l} \mid z^{j_l-1}), l = 1, \dots, L\}.$$

We further insist that decision maker j knows its own policy: $Q(dZ_j \mid z^{j-1}) \in \mathcal{L}_j$.

Definition 2.2.4 *We say decision maker j has complete system knowledge if*

$$\mathcal{K}_j = \{Q(dZ_{i_k} \mid z^{i_k-1}), k = 1, \dots, K\}.$$

And we say decision maker j has complete policy knowledge if

$$\mathcal{L}_j = \{Q(dZ_{j_l} \mid z^{j_l-1}), l = 1, \dots, L\}.$$

Definition 2.2.5 *We say a system is centrally designed if every decision maker $j_l \in J$ has complete system knowledge and complete policy knowledge. Otherwise call the system decentrally designed.*

In section 2.3 we will show that the centrally designed control problem can be formulated as a dynamic programming problem.

2.2.4 Partial Orderings on the Signal and System Knowledge

We have defined the information pattern and the system and policy knowledge. We now show that there exists a natural partial ordering on the information patterns and the system and policy knowledge.

Recall that an information pattern is defined by the set of sets: $\{J_{j_l} : l = 1, \dots, L\}$. Denote this by $\mathcal{J} \triangleq \{J_{j_l} : l = 1, \dots, L\}$.

Definition 2.2.6 *We say that $\mathcal{J} \preceq_I \tilde{\mathcal{J}}$ if*

$$J_{j_l} \subseteq \tilde{J}_{j_l} \quad \forall l = 1, \dots, L.$$

(The “I” in \preceq_I stands for “information pattern.”)

It is straightforward to show that \preceq_I is a partial order:

- (1) $\mathcal{J} \preceq_I \mathcal{J}$
- (2) $\mathcal{J} \preceq_I \tilde{\mathcal{J}}$ and $\tilde{\mathcal{J}} \preceq_I \mathcal{J}$ imply $\mathcal{J} = \tilde{\mathcal{J}}$
- (3) $\mathcal{J} \preceq_I \tilde{\mathcal{J}}$ and $\tilde{\mathcal{J}} \preceq_I \hat{\mathcal{J}}$ imply $\mathcal{J} \preceq_I \hat{\mathcal{J}}$

Furthermore there is a unique maximal element in this partial ordering. Specifically $\mathcal{J}^* = \{J_{j_l}^* : l = 1, \dots, L\}$. Where each $J_{j_l}^* = \{1, \dots, j_l - 1\}$.

It should be clear that when computing an optimal policy more information cannot increase the optimal cost (recall that an “optimal policy” is a policy that minimizes a given cost.) If $\mathcal{J} \preceq_I \tilde{\mathcal{J}}$ then the optimal cost associated with information pattern \mathcal{J} can be no less than the optimal cost associated with information pattern $\tilde{\mathcal{J}}$. Thus an optimal policy for the maximal information pattern will lead to the minimal cost. It may be the case, though, that there is a smaller, in the sense of \preceq_I , information pattern that leads to the same cost. This is an interesting, though, difficult, question. We discuss this in chapter three.

We can also define a partial ordering on the system and policy knowledge. Recall that the system knowledge and policy knowledge are defined by the two sets of sets: $\{\mathcal{K}_{j_l} : l = 1, \dots, L\}$ and $\{\mathcal{L}_{j_l} : l = 1, \dots, L\}$ respectively. Define $\mathcal{S} \triangleq \{\mathcal{K}_{j_l}, \mathcal{L}_{j_l} : l = 1, \dots, L\}$.

Definition 2.2.7 *We say that $\mathcal{S} \preceq_K \tilde{\mathcal{S}}$ if*

$$\mathcal{K}_{j_l} \subseteq \tilde{\mathcal{K}}_{j_l} \quad \text{and} \quad \mathcal{L}_{j_l} \subseteq \tilde{\mathcal{L}}_{j_l} \quad \forall l = 1, \dots, L.$$

(The “K” in \preceq_K stands for “system and policy knowledge.”)

It is straightforward to show that \preceq_K is a partial order. Note that the case where each decision maker has complete system knowledge and complete policy knowledge is the maximal element in this partial ordering. (See definition 2.2.4.)

2.2.5 Extensions

In this section we describe two extensions to our main formulation. The first extension involves conditions on when one can exchange variables in the given causal ordering. The second extension describes a specialization of the causal ordering on the variables of interest to a partial ordering on the variables of interest.

Conditions for Changing the Causal Ordering

Often there is flexibility in the choice of order in our causal ordering of variables.

Definition 2.2.8 *Let π be any permutation of $\{1, \dots, T\}$. Let $\tilde{Z}_t = Z_{\pi(t)}$. Then the models \mathcal{M} with ordering (Z_1, \dots, Z_T) and $\tilde{\mathcal{M}}$ with ordering $(\tilde{Z}_1, \dots, \tilde{Z}_T)$ are said to be equivalent if $\mathcal{M} = \tilde{\mathcal{M}}$. (Recall a model is a subset of the set of all measures on Z^T .)*

Proposition 2.2.1 *Assume we are given a model \mathcal{M} with causal ordering*


$$(Z_1^{t-1}, Z_t, Z_{t+1}, Z_{t+2}^T)$$

Furthermore assume that either the system specification or the information pattern satisfies the constraint $P(dZ_{t+1} | Z^t) = P(dZ_{t+1} | Z^{t-1})$. Then there exists an equivalent model $\tilde{\mathcal{M}}$ with causal ordering

$$(Z_1^{t-1}, Z_{t+1}, Z_t, Z_{t+2}^T)$$

Proof: Using the same idea as in lemma 2.2.1 one can show that for any measure $P \in \mathcal{M}$ the following $Z_{t+1} - Z^{t-1} - Z_t$ forms a Markov chain. Thus $P(Z_t | Z^{t-1}, Z_{t+1}) = P(Z_t | Z^{t-1})$. We can construct a new model $\tilde{\mathcal{M}}$ with causal ordering $(Z_1, \dots, Z_{t-1}, Z_{t+1}, Z_t, Z_{t+2}, \dots, Z_T)$. The system specification and information pattern stay the same. \square

The following diagram uses the notation of directed graphical models to show the result in proposition 2.2.1. [Pea] It shows that Z_t and Z_{t+1} can be interchanged in time since neither influences the other when conditioned on Z^{t-1} .



Corollary 2.2.1 *Let τ be a positive integer. Assume we are given a model \mathcal{M} with causal ordering*

$$(Z_1^{t-1}, Z_t^{t+\tau}, Z_{t+\tau+1}^T)$$

Furthermore assume that either the system specification or the information pattern satisfies the constraints

$$P(dZ_{t+i} | Z^{t+i-1}) = P(dZ_{t+i} | Z^{t-1}) \quad \forall i = 1, \dots, \tau.$$

Then there exists an equivalent model $\tilde{\mathcal{M}}$ with causal ordering

$$(Z_1^{t-1}, \tilde{Z}_0^\tau, Z_{t+\tau+1}^T)$$

where $(\tilde{Z}_0, \dots, \tilde{Z}_\tau)$ is any permutation of $(Z_t, \dots, Z_{t+\tau})$.

Proof: This follows from repeated use of proposition 2.2.1. \square

Proposition 2.2.1 and corollary 2.2.1 show that our model can treat systems where certain objects occur simultaneously. For example proposition 2.2.1 shows there is no difference in assuming Z_t or Z_{t+1} occurred first. And in fact they may occur simultaneously. Thus our model of a system can treat simultaneous events if appropriate conditional independencies are assumed. These results, though, are more naturally stated if we specialize our causal ordering to a partial ordering. We do that now.

Causal Ordering Versus Partial Ordering of Z^T


We mention one last extension to our model before presenting the dynamic programming formulation. Recall that our model is defined by a three-tuple: a causal order, a system specification, and an information pattern. We can specialize the causal order on (Z_1, \dots, Z_T) to a partial order. A partial order, as we will see, incorporates added information about different conditional independencies.

We say a partial order is a specialization of a linear order because for any given partial order we can “force” a linear order. Let \mathcal{N} be an at most countable index set with partial order \preceq . If $n \prec \tilde{n}$ then we say Z_n occurs before $Z_{\tilde{n}}$.

Lemma 2.2.2 *Given a partial order, \preceq , on \mathcal{N} we can define a linear order on \mathcal{N} that preserves the time ordering of the partial order.*

Proof: Let $A = \{n \in \mathcal{N} : n \text{ is a minimal element}\}$. Label the elements of A by $1, \dots, |A|$. Let $B = \{n \in \mathcal{N} : \text{the shortest path from } n \text{ to a minimal element is of length one}\}$. Label the elements of B by $|A|+1, \dots, |A|+|B|$. Define $C = \{n \in \mathcal{N} : \text{the shortest path from } n \text{ to a minimal element is of length two}\}$. Label the elements of C by $|A| + |B| + 1, \dots, |A| + |B| + |C|$. Now continue in this manner. We have constructed a linear order that preserves the time-ordering of the original partial order. \square

A picture may help elucidate the construction. Here $a \rightarrow b$ means $a \prec b$:



Note that the linear order we have constructed introduces many new order relations between the elements of \mathcal{N} . But it still preserves the original ordering of the partial order.

We now specialize definition 2.2.1.

Definition 2.2.9 *Let \mathcal{N} be an at most countable index set. Let $\{Z_n : n \in \mathcal{N}\}$ be the set of variables of interest. A model, \mathcal{M} , is a subset of the set of all measures on the variables $\{Z_n : n \in \mathcal{N}\}$. The model is defined by a three-tuple: partial ordering, system specification, and information pattern.*

- (1) *Partial Ordering: This is a specification of a partial-order, \preceq , on \mathcal{N} . If $n \prec \tilde{n}$ then we say Z_n occurs before $Z_{\tilde{n}}$.*

We further assume that our measure $P(\{dZ_n : n \in \mathcal{N}\})$ can be factored with respect to this partial ordering:

$$P(\{dZ_n : n \in \mathcal{N}\}) = \bigotimes_{n \in \mathcal{N}} P(dZ_n \mid \{z_{\tilde{n}} : \tilde{n} \prec n\}).$$

As before there are two kinds of factors $P(dZ_n \mid \{z_{\tilde{n}} : \tilde{n} \preceq n \text{ and } n \neq \tilde{n}\})$. Separate the set \mathcal{N} into these two disjoint sets:

(a) The system set: $I \subseteq \mathcal{N}$

(b) The decision set: $J = \mathcal{N} \setminus I$

(2) *System specification:* For each $n \in I$ we are given a stochastic kernel of the form $Q(dZ_n \mid \{z_{\tilde{n}} : \tilde{n} \prec n\})$. For all $n \in I$ the measure $P \in \mathcal{M}$ must satisfy

$$P(dZ_n \mid \{Z_{\tilde{n}} = z_{\tilde{n}} : \tilde{n} \prec n\}) = Q(dZ_n \mid \{z_{\tilde{n}} : \tilde{n} \prec n\}) \quad P(\{dZ_{\tilde{n}} : \tilde{n} \prec n\}) - a.s.$$

(3) *Information pattern:* For each $n \in J$ the information pattern of decision maker n is a subset $J_n \subseteq \{\tilde{n} : \tilde{n} \prec n\}$. The decision kernel $Q(dZ_n \mid \{z_{\tilde{n}} : \tilde{n} \prec n\})$ can only be a function of the information contained in $\{z_{\tilde{n}} : \tilde{n} \in J_n\}$. For all $n \in J$ the measure $P \in \mathcal{M}$ must satisfy

$$P(dZ_n \mid \{Z_{\tilde{n}} = z_{\tilde{n}} : \tilde{n} \prec n\}) = Q(dZ_n \mid \{z_{\tilde{n}} : \tilde{n} \prec n\}) \quad P(\{dZ_{\tilde{n}} : \tilde{n} \prec n\}) - a.s.$$

Summary

In this section we have defined a general model for treating distributed systems. We have defined the concepts of information pattern and system and policy knowledge. We showed that systems defined on partial orders are special cases of our formulation. We now discuss the dynamic programming formulation for centrally designed systems.

2.3 Dynamic Programming Formulation

In this section we describe the dynamic programming formulation for centrally designed systems. For the rest of this section we will assume that all the random variables $\{Z_t\}$ are finite valued random variables. This allows the interchange of infimizations and expectations in the Bellman recursion to be well-defined. See [BS] for more general conditions that allow this interchange.

By Bellman's principle of optimality we may write the optimization in equation (2.1) as

$$\begin{aligned} & \inf_{P \in \mathcal{M}} E_P (f(Z_1, \dots, Z_T)) \\ = & E_P \left[\inf_{\{Q(Z_{j_1} | z^{j_1-1})\}} E_P \left[\dots \inf_{\{Q(Z_{j_{L-1}} | z^{j_{L-1}-1})\}} \right. \right. \\ & \left. \left. E_P \left[\inf_{\{Q(Z_{j_L} | z^{j_L-1})\}} E_P \left[f(Z_1, \dots, Z_T) \mid z^{j_L-1} \right] \mid z^{j_{L-1}-1} \right] \dots \mid z^{j_1-1} \right] \right]. \end{aligned}$$

Let us examine the j_l th infimization. The decision maker at time j_l needs to compute

$$\inf_{\{Q(Z_{j_l} | z^{j_l-1})\}} E_P \left[\dots \inf_{\{Q(Z_{j_L} | z^{j_L-1})\}} E_P \left[f(Z_1, \dots, Z_T) \mid z^{j_L-1} \right] \dots \mid z^{j_l-1} \right].$$

Recall the information pattern for decision maker j_l is specified by the set J_{j_l} . The decision maker has knowledge only of the signals in $\{z_t : t \in J_{j_l}\}$. The j_l th infimization can be rewritten as

$$\inf_{\{Q(Z_{j_l} | z^{j_l-1})\}} E_P \left\{ E_P \left[\dots \inf_{\{Q(Z_{j_L} | z^{j_L-1})\}} E_P \left[f(Z_1, \dots, Z_T) \mid z^{j_L-1} \right] \dots \mid z^{j_l-1} \right] \mid \{z_t : t \in J_{j_l}\} \right\}.$$

In order to compute the inside expectations we need to know the conditional measure $P(Z^T \mid z^{j_l-1})$ and in order to compute the outside expectation we need to know the conditional measure $P(Z^{j_l-1} \mid \{z_t : t \in J_{j_l}\})$. We treat each case now:

- (1) $P(Z^T \mid z^{j_l-1}) = Q(Z_{j_l} \mid z^{j_l-1}) \otimes \left(\bigotimes_{t=j_l+1}^T P(Z_t \mid z^{t-1}) \right)$. This product consists of system kernels and decision kernels. The decision maker has complete system knowledge. Furthermore since dynamic programming is a backward recursion the decision maker knows what the future polices will be also. Thus the decision maker knows the conditional probability $P(Z^T \mid z^{j_l-1})$.
- (2) In order to compute $P(Z^{j_l-1} \mid \{z_t : t \in J_{j_l}\})$ we will first compute $P(Z^{j_l-1})$. Now $P(Z^{j_l-1}) = \bigotimes_{t=1}^{j_l-1} P(Z_t \mid z^{t-1})$. This product consists of past system kernels and past decision kernels. The decision maker has complete system knowledge. But unlike the case of future decision kernels it cannot compute the past decision kernels. To

get around this problem we assume that the decision maker has access to the past decision kernels. Specifically we assume that the state of the dynamic programming recursion at time j_l consists of the pair

$$\left(\{z_t : t \in J_{j_l}\}, \left\{ Q \left(Z_j \mid z^{j-1} \right) : j = j_1, \dots, j_{(l-1)} \right\} \right).$$

Now the decision maker can compute $P(Z^{j_l-1})$. Consequently it can compute $P(Z^{j_l-1} \mid \{z_t : t \in J_{j_l}\})$.

Note that at the end of the dynamic programming recursion decision maker j_1 will have chosen a policy kernel $Q^*(Z_{j_1} \mid z^{j_1-1})$. (Where * means optimal.) The optimal policy for decision maker j_2 depends on $Q^*(Z_{j_1} \mid z^{j_1-1})$. But this is part of j_2 's information. Thus we see that at the end of the backward dynamic programming recursion we need to make a forward pass to substitute in the past optimal policies.

We have just shown that centrally designed systems with arbitrary information patterns can be solved via dynamic programming. Conceptually the dynamic program is straightforward. In practice though it may be very complicated to implement. Furthermore even if the objective function is convex the intermediate optimizations in the above dynamic program may be non-convex. See for example Witsenhausen's counterexample paper. [Wit3]

If the information pattern of decision maker j_l contains the whole past, i.e. $J_{j_l} = \{1, \dots, j_l - 1\}$, then the decision maker does not need to know the past decision kernels. This follows because $P(Z^{j_l-1} \mid z^{j_l-1})$ is just a Dirac measure. This is a special case of Witsenhausen's work on policy independence of conditional expectations. [Wit2] This is also the traditional case, i.e. full state observation, dealt with in dynamic programming.

At the other extreme we can imagine a case where the information pattern for decision maker j_l is empty: $J_{j_l} = \emptyset$. This means that there is no signal feedback to the decision maker. In this case the decision maker needs to know all of the past decision kernels.

There are many cases in between these two extreme cases. Clearly there is a tradeoff between signal knowledge, as captured by the information pattern, and system and policy knowledge. We have already shown that there exist natural partial orders on the information pattern, system knowledge, and policy knowledge. We discuss some of the interactions between differing signal and system knowledge in chapter 3.

We now continue example 5.2.1 of the partially observed Markov decision problem.

Example 5.2.1 continued *In this case the information pattern for decision maker U_t is (Y^t, U^{t-1}) . Thus the infimization in the Bellman recursion requires us to compute*

$$P(X^t, Y^t, U^{t-1} \mid y^t, u^{t-1}).$$

As is well known, if the system dynamics are Markov and the cost additive then the probability of the current state given the past information is a sufficient statistic for the problem. Thus we need only compute $P(X_t \mid y^t, u^{t-1})$. But this is the usual filter for estimating the state:

$$P(x_{t+1} | y^{t+1}, u^t) = \frac{P(y_{t+1} | x_{t+1}) \sum_{x_t} P(x_{t+1} | x_t, u_t) P(x_t | y^t, u^{t-1})}{\sum_{x_{t+1}} P(y_{t+1} | x_{t+1}) \sum_{x_t} P(x_{t+1} | x_t, u_t) P(x_t | y^t, u^{t-1})}.$$

To compute this recursion the decision maker needs to know the values of the past controls, $\{u_t\}$, and the system kernels: $\{P(Y_t | x_t)\}$ and $\{P(X_{t+1} | x_t, u_t)\}$. In this case we do not need to augment the state of the dynamic programming recursion with information about the past policies. This is because the decision maker at time t has access to all the past controls u^{t-1} and hence does not need to integrate them out.

2.3.1 Difficulties with Decentralized Design

We have shown that dynamic programming can be used to design systems in a centralized way. Designing systems decentrally is much harder. This is due in part because there is no obvious analog of dynamic programming to apply here. In the centrally designed case the decision makers are coordinated in their design effort. There is an inherent game-theoretic issue when the decision makers don't have complete system or policy knowledge. We address a particular aspect of this issue in chapter three. It is our belief, though, that the above formulation may help in the characterization and solution of certain classes of decentrally designed distributed systems. But we leave this for future work.

There is the possibility that the different decision makers can learn the other decision makers' policies over time. This can occur, for example, if each decision maker knows the objective function and has the proper information pattern. The proper information pattern allows the decision maker the opportunity to learn another decision maker's policy through input/output data. The knowledge of the objective function and an assumption of rational design can also help the decision maker learn the other decision makers' policies. We leave the possibility of learning to future work.

There are three different kinds of knowledge in distributed systems. They are the information pattern, the system and policy knowledge, and the knowledge of the objective. Said succinctly they are knowledge of the signals, the system, and the goal. Each type of knowledge or lack of knowledge can contribute a different kind of complexity to the design.

Many distributed systems are designed in a centralized manner. Most of the systems in this thesis are centrally designed. We have shown that the issue of nontraditional information pattern has been, at least conceptually, taken care of provided one is allowed to centrally design the system. Thus one may choose to rethink our popular notion of a distributed system to not be one with nontraditional information pattern but to be one that is designed decentrally. That is a system where the decision makers do not have complete system or policy knowledge.

The third sort of complexity occurs if one assumes that each decision maker has different knowledge about the control objective or have different control objectives altogether. This added complexity will not be dealt with here.

In summary there are three kinds of knowledge: signal knowledge in the form of an information pattern, system and policy knowledge, and control objective knowledge.

2.4 Summary

In this chapter we introduced a general formulation for modeling distributed control problems. We examined the roles that information patterns and system and policy knowledge play in this framework. We showed that centrally designed systems with arbitrary information pattern can be solved, in principle, via dynamic programming.

We view a model as the set of all measures that complete a partially specified joint measure. And we view control as a selection of one of these consistent joint measures. Another way to view control is as the interconnection, under differing information patterns, of different stochastic kernels representing the system.

In this thesis we will apply this formulation to the problems of channel coding with feedback, sequential rate distortion, and control when there is a communication link connecting the sensor to the controller. For all these problems we examine the situation where centralized design is allowed. For the control problem we also examine some issues of decentralized design.

Chapter 3

Control of Deterministic Systems Under Communication Constraints

3.1 Introduction

In this chapter we examine the deterministic control problem under communication constraints. The reason we first examine the deterministic case instead of considering the stochastic case, covered in chapter 6, is that many salient issues can be brought out in this “simpler” setting. We examine traditional control properties, e.g. observability, stability, controllability, as well as some performance issues for this class of problems. We almost exclusively study the discrete time case. A discrete time model is more consistent with today’s digital communication links. The communication constraint we analyze in this chapter is a discrete time noiseless digital channel capable of transmitting R bits per time step.

In section 3.2 we discuss the problem setup in the context of a control problem where the information pattern as defined in the previous chapter can be quite general. In section 3.3 we provide lower bounds on the rate required to achieve observability, stability, and controllability. These bounds are independent of the information pattern chosen and the system and policy knowledge chosen.

In section 3.4 we provide necessary background for discussing schemes that upper bound the rate required to achieve the given control objectives. Here we describe the different encoder classes of interest to us and the primitive quantizer. We also provide the key technical lemma which relates a measure of the system’s growth to the channel rate. In section 3.5 we discuss encoder class one. Encoder class one, as we will show, has the “best” information pattern and system and policy knowledge. We will show that many of the lower bounds of section 3.3 are achievable under encoder class one. In section 3.6 we discuss encoder class two. This encoder class has a more realistic information pattern and system and policy knowledge. The rates, though, required to achieve the control objectives are larger than those in encoder class one.

In section 3.7 we discuss the case where there are multiple sensors. We relate this problem to the Slepian-Wolf coding problem. In section 3.8 we comment on the sampling of continuous time systems. In section 3.9 we discuss some performance criterion most

notably the linear quadratic cost criterion. We also introduce the concept of covering numbers. Finally we end with a discussion in section 3.10.

In summary there are two main contributions in this chapter. First we compute a lower bound on the rate required to achieve different control objectives. This lower bound is independent of the information patterns in place and depends only the plant. Second we give conditions on the information pattern for achieving this lower bound. We examine the rate in cases where we have different information patterns.

We end this introduction with three observations that will motivate our analysis. They are not meant to be rigorous statements but rather they are meant to guide our thinking.

Observation 1: Why feedback If there is no uncertainty in the initial position, no uncertainty in the plant dynamics, and there are no process disturbances then one can achieve most control objectives using an open loop controller. A closed loop controller for the same problem is often less complex to realize. Furthermore a closed loop controller can more robustly deal with the aforementioned uncertainties in initial position, plant dynamics and process disturbances. Thus the point of feedback, if we bar complexity considerations, is to transmit from the plant to the controller information about the state of the plant and the plant itself that the controller does not know. The question then becomes what information is relevant and what communication scheme should be used to transmit that information.

Observation 2: Full observation performance If the observation mechanism is instantaneous and lossless then we call the observation a full observation. We assume the control objective of interest is achievable under full observation. Clearly if an objective cannot be achieved under full observation it cannot be achieved under the rate constrained observation. Conversely if a control objective can be achieved under a rate constrained observation then it can be achieved under full observation.

Observation 3: Number of control sequences In a time horizon T the decoder will receive one of at most 2^{TR} channel symbol sequences. If the encoder, decoder, and controller are all deterministic then the number of different possible control sequences in this time must be smaller than or equal to 2^{TR} . Intuitively then a control objective under rate R can be achieved only if we can approximate well the control sequences for the full observation problem by one of only 2^{TR} control sequences. Thus in terms of the underlying quantization problem one may think of quantization as living in the control sequence space.

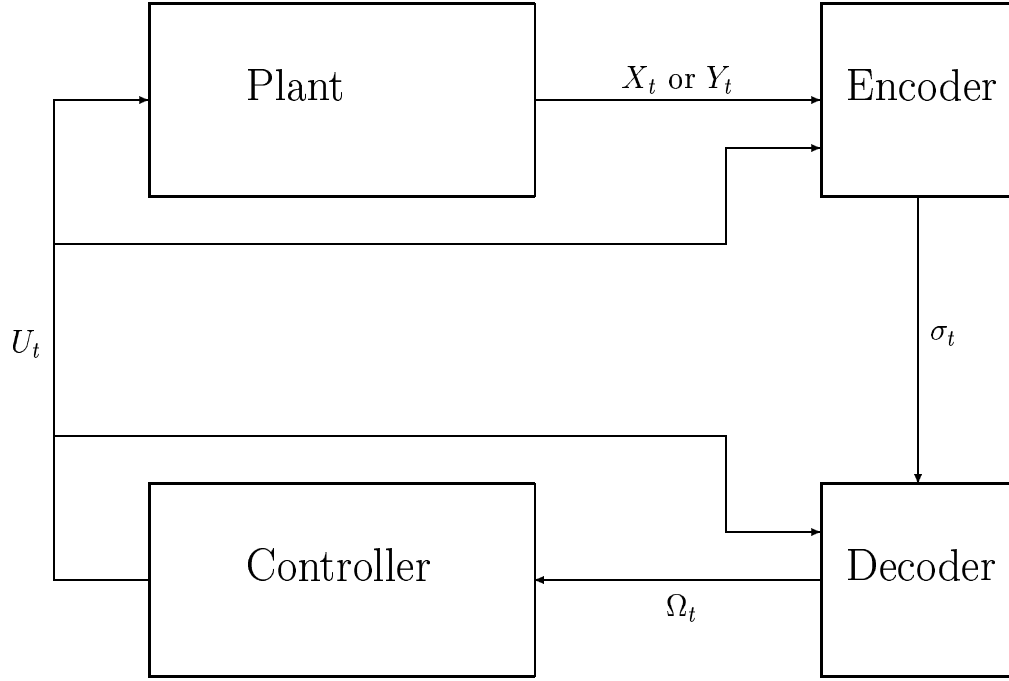


Figure 3-1: System

3.2 Problem Setup

Throughout this chapter we consider the following linear time-invariant system:

$$X_0 \in \Lambda_0, \quad X_{t+1} = AX_t + BU_t + W_t, \quad Y_t = CX_t \quad \forall t \geq 0 \quad (3.1)$$

where $\{X_t\}$ is a \mathbb{R}^d -valued state process, $\{U_t\}$ is a \mathbb{R}^m -valued control process, and $\{Y_t\}$ is a \mathbb{R}^l -valued observation process. The sequence $\{W_t\}$ is a \mathbb{R}^d -valued disturbance process with $\|W_t\|_2 \leq D \forall t$. We see that $A \in \mathbb{R}^{d \times d}$, $B \in \mathbb{R}^{d \times m}$, and $C \in \mathbb{R}^{l \times d}$. The initial position is $X_0 \in \Lambda_0$ where $\Lambda_0 \subseteq \mathbb{R}^d$. If $C = I$, where I is the identity matrix, then we have full state observation. See Figure 3-1.

In this chapter we consider a noiseless digital channel that can transmit at each time step one of 2^R symbols $\sigma \in \Sigma$. Note this is a hard rate constraint and not a time average rate constraint.

Convention 3.2.1 *Throughout this thesis we will allow R to take on real values.*

This convention allows for an easier analysis. To determine the rate in practice one can take the ceiling of R to get an integer.

3.2.1 Information Pattern and System Knowledge

The control problems we look at involve the design of the encoder, decoder, and controller. We must specify the information pattern of each component. Each component will implement a “policy.” Thus we further need to specify the system and policy knowledge of each component. These were discussed in chapter two. The information pattern can be considered “online information.” The *a priori* description of the policies can be considered “offline information.”

Endow \mathbb{R} with the usual Borel σ -field and Σ with the power set σ -field 2^Σ . Assume throughout this chapter that all maps between these measurable spaces are measurable.

We now provide a general description of the encoder, decoder, and controller.

Encoder:

The encoder at time t is a map

$$\mathcal{E}_t : \mathbb{R}^{l(t+1)} \times \Sigma^t \times \mathbb{R}^{mt} \rightarrow \Sigma$$

that takes

$$(Y^t, \sigma^{t-1}, U^{t-1}) \mapsto \sigma_t.$$

Note that we may restrict the inputs to the encoder. For example the encoder may not have access to the past controls. We will be more specific when we discuss the different encoder/decoder setups in section 3.4.

Decoder:

Let $\Omega_t \subset \mathbb{R}^d$. The decoder at time t is a map

$$\mathcal{D}_t : \Sigma^{t+1} \times \mathbb{R}^{mt} \rightarrow 2^{\mathbb{R}^d}$$

that takes

$$(\sigma^t, U^{t-1}) \mapsto \Omega_t.$$

The output of the decoder is a measurable set representing the uncertainty in the state estimate. In the parlance of partially observed control problems this is called the “information state.” We will formally define information state in the sequel. But for now it is sufficient to just consider it a set. Before we can describe how the information state is calculated we must specify what knowledge the decoder has of the encoder and controller. We will be more specific when we discuss the different encoder/decoder setups in section 3.4.

One may question why Ω_t , an information state, is the appropriate output for the decoder. Throughout this chapter we will assume that a separation structure between controller and decoder exists. This is a reasonable assumption as many partially observed control problems lend themselves to a separation theorem. (Note this separation theorem does not refer to the separation between source and channel coding but instead refers to the separation between the state estimator and controller.) We will show that for encoders in

encoder class one, to be defined in section 3.4, that there is no loss of generality in making this separation assumption.

Controller:

The controller at time t is a map

$$\mathcal{C}_t : 2^{\mathbb{R}^d} \rightarrow \mathbb{R}^m$$

that takes

$$\Omega_t U_t.$$

System and Policy Knowledge:

We assume that for all t the encoder, decoder, and controller at time t have knowledge of the dynamics of the plant. This knowledge is denoted $\mathcal{F} \triangleq \{A, B, C\}$. Similarly the knowledge of the other maps are denoted $\mathcal{E} \triangleq \{\mathcal{E}_t\}$, $\mathcal{D} \triangleq \{\mathcal{D}_t\}$, and $\mathcal{C} \triangleq \{\mathcal{C}_t\}$. We say, for example, that the encoder at time t has full decoder knowledge if it knows \mathcal{D} .

In this chapter one of our main interests will be in deciding whether the encoder should have knowledge of the controller or not. In fact this is one of the major distinctions between the two encoder classes that we define in section 3.4.

Control Objective Knowledge

The final piece of knowledge is the “control objective.” This knowledge becomes important when trying to optimize the design of the encoder, decoder, and controller. We assume that the encoder, decoder, and controller all have the same knowledge of the control objective.

3.3 Lower Bounds that are Independent of the Information Pattern

We now examine the control properties of observability, stability, and controllability under a rate constraint. Note that the usual algebraic conditions, e.g. certain Grammians having full rank, are necessary but no longer sufficient. Furthermore we will have to deal with asymptotic versions of observability and controllability.

In this subsection we provide lower bounds on the rate required to achieve the different control objectives. These lower bounds will be “universal” in the sense that they hold independently of the actual encoder, decoder, and controller used. That is they hold independently of the information pattern and system and policy knowledge chosen. One should note the analogy with Fano’s inequality used in converse theorems in information theory. Fano’s inequality holds independently of the actual encoder and decoder used. (For Fano’s inequality see the comments after lemma 4.4.6.)

We will show in section 3.5 that there exists an information pattern such that an encoder, decoder, and controller exist for which these lower bounds can be achieved. Thus the lower bounds are tight.

3.3.1 Observability

The purpose of any good observer is to distinguish points in the state space. In a time horizon of T we have at most 2^{TR} possible symbols arriving into the decoder. Thus we must be able to approximate the state by one of 2^{TR} points.

There are many examples of systems where the choice of control can effect the estimation error of a given observer. This is sometimes called the “dual effect” of control on state estimation. In other examples the ability to observe the state is independent of the control signal used.

First some definitions. Recall the information state Ω_t is the decoder output at time t .

Definition 3.3.1 *Let the state estimate at time t be $\hat{X}_t = \text{centroid}(\Omega_t)$. (Where “centroid” refers to the center of mass with respect to the uniform distribution.)*

Definition 3.3.2 *Let the error be $e_t = X_t - \hat{X}_t$ where \hat{X}_t is the state estimate.*

Definition 3.3.3 *System (3.1) is asymptotically observable if there exists a control sequence $\{u_t\}$ and an encoder and decoder such that*

1 *Stability: $\forall \epsilon > 0 \exists \delta(\epsilon)$ such that $\|X_0\|_2 \leq \delta(\epsilon)$ implies $\|e_t\|_2 \leq \epsilon \forall t \geq 0$.*

2 *Uniform attractivity: $\forall \epsilon > 0, \forall \delta > 0 \exists T(\epsilon, \delta)$ such that $\|X_0\|_2 \leq \delta$ implies $\|e_t\|_2 \leq \epsilon \ t \geq T$.*

Point one states that the error cannot grow without bound for bounded X_0 . The second point states that the error decreases to zero uniformly in X_0 . Note also that uniform attractivity is defined for all δ . Thus our definition of asymptotic observability is global.

Definition 3.3.4 *System (3.1) is uniform in control asymptotically observable if there exists an encoder and decoder independent of the control sequence applied such that the system is asymptotically observable.*

Note that these definitions differ from the usual definition of observability. The usual definition for deterministic systems without disturbances states that given enough time one can identify the initial condition exactly. The usual prescription is that once you know the initial condition and the controls one can compute the state at any time. In our case we can only distinguish between 2^{tR} initial positions in time t . It is for this reason that we introduce the definition of asymptotic observability.

In the proofs that follow we will provide conditions for $\text{diam}(\Omega_t) \rightarrow 0$. Clearly the diameter of Ω_t going to zero implies the error goes to zero. We choose to work with the information state because it is a natural object to work with in our approach to analyzing observability.

Now we are prepared to give a necessary condition on the rate required to achieve asymptotic observability.

Proposition 3.3.1 *Given system (3.1) a necessary condition on the rate required so that the system is uniform in control asymptotically observable is $R \geq \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$.*

Proof: Assume without loss of generality that the initial uncertainty contains the bounded set $\Lambda_0 = \{X : \|X\|_\infty \leq \frac{L}{2}\}$, that $A = \text{diag}[\lambda_1, \dots, \lambda_d]$ is a diagonal matrix, and that there are no disturbances. We will provide a lower bound on the rate required for asymptotic observability.

Fix an arbitrary control sequence U_0, \dots, U_{t-1} . The set of points that X_t can take contains the following set

$$A^t \Lambda_0 - \alpha_t = \left\{ X : X - \alpha_t \in \left[-\frac{L}{2}|\lambda_1|^t, \frac{L}{2}|\lambda_1|^t\right] \times \dots \times \left[-\frac{L}{2}|\lambda_d|^t, \frac{L}{2}|\lambda_d|^t\right] \right\}$$

where $\alpha_t = \sum_{j=0}^{t-1} A^{t-1-j} B U_j$.

To say that the system is uniform in control asymptotically observable means that for every $\epsilon > 0$ there is a $T(\epsilon, L)$ such that for $t \geq T(\epsilon, L)$ we have $\|e_t\|_2 \leq \epsilon \forall X_0 \in \Lambda_0, \forall \{U_t\}$.

A lower bound on the rate can be computed by counting the number of regions it takes to cover $A^t \Lambda_0 - \alpha_t$ by regions of diameter less than 2ϵ . If the diameter of the uncertainty set Ω_t goes to zero then the volume of Ω_t goes to zero. The converse, though, is not necessarily true. However the converse is true if we restrict ourselves to computing the volume of the projection of the uncertainty set onto the unstable subspace. The diameter of the uncertainty set goes to zero if and only if the volume of the uncertainty set projected on the unstable subspace goes to zero.

Let $\lambda_1, \dots, \lambda_u$ be the unstable eigenvalues. Given a region with diameter 2ϵ project it onto the unstable subspace. The largest volume this projected region can have is that given by a u -dimensional sphere of diameter 2ϵ . Specifically the projected volume is $\leq K_u \epsilon^u$ where K_u is the constant in the formula for the volume of a sphere.

Thus to cover $A^t \Lambda_0 - \alpha_t$ by regions of diameter 2ϵ we require at least

$$\begin{aligned}
R &\geq \frac{1}{t+1} \log \frac{L^u |\lambda_1 \dots \lambda_u|^t}{K_u \epsilon^u} \\
&= \sum_{i=1}^u \log |\lambda_i| + \frac{u}{t+1} \log \frac{L}{\epsilon} - \frac{1}{t+1} \log (\lambda_1 \dots \lambda_u) K_u \\
&= \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\} + \frac{u}{t+1} \log \frac{L}{\epsilon} - \frac{1}{t+1} \log |\lambda_1 \dots \lambda_u| K_u
\end{aligned}$$

the second term is positive and for t large the third term becomes negligible thus

$$R \geq \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}.$$

Note that if $\epsilon(t)$ is allowed to shrink sub-exponentially with t then the second term goes to zero. \square

3.3.2 Stability

In this section we discuss stability under a rate constraint. The lower bound uses a counting argument similar to that given in proposition 3.3.1. Assume that in system (3.1) the pair (A, B) is stabilizable.

First though some definitions. We combine both the traditional notions of stability and attractiveness in the following definition.

Definition 3.3.5 *System (3.1) is asymptotically stabilizable if there exists an encoder, decoder, and controller such that*

1 *Stability:* $\forall \epsilon > 0 \exists \delta(\epsilon)$ such that $\|X_0\|_2 \leq \delta(\epsilon)$ implies $\|X_t\|_2 \leq \epsilon \forall t \geq 0$.

2 *Uniform attractivity:* $\forall \epsilon > 0, \delta > 0 \exists T(\epsilon, \delta)$ such that $\|X_0\|_2 \leq \delta$ implies $\|X_t\|_2 \leq \epsilon \forall t \geq T$.

Point one states that the state cannot grow unbounded for bounded X_0 . The second point states that the state decreases to zero uniformly in X_0 . Note also that uniform attractivity is defined for all δ . Thus our definition of asymptotic stability is global.

Proposition 3.3.2 *We are given system (3.1). A necessary condition on the rate for asymptotic stability is $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$.*

Proof: Without loss of generality assume that A is a diagonal matrix with real eigenvalues λ_i and that there are no disturbances. We will provide a lower bound on the rate required

for uniform asymptotically stability in this case. Let $\{X : \|X\|_\infty \leq L\} \subset \Lambda_0$. For a given control sequence U_0, U_1, \dots, U_{t-1} we have

$$X_t = A^t X_0 + \sum_{i=0}^{t-1} A^{t-1-i} B U_i.$$

The condition for uniform stability states that for $\epsilon > 0$ there exists a $T(\epsilon)$ such that $\forall t \geq T(\epsilon)$ we have $\|X_t\|_\infty \leq \epsilon$ for all initial conditions $X_0 \in \Lambda_0$.

For $\epsilon > 0$ define the balls Γ , parameterized by the control sequences U_0, \dots, U_{t-1} , to be

$$\Gamma_{U_0^{t-1}} = \{X_0 : \|X_t\|_\infty \leq \epsilon\}.$$

A lower bound on the rate can be computed by counting how many Γ -boxes it takes to cover Λ_0 . Note that the dimensions of any $\Gamma_{U_0^{t-1}} = \{X_0 : \|X_t\|_\infty \leq \epsilon\}$ box are $\frac{2\epsilon}{|\lambda_1|^t}, \dots, \frac{2\epsilon}{|\lambda_u|^t}$.

Ignoring the stable subspace we get

$$\begin{aligned} R &\geq \frac{1}{t} \log \frac{(2L)^u}{\frac{2\epsilon}{|\lambda_1|^t} \times \dots \times \frac{2\epsilon}{|\lambda_u|^t}} \\ &= \sum_{j=1}^u \log(|\lambda_j|) + \frac{u}{t} \log\left(\frac{L}{\epsilon}\right) \end{aligned}$$

(Where $\lambda_1, \dots, \lambda_u$ are the unstable eigenvalues.) If $\epsilon(t)$ is growing sub-exponentially with t then the second addend above goes to zero. If we want to converge exponentially we require asymptotically an extra $\lim_{t \rightarrow \infty} \frac{u}{t} \log\left(\frac{L}{\epsilon}\right)$ bits. \square

3.3.3 Controllability

In this section we discuss controllability under a rate constraint. Assume that in system (3.1) the pair (A, B) is controllable.

There are differences in the definitions of reachability, controllability, null controllability, etcetera. Furthermore under a rate constraint one can only hope for an approximate controllability type result. The usual definition of controllability to a point P is that for any initial point X_0 there exists a control sequence such that one can drive X_0 to P . The usual definition of reachability from X_0 is that given any point P there exists a control sequence that drives X_0 to P . Neither is satisfactory for us since we cannot assume exact knowledge of X_0 . Thus we make the following definitions:

Definition 3.3.6 *System (3.1) is controllable under a rate constraint if for all $X_0, P \in \mathbb{R}^d$, $\forall \epsilon > 0$ there exists a controller, encoder and decoder and a $t(\epsilon, X_0)$ such that $\|X_t - P\| \leq \epsilon$.*

Definition 3.3.7 *System (3.1) with bounded Λ_0 is uniform in initial state controllable under a rate constraint if for all $P \in \mathbb{R}^d, X_0 \in \Lambda_0$, $\forall \epsilon > 0$ there exists a controller, encoder and decoder and a $t(\epsilon, \Lambda_0)$ such that $\|X_t - P\| \leq \epsilon$.*

Proposition 3.3.3 *For system (3.1) with bounded Λ_0 a necessary condition on the rate for uniform in initial state controllability under a rate constraint is $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$.*

Proof: By proposition 3.3.2, the rate condition is necessary to drive the system to the origin. (Driving the system to the origin is sometimes called null-controllability.) Thus the rate condition is necessary for the more general controllability problem. \square

3.4 Encoder Classes, Primitive Quantizers, and the Key Technical Lemma

In this section we provide background material for discussing schemes that upper bound the rate required to achieve the given control objectives. We first define the different encoder classes of interest to us. We then discuss the role of equi-memory. Next we define a primitive quantizer. We end with a statement of the key technical lemma.

3.4.1 Encoder Classes

Recall that the encoder at time t is a map \mathcal{E}_t that takes $(Y^t, \sigma^{t-1}, U^{t-1}) \mapsto \sigma_t$. In this case the encoder knows the past states, past channel symbols, and past controls. It may seem unreasonable to allow the encoder access to the past controls. For example the encoder may be geographically separated from the plant. It may even seem unreasonable to allow the encoder any memory whatsoever.

We can imagine a continuum where on one end we have an encoder with access to all the past information $(Y^t, \sigma^{t-1}, U^{t-1})$. On the other end we have an encoder with access to only Y_t . Of course there are many cases in between. In fact, from chapter 2, we know there is a partial order on the information patterns of the encoders. Our goal is to treat these two encoders plus another encoder that is in between the two.

We first make explicit the partial ordering on the set of all information patterns. We then describe the different encoders that we will examine in this chapter.

Partial Ordering of the Information Patterns

Recall our discussion in chapter two.

Definition 3.4.1 *The information pattern is defined as $\mathcal{I} = \{\{I_{\mathcal{E}_t}\}_{t=1}^T, \{I_{\mathcal{D}_t}\}_{t=1}^T\}$. Where $I_{\mathcal{E}_t} \subseteq \mathbb{R}^{l(t+1)} \times \Sigma^t \times \mathbb{R}^{mt}$ is the information available to the encoder at time t . Similarly $I_{\mathcal{D}_t} \subseteq \Sigma^{t+1} \times \mathbb{R}^{mt}$ is the information available to the decoder at time t .*

We define a partial ordering on information patterns.

Definition 3.4.2 *We say $\mathcal{I} \preceq \mathcal{J}$ if for all t we have $I_{\mathcal{E}_t} \subseteq J_{\mathcal{E}_t}$ and $I_{\mathcal{D}_t} \subseteq J_{\mathcal{D}_t}$.*

The unique maximal element in this partial ordering is $\mathcal{I}^* = \{\{I_{\mathcal{E}_t}^*\}_{t=1}^T, \{I_{\mathcal{D}_t}^*\}_{t=1}^T\}$. Where $I_{\mathcal{E}_t}^* = \mathbb{R}^{l(t+1)} \times \Sigma^t \times \mathbb{R}^{mt}$ and $I_{\mathcal{D}_t}^* = \Sigma^{t+1} \times \mathbb{R}^{mt}$.

As stated in chapter two when computing an optimal policy more information cannot increase the optimal cost (assume here that “optimal” means a policy that minimizes a given cost.) Thus an optimal policy for the maximal information pattern will lead to the minimal cost.

We now describe four encoder structures with nested, with respect to the partial order, information patterns. We distinguish between encoders that observe, and/or can compute, the control signals and one that does not. This distinction is important. While the encoder with access to the controls has lower rate requirements it is less practical in distributed settings. Among the encoders that do not observe the control we make a further distinction between those that have memory and those that do not.

Encoder class 1

In this class the encoder is a map, \mathcal{E}_t , that takes $(Y^t, \sigma^{t-1}, U^{t-1}) \mapsto \sigma_t$. The decoder at time t is a map, \mathcal{D}_t , that takes $(\sigma^t, U^{t-1}) \mapsto \Omega_t$. We assume that both the encoder and decoder have knowledge of the dynamics \mathcal{F} . Furthermore we assume that the encoder knows \mathcal{D} and the decoder knows \mathcal{E} . We do not assume that the encoder or decoder knows \mathcal{C} . Also we do not assume that the controller knows \mathcal{E} or \mathcal{D} .

Encoder class 1a

Encoder 1a is not allowed to observe the control signals. Thus it is a map, \mathcal{E}_t , that takes $(Y^t, \sigma^{t-1}) \mapsto \sigma_t$. The decoder at time t is a map, \mathcal{D}_t , that takes $(\sigma^t, U^{t-1}) \mapsto \Omega_t$. We assume that the encoder knows \mathcal{D} and the decoder knows \mathcal{E} . Furthermore we assume that the encoder knows \mathcal{C} . Since the channel is noiseless the encoder can simulate the actions of the decoder and controller and thus compute the control signals. Note because the encoder has policy knowledge it can reconstruct signal knowledge not directly available to it.

For the noiseless channel, encoder 1 and encoder 1a both effectively observe the control signal. In chapter six, where we discuss noisy channels, this will no longer be the case. The point of the distinction between encoder 1 and encoder 1a is that, though both encoders can observe/produce the same values $(Y^t, \sigma^{t-1}, U^{t-1})$, the “physical” realizations are different. In the former case we have a physical link connecting the previous U_i 's to the encoder. In the latter case the encoder has knowledge of the map that produces the U_i 's. There is an engineering tradeoff between signaling over a physical link versus incorporating knowledge and computation at a given component.

In both encoder 1 and 1a the optimal action for the encoder is to compute the optimal control signal, quantize it appropriately, and then treat the decoder and controller as table lookups. (In this case the “optimal control signal” is that control signal we would compute under full observation.) By assumption the control law will be a deterministic function of the state. Thus any partition of the control space induces a partition on the state space. In this chapter we choose to quantize the state space.

Encoder class 2 with memory

In this class the encoder is a map, \mathcal{E}_t , that takes $(Y^t, \sigma^{t-1}) \mapsto \sigma_t$. We assume that the encoder does not know the control law \mathcal{C} . (Thus this is not an encoder in encoder class 1a.) The decoder at time t is a map that takes $(\sigma^t, U^{t-1}) \mapsto \Omega_t$. Furthermore we assume that the encoder knows \mathcal{D} and the decoder knows \mathcal{E} . Note that this is an example of decentralized design. The encoder is not allowed to know the control policy.

In this case it is less obvious what the encoder should transmit. Though the encoder does not know the control signal or law we allow it to know the objective. For example if the objective is stability then the encoder knows that under a good control law the state will be converging to the origin. Thus the encoder can choose quantizers centered at the origin with shrinking dynamic range. The issue then is to determine at what rate should the dynamic range shrink. This will be discussed in section 3.6. Generally speaking, even

if the encoder does not observe the control signal, knowledge of the control objective will restrict the set of control signals available.

Encoder class 2 without memory

In this class the encoder is memoryless. It is a map, \mathcal{E}_t , taking $Y_t \mapsto \sigma_t$. This is the simplest encoder that we consider. In this class we assume that the encoder does not know the control law \mathcal{C} but we assume that it knows the control objective. The decoder at time t is a map that takes $(\sigma^t, U^{t-1}) \mapsto \Omega_t$. Finally we assume that the encoder knows \mathcal{D} and the decoder knows \mathcal{E} . Once again this is an example of decentralized design. This encoder is very simple and make the least demands on signal and system knowledge.

We summarize the properties of the different encoder classes in the following table:

<i>Encoder Class</i>	<i>Observations</i>	<i>Encoder's System Knowledge</i>
1	$Y^t, \sigma^{t-1}, U^{t-1}$	$\mathcal{F}, \mathcal{E}, \mathcal{D}$
1a	Y^t, σ^{t-1}	$\mathcal{F}, \mathcal{E}, \mathcal{D}, \mathcal{C}$
2 with memory	Y^t, σ^{t-1}	$\mathcal{F}, \mathcal{E}, \mathcal{D}$
2 without memory	Y_t	$\mathcal{F}, \mathcal{E}, \mathcal{D}$

We assume that all the encoders know the control objective. One could also envision a setup where the encoder does not know the control objective. In this case the encoder will have to learn the objective. We leave this scenario to future research.

For all the different encoder structures above we assume that the decoder, \mathcal{D}_t , is a map that takes $(\sigma^t, U^{t-1}) \mapsto \Omega_t$. We also assume that the decoder knows \mathcal{E} and \mathcal{F} . The encoder and decoder need to work together. Informally, the job of the decoder upon receiving σ_t is to invert the encoder operation. To do this it needs to know what the encoder operation is. Knowledge of \mathcal{E} is not enough to insure this. Thus we introduce the notion of equi-memory.

With the assumption of equi-memory and the definition of information state we can give a specification of how Ω_t is computed.

3.4.2 Equi-memory

We can define a state for the encoder and decoder at each time step. Specifically for encoder 1 let

$$I_t^1 = (Y^{t-1}, \sigma^{t-1}, U^{t-1}) \in \mathbb{R}^{lt} \times \Sigma^t \times \mathbb{R}^{mt} \quad \forall t \geq 0.$$

For encoder 2 with memory let

$$I_t^{2,\text{mem}} = (Y^{t-1}, \sigma^{t-1}) \in \mathbb{R}^{lt} \times \Sigma^t \quad \forall t \geq 0.$$

For encoder 2 without memory let

$$I_t^{2,\text{no mem}} = \emptyset \quad \forall t \geq 0.$$

For all the decoders let

$$J_t = (\sigma^{t-1}, U^{t-1}) \in \Sigma^t \times \mathbb{R}^{mt} \quad \forall t \geq 0.$$

Replace Y_t by X_t when the encoder can observe the state.

Use I_t^* to denote any of $I_t^1, I_t^{2,\text{mem}},$ or $I_t^{2,\text{no mem}}$. Similarly let \mathcal{E}_t^* represent any of the encoders. For a fixed I_t^* define $\bar{\mathcal{E}}_{I_t^*,t}^*(Y_t) \triangleq \mathcal{E}_t^*(Y_t, I_t^*)$ (where the latter is an abuse of notation.) Define *valid* (I_t^*, J_t) to mean any pair (I_t^*, J_t) capable of being produced by the system (i.e. the encoder, decoder, controller, and plant) at time t . Finally let $\bar{\mathcal{E}}_{I_t^*,t}^{-1,*}(\sigma) \triangleq \{Y_t : \mathcal{E}_t^*(Y_t, I_t^*) = \sigma\}$.

Definition 3.4.3 *An encoder/decoder pair are said to be equi-memory if for all valid (I_t^*, J_t) and $\sigma_t \in \Sigma$ the information (J_t, σ_t) is sufficient to determine the set $\bar{\mathcal{E}}_{I_t^*,t}^{-1,*}(\sigma_t)$. Specifically there exists a map*

$$\mathcal{D}_t^{\text{em}} : \Sigma^{t+1} \times \mathbb{R}^{mt} \rightarrow 2^{\mathbb{R}^l}.$$

taking

$$(\sigma^t, U^{t-1}) \mapsto \bar{\mathcal{E}}_{I_t^*,t}^{-1,*}(\sigma_t).$$

The superscript “em” represents “equi-memory.”

Implicit in this definition is that the information in J_t is sufficient for the decoder to invert the encoder map.

Assumption 3.4.1 *Throughout this chapter we will assume that the encoder and decoder are equi-memory.*

Definition 3.4.4 *An information state Ω_t is any set that contains X_t at time t .*

Note that Ω_t can be unbounded. This definition is different than the traditional definition of information state for deterministic systems. In the traditional definition Ω_t contains all X_t consistent with the observations of the decoder. Our weaker definition is sufficient for our purposes here.

We need to show that a nontrivial Ω_t can be computed (i.e. an $\Omega_t \neq \mathbb{R}^d$). By equi-memory the decoder can invert the encoder map. Thus at time t upon receipt of σ_t the decoder can determine that $Y_t \in \bar{\mathcal{E}}_{I_t^*,t}^{-1,*}(\sigma_t)$. Furthermore the decoder knows the observation matrix C . Thus it knows that $X_t \in \{X : CX \in \bar{\mathcal{E}}_{I_t^*,t}^{-1,*}(\sigma_t)\}$. Of course one can compute a smaller set Ω_t by incorporating the other information the decoder has. We will discuss this when describing specific schemes.

3.4.3 Primitive Quantizer

The space of encoders defined so far is still quite large. We will further restrict the encoders to use a primitive quantizer at each time step.




Figure 3-2: Primitive Quantizer

Definition 3.4.5 A primitive quantizer is a four-tuple $(C, \underline{R}, \underline{L}, \Phi)$ with $C \in \mathbb{R}^d$ representing the centroid, $\underline{R} = (R_1, \dots, R_d) \in \mathbb{R}^{d,+}$ representing the rate vector, $\underline{L} = (L_1, \dots, L_d) \in \mathbb{R}^{d,+}$ representing the side-lengths of the dynamic range, and Φ an invertible matrix representing a coordinate transformation. This quantizer partitions the region

$$\Lambda = \left\{ X \in \mathbb{R}^d : \Phi(X - C) \in \left[-\frac{L_1}{2}, \frac{L_1}{2} \right] \times \dots \times \left[-\frac{L_d}{2}, \frac{L_d}{2} \right] \right\}$$

into boxes with side lengths $\frac{L_i}{2^{R_i}}$. Let $R = \sum_{i=1}^d R_i$ be the total rate. Each of the 2^R boxes is represented by an element $\sigma \in \Sigma$. Upon observing X the $(C, \underline{R}, \underline{L}, \Phi)$ -quantizer subtracts off C , applies the coordinate transform Φ , determines which box it falls into, and then transmits the σ representing that box. If X falls outside the region Λ then the quantizer transmits a special symbol representing an overflow. Thus we have $2^R + 1$ symbols. The set Λ is called the dynamic range of the quantizer.

Figure 3-2 shows a two-dimensional primitive quantizer with $R_1 = 3$ and $R_2 = 2$.

The encoder based on its information I_t^* selects a $(C, \underline{R}, \underline{L}, \Phi)$ -quantizer. Upon observing Y_t it computes the appropriate σ_t and transmits it across the channel. Note that the decoder needs to know which quantizer was selected so that it may decode the received symbol σ_t appropriately. This is assured by equi-memory. The equi-memory condition forces the encoder and decoder to make decisions based on the same information.

Assumption 3.4.2 All the encoders in this chapter are restricted to using primitive quantizers.

One may ask why we have chosen boxes instead of more general polytopes to partition Λ . Clearly if one uses general polytopes one should achieve a lower rate than the rate one gets when restricting oneself to boxes. However the analysis for the boxes case is much easier. We postpone discussion of more general partitions to chapter five where we introduce the sequential rate distortion problem. In certain cases though we will show that schemes using boxes are sufficient to achieve the information theoretic lower bounds provided in section 3.3. A further reason for using boxes is their simplicity in practice.

3.4.4 Key Technical Lemma

Here we provide some notation and results that will be used throughout this chapter. In general the growth of the uncertainty in the state estimate can be characterized by the eigenvalues of a certain matrix. In order to understand this characterization we put forth the following definitions and lemmas.

We are interested in systems of the form $X_{t+1} = AX_t$ (ignore the control and disturbance terms for now.) Assume $X_t \in \left\{ \left[-\frac{L}{2}, \frac{L}{2}\right] \times \dots \times \left[-\frac{L}{2}, \frac{L}{2}\right] \right\}$. We would like a way to calculate the box that X_{t+1} lives in. This subsection provides a way to upper bound that box.

Let $A \in \mathbb{R}^{d \times d}$. Then A has a real Jordan canonical form.

Theorem 3.4.1 *For any real A there exists a real valued nonsingular matrix Φ and a real valued matrix Υ such that $\Phi A \Phi^{-1} = \Upsilon = \text{diag}[J_1, \dots, J_m]$. Where each J_j , $j = 1, \dots, m$, is a Jordan block of dimension (geometric multiplicity) d_j . Clearly $d_1 + \dots + d_m = d$. The Jordan block associated with a real eigenvalue λ takes the form*

$$\begin{bmatrix} \lambda & 1 & & \\ & \lambda & 1 & \\ & & \dots & \\ & & & \lambda \end{bmatrix}.$$

The Jordan block associated with the complex conjugate pair of eigenvalues $\lambda = \rho(\cos\theta \pm i \sin\theta)$ takes the form

$$\begin{bmatrix} D & I & & \\ & D & I & \\ & & \dots & \\ & & & D \end{bmatrix}$$

where $D = \rho r(\theta)$. Where $r(\theta) = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}$.

Proof: See theorem 2 of section 6.4 in [HS]. (Note $\rho \geq 0$.) \square

We now define a matrix H that serves to undo the rotation caused by each of the complex conjugate eigenvalue pairs.

Definition 3.4.6 Define $H = \text{diag}[H_1, \dots, H_m]$. Where each H_j is associated with one of the Jordan blocks J_j . Specifically $H_j = I$ if J_j is the Jordan block with real eigenvalue λ_j . And $H_j = \text{diag}[r(\theta)^{-1}, \dots, r(\theta)^{-1}]$ if J_j is the Jordan block associated with the complex conjugate eigenvalues $\rho(\cos \theta \pm i \sin \theta)$.

Note that if A has all real eigenvalues then $H = I$.

The following lemma shows that Υ and any power of the matrix H commute.

Lemma 3.4.1 $H^t \Upsilon H^{-t} = \Upsilon$

Proof: See section A.2 in the appendix. \square

We want a way to bound the growth of the operator $H\Upsilon$. Note that this is a block diagonal matrix: $H\Upsilon = \text{diag}[K_1, \dots, K_d]$. With $K_j = J_j$ if J_j is the Jordan block associated with a real eigenvalue. Otherwise

$$\begin{aligned} K_j &= H_j J_j \\ &= \begin{bmatrix} \rho I & r(\theta)^{-1} & & & \\ & \rho I & r(\theta)^{-1} & & \\ & & \dots & & \\ & & & \dots & \\ & & & & \rho I \end{bmatrix} \end{aligned}$$

if J_j is the Jordan block associated with a complex conjugate eigenvalue pair.. Note that the eigenvalues of the upper triangular matrix K_j are all equal to ρ . (See page 39 of [HJo].)

We will bound the growth of $H\Upsilon$ by introducing a new matrix $\overline{\Upsilon}$ that bounds $H\Upsilon$. Then we will bound $\overline{\Upsilon}$. To define $\overline{\Upsilon}$ we first need to define the following \overline{K} versions of the K matrices. For each Jordan block J_j associated with a real eigenvalue λ_j define

$$\overline{K}_j = \begin{bmatrix} |\lambda| & 1 & & & \\ & |\lambda| & 1 & & \\ & & \dots & & \\ & & & \dots & \\ & & & & |\lambda| \end{bmatrix}.$$

For each Jordan block J_j associated with a complex eigenvalue $\rho(\cos \theta + i \sin \theta)$ define

$$\overline{K}_j = \begin{bmatrix} \rho I & O & & & \\ & \rho I & O & & \\ & & \dots & & \\ & & & \dots & \\ & & & & \rho I \end{bmatrix}$$

where $O = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$. Finally define $\overline{\Upsilon} = \text{diag}[\overline{K}_1, \dots, \overline{K}_m]$. Note that the ij th entry of the matrix $\overline{\Upsilon}$ is nonnegative and greater than or equal to the absolute value of the ij th entry of $H\Upsilon$.

For a given rate vector \underline{R} define

$$F_{\underline{R}} = \begin{bmatrix} \frac{1}{2^{R_1}} & & & \\ & \frac{1}{2^{R_2}} & & \\ & & \dots & \\ & & & \frac{1}{2^{R_d}} \end{bmatrix}.$$

We are now finally in a position to prove the key technical lemma.

Lemma 3.4.2 *If for all i we have $R_i > \max\{0, \log |\lambda_i|\}$ then $\overline{\Upsilon}F_{\underline{R}}$ is stable. If there exists at least one i such that*

$$R_i < \max\{0, \log |\lambda_i|\}$$

then $\overline{\Upsilon}F_{\underline{R}}$ is unstable.

Proof: See section A.2 in the appendix. \square

We conclude this section with an important application of lemma 3.4.2. Let $X_{t+1} = AX_t$ and let $Z_t = H^t\Phi X_t$. Note that if A has real eigenvalues then $Z_t = \Phi X_t$. The H is needed to undue the rotation caused by the dynamics of the complex conjugate eigenvalue pairs:

$$\begin{aligned} Z_{t+1} &= H^{t+1}\Phi X_{t+1} \\ &= H^{t+1}\Phi AX_t \\ &= H^{t+1}\Phi A\Phi^{-1}H^{-t}Z_t \\ &= H^{t+1}\Upsilon H^{-t}Z_t \\ &= H\Upsilon Z_t \quad \text{by lemma 3.4.1} \end{aligned}$$

Lemma 3.4.3 *If Z_t is in the box determined by $\underline{L}(t)$ (i.e. $Z_t \in \{[-\frac{L_1}{2}, \frac{L_1}{2}] \times \dots \times [-\frac{L_d}{2}, \frac{L_d}{2}]\}$) then Z_{t+1} is in the box determined by $\overline{\Upsilon}\underline{L}(t)$.*

Proof: We know $Z_{t+1} = H\Upsilon Z_t$. By construction $\overline{\Upsilon}$ is a matrix whose ij th entry is nonnegative and greater than or equal to the absolute value of the ij th entry of $H\Upsilon$. Thus $\overline{\Upsilon}$ bounds the growth of each component of Z_t . \square

3.4.5 Comments

It is possible to analyze scenarios where we do not impose equi-memory. The decoder, though, may not know which quantizer the encoder is using. A game-theoretic formulation would be appropriate here. We leave this possibility for future research. In general though problems with differing information spaces, often called team problems, are notoriously difficult to analyze let alone optimize. However it is often the case that when the information spaces are “nested” the problem becomes tractable. Equi-memory is one way to enforce the nested information property. See Radner [Rad]. We take the view that part of the role of information transmission is to maintain a common global state. Of course in a truly decentralized theory one would have to understand how to deal with a lack of global state.

Our primitive quantizers have a dynamic range and “saturate” if the state falls outside this range. There is a large body of literature on control for saturated systems. Our problem is different in that the saturation bounds can be time-varying and the signals must be quantized.

In general we will require a time-varying encoder. Take for example a stability problem where one wants to drive the system to the origin. A time-invariant finite rate encoder will have a quantization region around the origin of a fixed size. The best we can do is drive the state to this region. Once inside this region we have no assurances that we are converging to the origin. We suggest three ways to deal with this. We can allow for an infinite number of regions. [EM] We can allow for ϵ -accuracy (e.g. notions of practical stability.) [EM]. Or we can allow for time-varying encoders. Roughly speaking the last case leads to the smallest rate requirements. And since that is what we are most interested in we will spend most of our time looking at this case.

It should be clear that under a rate constraint if the state estimation error increases with time in an unbounded fashion there will come a point when we can no longer satisfy the control objective. In the case where the error is unbounded we essentially have no information about the state (i.e. this is like having no useful feedback.) Thus unless the control objective can be achieved via an open loop controller we cannot hope to achieve the control objective. A guiding principle throughout this chapter is that the state estimation error should grow at a slower rate than the dynamics. Said another way we are interested in characterizing the largest tolerable level of state estimation error that still insures the control objective is satisfied. This will be discussed more extensively in section 3.6.4.

3.5 Encoder Class One

In this section we provide results on observability, stability, and controllability for encoder class one. We will show that we can achieve the lower bounds proved in section 3.3 by explicitly describing the encoder, decoder, and controller.

Since we are dealing with asymptotic versions observability and controllability it is of interest to determine the rate of convergence. Clearly this rate of convergence will depend on the channel rate used. The rate of convergence will depend on the difference between the channel rate used and the lower bound. The larger this difference the faster the convergence. We can make an analogy with the channel coding theorem in information theory. If the rate R is greater than the capacity C then one cannot transmit information reliably. If $R < C$ then one can. Furthermore the error exponent roughly depends on the difference $C - R$.

We now treat, in order, observability, stability, and controllability for encoder class one.

3.5.1 Observability

In this subsection we provide schemes that achieve asymptotic observability. In the following proposition we give a sufficient condition for uniform in control observability of system (3.1) under encoder class 1 when we observe the state (i.e. $Y_t = X_t$.)

Proposition 3.5.1 *Given system (3.1), encoder in encoder class 1, and $C = I$. Furthermore assume that the encoder knows a bound on Λ_0 . A sufficient condition on the rate for uniform in control asymptotic observability is $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$.*

Proof: Assume $\Lambda_0 \subset \{X : \|X\|_2 \leq L\}$. Let Φ diagonalize A into real Jordan canonical form: $\Phi A \Phi^{-1} = \Upsilon$. For $X \in \Lambda_0$ we have $\|\Phi X\|_2 \leq \|\Phi\| \|X\|_2 \leq \|\Phi\| L$. At time zero choose a $(C(0), \underline{R}, \underline{L}(0), \Phi(0))$ -quantizer where $C(0)$ is the origin and $\Phi(0) = \Phi$. Let $L_i(0) = \|\Phi\| L \forall i$ and choose any $R_i > \max\{0, \log |\lambda_i|\}$. Apply this quantizer to X_0 and transmit σ_0 . Note that σ_0 will not be the overflow symbol.

At time t let the state estimate, \hat{X}_t , be the centroid of the region defined by σ_t . Equivalently \hat{X}_t is the centroid of the decoder output set Ω_t . This equivalence holds because of our assumption of equi-memory. We update the quantizer parameters as follows. First the centroid of the $t + 1$ th quantizer is just the one step ahead state prediction (the encoder observes the controls):

$$C(t+1) = A\hat{X}_t + BU_t.$$

Second the coordinate transformation evolves as

$$\Phi(t+1) = H^{t+1}\Phi = H\Phi(t).$$

Third the size of the dynamic range of the $t + 1$ th quantizer will evolve according to:

$$\underline{L}(t+1) = \overline{\Upsilon} F_{\underline{R}} \underline{L}_t.$$

By lemma 3.4.2 $\overline{\Upsilon} F_{\underline{R}}$ is a stable matrix. Thus the dimensions of the dynamic range are decreasing in time.

The decoder upon receiving σ_t outputs Ω_t which is the set represented by σ_t . This completes our description of the encoder and decoder. By construction the state X_t never leaves the dynamic range of the t th quantizer Λ_t . See lemma 3.4.2. (One can think of Ω_t as containing the instantaneous estimation error of X_t and Λ_t as containing the one-step ahead estimation error of X_t .)

$$\Lambda_t = \left\{ X \in \mathbb{R}^d : \Phi(t)(X - C(t)) \in \{[-L_1(t), L_1(t)] \times \dots \times [-L_d(t), L_d(t)]\} \right\}.$$

Thus

$$\Omega_t = \left\{ X \in \mathbb{R}^d : \Phi(t)(X - \hat{X}_t) \in \left\{ \left[-\frac{L_1(t)}{2^{R_1}}, \frac{L_1(t)}{2^{R_1}} \right] \times \dots \times \left[-\frac{L_d(t)}{2^{R_d}}, \frac{L_d(t)}{2^{R_d}} \right] \right\} \right\}.$$

Now

$$\begin{aligned} \|e_t\|_2 &\leq \sup_{X \in \Omega_t} \|X - \hat{X}_t\|_2 \\ &= \sup_{X \in \Omega_t} \|\Phi(t)^{-1} \Phi(t)(X - \hat{X}_t)\|_2 \\ &\leq \sup_{X \in \Omega_t} \|\Phi(t)^{-1}\| \|\Phi(t)(X - \hat{X}_t)\|_2 \\ &\leq \|\Phi(t)^{-1}\| \|F_{\underline{R}} \underline{L}(t)\|_2 \\ &\leq \|\Phi(t)^{-1}\| \|F_{\underline{R}}\| \|(\overline{\Upsilon} F_{\underline{R}})^t\| \|\underline{L}(0)\|_2 \\ &\leq \sqrt{d} L \|\Phi\| \|\Phi(t)^{-1}\| \|F_{\underline{R}}\| \|(\overline{\Upsilon} F_{\underline{R}})^t\| \end{aligned}$$

The decay of $\|(\overline{\Upsilon} F_{\underline{R}})^t\|$ is determined by the largest eigenvalue of $\overline{\Upsilon} F_{\underline{R}}$. Specifically there exists a constant κ such that $\|(\overline{\Upsilon} F_{\underline{R}})^t\| \leq \kappa 2^{-t(\min_i (R_i - \log |\lambda_i(A)|))}$. \square

Note that one could optimize the constant in the upper bound by using more general polytopes than boxes. In systems with large state spaces the constant in front can be large.

In proposition 3.5.1 we used \underline{L}_t as a measure of the uncertainty. Furthermore we showed that \underline{L}_t decreased to zero. One may ask how this may be related to Lyapunov functions. We will explore this further in section 3.6.4.

Now we treat the case where the encoder does not know a bound on Λ_0 .

Corollary 3.5.1 *We are given system (3.1), an encoder in encoder class 1, and $C = I$. A sufficient condition on the rate for uniform in control asymptotic observability is $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$.*

Proof: For the case where the initial uncertainty, Λ_0 , is unknown one must first “capture” the state in the quantizer domain. Specifically let $(C(0), \underline{R}, \underline{L}(0), \Phi)$ -quantizer where $C(0)$ is the origin, $\Phi(t) = H^t \Phi$, and $L_i(0) = L \forall i$ for some arbitrary L . If upon observing X_n the quantizer at time transmits an overflow symbol then update the quantizer as follows: $C(t+1) = AC(t) + BU_t$ and $L_i(t+1) = L_i(t)2^{R_i}$. Since the L 's are growing faster than

the state eventually the quantizer will capture the state. At this point proceed as we did in proposition 3.5.1. \square

The idea of growing the quantizer range when the encoder does not know an *a priori* bound on the initial state works on all of the following results in this section. Thus we will only prove the semi-global statement (i.e. bounded Λ_0) with the understanding that the global result also holds.

We now consider the case of bounded additive disturbances.

$$X_{t+1} = AX_t + BU_t + W_t, \quad Y_t = X_t \quad t \geq 0 \quad (3.2)$$

where $\|W_t\|_2 \leq D$.

Proposition 3.5.2 *We are given system (3.2) with encoder in encoder class 1. If $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$ then there exists a scheme such that $\lim_{t \rightarrow \infty} \|e_t\|_2$ is bounded.*

Proof: Follow the same setup as in proposition 3.5.1 except update the L_i 's as follows:

$$\underline{L}(t+1) = \bar{\Upsilon} F_{\underline{R}} \underline{L}(t) + D \|\Phi(t+1)\| \begin{bmatrix} 1 \\ 1 \\ \dots \\ 1 \end{bmatrix}$$

Denote the second addend as $h(t)$. Essentially we are expanding the size of the dynamic range of the quantizer to take into account the disturbances. Since the disturbances are bounded we know that the state cannot leave the sets defined by the $\underline{L}(t)$'s.

Now we can write

$$\underline{L}(t) = (\bar{\Upsilon} F_{\underline{R}})^t \underline{L}(0) + \sum_{j=0}^{t-1} (\bar{\Upsilon} F_{\underline{R}})^{t-1-j} h(j)$$

Now

$$\begin{aligned} \|e_t\|_2 &\leq \|\Phi(t)^{-1}\| \|F_{\underline{R}} \underline{L}(t)\| \\ &\leq \|\Phi(t)^{-1}\| \|F_{\underline{R}}\| \left\{ \|(\bar{\Upsilon} F_{\underline{R}})^t \underline{L}(0)\| + \left\| \sum_{j=0}^{t-1} (\bar{\Upsilon} F_{\underline{R}})^{t-1-j} h(j) \right\| \right\} \\ &\leq \|\Phi(t)^{-1}\| \|F_{\underline{R}}\| \left\{ \|(\bar{\Upsilon} F_{\underline{R}})^t \underline{L}(0)\| + \sum_{j=0}^{t-1} D \|(\bar{\Upsilon} F_{\underline{R}})^{t-1-j}\| \|H^{j+1}\| \|\Phi\| \right\} \\ &\leq \|\Phi(t)^{-1}\| \|F_{\underline{R}}\| \left\{ \|(\bar{\Upsilon} F_{\underline{R}})^t \underline{L}(0)\| + D \|\Phi\| \sum_{j=0}^{t-1} \|(\bar{\Upsilon} F_{\underline{R}})^{t-1-j}\| \right\} \end{aligned}$$

Now there exists a constant κ such that $\|(\bar{\Upsilon}F_{\underline{R}})^t\| \leq \kappa 2^{-t(\min_i(R_i - \log |\lambda_i(A)|))}$. Thus

$$\lim_{t \rightarrow \infty} \sum_{j=0}^{t-1} \|(\bar{\Upsilon}F_{\underline{R}})^{t-1-j}\| \leq \frac{\kappa}{1 - 2^{-\min_i(R_i - \log |\lambda_i(A)|)}}.$$

Therefore $\lim_{t \rightarrow \infty} \|e_t\| \leq \frac{\kappa}{1 - 2^{-\min_i(R_i - \log |\lambda_i(A)|)}} \|\Phi(t)^{-1}\| \|\Phi\| \|F_{\underline{R}}\|$. Note that $\|F_{\underline{R}}\| = \max_i \frac{1}{2^{R_i}}$. Thus, as we expect, the bound goes to zero if we let the rate go to infinity. \square

Example 3.5.1 Take the scalar case: $X_{t+1} = aX_t + bU_t + W_t$. In this case the upper bound is $\lim_{t \rightarrow \infty} |e_t| \leq \frac{D}{2^{R-|a|}}$.

A tight lower bound on the rate is difficult to find. We do know, though, that for any finite rate the state estimation error cannot be driven to zero. There will always be a nonzero state estimation error in transmitting information regarding the new noise term W_t at time t . In chapter five, where we deal with stochastic disturbances, we will show that the sequential rate distortion function is a lower bound on the rate as a function of the estimation error (in a suitable time-average sense to be defined there.) Note that we could use a technique similar to the one used in proposition 3.3.1 to lower bound the rate. It turns out that the rate computed there is conservative. The partition induced by the covering cannot be achieved sequentially. This has to do with the fact that when there are no disturbances every control sequence and initial position X_0 determines a unique path X_1, \dots, X_T . However with disturbances there may be many paths. This is not accounted for in the counting argument. This will be discussed more extensively in chapter five where we discuss successive refinement.

We now discuss what happens when the magnitude of the disturbances shrink with time. First we require a technical lemma.

Lemma 3.5.1 Let A be a stable matrix. Let B_t be a set of matrices such that $\|B_t\| \leq L$ and the limit $\lim_{t \rightarrow \infty} B_t \rightarrow 0$. Let $S_t = \sum_{i=0}^{t-1} A^{t-1-i} B_i$ then $\lim_{t \rightarrow \infty} S_t = 0$.

Proof: See section A.2 in the appendix. \square

In the following proposition we allow the error to decay as $D_t \leq \alpha^t D$ where $0 \leq \alpha < 1$.

Proposition 3.5.3 Given system (3.2) with encoder in encoder class 1. The noise satisfies $\|W_t\|_2 \leq D_t$ where $D_t \leq \alpha^t D$ and $0 \leq \alpha < 1$. Then $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$ is sufficient for uniform in control asymptotic observability.

Proof: Follow the same setup as in proposition 3.5.2 except update the L_i 's as follows:

$$\underline{L}(t+1) = \bar{\Upsilon}F_{\underline{R}}\underline{L}(t) + \alpha^t D \|\Phi(t+1)\| \begin{bmatrix} 1 \\ 1 \\ \dots \\ 1 \end{bmatrix}$$

Denote the second addend as $\alpha^t h(t)$. Then

$$\underline{L}(t) = (\overline{\Upsilon} F_R)^t \underline{L}(0) + \sum_{j=0}^{t-1} (\overline{\Upsilon} F_R)^{t-1-j} \alpha^j h(j)$$

Since $\overline{\Upsilon} F_R$ is stable the first term goes to zero. Note that $\lim_{t \rightarrow \infty} \alpha^t h(t) = 0$. By lemma 3.5.1 the second term goes to zero. This implies $\lim_{t \rightarrow \infty} \|e_t\| = 0$. \square .

Now we consider the case with general observation equation $Y_t = C X_t$. Assume that the pair (A, C) is detectable.

Proposition 3.5.4 *We are given system (3.1) with encoder in encoder class one. If $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$ then there exists a scheme that is uniform in control asymptotically observable.*

Proof: At time t the encoder has received Y^t, U^{t-1} and needs to make a state estimate \bar{X}_t . We will run the Luenberger observer. (Which at time t only uses Y^{t-1}, U^{t-1} . This is sufficient for our purposes.) The observer at time t is

$$\bar{X}_t = A \bar{X}_{t-1} + B U_{t-1} + L(Y_{t-1} - C \bar{X}_{t-1}).$$

Where $A - LC$ is stable. Note that the error $\bar{e}_t = X_t - \bar{X}_t = (A - LC) \bar{e}_{t-1}$. Thus $\bar{e}_t = (A - LC)^t \bar{e}_0$ and $\|\bar{e}_t\| \leq \|(A - LC)^t\| \|\bar{e}_0\| \leq c \lambda^t \|\bar{e}_0\|$ for some constant c and $0 \leq \lambda < 1$.

Now $\bar{X}_t = A \bar{X}_{t-1} + B U_{t-1} + LC \bar{e}_{t-1}$. Use proposition 3.5.3 to show that we can asymptotically observe \bar{X}_t . Furthermore \bar{X}_t converges to the true state X_t . Thus we can asymptotically observe X_t . \square

The general prescription for observability in encoder class one is to transmit a finer and finer description of the zero control input response state trajectory. If there are no disturbances this is equivalent to successively refining the initial position. Note that if we allow the encoder “infinite” memory then it need only transmit a finer and finer description of X_0 . Furthermore such an encoder is completely independent of the controls being applied. If we allow the encoder such infinite memory and there are no disturbances then both encoder one and encoder two with memory are the same. This is similar to the statement that open loop and closed loop control are equivalent if the initial position is known exactly.

Assuming that the encoder can have in memory a perfect description of X_0 for all time n is unrealistic. Furthermore it is not robust to disturbances. For this reason we have proposed a recursive structure for the encoders in encoder class one. Specifically at time t we choose a primitive quantizer that quantizes \bar{X}_t . The quantizer essentially computes a state estimation error based on the difference between the current state and the one step ahead prediction. It is not evolving that estimation error from X_0 and the subsequent channel transmissions. We have shown that we can achieve the rate lower bound with encoders that use primitive quantizers. Thus there is no loss in generality restricting to these kinds of encoders. (Note this generality is true up to a constant in the error bound.)

3.5.2 Stability

For encoder class one we can combine the properties of asymptotic observability and full state feedback stability to get output feedback stability. Assume that the pair (A, B) are stabilizable.

Proposition 3.5.5 *Given system (3.1) with encoder in encoder class 1 and $C = I$. The encoder knows a bound on Λ_0 . The rate $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$ is a sufficient condition for asymptotic stability.*

Proof: Let K be a stabilizing controller, i.e. $A + BK$ is stable. Apply the certainty equivalent controller $U_t = K \hat{X}_t$ where \hat{X}_t is the centroid of the decoder output Ω_t . We need to show that under this controller the system is stable. Let $e_t = X_t - \hat{X}_t$. Then

$$X_t = (A + BK)^t X_0 - \sum_{j=0}^{t-1} (A + BK)^{t-1-j} B K e_j$$

By proposition 3.5.1 the system is asymptotically observable under any control sequence. Furthermore, using the notation in proposition 3.5.1, we have $\underline{L}(t+1) = \bar{\Upsilon} F_{\underline{R}} \underline{L}(t)$ with $\bar{\Upsilon} F_{\underline{R}}$ stable. This implies $\|e_t\| \leq \|\Phi(t)^{-1}\| \|F_{\underline{R}}\| \|(\bar{\Upsilon} F_{\underline{R}})^t\| \|\underline{L}(0)\|$.

Since $A + BK$ is stable the first addend in the above equation goes to zero. By lemma 3.5.1 so does the second. Hence $\lim_{t \rightarrow \infty} X_t = 0$. \square

We note that this result is related to a general result of Vidyasagar that states if a system is state feedback stabilizable and output detectable then it is output feedback stabilizable. [Vid] Furthermore the certainty equivalent controller applied to the state estimate is a stabilizing controller. The difference here is that the observation equation may depend on the past states and controls.

We can treat the case when the encoder does not have an a priori bound on Λ_0 .

Corollary 3.5.2 *Given system (3.1) with encoder in encoder class 1. The rate $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$ is a sufficient condition for asymptotic stability.*

Proof: Apply the zero control until the encoder, using the technique in corollary 3.5.1, “captures” the state. Then proceed as in proposition 3.5.5. \square

The idea of growing the quantizer range when the encoder does not know an a priori bound on the initial state works on all of the following results in this section. Thus we will only prove the semi-global statement (i.e. bounded Λ_0) with the understanding that the global result also holds.

Example 3.5.1 continued *Take the scalar system $X_{t+1} = aX_t + bU_t$, $a > 1$, $|X_0| \leq L$. Choose controller k such that $|a + bk| < 1$. Then under full state feedback the magnitude of the state is strictly decreasing to the origin $|X_t| = |a + bk|^t |X_0|$. Under a rate $R > \log a$ and the scheme proposed in the last proposition we see that*

$$|X_t| \leq |a + bk|^t |X_0| + \sum_{j=0}^{t-1} |a + bk|^{t-1-j} |bk| \frac{1}{2^{Rj}} L.$$

There can exist trajectories that initially are not strictly decreasing to the origin. One can consider this the price of learning the state under a rate constraint. We describe the case of trajectories that strictly decrease in section 3.9.3. Also see [EM].

Now we treat the case where there is an observation equation (i.e. $Y_t = CX_t$.) Assume (A, C) are detectable.

Proposition 3.5.6 *Given system (3.1) and encoder class 1. A sufficient condition on the rate for asymptotic stability is $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$.*

Proof: By proposition 3.5.4 we can asymptotically observe the state for any control sequence. At the encoder implement a Luenberger observer to obtain the dynamics

$$\bar{X}_{t+1} = A\bar{X}_t + BU_t + LC\bar{e}_t.$$

Using proposition 3.5.3 we can transmit \bar{X}_t with error, $e_t = \bar{X}_t - \hat{X}_t$. Using the notation there we know this error is bounded by \underline{L}_n and $\underline{L}(n+1) = \bar{\Upsilon}F_{\underline{R}}\underline{L}(n)$.

Let K be a stabilizing controller, i.e. $A + BK$ is stable. Apply the certainty equivalent controller $U_t = K\hat{X}_t$. We need to show that under this controller the system is stable. Let $\tilde{e}_t = X_t - \hat{X}_t = X_t - \bar{X}_t + \bar{X}_t - \hat{X}_t = \bar{e}_t + e_t$. Now

$$\begin{aligned} \|\tilde{e}_t\| &\leq \|\bar{e}_t\| + \|e_t\| \\ &\leq \|(A + LC)^t\|\|\bar{e}_0\| + \|\Phi(t)^{-1}\| \|F_{\underline{R}}\|\|\bar{\Upsilon}F_{\underline{R}}\|^t \|\underline{L}(0)\| \end{aligned}$$

Because $A + LC$ and $\bar{\Upsilon}F_{\underline{R}}$ are stable there exists a c and a $0 \leq \beta < 1$ such that $\|\tilde{e}_t\| \leq c\beta^t$. Now

$$X_t = (A + BK)^t X_0 - \sum_{j=0}^{t-1} (A + BK)^{t-1-j} BK\tilde{e}_j$$

Since $A + BK$ is stable the first addend in the above equation goes to zero. By lemma 3.5.1 so does the second. Hence $\lim_{t \rightarrow \infty} X_t = 0$. \square

3.5.3 Controllability

In this section we discuss controllability under a rate constraint. Assume that in system (3.1) the pair (A, B) is controllable.

By controllability we know that for every point P there exists a sequence of controls $\bar{U}_0, \dots, \bar{U}_{d-1}$ such that:

$$P = \sum_{j=0}^{d-1} A^{d-1-j} B\bar{U}_j.$$

The basic idea behind our scheme is to first drive the system to the origin. Once suitably close to the origin we will apply the appropriate controls to get to the point P .

Proposition 3.5.7 *For system (3.1), encoder in encoder class 1, $C = I$, and bounded Λ_0 the rate $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$ is a sufficient condition on the rate for uniform in initial state controllability under a rate constraint.*

Proof: By proposition 3.5.5, given $\epsilon > 0$ there exists a time $T(\epsilon)$ such that the state can be driven uniformly in initial state to within $\|X_T\| \leq \frac{\epsilon}{\|A^d\|}$. Once there apply the controls $\bar{U}_0, \dots, \bar{U}_{d-1}$. Then $X_{T+d} = A^d X_T + P$. Then $\|X_{T+d} - P\| = \|A^d X_T\| \leq \epsilon$. \square

Corollary 3.5.3 *For system (3.1), $C = I$, and encoder class 1 a sufficient condition on the rate for controllability under a rate constraint is $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$.*

Proof: As before apply the zero control until the encoder has captured the state and then apply proposition 3.5.7. \square

3.6 Encoder Class Two

In this section we examine observability, stability, and controllability for encoders in encoder class two. Recall that these encoders do not have access to the past controls. There are two types of encoders in encoder class two: those with memory and those without. Recall these encoders are also restricted to be equi-memory and to use primitive quantizers.

We also show that traditional Lyapunov synthesis methods can be used to design the encoder.

3.6.1 Observability

For encoders in encoder class one we were able to show that there exist encoders such that observability holds independently of the control signals chosen. Here we will show that observability for encoders in encoder class two will depend on the control signals chosen.

Proposition 3.6.1 *Given system (3.1) and encoder in encoder class two with memory. If we allow for arbitrary control signals then there is no finite rate such that the system is uniform in control asymptotically observable.*

Proof: Assume that at time t the uncertainty set is Ω_t . The question then becomes how do we update the quantizer parameters in a way independent of the control such that the dynamic range contains the next state. Clearly this cannot be done uniformly over all control signals. \square

Note that in the case of encoder class one we could subtract out the effect of the control and thus could allow for arbitrarily large control signals. Let us restrict the controls.

Proposition 3.6.2 *Given system (3.1) with encoder in encoder class two with memory and $C = I$. Furthermore let the controls satisfy $\|U_t\|_2 \leq D$. If $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$ then there exists a scheme such that $\lim_{n \rightarrow \infty} \|e_t\|$ is bounded uniformly in control.*

Proof: Treat the control as noise and proceed as in proposition 3.5.2. \square

Corollary 3.6.1 *We are given system (3.1) with encoder in encoder class two with memory and $C = I$. The controls satisfy $\|U_t\|_2 \leq D_t$ where $D_t \leq \alpha^t D$ $0 \leq \alpha < 1$. If $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$ then there exists a scheme that is uniform in control asymptotically observable.*

Proof: Treat the control as noise and proceed as in proposition 3.5.3. \square

Corollary 3.6.2 *Given system (3.1), unstable A , and encoder in encoder class two without memory. The controls are set to zero. There is no finite rate such that the error is bounded.*

Proof: Let $\Gamma_t = \{X_t : X_t = A^t X_0\}$. Then for A unstable $\lim_{t \rightarrow \infty} \text{diam}(\Gamma_t) = \infty$. To ensure bounded error for a fixed rate we require a bounded dynamic range. This cannot be done here. \square

In conclusion, encoders in encoder class two with memory cannot achieve asymptotic observability without limiting the controls. For encoders in encoder class two without memory we cannot even achieve bounded state estimation error for the all zero control.

These observations may appear to be bad news. However in most cases we are interested in observability when the controls being applied are in the loop. The statements so far have been about observability uniform over some set of controls. We will show in the next section that we can achieve asymptotic observability if we consider the control objective of stability.

3.6.2 Stability

Here we treat stability under encoder class two. First we discuss encoders in encoder class two with memory and then encoders in encoder class 2 without memory. Assume (A, B) are controllable.

Proposition 3.6.3 *Given system (3.1) with encoder in encoder class two with memory, $C = I$, and knowledge of a bound on Λ_0 . Then there exists a finite rate such that the system can be made asymptotically stable.*

Proof: We will keep track of two sets. One set will track the error. The other set will contain the state. We will show that if the rate is suitably large then both sets will decrease to zero. We use \underline{L} to represent the sets in which the error lives in and \underline{M} to represent the sets in which the state lives in. The encoder does not know the control. But we allow it to know \underline{M} . We can allow this because \underline{M} can be computed offline. (It can be thought of as knowledge of the objective.)

Choose K such that $A + BK$ is stable. Let Φ_{A+BK} diagonalize $A + BK$ into real Jordan canonical form: $\Phi_{A+BK}(A + BK)\Phi_{A+BK}^{-1} = \Upsilon_{A+BK}$. Similarly let Φ_A diagonalize A into real Jordan canonical form: $\Phi_A A \Phi_A^{-1} = \Upsilon_A$. Recall from subsection 3.4.4 the definition of $\overline{\Upsilon}_A$ and $\overline{\Upsilon}_{A+BK}$.

Assume $\Lambda_0 \subset \{X : \|X\|_2 \leq L\}$. For $X \in \Lambda_0$ we have $\|\Phi_A X\|_2 \leq \|\Phi_A\| \|X\|_2 \leq \|\Phi_A\| L$. At time zero choose a $(C(0), \underline{R}, \underline{L}(0), \Phi(0))$ -quantizer where $C(0)$ is the origin and $\Phi(0) = \Phi_A$. Let $L_i(0) = \|\Phi_A\| L \forall i$. R_i will be determined shortly. Apply this quantizer to X_0 and transmit σ_0 . Note that σ_0 will not be the overflow symbol.

For $X \in \Lambda_0$ we have $\|\Phi_{A+BK} X\|_2 \leq \|\Phi_{A+BK}\| \|X\|_2 \leq \|\Phi_{A+BK}\| L$. At time zero let $M_i(0) = \|\Phi_{A+BK}\| L \forall i$.

The encoder does not have access to the control value. So we let the encoder treat the control as noise. The bound on this “noise” can be determined offline. Specifically

$$X_{t+1} = A\hat{X}_t + Ae_t + BK\hat{X}_t.$$

Both the encoder and decoder know $A\hat{X}_t$ and $\underline{M}(t)$. The encoder further knows Ae_t but does not know $BK\hat{X}_t$. It does know, though, that X_t is in the box determined by $\underline{M}(t)$ and thus so is \hat{X}_t . Let the encoder update the primitive quantizer parameters as follows: $C(t+1) = A\hat{X}_t$ and $\Phi(t+1) = H_A\Phi(t)$. The dimensions of the range of the $t+1$ th quantizer

will evolve according to:

$$\underline{L}(t+1) = \overline{\Upsilon}_A F_{\underline{R}} \underline{L}(t) + \|\Phi(t+1)BK\Phi_{A+BK}^{-1}\| \underline{M}(t).$$

By construction the error does not leave the box determined by \underline{L} .

The state evolution can also be written as

$$X_{t+1} = (A + BK)X_t - BK e_t.$$

Let $\underline{M}(t)$ evolve as:

$$\underline{M}(t+1) = \overline{\Upsilon}_{A+BK} \underline{M}(t) + \|\Phi_{A+BK} BK \Phi(t)^{-1}\| F_{\underline{R}} \underline{L}(t).$$

By construction the state never leaves the box determined by \underline{M} .

We write the coupled equation

$$\begin{bmatrix} \underline{L}(t+1) \\ \underline{M}(t+1) \end{bmatrix} = \begin{bmatrix} \overline{\Upsilon}_A F_{\underline{R}} & \|\Phi(t+1)BK\Phi_{A+BK}^{-1}\| I \\ \|\Phi_{A+BK} BK \Phi(t)^{-1}\| F_{\underline{R}} & \overline{\Upsilon}_{A+BK} \end{bmatrix} \begin{bmatrix} \underline{L}(t) \\ \underline{M}(t) \end{bmatrix}$$

Note that the controller K was chosen in such a way so that $A + BK$ would be stable. Thus $\overline{\Upsilon}_{A+BK}$ has real stable eigenvalues. If $\min_i R_i = \infty$ then the above coupled equation is stable. To see this note that $F_{\underline{R}} = 0$ and hence the lefthand column in the matrix consists of zero entries. Now we know that the eigenvalues of a matrix vary smoothly with the components of the matrix. By this continuity we see that there exists a finite rate vector \underline{R} such that the above coupled equation is stable. Thus $\lim_{t \rightarrow \infty} \underline{L}(t) = 0$ and $\lim_{t \rightarrow \infty} \underline{M}(t) = 0$. The system is asymptotically stable with finite rate. \square

Computing the minimal rate is difficult in general. Furthermore it will depend on the controller K . Note that for encoder class one we needed only to keep track of where the error lived. For encoder class two with memory we needed to keep track of both bounds on the error and the state. The lower bound on the rate for encoder one came from counting the number of control sequences it takes to drive the state to a certain ball around the origin. One can think of that control sequence as a codeword representing a region of space that X_0 lives in. In proposition 3.5.5 we showed how to realize that codeword sequentially. For encoders in encoder class two we cannot use the same counting argument because we cannot realize those codewords sequentially. We require a larger rate to stabilize the system. To carry the analogy further one could call the resulting control sequence a “redundant” description of the region of space X_0 lives in.

Now we treat the case of encoder class two without memory. In this case we need only keep track of one set. This set will contain both the state and the error.

Proposition 3.6.4 *Given system (3.1) with encoder in encoder class two without memory, $C = I$, and knowledge of a bound on Λ_0 . Then there exists a finite rate such that the system is asymptotically stable.*

Proof: Choose K such that $A+BK$ is stable. Let Φ diagonalize $(A+BK)$ into real Jordan canonical form: $\Phi(A+BK)\Phi^{-1} = \Upsilon$. Assume $\Lambda_0 \subset \{X : \|X\|_2 \leq L\}$. For $X \in \Lambda_0$ we have $\|\Phi X\|_2 \leq \|\Phi\| \|X\|_2 \leq \|\Phi\|L$. At time zero choose a $(C(0), \underline{R}, \underline{L}(0), \Phi)$ -quantizer where $C(0)$ is the origin. Let $L_i(0) = \|\Phi\|L \forall i$. R_i will be determined shortly. Apply this quantizer to X_0 and transmit σ_0 . Note that σ_0 will not be the overflow symbol.

The encoder does not have access to the controls or the past channel symbols. Thus it can only evolve according to a schedule. Note that

$$X_{t+1} = (A+BK)X_t - BK e_t.$$

We will find bounds \underline{L} on the state. Update the quantizer parameters as follows: $C(t) = 0$. The dimensions of the range of the $t+1$ th quantizer will evolve according to:

$$\underline{L}(t+1) = \left\{ \overline{\Upsilon} + \|\Phi BK \Phi^{-1}\| F_{\underline{R}} \right\} \underline{L}(t)$$

Recall that Υ is stable. Thus $\overline{\Upsilon}$ is stable with real eigenvalues. We can then find a rate vector \underline{R} large enough so that $\overline{\Upsilon} + \|\Phi BK \Phi^{-1}\| F_{\underline{R}}$ is stable. Thus $\lim_{t \rightarrow \infty} \underline{L}(t) = 0$. The system is asymptotically stabilizable. \square

Example 3.5.1 continued Take the scalar system $X_{t+1} = aX_t + bU_t$. Let k be such that $|a+bk| < 1$. Then we get

$$L(t+1) = (|a+bk| + \frac{|bk|}{2R})L(t)$$

Letting $R > \max\{0, \log \frac{|bk|}{1-|a+bk|}\}$ is sufficient to ensure asymptotic stability. Furthermore note that if $a+bk = 0$ then the rate bound becomes $R > \max\{0, \log |a|\}$.

Because the encoder does not know the control signals we operate with the closed loop dynamics of $A+BK$ and not the open loop dynamics of A as we did in the encoder class one case.

Now we treat the case where there is an observation equation. For encoder class 1 we considered an encoder that applied a Luenberger observer to the observation and computed an estimate of the state. It then transmitted the state estimate. Here, since we do not have memory, we consider an encoder that only quantizes the observation Y_n . The decoder will then apply a Luenberger observer to the quantized observation.

Proposition 3.6.5 Given system (3.1) and encoder class 2 without memory and with knowledge of a bound on Λ_0 . Then there exists a finite rate such that the system is asymptotically stable.

Proof: By assumption (A, B) are controllable and (A, C) are observable. Choose K such that $A+BK$ is stable. Let Φ_{A+BK} diagonalize $A+BK$ into real Jordan canonical form:

$\Phi_{A+BK}(A+BK)\Phi_{A+BK}^{-1} = \Upsilon_{A+BK}$. Choose H such that $A+HC$ is stable. Let Φ_{A+HC} diagonalize $A+HC$ into real Jordan canonical form: $\Phi_{A+HC}(A+HC)\Phi_{A+HC}^{-1} = \Upsilon_{A+HC}$

Let $\overline{C\Phi_{A+BK}^{-1}}$ equal the matrix $C\Phi_{A+BK}^{-1}$ with all its components set to their absolute values. Let $\overline{\Phi_{A+HC}H}$ equal the matrix $\Phi_{A+HC}H$ with all of its components set to their absolute values.

Assume $\Lambda_0 \subset \{X : \|X\|_2 \leq L\}$. For $X \in \Lambda_0$ we have $\|\Phi_{A+BK}X\|_2 \leq \|\Phi_{A+BK}\| \|X\|_2 \leq \|\Phi_{A+BK}\|L$. Let $L_i(0) = \|\Phi_{A+BK}\|L \forall i$.

If $\Phi_{A+BK}X_t$ is bounded by the box centered at zero with ranges $\underline{L}(t)$ then Y_n is bounded by the box centered at zero with ranges $\overline{C\Phi_{A+BK}^{-1}}\underline{L}(t)$.

At time zero choose a $(C(0), \underline{R}, \underline{M}(0), I)$ -quantizer where $C(0)$ is the origin. R_i will be determined shortly. Note that \underline{R} is l -dimensional. Apply this quantizer to Y_0 and transmit σ_0 . Note that σ_0 will not be the overflow symbol. The decoder receives σ_0 and decodes it as \hat{Y}_0 . Where \hat{Y}_0 the centroid of the region represented by σ_0 . Further set $C(t) = 0$ for all n .

Let $e_t = X_t - \hat{X}_t$ and $f_t = Y_t - \hat{Y}_t$. We apply the following observer at the decoder

$$\begin{aligned}\hat{X}_{t+1} &= A\hat{X}_t + BU_t + H(\hat{Y}_t - C\hat{X}_t) \\ &= A\hat{X}_t + BU_t + H(Y_t - f_t - C\hat{X}_t) \\ &= A\hat{X}_t + BU_t + HCe_t - Hf_t\end{aligned}$$

Thus

$$e_{t+1} = (A+HC)e_t - Hf_t.$$

We now describe the bounds on the error. Let $\underline{E}(0) = \|\Phi_{A+HC}\|L$. We want $\Phi_{A+HC}e_t$ to be bounded by the box centered at zero with ranges $\underline{E}(t)$. Thus update $\underline{E}(t)$ as follows:

$$\underline{E}(t+1) = \overline{\Upsilon}_{A+HC} \underline{E}(t) + \overline{\Phi_{A+HC}H} \overline{F_R} \overline{C\Phi_{A+BK}^{-1}} \underline{L}(t).$$

Apply the certainty equivalent controller K to the state estimate \hat{X} . Then

$$X_{t+1} = (A+BK)X_t - BKe_t.$$

We will find bounds \underline{L} on the state:

$$\underline{L}(t+1) = \overline{\Upsilon}_{A+BK} \underline{L}(t) + \|\Phi_{A+BK}BK\Phi_{A+HC}^{-1}\| \underline{E}(t)$$

We write the coupled equation

$$\begin{bmatrix} \underline{L}(t+1) \\ \underline{E}(t+1) \end{bmatrix} = \begin{bmatrix} \overline{\Upsilon}_{A+BK} & \|\Phi_{A+BK}BK\Phi_{A+HC}^{-1}\| \\ \overline{\Phi_{A+HC}H} \overline{F_R} \overline{C\Phi_{A+BK}^{-1}} & \overline{\Upsilon}_{A+HC} \end{bmatrix} \begin{bmatrix} \underline{L}(t) \\ \underline{E}(t) \end{bmatrix}$$

By construction both $\overline{\Upsilon}_{A+BK}, \overline{\Upsilon}_{A+HC}$ are stable with real eigenvalues. Thus for \underline{R} large enough the above coupled equation is stable. Thus the system is asymptotically stable. \square .

We conclude this section with some general comments on stability under rate constraints. Our schemes involve keeping track of the uncertainty. For encoder class one this involves keeping track of the state estimation error. For encoder class two with memory this involves keeping track of both the state and the estimation error. This is also true for encoder class two without memory. However in this case they are the same regions. These regions essentially bound the reachable set of states at every time step.

It is natural to work with these regions. The definition of stability, in particular uniform attractivity, requires their existence. By computing these regions and explicitly stating how they evolve we are able to determine what rate vectors and primitive quantizers are needed.

3.6.3 Controllability

Here we discuss controllability under encoder class two. Recall that for any P there exist $\bar{U}_0, \dots, \bar{U}_{d-1}$ such that $P = \sum_{j=0}^{d-1} A^{d-1-j} B \bar{U}_j$.

Proposition 3.6.6 *For system (3.1) with encoder in encoder class two with memory or encoder class two without memory and bounded Λ_0 there exists a finite rate scheme such that the system is uniform in initial state controllable under a rate constraint.*

Proof: By propositions 3.6.3 and 3.6.4 we know given $\epsilon > 0$ there exists a time $T(\epsilon)$ such that the state can be driven uniformly in initial in initial state to within $\|X_T\| \leq \frac{\epsilon}{\|A^d\|}$. Once there apply the controls $\bar{U}_0, \dots, \bar{U}_{d-1}$. Then $X_{T+d} = A^d X_T + P$. However $\|X_{T+d} - P\| = \|A^d X_T\| \leq \epsilon$. \square

Corollary 3.6.3 *For system (3.1) and encoder class two with or without memory there exists a finite rate scheme such that the system is controllable under a rate constraint.*

Proof: As usual apply the zero control until the encoder has captured the state and then apply proposition 3.6.6. \square

In conclusion the general prescription for controllability is to drive the system to the origin and then drive it to the point of interest. It would be more reasonable to drive the initial state to the destination point directly. Since we need to determine where the state is it is best to drive the system to a known point, in this case, the origin. We drive the system to the origin until we have “learned” its position sufficiently well.

3.6.4 Lyapunov synthesis

A traditional method for designing controllers is to use Lyapunov theory. The prescription is here is to find a suitable Lyapunov function V such that along trajectories V is decreasing. For finite rate it is impossible to ensure V decreases at every time step. This is because there is always a region around the origin that is not under the influence of any control. Though Lyapunov theory is helpful for proving existence of schemes it is difficult to determine optimal rate vectors.




Figure 3-3: Lyapunov regions

We will first show that given any quadratic Lyapunov function one can find regions that contain the state and shrink with time. We then derive from this quadratic Lyapunov function a stabilizing scheme for encoder class two without memory.

Let $X_{t+1} = AX_t$ with A stable. Then there exist symmetric positive definite matrices P, Q such that

$$X_{t+1}'PX_{t+1} - X_t'PX_t = -X_t'QX_t.$$

Now assume that $X_t'PX_t \leq L_t$. We will compute an $L_{t+1} < L_t$ such that $X_{t+1}'PX_{t+1} \leq L_{t+1}$. Fix $0 < \Delta < L_t \frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)}$.

If $X_t'QX_t \geq \Delta$ then $X_{t+1}'PX_{t+1} \leq L_t - \Delta$. Else if $X_t'QX_t < \Delta$ then $X_t'X_t < \frac{\Delta}{\lambda_{\min}(Q)}$. Thus $X_{t+1}'PX_{t+1} < X_t'PX_t \leq X_t'X_t \lambda_{\max}(P) < \Delta \frac{\lambda_{\max}(P)}{\lambda_{\min}(Q)} < L_t$.

Choose $L_{t+1} = \max \left\{ L_t - \Delta, \Delta \frac{\lambda_{\max}(P)}{\lambda_{\min}(Q)} \right\}$. Then $L_{t+1} < L_t$. Thus the smallest $L_{t+1} = \frac{\lambda_{\max}(P)}{\lambda_{\min}(Q) + \lambda_{\max}(P)} L_t$. See figure 3-3.

Points, X_t , between the ellipse defined by L_t and L_{t+1} move inside the ellipse defined by L_{t+1} . Points, X_t , already inside the ellipse defined by L_{t+1} stay in that ellipse.

Now we will derive from a given quadratic Lyapunov function a scheme to stabilize the system (3.1) with encoder in encoder class two without memory and $C = I$.

Proposition 3.6.7 *Given a quadratic Lyapunov function for the full state observation system (3.1) one can derive a finite rate asymptotically stable scheme for encoder class two without memory.*

Proof: Let K be such that $A + BK$ is stable. Let P, Q satisfy the discrete time Lyapunov matrix equation

$$(A + BK)'P(A + BK) - P = -Q.$$

Now $X_{t+1} = (A + BK)X_t - BK e_t$. Assume that $X_t'PX_t \leq L_t$. Then $\|X_t\|^2 \leq \frac{L_t}{\lambda_{\min}(P)}$. Let the primitive quantizer at time n be tuned to the axis of P and use a uniform number of levels in each coordinate direction. Then $\|e_t\|^2 \leq \frac{L_t}{\lambda_{\min}(P)2^{\frac{2R}{d}}}$. Thus

$$\begin{aligned} & X_{t+1}'PX_{t+1} - X_t'PX_t \\ &= -X_t'QX_t + 2X_t'(A + BK)'PBK e_t + e_t'(BK)'PBK e_t \\ &\leq -X_t'QX_t + 2\|(A + BK)'PBK\|\|X_t\|\|e_t\| + \|(BK)'PBK\|\|e_t\|^2 \\ &\leq -X_t'QX_t + 2\|(A + BK)'PBK\|\frac{L_t}{\lambda_{\min}(P)}\frac{1}{2^{\frac{R}{d}}} \\ &\quad + \|(BK)'PBK\|\frac{L_t}{\lambda_{\min}(P)}\frac{1}{2^{\frac{2R}{d}}} \end{aligned}$$

Let $\beta(R) = 2\|(A + BK)'PBK\|\frac{1}{\lambda_{\min}(P)}\frac{1}{2^{\frac{R}{d}}} + \|(BK)'PBK\|\frac{1}{\lambda_{\min}(P)}\frac{1}{2^{\frac{2R}{d}}}$. Choose R large enough so that $\beta(R) < \frac{\lambda_{\min}(Q)}{\lambda_{\min}(Q) + \lambda_{\max}(P)}$. Then fix $0 < \Delta < L_t \left(\frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)} - \left(\frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)} + 1 \right) \beta(R) \right)$.

If $X_t'QX_t - \beta(R)L_t \geq \Delta$ then $X_{t+1}'PX_{t+1} \leq L_t - \Delta$.

Else if $X_t'QX_t - \beta(R)L_t < \Delta$ then $X_t'X_t < \frac{\Delta + \beta(R)L_t}{\lambda_{\min}(Q)}$. Thus

$$\begin{aligned} X_{t+1}'PX_{t+1} &< X_t'PX_t + \beta(R)L_t \\ &\leq X_t'X_t\lambda_{\max}(P) + \beta(R)L_t \\ &< \frac{\lambda_{\max}(P)}{\lambda_{\min}(Q)}(\Delta + \beta(R)L_t) + \beta(R)L_t \\ &< L_t. \end{aligned}$$

Choose $L_{t+1} = \max \left\{ L_t - \Delta, \frac{\lambda_{\max}(P)}{\lambda_{\min}(Q)}(\Delta + \beta(R)L_t) + \beta(R)L_t \right\}$. Then $L_{t+1} < L_t$. In general one would want to optimize the rate R over P, Q , and the controller K . \square

For related results see Liberzon and Brockett. [LB]

3.6.5 Minimum communication

In our discussion we constructed encoders, decoders, and controllers that would stabilize the system by incorporating an asymptotic observer. Specifically we showed that if the system is state feedback stabilizable and the state estimation error goes to zero then one could stabilize the system via output feedback.

The converse observation is of interest also. If a scheme is output stabilizable then there exists a state estimator with error going to zero. Clearly if we know the system is converging to the origin then a good asymptotic observer would also say the origin. Thus if there do not exist any asymptotic observers then the system cannot be stabilized. This suggests the question: what is the minimal requirement on the state error to insure stability?

For all three encoder classes we computed regions defined by the \underline{L}_t vectors. By choosing the rate suitably large we were able to show that the \underline{L}_t would shrink to zero. Note though that in these proofs all we required is that these regions shrink. The digital channel is just one of many ways to realize an “end-to-end” error bound on \underline{L}_t . One could also imagine a channel $Y_t = X_t + v_t$ where the disturbances v_t lie in the region defined by \underline{L}_t . This can be considered an analog channel with increasing SNR. In chapter five we will discuss in more detail the relationship between digital channels of a fixed rate and additive noise channels with fixed SNR.

In our schemes the size of \underline{L}_t depends solely on \underline{L}_0 and t (and of course A, B, K). Another approach would be to have \underline{L}_t shrink as X_t shrinks. This could be realized, for example, by a sector bounded nonlinearity. In the scalar case this would be $\underline{\phi}(X_t) \leq X_t \leq \bar{\phi}(X_t)$. Where $\underline{\phi}(\cdot), \bar{\phi}(\cdot)$ are class K functions. (A class K function is an increasing, continuous function that is zero at the origin. [Son]) In this case $\Lambda_t = [\underline{\phi}(X_t), \bar{\phi}(X_t)]$. [Eli]

Finally if we assume that we are going to use a certainty equivalent control scheme then it can be shown that

$$X_n = (A + BK)^n X_0 - \sum_{j=1}^{n-1} (A + BK)^{n-1-j} BK e_j.$$

By lemma 3.5.1 we know that if $e_n \rightarrow 0$ then $X_n \rightarrow 0$. Thus under the certainty equivalent controller $e_n \rightarrow 0$ implies $X_n \rightarrow 0$.

As stated before a necessary condition for stability is that there exist a state estimator such that the state estimation error goes to zero. We just showed that for the certainty equivalent scheme the error going to zero is also sufficient. Thus the minimal requirement on the state error to insure stability is that there exists a state estimator such that the state estimation error go to zero. It does not matter at what rate. Of course to insure a given rate of convergence one will need to determine a rate of convergence on the error.

Note that in the proof of the lower bound for stability under encoder class one, proposition 3.3.2, no mention was made of how the rate was distributed over time. Thus if we allow for time-varying rate we have another parameter that we can choose in our encoder design.

In conclusion the end-to-end error can be realized by many different forms of channels. For example three are: digital channels, power constrained analog channels, and sector bounded nonlinearities. The digital channel and the analog channel require coordination between the encoder and decoder. The example of the sector bounded nonlinearity suggests that there should be ties to robust control and specifically the theory of stability margins.

Thus in a certain sense we have shown that we can separate the communication part from the control part in our stability problem. The communication part has the job of delivering a state estimation error that decreases over time.

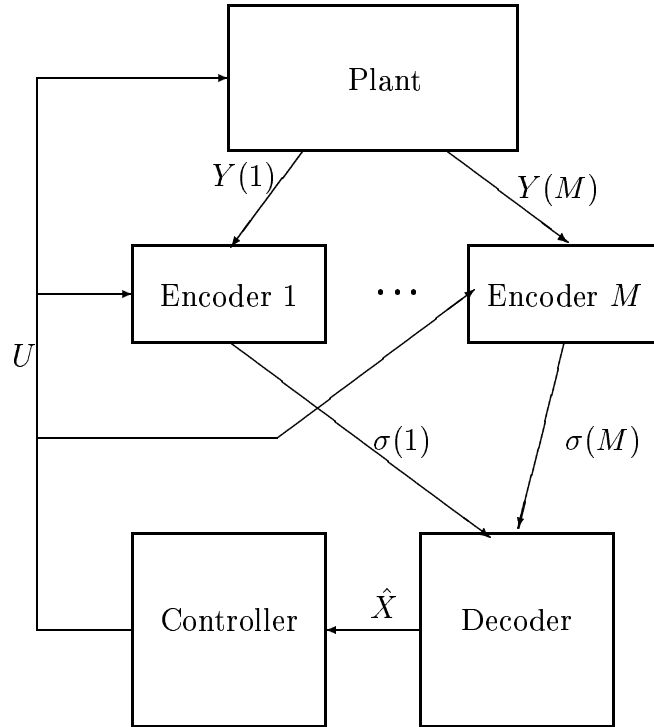


Figure 3-4: Distributed Sensor System

3.7 Systems with Multiple Sensors

In this section we examine a particular control problem with multiple sensors. We first formulate the problem. Then we relate it's solution to the Slepian-Wolf coding theorem. [CT]

3.7.1 Problem setup

Consider the linear system with M distributed observations and one controller:

$$X_{t+1} = AX_t + BU_t, \quad Y(i)_t = C_i X_t \quad i = 1, \dots, M \quad (3.3)$$

where $A \in \mathbb{R}^{d \times d}$, $B \in \mathbb{R}^{d \times m}$, $C_i \in \mathbb{R}^{l_i \times d}$, and $X_0 \in \Lambda_0$. See figure 3-4.

The encoder, decoder and controller are specified now. These are essentially the same definitions we have been using with appropriate generalization to the distributed setting. (Recall the definitions of encoder, decoder, and controller from section 3.2.1.)

Encoder Class: For each $i = 1, \dots, M$ encoder \mathcal{E}_i is a map that takes $(Y(i)^t, \sigma(i)^{t-1}, U^{t-1}) \mapsto \sigma(i)_t$. Note that encoder \mathcal{E}_i is allowed to depend only on its own past and current

observations and its own past channel signals. All the encoders are allowed to observe all the controls. Thus using our previous terminology these encoders are in encoder class one.

Decoder Class: The decoder \mathcal{D} is a map that takes $(\sigma(1)^t, \dots, \sigma(M)^t, U^{t-1}) \mapsto \Omega_t$. This is a centralized decoder in that it is allowed to observe all the different channel signals produced by the M encoders.

Controller: The controller \mathcal{C} is a map that takes $\Omega_t \mapsto U_t$. We will assume our usual certainty equivalent controller.

We will treat the problems of observability and stability. Recall the definitions of asymptotic observability and asymptotic stabilizability. (Definitions 3.3.3, 3.3.4, and 3.3.5.)

Let $C = [C'_1, \dots, C'_M]'$.

Assumption 3.7.1 *The pair (A, C) is detectable and the pair (A, B) is stabilizable.*

Though (A, C) is detectable it is not generally the case that (A, C_i) will be detectable. Define \mathcal{N}_i to be the unobservable subspace associated with (A, C_i) . Define \mathcal{O}_i to be the subspace orthogonal to \mathcal{N}_i . It can be shown that \mathcal{N}_i is indeed a subspace and furthermore that it is A -invariant. [Son] By assumption 3.7.1 we have $\cap_{i=1}^M \mathcal{N}_i = \emptyset$. Also any $x \in \mathcal{O}_i$ can be observed by encoder i .

Associate with each i the set $\Lambda_i = \{\lambda(A) : \text{those eigenvalues of } A \text{ corresponding to the subspace } \mathcal{O}_i\}$. We have $\cup_{i=1}^M \Lambda_i = \{\lambda(A)\}$. In general the Λ_i will not be disjoint. Because they are not disjoint we have freedom in determining what each encoder sends to the decoder.

We will show that this freedom can be captured by the following “rate region.” For each encoder \mathcal{E}_i define its rate vector to be $\underline{R}_i = (R_{i,1}, \dots, R_{i,\dim \mathcal{O}_i})$. (Where \underline{R}_i is the rate vector used in the primitive quantizer by \mathcal{E}_i .)

Definition 3.7.1 *For a given matrix A and observation matrices C_1, \dots, C_M define*

$$\mathcal{R} = \left\{ (\underline{R}_1, \dots, \underline{R}_M) : \sum_{i:\lambda(A) \in \Lambda_i} R_{i,j_{\lambda(A)}} \geq \max\{0, \log |\lambda(A)|\} \quad \forall \lambda(A) \right\}$$

where $j_{\lambda(A)}$ represents the index of the rate component associated with that eigenvalue.

Proposition 3.7.1 *Given system (3.3) a necessary condition on the rate so that the system is uniform in control asymptotically observable is that $(\underline{R}_1, \dots, \underline{R}_M) \in \mathcal{R}$.*

Proof: Follows directly from proposition 3.3.1. \square

We now show that for any rate vector in \mathcal{R} there exists an encoder such that the system is uniform in control asymptotically observable.

Proposition 3.7.2 *Given system (3.3) a sufficient condition on the rate so that the system is uniform in control asymptotically observable is that $(\underline{R}_1, \dots, \underline{R}_M) \in \mathcal{R}$.*

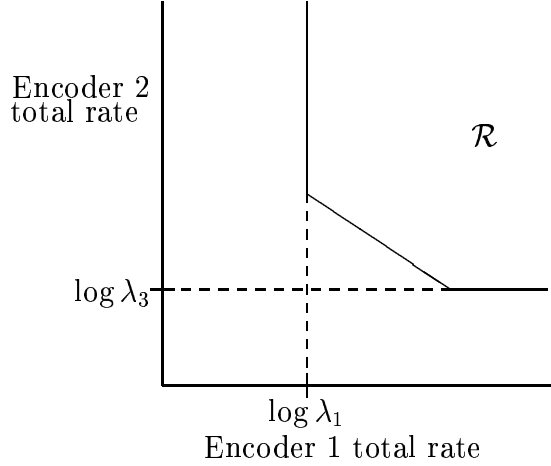


Figure 3-5: Rate region for example 3.7.1

Proof: We prove it for the case when $M = 2$. The proof for the more general case is straightforward. In this case there are three sets of interest: $\Lambda_1 \cap \Lambda_2$, $\Lambda_1 \setminus (\Lambda_1 \cap \Lambda_2)$, and $\Lambda_2 \setminus (\Lambda_1 \cap \Lambda_2)$. By proposition 3.5.1 we know that the decoder can asymptotically observe the subspace associated with $\Lambda_1 \setminus (\Lambda_1 \cap \Lambda_2)$ with $\sum_{\lambda \in \Lambda_1 \setminus (\Lambda_1 \cap \Lambda_2)} \max\{0, |\lambda|\}$ bits transmitted only by encoder 1. Similarly the decoder can asymptotically observe the subspace associated with $\Lambda_2 \setminus (\Lambda_1 \cap \Lambda_2)$ with $\sum_{\lambda \in \Lambda_2 \setminus (\Lambda_1 \cap \Lambda_2)} \max\{0, |\lambda|\}$ bits transmitted only by encoder two.

This leaves us with $\Lambda_1 \cap \Lambda_2$. Clearly we don't need both encoder one and encoder two sending the same information. Thus any splitting of the rate between encoder 1 and encoder 2 needed to describe $\Lambda_1 \cap \Lambda_2$ is sufficient for the decoder to asymptotically observe the subspace associated with $\Lambda_1 \cap \Lambda_2$. We only require that the combined rate used to describe $\Lambda_1 \cap \Lambda_2$ be greater than $\sum_{\lambda \in \Lambda_1 \cap \Lambda_2} \max\{0, \log |\lambda|\}$. For example encoder one can send coarse bits while encoder two sends fine bits (i.e. most significant digits and least significant digits.) \square

Example 3.7.1 Let $A = \text{diag}[\lambda_1, \lambda_2, \lambda_3]$ where $\lambda_i > 1$ $i = 1, 2, 3$. Let $C_1 = [1, 1, 0]$ and $C_2 = [0, 1, 1]$. The rate region \mathcal{R} is shown in figure 3-5.

Now we provide lower and upper bounds for asymptotic stability.

Proposition 3.7.3 Given system (3.3) a necessary condition on the rate so that the system is asymptotically stabilizable is that $(\underline{R}_1, \dots, \underline{R}_M) \in \mathcal{R}$.

Proof: Follows from proposition 3.3.2. \square

Proposition 3.7.4 Given system (3.3) a sufficient condition on the rate so that the system is asymptotically stabilizable is that $(\underline{R}_1, \dots, \underline{R}_M) \in \mathcal{R}$.

Proof: Follows from the propositions 3.5.6 and 3.7.2. \square

3.7.2 Relationship to Slepian-Wolf Coding

The distributed sensor problem we have set up is really a problem of distributed encoding for correlated sources. Thus it naturally falls under the purview of Slepian-Wolf coding theory.

The traditional Slepian-Wolf problem concerns itself with transmitting the random variable (X_1, \dots, X_M) losslessly. Encoder i observes X_i and transmits an encoding at rate R_i . For the case where $M = 2$ the rates must satisfy the following three inequalities:

- (1) $R_1 + R_2 \geq H(X_1, X_2)$
- (2) $R_1 \geq H(X_1|X_2)$
- (3) $R_2 \geq H(X_2|X_1)$

where $H(\cdot)$ and $H(\cdot|\cdot)$ are the discrete entropy and conditional discrete entropy respectively. See the appendix for definitions of these terms.

Under the correspondence

$$H(X_i) \iff \sum_{\lambda \in \Lambda_i} \max\{0, \log |\lambda|\}$$

and

$$H(X_i|X_j) \iff \sum_{\lambda \in \Lambda_i - (\Lambda_i \cap \Lambda_j)} \max\{0, \log |\lambda|\}$$

where $i, j = 1, 2$ we get that \mathcal{R} is indeed the region specified by the Slepian-Wolf region.

3.8 Sampling of Continuous Time Systems

In this section we discuss some issues of time-sampling. For the duration of section 3.8 we assume that t represents the continuous time index and that n represents the sampled time index.

The time-sampling we have in mind is a uniform sampling every T time units. That sample is then quantized by the encoder and transmitted to the decoder/controller. There the controller instantaneously computes a control signal and applies it. We use a sample and hold control sequence.

One can imagine many other sampling schemes. One particularly interesting scheme was proposed by Karl Astrom. He suggests that we should sample the state whenever an “interesting” event has occurred instead of sampling at regularly spaced time epochs. He calls uniform sampling “Riemann” sampling and his type of sampling “Lebesgue” sampling. [Ast] Furthermore this type of sampling can be related to event driven systems. In this section we restrict attention to uniform sampling.

We consider the following continuous time linear system:

$$\dot{X}(t) = AX(t) + BU(t) \quad (3.4)$$

where $X(t)$ is a \mathbb{R}^d -valued state process and $U(t)$ is a \mathbb{R}^m -valued control process. The initial position $X_0 \in \Lambda_0$ where $\Lambda_0 \subseteq \mathbb{R}^d$.

Let the sampling period be T . Define $X_n = X(nT)$. Let $U(t) = U_n$ $nT \leq t < (n+1)T$ for some control sequence $\{U_n\}$. Then for $t = nT + \delta$ we have

$$X(t) = e^{\delta A} X_n + \int_0^\delta e^{(\delta-\tau)A} B U_n d\tau.$$

We can rewrite this as

$$X_{n+1} = \bar{A}(T)X_n + \bar{B}(T)U_n$$

where $\bar{A}(T) = e^{TA}$ and $\bar{B}(T) = \int_0^T e^{(T-\tau)A} B d\tau$.

One can ask what is the best sampling period for achieving some control objective. Where by “best” we mean the minimum number of bits per time unit. Thus if we have a scheme that samples the system every T time units and transmits R bits then the number of bits per time unit is $\frac{R}{T}$. In general the rate R will depend on T .

3.8.1 Observability

First we repeat the definitions of asymptotic observability for continuous time systems. Let $e(t) = X(t) - \hat{X}(t)$.

Definition 3.8.1 *System (3.4) is asymptotically observable if there exists a sampling period T , a control sequence $\{u_n\}$, and an encoder and decoder such that*

$$1 \text{ Stability: } \forall \epsilon > 0 \quad \exists \delta(\epsilon) \text{ such that } \|X(0)\|_2 \leq \delta(\epsilon) \text{ implies } \|e(t)\|_2 \leq \epsilon \quad \forall t \geq 0.$$

2 *Uniform attractivity*: $\forall \epsilon > 0, \delta > 0 \exists M(\epsilon, \delta)$ such that $\|X(0)\|_2 \leq \delta$ implies $\|e(t)\|_2 \leq \epsilon \forall t \geq M$.

Point one states that the error cannot grow unbounded for bounded X_0 . The second point states that the error decreases to zero uniformly in X_0 . Note also that uniform attractivity is defined for all δ . Thus our definition of asymptotic observability is global.

Definition 3.8.2 *System (3.4) is uniform in control asymptotically observable if there exists a sampling period T , and an encoder and decoder independent of the control sequence applied such that the system is asymptotically observable.*

Proposition 3.8.1 *Given system (3.4) and encoder in encoder class one. Further assume that the encoder knows a bound on Λ_0 . Assume the sampling period is T . A sufficient condition on the rate for uniform in control asymptotic observability is $R > \sum_{\lambda(\bar{A})} \{0, \log |\lambda(\bar{A})|\}$.*

Proof: From proposition 3.5.1 we can show that $\lim_{n \rightarrow \infty} e(nT) = 0$. We need only show that the error in any intersample period stays bounded. Specifically let $t = nT + \delta$

$$\begin{aligned} e(t) &= X(t) - \hat{X}(t) \\ &= e^{\delta A} X_n + \int_0^\delta e^{(\delta-\tau)A} B U_n d\tau - \left(e^{\delta A} \hat{X}_n + \int_0^\delta e^{(\delta-\tau)A} B U_n d\tau \right) \\ &= e^{\delta A} e(nT) \end{aligned}$$

For finite T we have $\sup_{0 \leq \delta \leq T} \|e^{\delta A}\|$ is bounded. \square

Proposition 3.8.2 *Given system (3.4) and encoder in encoder class one. Assume the sampling period is T . A necessary condition for uniform in control asymptotic observability is $R > \sum_{\lambda(\bar{A})} \{0, \log |\lambda(\bar{A})|\}$.*

Proof: By proposition 3.3.1 the rate condition is necessary for $e(nT)$ to go to zero. Clearly if $e(nT)$ does not go to zero then $e(t)$, for general intersample times t , will not go to zero. \square

Let A have eigenvalues $\lambda_1, \dots, \lambda_d$. Then \bar{A} has eigenvalues $\bar{\lambda}_i = e^{T\lambda_i} \forall i = 1 \dots d$. Furthermore if λ_i is real then $\bar{\lambda}_i$ is real. If $\lambda_i = a \pm bi$ is complex then $\bar{\lambda}_i = e^a (\cos b \pm i \sin b)$ is complex. The rate is

$$\begin{aligned} \frac{R(T)}{T} &> \frac{1}{T} \sum_{\lambda(\bar{A})} \max \{0, \log |\lambda(\bar{A})|\} \\ &= \sum_{\lambda(A)} \max \left\{ 0, \frac{\text{Real}(\lambda(A))}{\ln 2} \right\} \end{aligned}$$

Thus for encoder class one the optimal number of bits per time unit is independent of the sampling period T . We are assuming, though, that at the start of every epoch we are

instantaneously transmitting $R(T)$ bits. So though the average rate $\frac{R(T)}{T}$ is independent of T the peak value is very much dependent on T .

For encoders in encoder class two there is little hope of a general theorem on uniform in control asymptotic observability. See our discussion of asymptotic observability for discrete time linear systems. We can find cases for asymptotic observability under feedback. We discuss this in the next section on stability.

3.8.2 Stability

Here we treat the problem of stability under sampling. Assume that (A, B) are controllable. Then by theorem 3.4.4 of [Son] we have

Theorem 3.8.1 *If (A, B) controllable and T is such that*

$$T(\lambda - \mu) \neq 2k\pi i \quad k = \pm 1, \pm 2, \dots$$

for every two eigenvalues of A then (\bar{A}, \bar{B}) are controllable.

Note that this condition need not be necessary.

Definition 3.8.3 *System (3.4) is asymptotically stabilizable if there exists a sampling period T , and an encoder, decoder, and controller such that*

- 1 *Stability:* $\forall \epsilon > 0 \quad \exists \delta(\epsilon)$ such that $\|X(0)\|_2 \leq \delta(\epsilon)$ implies $\|X(t)\|_2 \leq \epsilon \quad t \geq 0$.
- 2 *Uniform attractivity:* $\forall \epsilon > 0, \delta > 0 \quad \exists M(\epsilon, \delta)$ such that $\|X(0)\|_2 \leq \delta$ implies $\|X(t)\|_2 \leq \epsilon \quad \forall t \geq M$.

Encoder class 1

Proposition 3.8.3 *Given system (3.4) with encoder in encoder class one. The encoder knows a bound on Λ_0 . The rate $R > \sum_{\lambda(\bar{A})} \max\{0, \log |\lambda(\bar{A})|\}$ is a sufficient condition for asymptotic stability.*

Proof: Assume T satisfies the condition of theorem 2.8.1. Then (\bar{A}, \bar{B}) are controllable. By proposition 3.5.5 there exists a scheme such that both $\lim_{n \rightarrow \infty} X(nT) = 0$ and $\lim_{n \rightarrow \infty} e(nT) = 0$. We need only show that the intersample behavior of $X(t)$ is bounded. Specifically let $t = NT + \delta$

$$\begin{aligned} X(t) &= e^{\delta A} X_n + \int_0^\delta e^{(\delta-\tau)A} B K \hat{X}_n d\tau \\ &= e^{\delta A} X_n + \int_0^\delta e^{(\delta-\tau)A} B K X_n d\tau - \int_0^\delta e^{(\delta-\tau)A} B K e_n d\tau \end{aligned}$$

Thus $\|X(t)\| \leq \|e^{\delta A} + \int_0^\delta e^{(\delta-\tau)A} BK d\tau\| \|X_n\| + \|\int_0^\delta e^{(\delta-\tau)A} BK d\tau\| \|e_n\|$. For finite T this is bounded. \square

Proposition 3.8.4 *Given system (3.4) and encoder class one with knowledge of a bound on Λ_0 . The rate $R > \sum_{\lambda(\bar{A})} \max\{0, \log |\lambda(\bar{A})|\}$ is a necessary condition for asymptotic stability.*

Proof: By proposition 3.3.1 the rate condition is necessary for $X(nT)$ to go to zero. Clearly if $X(nT)$ does not go to zero then $X(t)$, for general intersample times t , will not go to zero. \square

Encoder class 2

Proposition 3.8.5 *Given system (3.4) with encoder in encoder class two without memory. Assume the encoder knows a bound on Λ_0 . Then there exists a finite rate such that the system is asymptotically stable.*

Proof: Assume T satisfies the condition of theorem 3.8.1. Then (\bar{A}, \bar{B}) are controllable. By proposition 3.6.4 there exists a scheme such that both $\lim_{n \rightarrow \infty} X(nT) = 0$ and $\lim_{n \rightarrow \infty} e(nT) = 0$. We need only show that the intersample behavior of $X(t)$ is bounded. This follows analogously to the proof of proposition 3.8.3. \square

Recall that the rate vector \underline{R} needs to be large enough so that the matrix

$$\left\{ \bar{\Upsilon} + \|\Phi \bar{B} K \Phi^{-1}\| F_{\underline{R}} \right\}$$

where $(\bar{A} + \bar{B}K) = \Phi^{-1} \Upsilon \Phi$ is the real canonical Jordan form. \underline{R} will in general depend on both \bar{A}, \bar{B} , and K . Now \bar{A}, \bar{B} , and K in turn depend on T . Thus for encoder class two without memory $\frac{R}{T}$ can have a complicated dependence on T .

Example 3.8.1 *Take the scalar system $\dot{X} = aX(t) + bU(t)$ $a > 0$. Then $\bar{a} = e^{Ta}$ and $\bar{b} = \frac{b(e^{Ta} - 1)}{a}$. Let k be such that $|\bar{a} + \bar{b}k| < 1$. Define $\alpha = \bar{a} + \bar{b}k$. Then $R > \max\{0, \log \frac{|\bar{b}k|}{1 - |\bar{a} + \bar{b}k|}\}$ is sufficient to insure asymptotic stability. Now*

$$\frac{R(T)}{T} > \frac{1}{T} \log \frac{|\bar{b}k|}{1 - |\bar{a} + \bar{b}k|}$$

For fixed T this lower bound is decreasing in α . For $T \rightarrow \infty$ this lower bound converges to $\frac{a}{\ln 2}$. Thus this suggests that if one wants to minimize the bits per time unit then one should use a large sampling period. Of course the excursions can be large for large sampling periods.

In conclusion under encoder class one the rate is essentially independent of the sampling rate. Whereas for encoder class two the rate is dependent on the sampling rate.

3.9 Performance

In this section we address two methods for determining the rate needed to achieve a given performance objective. We first describe a general optimal control setup. We then describe the LQ problem. We give an upper bound on the rate-cost function for both encoder class one and encoder class two. We then describe covering numbers. We apply this to a problem stability where the trajectories are strictly decreasing to the origin.

3.9.1 Problem setup

We assume the linear system (3.1): $X_{t+1} = AX_t + BU_t$ and a continuous cost $g : (X, U) \rightarrow \mathbb{R}$ with $g(0, 0) = 0$. And our goal is to minimize, for each initial position, the infinite horizon cost

$$\sum_{t=0}^{\infty} g(X_t, U_t)$$

over all control laws $U = k(X)$.

Let V be the optimal cost function. Then the control law k is optimal if it achieves the minimum of the Hamilton-Jacobi-Bellman equation: [Son]

$$V(AX + Bk(X)) - V(X) + g(X, k(X)) = \min_U \{V(AX + BU) - V(X) + g(X, U)\} = 0 \quad \forall X, U.$$

Under any other stabilizing control law \tilde{k} we have

$$V(AX + B\tilde{k}(X)) - V(X) + g(X, \tilde{k}(X)) \geq 0.$$

We can use Taylor's expansion to bound the difference between the cost under \tilde{k} and the optimal k

$$\begin{aligned} & \sum_{t=0}^{\infty} g(X_t^{\tilde{k}}, \tilde{k}(X_t^{\tilde{k}})) - V(X_0) \\ = & \sum_{t=0}^{\infty} [g(X_t^{\tilde{k}}, \tilde{k}(X_t^{\tilde{k}})) - g(X_t^k, k(X_t^k))] \\ = & \sum_{t=0}^{\infty} \left\{ \nabla_X g(X_t^k, k(X_t^k))(X_t^{\tilde{k}} - X_t^k) \right. \\ & \left. + \nabla_U g(X_t^k, k(X_t^k))(\tilde{k}(X_t^{\tilde{k}}) - k(X_t^k)) + \text{higher order terms} \right\} \end{aligned}$$

Thus we can bound the loss due to a suboptimal controller \tilde{k} by a measuring how much the trajectories diverge.

Under full state observation we can achieve the optimal cost. Under a rate constraint we will incur a larger cost. Our goal is to understand the tradeoff between the rate and the extra incurred cost. In the next section we use the above bound to analyze the extra cost incurred under a rate constraint for the LQ problem.

3.9.2 LQ

For the LQ problem the running cost has the form $g(X, U) = X'QX + U'TU$ where Q is positive semidefinite and T is positive definite. For the infinite horizon problem the optimal control law is linear and has the form $KX = -(B'PB + T)^{-1}B'PAX$ where P satisfies the Riccati equation

$$P = A'(P - PB(B'PB + T)^{-1}B'P)A + Q.$$

For a given rate R we can upper bound the extra cost incurred by assuming a certainty equivalent scheme applied to the state estimate. Let X_t, Y_t represent the state process under the optimal scheme and the certainty equivalent scheme respectively. Let \hat{Y}_t be the state estimate of Y_t and $e_t = Y_t - \hat{Y}_t$. Then

$$\begin{aligned} \Delta &= \sum_{t=0}^{\infty} g(Y_t, K\hat{Y}_t) - V(X_0) \\ &= \sum_{t=0}^{\infty} [Y_t'QY_t + (K\hat{Y}_t)'TK\hat{Y}_t - (X_t'QX_t + (KX_t)'TKX_t)] \\ &= \sum_{t=0}^{\infty} [(Y_t + X_t)'(Q + K'TK)(Y_t - X_t) - 2Y_t'K'TKe_t + e_tK'TKe_t] \end{aligned}$$

We can write a dynamics for the term $Y_t - X_t$:

$$\begin{aligned} Y_{t+1} - X_{t+1} &= AY_t + BK\hat{Y}_t - AX_t - BKX_t \\ &= (A + BK)(Y_t - X_t) - BKe_t \end{aligned}$$

with $Y_0 = X_0$. Thus

$$(Y_t - X_t) = -\sum_{i=0}^{t-1} (A + BK)^{t-i-1} BKe_i$$

Proposition 3.9.1 *Given system (3.1) with encoder in encoder class one, bounded Λ_0 and quadratic cost. The difference between the optimal cost for a given rate R and the optimal cost under full state observation is bounded for $R > \sum_{\lambda(A)} \max\{0, \log |\lambda(A)|\}$.*

Proof: By proposition 3.5.1 we know there exists an encoder/decoder such that $\|e_t\| \leq \|\Phi(t)^{-1}\| \|F_R\| \|(\bar{\Upsilon}_A F_R)^t\| \|\underline{L}(0)\|$. Which goes to zero exponentially. Now

$$\|(Y_t - X_t)\| \leq \sum_{i=0}^{t-1} \|(A + BK)^{t-i-1} BK\| \|\Phi(t)^{-1}\| \|F_R\| \|(\bar{\Upsilon}_A F_R)^i\| \|\underline{L}(0)\|.$$

By lemma 3.5.1 this converges to zero exponentially.

Since the original set Λ_0 is bounded we know there must exist a constant C such that $\|X_t\|, \|Y_t\| \leq C$ for all t . Thus, explicitly stating the dependence on the rate R , we have

$$\Delta(R) \leq \left| \sum_{t=0}^{\infty} [(Y_t + X_t)'(Q + K'TK)(Y_t - X_t) - 2Y_t'K'TKe_t + e_tK'TKe_t] \right|$$

$$\begin{aligned}
&\leq \sum_{t=0}^{\infty} \|Q + K'TK\| \|Y_t + X_t\| \|Y_t - X_t\| + 2\|K'TK\| \|Y_t\| \|e_t\| + \|K'TK\| \|e_t\|^2 \\
&\leq \sum_{t=0}^{\infty} 2C\|Q + K'TK\| \|Y_t - X_t\| + 2C\|K'TK\| \|e_t\| + \|K'TK\| \|e_t\|^2
\end{aligned}$$

Since $\|e_t\|$ and $\|Y_t - X_t\|$ are both converging to zero exponentially we see that this is a convergent series. Furthermore $\lim_{R \rightarrow \infty} \Delta(R) = 0$. \square

Proposition 3.9.2 *We are given system (3.1) with encoder in encoder class two, bounded Λ_0 and quadratic cost. The difference between the optimal cost for a given rate R and the optimal cost under full state observation is bounded for R large enough.*

Proof: By propositions 3.6.3 and 3.6.4 there exists a rate large enough so that the state estimate error decays exponentially. Thus an argument similar to that of proposition 3.9.1 shows that one can bound the loss in terms of R . \square

3.9.3 Covering numbers

The notion of covering number can help us compute the rate for cases where the performance objective is not the sum of a running cost. One control objective might be to ensure that the state is in a given region at a given time. Another might be that a given Lyapunov function V be strictly decreases along trajectories. First we give some definitions and then we treat an example.

Given the dynamics $X_{t+1} = f(X_t, U_t)$, $X_0 \in \Lambda_0$. Define the *objective set* to be

$$\Omega_T = \{(X_0^T, U_0^{T-1}) : X_{t+1} = f(X_t, U_t) \text{ and the objective is met}\}.$$

One example of an objective might be $\|X_t\| \leq \epsilon \forall t \geq T$.

Define for each control sequence U_0^{T-1} the Γ sets $\Gamma_{U_0^{T-1}} = \{X_0 : (X_0^T, U_0^{T-1}) \in \Omega_T\}$.

Definition 3.9.1 *The covering number is minimum number of Γ balls it takes to cover Λ_0 .*

Thus a lower bound on the rate to achieve some objective is

$$R \geq \limsup_T \frac{1}{T} \log (\text{the covering number at time } T).$$

Note that this is the same technique we used in computing the lower bound in propositions 3.3.1 and 3.3.2.

The following example treats the objective of strictly decreasing trajectories.

Example 3.9.1 *Take the scalar system $X_{t+1} = aX_t + bU_t$, $X_0 \in [-L, L]$. Assume that we want the system trajectories to strictly decrease to the origin. Specifically*

$$\Omega_T = \{(X_0^T, U_0^{T-1}) : X_{t+1} = aX_t + bU_t \text{ and } X_{t+1} < X_t \text{ } t = 0, \dots, T-1\}.$$

The corresponding Γ sets are

$$\Gamma_{U_0}^{T-1} = \{X_0 : X_{t+1} = aX_t + bU_t \text{ and } X_{t+1} < X_t \text{ } t = 0, \dots, T-1\}.$$

Take for example Γ_{U_0} and assume with loss of generality that $a > 1$ then

$$\Gamma_{U_0} = \left\{ X_0 : \begin{array}{ll} \frac{-U_0}{a-1} < X_0 < \frac{-U_0}{a+1} & \text{if } U_0 > 0 \\ \frac{-U_0}{a+1} < X_0 < \frac{-U_0}{a-1} & \text{if } U_0 < 0 \end{array} \right\}$$

Covering $[-L, L]$ by such sets induces a logarithmic partition. It takes a countable number of regions to cover this interval. For the general treatment of this problem see [EM].

3.9.4 Discussion

In this section we provided two means for determining the rate requirements for different control objectives. For running cost problems we showed how to upper bound the loss due to the rate constraint by computing the loss due to a certainty equivalent scheme. For more general control objectives we introduced the idea of covering number. We applied this method to a problem introduced by Elia and Mitter in [EM].

3.10 Summary

In this chapter we applied the general formulation of chapter two to the deterministic linear systems control problem with a noiseless digital communication link. We discussed the role of information patterns and system and policy knowledge in this context.

We first provided lower bounds on the rates required to achieve asymptotic observability, asymptotic stability, and controllability. These bounds hold independently of the information pattern chosen.

To compute upper bounds we explicitly described the encoder, decoder, and controller schemes. We characterized two different encoder structures based on whether the encoder observed the control signals or not. Under the added structural assumptions of equi-memory and use of a primitive quantizer we showed that encoders in encoder class one can achieve these lower bounds. Furthermore the information pattern used in encoder class one is not the maximal element in the partial ordering of information patterns. For encoders in encoder class two weaker bounds were provided. These schemes work by keeping track of the uncertainty set of both the error and the state. These relied heavily on the key technical lemma 3.4.2.

We extended these results to systems with multiple sensors. We showed that the problem reduces to a Slepian-Wolf coding problem.

We then analyzed the problem of time sampling under a rate constraint. We showed that for encoder class one the rate is independent of the sampling period. Whereas for encoder class two the rate is dependent on the sampling period.

Finally we addressed the problem of performance. We treated the LQ problem. We also introduced the notion of covering number. This covering number describes the minimum number of control sequences needed to achieve an objective. It is the basis for the lower bound results we proved earlier in the chapter.

Chapter 4

Channel Coding With Feedback

4.1 Introduction

The problem of channel coding goes back to Shannon's original work. [Sha1] In this chapter we examine the problem of channel coding for channels with feedback. The feedback channel coding problem goes back to early work by Shannon, Dobrushin, Wolfowitz, and others. [Sha2], [Dob1], [Wol] See figure 4-1. Due to increased demand for wireless communication and networked systems there is a renewed interest in this problem. Feedback can increase the capacity of a noisy channel, decrease the complexity of the encoder and decoder, and reduce the latency for a given probability of decoding error.

Recently Verdu and Han gave a very general formulation of the channel coding problem without feedback. [VH] Here we generalize that formulation to the case of channels with feedback. To that end we need to introduce the notion of code-functions. These are to be contrasted with codewords. The use of code-functions can be traced back to Shannon's work on transmitter side information. [Sha3] We show that we can convert the channel coding problem with feedback into a new channel coding problem without feedback. In the new channel, though, the channel inputs are now code-functions. We show how to interconnect a code-function distribution to the channel. We discuss the relationship between code-function distributions and channel input distributions. This relationship allows us to convert




Figure 4-1: Channel

an optimization problem over code-function distributions to an optimization over channel input distributions. Along the way we introduce the notion of directed mutual information and argue that it provides the correct measure of capacity. Directed mutual information was introduced by Massey who attributes it to Marko. [Mas], [Mar]

The Verdu-Han result is quite general. But that generality comes at a cost. It is very difficult to solve the capacity optimization problem in their formulation. The usual way to deal with this is to assume some sort of ergodicity in the channel model. Or to be more specific one makes suitable hypothesis so that the channel is information stable. To that end we examine the class of Markov channels. We show that the problem of feedback coding for Markov channels can be cast as a partially observed stochastic control problem. Consequently we can use the tools of dynamic programming to solve the mutual information optimization problem underlying the capacity problem.

In summary there are three main contributions in this chapter. First we give a rather general coding theorem for channels with feedback. Second we argue that directed mutual information is the appropriate notion of mutual information when calculating the capacity of a channel with or without feedback. Third for Markov channels we are able to convert the capacity problem into a stochastic control problem. Thus we can provide a dynamic programming determination for its solution. Furthermore the Markov formulation allows us to treat a large class of channel models with memory and different kinds of feedback in a unified way.

We now summarize each section in this chapter. In section 4.2 we introduce the notion of directed mutual information. In section 4.3 we formulate the channel coding problem with feedback.

In section 4.4 we examine channels with finite alphabets. Here we discuss the relationship between distributions on code-functions and channel input distributions. We provide a Verdu-Han like theorem for channels with feedback. We then show that the error exponents for channels with feedback are larger than the exponents for channels without feedback. In section 4.5 we focus on Markov channels. Here we formulate the equivalent partially observed stochastic control problem.

In section 4.6 we examine Gaussian channels. We provide a general coding theorem and error exponents. The finite alphabet channel and Gaussian channel are used most often in practice. In section 4.7 we list a few examples of these channels. We show how our formulation captures many existing results in the literature. We also point out new examples not existing, at least to the author's knowledge, in the literature.

In section 4.8 we provide a very general data processing inequality for channels with feedback. We also discuss the notion of channel realization. This section will be especially important for the results in chapter 5.

Finally in section 4.9 we conclude.

4.2 Directed Mutual Information

In this section we introduce the directed mutual information and its properties. Massey first defined directed mutual information in [Mas]. We will motivate its use in section 4.3.4. Here we just provide a few of its properties. See section A.1 of the appendix for a review of Polish spaces and stochastic kernels. See section A.3 for a review of divergence, mutual information, conditional mutual information, and the data processing inequality.

Let $\{X_t\}_{t=1}^T$ and $\{Y_t\}_{t=1}^T$ be random processes with each X_t taking values in the Polish space $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ and each Y_t taking values in the Polish space $(\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$. We use “log” to represent logarithm base 2 and “ln” to represent logarithm base e . Let $X^t \triangleq (X_1, \dots, X_t)$.

By the chain rule for mutual information, theorem A.3.4 (c), we know that $I(X^T; Y^T) = \sum_{t=1}^T I(X^T; Y_t | Y^{t-1})$.

Definition 4.2.1 *The directed mutual information is defined as*

$$I(X^T \rightarrow Y^T) \triangleq \sum_{t=1}^T I(X^t; Y_t | Y^{t-1}).$$

(Note that on the right hand side the superscript on X is “ t ” and not “ T ”.)

We now give an alternative characterization of the directed mutual information in terms of some “causal” stochastic kernels.

Definition 4.2.2 *Assume we are given an ordered set of random variables X_1, \dots, X_T . Furthermore assume we are given the following stochastic kernels: $P(dX_1), P(dX_2 | x_1), \dots, P(dX_T | x^{T-1})$. (By theorem A.1.1 the product of these kernels will give us a joint measure P_{X^T} .) Let $I = \{i_1, \dots, i_K\} \subset \{1, \dots, T\}$ and $i_1 < i_2 < \dots < i_K$. Let $I^c = \{1, \dots, T\} \setminus I$. Let $X^I = (X_{i_1}, \dots, X_{i_K})$. Define X^{I^c} similarly. Then we define the directed stochastic kernel to be*

$$\vec{P}_{X^I | X^{I^c}}(dX^I | x^{I^c}) \triangleq \bigotimes_{k=1}^K P(dX_{i_k} | x^{i_k-1}).$$

Recall our general model in definition 2.2.1. Definition 4.2.2 provides a time-ordering on the variables of interest and a “causal” factorization.

For the random variables (X^T, Y^T) representing the input and output of the channel over a time horizon T the natural time-ordering is

$$X_1, Y_1, X_2, Y_2, \dots, X_T, Y_T.$$

Assume we are given the stochastic kernels $\{P(dX_t | x^{t-1}, y^{t-1})\}_{t=1}^T$ and $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$. Then the directed stochastic kernels with respect to this time ordering are:

$$\vec{P}_{X^T | Y^T}(dX^T | y^T) = \bigotimes_{t=1}^T P(dX_t | x^{t-1}, y^{t-1})$$

and

$$\vec{P}_{Y^T | X^T}(dY^T | x^T) = \bigotimes_{t=1}^T P(dY_t | x^t, y^{t-1}).$$

Note that in the latter case Y_t is allowed to depend on x_t as well as (x^{t-1}, y^{t-1}) .

By theorem A.1.1 we can construct the joint measure

$$P_{X^T, Y^T}(dX^T, dY^T) = \vec{P}_{Y^T|X^T}(dY^T | x^T) \otimes \vec{P}_{X^T|Y^T}(dX^T | y^T). \quad (4.1)$$

We can define another measure on $\mathcal{X}^T \times \mathcal{Y}^T$ as follows

$$\vec{P}_{X^T|Y^T} P_{Y^T}(dX^T, dY^T) \triangleq \vec{P}_{X^T|Y^T}(dX^T | y^T) \otimes P_{Y^T}(dY^T). \quad (4.2)$$

The following proposition gives an alternative characterization of the directed mutual information in terms of the divergence between the measures described in equations (4.1) and (4.2).

Proposition 4.2.1

$$I_{P_{X^T, Y^T}}(X^T \rightarrow Y^T) = D(P_{X^T, Y^T} | \vec{P}_{X^T|Y^T} P_{Y^T}).$$

Proof:

$$\begin{aligned} I(X^T \rightarrow Y^T) &= \sum_{t=1}^T I(X^t; Y_t | Y^{t-1}) \\ &= \sum_{t=1}^T D(P_{X^t, Y^t} | P_{X^t|Y^{t-1}} P_{Y_t|Y^{t-1}} P_{Y^{t-1}}) \\ &= D(P_{X^T, Y^T} | \vec{P}_{X^T|Y^T} P_{Y^T}) \end{aligned}$$

Where in the last line we have made repeated use of theorem A.3.4 (b). \square

Corollary 4.2.1 *In the case when \mathcal{X} and \mathcal{Y} are spaces with a countable number of elements we have*

$$I_{P_{X^T, Y^T}}(X^T \rightarrow Y^T) = \sum_{x^T, y^T} P_{X^T, Y^T}(x^T, y^T) \log \frac{P_{X^T|Y^T}(x^T | y^T)}{\vec{P}_{X^T|Y^T}(x^T | y^T)}$$

and

$$I_{P_{X^T, Y^T}}(X^T \rightarrow Y^T) = \sum_{x^T, y^T} P_{X^T, Y^T}(x^T, y^T) \log \frac{\vec{P}_{Y^T|X^T}(y^T | x^T)}{P_{Y^T}(y^T)}$$

Proof: Let $E \subset \mathcal{A}$ and $F \subset \mathcal{B}$ be measurable sets. If $P_A(E), P_B(F) > 0$ then $\frac{P_{A,B}(E,F)}{P_A(E)P_B(F)} = \frac{P_{A|B}(E | F)}{P_A(E)} = \frac{P_{B|A}(F | E)}{P_B(F)}$. The lemma is a straightforward generalization of this fact. Note that $\lim_{z \rightarrow 0^+} -z \log z = 0$. Hence we do not have problems with dividing by zero inside the log. \square

The first characterization shows that the directed mutual information is the ratio between the posterior distribution and a ‘causal’ prior distribution.

Corollary 4.2.2 *In the case when P_{X^T, Y^T} admits a density $p(x^T, y^T)$ we have*

$$I_{P_{X^T, Y^T}}(X^T \rightarrow Y^T) = \int_{\mathcal{X}^T \times \mathcal{Y}^T} p(dx^T, dy^T) \log \frac{p(x^T | y^T)}{\bar{p}(x^T | y^T)}$$

and

$$I_{P_{X^T, Y^T}}(X^T \rightarrow Y^T) = \int_{\mathcal{X}^T \times \mathcal{Y}^T} p(dx^T, dy^T) \log \frac{\bar{p}(y^T | x^T)}{p(y^T)}$$

Proof: Follows from repeated application of the following fact. Let $p(a, b)$ be the density of $P_{A, B}$. Then we have $\frac{p(a, b)}{p(a)p(b)} = \frac{p(a | b)}{p(a)} = \frac{p(b | a)}{p(b)}$. \square

We now compare the directed mutual information to the usual mutual information.

Proposition 4.2.2 $I(X^T; Y^T) \geq I(X^T \rightarrow Y^T)$ with equality if and only if $Y_t - (X^t, Y^{t-1}) - X_{t+1}^T$ forms a Markov chain.

Proof:

$$\begin{aligned} I(X^T; Y^T) - I(X^T \rightarrow Y^T) &= \sum_{t=1}^T I(X^T; Y_t | Y^{t-1}) - I(X^t; Y_t | Y^{t-1}) \\ &= \sum_{t=1}^T I(X_{t+1}^T; Y_t | X^t, Y^{t-1}) \end{aligned}$$

Each term in the sum is nonnegative. By proposition A.3.1 (d) each term equals zero if and only if $Y_t - (X^t, Y^{t-1}) - X_{t+1}^T$ forms a Markov chain for each t . \square

The following corollary states that if the future X 's are not influenced by the past Y 's when conditioned on the past X 's then the directed mutual information equals the regular mutual information. In the context of channels this states that if there is no feedback then the two different mutual information measures are equal.

Corollary 4.2.3 *If $X_{t+1}^T - X^t - Y^t$ is a Markov chain then $I(X^T; Y^T) = I(X^T \rightarrow Y^T)$*

Proof: Clearly if $X_{t+1}^T - X^t - Y^t$ is a Markov chain then $Y_t - (X^t, Y^{t-1}) - X_{t+1}^T$ is a Markov chain. Pearl call this the “weak union” property of conditional independence. [Pea] The result then follows from proposition 4.2.2. \square

4.3 Channel Capacity

In this section we define the channel coding problem for finite alphabet channels. We first define a channel, encoder, and decoder. Next we describe the interconnection between the code-functions and the channel. Then we discuss the role of directed mutual information. Finally we end with two definitions of capacity: operational channel capacity and channel capacity.

Figure 4-2 shows diagrammatically the order of events that we have in mind for the channel coding problem. Briefly, at time 0 we choose a message. This message is assigned a code-function. Then for times $1, \dots, T$ we use the channel sequentially. Note that the channel input symbol at time t is allowed to depend on the past channel output symbols. At time $T + 1$ we decode the message. Note that this is a one-shot scheme. We now give the details.

4.3.1 Channels

Let \mathcal{A}, \mathcal{B} be spaces with a finite number of elements. (We use the symbols \mathcal{A}, \mathcal{B} instead of \mathcal{X}, \mathcal{Y} to emphasize that we are working with finite valued random alphabets.) Let A_t, B_t be measurable random variables taking values in \mathcal{A}, \mathcal{B} respectively. The product spaces \mathcal{A}^T and \mathcal{B}^T represent the input and output spaces for the channel for a time horizon T .

Definition 4.3.1 *A channel is a family of stochastic kernels $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$. (T may be infinity.)*

Note that the specification of the conditioning includes only A^t, B^{t-1} . Thus our channels are nonanticipative channels.

Before we can compute any of the “information” measures of the last section we need to determine the joint measure P_{A^T, B^T} . In general any measure, P_{A^T, B^T} , can be factored as

$$P_{A^T, B^T}(A^T, B^T) = \bigotimes_{t=1}^T P_{A_t | A^{t-1}, B^{t-1}}(A_t | a^{t-1}, b^{t-1}) \otimes P_{B_t | A^t, B^{t-1}}(B_t | a^t, b^{t-1}). \quad (4.3)$$

In order to complete the description of the joint measure we need to specify the kernels $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. These kernels are determined by specifying an encoder. Note that this is consistent with the formulation presented in chapter two.

In summary a channel is a sequence of stochastic kernels. We can “interconnect” them and create a joint measure by specifying another sequence of stochastic kernels. This latter sequence, as we will see, will be determined by the encoder.

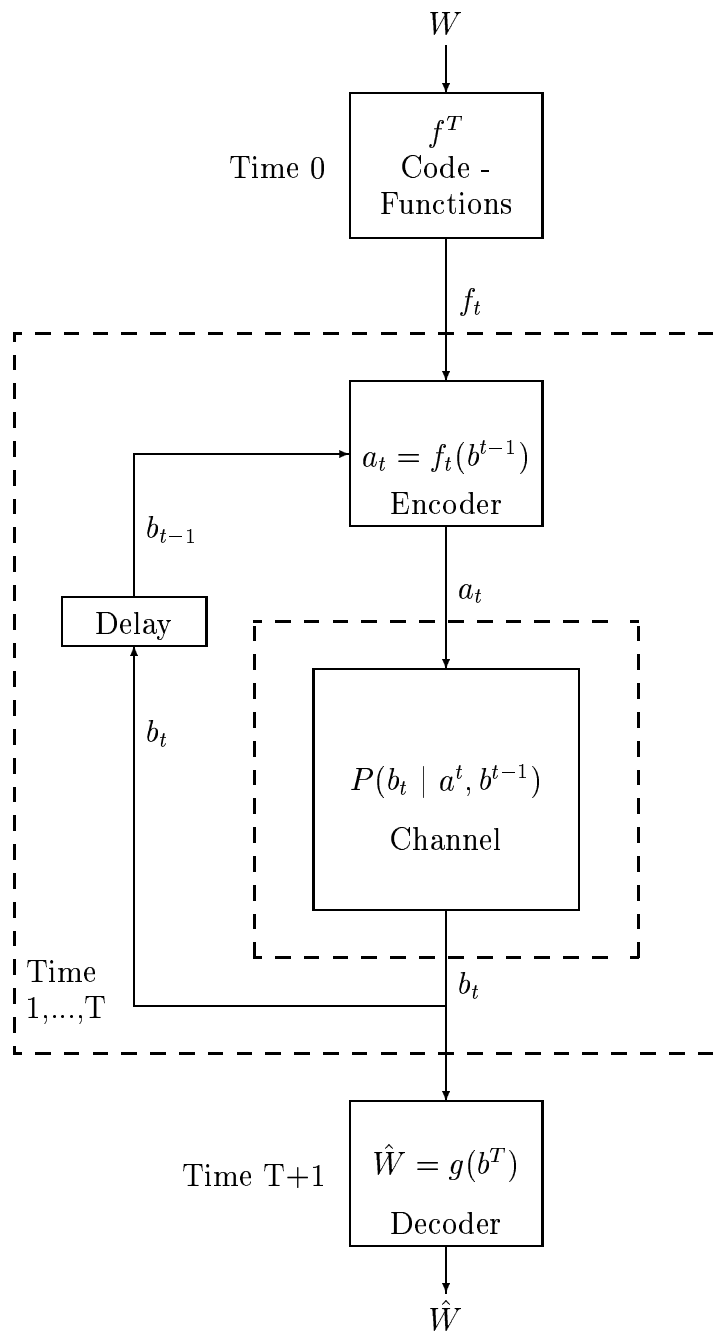


Figure 4-2: Interconnection

4.3.2 Message Set, Encoder, and Decoder

In this subsection we define a message set, an encoder, and a decoder.

Definition 4.3.2 A message set is a set $\mathcal{W} = \{1, \dots, M\}$.

Definition 4.3.3 A channel code-function is a sequence of T deterministic measurable maps $\{f_t\}_{t=1}^T$ such that $f_t : \mathcal{B}^{t-1} \rightarrow \mathcal{A}$ which takes $b^{t-1} \mapsto a_t$. Let $f^T \triangleq \{f_t\}_{t=1}^T$. Denote the set of all code-functions by $\mathcal{F}_T \triangleq \{f^T : f^T \text{ is a code-function}\}$.

Definition 4.3.4 A channel encoder, or channel code, is a set of M channel code-functions. Denote them by $f^T[w]$, $w = 1, \dots, M$.

For message w at time t with channel feedback b^{t-1} the channel encoder outputs $f_t[w](b^{t-1})$.

The following is a special case of the channel encoder when there is no feedback to the encoder:

Definition 4.3.5 A channel codeword is a channel code-function, f^T , where each f_t is independent of b^{t-1} . Thus any codeword f^T can be associated to a vector a^T . The set of all codewords can be represented by the space \mathcal{A}^T .

Definition 4.3.6 A channel encoder without feedback, or channel code without feedback, is a set of M channel codewords. Denote them by $a^T[w]$ $w = 1, \dots, M$.

For message w at time t the channel encoder outputs $a_t[w]$.

Definition 4.3.7 A channel decoder is a map $g : \mathcal{B}^T \rightarrow \mathcal{W}$ taking $b^T \mapsto w$.

Note that the decoder is allowed to wait till it observes all the channel outputs before reconstructing the input message. We will relax this condition in chapter four when we discuss transmission of a process, instead of a message set, over a channel.

4.3.3 Interconnection of Code-Functions to the Channel

Now we are in a position to connect the pieces: channel, code-functions, encoder, and decoder. See figure 4-2.

Let $P_{\mathcal{F}_T}$ be a distribution on \mathcal{F}_T . For example $P_{\mathcal{F}_T}$ may be a distribution that places mass $1/M$ on each of M different code-functions. (I.e. this could be a uniform distribution on the code-functions that make up a particular channel code.) But generally we allow $P_{\mathcal{F}_T}$ to be any distribution on \mathcal{F}_T .

We are given a distribution on code-functions $P_{\mathcal{F}_T}$, a channel $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$, and the deterministic relations: $a_t = f_t(b^{t-1})$. From these we want to construct a new channel that connects the random variable F^T to the random variable B^T . To this end we need to

define the following set of stochastic kernels $\{Q(B_t | f^t, b^{t-1})\}_{t=1}^T$. We use “ Q ” to denote the new joint measure, $Q(F^T, A^T, B^T)$, that we will construct in the course of this subsection.

The kernels $\{Q(B_t | f^t, b^{t-1})\}_{t=1}^T$ need to be defined in such a way so that the following three properties hold:

- (1) *There is no feedback to the code-functions in the new channel.*

The measure on \mathcal{F}_T is chosen at time 0. Thus it cannot causally depend on the B_t 's. Specifically we require that $F_t - F^{t-1} - B^{t-1}$ be a Markov chain under Q . Thus

$$Q(F_t | F^{t-1} = f^{t-1}, B^{t-1} = b^{t-1}) = P(F_t | f^{t-1}) \quad Q(F^{t-1}, B^{t-1}) - a.s.$$

- (2) *The channel input is a function of the past outputs: $a_t = f_t(b^{t-1})$*

We require that $A_t = F_t(B^{t-1}) \quad Q - a.s.$

- (3) *The new channel preserves the properties of the underlying channel $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$.*

Thus we require

$$Q(B_t | F^t = f^t, A^t = a^t, B^{t-1} = b^{t-1}) = P(B_t | a^t, b^{t-1}) \quad Q(F^t, A^t, B^{t-1}) - a.s.$$

Definition 4.3.8 *We are given a code-function distribution, P_{FT} , the relations $a_t = f_t(b^{t-1})$, and a channel, $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$. A measure $Q(F^T, A^T, B^T)$ is said to be consistent with the channel and the code-function distribution if it satisfies the three properties above.*

An obvious question to ask is does such a measure Q on $\mathcal{F}_T \times \mathcal{A}^T \times \mathcal{B}^T$ exist satisfying these requirements? We will show that there exists a unique measure Q .

Lemma 4.3.1 *Given P_{FT} , the channel $\{P(b_t | a^t, b^{t-1})\}_{t=1}^T$, and the relations $a_t = f_t(b^{t-1})$ there exists a unique measure $Q(F^T, A^T, B^T)$ on $\mathcal{F}_T \times \mathcal{A}^T \times \mathcal{B}^T$ satisfying the above three properties. Furthermore the channel from \mathcal{F}_T to \mathcal{B}^T for each $t = 1, \dots, T$ is*

$$Q(B_t | F^t = f^t, B^{t-1} = b^{t-1}) = P(B_t | f^t(b^{t-1}), b^{t-1}) \quad Q(F^t, B^{t-1}) - a.s. \quad (4.4)$$

Proof: Any measure satisfying properties (1) and (2) must be of the form

$$\begin{aligned} Q(F^T, A^T, B^T) &= \left\{ \bigotimes_{t=1}^T Q(B_t | f^t, b^{t-1}) \otimes Q(F_t | f^{t-1}, b^{t-1}) \right\} \otimes Q(A^T | f^T, b^T) \\ &= \left\{ \bigotimes_{t=1}^T Q(B_t | f^t, b^{t-1}) \otimes P(F_t | f^{t-1}) \right\} \otimes \delta_{\{A^T = f^T(b^{T-1})\}} \end{aligned}$$

where $f^T(b^{T-1}) \triangleq (f_1, f_2(b_1), \dots, f_T(b^{T-1}))$. Thus we need only identify the “ $\mathcal{F}_T - \mathcal{B}^T$ ” channel: $\{Q(B_t | f^t, b^{t-1})\}_{t=1}^T$.

$$\begin{aligned} &Q(B_t | F^t = f^t, B^{t-1} = b^{t-1}) \\ &= Q(B_t | F^t = f^t, A^t = f^t(b^{t-1}), B^{t-1} = b^{t-1}) \quad Q(F^t, A^t, B^{t-1}) - a.s. \\ &= P(B_t | f^t(b^{t-1}), b^{t-1}) \quad Q(F^t, A^t, B^{t-1}) - a.s. \end{aligned}$$

where the first equality follows from property (2) and the second equality follows from property (3). Thus the new channel can be uniquely defined as

$$Q(B_t | F^t = f^t, B^{t-1} = b^{t-1}) = P(B_t | f^t(b^{t-1}), b^{t-1}) \quad Q(F^t, B^{t-1}) - a.s.$$

□

Corollary 4.3.1 *A distribution, $P(W)$, on \mathcal{W} , a channel code $\{f^T[w]\}_{w=1}^M$, and the channel $\{P(B_t|a^t, b^{t-1})\}_{t=1}^T$ uniquely define a measure $Q(W, A^T, B^T)$ on $\mathcal{W} \times \mathcal{A}^T \times \mathcal{B}^T$.*

Proof: P_W induces a measure $P(F^T)$ on \mathcal{F}_T . Now apply lemma 4.3.1 to get $Q(F^T, A^T, B^T)$. By the correspondence between the random variables F^T and W we get the measure $Q(W, A^T, B^T)$. □

4.3.4 Directed Mutual Information: The Intuition

It turns out that mutual information is insufficient for computing the channel capacity of a feedback channel. [Mas] To appreciate this fact we need to distinguish between causal independence and probabilistic independence. We discuss the role of directed mutual information. (Recall the definitions in section 4.2.)

As was shown in corollary 4.2.3 if $X_{t+1}^T - X^t - Y^t$ is a Markov chain then $I(X^T; Y^T) = I(X^T \rightarrow Y^T)$. Thus when there is no feedback the two measures of mutual information are the equal.

The traditional mutual information measure is the incorrect one to use when computing the capacity of channels with feedback. We now give some intuition for why this is the case. Assume we are given a joint measure $P(A^T, B^T)$. Then we have via Bayes' law:

$$P(b_t | a^T, b^{t-1}) = \left(\frac{P(a_{t+1}^T | a^t, b^t)}{P(a_{t+1}^T | a^t, b^{t-1})} \right) P(b_t | a^t, b^{t-1}).$$

Thus, in general, $P(b_t | a^T, b^{t-1}) \neq P(b_t | a^t, b^{t-1})$. Even though A_{t+1} occurs after B_t it still has a probabilistic influence on it. To quote Massey, "statistical dependence, unlike causality, has no inherent directivity." Now $I(A^T; B^T)$ depends on $P(B_t | a^T, b^{t-1})$ whereas we would like a measure that is only dependent on $P(B_t | a^t, b^{t-1})$. That measure turns out to be the directed mutual information.

Assume the ordering $A_1, B_1, \dots, A_T, B_T$. Figure 4-3 shows two directed graphs, sometimes called Bayesian networks, for this particular ordering. [Pea] The first graph holds when there is no feedback. The second holds when there is feedback. One can use the graphical modeling principle of d-separation to show that under when there is no feedback, A_{t+1}^T and B_t are independent given (A^t, B^{t-1}) .

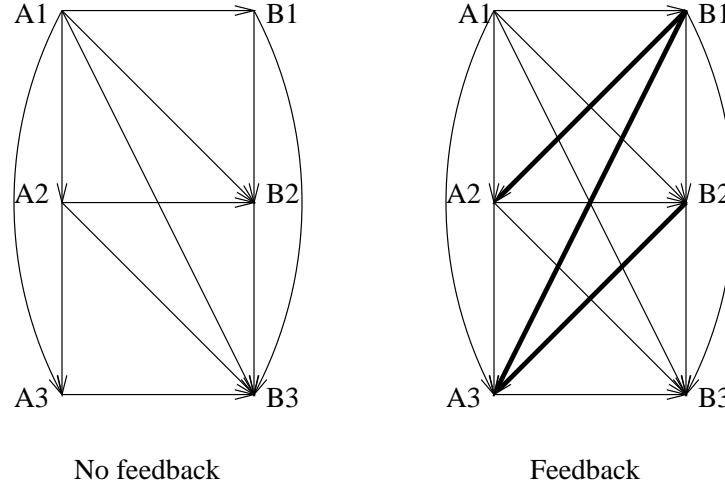


Figure 4-3: Graphical Model Representation

We end this subsection with a simple proposition that extends corollary 4.2.3 to the case of a channel code without feedback.

Proposition 4.3.1 *If we are using a channel encoder without feedback or a channel input distribution without feedback then $I(A^T; B^T) = I(A^T \rightarrow B^T)$.*

Proof: A channel encoder without feedback induces a channel input distribution without feedback. To see this note that since we are using an encoder without feedback $W - A^t - B^t$ is a Markov chain. Now A_{t+1}^T is just a function of W thus $A_{t+1}^T - A^t - B^t$. The proposition then follows from corollary 4.2.3. \square

4.3.5 Operational Channel Capacity

We are now ready to define the operational channel capacity. Take the distribution, P_W , on the message set \mathcal{W} to be the uniform distribution.

Definition 4.3.9 *An (T, M, ϵ) channel code over time horizon T has a channel encoder with M code-functions, a channel decoder g , and an error probability $\frac{1}{M} \sum_{w=1}^M \Pr(w \neq g(b^T)|w) \leq \epsilon$.*

Definition 4.3.10 *An (T, M, ϵ) channel code without feedback over time horizon T has a channel encoder with M codewords, a channel decoder g , and an error probability $\frac{1}{M} \sum_{w=1}^M \Pr(w \neq g(b^T)|w) \leq \epsilon$.*

In what follows the superscript “o” and “nfb” represent the “operational” and “no feedback” respectively.

Definition 4.3.11 R is an ϵ -achievable rate if, for every $\delta > 0$ there exists, for sufficiently large T , an (T, M, ϵ) channel code with rate $\frac{\log M}{T} > R - \delta$. The maximum ϵ -achievable rate is called the ϵ -capacity and denoted C_ϵ^o . The operational channel capacity is defined as the maximal rate that is achievable for all $0 < \epsilon < 1$ and is denoted C^o .

Definition 4.3.12 R is an ϵ -achievable rate without feedback if, for every $\delta > 0$ there exists, for sufficiently large T , an (T, M, ϵ) channel code without feedback with rate $\frac{\log M}{T} > R - \delta$. The maximum ϵ -achievable rate without feedback is called the ϵ -capacity without feedback and is denoted $C_\epsilon^{o, nfb}$. The operational channel capacity without feedback is defined as the maximal rate that is achievable for all $0 < \epsilon < 1$ and is denoted $C^{o, nfb}$.

4.3.6 Channel Capacity

For the case when there is no feedback we know that the operational channel capacity can be characterized by a particular mutual information optimization problem. [Sha1] [VH] In this section we state that optimization problem. We then introduce another optimization problem for the case of channels with feedback. In section 4.4 we show that this is the correct optimization problem to solve.

The following two definitions define distributions on the channel input space.

Definition 4.3.13 A channel input distribution is a sequence of stochastic kernels $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$.

Definition 4.3.14 A channel input distribution without feedback is a channel input distribution with the further condition that for each t the kernel $P(A_t | a^{t-1}, b^{t-1})$ is independent of b^{t-1} . (Specifically $P(A_t | a^{t-1}, b^{t-1}) = P(A_t | a^{t-1}, \tilde{b}^{t-1}) \quad \forall b^{t-1}, \tilde{b}^{t-1}$.)

When computing the capacity of a channel it will turn out that we are most interested in the convergence properties of the random variables $\frac{1}{T} \log \frac{P_{A^T, B^T}(A^T, B^T)}{\tilde{P}_{A^T | B^T} P_{B^T}(A^T, B^T)}$. If there are reasonable regularity properties, like information stability, then these random variables will converge in probability to a deterministic limit. In the absence of any such structure we are forced to follow Verdu and Han's lead and define the following "floor" and "ceiling" limits. [VH]

Definition 4.3.15 The limsup in probability of a sequence of random variables $\{X_t\}$ is defined as the smallest extended real number α such that $\forall \epsilon > 0$

$$\lim_{t \rightarrow \infty} \Pr[X_t \geq \alpha + \epsilon] = 0.$$

Denote this number α by \limsup in prob X_t .

Definition 4.3.16 The liminf in probability of a sequence of random variables $\{X_t\}$ is defined as the largest extended real number α such that $\forall \epsilon > 0$

$$\lim_{t \rightarrow \infty} \Pr[X_t \leq \alpha - \epsilon] = 0.$$

Denote this number α by \liminf in prob X_t .

Definition 4.3.17 Let $\vec{i}(a^T; b^T) \triangleq \log \frac{P_{A^T, B^T}(a^T, b^T)}{\bar{P}_{A^T|B^T} P_{B^T}(a^T, b^T)}$.

Definition 4.3.18 For a sequence of joint measures $\{P_{A^T, B^T}\}_{T=1}^\infty$ let

$$\underline{I}(A \rightarrow B) \triangleq \liminf_{\text{in prob}} \frac{1}{T} \vec{i}(A^T; B^T) \quad \text{and} \quad \bar{I}(A \rightarrow B) \triangleq \limsup_{\text{in prob}} \frac{1}{T} \vec{i}(A^T; B^T).$$

Lemma 4.3.2 For any sequence of joint measures $\{P_{A^T, B^T}\}_{T=1}^\infty$ we have

$$\underline{I}(A \rightarrow B) \leq \liminf_{T \rightarrow \infty} \frac{1}{T} I(A^T \rightarrow B^T) \leq \limsup_{T \rightarrow \infty} \frac{1}{T} I(A^T \rightarrow B^T) \leq \bar{I}(A \rightarrow B)$$

Proof: See the end of section A.3 in the appendix. \square

If $\underline{I}(A \rightarrow B) = \bar{I}(A \rightarrow B)$ then we say that the process $\{P_{A^T, B^T}\}_{t=1}^\infty$ is *information stable*. Note that this is a generalization of Dobrushin's definition of information stability. [Dob2] The preceding lemma states that if the process is information stable, $\underline{I}(A \rightarrow B) = \bar{I}(A \rightarrow B)$, then $\lim_{T \rightarrow \infty} \frac{1}{T} I(A^T \rightarrow B^T)$ exists and equals $\underline{I}(A \rightarrow B)$. Thus, in this case, we can work directly with $I(A^T \rightarrow B^T)$.

Definition 4.3.19 Let $\mathcal{S}_T = \{\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T\}$ be the set of all channel input distributions.

Definition 4.3.20 Let $\mathcal{S}_T^{nfb} = \{\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T\}$ be the set of all channel input distributions without feedback.

We are now ready to define the mutual information optimization problems. Recall a channel input distribution and a channel define the joint measure.

Definition 4.3.21 For finite T let

$$C_T = \sup_{S \in \mathcal{S}_T} \frac{1}{T} I(A^T \rightarrow B^T)$$

and

$$C_T^{nfb} = \sup_{S \in \mathcal{S}_T^{nfb}} \frac{1}{T} I(A^T \rightarrow B^T) = \sup_{S \in \mathcal{S}_T^{nfb}} \frac{1}{T} I(A^T; B^T).$$

For the infinite horizon case let

$$C = \sup_{S \in \mathcal{S}_\infty} \underline{I}(A \rightarrow B)$$

and

$$C^{nfb} = \sup_{S \in \mathcal{S}_\infty^{nfb}} \underline{I}(A \rightarrow B) = \sup_{S \in \mathcal{S}_\infty^{nfb}} \underline{I}(A; B).$$

Verdu and Han proved the following theorem.

Theorem 4.3.1 *For channels without feedback $C^{o,nfb} = C^{nfb}$.*

Proof: See theorems 1 and 5 of [VH]. \square

In a certain sense we already have the solution to the coding problem for channels with feedback. Specifically lemma 4.3.1 tells us that the feedback channel problem is equivalent to a new channel coding problem without feedback. This new channel is from \mathcal{F}_T to \mathcal{B}^T and has channel kernels defined by equation 4.4. Thus we can directly apply theorem 4.3.1 to this new channel. It turns out, though, that this is a very complicated problem to solve. We would have to optimize the mutual information over distributions on code functions. Note that in definition 4.3.21 the optimization is over channel input distributions. We will show that we can simplify the optimization problem over distributions on code-functions to that of an optimization problem over channel input distributions. Specifically we don't have to work on $\mathcal{F}_T \times \mathcal{B}^T$ space but instead can work on the original $\mathcal{A}^T \times \mathcal{B}^T$ space. This next section proves the analogous result to theorem 4.3.1 for channels with feedback: $C^O = C$.

To see that this might be possible we show that the directed mutual information is the same for both channels.

Proposition 4.3.2 *Assume we are given a joint measure $Q(F^T, A^T, B^T)$ consistent with a channel and a code-function distribution. Then $I(F^T; B^T) = I(A^T \rightarrow B^T)$.*

Proof:

$$\begin{aligned} I(F^T; B^T) &= I(F^T \rightarrow B^T) \quad \text{property (1) of consistency} \\ &= \sum_{t=1}^T I(F^t; B_t | B^{t-1}) \\ &= \sum_{t=1}^T I(F^t, A^t; B_t | B^{t-1}) \quad \text{property (2) of consistency} \\ &= \sum_{t=1}^T I(A^t; B_t | B^{t-1}) \quad \text{property (3) of consistency} \\ &= I(A^T \rightarrow B^T) \end{aligned}$$

\square

In fact, as we will see in section 4.4.1, this proposition is the basis for the converse and direct parts of the feedback channel coding problem.

4.4 Coding Theorem for Finite Alphabet Channels

In this section we treat the coding problem for finite alphabet channels with feedback. There are three main parts to this section. First we provide the necessary technical lemmas to show the relationship between code-function distributions and channel input distributions. Second we prove the feedback channel coding theorem. Third we compute error exponents.

Theorem 4.4.1 *For channels with feedback $C^0 = C$.*

Proof: This will be proved in section 4.4.1 and 4.4.2. \square

Before starting with the proof we give a high-level summary of the issues involved. The converse part is rather straightforward. For any channel code and channel we know by lemma 4.3.1 that there is a unique consistent measure $Q(F^T, A^T, B^T)$. From this measure we can compute the induced channel input distribution $\{Q(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. Now $\{Q(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T \in \mathcal{S}_T$ but it need not be the supremizing channel input distribution. Thus the directed mutual information under the induced channel input distribution may be less than the directed mutual information under the supremizing channel input distribution. This is how we will show $C^0 \leq C$.

The direct part is the interesting part of the theorem 4.4.1. Here we take the optimizing channel input distribution $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$ and construct a distribution on code-functions P_{F^T} . We then prove the direct part of the coding theorem for the channel from \mathcal{F}_T to \mathcal{B}^T by the usual techniques for channels without feedback. By a suitable construction of P_{F^T} it can be shown that the induced channel input distribution equals the original channel input distribution. Thus a generalization of proposition 4.3.2 shows that the directed mutual information measures are the same for both the “ $\mathcal{F}_T - \mathcal{B}^T$ ” channel and the “ $\mathcal{A}^T - \mathcal{B}^T$ ” channel. This is how we will show all rates less than C are achievable.

4.4.1 Main Technical Lemmas

In this subsection we present the main technical lemmas we need for the feedback channel coding theorem. Given P_{F^T} and $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$ we know by lemma 4.3.1 that there exists a unique consistent measure $Q(F^T, A^T, B^T)$. Furthermore the new channel is defined as $Q(B_t | f^t, b^{t-1}) = P(B_t | f^t(b^{t-1}), b^{t-1})$. (See equation (4.4).)

Equivalence of Ratios

The following lemma allows us to generalize proposition 4.3.2 to the infinite horizon case.

Lemma 4.4.1 *We are given $P(F^T)$ and $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$ and the consistent joint measure $Q(F^T, A^T, B^T)$. Then with Q -probability one we have*

$$\frac{Q_{F^T, B^T}(F^T, B^T)}{Q_{F^T} Q_{B^T}(F^T, B^T)} = \frac{Q_{A^T, B^T}(A^T, B^T)}{\vec{Q}_{A^T|B^T} Q_{B^T}(A^T, B^T)}$$

Proof: For every (f^T, a^T, b^T) such that $a_t = f_t(b^{t-1})$ for $t = 1, \dots, T$ we have the following

$$\begin{aligned}
\frac{Q_{F^T, B^T}(f^T, b^T)}{Q_{F^T} Q_{B^T}(f^T, b^T)} &= \frac{Q_{B^T|F^T}(b^T|f^T)}{Q_{B^T}(b^T)} \\
&= \frac{\prod_{t=1}^T Q_{B_t|B^{t-1}, F^t}(b_t|b^{t-1}, f^t)}{Q_{B^T}(b^T)} \\
&= \frac{\prod_{t=1}^T Q_{B_t|B^{t-1}, A^t}(b_t|a^t, b^{t-1})}{Q_{B^T}(b^T)} \\
&= \frac{\vec{Q}_{B^T|A^T}(b^T|a^T) \vec{Q}_{A^T|B^T}(a^T|b^T)}{Q_{B^T}(b^T) \vec{Q}_{A^T|B^T}(a^T|b^T)} \\
&= \frac{Q_{A^T, B^T}(a^T, b^T)}{\vec{Q}_{A^T|B^T} Q_{B^T}(a^T, b^T)}
\end{aligned}$$

□

Corollary 4.4.1 $\underline{I}(F^T; B^T) = \underline{I}(A^T \rightarrow B^T)$.

Proof: Follows from lemma 4.4.1 and the definition of \underline{I} . □

Induced Channel Input Distribution

We now discuss the induced channel input distribution. This is the channel input distribution induced by a given code-function distribution. First some definitions.

Definition 4.4.1 Define the $\text{graph}(f_t) \triangleq \{(b^{t-1}, a_t) : f_t(b^{t-1}) = a_t\} \subset \mathcal{B}^{t-1} \times \mathcal{A}$.

Definition 4.4.2 Let

$$\Gamma_t(b^{t-1}, a_t) \triangleq \{f_t : (b^{t-1}, a_t) \in \text{graph}(f_t)\}$$

and

$$\Gamma^t(b^{t-1}, a^t) \triangleq \{f^t : (b^{j-1}, a_j) \in \text{graph}(f_j) \quad j = 1, \dots, t\}.$$

The following lemma characterizes the induced channel input distribution.

Lemma 4.4.2 We are given P_{F^T} , a channel $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$, and a consistent joint measure by $Q(F^T, A^T, B^T)$. The induced channel distribution can be determined as follows. For every (a^t, b^{t-1}) such that $\Gamma^{t-1}(b^{t-2}, a^{t-1}) \neq \emptyset$ we have

$$Q(a_t | a^{t-1}, b^{t-1}) = P_{F_t|F^{t-1}} \left(\Gamma_t(b^{t-1}, a_t) | \Gamma^{t-1}(b^{t-2}, a^{t-1}) \right).$$




Figure 4-4: Channel Input Distributions

Proof: Note that (a^{t-1}, b^{t-1}) uniquely specifies $(\Gamma^{t-1}(b^{t-2}, a^{t-1}), b^{t-1})$ and vice-versa. Thus it must be the case that

$$Q(a_t | a^{t-1}, b^{t-1}) = Q(a_t | \Gamma^{t-1}(b^{t-2}, a^{t-1}), b^{t-1}).$$

Note that (b^{t-1}, a_t) uniquely specifies $(\Gamma_t(b^{t-1}, a_t), b^{t-1})$ and vice-versa. Thus it must be the case that

$$Q(a_t | \Gamma^{t-1}(b^{t-2}, a^{t-1}), b^{t-1}) = Q(\Gamma_t(b^{t-1}, a_t) | \Gamma^{t-1}(b^{t-2}, a^{t-1}), b^{t-1}).$$

Now by property one of consistency we have

$$Q(\Gamma_t(b^{t-1}, a_t) | \Gamma^{t-1}(b^{t-2}, a^{t-1}), b^{t-1}) = P(\Gamma_t(b^{t-1}, a_t) | \Gamma^{t-1}(b^{t-2}, a^{t-1})).$$

Combining the above equalities we get

$$Q(a_t | a^{t-1}, b^{t-1}) = P(\Gamma_t(b^{t-1}, a_t) | \Gamma^{t-1}(b^{t-2}, a^{t-1})).$$

□

Note that lemma 4.4.2 implies that the induced channel input distribution depends only on the code-function distribution. It does not depend on the particular channel given. Also note that many different code-function distributions may induce the same channel input distribution.

Conditions for the Induced Channel Input Distribution to Equal the Original Channel Input Distribution

So far we have shown how a code-function distribution induces a channel input distribution. As we discussed in the introduction to this section, we would like to start with a channel input distribution and construct a code-function distribution such that the resulting induced channel input distribution equals the original channel input distribution. This is shown pictorially in figure 4-4. In the figure we want the two channel input distributions to be the same. The first arrow represents the construction of the code-function distribution. And the second arrow is described by the result in lemma 4.4.2. Corollary 4.4.1 states

that $\underline{I}_Q(F^T; B^T) = \underline{I}_Q(A^T \rightarrow B^T)$. If we show that the induced channel input distribution equals the original channel input distribution then we have $\underline{I}_Q(A^T \rightarrow B^T) = \underline{I}_P(A^T \rightarrow B^T)$. Consequently $\underline{I}_Q(F^T; B^T) = \underline{I}_P(A^T \rightarrow B^T)$.

Definition 4.4.3 We call a code-function distribution $P(F^T)$ good with respect to the channel input distribution $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$ if the following holds for all (b^{t-1}, a_t)

$$P_{F_t | F^{t-1}}(\Gamma_t(b^{t-1}, a_t) | f^{t-1}) = P(a_t | f^{t-1}(b^{t-2}), b^{t-1}).$$

The first question to ask is does such a good code-function distribution exist? And if so is it unique? We now show that there do exist good code-function distributions but they are not unique.

Lemma 4.4.3 There exists a code-function distribution $P(F^T)$ good with respect to the channel input distribution $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$.

Proof: For all f^t define $P(f_t | f^{t-1})$ as follows

$$P(f_t | f^{t-1}) = \prod_{(b^{t-1}, a_t) \in \text{graph}(f_t)} P(a_t | f_1, \dots, f_{t-1}(b^{t-2}), b^{t-1}).$$

It is a tedious but straightforward exercise to show that $\sum_{f_t} P(f_t | f^{t-1}) = 1$ for all $t = 1, \dots, T$. Thus $P(F^T)$ is a code-function distribution. Clearly this construction is good with respect to the channel input distribution $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. \square

A function f_t is defined by its graph. In the above construction we have enforced independence. In the case where $(b^{t-1}, a_t) \neq (\tilde{b}^{t-1}, \tilde{a}_t)$ we have

$$\begin{aligned} & P\left(\left\{f_t : (b^{t-1}, a_t), (\tilde{b}^{t-1}, \tilde{a}_t) \in \text{graph}(f_t)\right\} | f^{t-1}\right) \\ &= P\left(\left\{f_t : (b^{t-1}, a_t) \in \text{graph}(f_t)\right\} | f^{t-1}\right) \times P\left(\left\{f_t : (\tilde{b}^{t-1}, \tilde{a}_t) \in \text{graph}(f_t)\right\} | f^{t-1}\right) \end{aligned}$$

We do not need to assume this independence. In fact there are many good code-function distributions without this independence. Some are simpler than others. A particularly simple one will be used in section 4.6 where we deal with Gaussian channels.

Now we show that we can achieve the program outlined in figure 4-4.

Lemma 4.4.4 We are given a channel input distribution $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. The induced channel input distribution equals the original channel input distribution if and only if the code-function distribution $P(F^T)$ is good with respect to the original channel input distribution.

Proof: By lemma 4.4.2 we have

$$Q(a_t | a^{t-1}, b^{t-1}) = P\left(\Gamma_t(b^{t-1}, a_t) | \Gamma^{t-1}(b^{t-2}, a^{t-1})\right)$$




Figure 4-5: Code-Function Distributions

Now we want

$$P\left(\Gamma_t(b^{t-1}, a_t) \mid \Gamma^{t-1}(b^{t-2}, a^{t-1})\right) = P(a_t \mid a^{t-1}, b^{t-1}).$$

But this is precisely the definition of P_{F^T} being good with respect to the original channel input distribution. \square

We have just shown that any code-function distribution good with respect to the channel input distribution satisfies figure 4-4. We can also ask if the outline in figure 4-5 is possible. That is we start with a code-function distribution $P(F^T)$ and compute its induced channel input distribution $\{Q(A_t \mid a^{t-1}, b^{t-1})\}_{t=1}^T$. Then we compute a code-function distribution $Q(F^T)$ good with respect to $\{Q(A_t \mid a^{t-1}, b^{t-1})\}_{t=1}^T$. Is it possible for $Q(F^T) = P(F^T)$? As stated in the remarks after lemma 4.4.3 there are many ways to go from $\{Q(A_t \mid a^{t-1}, b^{t-1})\}_{t=1}^T$ to $Q(F^T)$. Thus the answer is no in general. But we can show the following:

Lemma 4.4.5 *We are given a code-function distribution $P(F^T)$. If $Q(F^T)$ is a code-function distribution good with respect to the induced channel input distribution $\{Q(A_t \mid a^{t-1}, b^{t-1})\}_{t=1}^T$ then for all (b^{t-1}, a_t)*

$$Q(\Gamma_t(b^{t-1}, a_t) \mid f^{t-1}) = P(\Gamma_t(b^{t-1}, a_t) \mid f^{t-1}).$$

Proof:

$$\begin{aligned} Q(\Gamma_t(b^{t-1}, a_t) \mid f^{t-1}) &= Q(a_t \mid f^{t-1}(b^{t-2}), b^{t-1}) \\ &= P\left(\Gamma_t(b^{t-1}, a_t) \mid \Gamma^{t-1}(f^{t-1}(b^{t-2}), b^{t-1})\right) \\ &= P\left(\Gamma_t(b^{t-1}, a_t) \mid f^{t-1}\right) \end{aligned}$$

where the first equality follows from definition 4.4.3 and the second equality follows from lemma 4.4.2. \square

The space of good code-function distributions can be quite large. Thus another reason for reducing the problem to an optimization of channel input distributions instead of one over code-function distributions is that we reduce the number of extrema.

4.4.2 Feedback Channel Coding Theorem

Now we can prove the feedback channel coding theorem 4.4.1. We first prove the converse part. Then we prove the direct part.

Converse Theorem

Choose a (T, M, ϵ) channel code $\{f^T[w]\}_{w=1}^M$. Place a prior probability $\frac{1}{M}$ on each code-function $f^T(w)$. By lemma 3.3.1 this defines a consistent measure $Q(W, A^T, B^T)$. The following is a generalization of the Verdú-Han converse presented in [VH].

Lemma 4.4.6 *Every (T, M, ϵ) channel code satisfies*

$$\epsilon \geq Q_{A^T, B^T} \left(\frac{1}{T} \log \frac{Q_{A^T, B^T}(A^T, B^T)}{Q_{A^T|B^T} Q_{B^T}(A^T, B^T)} \leq \frac{1}{T} \log M - \gamma \right) - 2^{-\gamma T} \quad \forall \gamma > 0$$

Proof: Choose a $\gamma > 0$. Let $D_w \subset \mathcal{B}^T$ be the decoding region for message w . The only restriction we place on the decoding regions is that they do not intersect: $D_w \cap D_{\tilde{w}} = \emptyset \quad \forall \tilde{w} \neq w$. (This is always true when using a channel decoder: $D_w = \{w : g(b^T) = w\}$.)

The probability of error is

$$\sum_{w=1}^M \Pr(w, D_w^c).$$

Define

$$\Omega = \{(w, b^T) : \frac{1}{T} \log \frac{Q_{W, B^T}(w, b^T)}{Q_W Q_{B^T}(w, b^T)} \leq \frac{1}{T} \log M - \gamma\}$$

and

$$\Omega_w = \{b^T : \frac{1}{T} \log \frac{Q_{W, B^T}(w, b^T)}{Q_W Q_{B^T}(w, b^T)} \leq \frac{1}{T} \log M - \gamma\}.$$

Note that

$$\begin{aligned} \Omega_w &= (\Omega_w \cap D_w) \cup (\Omega_w \cap D_w^c) \\ &\subseteq (\Omega_w \cap D_w) \cup D_w^c. \end{aligned}$$

Thus

$$\begin{aligned} \Pr(w, D_w^c) &\geq \Pr(w, \Omega_w) - \Pr(w, \Omega_w \cap D_w) \\ &= \Pr(w, \Omega_w) - \sum_{b^T \in \Omega_w \cap D_w} Q_{W, B^T}(w, b^T) \\ &= \Pr(w, \Omega_w) - \sum_{b^T \in \Omega_w \cap D_w} \frac{Q_{W, B^T}(w, b^T)}{Q_W Q_{B^T}(w, b^T)} Q_W Q_{B^T}(w, b^T) \\ &\geq \Pr(w, \Omega_w) - \sum_{b^T \in \Omega_w \cap D_w} M 2^{-T\gamma} Q_W Q_{B^T}(w, b^T) \\ &\geq \Pr(w, \Omega_w) - \sum_{b^T \in D_w} M 2^{-T\gamma} Q_W Q_{B^T}(w, b^T) \end{aligned}$$

$$\begin{aligned}
&\geq \Pr(w, \Omega_w) - M2^{-T\gamma} Q_W(w) \Pr(D_w) \\
&= \Pr(w, \Omega_w) - 2^{-T\gamma} \Pr(D_w)
\end{aligned}$$

Thus

$$\begin{aligned}
\sum_{w=1}^M \Pr(w, D_w^c) &\geq \sum_{w=1}^M \left(\Pr(w, \Omega_w) - 2^{-T\gamma} \Pr(D_w) \right) \\
&\geq \Pr(\Omega) - 2^{-T\gamma} \sum_{w=1}^M \Pr(D_w) \\
&\geq \Pr(\Omega) - 2^{-T\gamma}
\end{aligned}$$

By lemma 4.4.1 we know that the following holds for all (f^T, a^T, b^T) such that $a_t = f_t(b^{t-1})$

$$\frac{Q_{W, B^T}(w, b^T)}{Q_W P_{B^T}(w, b^T)} = \frac{Q_{A^T, B^T}(a^T, b^T)}{\vec{Q}_{A^T|B^T} Q_{B^T}(a^T, b^T)}$$

From this we can conclude

$$\epsilon \geq Q_{A^T, B^T} \left(\frac{1}{T} \log \frac{Q_{A^T, B^T}(A^T, B^T)}{\vec{Q}_{A^T|B^T} Q_{B^T}(A^T, B^T)} \leq \frac{1}{T} \log M - \gamma \right) - 2^{-\gamma T} \quad \forall \gamma > 0$$

□

Note that in the proof of lemma 4.4.6 the only property of the decoder we used is the restriction that the decoding regions not overlap. Thus the lemma holds independently of the decoder that one uses. Thus the lemma is quite general.

A weaker form of the converse lemma can be found in [Mas]. Specifically he shows $I(W; B^T) \leq I(A^T \rightarrow B^T)$. This combined with Fano's inequality, [Gal], gives us

$$H(\epsilon) + \epsilon \log M \geq H(W) - I(W; B^T)$$

where $H(\epsilon)$ is the entropy of a random coin with bias ϵ . This implies

$$\frac{H(\epsilon)}{\log M} + \epsilon \geq 1 - \frac{\frac{1}{T} I(A^T \rightarrow B^T)}{\frac{1}{T} \log M} \geq 1 - \frac{C_T}{\frac{1}{T} \log M}$$

Thus the rate $R = \frac{1}{T} \log M \leq C_T$ is a necessary condition for the error to go to zero. We state the more general converse theorem now.

Theorem 4.4.2 *The channel capacity $C^o \leq C$.*

Proof: Assume there exists a sequence of (T, M_T, ϵ_T) channel codes with $\epsilon_T \rightarrow 0$ as $T \rightarrow \infty$. Assume towards a contradiction that $\lim_{T \rightarrow \infty} \frac{1}{T} \log M_T > C + 2\gamma$. By the previous lemma

we know

$$\begin{aligned}\epsilon_T &\geq Q_{A^T, B^T} \left(\frac{1}{T} \log \frac{Q_{A^T, B^T}(A^T, B^T)}{\bar{Q}_{A^T|B^T} Q_{B^T}(A^T, B^T)} \leq \frac{1}{T} \log M_T - \gamma \right) - 2^{\gamma T} \\ &\geq Q_{A^T, B^T} \left(\frac{1}{T} \log \frac{Q_{A^T, B^T}(A^T, B^T)}{\bar{Q}_{A^T|B^T} Q_{B^T}(A^T, B^T)} \leq C + \gamma \right) - 2^{\gamma T}\end{aligned}$$

But by the definition of C and for T large enough the mass below $C + \gamma$ has nonzero probability. Therefore the right hand side in the inequality is greater than zero. Thus contradicting $\epsilon_T \rightarrow 0$. \square

The converse for the feedback channel coding theorem is really the easy part. This is due to the fact that any channel code induces a measure on the $\mathcal{A}^T \times \mathcal{B}^T$ space. The direct theorem is harder to prove. We want to use a random coding argument. To that end we need to first generate a channel code distribution from the channel input distribution. Then we need to show that at least one of these randomly drawn codes achieves capacity. We show that this can be done in the next section.

Direct Theorem

We will prove the direct theorem via a random coding argument. The following is a generalization of Feinstein's lemma [Fei].

Lemma 4.4.7 *Fix a time T and $0 < \epsilon < 1$. Fix a channel $\{P(B_t | b^{t-1}, a^t)\}_{t=1}^T$. Then for all $\gamma > 0$ and channel input distributions $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$, there exists an (T, M, ϵ) channel code for the channel that satisfies*

$$\epsilon \leq P_{A^T, B^T} \left(\frac{1}{T} \log \frac{P_{A^T, B^T}(A^T, B^T)}{\bar{P}_{A^T|B^T} P_{B^T}(A^T, B^T)} \leq \frac{1}{T} \log M + \gamma \right) + 2^{-\gamma T}.$$

Proof: Let $P(F^T)$ be any code-function distribution good with respect to the channel input distribution $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. Let $Q(F^T, A^T, B^T)$ be the consistent joint measure induced by $P(F^T)$ and the channel. A channel code is selected at random by drawing M code-functions from $P(F^T)$. Choose a $\gamma > 0$.

Define

$$\Omega = \{(f^T, b^T) : \frac{1}{T} \log \frac{Q_{F^T, B^T}(f^T, b^T)}{Q_{F^T} Q_{B^T}(f^T, b^T)} \leq \frac{1}{T} \log M + \gamma\}.$$

Recall the inequality $1 - (1 - x)^k \leq kx \quad \forall x \in (0, 1)$.

Then via the usual random coding arguments we will compute the average codebook error averaged over all codebooks. Let $\Pr(\text{error} | f^T, b^T)$ be the probability of error if we

have chosen one codeword to be f^T and the received channel output is b^T . Then

$$\begin{aligned}
& E_{\text{codebook}(T,M)}(\text{Error}) \\
&= \sum_{\Omega} \Pr(\text{error}|f^T, b^T) Q_{F^T, B^T}(f^T, b^T) + \sum_{\Omega^c} \Pr(\text{error}|f^T, b^T) Q_{F^T, B^T}(f^T, b^T) \\
&\leq Q_{F^T, B^T}(\Omega) + \sum_{\Omega^c} \Pr(\text{error}|f^T, b^T) Q_{F^T, B^T}(f^T, b^T) \\
&= Q(\Omega) + \sum_{\Omega^c} Q_{F^T, B^T}(f^T, b^T) \left[1 - \left(1 - Q_{F^T} \left\{ \tilde{f}^T : \frac{Q_{B^T|\tilde{F}^T}(b^T|\tilde{f}^T)}{Q_{B^T|F^T}(b^T|f^T)} > 1 \right\} \right)^{M-1} \right] \\
&\leq Q(\Omega) + \sum_{\Omega^c} Q_{F^T, B^T}(f^T, b^T) (M-1) Q_{F^T} \left\{ \tilde{f}^T : \frac{Q_{B^T|\tilde{F}^T}(b^T|\tilde{f}^T)}{Q_{B^T|F^T}(b^T|f^T)} > 1 \right\} \\
&= Q(\Omega) + (M-1) \sum_{\Omega^c} Q_{F^T, B^T}(f^T, b^T) \left(\sum_{\tilde{f}^T} \frac{Q_{B^T|\tilde{F}^T}(b^T|\tilde{f}^T)}{Q_{B^T|F^T}(b^T|f^T)} Q_{F^T}(\tilde{f}^T) \right) \\
&= Q(\Omega) + (M-1) \sum_{\Omega^c} Q_{F^T, B^T}(f^T, b^T) \left(\frac{Q_{B^T}(b^T)}{Q_{B^T|F^T}(b^T|f^T)} \right) \\
&\leq Q(\Omega) + (M-1) \sum_{\Omega^c} \frac{1}{M} 2^{-\gamma T} Q_{F^T, B^T}(f^T, b^T) \\
&\leq Q(\Omega) + 2^{-\gamma T}
\end{aligned}$$

By lemma 4.4.1 the following equality holds for all (f^T, a^T, b^T) such that $a_t = f_t(b^{t-1})$

$$\frac{Q_{F^T, B^T}(f^T, b^T)}{Q_{F^T} P_{B^T}(f^T, b^T)} = \frac{Q_{A^T, B^T}(a^T, b^T)}{\vec{Q}_{A^T|B^T} Q_{B^T}(a^T, b^T)}$$

Thus we have shown

$$E_{\text{codebooks}(T,M)}(\text{Error}) \leq Q_{A^T, B^T} \left(\frac{1}{T} \log \frac{Q_{A^T, B^T}(A^T, B^T)}{\vec{Q}_{A^T|B^T} Q_{B^T}(A^T, B^T)} \leq \frac{1}{T} \log M + \gamma \right) + 2^{-\gamma T}.$$

By lemma 4.4.4 and consistency we have

$$Q(A^T, B^T) = \vec{Q}(A^T | b^T) \otimes \vec{P}(B^T | a^T) = \vec{P}(A^T | b^T) \otimes \vec{P}(B^T | a^T) = P(A^T, B^T)$$

thus

$$E_{\text{codebooks}(T,M)}(\text{Error}) \leq P_{A^T, B^T} \left(\frac{1}{T} \log \frac{P_{A^T, B^T}(A^T, B^T)}{\vec{P}_{A^T|B^T} P_{B^T}(A^T, B^T)} \leq \frac{1}{T} \log M + \gamma \right) + 2^{-\gamma T}.$$

Since the error averaged over codebooks is less than the right hand side we know there must exist at least one code with error

$$\epsilon \leq P_{A^T, B^T} \left(\frac{1}{T} \log \frac{P_{A^T, B^T}(A^T, B^T)}{\bar{P}_{A^T|B^T} P_{B^T}(A^T, B^T)} \leq \frac{1}{T} \log M + \gamma \right) + 2^{-\gamma T}.$$

Thus the lemma holds. \square

The following follows [VH].

Theorem 4.4.3 *All rates less than C are achievable.*

Proof: Fix $\epsilon > 0$. We will show that C is an ϵ -achievable rate by demonstrating for every $\delta > 0$ and large enough T that there exists a $(T, M, 2^{-\frac{T\delta}{4}} + \frac{\epsilon}{2})$ code with rate

$$C - \delta \leq \frac{\log M}{T} \leq C - \frac{\delta}{2}.$$

If in the previous lemma we choose $\gamma = \frac{\delta}{4}$, then we get

$$\begin{aligned} & P_{A^T, B^T} \left(\frac{1}{T} \log \frac{P_{A^T, B^T}(A^T, B^T)}{\bar{P}_{A^T|B^T} P_{B^T}(A^T, B^T)} \leq \frac{1}{T} \log M + \frac{\delta}{4} \right) \\ & \leq P_{A^T, B^T} \left(\frac{1}{T} \log \frac{P_{A^T, B^T}(A^T, B^T)}{\bar{P}_{A^T|B^T} P_{B^T}(A^T, B^T)} \leq C - \frac{\delta}{4} \right) \\ & \leq \frac{\epsilon}{2} \end{aligned}$$

where the second inequality holds by for T large enough. To see this note that by the definition of C and T large enough the mass below $C - \frac{\delta}{4}$ has probability zero. \square

We have shown that C is the feedback channel capacity. It should be clear that if we restrict ourselves to channels without feedback then we recover the original coding theorem in [VH].

The direct theorem is an asymptotic theorem. We now provide error exponents to show the rate at which the random coding error decreases with T .

4.4.3 Error Exponents

In this section we provide upper bounds on the random coding probability of error. We show that at least in terms of the random coding error, the feedback coding error exponent is no smaller than the no feedback coding error exponent.

In the direct part of the coding theorem we constructed a distribution, $P(F^T)$, from a channel input distribution $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. By lemma 4.3.1 we know the channel $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$ and the code-function distribution $P(F^T)$ uniquely define a channel

from \mathcal{F}_T to \mathcal{B}^T denoted by $\{P(B_t | f^t, b^{t-1})\}_{t=1}^T$. Thus we can directly apply Gallager's random coding error exponent theorem to the " $\mathcal{F}_T - \mathcal{B}^T$ " channel. [Gal]

Definition 4.4.4 Given a channel $\{P(B_t | f^t, b^{t-1})\}_{t=1}^T$ define the error exponent to be

$$E_T(R, P(F^T)) \triangleq \max_{0 \leq \rho \leq 1} \left(-\rho R - \frac{1}{T} \ln \sum_{b^T} \left[\sum_{f^T} P(f^T) \{P(b^T | f^T)\}^{\frac{1}{1+\rho}} \right]^{1+\rho} \right).$$

Theorem 4.4.4 The average random coding error over (T, e^{TR}) channel codes drawn according to $P(F^T)$ can be upperbounded as

$$E_{\text{codebooks}(T, e^{TR})}(\text{error}) \leq e^{-TE_T(R, P(F^T))}.$$

Proof: This is theorem 5.6.1 in Gallager's text. [Gal] \square

Definition 4.4.5 The optimal error exponent is

$$E_T(R) \triangleq \sup_{P(F^T)} E_T(R, P(F^T)).$$

The optimization in definition 4.4.5 is rather difficult to compute in general. We now show that we can simplify the optimization by rewriting it as an optimization over channel input distributions defined on $\mathcal{A}^T \times \mathcal{B}^T$.

Definition 4.4.6 Given a channel $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$ define the directed error exponent to be

$$\vec{E}_T(R, \{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T) \triangleq \max_{0 \leq \rho \leq 1} \left(-\rho R - \frac{1}{T} \ln \sum_{b^T} \left[\sum_{a^T} \vec{P}(a^T | b^T) \{\vec{P}(b^T | a^T)\}^{\frac{1}{1+\rho}} \right]^{1+\rho} \right).$$

Definition 4.4.7 The optimal directed error exponent is

$$\vec{E}_T(R) \triangleq \sup_{\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T \in \mathcal{S}_T} \vec{E}_T(R, \{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T).$$

Where \mathcal{S}_T was defined in definition 4.3.19.

Proposition 4.4.1 Fix a channel $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$. Let the code-function distribution $P(F^T)$ be good with respect to the channel input distribution $\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. Then

$$\vec{E}_T(R, \{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T) = E_T(R, P(F^T)).$$

Proof: By lemma 4.3.1 we know the code-function distribution and the channel define a unique consistent measure $Q(F^T, A^T, B^T)$. Note $Q(f^T, a^T, b^T) > 0$ implies that $a_t = f_t(b^{t-1})$.

$$\begin{aligned}
\sum_{b^T} \left[\sum_{f^T} P(f^T) \{P(b^T|f^T)\}^{\frac{1}{1+\rho}} \right]^{1+\rho} &= \sum_{b^T} \left[\sum_{f^T} P(f^T) P(b^T|f^T) \{P(b^T|f^T)\}^{\frac{-\rho}{1+\rho}} \right]^{1+\rho} \\
&= \sum_{b^T} \left[\sum_{f^T} Q(f^T, b^T) \{Q(b^T|f^T)\}^{\frac{-\rho}{1+\rho}} \right]^{1+\rho} \\
&= \sum_{b^T} \left[\sum_{f^T, a^T} Q(f^T, a^T, b^T) \{Q(b^T|f^T)\}^{\frac{-\rho}{1+\rho}} \right]^{1+\rho} \\
&= \sum_{b^T} \left[\sum_{f^T, a^T} Q(f^T, a^T, b^T) \{\tilde{Q}(b^T|a^T)\}^{\frac{-\rho}{1+\rho}} \right]^{1+\rho} \\
&= \sum_{b^T} \left[\sum_{a^T} Q(a^T, b^T) \{\tilde{Q}(b^T|a^T)\}^{\frac{-\rho}{1+\rho}} \right]^{1+\rho} \\
&= \sum_{b^T} \left[\sum_{a^T} \tilde{Q}(a^T|b^T) \tilde{Q}(b^T|a^T) \{\tilde{Q}(b^T|a^T)\}^{\frac{-\rho}{1+\rho}} \right]^{1+\rho} \\
&= \sum_{b^T} \left[\sum_{a^T} \tilde{Q}(a^T|b^T) \{\tilde{Q}(b^T|a^T)\}^{\frac{1}{1+\rho}} \right]^{1+\rho} \\
&= \sum_{b^T} \left[\sum_{a^T} \tilde{P}(a^T|b^T) \{\tilde{P}(b^T|a^T)\}^{\frac{1}{1+\rho}} \right]^{1+\rho}
\end{aligned}$$

where the last line follows from lemma 4.4.4. Thus the proposition follows. \square

Corollary 4.4.2 Fix a channel $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$. Then

$$\vec{E}_T(R) = E_T(R).$$

Proof: Denote the channel input distribution that supremizes $\vec{E}_T(R, \{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T)$ by $\{P^*(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. Let $P(F^T)$ be any code-function distribution good with respect to this optimal channel input distribution. Then by proposition 4.4.1 must have

$$\vec{E}_T(R) = E_T(R, P(F^T)) \leq E_T(R).$$

Now let $P^*(F^T)$ supremize $E_T(R, P(F^T))$. By lemma 4.4.2 we can compute the induced channel input distribution $\{Q(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. Clearly $P^*(F^T)$ is a code-function distribution good with respect to $\{Q(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T$. Thus by proposition 4.4.1 we must

have

$$E_T(R) = \vec{E}_T \left(R, \{Q(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T \right) \leq \vec{E}_T(R).$$

Thus the corollary is proved. \square

We have reduced the calculation of the error exponent from an optimization problem over $\mathcal{F}_T \times \mathcal{B}^T$ to one over $\mathcal{A}^T \times \mathcal{B}^T$.

Corollary 4.4.3 *If $0 \leq R < C_T$ then $\vec{E}_T(R) > 0$.*

Proof: By theorem 5.6.4 in Gallager's text we know if $0 \leq R < C_T$ then $E_T(R) > 0$. [Gal] By corollary 4.4.2 the result holds. \square

We end this subsection by showing that the feedback channel error exponent can be no worse than the no feedback channel error exponent. The following definition is essentially definition 4.4.4 rewritten in terms of a^T instead of f^T .

Definition 4.4.8 *The optimal no feedback directed error exponent is*

$$\vec{E}_T^{nfb}(R) \triangleq \sup_{\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T \in \mathcal{S}_T^{nfb}} \vec{E}_T \left(R, \{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T \right).$$

Where \mathcal{S}_T^{nfb} was defined in definition 4.3.20.

Proposition 4.4.2 *For a given channel $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$ we have $\vec{E}_T(R) \geq \vec{E}_T^{nfb}(R)$.*

Proof:

$$\begin{aligned} \vec{E}_T(R) &= \sup_{\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T \in \mathcal{S}_T} \vec{E}_T \left(R, \left(\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T \right) \right) \\ &\geq \sup_{\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T \in \mathcal{S}_T^{nfb}} \vec{E}_T \left(R, \left(\{P(A_t | a^{t-1}, b^{t-1})\}_{t=1}^T \right) \right) \\ &= \vec{E}_T^{nfb}(R) \end{aligned}$$

\square

We have shown the feedback random coding error exponent cannot be smaller than the no feedback random coding error exponent. Thus not only is the feedback capacity larger than or equal to the no feedback capacity but also the error exponents can be better. Unfortunately the result does not tell us how much better the exponent can be. This is an open problem. But we will shed some light on this issue in chapter five where we relate the successive refinement problem to the feedback channel coding problem.

Finally note that we have defined the exponent for each time horizon T . It is not true in general that $\lim_{T \rightarrow \infty} E_T(R)$ exists or is bounded away from zero. Suitable regularity assumptions need to be made in order for this to occur.

4.5 Markov Channels and the Dynamic Programming Formulation

In this section we discuss channels with state. We provide a coding theorem for these channels. Furthermore we reduce the optimization problem to a partially observed stochastic control problem.

We have shown that the supremization over channel input distributions of the directed mutual information is a measure of the channel capacity. This optimization is much easier than determining the optimal distribution on the space of code-functions. But computing the optimal channel input distribution is still difficult. There are many approximation approaches that one can use. For example one could use the Blahut-Arimoto algorithm. Here, though, we would like to use the tools of dynamic programming to convert the optimization into a series of “simpler” optimizations.

For a dynamic programming approach to work, though, we require that the objective function be a summation of costs. This allows us to embed the general optimization problem into a series of simpler ones. A straightforward computation shows that the directed information can be written as the sum

$$I(A^T \rightarrow B^T) = \sum_{t=1}^T I(A^t; B_t | B^{t-1}).$$

Thus in principle one could apply dynamic programming ideas here. Unfortunately given the dependence on both A^t and B^t we see that the state space will be growing with t . The usual trick when dealing with a growing state space is to seek out a sufficient statistic. To that end we examine “Markov” channels.

4.5.1 Setup

In this subsection we define the Markov channel, the code-functions, the interconnection between channel code and channel, and the channel capacity. This material follows from section 4.3 Let \mathcal{Z} be a finite set representing the space where the state lives. Let $Z_t \in \mathcal{Z}$ be the state process. See figure 4-6.

Definition 4.5.1 *A Markov channel consists of two sequences of stochastic kernels. One sequence governing the state evolution: $\{P(Z_1), P(Z_{t+1} | z_t, a_t), t = 1, \dots, T-1\}$ and one sequence governing the channel output: $\{P(B_t | z_t, a_t)\}_{t=1}^T$. If the stochastic kernel $P(Z_{t+1} | z_t, a_t)$ is independent of a_t for $t = 1, \dots, T-1$ then we say the channel is a Markov channel without ISI (Intersymbol Interference.)*

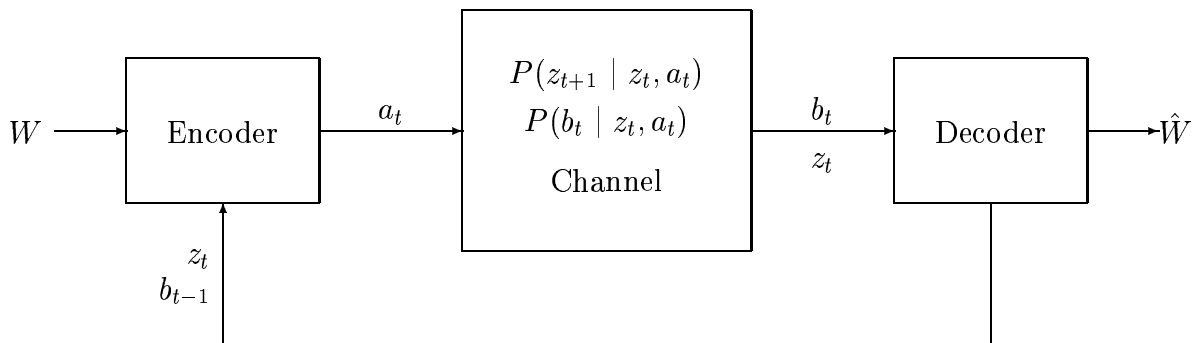


Figure 4-6: Markov Channel

Lemma 4.5.1 *For Markov channels the following is a Markov chain:*

$$(B_t, Z_{t+1}) - (A_t, Z_t) - (A^{t-1}, B^{t-1}, Z^{t-1})$$

Proof:

$$\begin{aligned} P(B_t, Z_{t+1} | a^t, b^{t-1}, z^t) &= P(Z_{t+1} | a^t, b^t, z^t) \otimes P(B_t | a^t, b^{t-1}, z^t) \\ &= P(Z_{t+1} | a_t, z_t) \otimes P(B_t | a_t, z_t) \end{aligned}$$

□

Definition 4.5.2 *A channel code-function is a sequence of T deterministic measurable functions $\{f_t\}_{t=1}^T$ such that $f_t : \mathcal{Z}^t \times \mathcal{B}^{t-1} \rightarrow \mathcal{A}$ which takes $(z^t, b^{t-1}) \mapsto a_t$. Two special cases are*

- (1) *There is no feedback to the encoder. The code-function, f_t , is independent of z^t, b^{t-1} .*
- (2) *The code-function has access only to the channel output. That is $f_t : \mathcal{B}^{t-1} \rightarrow \mathcal{A}$ which takes $b^{t-1} \mapsto a_t$.*

As before a channel code is a set of M code-functions. A message set is defined as in definition 4.3.2.

Definition 4.5.3 *A channel decoder is a map $g : \mathcal{Z}^T \times \mathcal{B}^T \rightarrow \mathcal{W}$ taking $(z^T, b^T) \mapsto w$.*

Note that the decoder is allowed to observe the state. In the literature this is sometimes called receiver CSI (channel state information.) We discuss relaxing this assumption in section 4.5.4.

Interconnection

The natural time ordering for the random variables involved is

$$F^T, Z_1, A_1, B_1, Z_2, \dots, Z_t, A_t, B_t, \dots, Z_T, A_T, B_T. \quad (4.5)$$

Given a measure $P(F^T)$ we can define an interconnection as we did in in section 4.3.3. We will define a new Markov channel without feedback. We want to compute a joint measure $Q(F^T, Z^T, A^T, B^T)$ such that the following all hold almost surely Q

- (1) There is no feedback to the code-functions:

$$Q(F_t | F^{t-1} = f^{t-1}, Z^{t-1} = z^{t-1}, B^{t-1} = b^{t-1}) = P(F_t | f^{t-1}) \\ Q(F^{t-1}, Z^{t-1}, B^{t-1}) - a.s.$$

- (2) Q preserves the properties of the underlying channel:

$$Q(Z_t | F^T = f^t, Z^{t-1} = z^{t-1}, A^{t-1} = a^{t-1}, B^{t-1} = b^{t-1}) = P(Z_t | z_{t-1}, a_{t-1}) \\ \text{or } Q(Z_t | F^t = f^t, Z^{t-1} = z^{t-1}, A^{t-1} = a^{t-1}, B^{t-1} = b^{t-1}) = P(Z_t | z_{t-1}) \\ Q(F^t, Z^{t-1}, A^{t-1}, B^{t-1}) - a.s.$$

depending on whether there is ISI or not. And

$$Q(B_t | F^T = f^t, Z^t = z^t, A^t = a^t, B^{t-1} = b^{t-1}) = P(B_t | z_t, a_t) \\ Q(F^t, Z^t, A^t, B^{t-1}) - a.s.$$

- (3) The channel input is a function of the past outputs and/or states:

$A_t = F_t(Z^t, B^{t-1})$ Q -a.s. or $A_t = F_t(B^{t-1})$ Q -a.s. depending on whether there is state feedback or not. More generally F_t can be any $\sigma(Z^t, B^{t-1})$ -measurable function.

Definition 4.5.4 *We call any measure satisfying the above four properties a consistent measure for the Markov channel.*

Lemma 4.5.2 *Given $P(F^T)$, a Markov channel $\{P(Z_1), P(Z_{t+1} | a_t, z_t), t = 1, \dots, T-1\}$ and $\{P(B_t | a_t, z_t)\}_{t=1}^T$, and the relations $a_t = f_t(z^t, b^{t-1})$, there exists a unique consistent measure $Q(F^T, Z^T, A^T, B^T)$. Similarly, there exists a unique consistent measure for the case without ISI and the case where the code-functions do not observe the state.*

Proof: We prove the case when there is ISI and the code-functions observe the state. The other cases are proved similarly. This lemma is a generalization of lemma 4.3.1.

Any measure satisfying (1) and (3) must be of the form

$$\begin{aligned}
& Q(F^T, Z^T, A^T, B^T) \\
&= \left\{ \bigotimes_{t=1}^T Q(B_t | f^t, z^t, b^{t-1}) \otimes Q(Z_t | f^t, z^{t-1}, b^{t-1}) \otimes Q(F_t | f^{t-1}, z^{t-1}, b^{t-1}) \right\} \\
&\quad \otimes Q(A^T | f^T, z^T, b^T) \\
&= \left\{ \bigotimes_{t=1}^T Q(B_t | f^t, z^t, b^{t-1}) \otimes Q(Z_t | f^t, z^{t-1}, b^{t-1}) \otimes P(F_t | f^{t-1}) \right\} \otimes \delta_{\{A^T = f^T(z^T, b^{T-1})\}}
\end{aligned}$$

Where $f^T(z^T, b^{T-1}) \triangleq (f_1(z_1), f_2(z^2, b_1), \dots, f_T(z^T, b^{T-1}))$. Thus we need only identify: $\{Q(B_t | f^t, z^t, b^{t-1})\}_{t=1}^T$ and $\{Q(Z_t | f^t, z^{t-1}, b^{t-1})\}_{t=1}^T$.

We know that (f^t, z^t, b^{t-1}) uniquely identifies a^t . Thus it must be the case that

$$\begin{aligned}
& Q(B_t | F^t = f^t, Z^t = z^t, B^{t-1} = b^{t-1}) \\
&= Q(B_t | F^t = f^t, Z^t = z^t, A^t = f^t(z^t, b^{t-1}), B^{t-1} = b^{t-1}) \quad Q(F^t, Z^t, A^t, B^{t-1}) - a.s. \\
&= P(B_t | z_t, f_t(z^t, b^{t-1})) \quad Q(F^t, Z^t, A^t, B^{t-1}) - a.s.
\end{aligned}$$

and

$$\begin{aligned}
& Q(Z_t | F^t = f^t, Z^{t-1} = z^{t-1}, B^{t-1} = b^{t-1}) \\
&= Q(Z_t | F^t = f^t, Z^{t-1} = z^{t-1}, A^{t-1} = f^{t-1}(z^{t-1}, b^{t-2}), B^{t-1} = b^{t-1}) \\
&\quad Q(F^t, Z^{t-1}, A^{t-1}, B^{t-1}) - a.s. \\
&= P(Z_t | z_{t-1}, f_{t-1}(z^{t-1}, b^{t-2})) \quad Q(F^t, Z^{t-1}, A^{t-1}, B^{t-1}) - a.s.
\end{aligned}$$

□

Operational Channel Capacity

We define the operational channel capacity just as in subsection 3.3.5.

Definition 4.5.5 *R is an ϵ -achievable rate if, for every $\delta > 0$ there exists, for sufficiently large T, an (T, M, ϵ) channel code with rate $\frac{\log M}{T} > R - \delta$. The maximum ϵ -achievable rate is called the ϵ -capacity and denoted C_ϵ^o . The operational channel capacity is defined as the maximal rate that is achievable for all $0 < \epsilon < 1$ and is denoted C^o .*

Let $C^o, \text{out}, C^o, \text{nfb}$ represent the operational capacity under output feedback alone and no feedback respectively.

Channel Capacity

We define channel input distributions and the channel capacity just as in section 3.3.6.

Definition 4.5.6 For all T Let

- (1) $\mathcal{S}_T = \{\{P(A_t | z^t, a^{t-1}, b^{t-1})\}_{t=1}^T\}$.
- (2) $\mathcal{S}_T^{out} \subset \mathcal{S}_T$ be the set of all channel input distributions without state feedback.
- (3) $\mathcal{S}_T^{nfb} \subset \mathcal{S}_T^{out}$ be the set of all channel input distributions without any feedback.

Definition 4.5.7 For finite T let

- (1) $C_T = \sup_{S \in \mathcal{S}_T} \frac{1}{T} I(A^T \rightarrow (Z^T, B^T))$
- (2) $C_T^{out} = \sup_{S \in \mathcal{S}_T^{out}} \frac{1}{T} I(A^T \rightarrow (Z^T, B^T))$
- (3) $C_T^{nfb} = \sup_{S \in \mathcal{S}_T^{nfb}} \frac{1}{T} I(A^T \rightarrow (Z^T, B^T))$

For the infinite horizon case let

- (1) $C = \sup_{S \in \mathcal{S}_\infty} \underline{I}(A \rightarrow (Z, B))$
- (2) $C^{out} = \sup_{S \in \mathcal{S}_\infty^{out}} \underline{I}(A \rightarrow (Z, B))$
- (3) $C^{nfb} = \sup_{S \in \mathcal{S}_\infty^{nfb}} \underline{I}(A \rightarrow (Z, B))$

In the next subsection we will show that $C^o = C$, $C^{o, out} = C^{out}$, and $C^{o, nfb} = C^{nfb}$.

4.5.2 Coding Theorem

We now prove the feedback Markov channel coding theorem. The Markov channel is a special case of the finite alphabet channel we examined in section 4.3. Thus we can apply theorem 4.4.1 directly.

Theorem 4.5.1 $C^o = C$, $C^{o, out} = C^{out}$, and $C^{o, nfb} = C^{nfb}$.

Proof: A Markov channel is just a special case of the channel we examined in section 4.3. Let $V_0 = Z_0$, $V_t = (B_t, Z_{t+1})$, $t = 1, \dots, T-1$, and $V_T = B_T$. Then we can define a new

“ $A - V$ ” channel. Specifically this channel is $\{P(V_0), P(V_t | a^t, v^{t-1}), t = 1, \dots, T\}$. The time ordering of the random variables involved is, as before,

$$F^T, V_0, A_1, V_1, \dots, A_T, V_T.$$

There are two differences between this channel and the channels defined in section 4.3. First there is the addition of the initial state: $P(V_0)$. Second the spaces in which the V_t live in are changing with t . After straightforward corrections for these changes we can apply theorem 4.4.1 directly. \square

4.5.3 Stochastic Control Formulation

In this subsection we introduce the stochastic control formulation. The optimization problem consists of supremizing over \mathcal{S}_T the directed mutual information between the channel input and the channel output and state.

$$\frac{1}{T} I(A^T \rightarrow (Z^T, B^T)).$$

The following lemma describes how to write this as a sum of “costs.”

Lemma 4.5.3 *For Markov channels we have*

$$I(A^T \rightarrow (Z^T, B^T)) = \sum_{t=1}^{T-1} I(A_t; (B_t, Z_{t+1}) | B^{t-1}, Z^t) + I(A_T; B_T | B^{T-1}, Z^T)$$

Proof: By proposition 4.2.1 and using the time ordering described in equation (4.5) we see that

$$\begin{aligned} I(A^T \rightarrow (Z^T, B^T)) &= D(P_{Z^T, A^T, B^T} | \vec{P}_{A^T | Z^T, B^T} P_{Z^T, B^T}) \\ &= \sum_{z^T, a^T, b^T} P(z^T, a^T, b^T) \log \frac{\vec{P}(z^T, b^T | a^T)}{P(z^T, b^T)} \\ &= \sum_{t=1}^{T-1} I(A^t; (B_t, Z_{t+1}) | B^{t-1}, Z^t) + I(A^T; B_T | B^{T-1}, Z^T) \\ &= \sum_{t=1}^{T-1} I(A_t; (B_t, Z_{t+1}) | B^{t-1}, Z^t) + I(A_T; B_T | B^{T-1}, Z^T) \end{aligned}$$

The third equality follows by noting definition 4.2.2 and

$$\begin{aligned} \vec{P}(Z^T, B^T | a^T) &= P(Z_1) \otimes P(B_1 | z_1, a_1) \otimes P(Z_2 | z_1, a_1, b_1) \otimes \dots \otimes P(B_T | z^T, a^T, b^{T-1}) \\ &= P(Z_1) \otimes P(B_1, Z_2 | z_1, a_1) \otimes \dots \otimes P(B_{T-1}, Z_T | z^{T-1}, a^{T-1}, b^{T-2}) \\ &\quad \otimes P(B_T | z^T, a^T, b^{T-1}). \end{aligned}$$

To see the fourth equality holds note

$$\begin{aligned}
& I(A^t; (B_t; Z_{t+1}) \mid B^{t-1}, Z^t) \\
&= I(A_t; (B_t; Z_{t+1}) \mid B^{t-1}, Z^t) + I(A^{t-1}; (B_t; Z_{t+1}) \mid B^{t-1}, Z^t, A_t) \\
&= I(A_t; (B_t; Z_{t+1}) \mid B^{t-1}, Z^t)
\end{aligned}$$

where the last line follows because, by lemma 4.5.1, $(Z_{t+1}, B_t) - (A_t, B^{t-1}, Z^t) - A^{t-1}$ forms a Markov chain. We can prove $I(A^T; B_T \mid B^{T-1}, Z^T) = I(A_T; B_T \mid B^{T-1}, Z^T)$ holds in a similar manner. \square

Corollary 4.5.1 *For Markov channels without ISI we have*

$$I(A^T \rightarrow (Z^T, B^T)) = \sum_{t=1}^T I(A_t; B_t \mid B^{t-1}, Z^t)$$

Proof:

$$\begin{aligned}
I(A_t; (B_t, Z_{t+1}) \mid B^{t-1}, Z^t) &= I(A_t; Z_{t+1} \mid B^t, Z^t) + I(A_t; B_t \mid B^{t-1}, Z^t) \\
&= I(A_t; B_t \mid B^{t-1}, Z^t)
\end{aligned}$$

where the last line follows because $Z_{t+1} - (B^t, Z^t) - A_t$ forms a Markov chain when there is no ISI. \square

Note that the mutual information $I(A^T \rightarrow (Z^T, B^T))$ can be decomposed into two sums. The first sum represents the contribution of the channel input and the channel output and the second sum represents the contribution of the channel input and the next channel state:

$$\begin{aligned}
I(A^T \rightarrow (Z^T, B^T)) &= \sum_{t=1}^{T-1} I(A_t; (B_t, Z_{t+1}) \mid B^{t-1}, Z^t) + I(A_T; B_T \mid B^{T-1}, Z^T) \\
&= \sum_{t=1}^T I(A_t; B_t \mid B^{t-1}, Z^t) + \sum_{t=1}^{T-1} I(A_t; Z_{t+1} \mid B^t, Z^t)
\end{aligned}$$

When there is no ISI the second sum in the last line equals zero.

Without loss of generality restrict we can restrict attention to channel input distributions that are independent of a^{t-1} . This is a consequence of the decoder observing the current state. We will discuss in subsection 4.5.4 why there is a dependence on a^{t-1} in the case when the decoder does not observe the state.

We prove two final simplifications before presenting the details of the stochastic control problem. Define G_t to be the information available to the encoder at time t . Recall at one extreme: $G_t = (B^{t-1}, Z^t)$. One intermediate case is if there is no state feedback: $G_t = B^{t-1}$.

And the other extreme occurs if there is no feedback at all: $G_t = \emptyset$. There are of course many other information patterns in between.

Lemma 4.5.4

$$I(A_t; (B_t, Z_{t+1}) \mid B^{t-1}, Z^t) = I(A_t; (B_t, Z_{t+1}) \mid Z_t, G_t) \quad t = 1, \dots, T-1$$

and

$$I(A_T; B_T \mid B^{T-1}, Z^T) = I(A_T; B_T \mid Z_T, G_T)$$

Proof: By lemma 4.5.1 we know $(B_t, Z_{t+1}) - (A_t, Z_t) - (B^{t-1}, Z^{t-1})$ forms a Markov chain. Also $A_t - G_t - (B^{t-1}, Z^t)$ forms a Markov chain. Thus $(B_t, Z_{t+1}) - (G_t, Z_t) - (B^{t-1}, Z^{t-1})$ forms a Markov chain. Now

$$\begin{aligned} I(A_t; (B_t, Z_{t+1}) \mid B^{t-1}, Z^t) &= I(A_t; (B_t, Z_{t+1}) \mid B^{t-1}, Z^t, G_t) \\ &= I((B_t, Z_{t+1}); A_t, B^{t-1}, Z^t, G_t) - I((B_t, Z_{t+1}); B^{t-1}, Z^t, G_t) \\ &= I((B_t, Z_{t+1}); A_t, Z_t, G_t) - I((B_t, Z_{t+1}); B^{t-1}, Z^t, G_t) \\ &= I((B_t, Z_{t+1}); A_t, Z_t, G_t) - I((B_t, Z_{t+1}); Z_t, G_t) \\ &= I(A_t; (B_t, Z_{t+1}) \mid Z_t, G_t) \end{aligned}$$

The other equality in the lemma statement is proved analogously. \square

Let $\pi_t(g_t) = P(Z_t \mid g_t)$ be the conditional probability of the current state given the information available to the encoder at time t . Note that Π_t is well defined once we have fixed the channel input distribution $\{P(A_t \mid g_t)\}_{t=1}^T$. We now show that Π_t is a sufficient statistic.

Lemma 4.5.5

$$\sup_{P(A_t \mid g_t)} I(A_t; (B_t, Z_{t+1}) \mid Z_t, G_t) = \sup_{P(A_t \mid \pi_t)} I(A_t; (B_t, Z_{t+1}) \mid Z_t, \Pi_t) \quad t = 1, \dots, T-1$$

and

$$\sup_{P(A_T \mid g_T)} I(A_T; B_T \mid Z_T, G_T) = \sup_{P(A_T \mid \pi_T)} I(A_T; B_T \mid Z_T, \Pi_T)$$

Proof: One direction is obvious:

$$\sup_{P(A_t \mid g_t)} I(A_t; (B_t, Z_{t+1}) \mid Z_t, G_t) \geq \sup_{P(A_t \mid \pi_t)} I(A_t; (B_t, Z_{t+1}) \mid Z_t, \Pi_t).$$

This is because the set on the left hand side that we are supremizing over is larger. (Recall π_t is a function of g_t .) We now show the inequality in the other direction holds. Note that

$Z_t - \Pi_t - G_t$ forms a Markov chain.

$$\begin{aligned}
I(A_t; (B_t, Z_{t+1}) | Z_t, G_t) &= I(A_t; (B_t, Z_{t+1}) | Z_t, G_t, \Pi_t) \\
&= \sum_{z_t, g_t, \pi_t} P(z_t, g_t, \pi_t) I(A_t; (B_t, Z_{t+1}) | z_t, g_t, \pi_t) \\
&= \sum_{z_t, \pi_t} P(z_t, \pi_t) \left\{ \sum_{g_t} P(g_t | \pi_t) I(A_t; (B_t, Z_{t+1}) | z_t, g_t, \pi_t) \right\} \\
&\leq \sum_{z_t, \pi_t} P(z_t, \pi_t) I(A_t; (B_t, Z_{t+1}) | z_t, \pi_t) \\
&= I(A_t; (B_t, Z_{t+1}) | Z_t, \Pi_t)
\end{aligned}$$

Where the inequality follows from the convexity of mutual information with respect to the input distribution: $P(a_t | \pi_t) = \sum_{g_t} P(a_t | g_t) P(g_t | \pi_t)$. Thus we see that

$$\sup_{P(A_t | g_t)} I(A_t; (B_t, Z_{t+1}) | Z_t, G_t) \leq \sup_{P(A_t | \pi_t)} I(A_t; (B_t, Z_{t+1}) | Z_t, \Pi_t).$$

The second equality in the lemma statement is proved analogously. \square

We can conclude from this lemma that the receiver should feedback information only helpful in estimating the state of the channel.

The Dynamic Programming Equation

We are now ready to present the dynamic programming formulation. This formulation is rather powerful. It allows us to unify many existing results in the literature and provides many new coding results. Furthermore we are now in a position to apply many of the exact and approximate solution methods found in the dynamic programming literature. Recall the optimization problem is

$$\sup \sum_{t=1}^{T-1} I(A_t; (B_t, Z_{t+1}) | Z_t, G_t) + I(A_T; B_T | Z_T, G_T).$$

Where the supremum is over all admissible channel input distributions. We now translate the elements of this problem into a traditional control problem.

- *State*

Let the state process be Z_t . The dynamics of the state are determined by

$$P(Z_{t+1} | z_t, u_t) = \sum_{a_t} P(Z_{t+1} | z_t, a_t) u_t(a_t).$$

- *Control*

The control, U_t , takes values in $\mathcal{P}(\mathcal{A})$. It is allowed to be a function of the past controls and its observations G^t . At time t the controller draws an a_t from the distribution u_t and inputs it into the channel.

- *Observation*

The observation is G_t . The channel output follows

$$P(B_t | z_t, u_t) = \sum_{a_t} P(B_t | z_t, a_t) u_t(a_t).$$

- *Running and terminal cost*

For $t = 1, \dots, T - 1$

$$c_t(z_t, u_t) = \sum_{a_t, b_t, z_{t+1}} P(b_t, z_{t+1} | z_t, a_t) u_t(a_t) \log \frac{P(b_t, z_{t+1} | z_t, a_t)}{\sum_{\tilde{a}_t} P(b_t, z_{t+1} | z_t, \tilde{a}_t) u_t(\tilde{a}_t)}$$

$$c_T(z_T, u_T) = \sum_{a_T, b_T} P(b_T | z_T, a_T) u_T(a_T) \log \frac{P(b_T | z_T, a_T)}{\sum_{\tilde{a}_T} P(b_T | z_T, \tilde{a}_T) u_T(\tilde{a}_T)}.$$

By lemma 4.5.5 we know that Π_t is a sufficient statistic. Unfortunately given a general observation G_t it does not follow that Π_t will be a controlled Markov process (i.e. $P(\Pi_{t+1} | \pi^t, u^t) = P(\Pi_{t+1} | \pi_t, u_t)$.) We need to place some restrictions on the form of the observation.

Lemma 4.5.6 *For the following cases*

- (1) $G_t = (B^{t-1}, Z^t)$ (full observation)
- (2) $G_t = B^{t-1}$ (channel output feedback)
- (3) $G_t = \emptyset$ (no feedback at all)

the process Π_t is a controlled Markov process.

Proof: Case one is obvious. We prove case two:

$$\begin{aligned} P(z_{t+1} | g^{t+1}, u^t) &= P(z_{t+1} | b^t, u^t) \\ &= \sum_{z_t} P(z_{t+1} | z_t, b^t, u^t) P(z_t | b^t, u^t) \\ &= \sum_{z_t} P(z_{t+1} | z_t, u_t) \frac{P(b_t | z_t, u^t, b^{t-1}) P(z_t | b^{t-1}, u^t)}{\sum_{\tilde{z}_t} P(b_t | \tilde{z}_t, u^t, b^{t-1}) P(\tilde{z}_t | b^{t-1}, u^t)} \\ &= \sum_{z_t} P(z_{t+1} | z_t, u_t) \frac{P(b_t | z_t, u_t) P(z_t | b^{t-1}, u^{t-1})}{\sum_{\tilde{z}_t} P(b_t | \tilde{z}_t, u_t) P(\tilde{z}_t | b^{t-1}, u^{t-1})} \\ &= \Phi \left(P(z_t | g^t, u^{t-1}), b_t, u_t \right) \end{aligned}$$

for some function Φ . Thus $\pi_{t+1} = \Phi(\pi_t, b_t, u_t)$. Case three follows analogously. \square

We can now convert the above partially observed Markov decision problem into a fully observed Markov decision problem. The new state process is Π_t and the new running cost is

$$\bar{c}(\pi_t, u_t) = \sum_{z_t} c(z_t, u_t) \pi_t(z_t).$$

Theorem 4.5.2 *The optimal cost and control can be determined by solving the following dynamic programming equations*

$$J_T(\pi) = \sup_{u_T} \bar{c}(\pi, u_T)$$

and

$$J_t(\pi) = \sup_{u_t} \bar{c}(\pi, u_t) + \int J_{t+1}(\tilde{\pi}) P(d\tilde{\pi} | \pi, u_t)$$

where u_t is a function of π_t .

Proof: By lemma 4.5.5 and lemma 4.5.6 we have shown that Π_t is a sufficient statistic and a controlled Markov process. The theorem then follows from theorem 5.4.8 of [Str]. \square

Infinite Horizon Average Cost Problem

So far we have dealt with the finite horizon problem. We can also treat the average cost infinite horizon problem:

$$\sup \liminf_{T \rightarrow \infty} \frac{1}{T} \left\{ \sum_{t=1}^{T-1} I(A_t; (B_t, Z_{t+1}) | Z_t, G_t) + I(A_T; B_T | Z_T, G_T) \right\}$$

where the supremum is over all admissible input distributions.

Assumption 4.5.1 *Assume that under all stationary control policies the process Z_t is an ergodic and aperiodic Markov chain.*

Note if there is no ISI then assumption 4.5.1 states that, independent of the policy, the Z_t process is an ergodic, aperiodic, Markov chain.

Theorem 4.5.3 *Under assumption 4.5.1 the optimal control is the solution to the following average cost Bellman equation*

$$J^* + h(\pi) = \max_u \bar{c}(\pi, u) + \int h(\tilde{\pi}) P(d\tilde{\pi} | \pi, u)$$

Proof: See theorem 4.1 of [Bor]. \square

There are many ways to solve the Bellman equation. Traditional methods include value and policy iteration. We can also apply linear programming tools. We state now the linear program:

$$\min J^*$$

subject to

$$J^* + h(\pi) \geq \bar{c}(\pi, u) + \sum_{\tilde{\pi} \in T(\pi, u)} h(\tilde{\pi})P(\tilde{\pi}|\pi, u) \quad \forall \pi, u$$

where $T(\pi, u) = \{\tilde{\pi} : P(\tilde{\pi}|\pi, u) > 0\}$.

4.5.4 Extensions and Limitations of our Formulation

Extensions

- (1) We can treat the case of delayed feedback by the usual state augmentation techniques. See [Bert] for examples of state augmentation. For the case of channels without ISI we do not need to augment the state. See [Vis] for similar results.
- (2) The error exponents in section 4.4.3 can also be described as a partially observed stochastic control problem. Specifically the optimization problem defined in definition 4.4.7 can be written as a dynamic program. There is, though, a second optimization over ρ .

Limitations

- (1) One strong assumption in our model is that the state is known to the decoder. As shown in lemma 4.5.3, this allows for the channel input distribution to be independent of the past a^{t-1} . In the case when the decoder does not know the channel state the optimization problem is much harder. It is not clear that we can decompose the cost into a sum of costs. See theorem 4.6.1 of [Gal] for discussion on how to compute the optimal channel input distribution without feedback. Also see [GV1] for another method of computing the channel input distribution without feedback. The issue of dependence on the past a 's is taken up in [HJ]. There they discuss the role of mixing in finite state channels. If the channel mixes quickly then intuitively there is no need for the channel input distribution to depend on all the past a 's. They quantify this dependence.
- (2) Another assumption in our model is that the receiver knows what information the transmitter has. Specifically the receiver knows G_t . We can treat the case where the transmitter uses information the receiver does know by the usual trick of code-functions. [Sha3] But it is not clear that the simplification of the optimization over code-function distributions to the optimization over channel input distributions will continue to hold.

4.6 Coding Theorem for Gaussian Channels

In this section we address Gaussian channels with feedback. We define the Gaussian channel and provide a coding theorem. We will follow the same steps as we did in sections 4.3 and 4.4 and make appropriate extensions for the Gaussian case as needed.

4.6.1 Setup

Let the channel input alphabet be \mathcal{X} and channel output alphabet be \mathcal{Y} . Where $\mathcal{X} = \mathcal{Y} = \mathbb{R}$. Endow \mathcal{X} and \mathcal{Y} with the usual Lebesgue measure.

Definition 4.6.1 A Gaussian channel is a sequence of stochastic kernels, $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$, such that there exist vectors $\alpha_t \in \mathbb{R}^t$, $t = 1, \dots, T$ and $\beta_t \in \mathbb{R}^{t-1}$, $t = 2, \dots, T$ and independent Gaussian random variables $V_t \sim \mathcal{N}(0, K_{V_t})$ such that for all (x^t, y^{t-1})

$$Y_t = \alpha_t' x^t + \beta_t' y^{t-1} + V_t.$$

The mean of the Gaussian measure on Y_t is dependent on (x^t, y^{t-1}) but the variance is not. This fact, as we will see, allows for a simplification in the design of the code-functions.

Note that we have defined the stochastic kernel, $P(dY_t | x^t, y^{t-1})$, in terms of the recurrence $Y_t = \alpha_t' x^t + \beta_t' y^{t-1} + V_t$. The two approaches are equivalent.

Message Set, Encoder, and Decoder

We use the same definitions of message set, encoder, and decoder as given in section 4.3.2 except for the following change. A code-function, f_t , is now a measurable map from $\mathcal{Y}^{t-1} \rightarrow \mathcal{X}$ taking $y^{t-1} \mapsto x_t$. The space of code-functions is still denoted \mathcal{F}_T . Note that this can be quite a complicated space. We will show, though, that without loss of generality we can restrict our attention to functions, f_t , that are affine in y^{t-1} .

Interconnection

Given a measure $P(dF^T)$ we can define an interconnection as we did in section 4.3.3. Specifically we want to compute a joint measure $Q(dF^T, dX^T, dY^T)$ such that the following all hold Q almost surely:

- (1) $Q(dF_t | F^{t-1} = f^{t-1}, Y^{t-1} = y^{t-1}) = P(dF_t | f^{t-1})$
- (2) $X_t = F_t(Y^{t-1})$
- (3) $Q(dY_t | F^t = f^t, X^t = x^t, Y^{t-1} = y^{t-1}) = P(dY_t | x^t, y^{t-1})$

As before we call any measure satisfying the above three properties a *consistent* measure.

Lemma 4.6.1 *Given $P(dF^T)$, a Gaussian channel, $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$, and the relations $x_t = f_t(y^{t-1})$, there exists a unique consistent measure $Q(dF^T, dX^T, dY^T)$.*

Proof: This is a generalization of lemma 4.3.1. Any measure satisfying (1) and (2) must be of the form

$$Q(dF^T, dX^T, dY^T) = Q(dX^T | f^T, y^T) \otimes \left\{ \bigotimes_{t=1}^T Q(dY_t | f^t, y^{t-1}) \otimes P(dF_t | f^{t-1}) \right\}$$

where $Q(dX^T | f^T, y^T) = \delta_{(X^T = f^T(y^{T-1}))}$ is a Dirac measure at the point $f^T(y^{T-1})$. Recall $f^T(y^{T-1}) = (f_1, f_2(y_1), \dots, f_T(y^{T-1}))$.

We now need to identify the channel $\{Q(dY_t | f^t, y^{t-1})\}_{t=1}^T$. The following hold almost surely $Q(dF^t, dX^t, dY^{t-1})$

$$\begin{aligned} Q(dY_t | F^t = f^t, Y^{t-1} = y^{t-1}) &= Q(dY_t | F^T = f^t, X^t = f^t(y^{t-1}), Y^{t-1} = y^{t-1}) \\ &= P(dY_t | f^t(y^{t-1}), y^{t-1}) \end{aligned}$$

□

Operational Channel Capacity

Recall a channel code is M code-functions $f^T[w]$, $w = 1, \dots, M$. For a given channel code let P_{F^T} place mass $\frac{1}{M}$ on each code function $f^T[w]$, $w = 1, \dots, M$.

Definition 4.6.2 *A channel code and a channel over a horizon T are said to satisfy a power constraint at power K if under the consistent measure $Q(dF^T, dX^T, dY^T)$ we have $\frac{1}{T}E_Q(\sum_{t=1}^T X_t^2) \leq K$.*

The definition of operational capacity is the same as definition 4.3.11 except we add the power constraint.

Definition 4.6.3 *R is an ϵ -achievable rate at power K if, for every $\delta > 0$ there exists, for sufficiently large T , an (T, M, ϵ) channel code with rate $\frac{\log M}{T} > R - \delta$ and $\frac{1}{T}E_Q(\sum_{t=1}^T X_t^2) \leq K + \epsilon$. The maximum ϵ -achievable rate at power K is called the ϵ -capacity at power K and denoted $C_\epsilon^O(K)$. The operational channel capacity at power K is defined as the maximal rate that is achievable for all $0 < \epsilon < 1$ and is denoted $C^O(K)$.*

Channel Capacity

Just as in definition 4.3.13 we define a channel input distribution to be a sequence of stochastic kernels $\{P(dX_t | x^{t-1}, y^{t-1})\}_{t=1}^T$. We also define the Gaussian channel input distribution:

Definition 4.6.4 A Gaussian channel input distribution, $\{P(dX_t | x^{t-1}, y^{t-1})\}_{t=1}^T$, is a sequence of stochastic kernels such that there exist vectors $\gamma_t, \eta_t \in \mathbb{R}^{t-1}, t = 2, \dots, T$ and independent Gaussian random variables $W_t \sim \mathcal{N}(0, K_{W_t})$ such that for all (x^{t-1}, y^{t-1})

$$X_t = \gamma_t' x^{t-1} + \eta_t' y^{t-1} + W_t.$$

Note that we have defined the stochastic kernel, $P(dX_t | x^{t-1}, y^{t-1})$, in terms of the recurrence $X_t = \gamma_t' x^{t-1} + \eta_t' y^{t-1} + W_t$. As stated before the two approaches are equivalent.

Definition 4.6.5 For a sequence of joint measures $\{P(dX^T, dY^T)\}_{T=1}^\infty$ let $\underline{I}(X \rightarrow Y) \triangleq \liminf_{in\ prob} \frac{1}{T} \log \frac{dP_{X^T, Y^T}}{d\bar{P}_{X^T|Y^T} P_{Y^T}}(X^T, Y^T)$.

As before let \mathcal{S}_T be the set of all channel input distributions. We now define the channel capacity.

Definition 4.6.6 For finite T let

$$C_T(K) = \sup_{S \in \mathcal{S}_T} \frac{1}{T} I(X^T \rightarrow Y^T) \text{ such that } E \left(\frac{1}{T} \sum_{t=1}^T X_t^2 \right) \leq K$$

and for the infinite horizon case let

$$C(K) = \sup_{S \in \mathcal{S}_\infty} \underline{I}(X \rightarrow Y) \text{ such that } \limsup_{T \rightarrow \infty} E \left(\frac{1}{T} \sum_{t=1}^T X_t^2 \right) \leq K.$$

In the next subsection we will prove that $C(K) = C^0(K)$. Furthermore we will show that the optimization can, without loss of generality, be restricted to Gaussian channel input distributions.

4.6.2 Main Theorem

Before proceeding to the coding theorem we prove the following structure result.

Structure Result

In the following if a measure P admits a density with respect to the Lebesgue measure then denote its density with the lower case p .

Lemma 4.6.2 Let $\{G(dY_t | x^t, y^{t-1})\}_{t=1}^T$ be a Gaussian channel with density g . Let $\{P(dX_t | x^{t-1}, y^{t-1})\}_{t=1}^T$ be a channel input distribution. The joint measure is $P(dX^T, dY^T) = \otimes_{t=1}^T P(dX_t | x^{t-1}, y^{t-1}) \otimes G(dY_t | x^t, y^{t-1})$. Furthermore assume that $P(dX^T, dY^T)$ admits a density p . Let $G(dX^T, dY^T)$ be a jointly Gaussian measure with the same second order properties as $P(dX^T, dY^T)$. Then

- (a) $\{G(dX_t | x^{t-1}, y^{t-1})\}_{t=1}^T$ is a Gaussian channel input distribution.
- (b) $G(dX^T, dY^T)$ has the same independence properties as $P(dX^T, dY^T)$. Hence they have the same power.
- (c) $I_G(X^T \rightarrow Y^T) \geq I_P(X^T \rightarrow Y^T)$.

Proof: Part (a) follows from the fact that $G(dX^T, dY^T)$ is jointly Gaussian. Part (b) follows from noting that independence or conditional independence of some set of random variables implies that those same random variables are uncorrelated or conditionally uncorrelated. $G(dX^T, dY^T)$ has the same second order properties as $P(dX^T, dY^T)$ thus it inherits the same independence properties. For part (c) note

$$\begin{aligned}
& I_G(X^T \rightarrow Y^T) - I_P(X^T \rightarrow Y^T) \\
&= D(G_{X^T, Y^T} | \vec{G}_{X^T | Y^T} G_{Y^T}) - D(P_{X^T, Y^T} | \vec{P}_{X^T | Y^T} P_{X^T}) \\
&= \int g_{X^T, Y^T}(x^T, y^T) \log \frac{g_{X^T, Y^T}(x^T, y^T)}{\vec{g}_{X^T | Y^T} g_{Y^T}(x^T, y^T)} dx^T dy^T \\
&\quad - \int p_{X^T, Y^T}(x^T, y^T) \log \frac{p_{X^T, Y^T}(x^T, y^T)}{\vec{p}_{X^T | Y^T} p_{Y^T}(x^T, y^T)} dx^T dy^T \\
&= \int p_{X^T, Y^T}(x^T, y^T) \log \frac{g_{X^T, Y^T}(x^T, y^T)}{\vec{g}_{X^T | Y^T} g_{Y^T}(x^T, y^T)} dx^T dy^T \\
&\quad - \int p_{X^T, Y^T}(x^T, y^T) \log \frac{p_{X^T, Y^T}(x^T, y^T)}{\vec{p}_{X^T | Y^T} p_{Y^T}(x^T, y^T)} dx^T dy^T \\
&= \int p_{X^T, Y^T}(x^T, y^T) \log \frac{\vec{g}_{Y^T | X^T}(y^T | x^T)}{g_{Y^T}(y^T)} dx^T dy^T \\
&\quad - \int p_{X^T, Y^T}(x^T, y^T) \log \frac{\vec{p}_{Y^T | X^T}(y^T | x^T)}{p_{Y^T}(y^T)} dx^T dy^T \\
&= \int p_{X^T, Y^T}(x^T, y^T) \log \left(\frac{\vec{g}_{Y^T | X^T}(y^T | x^T)}{g_{Y^T}(y^T)} \frac{p_{Y^T}(y^T)}{\vec{p}_{Y^T | X^T}(y^T | x^T)} \right) dx^T dy^T \\
&= \int p_{X^T, Y^T}(x^T, y^T) \log \frac{p_{Y^T}(y^T)}{g_{Y^T}(y^T)} dx^T dy^T \\
&= D(P_{Y^T} | G_{Y^T}) \\
&\geq 0
\end{aligned}$$

Where the third equality comes from the fact that the G has the same second order properties as P . The sixth equality follows because $\vec{p}(y^T | x^T) = \vec{g}(y^T | x^T)$. \square

For Gaussian channels with a power constraint the supremizing channel input distribution can be taken to be a Gaussian channel input distribution.

The Asymptotic Equi-partition Property for Gaussian Processes

The following theorem from [CP] is a generalization of the AEP to non-ergodic Gaussian processes.

Theorem 4.6.1 *Let $\{Z_t\}$ be an arbitrary Gaussian stochastic process with density p . Define the differential entropies $h_T \triangleq -\frac{1}{T} \int p(z^T) \log p(z^T) dz^T$. Then*

$$\frac{1}{T} \log p(Z^T) - h_T \rightarrow 0 \quad \text{w.p.1.}$$

Proof: See theorem 5 of [CP]. Also [Pin]. \square

We have shown that the optimal channel input distribution can be taken to be a Gaussian channel input distribution. This along with a Gaussian channel defines a jointly Gaussian measure on (X^T, Y^T) . In section 4.4 we proved a coding theorem for finite alphabet channels. The operational capacity was shown to equal a particular optimization of \underline{I} . By using theorem 4.6.1 we will show that for Gaussian channels we can work directly with I . This is because the process is information stable. We can often explicitly compute the value of I .

Corollary 4.6.1 *For any sequence of jointly Gaussian measures $\{P(dX^T, dY^T)\}_{t=1}^{\infty}$ we have*

$$\underline{I}(X \rightarrow Y) = \liminf_{T \rightarrow \infty} \frac{1}{T} I(X^T \rightarrow Y^T).$$

Proof: From theorem 4.6.1 we see

$$\begin{aligned} & \frac{1}{T} \log \vec{i}(X^T, Y^T) - \frac{1}{T} I(X^T \rightarrow Y^T) \\ = & \frac{1}{T} \log p_{X^T, Y^T}(X^T, Y^T) - \frac{1}{T} H_{P_{X^T, Y^T}}(X^T, Y^T) \\ & - \left(\frac{1}{T} \log \vec{p}_{X^T|Y^T} p_{Y^T}(X^T, Y^T) - \frac{1}{T} H_{\vec{P}_{X^T|Y^T} P_{Y^T}}(X^T, Y^T) \right) \\ \rightarrow & 0 - 0 \quad \text{w.p.1.} \end{aligned}$$

\square

Converse Theorem

Theorem 4.6.2 *The operational channel capacity for a Gaussian channel is less than or equal to the channel capacity: $C^o(K) \leq C(K)$.*

Proof: Assume there exists a sequence of (T, M_T, ϵ_T) channel codes satisfying the power constraint, K , with $\epsilon_T \rightarrow 0$. By Fano's inequality we know

$$\frac{H(\epsilon_T)}{\log M_T} + \epsilon_T \geq 1 - \frac{\frac{1}{T} I(F^T; Y^T)}{\frac{1}{T} \log M_T}.$$

Where $P(F^T)$ is the distribution on code-functions in the T th codebook. (Here $H(\cdot)$ is the binary entropy function: $H(\epsilon) = -\epsilon \log \epsilon - (1 - \epsilon) \log(1 - \epsilon)$.)

Note that proposition 4.3.2 continues to hold for the general alphabet case. Specifically

$$I(F^T; Y^T) = I(X^T \rightarrow Y^T).$$

Thus

$$\frac{H(\epsilon_T)}{\log M_T} + \epsilon_T \geq 1 - \frac{\frac{1}{T} I(X^T \rightarrow Y^T)}{\frac{1}{T} \log M_T} \geq 1 - \frac{C_T(K)}{\frac{1}{T} \log M}.$$

Thus $C^0(K) \leq \liminf_T C_T(K) = C(K)$ is a necessary condition for $\epsilon_T \rightarrow 0$. \square

Direct Theorem

We prove the direct theorem exactly as before. First we construct a code-function distribution so that the induced channel input distribution equals the supremizing channel input distribution. We then use a random coding argument to prove the direct theorem.

There are two ways we can proceed. One approach is to assume finite partitions on \mathcal{X} and \mathcal{Y} . Then consider the Gaussian channel restricted to this partition. We have then reduced the problem to a finite alphabet coding problem and thus can use the results of the previous sections. To prove the Gaussian channel coding theorem we would have to then take a limit over finer and finer partitions.

The second approach is to take advantage of the linear structure of the supremizing channel input distribution. This is how we will proceed here. By lemma 4.6.2 we know that we can restrict our attention to Gaussian channel input distributions.

Definition 4.6.7 *A code-function, f^T , is an affine code-function if for every $t = 1, \dots, T$ it is of the form*

$$f_t(y^{t-1}) = L'_t f^{t-1}(y^{t-2}) + M'_t y^{t-1} + u_t$$

for some vectors L_t, M_t , $t = 2, \dots, T$ of the appropriate dimensions and numbers u_t .

Note that for an affine code-function, f^T , we can find vectors \bar{M}_t, \bar{L}_t such that

$$f_t(y^{t-1}) = \bar{L}'_t y^{t-1} + \bar{M}_t u^t.$$

Any affine code-function, f^T , can be uniquely parameterized by the vectors $\bar{L}_t, t = 2, \dots, T$ and $\bar{M}_t, t = 1, \dots, T$ and the vector u^T . Thus any distribution on affine code-functions can be associated with a distribution on the vectors \bar{L}_t, \bar{M}_t and u^T .

Now we construct the code-function distribution. We are given a channel input distribution defined by $\{X_t = \gamma'_t x^{t-1} + \eta'_t y^{t-1} + W_t\}_{t=1}^T$. We construct a distribution on affine code-functions by choosing $L_t = \gamma_t$ and $M_t = \eta_t$ for all $t = 1, \dots, T$. We let $U_t \sim \mathcal{N}(0, K_{W_t})$, $t = 1, \dots, T$. Clearly the code-function distribution just constructed is *good* with respect to the channel input distribution. (Recall definition 4.4.3.) Note that this construction is very different from the one given in lemma 4.4.3.

We need to show that the induced channel input distribution equals the original channel input distribution.

Lemma 4.6.3 *We are given a channel input distribution defined by $\{X_t = \gamma'_t x^{t-1} + \eta'_t y^{t-1} + W_t\}_{t=1}^T$. Let $P(F^T)$ be a distribution on affine code-functions where $L_t = \gamma_t$ and $M_t = \eta_t$ for all $t = 1, \dots, T$. And $U_t \sim \mathcal{N}(0, K_{W_t})$ for all $t = 1, \dots, T$. Then the induced channel input distribution equals the original channel input distribution.*

Proof: We are given a Gaussian channel. Let $Q(dF^T, dX^T, dY^T)$ be the consistent measure. Because $P(dF^T)$ can be associated with a Gaussian distribution on U^T we know that Q will admit a density with respect to the Lebesgue measure. We can use the same arguments as in lemma 4.4.2 and lemma 4.4.4, with probabilities replaced with densities, to show that the lemma holds. \square

We now show that lemma 4.4.1 on the equivalence of Radon-Nikodym derivatives continues to hold for the case of Gaussian channels and affine code-functions.

Lemma 4.6.4 *We are given a distribution on affine code-functions, $P(dF^T)$, and a Gaussian channel $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$. Let $Q(dF^T, dX^T, dY^T)$ be the consistent joint measure. Then with Q -probability one we have*

$$\frac{dQ_{F^T, Y^T}}{dQ_{F^T} Q_{Y^T}}(F^T, Y^T) = \frac{dQ_{X^T, Y^T}}{d\vec{Q}_{X^T|Y^T} Q_{Y^T}}(X^T, Y^T).$$

Proof: Since $P(dF^T)$ can be associated with a Gaussian distribution over U^T we know that $Q(dF^T, dX^T, dY^T)$ admits a density with respect to the Lebesgue measure. Thus we can follow the steps in lemma 4.4.1 excepting that now we work with densities. Denote the joint density by $q(F^T, X^T, Y^T)$. Then for all (f^T, x^T, y^T) such that $x^T = f^T(y^{T-1})$ we have

$$\begin{aligned} \frac{dQ_{F^T, Y^T}}{dQ_{F^T} Q_{Y^T}}(f^T, y^T) &= \frac{q_{F^T, Y^T}(f^T, y^T)}{q_{F^T} q_{Y^T}(f^T, y^T)} \\ &= \frac{q_{Y^T|F^T}(y^T | f^T)}{q_{Y^T}(y^T)} \\ &= \frac{\prod_{t=1}^T q_{Y_t|Y^{t-1}, F^t}(y_t | f^t, y^{t-1})}{q_{Y^T}(y^T)} \\ &= \frac{\prod_{t=1}^T q_{Y_t|Y^{t-1}, X^t}(y_t | x^t, y^{t-1})}{q_{Y^T}(y^T)} \\ &= \frac{\vec{q}_{Y^T|X^T}(y^T | x^T) \vec{q}_{X^T|Y^T}(x^T | y^T)}{q_{Y^T}(y^T) \vec{q}_{X^T|Y^T}(x^T | y^T)} \\ &= \frac{q_{X^T, Y^T}(x^T, y^T)}{\vec{q}_{X^T|Y^T} q_{Y^T}(x^T, y^T)} \\ &= \frac{dQ_{X^T, Y^T}}{d\vec{Q}_{X^T|Y^T} Q_{Y^T}}(x^T, y^T) \end{aligned}$$

\square

Now we have the necessary lemmas to prove the direct part of the coding theorem. The proof follows lemma 4.4.7 rather closely.

Theorem 4.6.3 *All rates less than $C(K)$ are achievable.*

Proof: Let $P(dF^T)$ be the distribution on affine code-functions constructed as in lemma 4.6.3 for the supremizing Gaussian channel input distribution $\{P(dX_t|x^{t-1}, y^{t-1})\}_{t=1}^T$. Let $Q(dF^T, dX^T, dY^T)$ be the consistent joint measure induced by $P(dF^T)$ and the channel. Let $q(F^T, X^T, Y^T)$ be its density. A channel code is selected at random by drawing M code-functions from $P(dF^T)$. Choose a $\gamma > 0$.

Define

$$\Omega = \{(f^T, b^T) : \frac{1}{T} \log \frac{q_{F^T, B^T}(f^T, b^T)}{q_{F^T} q_{B^T}(f^T, b^T)} \leq \frac{1}{T} \log M + \gamma\}.$$

Then via arguments very similar to those of lemma 4.4.7 with probabilities replaced with densities we have:

$$E_{\text{codebook}(T, M)}(\text{Error}) \leq Q(\Omega) + 2^{-\gamma T}$$

By lemma 4.6.4 the following equality holds for all (f^T, x^T, y^T) such that $x^T = f^T(y^{T-1})$

$$\frac{q_{F^T, Y^T}(f^T, y^T)}{q_{F^T} q_{Y^T}(f^T, y^T)} = \frac{q_{X^T, Y^T}(x^T, y^T)}{\vec{q}_{X^T|Y^T} q_{Y^T}(x^T, y^T)}$$

Thus we have shown

$$E_{\text{codebooks}(T, M)}(\text{Error}) \leq Q_{X^T, Y^T} \left(\frac{1}{T} \log \frac{q_{X^T, Y^T}(X^T, Y^T)}{\vec{q}_{X^T|Y^T} q_{Y^T}(X^T, Y^T)} \leq \frac{1}{T} \log M + \gamma \right) + 2^{-\gamma T}.$$

By lemma 4.6.3 and consistency we have

$$q(X^T, Y^T) = \vec{q}(X^T|Y^T) \vec{p}(Y^T|X^T) = \vec{p}(X^T|Y^T) \vec{p}(Y^T|X^T) = p(X^T, Y^T)$$

thus

$$E_{\text{codebooks}(T, M)}(\text{Error}) \leq P_{X^T, Y^T} \left(\frac{1}{T} \log \frac{p_{X^T, Y^T}(X^T, Y^T)}{\vec{p}_{X^T|Y^T} p_{Y^T}(X^T, Y^T)} \leq \frac{1}{T} \log M + \gamma \right) + 2^{-\gamma T}.$$

By theorem 4.6.1 we know that the AEP holds for Gaussian processes. Thus we have

$$\frac{1}{T} \log \frac{p_{X^T, Y^T}(X^T, Y^T)}{\vec{p}_{X^T|Y^T} p_{Y^T}(X^T, Y^T)} - C_T(K) \rightarrow 0 \text{ w.p.1.}$$

If we choose M_T such that $\frac{1}{T} \log M_T + \gamma < C_T(K)$, $\forall T$ then we see

$$\lim_{T \rightarrow \infty} E_{\text{codebooks}(T, M_T)}(\text{Error}) = 0.$$

Now we examine the power constraint. By construction the average power satisfies

$$E_{\text{codebooks}(T,M)} \left(E_Q \left(\frac{1}{T} \sum X_t^2 \right) \right) \leq K$$

where the inner E_Q is with respect to the Q distribution consistent with the particular channel code chosen. To see this note

$$\begin{aligned} & E_{\text{codebooks}(T,M)} \left(E_Q \left(\frac{1}{T} \sum X_t^2 \right) \right) \\ &= \int \prod_{w=1}^M p_{F^T}(f^T[w]) \frac{1}{M} \sum_{\tilde{w}=1}^M E \left(\frac{1}{T} \sum_{t=1}^T X_t^2 \mid f^T[\tilde{w}] \right) df^T[1], \dots, df^T[M] \\ &= \frac{1}{M} \sum_{\tilde{w}=1}^M \int \prod_{w=1}^M p_{F^T}(f^T[w]) E \left(\frac{1}{T} \sum_{t=1}^T X_t^2 \mid f^T[\tilde{w}] \right) df^T[1], \dots, df^T[M] \\ &= \frac{1}{M} \sum_{\tilde{w}=1}^M \int p_{F^T}(f^T[\tilde{w}]) E \left(\frac{1}{T} \sum_{t=1}^T X_t^2 \mid f^T[\tilde{w}] \right) df^T[\tilde{w}] \\ &= \int p_{F^T}(f^T) E \left(\frac{1}{T} \sum_{t=1}^T X_t^2 \mid f^T \right) df^T \\ &\leq K \end{aligned}$$

By use of the union bound and Markov's inequality we have

$$\begin{aligned} & P_{\text{codebooks}(T,M)} (\text{ave. error} > \epsilon \text{ or ave. power} > K + \epsilon) \\ &\leq P_{\text{codebooks}(T,M)} (\text{ave. error} > \epsilon) + P_{\text{codebooks}(T,M)} (\text{ave. power} > K + \epsilon) \\ &\leq \frac{E_{\text{codebooks}(T,M)}(\text{error})}{\epsilon} + \frac{K}{K + \epsilon} \end{aligned}$$

If $\frac{1}{T} \log M_T + \gamma < C_T$, $\forall T$ then there exists a T large enough so that the last line in the above series of inequalities is less than one.

From this we can conclude that there exists a T and a channel code with M_T code-functions such that the probability of error is less than ϵ and the average power is less than $K + \epsilon$. Since this holds for all $\epsilon > 0$ and all $\gamma > 0$ we have proved the theorem. \square

We end this subsection with some comments on the construction of affine code-functions. First note that the channel code is determined by choosing for each code-function a vector u^T . The code-function at time t can be thought of linearly modulating u^t and the received y^{t-1} . Since both the transmitter and the receiver observe y^{t-1} the design of the code comes down to the determination of u^T . See Klein for more discussion on the structure of the code-functions. [Kle]

4.6.3 Error Exponents

For Gaussian channels there is an error exponent analysis analogous to that given in section 4.4.3. We quickly cite the relevant results. We skip the proofs. They are essentially the same as in subsection 4.4.3 with probabilities replaced with densities.

Definition 4.6.8 Given a Gaussian channel $\{P(dY_t|f^t, y^{t-1})\}_{t=1}^T$ and an affine code-function distribution $P(dF^T)$ define the error exponent to be

$$E_T(R, P(dF^T)) \triangleq \max_{0 \leq \rho \leq 1} \left(-\rho R - \frac{1}{T} \ln \int_{\mathcal{Y}^T} \left[\int_{\mathcal{F}^T} p(f^T) \{p(y^T|f^T)\}^{\frac{1}{1+\rho}} df^T \right]^{1+\rho} dy^T \right).$$

Theorem 4.6.4 The average random coding error over (T, e^{TR}) channel codes drawn according to $P(dF^T)$ can be upperbounded as

$$E_{\text{codebooks}(T, e^{TR})}(\text{error}) \leq e^{-TE_T(R, P(dF^T))}.$$

Definition 4.6.9 The optimal error exponent is

$$E_T(R) \triangleq \sup_{P(F^T)} E_T(R, P(F^T)).$$

Definition 4.6.10 Given a Gaussian channel $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$ and a Gaussian channel input distribution $\{P(dX_t|x^{t-1}, y^{t-1})\}_{t=1}^T$ define the directed error exponent to be

$$\begin{aligned} & \vec{E}_T(R, \{P(dX_t|x^{t-1}, y^{t-1})\}_{t=1}^T) \\ & \triangleq \max_{0 \leq \rho \leq 1} \left(-\rho R - \frac{1}{T} \ln \int_{\mathcal{Y}^T} \left[\int_{\mathcal{X}^T} \vec{p}(x^T|y^T) \{\vec{p}(y^T|x^T)\}^{\frac{1}{1+\rho}} dx^T \right]^{1+\rho} dy^T \right). \end{aligned}$$

Definition 4.6.11 The optimal directed error exponent is

$$\vec{E}_T(R) \triangleq \sup_{\{P(dX_t|x^{t-1}, y^{t-1})\}_{t=1}^T \in \mathcal{S}_T} \vec{E}_T(R, \{P(dX_t|x^{t-1}, y^{t-1})\}_{t=1}^T).$$

Where \mathcal{S}_T is the set of all Gaussian channel input distributions.

Proposition 4.6.1 Fix a Gaussian channel $\{P(dY_t|x^t, y^{t-1})\}_{t=1}^T$. Let the affine code-function distribution $P(dF^T)$ be good with respect to the Gaussian channel input distribution $\{P(dX_t|x^{t-1}, y^{t-1})\}_{t=1}^T$. Then

$$\vec{E}_T(R, \{P(dX_t|x^{t-1}, y^{t-1})\}_{t=1}^T) = E_T(R, P(dF^T)).$$

Corollary 4.6.2 Fix a Gaussian channel $\{P(dY_t|x^t, y^{t-1})\}_{t=1}^T$. Then

$$\vec{E}_T(R) = E_T(R).$$

Definition 4.6.12 *The optimal no feedback directed error exponent is*

$$\vec{E}_T^{nfb}(R) \triangleq \sup_{\{P(dX_t|x^{t-1}, y^{t-1})\}_{t=1}^T \in \mathcal{S}_T^{nfb}} \vec{E}_T \left(R, \{P(dX_t|x^{t-1}, y^{t-1})\}_{t=1}^T \right).$$

Where \mathcal{S}_T^{nfb} is the set of all Gaussian channel input distributions with out feedback. (That is the stochastic kernel $P(dX_t|x^{t-1}, y^{t-1})$ is a function independent of y^{t-1} .)

Proposition 4.6.2 *For a given Gaussian channel $\{P(dY_t|x^t, y^{t-1})\}_{t=1}^T$ we have $\vec{E}_T(R) \geq \vec{E}_T^{nfb}(R)$.*

4.7 Examples

In this section we discuss how our formulation captures and complements some of the existing results in the literature. We also discuss new results. The primary importance of this work, though, is the introduction of a unified formulation for treating a large class of feedback channel coding problems with different forms of memory.

- (1) Kramer uses the tools of directed mutual information and graphical methods to examine the multiple access channel and two-way channel problems. [Kra] We treat the single user case exclusively. In regards to the single user case he only comments on Massey's converse.
- (2) Cover and Pombra prove a feedback coding theorem for Gaussian channels of the form

$$Y_t = x_t + Z_t$$

where Z_t is an arbitrary Gaussian process. [CP] The Gaussian channel defined in definition 4.6.1 is more general than the one proposed by them. The following proposition shows this.

Proposition 4.7.1 *Gaussian channels of the form $Y_t = \alpha_t x^t + Z_t$ where $\{Z_t\}$ is any Gaussian process are contained in definition 3.6.1.*

Proof: Since $\{Z_t\}$ is a Gaussian process we can find vectors $\gamma_t \in \mathbb{R}^{t-1}$, $t > 1$ such that $Z_t = \gamma_t Z^{t-1} + V_t$ for some independent Gaussians V_t .

Let $\bar{\alpha}_t = \alpha_t - \gamma_t[\alpha_1 x_1, \alpha_2 x_2, \dots, \alpha_{t-1} x^{t-1}]'$ and let $\bar{\beta}_t = \gamma_t$. One can verify that upon substitution

$$Y_t = \alpha_t x^t + Z_t = \bar{\alpha}_t x^t + \bar{\beta}_t y^{t-1} + V_t.$$

□

We generalize their coding theorem for Gaussian channels with colored Gaussian noise to Gaussian channels with ISI and colored Gaussian noise. It is an open question as to whether their $\frac{1}{2}$ bit bound for the gain in feedback capacity continues to hold for this ISI channel.

- (3) Goldsmith and Varaiya examine a Gaussian channel with fading modulated by a stationary and ergodic gain sequence. [GV2] In their notation the channel is of the form

$$Y(i) = \sqrt{g(i)} x(i) + N(i)$$

where $\{\sqrt{g(i)}\}$ is a stationary and ergodic process and $\{N(i)\}$ an IID sequence of Gaussian noise.

If we restrict their model to channels where the $\{\sqrt{g(i)}\}$ gain process is Markov and ergodic then after a straightforward combination of the results in section 4.5 and 4.6

we can recover the coding theorem they provide. Specifically we can prove a coding theorem for Markov channels where conditioned on the state the channel looks like an AWGN channel.

We can generalize the coding theorem, for an ergodic Markov gain sequence, to cases where we feedback only channel output information. It is an open question, though, what form water-filling takes in this more general setting.

- (4) Viswanathan examines the case of discrete alphabet Markov channels without ISI. [Vis] In his notation the channel is of the form

$$P(Y_t | x_t, s_t) \quad \text{and} \quad P(S_t | s_{t-1})$$

where $\{S_t\}$ is a finite alphabet Markov state process. He solves the channel coding problem for the case where the encoder observes the state, though possibly delayed. We generalize his result to discrete alphabet Markov channels that include channel output feedback and ISI.

- (5) Shamai and Caire examine finite alphabet channels with memory. [CS] In their proposition 2 they provide the following coding theorem:

Proposition *Assume i) perfect channel state information, S_t , at the receiver; ii) that the receiver knows the channel state information, U_t , at the transmitter: $U_t = g_t(S^t)$; iii) that $P(S_t | U^t) = P(S_t | U_t)$; and iv) that $\{S_t\}, \{U_t\}$ are jointly stationary and ergodic. Then $C = \sum_u P(u) \max_{Q(x|u)} I(X; Y | S, u)$.*

Hypothesis iii) is a bit nonintuitive. Note that our model of a Markov channel satisfies i), ii), and iv) above. In place of iii) we assume that the channel is a Markov channel. We then provide a coding theorem. We believe the state formulation provided in this chapter is the correct formulation for computing capacity of channels with memory.

There are, of course, many other results in the literature. One thing lacking, though, is a general formulation for analyzing channels with memory and differing forms of feedback. We believe that the appropriate way to approach channels with memory and differing forms of feedback is to first *realize* the channel as a Markov channel and then use the techniques developed here to compute the capacity. There is still much work to be done. This chapter is a step towards such a general formulation.




Figure 4-7: Joint Source and Channel

4.8 Channel Realizations and the Directed Data Processing Inequality

In chapter five we will examine the problem of transmitting a process over a channel. Specifically we are interested in combined source-channel encoding and decoding. To that end we define the source-channel encoder and the channel-source decoder. Furthermore we provide a directed version of the data-processing inequality.

Source-Channel Encoder

Let $\{X_t\}_{t=1}^T$ be a source process taking values in \mathbb{R}^d with measure $P(X^T)$. Let $\{Y_t\}_{t=1}^T$ be the source reconstruction process taking values in \mathbb{R}^d . We now discuss the interconnection of the source, the channel and the source reconstruction. As before our channel is $\{P(B_t | a^t, b^{t-1})\}_{t=1}^T$ where now \mathcal{A} and \mathcal{B} can be general Polish spaces. We will usually take these alphabets to be finite spaces or \mathbb{R}^d . See figure 4-7.

We assume the following time ordering on the random variables (X^T, A^T, B^T, Y^T)

$$X_1, A_1, B_1, Y_1, X_2, \dots, X_T, A_T, B_T, Y_T.$$

Definition 4.8.1 A combined source-channel encoder is any family of stochastic kernels $\{Q(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1})\}_{t=1}^T$ such that $Q(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1})$ is independent of y^{t-1} .

Here we allow the channel input symbol A_t to depend on the past source symbols, channel input symbols, and channel output symbols but we do not allow it to depend on the past decoder outputs. Specifically under the joint measure the following is a Markov chain: $A_t - (X^t, A^{t-1}, B^{t-1}) - Y^{t-1}$.

Definition 4.8.2 A combined source-channel encoder without feedback is any family of stochastic kernels $\{Q(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1})\}_{t=1}^T$ such that $Q(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1})$ is independent of (b^{t-1}, y^{t-1}) .

In this case under the joint measure the following is a Markov chain: $A_t - (X^t, A^{t-1}) - (B^{t-1}, Y^{t-1})$.

Channel-Source Decoder

Definition 4.8.3 A combined channel-source decoder is any family of stochastic kernels $\{Q(dY_t | x^t, a^t, b^t, y^{t-1})\}_{t=1}^T$ such that $Q(dY_t | x^t, a^t, b^t, y^{t-1})$ is independent of (x^t, a^t) .

Note that both the source-channel encoder and the channel-source decoder are nonanticipative. Further note the Markov structure of the different elements. These Markov relationships are represented diagrammatically in figure 4-7.

Channel Realization

We are interested in the joint measure $Q(dX^T, dA^T, dB^T, dY^T)$. This measure Q must preserve the underlying channel

$$Q(dB_t | X^t = x^t, A^t = a^t, B^{t-1} = b^{t-1}, Y^{t-1} = y^{t-1}) = P(dB_t | a^t, b^{t-1}) \\ Q(dX^t, dA^t, dB^{t-1}, dY^{t-1}) - a.s.$$

and the underlying source

$$Q(dX_t | X^{t-1} = x^{t-1}, A^{t-1} = a^{t-1}, B^{t-1} = b^{t-1}, Y^{t-1} = y^{t-1}) = P(dX_t | x^{t-1}) \\ Q(dX^{t-1}, dA^{t-1}, dB^{t-1}, dY^{t-1}) - a.s.$$

By theorem A.1.1 the measure Q can be factored as

$$Q(dX^T, dA^T, dB^T, dY^T) \\ = \bigotimes_{t=1}^T Q(dY_t | x^t, a^t, b^t, y^{t-1}) \otimes Q(dB_t | x^t, a^t, b^{t-1}, y^{t-1}) \\ \otimes Q(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1}) \otimes Q(dX_t | x^{t-1}, a^{t-1}, b^{t-1}, y^{t-1}) \\ = \bigotimes_{t=1}^T Q(dY_t | b^t, y^{t-1}) \otimes P(dB_t | a^t, b^{t-1}) \otimes Q(dA_t | x^t, a^{t-1}, b^{t-1}) \otimes P(dX_t | x^{t-1})$$

One term in the above product consists of, going left to right, the channel-source decoder, the channel, the source-channel encoder, and the source.

From this joint measure we can determine the marginal $Q(dX^T, dY^T)$. By theorem A.1.2 we can disintegrate it as follows:

$$Q(dX^T, dY^T) = \bigotimes_{t=1}^T Q(dY_t | x^t, y^{t-1}) \otimes Q(dX_t | x^{t-1})$$

where $Q(dY_t | x^t, y^{t-1})$ is determined $Q(dX^t, dY^{t-1})$ —almost surely. Note that the stochastic kernels, $\{Q(dY_t | x^t, y^{t-1})\}_{t=1}^T$, can also be considered a channel between the source and the reconstruction. The next definition states in what sense one channel can be used to realize another channel.

Definition 4.8.4 We say the channel $\{Q(dB_t | a^t, b^{t-1})\}_{t=1}^T$ is a realization of the channel $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$ if there exists

- (a) a combined source-channel encoder, $\{Q(dA_t | x^t, a^{t-1}, b^{t-1})\}_{t=1}^T$,
- (b) a combined channel-source decoder, $\{Q(dY_t | b^t, y^{t-1})\}_{t=1}^T$,

such that

$$\vec{P}(dY^T | x^T) = \int_{\mathcal{A}^T \times \mathcal{B}^T} \vec{Q}(da^T, db^T, dY^T | x^T).$$

Recall

$$\vec{Q}(dA^T, dB^T, dY^T | x^T) = \bigotimes_{t=1}^T Q(dY_t | b^t, y^{t-1}) \otimes Q(dB_t | a^t, b^{t-1}) \otimes Q(dA_t | x^t, a^{t-1}, b^{t-1}).$$

We say that a realization $\{Q(dB_t | a^t, b^{t-1})\}_{t=1}^T$ of $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$ is a *realization without feedback* if the source-channel encoder used is a source-channel encoder without feedback.

Directed Data Processing Inequality

We now prove a “directed” version of the data processing inequality. As before X_t is the source, A_t is the channel input, B_t is the channel output, and Y_t is the reconstruction.

Lemma 4.8.1 $I(X^T; Y^T) = I(X^T \rightarrow Y^T) \leq I(A^T \rightarrow B^T)$.

Proof: X^T is the source and hence not effected by feedback. The equality follows then from corollary 3.2.3. We prove the inequality in two steps. We first show $I(X^T \rightarrow Y^T) \leq I(X^T \rightarrow B^T)$ and then we show $I(X^T \rightarrow B^T) \leq I(A^T \rightarrow B^T)$.

First note that $X^T - B^T - Y^T$ forms a Markov chain. This is because $Y_t - (B^t, Y^{t-1}) - (X^t, A^t)$ and there is no feedback to the source X^T . Thus

$$I(X^T \rightarrow Y^T) = I(X^T; Y^T) \leq I(X^T; B^T) = I(X^T \rightarrow B^T)$$

where the inequality follows from the *regular* data processing inequality given in proposition A.2. Second note that

$$\begin{aligned} I(X^T \rightarrow B^T) &= \sum_{t=1}^T I(X^t; B_t | B^{t-1}) \\ &\leq \sum_{t=1}^T I(A^t; B_t | B^{t-1}) \\ &= I(A^T \rightarrow B^T) \end{aligned}$$

where we have used the fact that conditioned on B^{t-1} the following $X^t - A^t - B_t$ is a Markov chain. Thus by the data processing inequality, proposition A.3.2, we have $I(X^t; B_t | B^{t-1}) \leq I(A^t; B_t | B^{t-1})$. \square

The directed data processing inequality will be used extensively in chapter five where we deal with the sequential rate distortion problem and the joint source-channel coding problem. It is primarily used to prove the converse theorems.

4.9 Summary

In this chapter we treated the feedback channel coding problem. By using our general formulation from chapter two we were able to prove a general coding theorem for finite alphabet and Gaussian alphabet channels with differing forms of feedback.

To prove these coding theorems we showed that one can convert the feedback channel coding problem to a non-feedback channel coding problem where the input space is now the space of code-functions. We then reduced the underlying optimization problem from one living on code-function space to one living on input distribution space. We showed that the controller, in this case an encoder, completes the joint measure on channel input/output space.

We showed that the directed mutual information, as introduced by Massey, is the correct notion of capacity for channels with and without feedback. The directed mutual information, as opposed to the traditional mutual information, depends on a “causal” factorization of the underlying joint measure.

We provided random coding error exponents for the case of feedback and showed that they can be no worse than the non-feedback error exponents.

For Markov channels we showed that one can solve the capacity optimization problem via dynamic programming. We treat channels with ISI. We argued that Markov channels are the correct way to treat channels with memory.

We compared our results to some of the existing results in the literature. Furthermore we showed how our formulation captures many new results. We argued that there is a need for a general formulation for analyzing channels with memory and differing forms of feedback. This chapter is a step towards such a general formulation.

We concluded with a discussion of channel realizations and provided a causal generalization of the data processing inequality.

Chapter 5

Sequential Rate Distortion

5.1 Introduction

In this chapter we introduce the sequential rate distortion function and its variants. In chapter three we learned that an important goal of the encoder and decoder in control systems under communication constraints is to maintain an error schedule for the state estimation error. For the most part in chapter three we were interested in driving the error to zero. But in cases with additive disturbance, as in proposition 3.5.2, we showed that the objective was to bound the state estimation error. We are now interested in the case where the source and the channel can be stochastic.

Given a stochastic process, $\{X_t\}$, it is natural to ask what is the best approximate representation one can achieve under a rate constraint R . Or conversely for a given approximation what is the minimum rate required to achieve it. Information theory provides a methodology for computing lower bounds on the rate given a distortion constraint. This is called the rate distortion function. Furthermore information theory gives conditions on when one can achieve this lower bound.

Unfortunately traditional information theory is not causal. This chapter is an attempt towards developing a sequential information theory. There are many issues to deal with. One has to be precise as to what assumptions we make on the availability and timing of encoder and decoder knowledge. Furthermore the role of feedback turns out to be very important especially when dealing with unstable processes. To this end we introduce the concept of a sequential rate distortion function. The sequential rate distortion theory tells us how much channel capacity we will need to transmit a process, say video, over a channel so that a distortion criterion at each time step is met.

Computing the rate distortion function entails minimizing a mutual information term over all conditional laws between the input and the output such that a distortion criterion is maintained. For noiseless channels we are actually interested in a deterministic scheme between the input and the output. Thus the mutual information infimization problem can be considered to be a “relaxation” of the problem. There are conditions, most notably large blocklengths, where this “relaxed” solution can actually be achieved. This is Shannon’s rate distortion theorem. We will discuss another condition which we call “channel matching” under which the “relaxed” solution is achievable. Specifically there can exist channels with

capacity R such that the end to end distortion is the distortion rate value. Furthermore we discuss what can happen when the “channel matching” condition does not hold. Finally we will see that the theory of sequential rate distortion is intimately related to the work on feedback coding in chapter four.

In summary there are three main contributions in this chapter. First we prove the direct and converse parts of the sequential rate distortion theorem for noiseless digital channels. Second we show that for Markov sources we can formulate the mutual information optimization problem as a dynamic programming problem with running cost. Third we discuss the notion of *matching* between the source and the channel. This notion was first introduced by Pilc but we expand this idea in many directions. [Pil] We provide a general converse theorem for transmitting a process over a noisy channel. We then find conditions on the noisy channel so that the direct theorem still holds. We discuss the different ways in which one can *realize* the matched channel.

In section 5.2 we state background results on rate distortion. In section 5.3 we introduce the sequential rate distortion problem. Then in section 5.4 we prove a coding theorem for noiseless digital channels. We also provide a dynamic programming formulation for computing the infimizing conditional law. An important simplification of the sequential rate distortion problem leads to the so called successive refinement problem. We discuss successive refinement in section 5.5. We also relate it to channel coding with feedback. In section 5.6 we examine the sequential rate distortion problem for a few cases and discuss some high rate approximations. Finally in section 5.7 we conclude.




Figure 5-1: Traditional Rate Distortion Setup

5.2 Rate Distortion

In this section we discuss background results in rate distortion theory. We discuss the issue of delay in subsection 5.2.2. In subsection 5.2.3 we introduce the idea of induced channel. This will turn out to be important for the rest of this chapter. In subsection 5.2.4 we give some examples. Then in subsection 5.2.5 we discuss the high rate approximation. Finally in subsection 5.2.6 we discuss regimes where we can achieve the rate distortion bound over noiseless digital channels.

5.2.1 Review

We assume that our source takes values in a Polish space \mathcal{X} . See sections A.1 and A.3 of the appendix for a review of Polish spaces and information theoretic quantities. In this chapter our focus will be on the cases where \mathcal{X} is a countable set or \mathbb{R}^d . Specifically let $\{X_t : t = 1, \dots, T\}$ be a set random variables taking values in \mathcal{X} and defined on the probability space $(\Omega, \mathcal{F}, \mathcal{P})$. This process $\{X_t\}$ represents our source. The index t represents a discrete time index.

Definition 5.2.1 A distortion measure $d(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$ is a map taking (x, y) into the nonnegative reals. Let $d_t(x^t, y^t) \triangleq \frac{1}{t} \sum_{i=1}^t d(x_i, y_i)$.

We will need the following boundedness property.

Assumption 5.2.1 We assume there exists an element $x_0 \in \mathcal{X}$ and a positive number D_{\max} such that $E_{P_{X_t}} d(X_t, x_0) < D_{\max} \quad \forall t = 1, \dots, T$.

For the case when $\mathcal{X} = \mathbb{R}^d$ we will often take $x_0 = 0$. Where 0 is the zero vector in \mathbb{R}^d . (Note if $x_0 \neq 0$ and the distortion measure is translation invariant then we can without loss of generality recenter the distribution by translating x_0 to the origin.)

The traditional goal of rate distortion theory is to compress the signals into a smaller set of signals such that a distortion criterion is satisfied. The idea is to quantize the source and then transmit it over a noiseless digital channel at the minimal rate. See figure 5-1.

Definition 5.2.2 A blocklength T quantizer is any measurable map $f : \mathcal{X}^T \rightarrow \Omega \subset \mathcal{X}^T$ where Ω is at most countable.

The superscript “o” in the following represents the word “operational.”

Definition 5.2.3 The operational rate distortion function for horizon T is

$$R_T^o(D) \triangleq \inf_{f \in \mathcal{F}} \frac{1}{T} H(f(X^T))$$

where $\mathcal{F} = \{f : E_{P_{X^T}} d_T(X^T, f(X^T)) \leq D \text{ and } f \text{ is a blocklength quantizer}\}$.

Definition 5.2.4 The operational rate distortion function is

$$R^o(D) \triangleq \limsup_{T \rightarrow \infty} R_T^o(D).$$

We have defined the operational rate distortion functions. Now we define the rate distortion function. In the following let $\mathcal{Y} = \mathcal{X}$. Let $\mathcal{K}_{\mathcal{Y}^T | \mathcal{X}^T}$ represent the space of all stochastic kernels from \mathcal{X}^T to \mathcal{Y}^T . (See section A.1 of the appendix for the definition of stochastic kernels.)

Definition 5.2.5 The rate distortion function for horizon T is

$$R_T(D) \triangleq \inf_{Q \in \Lambda} \frac{1}{T} I_{P_{X^T} \otimes Q}(X^T; Y^T)$$

where $\Lambda = \{Q : Q \in \mathcal{K}_{\mathcal{Y}^T | \mathcal{X}^T} \text{ and } E_{P_{X^T} \otimes Q} d_T(X^T, Y^T) \leq D\}$.

Definition 5.2.6 The rate distortion function is $R(D) \triangleq \limsup_{T \rightarrow \infty} R_T(D)$

We now state Shannon's rate distortion theorem.

Theorem 5.2.1 If the process X_t is stationary and ergodic and satisfies assumption 5.2.1 then for any $\epsilon > 0$ and any $D \geq 0$ such that $R(D) < \infty$ there exists a T large enough and a quantizer f such that $E_{P_{X^T}} d_T(X^T, f(X^T)) \leq D$ and $\frac{1}{T} H(f(X^T)) \leq R(D) + \epsilon$. Conversely, if $R < R(D)$ then there is no T and no quantizer f with $\frac{1}{T} H(f(X^T)) \leq R$ such that $E_{P_{X^T}} d_T(X^T, f(X^T)) \leq D$.

Proof: For the direct part see theorem 7.2.4 of [Berg]. For the converse part see theorem 7.2.5 of [Berg]. \square

A much more general statement of the rate distortion theorem without the ergodicity assumption can be found in [SV] and [Han].

5.2.2 Issue of Delay

Theorem 5.2.1 shows that there exists a coding scheme over noiseless digital channels that achieves the rate distortion bound. But to achieve this bound the delay in computing Y_t from X_t is $2T - t$. At time t it takes $T - t$ steps to observe the remainder of the source, X_{t+1}, \dots, X_T , and then it takes another T steps to transmit the quantized source over a digital noiseless channel. The time-ordering on the random variables is

$$X_1, \dots, X_T, Y_1, \dots, Y_T. \tag{5.1}$$

In control situations this delay may be unacceptable. We seek the minimal rate required to achieve a given distortion without delay. The time-ordering we would like is

$$X_1, Y_1, X_2, \dots, X_T, Y_T. \quad (5.2)$$

The sequential rate distortion problem that we formulate in section 5.3 answers this question. (Note for some control systems some delay or varying delay may be acceptable. See [Sah] for discussion of coding for varying delays.)

We will show that we can often get around this delay if we do not insist on using a noiseless digital channel. Recall our main goal is to transmit the source over a given channel while maintaining a given end-to-end distortion. In general the channel between the source and the recipient is a noisy channel and not a digital noiseless channel. Shannon's separation theorem shows that if the blocklengths are long enough one can transmit the quantization with very small probability of channel error. Specifically, under rather broad conditions, one can design the source encoder and channel encoder separately so as to achieve a given end-to-end distortion. [VVS] But for small delays this separation no longer holds. The following is an example of this.

Example 5.2.1 *Let our source X be a normal $\mathcal{N}(0, 1)$ random variable and the distortion measure be squared error. The rate distortion function for this random variable is $R(D) = \frac{1}{2} \log \frac{1}{D}$ $D \leq 1$. Assume we want to achieve a distortion $D = \frac{1}{4}$. Then $R(D) = 1$.*

If we are given a noiseless digital channel capable of transmitting one bit then the best distortion we can hope to achieve is $\frac{\pi-2}{\pi} > \frac{1}{4}$. (The optimal quantizer can easily be seen to be that which reproduces $\pm\sqrt{\frac{2}{\pi}}$ depending on whether $X > 0$ or not.)

Now assume we are given an AWGN channel of the form $B = A+W$ where W is normal $\mathcal{N}(0, \frac{1}{3})$ and that there is a power constraint of the form $E(A^2) \leq 1$. One can compute the capacity of this channel to be 1.

We will now show that we can transmit our source over this channel and achieve a distortion $\frac{1}{4}$. Let $A = X$. Then upon observing B the best estimate of X can be seen to be $\hat{X} = \frac{3}{4}B$ and the expected square error can be seen to be $\frac{1}{4}$.

We have shown two channels with capacity 1. On one we can achieve the distortion $\frac{1}{4}$ on the other we cannot. The separation result does not hold in this case. Furthermore we were able to achieve the distortion rate value by transmitting the source over a noisy channel without any quantization. Thus in figure 5-1 the digital noiseless channel would be replaced with, in this case, the additive white Gaussian noise channel.

It turns out that the AWGN channel used in the above example is in fact the rate distortion infimizing conditional law. In the next subsection we discuss this induced channel.

5.2.3 Induced Channel and Approximate Factorization

In the last subsection we showed that a noisy channel can be better than a digital noiseless channel when one is interested only in end-to-end distortion. In this subsection we state a

necessary condition on the capacity of the channel to achieve a given end-to-end distortion. We then discuss the induced channel and approximate factorizations.

Theorem 5.2.2 *A necessary condition to achieve end-to-end distortion D over a given channel over a horizon T is $R_T(D) \leq C$.*

Proof: Assume we are given a channel $Q(dB^T | a^T)$. (Recall a channel is a stochastic kernel.) Furthermore assume that this channel has no feedback capacity C . Then for all joint source-channel encoders and joint channel-source decoders such that $E_{P_{X^T, Y^T}} d_T(X^T, Y^T) \leq D$ we have $R_T(D) \leq \frac{1}{T} I(X^T, Y^T) \leq \frac{1}{T} I(A^T; B^T) \leq C$. The second inequality follows from the data processing inequality (proposition A.3.2). \square

Induced Channel

The rate distortion infimizing stochastic kernel, $Q(dY^T | x^T)$, can be viewed as a channel.

Definition 5.2.7 *We call the rate distortion infimizing stochastic kernel the induced channel.*

Recall from chapter 4 that \mathcal{A} and \mathcal{B} are the channel input and output spaces. In this case we let $\mathcal{A} = \mathcal{X}$ and $\mathcal{B} = \mathcal{Y}$ and the source-channel encoder and channel-source decoder are identity maps. (Recall figure 4-7 of chapter 4.) Furthermore the capacity of this induced channel is greater than or equal to $R_T(D)$. To see this note $R_T(D) = I(X^T; Y^T) \leq C$.

The infimizing stochastic kernel can be factored as

$$Q(dY^T | x^T) = \otimes_{t=1}^T Q(dY_t | y^{t-1}, x^T)$$

which has the unfortunate property of being anticipative (with respect to the time-ordering in equation (5.2).) To physically realize this anticipative channel we must introduce delay (i.e. the time-ordering in equation (5.1).) For the channel $Q(dY^T | x^T) = \otimes_{t=1}^T Q(dY_t | y^{t-1}, x^T)$ the delay for determining Y_t from X_t is T . Y_t cannot be produced until X_1, \dots, X_T and Y_1, \dots, Y_{t-1} occur. We comment here that in general $I(X^T; Y^T) \neq I(X^T \rightarrow Y^T)$.

Approximate Factorization

We know that the induced channel achieves the rate distortion bound. We also know that in the limit of large blocklength the noiseless digital channel can be used to achieve the rate distortion bound. We want a way to characterize all channels that achieve the rate distortion bound to within some ϵ .

Definition 5.2.8

$$A_{\epsilon, T}(P_{X^T}) = \left\{ Q(dY^T | x^T) : E_{P_{X^T} \otimes Q} d_T(X^T, Y^T) \leq D \text{ and } \frac{1}{T} I_{P_{X^T} \otimes Q}(X^T; Y^T) \leq R_T(D) + \epsilon \right\}$$

This set is nonempty because the rate distortion infimizing law is a member of it. Shannon's rate distortion theorem states that under suitable conditions and T large enough the set $A_{\epsilon,T}$ will contain a noiseless digital channel of rate $R(D) + \epsilon$.

Definition 5.2.9 A source, P_{X^T} , and a channel, $Q(dB^T | a^T)$, are said to be matched if $Q(dB^T | a^T)$ is a realization of the rate distortion infimizing stochastic kernel $Q(dY^T | x^T)$. The source and channel are said to be approximately matched if $Q(dB^T | a^T) \in A_{\epsilon,T}(P_{X^T})$. (Recall definition 4.8.5.)

We call a matched channel, $Q(dB^T | a^T)$, a *factor*, because as we have shown in section 4.8 the induced channel $P(dY^T | x^T)$ can be factored as

$$P(dY^T | x^T) = \int_{A^T \times B^T} Q(dY^T | b^T) Q(dB^T | a^T) Q(dA^T | x^T)$$

where $Q(dA^T | x^T)$ and $Q(dY^T | b^T)$ are the encoder and decoder respectively. By analogy an approximately matched channel is called an *approximate factor*.

There can be many realizations and approximate realizations. Often times we can realize a particular channel with much simpler channels. We discuss this in the examples section. Furthermore our interest in matched channels comes from the fact that they will help us compute closed form solutions for some classes of sources. Determining a matched channel is a crucial part in the computation of the sequential rate distortion function for Gaussian sources.

It is unfortunate that in general the induced channel forces a delay of T . But for the case when the source is independent (i.e. $P(X^T) = \otimes_{t=1}^T P(X_t)$) it can be shown that the infimizing law is nonanticipative. Specifically

Lemma 5.2.1 If $\{X_t\}_{t=1,\dots,T}$ are independent then the stochastic kernel that infimizes $R_T(D)$ has the form $Q(dY^T | x^T) = \otimes_{t=1}^T Q(dY_t | x_t)$.

Proof:

$$\begin{aligned} I(X^T; Y^T) &= \sup_{\Pi} \sum_{\pi \in \Pi} P_{X^T, Y^T}(\pi) \log \frac{P_{X^T, Y^T}(\pi)}{P_{X^T} P_{Y^T}(\pi)} \\ &\geq \sum_{\pi \in \Pi} P_{X^T, Y^T}(\pi) \log \frac{P_{X^T, Y^T}(\pi)}{P_{X^T} P_{Y^T}(\pi)} \\ &\geq \sum_{\pi \in \Pi} \sum_{t=1}^T P_{X_t, Y_t}(\pi) \log \frac{P_{X_t, Y_t}(\pi)}{P_{X_t} P_{Y_t}(\pi)} \end{aligned}$$

Where the first equality follows from theorem A.3.2. The first inequality follows because Π can be any partition (not necessarily the supremizing partition.) The third inequality follows from theorem 9.2.1 of Gallager's text. [Gal] Since the inequality holds for all partitions it

must be the case that

$$I(X^T; Y^T) \geq \sup_{\Pi} \sum_{\pi \in \Pi} \sum_{t=1}^T P_{X_t, Y_t}(\pi) \log \frac{P_{X_t, Y_t}(\pi)}{P_{X_t} P_{Y_t}(\pi)} = \sum_{t=1}^T I(X_t; Y_t).$$

Equality is achieved if $Q(dY^T | x^T) = \otimes_{t=1}^T Q(dY_t | x_t)$. Furthermore $Ed(X_t, Y_t)$ depends only on the marginal $P(dX_t, dY_t)$. Thus given any measure we can replace it with its product form. By doing so the distortion remains the same and the mutual information cannot increase. \square

5.2.4 Examples

In this subsection we examine two sources.

- (a) $\{X_t, t = 1, \dots, T\}$ is Gaussian $\mathcal{N}(0, \Lambda)$.
- (b) $\{X_t, t = 1, \dots, T\}$ is uniformly distributed over the box $[-L, L]^T$.

Definition 5.2.10 Given a distortion measure d its semi-faithful version for distortion D is defined as

$$\bar{d}_D(X, Y) = \begin{cases} d(X, Y) & \text{if } d(X, Y) \leq D \\ +\infty & \text{if } d(X, Y) > D \end{cases}$$

For the Gaussian source we use the following *squared error* distortion measure:

$$d(x^T, y^T) = \|x^T - y^T\|_M^2 = (x^T - y^T)' M (x^T - y^T)$$

where M is a positive definite weighting matrix. For the uniform source we use the semi-faithful version of the squared difference distortion measure.

Before proceeding we prove some simplifying properties.

Lemma 5.2.2 Let $\tilde{X}^T = M^{\frac{1}{2}} X^T, \tilde{Y}^T = M^{\frac{1}{2}} Y^T$. Then

- (1) $I(\tilde{X}^T; \tilde{Y}^T) = I(X^T; Y^T)$
- (2) $\|\tilde{x}^T - \tilde{y}^T\|_I^2 = \|x^T - y^T\|_M^2$ where I is the identity matrix.

Proof: (1) holds because mutual information is invariant under injective transformations (proposition A.3.1 (c).) To see (2) note

$$\begin{aligned} \|\tilde{x}^T - \tilde{y}^T\|_I^2 &= (M^{\frac{1}{2}} x^T - M^{\frac{1}{2}} y^T)' (M^{\frac{1}{2}} x^T - M^{\frac{1}{2}} y^T) \\ &= (x^T - y^T)' M (x^T - y^T) \\ &= \|x^T - y^T\|_M^2 \end{aligned}$$

\square

Thus without loss of generality we can restrict our attention to the case where $M = I$. In practice the encoder would preprocess the observation by applying $M^{\frac{1}{2}}$ to it. Similarly the decoder would postprocess its output by $M^{-\frac{1}{2}}$.

Lemma 5.2.3 *Let U be the unitary matrix that diagonalizes $E_{P_{X^T}} X^T X^{T'} = U' \Gamma U$. Let $\tilde{X}^T = UX^T, \tilde{Y}^T = UY^T$. Then*

- (1) $I(\tilde{X}^T; \tilde{Y}^T) = I(X^T; Y^T)$
- (2) $\|\tilde{x}^T - \tilde{y}^T\|_I^2 = \|x^T - y^T\|_I^2$.

Proof: Both (1) and (2) hold because mutual information and the squared error distortion with weight matrix I are invariant under unitary transformations. \square

Thus, without loss of generality, we can restrict our attention to the case where the source covariance $E_{P(X^T)} X^T X^{T'}$ is diagonal. In practice the encoder will preprocess the observation by applying U to it. Similarly the decoder will postprocess its output by U'

Gaussian Source

It is straightforward to compute the rate distortion function for the Gaussian source. By lemmas 5.2.2 and 5.2.3 we can, without loss of generality, restrict ourselves to Gaussians with covariance $\Lambda = \text{diag}[\lambda_1, \dots, \lambda_T]$ and a squared error distortion measure with weight matrix I .

By equation 9.7.41 of [Gal] we have

$$R(D) = \frac{1}{T} \sum_{i=1}^T \frac{1}{2} \log \frac{\lambda_i}{\delta_i}$$

where

$$\delta_i = \begin{cases} \eta & \text{if } \eta \leq \lambda_i \\ \lambda_i & \text{if } \eta > \lambda_i \end{cases}$$

where η is chosen such that $\sum_{i=1}^T \delta_i = D$. This is the so called water-filling solution. For D small enough one can show the above formula reduces to $R(D) = \frac{1}{2} \log \frac{T|\Lambda|^{\frac{1}{T}}}{D}$.

We now characterize the infimizing law. Following equation 9.7.16 of [Gal] we see that the backward channel has the form

$$X^T = Y^T + V^T$$

where V^T is distributed normally with mean zero and covariance $\text{diag}[\delta_1, \dots, \delta_T]$. Thus the forward channel has the form

$$Y^T = HX^T + W^T$$

where $H = E(Y^T X^{T'})E(X^T X^{T'})^{-1}$ and W^T is a zero mean Gaussian vector with covariance $E(Y^T Y^{T'}) - E(Y^T X^{T'})E(X^T X^{T'})^{-1}E(X^T Y^{T'})$. This forward channel is the infimizing law.

There are many ways to realize the channel $Y^T = HX^T + W^T$. For example let G be any invertible transformation. Define $B^T = GY^T$, $A^T = GX^T$. Then the channel $B^T = GHG^{-1}A^T + GW^T$ also realizes the infimizing law with encoder $A^T = GX^T$ and decoder $Y^T = G^{-1}B^T$. Every such G represents a coordinate transformation and thus induces a different matched channel to the Gaussian source. Note that these realizations are a realizations without feedback (recall definition 4.8.5 of chapter four.)

Uniform Source

We assume that we have a uniform source with the semi-faithful version of the mean squared distortion as its distortion. It is difficult to compute the infimizing law in this case. We can, though, compute an upper bound on the rate for a noiseless digital channel. This in turn is an upper bound to the rate distortion function for the uniform source.

We can upper bound the rate distortion function as the logarithm of the ratio of the total volume of the box to the volume of a box of side-length $\frac{2D}{\sqrt{T}}$. Then

$$R(D) \leq \frac{1}{T} \log \frac{(2L)^T}{\left(\frac{2D}{\sqrt{T}}\right)^T} = \log(\sqrt{T} \frac{L}{D}).$$

5.2.5 High Rate Approximation

Now we show that as the distortion goes to zero we can transmit the Gaussian source over a digital noiseless channel at essentially the rate distortion limit for arbitrary T . We use the uniform source results to approximate the Gaussian source. As $D \rightarrow 0$ we can approximate the infimizing law by a digital law.

We first analyze the scalar case and then generalize it to the vector case. Note the approximations we make now are in no way optimal. They serve only to show that there exists a coding scheme such that the ratio of the rate distortion function and the rate of this particular scheme goes to one. To get better results on the optimal rate of convergence one should apply Bennett's distortion integral. [Ben]

Let $X \sim \mathcal{N}(0, \lambda)$. Recall we are using the squared error distortion criterion. The rate distortion function for this source is $R(D) = \max\{0, \frac{1}{2} \log \frac{\lambda}{D}\}$. There is nothing to do if $D \geq \lambda$. Thus we assume $D < \lambda$. We propose the following scheme: choose a finite interval and partition it into regions of size $2\sqrt{D}$. If X falls inside this region we will achieve a distortion $\leq D$. We need to bound the average distortion when conditioned on falling outside this region. In particular let $[-L, L]$ be the region in which we quantize the source. L represents the dynamic range. If X falls outside of this region then send the zero signal.

Let the superscript s represent “scheme.” Thus

$$\begin{aligned}
D^s &\leq \left(\frac{L}{2R}\right)^2 + 2 \int_L^\infty x^2 p(x) dx \\
&= \left(\frac{L}{2R}\right)^2 + 2\sqrt{\frac{\lambda}{2\pi}} L e^{-\frac{L^2}{2\lambda}} + 2\lambda \Pr(X \geq L) \\
&\leq \left(\frac{L}{2R}\right)^2 + 2\sqrt{\frac{\lambda}{2\pi}} L e^{-\frac{L^2}{2\lambda}} + 2\lambda e^{-\frac{L^2}{2\lambda}} \\
&= \left(\frac{L}{2R}\right)^2 + \left(2\sqrt{\frac{\lambda}{2\pi}} L + 2\lambda\right) e^{-\frac{L^2}{2\lambda}}
\end{aligned}$$

where the second line follows by integration by parts (see lemma A.4.1) and the third line follows from the Chernoff bound.

We need to choose L and R such that this distortion is always less than D and that the rate is close to the rate distortion rate in the limit of small distortion. We will do this by choosing an L such that the second term is less than $\frac{D}{2}$. Then for this L we will choose an R such that the first term is less than $\frac{D}{2}$.

It should be clear that if $D \rightarrow 0$ then it must be the case that the dynamic range $L \rightarrow \infty$. Furthermore there exists an \tilde{L} large enough so that for all $L > \tilde{L}$ we have $\left(2\sqrt{\frac{\lambda}{2\pi}} L + 2\lambda\right) e^{-\frac{L^2}{2\lambda}} \leq e^{-\frac{L^2}{4\lambda}}$. Let $L^s(D) = \sqrt{4\lambda \ln\left(\frac{2}{D}\right)}$. Assume that D is small enough so that $L^s(D) > \tilde{L}$. Then $\left(2\sqrt{\frac{\lambda}{2\pi}} L^s(D) + 2\lambda\right) e^{-\frac{L^s(D)^2}{2\lambda}} \leq e^{-\frac{L^s(D)^2}{4\lambda}} = \frac{D}{2}$.

Then choose R so that $\left(\frac{L^s(D)}{2R}\right)^2 \leq \frac{D}{2}$. One choice is $R^s(D) = \frac{1}{2} \log\left(\frac{2(L^s(D))^2}{D}\right)$. (As usual we ignore the fact that R should be an integer.) By substituting for $L^s(D)$ we can compute $R^s(D) = \frac{1}{2} \log\left(\frac{8\lambda \ln\left(\frac{2}{D}\right)}{D}\right)$.

For D small enough we have by construction $D^s \leq D$. Now we need to show that $\lim_{D \rightarrow 0} \frac{R^s(D)}{R(D)} = 1$.

$$\begin{aligned}
\frac{R^s(D)}{R(D)} &= \frac{\frac{1}{2} \log\left(\frac{8\lambda \ln\left(\frac{2}{D}\right)}{D}\right)}{\frac{1}{2} \log \frac{\lambda}{D}} \\
&= 1 + \frac{\log\left(8 \ln\left(\frac{2}{D}\right)\right)}{\log \frac{\lambda}{D}}
\end{aligned}$$

Because of the $\log \ln$ term in the numerator the second term goes to zero as $D \rightarrow 0$. Thus in the limit of small D we can approximate the infimizing law by a digital channel with the same distortion and rate.

Now we generalize to the vector source case. Let X_1, \dots, X_N be distributed $\mathcal{N}(0, \Lambda)$. Assume without loss of generality that $\Lambda = \text{diag}[\lambda_1, \dots, \lambda_N]$. For D small enough $\delta_i =$

$\eta = \frac{D}{N}$ in the rate distortion formula. For each coordinate let $L_n^s(D) = \sqrt{4\lambda_n \ln(\frac{2N}{D})}$ and $R_n^s(D) = \frac{1}{2} \log(\frac{8N\lambda_n}{D} \ln(\frac{2N}{D}))$.

The scheme is to partition the box $\{[-L_1^s(D), L_1^s(D)] \times \dots \times [-L_N^s(D), L_N^s(D)]\}$ into smaller boxes (i.e. a primitive quantizer in the notation of chapter three.) In the n th coordinate direction we divide the interval into $2^{R_n^s(D)}$ smaller intervals. By our results for the scalar case the average distortion along any coordinate axis will be less than $\frac{D}{N}$. Thus the total distortion is less than D .

The ratio of the rates goes to one as

$$\begin{aligned} \frac{R^s(D)}{R(D)} &= \frac{\frac{1}{2N} \sum_{n=1}^N \log(\frac{8N\lambda_n}{D} \ln(\frac{2N}{D}))}{\frac{1}{2N} \sum_{n=1}^N \log \frac{N\lambda_n}{D}} \\ &= 1 + \frac{\sum_{n=1}^N \log(8 \ln(\frac{2N}{D}))}{\sum_{n=1}^N \log \frac{N\lambda_n}{D}} \end{aligned}$$

The second term goes to zero as $D \rightarrow 0$.

Note that similar results hold for other real-valued random variables. We do not have to limit ourselves to Gaussian sources. We only require that the source have mass at the tail, $\int_L^\infty x^2 p(x) dx$, decreasing sufficiently fast with respect to L .

5.2.6 Summary

In this section we have reviewed the rate distortion problem. We showed that if delay is an issue then the noiseless digital channel may not be the optimal channel over which to transmit the source.

We have shown two regimes in which the noiseless digital channel is approximately matched to the induced channel:

- (1) D fixed, $T \rightarrow \infty$, and the process $\{X_t\}$ ergodic. This is Shannon's rate distortion theorem. There is a delay of T units.
- (2) T fixed and $D \rightarrow 0$. This was shown in subsection 5.2.5. There is no delay in this case.

In terms of the time-ordering described in equation (5.2) the infimizing law may be anticipative. We address this problem in the next section by introducing the sequential rate distortion problem.

5.3 Sequential Rate Distortion

In this section we formulate the sequential rate distortion problem. The basic idea is to transmit a lossy version of a process, say video, over a channel while maintaining some distortion criterion. A typical control application would be a sensor that is a video camera connected to a controller via a noisy communication channel.

We first define the sequential rate distortion function. And then provide a general converse theorem. In section 5.4 we provide the direct part of the coding theorem.

5.3.1 Setup

We define the following two-parameter source. (As before \mathcal{X} is a Polish space.) Let $\{X_{t,n} : t = 1, \dots, T, n = 1, \dots, N\}$ be a set random variables taking values in \mathcal{X} and defined on the probability space $(\Omega, \mathcal{F}, \mathcal{P})$. This process $\{X_{t,n}\}$ represents our source. The index t represents a discrete time index. The index n represents a spatial index. Thus at time t we observe the random variables $X_{t,1}, \dots, X_{t,N}$. For example the index (t, n) can represent a particular element, n , of a raster scan of the t th image in a video stream. We use the notation $X_t^n \triangleq (X_{t,1}, \dots, X_{t,n})$, $X_n^t \triangleq (X_{1,n}, \dots, X_{t,n})$, and $X^{t,n} \triangleq \{X_{i,j}\}_{i=1, \dots, t, j=1, \dots, n}$. (Note that there is a natural ordering in time but not in space.)

We will generalize assumption 5.2.1 for this new source.

Assumption 5.3.1 *We assume there exists an element $x_0 \in \mathcal{X}$ and a positive number D_{\max} such that $E_{P_{X_{n,t}}} d(X_{n,t}, x_0) < D_{\max} \quad \forall n = 1 \dots N, t = 1 \dots T$.*

Sequential Rate Distortion Quantizer

Definition 5.3.1 *A sequential rate distortion quantizer is a sequence of measurable functions f_t such that*

$$f_t : \mathcal{X}^{t,N} \times \mathcal{Y}^{(t-1),N} \rightarrow \mathcal{Y}^N$$

where the range of each function is at most countable. Specifically f_t takes $(x^{t,N}, y^{t-1,N}) \mapsto y_t^N$.

Lemma 5.3.1 *A source distribution, $P_{X^{T,N}}$ and a sequential rate distortion quantizer, f_1, \dots, f_T , specifies a unique measure $P_{X^{T,N}, Y^{T,N}}$.*

Proof: Note that the stochastic kernel $Q(dY_t^N | x^{T,N}, y^{t-1,N}) = \delta_{\{Y_t^N = f_t(x^{t,N}, y^{t-1,N})\}}$. Then by theorem A.1.1 the following measure exists

$$P(dX^{T,N}, dY^{T,N}) = P(dX^{T,N}) \otimes \left\{ \bigotimes_{t=1}^T Q(dY_t^N | x^{T,N}, y^{t-1,N}) \right\}.$$

□

Operational Sequential Rate Distortion

We formulate two forms of the operational sequential rate distortion function for noiseless digital channels. In the first one we use a distortion schedule. In the second one we use a time-averaged distortion. The superscript “SRD” represents “sequential rate distortion” and the “o” represents “operational.”

Definition 5.3.2 Formulation 1

The operational sequential rate distortion function under a distortion schedule is

$$R_{T,N}^{SRD, o}(D_1, \dots, D_T) = \inf_{(f_1, \dots, f_T) \in \mathcal{F}} \frac{1}{NT} H(Y_1^N, \dots, Y_T^N)$$

where $\mathcal{F} = \{(f_1, \dots, f_T) : E_{P_{X^t, N}} d_N(X_t^N, Y_t^N) \leq D_t \quad t = 1, \dots, T\}$.

Definition 5.3.3 Formulation 2:

The operational sequential rate distortion under a time-average distortion constraint is

$$R_{T,N}^{SRD, o, ave}(D) = \inf_{(f_1, \dots, f_T) \in \mathcal{F}} \frac{1}{NT} H(Y_1^N, \dots, Y_T^N)$$

where $\mathcal{F} = \{(f_1, \dots, f_T) : \frac{1}{T} \sum_{t=1}^T E_{P_{X^t, N}} d_N(X_t^N, Y_t^N) \leq D\}$.

In both formulations the expectation of the distortion depends on $P_{X^t, N}$ and not just $P_{X_t^N}$. This is because Y_t^N is a function of $X^{t, N}$. (Note, though, that $E_{P_{X^t, N}} d_N(X_t^N, Y_t^N) = E_{P_{X_t^N, Y_t^N}} d_N(X_t^N, Y_t^N)$.)

Sequential Rate Distortion

As discussed in the last section the rate distortion problem consists of minimizing the mutual information over all stochastic kernels satisfying some distortion constraint. For the sequential rate distortion problem we will take the minimization to be over a “causal” sequence of stochastic kernels. We define that now.

Definition 5.3.4 A sequence of stochastic kernels $\{Q(dY_t^N | y^{t-1, N}, x^{t, N})\}_{t=1}^N$ is called a causal sequence of stochastic kernels if $\forall t = 1, \dots, T$

$$Q(dY_t^N | y^{t-1, N}, x^{t, N}, x_{t+1}^N, \dots, x_T^N) = Q(dY_t^N | y^{t-1, N}, x^{t, N}, \tilde{x}_{t+1}^N, \dots, \tilde{x}_T^N)$$

$\forall x_{t+1}^N, \dots, x_T^N, \tilde{x}_{t+1}^N, \dots, \tilde{x}_T^N$. A causal sequence of stochastic kernels is a channel and thus denoted $\{Q(dY_t^N | y^{t-1, N}, x^{t, N})\}_{t=1}^N$.

Lemma 5.3.2 A source, $P(dX^{T, N})$, and a causal sequence of stochastic kernels, $\{Q(dY_t^N | y^{t-1, N}, x^{t, N})\}_{t=1}^N$, uniquely determines a joint measure $P(dX^{T, N}, dY^{T, N})$.

Proof: Follows from theorem A.1.1. Specifically

$$\begin{aligned} P(dX^{T,N}, dY^{T,N}) &= \bigotimes_{t=1}^T P(dY_t^N | x^{t,N}, y^{t-1,N}) \otimes P(dX_t^N | x^{t-1,N}, y^{t-1,N}) \\ &= \bigotimes_{t=1}^T Q(dY_t^N | x^{t,N}, y^{t-1,N}) \otimes P(dX_t^N | x^{t-1,N}) \end{aligned}$$

Where the last line follows because the source is not affected by the past reconstructions. \square

We now state the sequential rate distortion function formulation. This is an optimization of the mutual information over all infimizing “channels.” We want the infimizing stochastic kernel to be nonanticipative.

Definition 5.3.5 Formulation 1:

The sequential rate distortion function under a distortion schedule *is*

$$R_{T,N}^{SRD}(D_1, \dots, D_T) = \inf_{\mathcal{F}} \frac{1}{NT} I_{P_{X^T,N}, Y^T,N}(X_1^N, \dots, X_T^N; Y_1^N, \dots, Y_T^N)$$

where $\mathcal{F} = \{\{Q(dY_t^N | y^{t-1,N}, x^{t,N})\}_{t=1}^T : E_{P_{X_t^N, Y_t^N}} d_N(X_t^N, Y_t^N) \leq D_t \ t = 1, \dots, T \text{ and it is a causal sequence of stochastic kernels}\}$.

Definition 5.3.6 Formulation 2:

The sequential rate distortion function under a time-average distortion constraint *is*

$$R_{T,N}^{SRD, ave}(D) = \inf_{\mathcal{F}} \frac{1}{NT} I_{P_{X^T,N}, Y^T,N}(X_1^N, \dots, X_T^N; Y_1^N, \dots, Y_T^N)$$

where $\mathcal{F} = \{\{Q(dY_t^N | y^{t-1,N}, x^{t,N})\}_{t=1}^T : \frac{1}{T} \sum_{t=1}^T E_{P_{X_t^N, Y_t^N}} d_N(X_t^N, Y_t^N) \leq D \text{ and it is a causal sequence of stochastic kernels}\}$.

A similar formulation of the sequential rate distortion problem can be found in the work of Pinsker. [GP] Our work differs in many ways. In particular we will give an operational meaning to the sequential rate distortion function. We will provide a coding theorem and high rate asymptotics.

Sequential Rate Distortion: The Intuition

The notion of causality makes sense only after we have defined a specific time-ordering on the random variables of interest. As already stated in section 4.3.4 the joint measure tells us nothing about causality. The time-ordering we are interested in is

$$X_1^N, Y_1^N, X_2^N, \dots, X_T^N, Y_T^N.$$

Now given the joint measure $P(dX^{T,N}, dY^{T,N})$ we can factor it as

$$P(dX^{T,N}, dY^{T,N}) = \bigotimes_{t=1}^T P(dY_t^N | x^{t,N}, y^{t-1,N}) \otimes P(dX_t^N | x^{t-1,N}, y^{t-1,N}).$$

The source, though, is not allowed to be affected by the past reconstructions. That is $X_t^N - X^{t-1,N} - Y^{t-1,N}$ forms a Markov chain. Thus $P(dX_t^N | x^{t-1,N}, y^{t-1,N})$ is independent of $y^{t-1,N}$. Therefore we see that the induced channel, $\{P(dY_t^N | x^{t,N}, y^{t-1,N})\}_{t=1}^T$, is a channel used without feedback. Thus

$$I(X^{T,N} \rightarrow Y^{T,N}) \triangleq \sum_{t=1}^T I(X^{t,N}; Y_t^N | Y^{t-1,N}) = I(X^{T,N}; Y^{T,N}).$$

Note that though the induced channel, $\{P(dY_t^N | x^{t,N}, y^{t-1,N})\}_{t=1}^T$, is used without feedback it may be the case that a realization of this channel does use feedback. In fact many channels with memory can be realized by simpler channels, i.e. memoryless channels, with feedback. This is discussed in more detail in section 5.6.

5.3.2 Converse Theorem

In this section we provide converse theorems for the sequential rate distortion problem.

Theorem 5.3.1 $R_{T,N}^{SRD, o}(D_1, \dots, D_T) \geq R_{T,N}^{SRD}(D_1, \dots, D_T)$ and $R_{T,N}^{SRD, o, ave}(D) \geq R_{T,N}^{SRD, ave}(D)$.

Proof: We prove the theorem for formulation 1. Formulation 2 follows analogously. Let f_1, \dots, f_T be any sequential quantizer such that $E_{P_{X_t^N, N}} d_N(X_t^N, Y_t^N) \leq D_t$ $t = 1, \dots, T$. Then

$$\begin{aligned} \frac{1}{NT} H(Y_1^N, \dots, Y_T^N) &\geq \frac{1}{NT} I(X_1^N, \dots, X_T^N; Y_1^N, \dots, Y_T^N) \\ &\geq R_{T,N}^{SRD}(D_1, \dots, D_T) \end{aligned}$$

□

Theorem 5.3.1 gives us a necessary condition on the capacity of a digital channel in order to achieve the operational rate distortion rate. The next lemma shows that the necessary condition continues to hold for all channels.

Assume we are given a channel $\{Q(dB_t | a^t, b^{t-1})\}_{t=1}^T$. We provide a necessary condition on the capacity of the channel, C_T , to achieve an end-to-end distortion criterion.

Theorem 5.3.2 *A necessary condition to achieve end-to-end distortion (D_1, \dots, D_T) in formulation one over a given channel over a time horizon T is $R_{T,N}^{SRD}(D_1, \dots, D_T) \leq \frac{1}{N} C_T$.*

Similarly a necessary condition to achieve end-to-end distortion D in formulation two over a given channel over a time horizon T is $R_{T,N}^{\text{SRD}, \text{ave}}(D) \leq \frac{1}{N}C_T$.

Proof: We are given a source $P(dX^{T,N})$. Now given any channel, $\{Q(dB_t | b^{t-1}, a^t)\}_{t=1}^T$, source-channel encoder, $\{Q(dA_t | x^{t,N}, a^{t-1}, b^{t-1}, y^{t-1,N})\}_{t=1}^T$, and channel-source decoder, $\{Q(dY_t^N | b^t, y^{t-1,N})\}_{t=1}^T$, such that $E_{P_{X_t, Y_t}} d_N(X_t^N, Y_t^N) \leq D_t \quad t = 1, \dots, T$ we have

$$\begin{aligned} R_{T,N}^{\text{SRD}}(D_1, \dots, D_T) &\leq \frac{1}{NT} I(X^{T,N}; Y^{T,N}) \\ &= \frac{1}{NT} I(X^{T,N} \rightarrow Y^{T,N}) \\ &\leq \frac{1}{NT} I(A^T \rightarrow B^T) \\ &\leq \frac{1}{N} C_T \end{aligned}$$

where we have used the directed information data-processing inequality (lemma 4.8.1 of chapter four.) The proof for formulation two follows analogously. (The $\frac{1}{N}C_T$ comes from the fact that the sequential rate distortion function measure rate per spatial dimension.) \square

The idea of using the directed data-processing inequality for computing the lower bound for the sequential rate distortion function is a special example of a more general lower bounding technique. This technique, discussed in [Mit], shows that the lower bound can be computed as a maximization of a mutual information representing a channel capacity calculation and a minimization of a mutual information representing a rate distortion calculation.

5.4 Coding Theorem for the Noiseless Digital Channel

In this section we provide a coding theorem for transmitting sources over a noiseless digital channel. We also provide a dynamic programming formulation for computing the infimizing law.

In the last section we proved a rather general converse theorem. We prove the direct theorem under the following IID assumption.

Assumption 5.4.1 *The source $\{X_{t,n}\}_{n=1,\dots,N,t=1,\dots,T}$ has the following independence structure:*

$$P(dX^{T,N}) = \otimes_{n=1}^N P(dX_n^T).$$

Furthermore the $\{X_n^T\}_{n=1}^N$ are all identically distributed.

This states that for each n the trajectory in time, $\{X_{1,n}, \dots, X_{T,n}\}$, is independent of all other trajectories $\{X_{1,\tilde{n}}, \dots, X_{T,\tilde{n}}\}$ $\tilde{n} \neq n$. Furthermore each trajectory is identically distributed.

Lemma 5.4.1 *Under assumption 5.4.1, the causal sequence of stochastic kernels that infimizes $R_{T,N}^{SRD}(D_1, \dots, D_T)$ and $R_{T,N}^{SRD, ave}(D)$ factors as*

$$Q(dy_t^N | y^{t-1,N}, x^{t,N}) = \otimes_{n=1}^N Q(dy_{t,n} | y_n^{t-1}, x_n^t) \quad \forall t = 1, \dots, T.$$

Consequently $R_{T,1}^{SRD}(D_1, \dots, D_T) = R_{T,N}^{SRD}(D_1, \dots, D_T)$ and $R_{T,1}^{SRD, ave}(D) = R_{T,N}^{SRD, ave}(D)$.

Proof:

$$\begin{aligned} I(X^{T,N}; Y^{T,N}) &= \sum_{t=1}^T I(X^{t,N}; Y_t^N | Y^{t-1,N}) \\ &\geq \sum_{t=1}^T \sum_{n=1}^N I(X_n^t; Y_{t,n} | Y_n^{t-1}) \\ &= \sum_{n=1}^N I(X_n^T; Y_n^T) \end{aligned}$$

where the inequality follows from a straightforward extension of lemma 5.2.1. We have equality when $Q(dy_t^N | y^{t-1,N}, x^{t,N}) = \otimes_{n=1}^N Q(dy_n^t | y_n^{t-1}, x_n^t) \quad \forall t = 1, \dots, T$. Furthermore $Ed(X_{n,t}, Y_{n,t})$ only depends on the measure $P_{X_{t,n}, Y_{t,n}}$. Thus given any measure we can replace it with its product form. By doing so the distortion remains the same and the mutual information cannot increase. \square

By lemma 5.4.1 we see that we can restrict our attention to the single-letter forms of $R_T^{SRD}(D_1, \dots, D_T)$ and $R_T^{SRD, ave}(D)$. (Where the “single-letter” is with respect to the spatial index n .) Thus when computing the infimizing conditional law we need only look at sources of the form P_{X^T} (i.e. $N = 1$.)

5.4.1 Finite Alphabet Source Coding Theorem

In this subsection we prove the direct coding theorem for finite alphabet sources under assumption 5.4.1. Specifically let \mathcal{X} be a finite alphabet and our source $P(X^{T,N}) = \otimes_{n=1}^N P(X_n^T)$.

Theorem 5.4.1 *For finite alphabet sources satisfying assumption 5.4.1 we have for both formulations*

Formulation 1: Distortion Schedule

For any $\epsilon > 0$ and finite T one can find an $N(\epsilon, T)$ such that for $N \geq N(\epsilon, T)$

$$R_{T,N}^{SRD, o}(D_1 + \epsilon, \dots, D_T + \epsilon) \leq R_{T,1}^{SRD}(D_1, \dots, D_T) + \epsilon.$$

Formulation 2: Time-Average Distortion

For any $\epsilon > 0$ and finite T one can find an $N(\epsilon, T)$ such that for $N \geq N(\epsilon, T)$

$$R_{T,N}^{SRD, o, ave}(D + \epsilon) \leq R_{T,1}^{SRD, ave}(D) + \epsilon.$$

Proof: We prove formulation 1. Formulation 2 follows analogously. The proof uses a random coding argument and has many steps. We outline the steps before proceeding with the proof.

- (1) Test channel measure
- (2) Typical sets
- (3) Random codebook generation
- (4) Fixing the rate
- (5) Large deviation bound
- (6) Existence of sequential quantizer

Step 1: Test channel measure

Given the source and the infimizing law we know, by lemma 5.3.2, that there exists a joint measure $P(X^{T,N}, Y^{T,N})$. By assumption 5.4.1 and lemma 5.4.1 we see that this joint measure factors as $P(X^{T,N}, Y^{T,N}) = \otimes_{n=1}^N P(X_n^T, Y_n^T)$. Thus we can work with the single-letter characterization under the measure $P(X^T, Y^T)$. Since we are using the infimizing law we know $E_P d(X_t, Y_t) \leq D_t \quad \forall t = 1, \dots, T$. Furthermore $I_P(X^T; Y^T) = R_{T,1}^{SRD}(D_1, \dots, D_T)$.

Step 2: Typical sets

The following definitions and results follow from Orlicsky and Roche. [OR] Let A, B, C be generic finite valued random variables taking values in $\mathcal{A}, \mathcal{B}, \mathcal{C}$ respectively. Let their

joint distribution be denoted $P(A, B, C)$. Given any sequence a^N define its empirical measure to be $\nu_{a^N}(a) \triangleq \frac{|\{n: a_n = a\}|}{N}$. This is extended to multiple sequences in the natural way: $\nu_{(a^N, b^N)}(a, b) \triangleq \frac{|\{n: (a_n, b_n) = (a, b)\}|}{N}$.

Define the δ -robustly typical set

$$\Omega_A^{\delta, N} \triangleq \{a^N : |\nu_{a^N}(a) - P(a)| \leq \delta P(a), \quad \forall a \in \mathcal{A}\}.$$

We can define $\Omega_{A, B}^{\delta, N}$ similarly. Also define the conditional δ -robustly typical set

$$\Omega_{A, B}^{\delta, N}(a^N) \triangleq \{b^N : (a^N, b^N) \in \Omega_{A, B}^{\delta, N}\}.$$

Let

$$\mu_A \triangleq \min_{\{a \in \mathcal{A} : P(a) \neq 0\}} P(a).$$

We can define $\mu_{A, B}$ similarly. Finally define

$$\epsilon_A^{\delta, N} \triangleq 2|\mathcal{A}|e^{-N\frac{\delta^2\mu_A}{3}}$$

and for $\delta_2 > \delta_1$ let

$$\epsilon_{A, B}^{\delta_1, \delta_2, N} \triangleq 2|\mathcal{A} \times \mathcal{B}|e^{-N\frac{(\delta_2 - \delta_1)^2\mu_{A, B}}{3(1 + \delta_1)}}$$

The following three lemmas contain the technical results we will need to prove the theorem. The results can be found in [OR].

Lemma A.4.2 *The following hold*

- (a) $P(a^N \in \Omega_A^{\delta, N}) \geq 1 - \epsilon_A^{\delta, N}$.
- (b) For all $(a^N, b^N) \in \Omega_{A, B}^{\delta, N}$ we have $2^{-N(1+\delta)H(B|A)} \leq P(b^N | a^N) \leq 2^{-N(1-\delta)H(B|A)}$.
- (c) Let $a^N \in \Omega_A^{\delta_1, N}$. Let b^N be drawn from $P(B^N | a^N)$. Then

$$P(b^N \in \Omega_{A, B}^{\delta_2, N}(a^N) | a^N = a^N) \geq 1 - \epsilon_{A, B}^{\delta_1, \delta_2, N}.$$

- (d) For every $a^N \in \Omega_A^{\delta_1, N}$ we have $|\Omega_{A, B}^{\delta_2, N}(a^N)| \geq (1 - \epsilon_{A, B}^{\delta_1, \delta_2, N})2^{N(1-\delta_2)H(B|A)}$.

Proof: See lemma A.4.2 of the appendix. \square

The following is the key technical lemma in the proof of the coding theorem.

Lemma A.4.3 *Let $(x^{t, N}, y^{t-1, N}) \in \Omega_{X^t, Y^{t-1}}^{\delta_1, N}$. Let y_t^N be drawn from $P(Y_t^N | y^{t-1, N})$. Then*

$$\begin{aligned} & P(y_t^N \in \Omega_{X^t, Y^{t-1}}^{\delta_2, N}(x^{t, N}, y^{t-1, N}) | x^{t, N}, y^{t-1, N}) \\ & \geq (1 - \epsilon_{X^t, Y^t}^{\delta_1, \delta_2, N})2^{-N(I(X^t; Y_t | Y^{t-1}) + 2\delta_2 H(Y_t | Y^{t-1}))}. \end{aligned}$$

Proof: See lemma A.4.3 in the appendix. \square

In words, this lemma states that the probability that Y_t^N , drawn from $P(Y_t^N|y^{t-1,N})$, looks typical with respect to the measure $P(Y_t^N|x^{t,N}, y^{t-1,N})$ is approximately $2^{-NI(X^t; Y_t|Y^{t-1})}$.

Finally we show that if (x_t^N, y_t^N) are δ -robustly typical then their distortion is close to the average distortion.

Lemma A.4.4 *If $(x_t^N, y_t^N) \in \Omega_{X_t, Y_t}^{\delta, N}$ then $d_N(x_t^N, y_t^N) \leq (1 + \delta)E(d(X_t, Y_t))$.*

Proof: See lemma A.4.4 in the appendix. \square

Step 3: Random codebook generation

We will use the test channel measure, $P(X^{T,N}, Y^{T,N})$, discussed in step 1, to produce a distribution on sequential quantizers f_1, \dots, f_T . First we need some notation. Let M_t be the number of codewords available to the quantizer at time t . We will specify their size in step 4. Now let

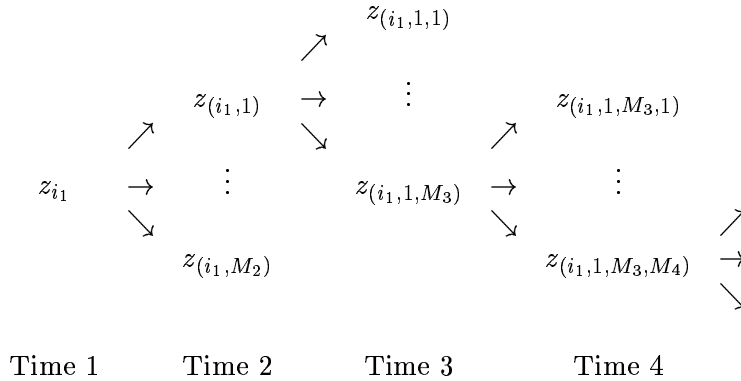
$$\gamma_1 = \{z_{i_1}, i_1 = 1, \dots, M_1\}$$

denote the range of f_1 . We call z_{i_1} the i_1 -codeword used at time one. In general let

$$\gamma_t = \{z_{(i_1, \dots, i_t)}, i_1 = 1, \dots, M_1, \dots, i_t = 1, \dots, M_t\}$$

denote the range of f_t . We call $z_{(i_1, \dots, i_t)}$ the (i_1, \dots, i_t) -codeword used at time t . Each codeword $z \in \mathcal{Y}^N$. Note that every element in the range of f^t is uniquely labeled by an (i_1, \dots, i_t) sequence. (Recall $f^t = (f_1, \dots, f_t)$.)

The following example may help understand this tree structure. Starting with the codeword z_{i_1} we have M_2 choices of codewords at the next step. This set of M_2 codewords can be different for each $i_1 = 1, \dots, M_1$. The diagram illustrates a typical path: $z_{i_1}, z_{(i_1,1)}, z_{(i_1,1,M_3)}, z_{(i_1,1,M_3,M_4)}, \dots$



We can define a distribution on $\Gamma_1, \dots, \Gamma_T$ as follows

$$P(\gamma^T) = \prod_{i_1=1}^{M_1} P_{Y_1^N}(z_{i_1}) \prod_{i_2=1}^{M_2} P_{Y_2^N|Y_1^N}(z_{i_1, i_2}|z_{i_1}) \times \dots \\ \times \prod_{i_T=1}^{M_T} P_{Y_T^N|Y^{T-1, N}}(z_{i_1, i_2, \dots, i_T}|z_{i_1}, z_{i_1, i_2}, \dots, z_{i_1, \dots, i_{T-1}}).$$

Note that we have only defined the ranges of the quantizers. We still need to define the quantizer maps. The idea is to use joint typicality encoders. Before doing so we add one special symbol, the error symbol, to the range of each quantizer. Thus the cardinality of the range of each quantizer at time t is now $M_t + 1$.

Given $\gamma_1, \dots, \gamma_T$ define

$$f_{\delta, t}(x^{t, N}) = \begin{cases} z_{(i_1, \dots, i_t)} & \text{if } (x^{t, N}, z_{i_1}, \dots, z_{(i_1, \dots, i_t)}) \in \Omega_{X^t, Y^t}^{\delta, N} \\ & \text{and } (i_1, \dots, i_{t-1}) \text{ is the labeling of } f_{\delta}^{t-1}(x^{t-1, N}) \\ \text{error symbol} & \text{if any of } z_{i_1}, \dots, z_{(i_1, \dots, i_{t-1})} = \text{error symbol} \\ & \text{or there are no typical elements in } \gamma_t \end{cases}$$

If there is more than one typical element then use some priority ordering to choose one of them. In summary the quantizer at time t looks to see if there are any codewords in the range of γ_t that are typical with respect to the source and consistent with the previous labeling of the codewords. If there are none then it produces an error symbol. Once a quantizer produces an error symbol it will continue to produce error symbols for all subsequent time steps.

The decoder upon receiving $z_{(i_1, \dots, i_t)}$ decodes it as $z_{(i_1, \dots, i_t)}$. The decoder upon receiving the error symbol decodes it as (x_0, x_0, \dots, x_0) . Recall assumption 5.3.1 where x_0 is an element such that $E_{P_{X_t}} d(X_t, x_0) < D_{\max}$. This assumption will allow us to bound the distortion caused by an error symbol. More on this in steps 6 and 7.

Step 4: Fixing the rate

Choose the M_t 's such that

$$\frac{1}{N} \log(M_t + 1) \leq I_P(X^t; Y_t | Y^{t-1}) + \epsilon, \quad t = 1, \dots, T$$

We use $M_t + 1$ instead of M_t to account for the error symbol. Any sequential quantizer drawn from $P(\Gamma^T)$ with M_t 's satisfying the inequalities above will satisfy the rate part of the theorem statement. Specifically for all such γ^T and their corresponding quantizer f^T we have

$$\frac{1}{NT} H_P(f_1(X_1^N), \dots, f_T(X^T, N)) \leq \frac{1}{NT} \sum_{t=1}^T H_P(f_T(X^t, N))$$

$$\begin{aligned}
&\leq \frac{1}{NT} \sum_{t=1}^T \log(M_t + 1) \\
&\leq \frac{1}{T} \sum_{t=1}^T \left(I_P(X^t; Y_t | Y^{t-1}) + \epsilon \right) \\
&= R_{T,1}^{\text{SRD}}(D_1, \dots, D_T) + \epsilon
\end{aligned}$$

Step 5: Large deviation bound

Recall the construction of the quantizer in step 3. If at time t we receive a non-error message then it must be the case that $(x^{t,N}, z_{i_1}, \dots, z_{(i_1, \dots, i_t)}) \in \Omega_{X^t, Y^t}^{\delta, N}$. By lemma A.4.4 this assures us that

$$d_N(x_t^N, z_{(i_1, \dots, i_t)}) \leq (1 + \delta) E P d(X_t, Y_t).$$

Our goal, then, is to show that the probability of producing an error signal can be made arbitrarily small. Let

$$0 < \delta_{(1,1)} < \delta_{(1,2)} < \delta_{(2,1)} < \dots < \delta_{(t,1)} < \delta_{(t,2)} < \dots < \delta_{(T,1)} < \delta_{(T,2)}$$

to be specified later.

We will bound the probability of error at time one. Then we will find a recursive bound on the probability of error for general times t .

$$\begin{aligned}
&P(\text{error symbol at time 1}) \\
&\leq P(y_1^N \text{ not typical given typical } x_1^N) + P(x_1^N \text{ not typical}) \\
&= \sum_{x_1^N \in \Omega_{X_1}^{N, \delta_{(1,1)}}} P(x_1^N) \left[P\left(y_1^N \notin \Omega_{X_1, Y_1}^{\delta_{(1,2)}, N}(x_1^N)\right) \right]^{M_1} + P(x_1^N \notin \Omega_{X_1}^{N, \delta_{(1,1)}}) \\
&\leq \sum_{x_1^N} P(x_1^N) \left[1 - P\left(y_1^N \in \Omega_{X_1, Y_1}^{\delta_{(1,2)}, N}(x_1^N)\right) \right]^{M_1} + \epsilon_{X_1}^{\delta_{(1,1)}, N} \\
&\leq \sum_{x_1^N} P(x_1^N) \exp \left[-M_1 P\left(y_1^N \in \Omega_{X_1, Y_1}^{\delta_{(1,2)}, N}(x_1^N)\right) \right] + \epsilon_{X_1}^{\delta_{(1,1)}, N} \\
&\leq \sum_{x_1^N} P(x_1^N) \exp \left[-M_1 (1 - \epsilon_{X^t, Y^t}^{\delta_{(1,1)}, \delta_{(1,2)}, N}) 2^{-N(I(X_1; Y_1) + 2\delta_{(1,2)} H(Y_1))} \right] + \epsilon_{X_1}^{\delta_{(1,1)}, N} \\
&\leq \exp \left[-M_1 (1 - \epsilon_{X^t, Y^t}^{\delta_{(1,1)}, \delta_{(1,2)}, N}) 2^{-N(I(X_1; Y_1) + 2\delta_{(1,2)} H(Y_1))} \right] + \epsilon_{X_1}^{\delta_{(1,1)}, N}
\end{aligned}$$

where the fourth line follows from lemma A.4.2 (a). The fifth line follows from the useful inequality $(1 - x)^M \leq e^{-Mx}$ for $0 < x < 1$. The sixth line follows from lemma A.4.4.

Now we show the general step:

$$P(\text{error symbol at time } t) = P(\text{error symbol before } t) + P(\text{first error symbol at time } t)$$

Recall that once an error symbol occurs all subsequent quantizers will output an error symbol. Thus $P(\text{error symbol before } t) = P(\text{error symbol at } t-1)$.

Now if no error symbol has occurred at time $t-1$ then it must be the case that

$$(x^{t-1,N}, y^{t-1,N}) \in \Omega_{X^{t-1}, Y^{t-1}}^{\delta_{(t-1,2),N}}$$

We will show there exists a bound on the probability that the first error occurs at time t that holds uniformly for every $(x^{t-1,N}, y^{t-1,N}) \in \Omega_{X^{t-1}, Y^{t-1}}^{\delta_{(t-1,2),N}}$.

$$\begin{aligned} & P\left(\text{error symbol at time } t \mid (x^{t-1,N}, y^{t-1,N}) \in \Omega_{X^{t-1}, Y^{t-1}}^{\delta_{(t-1,2),N}}\right) \\ & \leq \sum_{x_t^N \in \Omega_{X^{t-1}}^{\delta_{(t,1),N}}(x^{t-1,N})} P(x_t^N \mid x^{t-1,N}) \left[P\left(y_t^N \notin \Omega_{X^t, Y^t}^{\delta_{(t,2),N}}(x^{t,N}, y^{t-1,N}) \mid y^{t-1,N}\right) \right]^{M_t} \\ & \quad + P(x_t^N \notin \Omega_{X^{t-1}}^{\delta_{(t,1),N}}(x^{t-1,N}) \mid x^{t-1,N}) \\ & \leq \sum_{x_t^N \in \Omega_{X^{t-1}}^{\delta_{(t,1),N}}(x^{t-1,N})} P(x_t^N \mid x^{t-1,N}) \exp\left[-M_t P\left(y_t^N \in \Omega_{X^t, Y^t}^{\delta_{(t,2),N}}(x^{t,N}, y^{t-1,N}) \mid y^{t-1,N}\right)\right] \\ & \quad + \epsilon_{X_t, X^{t-1}}^{\delta_{(t-1,2), \delta_{(t,1),N}}} \\ & \leq \sum_{x_t^N \in \Omega_{X^{t-1}}^{\delta_{(t,1),N}}(x^{t-1,N})} P(x_t^N \mid x^{t-1,N}) \exp\left[-M_t (1 - \epsilon_{X^t, Y^t}^{\delta_{(t,1), \delta_{(t,2),N}}}) 2^{-N(I(X^t; Y_t | Y^{t-1}) + 2\delta_{(t,2)} H(Y_t | Y^{t-1}))}\right] \\ & \quad + \epsilon_{X_t, X^{t-1}}^{\delta_{(t-1,2), \delta_{(t,1),N}}} \\ & \leq \exp\left[-M_t (1 - \epsilon_{X^t, Y^t}^{\delta_{(t,1), \delta_{(t,2),N}}}) 2^{-N(I(X^t; Y_t | Y^{t-1}) + 2\delta_{(t,2)} H(Y_t | Y^{t-1}))}\right] + \epsilon_{X_t, X^{t-1}}^{\delta_{(t-1,2), \delta_{(t,1),N}}} \end{aligned}$$

Thus we have

$$\begin{aligned} & P(\text{error symbol at time } t) \\ & = P(\text{error symbol at } t-1) + P(\text{first error symbol at time } t) \\ & \leq P(\text{error symbol at } t-1) + \\ & \quad \exp\left[-M_t (1 - \epsilon_{X^t, Y^t}^{\delta_{(t,1), \delta_{(t,2),N}}}) 2^{-N(I(X^t; Y_t | Y^{t-1}) + 2\delta_{(t,2)} H(Y_t | Y^{t-1}))}\right] + \epsilon_{X_t, X^{t-1}}^{\delta_{(t-1,2), \delta_{(t,1),N}}} \quad (5.3) \end{aligned}$$

Thus we have a recursive bound for $P(\text{error symbol at time } t)$.

Step 6: Existence of sequential quantizer

Now we show the existence of a sequential quantizer. By lemma A.4.4 and step 5 we know the expected distortion average over the quantizers at time t is less than or equal to

$$(1 + \delta_{(t,2)})E_P(d(X_t, Y_t)) + D_{\max}P(\text{error symbol at time } t).$$

Choose $\delta_{(t,2)} < \frac{\epsilon}{2D_t}$. Then the first addend is less than $D_t + \frac{\epsilon}{2}$. We now show that we can make the second addend less than $\frac{\epsilon}{2}$.

In step 4 we have chosen the rates such that

$$M_t \leq 2^{N(I_P(X^t; Y_t | Y^{t-1}) + \epsilon)}.$$

Thus we should further choose $\delta_{(t,2)}$ to be small enough so that exponent in the second addend in (5.3) is negative. (Note we still need $\delta_{(t,2)} < \frac{\epsilon}{2D_t}$ also.) Thus the second addend is decreasing to zero with N . By our assumption that $\delta_{(t-1,2)} < \delta_{(t,1)}$ we see that the exponent in the third addend in (5.3) is negative. Thus the third addend is decreasing to zero with N . Since T is finite we can find an N large enough so that $P(\text{error symbol at time } t) < \frac{\epsilon}{2D_{\max}}$ holds for all $t = 1, \dots, T$.

Thus we have shown that the expected distortion at time t averaged over all quantizers is $\leq D_t + \epsilon$. Since this is an average statement there must exist at least one quantizer that achieves the required distortion schedule. The theorem is proved. \square

Note that we have proved something slightly stronger. In the formulation of the sequential rate distortion problem with a distortion criterion our goal is to minimize the average rate over time. But in the course of our proof we were able to show exactly what the rate at each time instant will be. The rate at time t is less than or equal to $I_P(X^t; Y_t | Y^{t-1})$ for all $t = 1, \dots, T$.

5.4.2 Stochastic Control Formulation for Finite Alphabet Source Coding

In this section we continue to assume that the source alphabet is finite and that assumption 5.4.1 holds. Our goal in this subsection is to pose the SRD problem as a constrained Markov decision problem. To that end we assume that the source is Markov in time.

Assumption 5.4.2 *Our source is Markov in time. Specifically there exist stochastic kernels such that*

$$P(X^T) = P(X_1) \otimes_{t=2}^T Q(X_t | x_{t-1}).$$

Structure Results

Recall our objective is to minimize $I(X^T; Y^T)$ over all causal channels satisfying the distortion criterion. Under a causal sequence of stochastic kernels the mutual information

decomposes in the following manner

$$I(X^T; Y^T) = I(X^T \rightarrow Y^T) = \sum_{t=1}^T I(X^t, Y_t | Y^{t-1}).$$

Note the dependence on X^t in the addend $I(X^t, Y_t | Y^{t-1})$. We will show that under assumption 5.4.2 we can, without loss of generality, restrict our attention to addends of the form $I(X_t, Y_t | Y^{t-1})$. By doing so we will be able to formulate a stochastic control problem with running cost equal to these addends.

Definition 5.4.1 *A causal sequence of stochastic kernels, $\{Q(Y_t | x^t, y^{t-1}, \cdot)\}_{t=1}^T$, is called a simplified causal sequence of stochastic kernels if $\forall t = 1, \dots, T$ we have*

$$Q(Y_t | x^{t-1}, x_t, y^{t-1}) = Q(Y_t | \tilde{x}^{t-1}, x_t, y^{t-1}) \quad \forall x^{t-1}, \tilde{x}^{t-1}.$$

We denote a simplified causal sequence of stochastic kernels by $\{Q(Y_t | x_t, y^{t-1})\}_{t=1}^T$.

In words this states that the distribution of Y_t is a function of only x_t, y^{t-1} and is independent of x^{t-1} . The following two lemmas show that we can restrict our attention to a simplified causal sequence of stochastic kernels.

Lemma 5.4.2 *We are given a Markov $P(X^T)$ and a causal sequence of stochastic kernels, $\{P(Y_t | x^t, y^{t-1})\}_{t=1}^T$. Denote the resulting joint measure by $P(X^T, Y^T) = P(X^T) \otimes \left(\otimes_{t=1}^T P(Y_t | x^t, y^{t-1})\right)$. Then there exists a simplified causal sequence of stochastic kernels, $\{Q(Y_t | x_t, y^{t-1})\}_{t=1}^T$, such that for the resulting joint measure $Q(X^T, Y^T) = P(X^T) \otimes \left(\otimes_{t=1}^T \tilde{Q}(Y_t | x_t, y^{t-1})\right)$ the following marginals hold for all $t = 1, \dots, T$*

$$Q(X_t, Y^t) = P(X_t, Y^t).$$

Proof: We can decompose $P(X_t, Y^t) = P(Y_t | x_t, y^{t-1}) \otimes P(X_t, Y^{t-1})$. Let

$$Q(Y_t | x_t, y^{t-1}) = P(Y_t | x_t, y^{t-1}).$$

It is straightforward to verify that $Q(X_1, Y_1) = P(X_1, Y_1)$. Assume the result holds for all $t \leq k$. We now prove the induction step.

$$\begin{aligned} Q(X_{k+1}, Y^{k+1}) &= Q(Y_{k+1} | x_{k+1}, y^k) \otimes Q(X_{k+1}, Y^k) \\ &= P(Y_{k+1} | x_{k+1}, y^k) \otimes \left(\sum_{x_k \in \mathcal{X}} Q(x_k, X_{k+1}, Y^k) \right) \\ &= P(Y_{k+1} | x_{k+1}, y^k) \otimes \left(\sum_{x_k \in \mathcal{X}} Q(X_{k+1} | x_k, y^k) \otimes Q(x_k, Y^k) \right) \\ &= P(Y_{k+1} | x_{k+1}, y^k) \otimes \left(\sum_{x_k \in \mathcal{X}} P(X_{k+1} | x_k) \otimes P(x_k, Y^k) \right) \end{aligned}$$

$$\begin{aligned}
&= P(Y_{k+1} \mid x_{k+1}, y^k) \otimes P(X_{k+1}, Y^k) \\
&= P(X_{k+1}, Y^{k+1})
\end{aligned}$$

Where in the fourth equality we have used the Markov property of the source and the induction hypothesis. This proves the induction step and thus the lemma. \square

We now show that the mutual information is not increased by this Q channel.

Lemma 5.4.3 *For the measures P_{X^T, Y^T} and Q_{X^T, Y^T} of the previous lemma we have*

$$I_Q(X^T; Y^T) \leq I_P(X^T; Y^T).$$

Proof: We know that under Q the following is a Markov chain: $Y_t - (X_t, Y^{t-1}) - X^{t-1}$. Thus $I_Q(X^t; Y_t \mid Y^{t-1}) = I_Q(X_t; Y_t \mid Y^{t-1})$.

$$\begin{aligned}
I_P(X^T; Y^T) &= I_P(X^T \rightarrow Y^T) \\
&= \sum_{t=1}^T I_P(X^t; Y_t \mid Y^{t-1}) \\
&= \sum_{t=1}^T I_P(Y_t; X^t, Y^{t-1}) - I_P(Y_t; Y^{t-1}) \\
&\geq \sum_{t=1}^T I_P(Y_t; X_t, Y^{t-1}) - I_P(Y_t; Y^{t-1}) \\
&= \sum_{t=1}^T I_Q(Y_t; X_t, Y^{t-1}) - I_Q(Y_t; Y^{t-1}) \text{ by lemma 4.4.2} \\
&= \sum_{t=1}^T I_Q(X_t; Y_t \mid Y^{t-1}) \\
&= \sum_{t=1}^T I_Q(X^t; Y_t \mid Y^{t-1}) \\
&= I_Q(X^T \rightarrow Y^T) \\
&= I_Q(X^T; Y^T)
\end{aligned}$$

\square

By lemma 5.4.2 the marginals are equal: $P_{X_t, Y_t} = Q_{X_t, Y_t}$. Thus the average distortion under the causal sequence of stochastic kernels $\{P(Y_t \mid x^t, y^{t-1})\}_{t=1}^T$ equals the average distortion under the simplified causal sequence of stochastic kernels $\{Q(Y_t \mid x_t, y^{t-1})\}_{t=1}^T$.

Furthermore the mutual information is not increased by using the sequence of causal stochastic kernels. Thus we can restrict the infimization in definitions 5.3.5 and 5.3.6 to simplified causal sequences of stochastic kernels.

Dynamic Programming Formulation

We are now in a position to translate the elements of this problem into a traditional control problem. By lemmas 5.4.2 and 5.4.3 we know the optimization problem reduces to the following: Infimize $\sum_{t=1}^T I(X_t; Y_t | Y^{t-1})$ over all *simplified* causal sequences of stochastic kernels (e.g. channels) of the form $\{Q(Y_t | x_t, y^{t-1})\}_{t=1}^T$ while maintaining a given distortion criterion.

We view the sequential rate distortion problem as a control problem where the decoder acts as the controller. Specifically the decoder at time t chooses the encoder at time $t + 1$. In this case the encoder is a stochastic kernel. Let $\mathcal{K}_{\mathcal{Y}|\mathcal{X}}$ be the space of stochastic transition matrices from \mathcal{X} to \mathcal{Y} .

The components of the control problem are:

- *Control*

The control takes values in the space of stochastic kernels $u_t \in \mathcal{K}_{\mathcal{Y}|\mathcal{X}}$. A control policy is a sequence of measurable functions $\mu_t : \mathcal{K}_{\mathcal{Y}|\mathcal{X}}^{t-1} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{K}_{\mathcal{Y}|\mathcal{X}}$ taking $(u^{t-1}, y^{t-1}) \mapsto u_t$.

- *State*

Let the state be $z_t = (u^{t-1}, y^{t-1})$ where $z_t \in \mathcal{K}_{\mathcal{Y}|\mathcal{X}}^{t-1} \times \mathcal{Y}^{t-1}$.

- *Time Ordering and Joint Measure*

$$X_1, U_1, Y_1, \dots, X_T, U_T, Y_T$$

The source and any policy, μ^T , defines a unique measure, P^{μ^T} , on $\mathcal{X}^T \times \mathcal{K}_{\mathcal{Y}|\mathcal{X}}^T \times \mathcal{Y}^T$.

$$\begin{aligned} P^{\mu^T}(X^T, dU^T, Y^T) &= P(X^T) \otimes \left\{ \bigotimes_{t=1}^T P^{\mu^T}(Y_t | x^t, u^t, y^{t-1}) \otimes P^{\mu^T}(dU_t | x^t, u^{t-1}, y^{t-1}) \right\} \\ &= P(X^T) \otimes \left\{ \bigotimes_{t=1}^T u_t(Y_t | x_t) \otimes \delta_{\{U_t = \mu_t(u^{t-1}, y^{t-1})\}} \right\} \end{aligned}$$

- *State Evolution*

Note that

$$\begin{aligned} P^{\mu^t}(Y_t | u^t, y^{t-1}) &= \sum_{x_t} P^{\mu^t}(Y_t | x_t, u^t, y^{t-1}) P^{\mu^t}(x_t | u^t, y^{t-1}) \\ &= \sum_{x_t} u_t(Y_t | x_t) P(x_t | u^{t-1}, y^{t-1}) \\ &= \sum_{x_t} u_t(Y_t | x_t) P(x_t | z^t) \end{aligned}$$

and is independent of the policy μ^t . Thus the state evolution is

$$P(Z_{t+1} | z_t, u_t) = \begin{cases} 0 & \text{if } (U^t, Y^{t-1}) \neq (u^t, y^{t-1}) \\ P(Y_t | u^t, y^{t-1}) & \text{else.} \end{cases}$$

- *Running cost*

Define the following two measures

$$Q_1(X_t, Y_t | z_t, u_t) \triangleq u_t(Y_t | x_t) \otimes P(X_t | z_t)$$

and

$$Q_2(X_t, Y_t | z_t, u_t) \triangleq \left(\sum_{\tilde{x}} u_t(Y_t | \tilde{x}_t) P(\tilde{x}_t | z_t) \right) \otimes P(X_t | z_t)$$

The running cost is

$$c(z_t, u_t) \triangleq D(Q_1(X_t, Y_t | z_t, u_t) | Q_2(X_t, Y_t | z_t, u_t)) = I(X_t; Y_t | z_t, u_t).$$

- *Distortion*

$$d(z_t, u_t) \triangleq E(d(X_t, Y_t) | z_t, u_t) = \sum_{x_t, y_t} d(x_t, y_t) u_t(y_t | x_t) P(x_t | z_t)$$

(Note that we have used $d(\cdot, \cdot)$ in two different ways.)

- *Objective*

Infimize $\sum_{t=1}^T E(c(Z_t, U_t))$ over all control policies, μ^T , while maintaining either distortion constraint:

- (1) $E(d(z_t, u_t)) \leq D_t \quad \forall t = 1, \dots, T$
- (2) $\frac{1}{T} \sum_{t=1}^T E(d(z_t, u_t)) \leq D$

In the terminology of [BS] we have just defined a non-stationary stochastic optimal control problem with constraint. The qualifier “non-stationary” comes from the fact that the state space $\mathcal{Y}^{t-1} \times \mathcal{K}_{\mathcal{Y}|\mathcal{X}}^{t-1}$ is changing in time. In fact it is growing.

We now prove one final simplification. We will show that the distribution of the state given z_t is a sufficient statistic for the problem. Consequently we can use this distribution as the state.

We will use the conditional probability of X_t given z_t as our sufficient statistic. Specifically let $\eta_t : \mathcal{K}_{\mathcal{Y}|\mathcal{X}}^{t-1} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{P}(\mathcal{X})$ taking $z_t \mapsto P(X_t | z_t)$. Let $\Pi_t = \eta_t(Z_t)$. We will show that (Π_1, \dots, Π_T) is a sufficient statistic for control.

Note that in the definitions of $c(z_t, u_t)$ and $d(z_t, u_t)$ the term z_t enters the cost and distortion only through $P(X_t|z_t)$. Thus we define

$$\bar{c}(\pi_t, u_t) \triangleq E(c(Z_t, u_t) | \pi_t).$$

Note that if $\pi_t = \eta_t(z_t)$ then $\bar{c}(\pi_t, u_t) = c(z_t, u_t)$. Similarly define

$$\bar{d}(\pi_t, u_t) \triangleq E(d(Z_t, u_t) | \pi_t).$$

Similarly if $\pi_t = \eta_t(z_t)$ then $\bar{d}(\pi_t, u_t) = d(z_t, u_t)$.

Lemma 5.4.4 *The process Π_t is a controlled Markov process.*

Proof:

$$\begin{aligned} P(X_{t+1} | y^t, u^t) &= \sum_{x_t} P(X_{t+1} | x_t) P(x_t | y^t, u^t) \\ &= \sum_{x_t} P(X_{t+1} | x_t) \frac{P(y_t | x_t, y^{t-1}, u^t) P(x_t | y^{t-1}, u^t)}{\sum_{\tilde{x}} P(y_t | \tilde{x}_t, y^{t-1}, u^t) P(\tilde{x}_t | y^{t-1}, u^t)} \\ &= \sum_{x_t} P(X_{t+1} | x_t) \frac{u_t(y_t | x_t) P(x_t | y^{t-1}, u^{t-1})}{\sum_{\tilde{x}} u_t(y_t | \tilde{x}_t) P(\tilde{x}_t | y^{t-1}, u^{t-1})} \\ &= \Phi(\pi_t, u_t, y_t) \end{aligned}$$

for some function Φ . Thus

$$P(\Pi_{t+1} | \pi^t, u^t) = P(\Pi_{t+1} | \pi_t, u_t).$$

□

It is straightforward to verify that

$$I_{P^\mu}(X^T; Y^T) = \sum_{t=1}^T I_{P^\mu}(X_t; Y_t | Y^{t-1}) = \sum_{t=1}^T E_{P^\mu}(\bar{c}(\Pi_t, U_t))$$

and

$$E_{P^\mu}(d(X_t, Y_t)) = E_{P^\mu}(\bar{d}(\Pi_t, U_t)) \quad t = 1, \dots, T.$$

Now we can pose the dynamic programming equations. As stated before the problem is a constrained Markov decision problem. Such problems are difficult to solve. We simplify the problem by strengthening the form of the constraint. The resulting optimal rate under formulation 1a and 2a defined below will be an upper bound on the optimal rate in formulation 1 and 2 respectively.

Formulation 1a: Strong Distortion Schedule

Strengthen the constraint in formulation 1 by using

$$\mathcal{F} = \{\{Q(Y_t|x_t, \pi_t)\}_{t=1}^T : \bar{d}(\pi_t, u_t) \leq D_t \quad t = 1, \dots, T\}.$$

Formulation 2a: Time-average Strong Distortion

Strengthen the constraint in formulation 2 by using

$$\mathcal{F} = \{\{Q(Y_t|x_t, \pi_t)\}_{t=1}^T : \frac{1}{T} \sum_{t=1}^T \bar{d}(\pi_t, u_t) \leq D\}.$$

Definition 5.4.2 Let J_1, \dots, J_T be functions on $\mathcal{P}(\mathcal{X})$ defined backwards starting with T :

Formulation 1a: Strong Distortion Schedule

where

$$J_T(\pi) = \inf_{u \in \Omega_T(\pi)} \bar{c}(\pi, u)$$

and

$$J_t(\pi) = \inf_{u \in \Omega_t(\pi)} \bar{c}(\pi, u) + \int J_{t+1}(\tilde{\pi}) P(d\tilde{\pi}|\pi, u)$$

where $\Omega_t(\pi) = \{u \in \mathcal{K}_{\mathcal{Y}|\mathcal{X}} \text{ such that } \bar{d}(\pi, u) \leq D_t\}$.

Formulation 2a: Time-Average Strong Distortion

Here we expand the state to (π_t, δ_t) where δ_t represents the distortion accrued up to time $t - 1$. Let

$$J_T(\pi, \delta) = \inf_{u \in \Omega_T(\pi, \delta)} \bar{c}(\pi, u)$$

and

$$J_t(\pi, \delta) = \inf_{u \in \Omega_t(\pi, \delta)} \bar{c}(\pi, u) + \int J_{t+1}(\tilde{\pi}, \tilde{\delta}) P(d\tilde{\pi}, d\tilde{\delta}|\pi, \delta, u)$$

where $\Omega_t(\pi, \delta) = \{u \in \mathcal{K}_{\mathcal{Y}|\mathcal{X}} \text{ such that } \bar{d}(\pi, u) \leq D - \delta\}$. Also $\delta_{t+1} = \delta_t + \bar{d}(\pi_t, u_t)$.

Theorem 5.4.2 In both cases if the infimization is achieved by a policy $\mu^T = (\mu_1, \dots, \mu_T)$ then μ^T is optimal. Furthermore μ^T can be chosen to be a deterministic function of the $\{\pi_t\}$ or $\{\pi_t, \delta_t\}$ processes.

Proof: Theorem 3.2.1 of [HLL]. \square

5.4.3 Gaussian Source Coding Theorem

In this subsection we provide a source coding theorem for Gaussian sources. We first give a structure result. Then we discuss source mismatch. And we conclude with a coding theorem.

Gaussian Structure Result

We show that for a Gaussian source the Gaussian channel is the best channel in terms of minimizing the mutual information.

Lemma 5.4.5 *Let $G(dX^T)$ be a jointly Gaussian source admitting a density g . Let $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$ be a channel. Call the resulting joint measure $P(dX^T, dY^T)$. Furthermore assume that $P(dX^T, dY^T)$ admits a density with respect to the Lebesgue measure. Let $G(dX^T, dY^T)$ be a jointly Gaussian measure with the same second order properties as $P(dX^T, dY^T)$. Then*

- (a) $\{G(dY_t | x^t, y^{t-1})\}_{t=1}^T$ is a Gaussian channel.
- (b) $G(dX^T, dY^T)$ has the same independence properties as $P(dX^T, dY^T)$.
- (c) $I_G(X^T; Y^T) \leq I_P(X^T; Y^T)$.

Proof: Part (a) follows from the fact that $G(dX^T, dY^T)$ is jointly Gaussian. Part (b) follows from noting that independence or conditional independence of some random variables implies that those same random variables are uncorrelated or conditionally uncorrelated. $G(dX^T, dY^T)$ has the same second order properties as $P(dX^T, dY^T)$ thus it inherits the same independence properties. For part (c) note

$$\begin{aligned}
& I_P(X^T; Y^T) - I_G(X^T; Y^T) \\
&= \int p_{X^T, Y^T}(x^T, y^T) \log \frac{p_{X^T, Y^T}(x^T, y^T)}{p_{X^T}(x^T)p_{Y^T}(y^T)} dx^T dy^T \\
&\quad - \int g_{X^T, Y^T}(x^T, y^T) \log \frac{g_{X^T, Y^T}(x^T, y^T)}{g_{X^T}(x^T)g_{Y^T}(y^T)} dx^T dy^T \\
&= \int p_{X^T, Y^T}(x^T, y^T) \log \frac{p_{X^T|Y^T}(x^T | y^T)}{g_{X^T}(x^T)} dx^T dy^T \\
&\quad - \int p_{X^T, Y^T}(x^T, y^T) \log \frac{g_{X^T|Y^T}(x^T | y^T)}{g_{X^T}(x^T)} dx^T dy^T \\
&= \int p_{X^T, Y^T}(x^T, y^T) \log \frac{p_{X^T|Y^T}(x^T | y^T)}{g_{X^T|Y^T}(x^T | y^T)} dx^T dy^T \\
&= \int p_{Y^T}(y^T) D(P_{X^T|y^T} | G_{X^T|y^T}) dy^T \\
&\geq 0
\end{aligned}$$

Where the second equality comes from the fact that the G has the same second order properties as P and because $P(dX^T) = G(dX^T)$. \square

If we are using the squared error distortion measure we see that the distortion is the same under P or G . Thus for the Gaussian source the best channel that minimizes mutual information while maintaining a given squared error distortion is the Gaussian channel.

Mismatch

What happens if one uses a channel not matched to a given source?

Corollary 5.4.1 *Let $P(dX^T)$ be a source admitting a density. Let $G(dX^T)$ be a Gaussian source with the same second order statistics as $P(dX^T)$. Let $\{G(dY_t | x^t, y^{t-1})\}_{t=1}^T$ be a Gaussian channel. Then both joint measures $P(dX^T) \otimes \left(\bigotimes_{t=1}^T G(dY_t | x^t, y^{t-1})\right)$ and $G(dX^T) \otimes \left(\bigotimes_{t=1}^T G(dY_t | x^t, y^{t-1})\right)$ have the same second order statistics.*

Proof: The result follows from lemma 5.4.5. \square

Thus if we use the Gaussian channel $\{G(dY_t | x^t, y^{t-1})\}_{t=1}^T$ for the source P_{X^T} then the resulting distortions D_1, \dots, D_T will equal that of the distortions under G_{X^T} . This is related to the rate distortion mismatch problem. [Lap] The mismatch problem pertains to computing the resulting distortion when one does not use the appropriate infimizing channel.

Gaussian Coding Theorem

We now prove a direct coding theorem for transmitting Gaussian sources over a noiseless digital channel.

In section 5.3.2 we proved a general converse theorem. We still assume assumption 5.4.1.

Theorem 5.4.3 *For Gaussian sources satisfying assumption 5.4.1 and the squared error distortion measure we have:*

Formulation 1: Distortion Schedule

For any $\epsilon > 0$ and finite T one can find an $N(\epsilon, T)$ such that for $N \geq N(\epsilon, T)$

$$R_{T,N}^{SRD, o}(D_1 + \epsilon, \dots, D_T + \epsilon) \leq R_{T,1}^{SRD}(D_1, \dots, D_T) + \epsilon.$$

Formulation 2: Time-Average Distortion

For any $\epsilon > 0$ and finite T one can find an $N(\epsilon, T)$ such that for $N \geq N(\epsilon, T)$

$$R_{T,N}^{SRD, o, ave}(D + \epsilon) \leq R_{T,1}^{SRD, ave}(D) + \epsilon.$$

Proof: We prove formulation 1. Formulation 2 follows analogously. We reduce this problem to the problem we solved in theorem 5.4.1 by partitioning the alphabet space.

Let $P(dX^T, dY^T)$ be the interconnection between the Gaussian source and the infimizing channel. We know this is a jointly Gaussian measure.

We will choose an appropriate “rectangular” partition, Π , of $\mathcal{X}^T \times \mathcal{Y}^T$. By “rectangular” we mean that every element $\pi \in \Pi$ is of the form $\pi = E_1 \times \dots \times E_T \times F_1 \times \dots \times F_T$ where $E_t \in \mathcal{B}(\mathcal{X})$ and $F_t \in \mathcal{B}(\mathcal{Y})$.

With this partition, Π , we define new “quantized” random variables \tilde{X}^T, \tilde{Y}^T . Specifically if $X_t \in E_t$ then let \tilde{X}_t be any fixed representative point in E_t . Similarly if $Y_t \in F_t$ then let \tilde{Y}_t be any fixed representative point in F_t .

Now choose a finite partition Π_1^δ so that

$$|I(X^t; Y_t | Y^{t-1}) - I(\tilde{X}^t; \tilde{Y}_t | \tilde{Y}^{t-1})| \leq \delta \quad \forall t = 1, \dots, T.$$

By the definition of mutual information such a partition exists. (See definition A.3.2 in the appendix.)

Similarly choose a finite partition Π_2^δ so that

$$E(d(X_t, \tilde{X}_t)) \leq \delta \quad \text{and} \quad E(d(Y_t, \tilde{Y}_t)) \leq \delta \quad \forall t = 1, \dots, T.$$

Since the variance of X_t and Y_t is finite for all t we know that such a finite partition exists.

Let Π^δ be the common refinement of Π_1^δ and Π_2^δ . Now note that \tilde{X}_t, \tilde{Y}_t are finite valued random variables. Thus we can use theorem 5.4.1 to show that there exists a sequential quantizer with rates

$$R_t \leq I(\tilde{X}^t; \tilde{Y}_t | \tilde{Y}^{t-1}) + \frac{\epsilon}{2}$$

and average distortion

$$D_t \leq E(d(\tilde{X}_t, \tilde{Y}_t)) + \frac{\epsilon}{2}$$

for all $t = 1, \dots, T$.

We can find a δ small enough so that

$$I(\tilde{X}^t; \tilde{Y}_t | \tilde{Y}^{t-1}) \leq I(X^t; Y_t | Y^{t-1}) + \frac{\epsilon}{2}$$

and

$$\begin{aligned} E(d(\tilde{X}_t, \tilde{Y}_t)) &\leq E(d(X_t, \tilde{X}_t)) + E(d(X_t, Y_t)) + E(d(Y_t, \tilde{Y}_t)) \\ &\leq E(d(X_t, Y_t)) + \frac{\epsilon}{2} \end{aligned}$$

for all $t = 1, \dots, T$. For this choice of δ we see that

$$R_t \leq I(X^t; Y_t | Y^{t-1}) + \epsilon$$

and average distortion

$$D_t \leq E(d(X_t, Y_t)) + \epsilon$$

for all $t = 1, \dots, T$. \square

5.4.4 Stochastic Control Formulation for Gaussian Source Coding

Our goal in this subsection is to pose the SRD problem as a constrained Markov decision problem. To that end we assume that the source is Markov in time. This is just assumption 5.4.2 specialized to the Gaussian source.

Assumption 5.4.3 *Our Gaussian source is Markov in time. Specifically there exist gains A_t and independent Gaussians $W_t \sim \mathcal{N}(0, K_{W_t})$ such that*

$$X_1 \sim \mathcal{N}(0, K_{X_1}) \quad \text{and} \quad X_{t+1} = A_t x_t + W_t, \quad t = 1, \dots, T-1$$

(Without loss of generality we assume the process is zero mean.)

Here we have specified the source stochastic kernels in terms of a recurrence. The two approaches are equivalent.

Structure Results

Recall a Gaussian channel $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$ can be realized by the recursive description $Y_t = \alpha_t x^t + \beta_t y^{t-1} + V_t$ (for some gains α_t, β_t and independent Gaussian random variables V_t .)

Analogous to definition 5.4.1 we define

Definition 5.4.3 *A causal sequence of Gaussian stochastic kernels, $\{Y_t = \alpha_t x^t + \beta_t y^{t-1} + V_t\}_{t=1}^T$, is called a simplified causal sequence of Gaussian stochastic kernels if $\forall t = 1, \dots, T$ we have $\alpha_t x^t = \alpha_{t,t} x_t$ (where $\alpha_t = (\alpha_{t,1}, \dots, \alpha_{t,t})$.) We denote a simplified causal sequence of Gaussian stochastic kernels by $\{Y_t = \alpha_{t,t} x_t + \beta_t y^{t-1} + V_t\}_{t=1}^T$.*

We also prove lemmas analogous to lemma 5.4.2 and lemma 5.4.3.

Lemma 5.4.6 *We are given a Gauss-Markov source $P(dX^T)$ and a causal sequence of Gaussian stochastic kernels, $\{P(dY_t | x^t, y^{t-1})\}_{t=1}^T$. Denote the resulting joint measure by $P(dX^T, dY^T) = P(dX^T) \otimes \left(\bigotimes_{t=1}^T P(dY_t | x^t, y^{t-1}) \right)$. Then there exists a simplified causal sequence of stochastic kernels, $\{Q(dY_t | x_t, y^{t-1})\}_{t=1}^T$, such that for the resulting joint measure $Q(dX^T, dY^T) = P(dX^T) \otimes \left(\bigotimes_{t=1}^T \tilde{Q}(dY_t | x_t, y^{t-1}) \right)$ the following marginals hold for all $t = 1, \dots, T$*

$$Q(dX_t, dY^t) = P(dX_t, dY^t).$$

Proof: Let p and q represent the densities of P and Q respectively. We can decompose $p(x_t, y^t) = p(y_t | x_t, y^{t-1})p(x_t, y^{t-1})$. Let

$$q(y_t | x_t, y^{t-1}) = p(y_t | x_t, y^{t-1}).$$

It is straightforward to verify that $q(x_1, y_1) = p(x_1, y_1)$ holds for all (x_1, y_1) . Assume the result holds for all $t \leq k$. We now prove the induction step. For any (x_{k+1}, y^{k+1}) we have

$$\begin{aligned} q(x_{k+1}, y^{k+1}) &= q(y_{k+1} \mid x_{k+1}, y^k) \left(\int q(x_{k+1} \mid x_k, y^k) q(x_k, y^k) dx_k \right) \\ &= q(y_{k+1} \mid x_{k+1}, y^k) \left(\int p(x_{k+1} \mid x_k) p(x_k, y^k) dx_k \right) \\ &= p(x_{k+1}, y^{k+1}) \end{aligned}$$

This proves the induction step and thus the lemma. \square

We now show that the mutual information is not increased by this Q channel.

Lemma 5.4.7 *For the measures P_{X^T, Y^T} and Q_{X^T, Y^T} of the previous lemma we have*

$$I_Q(X^T; Y^T) \leq I_P(X^T; Y^T).$$

Proof: Lemma 5.4.3 holds for general alphabets.

By lemma 5.4.6 the marginals are equal: $P(dX_t, dY_t) = Q(dX_t, dY_t)$. Thus the average distortion under the causal sequence of Gaussian stochastic kernels $\{P(dY_t \mid x^t, y^{t-1})\}_{t=1}^T$ equals the average distortion under the simplified causal sequence of Gaussian stochastic kernels $\{Q(dY_t \mid x_t, y^{t-1})\}_{t=1}^T$.

Furthermore the mutual information is not increased by using the causal sequence of Gaussian stochastic kernels. Thus we can restrict the infimization in definitions 5.3.5 and 5.3.6 to simplified causal sequences of Gaussian stochastic kernels.

Dynamic Programming Formulation

It is straightforward to extend the dynamic programming formulation for finite alphabet sources described in subsection 5.4.2. to the Gaussian source case. We quickly specify the components of the control problem.

- *Control*

The control takes values in the space of Gaussian stochastic kernels $u_t \in \mathcal{K}_{\mathcal{Y}|\mathcal{X}}^G$ (where we have used the superscript “G” to represent Gaussian.) A control policy is a sequence of measurable functions $\mu_t : \mathcal{K}_{\mathcal{Y}|\mathcal{X}}^{G, t-1} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{K}_{\mathcal{Y}|\mathcal{X}}^G$ taking $(u^{t-1}, y^{t-1}) \mapsto u_t$. Recall a simplified Gaussian channel has the form $Y_t = \alpha_t x_t + \beta_t y^{t-1} + V_t$. Thus we need only specify α_t, β_t and the Gaussian V_t .

- *State*

Let the state be $z_t = (u^{t-1}, y^{t-1})$ where $z_t \in \mathcal{K}_{\mathcal{Y}|\mathcal{X}}^{G, t-1} \times \mathcal{Y}^{t-1}$.

- *Joint Measure*

$$P^{\mu^T}(dX^T, dU^T, dY^T) = P(dX^T) \otimes \left\{ \bigotimes_{t=1}^T u_t(dY_t | x_t) \otimes \delta_{\{U_t = \mu_t(u^{t-1}, y^{t-1})\}} \right\}$$

- *State Evolution*

As before the state evolution can be shown to be

$$P(dZ_{t+1} | z_t, u_t) = \begin{cases} 0 & \text{if } (U^t, Y^{t-1}) \neq (u^t, y^{t-1}) \\ P(dY_t | u^t, y^{t-1}) & \text{else.} \end{cases}$$

- *Running cost and distortion*

The running cost is

$$c(z_t, u_t) \triangleq I(X_t; Y_t | z_t, u_t) = \frac{1}{2} \log \frac{|\text{cov}(X_t | Y^{t-1})|}{|\text{cov}(X_t | Y^t)|}.$$

And the distortion is

$$d(z_t, u_t) \triangleq E(\|X_t - Y_t\|^2 | z_t, u_t) = \text{trace}(\text{cov}(X_t | Y^t))$$

where the distribution on (X_t, Y^t) is well-defined given u^t .

- *Objective*

Infimize $\sum_{t=1}^T E(c(Z_t, U_t))$ over all control policies, μ^T , while maintaining either distortion constraint:

- (1) $E(d(z_t, u_t)) \leq D_t \quad \forall t = 1, \dots, T$
- (2) $\frac{1}{T} \sum_{t=1}^T E(d(z_t, u_t)) \leq D$

We will use the conditional probability of X_t given z_t as our sufficient statistic. Specifically let $\eta_t : \mathcal{K}_{\mathcal{Y} | \mathcal{X}}^{\mathbf{G}, t-1} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{P}(\mathcal{X})$ taking $z_t \mapsto P(dX_t | z_t) = \mathcal{N}(E(X_t | z_t), \text{cov}(X_t | z_t))$. Let $\Pi_t = \eta_t(Z_t)$. As before one can show that (Π_1, \dots, Π_T) is a sufficient statistic for control. Define

$$\bar{c}(\pi_t, u_t) \triangleq E(c(Z_t, u_t) | \pi_t) \quad \text{and} \quad \bar{d}(\pi_t, u_t) \triangleq E(d(Z_t, u_t) | \pi_t).$$

Lemma 5.4.8 *The process Π_t is a controlled Markov process.*

Proof: For all (x_{t+1}, y^t, u^t) we have

$$\begin{aligned} p(x_{t+1} | y^t, u^t) &= \int p(x_{t+1} | x_t) p(x_t | y^t, u^t) dx_t \\ &= \int p(x_{t+1} | x_t) \frac{p(y_t | x_t, y^{t-1}, u^t) p(x_t | y^{t-1}, u^t)}{\int p(y_t | \tilde{x}_t, y^{t-1}, u^t) p(\tilde{x}_t | y^{t-1}, u^t) d\tilde{x}_t} dx_t \end{aligned}$$

$$\begin{aligned}
&= \int p(x_{t+1} | x_t) \frac{u_t(y_t | x_t) p(x_t | y^{t-1}, u^{t-1})}{\int u_t(y_t | \tilde{x}_t) p(\tilde{x}_t | y^{t-1}, u^{t-1}) d\tilde{x}_t} dx_t \\
&= \Phi(\pi_t, u_t, y_t)
\end{aligned}$$

for some function Φ . Thus

$$P(d\Pi_{t+1} | \pi^t, u^t) = P(d\Pi_{t+1} | \pi_t, u_t).$$

□

As stated before the problem is a constrained Markov decision problem. We simplify the problem by strengthening the form of the constraint. The resulting optimal rate under formulation 1a and 2a defined below will be an upper bound on the optimal rate in formulation 1 and 2 respectively.

Definition 5.4.4 Let J_1, \dots, J_T be functions on the space of Gaussians random variables on \mathbb{R}^d defined backwards starting with T :

Formulation 1a: Strong Distortion Schedule

where

$$J_T(\pi) = \inf_{u \in \Omega_T(\pi)} \bar{c}(\pi, u)$$

and

$$J_t(\pi) = \inf_{u \in \Omega_t(\pi)} \bar{c}(\pi, u) + \int J_{t+1}(\tilde{\pi}) P(d\tilde{\pi} | \pi, u)$$

where $\Omega_t(\pi) = \{u \in \mathcal{K}_{\mathcal{Y}|\mathcal{X}} \text{ such that } \bar{d}(\pi, u) \leq D_t\}$.

Formulation 2a: Time-Average Strong Distortion

Here we expand the state to (π_t, δ_t) where δ_t represents the distortion accrued up to time $t - 1$. Let

$$J_T(\pi, \delta) = \inf_{u \in \Omega_T(\pi, \delta)} \bar{c}(\pi, u)$$

and

$$J_t(\pi, \delta) = \inf_{u \in \Omega_t(\pi, \delta)} \bar{c}(\pi, u) + \int J_{t+1}(\tilde{\pi}, \tilde{\delta}) P(d\tilde{\pi}, d\tilde{\delta} | \pi, \delta, u)$$

where $\Omega_t(\pi, \delta) = \{u \in \mathcal{K}_{\mathcal{Y}|\mathcal{X}} \text{ such that } \bar{d}(\pi, u) \leq D - \delta\}$. Also $\delta_{t+1} = \delta_t + \bar{d}(\pi_t, u_t)$.

Theorem 5.4.4 In both cases if the infimization is achieved by a policy $\mu^T = (\mu_1, \dots, \mu_T)$ then μ^T is optimal. Furthermore μ^T can be chosen to be a deterministic function of the $\{\pi_t\}$ or $\{\pi_t, \delta_t\}$ processes.

Proof: Theorem 3.2.1 of [HLL]. □

5.4.5 Summary and Extensions

In this section we have proved a sequential source coding theorem over noiseless digital channel for both the finite alphabet and Gaussian case. When the source is Markov we showed that the mutual information optimization problem can be converted into a constrained dynamic programming problem.

Some possible extensions include

- (1) The calculation of error exponents.
- (2) Generalize the sequential rate distortion problem to sources that are trees. Trees are a natural extension of Markov processes.
- (3) Allow for memory in the spatial direction. For example we might be able to treat processes that are Markov random fields evolving in time.

5.5 Successive Refinement

In this section we describe the successive refinement problem. The basic problem is to observe a static source and then successively transmit more information about it over time. At each time step the decoder outputs a better and better reconstruction. Note that this is essentially what we were doing in chapter three. There the initial state uncertainty was successively refined over time.

The problem of successive refinement was first introduced by Equitz and Cover. [EC] We show here that successive refinement problem is just a special case of the sequential rate distortion problem.

5.5.1 Setup and Results

In this section our source is X^N with distribution P_{X^N} .

Definition 5.5.1 *A successive refinement quantizer is a sequence of measurable quantizers f_t such that*

$$f_t : \mathcal{X}^N \times \mathcal{Y}^{t-1,N} \rightarrow \mathcal{Y}^N$$

where the range of each function is at most countable. Specifically f_t takes $(x^N, y^{t-1,N}) \mapsto y_t^N$.

At time t the quantizer f_t has access to the observation X^N and all the previous reconstructions Y_1^N, \dots, Y_{t-1}^N . Just as in the rate distortion case if the channel is noiseless then both the encoder and decoder have access to the past Y_t^N 's. If the channel is noisy then we need to be explicit about whether there is a feedback link or not.

We now formulate the successive refinement problem. The superscript “SR” represents “successive refinement.”

Definition 5.5.2 *The operational successive refinement function is*

$$R_{T,N}^{SR, o}(D_1, \dots, D_T) = \inf_{(f_1, \dots, f_T) \in \mathcal{F}} \frac{1}{NT} H(Y_1^N, \dots, Y_T^N)$$

where $\mathcal{F} = \{(f_1, \dots, f_T) : E_{P_{X^N}} d_N(X^N, Y_t^N) \leq D_t \quad t = 1, \dots, T\}$. (Note we assume $D_1 \geq D_2 \geq \dots \geq D_T$.)

We are interested in minimizing the time-average entropy. One could also ask for the set of all rates that satisfy the distortion schedule. This is more in tune with the original formulation of the successive refinement problem [EC], [Rim]. Computing the acceptable rate region is rather difficult in general. But reducing a criterion on the set of rates to an average rate makes the characterization of the problem easier. Also for the control applications we have in mind we are interested in the average rate needed to achieve some goal.

Now we define the successive refinement function.

Definition 5.5.3 *The successive refinement function is*

$$R_{T,N}^{SR}(D_1, \dots, D_T) = \inf_{\mathcal{F}} \frac{1}{NT} I_{P_{X^N, Y^{T,N}}}(X^N; Y^{T,N})$$

where $\mathcal{F} = \{\{Q(dY_t^N | x^N, y^{t-1,N})\}_{t=1}^T : E_{P_{X^N, Y^N}} d_N(X^N, Y_t^N) \leq D_t \quad t = 1, \dots, T\}$.

It should be clear that the successive refinement problem is just a special case of the sequential rate distortion problem with distortion schedule. Thus the converse theorem, the direct theorem, the dynamic programming formulation, and the extensions to the Gaussian source all continue to hold.

For the sake of convenience we state the coding theorem.

Theorem 5.5.1 *For any $\epsilon > 0$ one can find an $N(\epsilon)$ such that for $N \geq N(\epsilon)$*

$$R_{T,N}^{SR, o}(D_1 + \epsilon, \dots, D_T + \epsilon) \leq R_{T,N}^{SR}(D_1, \dots, D_T) + \epsilon.$$

Furthermore a necessary condition to achieve (D_1, \dots, D_T) over a given channel of capacity C_T is $R_{T,N}^{SR}(D_1, \dots, D_T) \leq \frac{1}{N} C_T$.

Proof: Follows from sections 5.3 and 5.4. \square

5.5.2 Examples

In this section we examine two sources.

- (1) X^N is an N dimensional Gaussian $\mathcal{N}(0, \Lambda_{X^N})$
- (2) X^N is uniformly distributed over the box $[-L, L]^N$.

For the Gaussian source we use the weighted squared error distortion measure. For the uniform source we use the semi-faithful version of the squared error measure.

We want to successively refine both sources with the following distortion schedule: $D_1 \geq D_2 \geq \dots \geq D_T$. For the Gaussian source we compute the infimizing channel and discuss how that channel may be realized. For the uniform source we upper bound the successive refinement rate by computing the rate required for a digital channel. Then as we did in section 5.2.5 we will show how the uniform source result can be used to determine a high rate approximation for the Gaussian source.

Gaussian Source

In this subsection we characterize the infimizing channel for the successive refinement of a Gaussian source.

We first state a preliminary lemma.

Lemma 5.5.1 *Assume that U, V, W are jointly Gaussian vectors. Then*

$$I(U; V|W) = \frac{1}{2} \log \frac{|\Lambda_{U|W}|}{|\Lambda_{U|W,V}|}$$

where $\Lambda_{U|W}$ and $\Lambda_{U|W,V}$ are the conditional covariances.

Proof: This is a straightforward calculation. \square

Assume we have chosen some nondegenerate Gaussian measure $P_{X^N, Y^{T,N}}$. Let $\Lambda_{X^N|Y^{t,N}}$ be the covariance of X^N given observation $Y^{t,N}$. Then from lemma 5.5.1 we can show the following sum is a telescoping sum:

$$\begin{aligned} I(X^N; Y^{T,N}) &= \sum_{t=1}^T I(X^N; Y_t^N | Y^{t-1,N}) \\ &= \sum_{t=1}^T \frac{1}{2} \log \frac{|\Lambda_{X^N|Y^{t-1,N}}|}{|\Lambda_{X^N|Y^{t,N}}|} \\ &= \frac{1}{2} \log \frac{|\Lambda_{X^N}|}{|\Lambda_{X^N|Y^{T,N}}|} \end{aligned}$$

This tells us that the mutual information only depends on the initial covariance and the final covariance. This observation greatly simplifies the successive refinement problem.

Before computing the infimizing channel we consider some other simplifications. By lemmas 5.2.2 and 5.2.3 we can without loss of generality assume that the weight matrix in the squared error distortion measure is identity and that the covariance Λ_{X^N} is diagonal.

Since Λ_{X^N} is diagonal we know by lemma 5.4.1 that we can choose the infimizing Gaussian channel to factor as $P(dY_t^N | x^N, y^{t-1,N}) = \otimes_{n=1}^N P(dY_{t,n} | x_n, y_n^{t-1})$. Thus we can restrict the infimizing channel to be a Gaussian channel consisting of N parallel independent channels.

The successive refinement problem defined in section 5.5.1 can be restated with the above simplifications as:

$$R_{T,N}^{\text{SR}}(D_1, \dots, D_T) = \inf_{\mathcal{F}} \frac{1}{2T} \log \frac{|\Lambda_{X^N}|}{|\Lambda_{X^N|Y^{T,N}}|}$$

where $\mathcal{F} = \{ \{P(dY_t^N | y^{t-1,N}, x^N)\}_{t=1}^T : E_{P_{X^N, Y_t^N}} d_N(X^N, Y_t^N) \leq D_t \quad t = 1, \dots, T$
and the channels are independent parallel Gaussian channels}.




Figure 5-2: Successive Refinement

Because the objective $\frac{1}{2T} \log \frac{|\Lambda_{X^N}|}{|\Lambda_{X^N|Y^{T,N}}|}$ only depends on the initial covariance and the final covariance we see that this problem is essentially the rate distortion problem discussed in section 5.2.4. The difference, though, is that we need to achieve the intermediate distortions also.

There is not a unique solution to this optimization. One solution is to choose a channel $P(dY_1^N | X^N)$ such that $ED(X^N, Y_1^N) \leq D_T$. And then not transmit anything over the next $T - 1$ time steps. Note that in this case all the rate occurs in the first step and no rate occurs in the subsequent steps. We can make the problem more realistic by imposing a peak rate constraint. We will not, though, formulate this problem here.

The non-uniqueness of the solution is related to the same phenomena we saw in chapter three. For the case where we only had uncertainty in the initial condition the role of the encoder and decoder is to essentially deliver a better and better description of the initial state. Over a time horizon T this can be done in many ways while still maintaining the same error at time T . For example we could transmit all RT bits in the first step and zero for all the remaining steps or we could send R bits every step.

Figure 5-2 shows one possible way to successively refine a 5-dimensional Gaussian source with covariance $\text{diag}[\lambda_1, \dots, \lambda_5]$ over a time horizon of 3. The area under η_1 should be $\leq D_1$, the area under η_2 should be $\leq D_2$, and the area under η_3 should equal D_3 .

We present one solution that is based on achieving the distortion D_t exactly at every time step (as opposed to $< D_t$ at every time step.) This solution is based on the basic rate distortion solution described in section 5.2.4.

Recall that $\Lambda = \text{diag}[\lambda_1, \dots, \lambda_N]$ is diagonal. To achieve a distortion D_1 in the first time step we apply directly the results of section 5.2.4.

Thus we get

$$I(X^N; Y_1^N) = \frac{1}{N} \sum_{n=1}^N \frac{1}{2} \log \frac{\lambda_n}{\delta_{1,n}}$$

where

$$\delta_{1,n} = \begin{cases} \eta_1 & \text{if } \eta_1 \leq \lambda_n \\ \lambda_n & \text{if } \eta_1 > \lambda_n \end{cases}$$

where η_1 is chosen such that $\frac{1}{N} \sum_{n=1}^N \delta_{1,n} = D_1$. The backward channel has the form

$$X^N = Y_1^N + V_1^N$$

where V_1^N is distributed normally with mean zero and covariance $\text{diag}[\delta_{1,1}, \dots, \delta_{1,N}]$.

The residual uncertainty in X is captured by V_1^N . Thus at time 2 we refine V_1^N . Thus we get

$$I(X^N; Y_2^N | Y_1^N) = \frac{1}{N} \sum_{n=1}^N \frac{1}{2} \log \frac{\delta_{1,n}}{\delta_{2,n}}$$

where

$$\delta_{2,n} = \begin{cases} \eta_2 & \text{if } \eta_2 \leq \delta_{1,n} \\ \delta_{1,n} & \text{if } \eta_2 > \delta_{1,n} \end{cases}$$

where η_2 is chosen such that $\frac{1}{N} \sum_{n=1}^N \delta_{2,n} = D_2$. The backward channel has the form

$$V_1^N = (Y_2^N - Y_1^N) + V_2^N$$

where V_2^N is distributed normally with mean zero and covariance $\text{diag}[\delta_{2,1}, \dots, \delta_{2,N}]$. Note this implies

$$X^N = Y_2^N + V_2^N.$$

This procedure is repeated for $t = 3, \dots, T$. In general we have

$$V_{t-1}^N = (Y_t^N - Y_{t-1}^N) + V_t^N$$

where V_t^N is distributed normally with mean zero and covariance $\text{diag}[\delta_{t,1}, \dots, \delta_{t,N}]$. This implies

$$X^N = Y_t^N + V_t^N.$$

Now we can convert the backward channels to forward channels as follows

$$Y_1^N = H_1 X^N + W_1^N$$

where $H_1 = E(Y_1^N X^{N'}) E(X^N X^{N'})$ and W_1^N is a zero mean Gaussian vector with covariance $E(Y_1^N Y_1^{N'}) - E(Y_1^N X^{N'}) E(X^N X^{N'})^{-1} E(X^N Y_1^{N'})$. For $t > 1$ we have

$$Y_t^N - Y_{t-1}^N = H_t V_{t-1}^N + W_t$$

where where $H_t = E((Y_t^N - Y_{t-1}^N)V_{t-1}^{N'})E(V_{t-1}^N V_{t-1}^{N'})$ and W_t^N is a zero mean Gaussian vector with covariance $E((Y_t^N - Y_{t-1}^N)(Y_t^{N'} - Y_{t-1}^{N'})) - E((Y_t^N - Y_{t-1}^N)V_{t-1}^{N'})E(V_{t-1}^N V_{t-1}^{N'})^{-1}E(V_{t-1}^N(Y_t^{N'} - Y_{t-1}^{N'}))$. This forward channel can be rewritten as

$$Y_t^N = H_t(X^N - Y_{t-1}^N) + Y_{t-1}^N + W_t.$$

The second addend is the minimum mean squared estimate of the state given $Y^{t-1,N}$. The first addend is a suitably scaled innovation.

There are many ways to realize these successive refinement channels. We will discuss one particular realization of the infimizing channel: the memoryless additive white Gaussian noise channel with noiseless feedback. This realization will be important in section 5.5.4 where we discuss the Schalkwijk-Kailath feedback channel coding scheme.

Proposition 5.5.1 *We are given a scalar Gaussian source $X \sim \mathcal{N}(0, \lambda)$ that we want to successively refine according to the distortion schedule $D_t = \alpha^t \lambda$ where $0 < \alpha < 1$. We can realize the successive refinement channel by an AWGN memoryless channel with noiseless feedback and capacity $C = \frac{1}{2} \log \frac{1}{\alpha}$.*

Proof: Let the AWGN channel have the form $B_t = A_t + V_t$ where $V_t \sim \mathcal{N}(0, 1)$ and power constraint P . We will specify the power P needed to achieve the distortion schedule.

The successive refinement channel is of the form $Y_t = h_t(X - Y_{t-1}) + Y_{t-1} + W_t$ where $h_t = (1 - \alpha)$ and $W_t \sim \mathcal{N}(0, (1 - \alpha)D_t)$.

We now construct a source-channel encoder and channel-source decoder and realize this channel over the $B_t = A_t + V_t$ channel. Let $A_t = g_t(X - Y_{t-1})$ where $g_t = \sqrt{\frac{1-\alpha}{D_t}}$. Note that this source-channel encoder uses channel output feedback Y_{t-1} . And let $Y_t = Y_{t-1} + \sqrt{(1 - \alpha)D_t}B_t$.

We now show that this source-channel encoder and channel-source decoder realize the channel:

$$\begin{aligned} Y_t &= Y_{t-1} + \sqrt{(1 - \alpha)D_t} B_t \\ &= Y_{t-1} + \sqrt{(1 - \alpha)D_t} (A_t + V_t) \\ &= Y_{t-1} + \sqrt{(1 - \alpha)D_t} (g_t(X - Y_{t-1}) + V_t) \\ &= Y_{t-1} + (1 - \alpha)(X - Y_{t-1}) + \sqrt{(1 - \alpha)D_t} V_t \\ &= Y_{t-1} + (1 - \alpha)(X - Y_{t-1}) + W_t \end{aligned}$$

To compute the capacity of the $B_t = A_t + V_t$ channel note that $P = g_t^2 D_{t-1} = \frac{1-\alpha}{\alpha}$. This implies $C = \frac{1}{2}(1 + P) = \frac{1}{2} \log \frac{1}{\alpha}$. \square

In summary this proposition states that a memoryless AWGN channel with noiseless feedback and capacity C can be used to successively refine a source $\mathcal{N}(0, \lambda)$ with distortion schedule $D_t = \lambda 2^{-2CT}$.

Uniform Source

Assume we are given a uniform source with the semi-faithful version of the mean squared error distortion. It is difficult to compute the infimizing law in this case. We can, though, compute an upper bound on the rate for a noiseless digital channel to achieve a distortion schedule $D_1 \geq \dots \geq D_T$. In fact we have already done this in chapter three.

In the first step we can upper bound the number of bits needed to achieve D_1 by $\log(\sqrt{N} \frac{L}{D_1})$. (See section 5.2.4.) For time step two we know the error is in a ball of size at most D_1 . Thus we can upper bound the number bits needed to achieve D_2 by $\log(\sqrt{N} \frac{D_1}{D_2})$. And so on.

The average total number of bits for this scheme is

$$\frac{1}{T} \sum_{t=1}^T \log(\sqrt{N} \frac{D_{t-1}}{D_t}) = \frac{1}{T} \log(\sqrt{N} \frac{L}{D_T})$$

where $D_{-1} \triangleq L$. The peak rate for this scheme is $\max_t \log(\sqrt{N} \frac{D_{t-1}}{D_t})$.

High Rate Approximation

We now show that as the distortion D_1 goes to zero we can transmit the Gaussian source over a digital noiseless channel at essentially the successive refinement rate. We will use the uniform source results to approximate the Gaussian source. Just as we did in section 5.2.5 we will show that as $D_1 \rightarrow 0$ we can approximate the infimizing Gaussian channel by a digital channel.

We first treat the scalar case: $X \sim \mathcal{N}(0, \lambda)$. The vector case is a straightforward extension. Assume that we are given the Gaussian channels $\{Q(dY_t | x, y^{t-1})\}_{t=1}^T$ with $R_t = I(X; Y_t | Y^{t-1})$ and $D_t = E(d(X, Y_t))$ for all $t = 1, \dots, T$.

We will find an L representing the dynamic range of all the quantizers in the successive refinement scheme. For the first time step we use the exact scheme described in section 5.2.5. Then

$$D_1^s \leq \left(\frac{L}{2^{R_1^s}} \right)^2 + 2 \int_L^\infty x^2 p(x) dx.$$

At time two we send a refinement of X . In particular, if X falls into $[-L, L]$ at time one then we know that the distortion after one step will be $\leq \frac{L}{2^{R_1^s}}$. We need only refine this using R_2^s bits. If X did not fall into $[-L, L]$ then we send zero. Thus the distortion at time two is

$$D_2^s \leq \left(\frac{L}{2^{R_1^s} 2^{R_2^s}} \right)^2 + 2 \int_L^\infty x^2 p(x) dx.$$

By repeating this scheme we see that

$$D_t^s \leq \left(\frac{L}{\prod_{i=1}^t 2^{R_i^s}} \right)^2 + 2 \int_L^\infty x^2 p(x) dx \quad \forall t = 1, \dots, T.$$

We need to choose L and R_t^s $t = 1, \dots, T$ such that $D_t^s \leq D_t$ and R_t^s is close to R_t for all t .

By section 5.2.5, let $L^s = \sqrt{4\lambda \ln\left(\frac{2}{D_T}\right)}$. Then for D_T small enough we will have $2 \int_L^\infty x^2 p(x) dx \leq \frac{D_T}{2}$.

Now choose R_t^s such that $\left(\frac{L}{\prod_{i=1}^t 2^{R_i^s}}\right)^2 \leq \frac{D_t}{2}$. Specifically choose $R_1^s = \frac{1}{2} \log\left(\frac{8\lambda \ln\frac{2}{D_T}}{D_1}\right)$. And for $t = 2, \dots, T$ let $R_t^s = \frac{1}{2} \log\frac{D_{t-1}}{D_t}$.

By construction $D_t^s \leq \frac{D_t}{2} + \frac{D_T}{2} \leq D_t$. Now

$$\frac{R_1^s}{R_1} = 1 + \frac{\log(8 \ln(\frac{2}{D_T}))}{\log \frac{\lambda}{D_1}}$$

and for $t = 2, \dots, T$ we have

$$\frac{R_t^s}{R_t} = \frac{\frac{1}{2} \log \frac{D_{t-1}}{D_t}}{\frac{1}{2} \log \frac{D_{t-1}}{D_t}} = 1.$$

Now we will take $D_1 \rightarrow 0$. Because $D_1 \geq D_2 \geq \dots \geq D_T$ we see that $D_2, \dots, D_T \rightarrow 0$. Let us assume that D_T goes to zero with respect to D_1 at a rate slow enough to insure that $\frac{\log \ln \frac{1}{D_T}}{\log \frac{1}{D_1}} \rightarrow 0$. Note that this is not a stringent assumption. (In section 5.5.3 we shed light on this assumption.) Under this assumption we see that all the ratios of the rates converge to 1. Thus we can achieve the successive refinement rate over a digital channel in the limit of high rate.

5.5.3 Rate of Convergence

In the successive refinement formulation in section 5.5.1 we start with a distortion schedule and ask what is the minimum rate required to achieve those distortions? We can also ask the opposite question. For a given rate how fast can the distortions in the distortion schedule decrease? We answer this question here as well as make some connections to chapter three.

In chapter three we examined the case where the only uncertainty in the system occurred in the initial condition. The basic idea there was to send a better and better description of the initial condition. Here we ask how fast can the D_1, D_2, \dots converge to zero while maintaining a finite rate per time step. Specifically how fast can D_t converge to zero while maintaining: $\limsup_{T \rightarrow \infty} \frac{1}{T} I(X; Y_1, \dots, Y_T) < \infty$. The basic result is that for a finite rate the distortion can not decrease faster than exponentially in t .

Lemma 5.5.2 *Let X be a \mathbb{R} valued random variable admitting a density. Under the square error distortion measure the distortion rate function is bounded as*

$$\frac{2^{2h(X)}}{2\pi e} 2^{-2R} \leq D(R) \leq \text{cov}(X) 2^{-2R}$$

with equality when X is Gaussian. (Where $h(X)$ is the differential entropy.)

Proof: This follows from equation 4.3.15 and theorem 4.3.3 of [Berg]. \square

Suppose we are given a channel $\{Q(dB_t | a^t, b^{t-1})\}_{t=1}^{\infty}$, a source-channel encoder, and channel-source decoder such that $\sup_t \frac{1}{t} I(X; Y_1, \dots, Y_t) = R < \infty$. Then by lemma 5.5.2 we have

$$D_t \geq \frac{2^{2h(X)}}{2\pi e} 2^{-2tR} \quad \forall t.$$

Thus for any channel with finite capacity the end-to-end distortion across that channel cannot decrease faster than exponentially in t .

Rate of Convergence of the Chapter Two Tracking Scheme

We now re-formulate the basic problem of chapter two as a successive refinement problem. The basic problem there concerns itself with finding conditions so that the estimation error goes to zero. We know we cannot drive the distortion to zero faster than exponentially in time. Furthermore the system dynamics are increasing the size of the error at each time step by the matrix A . Thus the rate has to be larger than a certain measure on the eigenvalues of the matrix A . In proposition 3.5.1 of chapter three a scheme was proposed with error decaying exponentially.

We now convert that problem into a successive refinement problem. For a given rate we find an upperbound on how fast the distortion can decrease in the distortion schedule. This upperbound is computed for the scheme described in chapter three. Here, though, we allow the initial state to be distributed according to a given distribution.

The scheme presented in chapter three first grew the dynamic range of the quantizer until it “captured” the state. The dynamic range has to grow faster than the dynamics. After the state is “captured” the scheme starts sending quantization information.

Let us look at the scalar version of that scheme. Let P_X be a measure on \mathbb{R} . The dynamics of the system are $X_{t+1} = aX_t$ and X_0 is distributed according to P_X . We want $|X_k - \hat{X}_t|$ to go to zero. This is essentially the same thing as asking $|a^t X_0 - a^t \hat{X}_0(t)|$ to go to zero. This, though, is nothing more than a successive refinement problem where the distortion schedule D_t has to go to zero faster than a^{-t} .

We showed that if we know a bound on L where $X_0 \in [-L, L]$ then D_t decreases as $D_t \leq L(\frac{a}{2^R})^t$ (where $R > \log a$.) Now we give the rate of convergence of the expected distortion when we don't have known bounds on the support. We give an upper bound by analyzing the scalar version of the quantization scheme. The dynamic range growing scheme is as follows: first start the dynamic range at some L and grow it until we capture the state. Let the initial position, x , be drawn from P_X . Define $D_t(x)$ to be the distortion accrued at time t when the initial position is x . We can bound this by

$$D_t(x) \leq (L2^{c(x)R})\left(\frac{a}{2^R}\right)^{t-c(x)}$$

where $c(x)$ represents the first time we capture the state. Specifically $c(x)$ is the smallest integer t such that $a^t x < L2^{tR}$. The first term in the product above represents the “effective”

initial dynamic range. The second term represents the decrease in distortion once we have captured the state. Clearly $R > \log a$ is a necessary condition for convergence.

We are interested in

$$\begin{aligned} D_t &= \int D_t(x) P(dx) \\ &\leq \int L 2^{c(x)R} \left(\frac{a}{2^R}\right)^{t-c(x)} P(dx) \\ &= L \int \left(\frac{a}{2^R}\right)^t \left(\frac{2^{2R}}{a}\right)^{c(x)} P(dx) \end{aligned}$$

Now $c(x) \leq \frac{\log \frac{x}{L}}{\log \frac{2^R}{a}} + 1$. Thus

$$\begin{aligned} D_t &\leq L \int \left(\frac{a}{2^R}\right)^t \left(\frac{2^{2R}}{a}\right)^{c(x)} P(dx) \\ &\leq L \int \left(\frac{a}{2^R}\right)^t \frac{2^{2R}}{a} \left(\frac{x}{L}\right)^{1+\frac{R}{R-\log a}} P(dx) \\ &= \alpha(R, L, t) \int x^{(1+\frac{R}{R-\log a})} P(dx) \end{aligned}$$

Where

$$\alpha(R, L, t) = \frac{2^{2R}}{a L^{\frac{R}{R-\log a}}} \left(\frac{a}{2^R}\right)^t$$

is a function that decreases exponentially to zero in t for fixed L, R and $R > \log a$.

Thus if the $1 + \frac{R}{R-\log a}$ moment of X exists then D_t decreases at a rate $2^{-t(R-\log a)}$. Note that

$$\lim_{R \rightarrow \log a} 1 + \frac{R}{R - \log a} = \infty$$

and

$$\lim_{R \rightarrow \infty} 1 + \frac{R}{R - \log a} = 2.$$

Thus this scheme requires at least finite variance to work. The generalization to the vector case is straightforward.

In chapter three we showed that for any initial condition we could drive the error to zero. Here we have refined those results. First we showed that the problem can be posed a successive refinement problem. And second we showed that if an appropriate moment condition holds then the average distortion, D_t , decreases at a rate $2^{-t(R-\log a)}$.

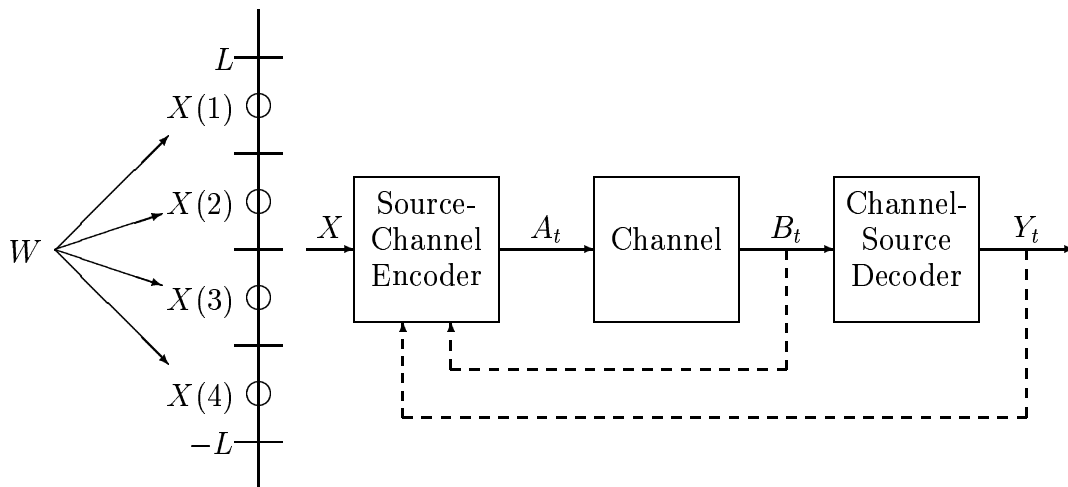


Figure 5-3: General Schalkwijk-Kailath Scheme

5.5.4 The Relationship between Successive Refinement and Feedback Channel Capacity

In this section we show that the Schalkwijk-Kailath feedback channel coding scheme is a special case of the successive refinement problem. [SK] We give new insight into this feedback channel coding problem and the double exponent error property. Recall from chapter four that we left open the question of how much better the channel coding error exponent can be under feedback as compared to the exponent without feedback.

First we describe a generalized form of the Schalkwijk-Kailath scheme. Assume that we want to send one of 2^{RT} messages over a given channel over a time horizon T with arbitrary small probability of decoding error. The idea is to divide an interval $[-L, L]$ into 2^{RT} uniform regions. We map each message w to the centroid of one of the regions.

Let the messages w be drawn uniformly with probability $\frac{1}{2^{RT}}$. Then let X be the random variable taking values in $[-L, L]$ according to whichever message was chosen. This P_X is our source and we will successively refine it. Now assume we are given a channel $\{P(dB_t | a^t, b^{t-1})\}_{t=1}^T$. The generalized Schalkwijk-Kailath scheme consists of designing a source-channel encoder and channel-source decoder so as to minimize the distortion in the reconstruction of X at time T . See figure 5-3. (The dotted lines in the figure represent possible feedback links.)

The probability of decoding error is the probability that Y_T falls outside the region the source X lives in. Specifically Chebychev's inequality shows us

$$\begin{aligned} \Pr(\text{decoding error}) &= \Pr\left(|Y_T - X| > \frac{L}{2^{RT}}\right) \\ &\leq \frac{2^{2RT} D_T}{L^2} \end{aligned}$$

where $D_T = E((Y_T - X)^2)$.

Given a particular channel, $\{Q(dB_t | a^t, b^{t-1})\}_{t=1}^T$, it is in general rather difficult to determine the optimal encoding and decoding scheme to minimize D_T . We can, though, show the following:

Proposition 5.5.2 *There exists a channel of capacity C such that*

$$\Pr(\text{decoding error}) \leq 2^{-2T(C-R)}$$

Proof: From lemma 5.5.2 we know there exists a channel with capacity C such that $D(CT) \leq \text{cov}(X)2^{-CT}$. Thus

$$\begin{aligned} \Pr(\text{decoding error}) &\leq \Pr\left(|Y_T - X| > \frac{L}{2^{RT}}\right) \\ &\leq \frac{2^{2RT} D_T}{L^2} \\ &\leq \frac{\text{cov}(X)}{L^2} 2^{-2T(C-R)} \\ &\leq 2^{-2T(C-R)} \end{aligned}$$

□

By lemma 5.5.2 we also know that D_T cannot decrease faster than exponentially. Thus the bound $\frac{2^{2RT} D_T}{L^2}$ cannot decrease faster than exponentially. How then does one recover the double exponential exponent found in the work of Schalkwijk-Kailath? By using Chernoff's inequality instead of Chebychev's inequality. Specifically:

$$\begin{aligned} \Pr(\text{decoding error}) &= \Pr\left(|Y_T - X| > \frac{L}{2^{RT}}\right) \\ &\leq e^{-\sup_{\theta \geq 0} \left(\theta \frac{L^2}{2^{2RT}} - \log E(e^{\theta(Y_T - X)^2})\right)} \end{aligned}$$

In general this large deviation exponent is difficult to compute. But we can say something for the case of Gaussian channels with affine encoders.

If we restrict our channels to be Gaussian channels with or without noiseless feedback and our source-channel encoder to be affine then we can show that $P(dY_t | X = x)$ is a Gaussian distribution for all x . Thus

$$\Pr\left(|Y_T - X| > \frac{L}{2^{RT}}\right) = \sum_x \Pr\left(|Y_T - x| > \frac{L}{2^{RT}} | X = x\right) P(x). \quad (5.4)$$

Furthermore $E((Y_T - x)^2 | X = x) = D_T$ independent of x . To see this note that by corollary 5.4.1 we can treat the source as mismatched to the Gaussian channel. In this case the conditional distributions of the channel outputs Y_t conditioned on X are Gaussian.

The following lemma tells us how to bound the deviations of a Gaussian random variable. Note that we also have a lower bound.

Lemma 5.5.3 *Let Z be a scalar zero mean Gaussian random variable $\mathcal{N}(0, \lambda)$. Then*

(1)

$$\Pr(|Z| > M) \leq e^{-\frac{M^2}{2\lambda}}$$

(2) *If $\frac{M}{\sqrt{\lambda}} \geq 1$ then*

$$\frac{\sqrt{\lambda}}{M} e^{-\frac{M^2}{2\lambda}} \leq \Pr(|Z| > M)$$

Proof: This is proposition 2.2.1 of [Dud]. \square

The following proposition shows that over AWGN memoryless channels with feedback the error exponent decreases at a doubly exponential rate. (Recall that the capacity of a memoryless channel is not increased under feedback.) We also show that for the Schalkwijk-Kailath scheme the exponent cannot decrease faster than double exponentially.

Proposition 5.5.3 *For a memoryless AWGN channel of capacity C and noiseless feedback the probability of channel error under the Schalkwijk-Kailath scheme can be bounded as*

$$2^{-T(C-R)} e^{-\frac{2^{2T(C-R)}}{2}} \leq \Pr(\text{decoding error}) \leq e^{-\frac{2^{2T(C-R)}}{2}}$$

where $R < C$ and the lower bound holds for T large enough.

Proof: By proposition 5.5.1 and the mismatch corollary 5.4.1 we know that we can use the memoryless AWGN channel with feedback to successively refine the source P_X with distortion schedule $D_t = E_P(X^2)2^{-2Ct}$. Where $E_P(X^2) = \sum_{i=1}^{2^R} \frac{1}{2^R} \left((-1 - \frac{1}{2^R} + \frac{i}{2^{R-1}})L \right)^2 < L^2$.

Now apply lemma 4.5.3 with $M = \frac{L}{2^{RT}}$, $\lambda = D_T = E_P(X^2)2^{-2CT}$, $C > R$, and T large enough we have

$$\frac{\sqrt{E_P(X^2)}}{L} 2^{-T(C-R)} e^{-\frac{L^2}{2E_P(X^2)} 2^{2T(C-R)}} \leq \Pr\left(|Y_T - x| \geq \frac{L}{2^{RT}} | X = x \right) \leq e^{-\frac{L^2}{2E_P(X^2)} 2^{2T(C-R)}}$$

Then by equation (5.4) the proposition follows. \square

We have reproduced the double exponent channel error of the Schalkwijk-Kailath scheme. Furthermore we have shown that one cannot do better than the double exponent on Gaussian channels with affine encoders. Finally we have shown that the generalized Schalkwijk-Kailath scheme is really a special case of the successive refinement problem.

5.6 Solution Characterization and High Rate Approximation

In this section we will compute the sequential rate distortion function for Gauss-Markov sources under differing information patterns. We discuss what happens when the Gauss-Markov source is unstable. At the the end we discuss the high rate case.

Recall a Gauss-Markov source can be defined recursively as

$$X_1 \sim \mathcal{N}(0, \Lambda_{X_1}) \quad \text{and} \quad X_{t+1} = AX_t + Z_t \quad (5.5)$$

where $Z_t \sim \mathcal{N}(0, \Lambda_{Z_t})$. Furthermore assume that we are using the squared error distortion measure with identity as the weight matrix.

5.6.1 Sequential Rate Distortion for Gauss-Markov Sources

When computing the sequential rate distortion infimizing channel we can, by lemma 4.4.6, restrict our attention to simplified causal sequence of Gaussian stochastic kernels. We now show that Gaussian channels of the form, $\{Y_t = \alpha_t x_t + \beta_t y^{t-1} + W_t\}_{t=1}^T$, can be realized over memoryless AWGN channels with noiseless feedback. This result is a generalization of the construction given in proposition 5.5.1.

Proposition 5.6.1 *Gaussian channels of the form, $\{Y_t = \alpha_t x_t + \beta_t y^{t-1} + W_t\}_{t=1}^T$, can be realized over memoryless AWGN channels with noiseless feedback.*

Proof: Assume we are given a Gauss-Markov source P_{X^T} . Let the source-channel encoder have the form

$$A_t = K_{W_t}^{-\frac{1}{2}} \alpha_t (x_t - E(X_t|y^{t-1})).$$

Note that it is a function of the past y^{t-1} . Let the channel-source decoder have the form

$$Y_t = \alpha_t E(X_t|y^{t-1}) + \beta_t y^{t-1} + K_{W_t}^{\frac{1}{2}} b_t.$$

And let the AWGN channel at time t have the form

$$B_t = A_t + V_t$$

where $V_t \sim \mathcal{N}(0, I)$ with an average power constraint

$$P_t = E(A_t' A_t) = \text{trace} \left(K_{W_t}^{-\frac{1}{2}} \alpha_t \Lambda_{X_t|Y^{t-1}} \alpha_t' K_{W_t}^{-\frac{1}{2}} \right)$$

where $\Lambda_{X_t|Y^{t-1}} = E((x_t - E(X_t|y^{t-1}))(x_t - E(X_t|y^{t-1}))')$ is the error covariance in estimating X_t from Y^{t-1} .

We now show that these source-channel encoders, memoryless AWGN channels, and channel-source decoders can be used to realize the original channel.

$$\begin{aligned} Y_t &= \alpha_t x_t + \beta_t y^{t-1} + W_t \\ &= \alpha_t (x_t - E(X_t|y^{t-1})) + \alpha_t E(X_t|y^{t-1}) + \beta_t y^{t-1} + W_t \end{aligned}$$

$$\begin{aligned}
&= \alpha_t E(X_t|y^{t-1}) + \beta_t y^{t-1} + K_{W_t}^{\frac{1}{2}} \left(K_{W_t}^{-\frac{1}{2}} \alpha_t (x_t - E(X_t|y^{t-1})) + V_t \right) \\
&= \alpha_t E(X_t|y^{t-1}) + \beta_t y^{t-1} + K_{W_t}^{\frac{1}{2}} (a_t + V_t)
\end{aligned}$$

□

Proposition 5.6.2 *For the channel realization above the mutual information separates into a sum of the mutual informations for each memoryless AWGN channel: $I(X^T, Y^T) = \sum_{t=1}^T I(A_t; B_t)$.*

Proof:

$$\begin{aligned}
I(X^T, Y^T) &= \sum_{t=1}^T I(X^t; Y_t | Y^{t-1}) \\
&= \sum_{t=1}^T I(X_t; Y_t | Y^{t-1}) \\
&= \sum_{t=1}^T I(A_t; B_t | Y^{t-1}) \\
&= \sum_{t=1}^T I(A_t; B_t)
\end{aligned}$$

Where the second equality comes from using a simplified causal sequence of Gaussian stochastic kernels. The third equality comes from the fact that A_t is a function of X_t and Y^{t-1} . Similarly Y_t is a function of B_t and Y^{t-1} . The fourth equality comes from noting that $(X_t - E(X_t|Y^{t-1}))$ is independent of Y^{t-1} . □

Sequential Rate Distortion Solution

We first solve the sequential rate distortion function in closed form for the scalar Gauss-Markov process. We then solve the vector case in the high rate regime.

For the scalar case our source is of the form:

$$X_1 \sim \mathcal{N}(0, \lambda_{X_1}) \quad \text{and} \quad X_{t+1} = aX_t + Z_t \quad (5.6)$$

where $Z_t \sim \mathcal{N}(0, \lambda_{Z_t})$.

Let $\{B_t = A_t + V_t\}_{t=1}^T$ be a sequence of memoryless AWGN channels that realize, under noiseless feedback, the sequential rate distortion infimizing channel. Let $R_t = I(A_t; B_t)$. From subsection 5.2.4 we know that we can reconstruct a Gaussian source with variance λ over a Gaussian channel of rate R with distortion $\lambda 2^{-2R}$.

Now the source-channel encoder computes the innovation $x_t - E(X_t|y^{t-1}) = ax_{t-1} + z_{t-1} - ay_{t-1}$. It then scales it and transmits it over the $A_t - B_t$ channel. Thus the reconstruction has distortion for $t \geq 1$:

$$D_t = \left(a^2 D_{t-1} + \lambda_{Z_{t-1}} \right) 2^{-2R_t} \quad \text{with } D_0 = 0, \lambda_{Z_0} = \lambda_{X_1}.$$

Thus to achieve the distortion schedule $\{D_t\}_{t=1}^T$ we require

$$R_T^{\text{SRD}}(D_1, \dots, D_T) = \frac{1}{T} \sum_{t=1}^T R_t = \frac{1}{T} \sum_{t=1}^T \max \left\{ 0, \frac{1}{2} \log \left(\frac{a^2 D_{t-1} + \lambda_{Z_{t-1}}}{D_t} \right) \right\}.$$

In the case where $D_t = D$, $\forall t$ and $\lambda_{Z_t} = \lambda_Z$, $\forall t$ we have

$$\lim_{T \rightarrow \infty} R_T^{\text{SRD}}(D, D, \dots, D) = \max \left\{ 0, \frac{1}{2} \log \left(a^2 + \frac{\lambda_Z}{D} \right) \right\}.$$

We comment here that even if the source is unstable we can still achieve the distortion schedule as long as the rate is large enough.

Now we treat the case of an N -dimensional vector valued Gauss-Markov source. We treat only the high rate (low distortion) regime. In this high rate regime we know, by the results in subsection 5.2.4, that we can reconstruct a Gaussian source with covariance Λ over a matched Gaussian channel of rate R with an error covariance $\Lambda 2^{-\frac{2R}{N}}$. The resulting distortion is $2^{-\frac{2R}{N}} \text{trace}(\Lambda)$

Thus, just as before we get, the recursion

$$\Lambda_t = (A \Lambda_{t-1} A' + \Lambda_{Z_{t-1}}) 2^{-\frac{2R_t}{N}} \quad \text{with } \Lambda_0 = 0, \Lambda_{Z_0} = \Lambda_{X_1}.$$

Where Λ_t represents the error covariance matrix. The distortion is $D_t = \text{trace}(\Lambda_t)$.

In the case when $\Lambda_{Z_t} = \Lambda_Z$ and $R_t = R$ we get

$$\Lambda_T = (2^{-\frac{R}{N}} A)^T \Lambda_{X_1} (2^{-\frac{R}{N}} A)^{T'} + \sum_{t=1}^T (2^{-\frac{R}{N}} A)^{t-1} \Lambda_Z (2^{-\frac{R}{N}} A)^{t-1'}.$$

Since we are in the high rate regime we can assume that R is large enough so that $2^{-\frac{R}{N}} A$ is a stable matrix. Thus in steady state the distortion per time step is equal to

$$\text{trace} \left(\sum_{t=1}^{\infty} (2^{-\frac{R}{N}} A)^{t-1} \Lambda_Z (2^{-\frac{R}{N}} A)^{t-1'} \right).$$

Unfortunately we cannot go much further than this. It is very difficult to get closed form solutions for the SRD function for the vector-valued Gauss-Markov case.

Source-Channel Mismatch

So far we have been assuming that the “ $A - B$ ” channel with noiseless feedback is matched to the SRD infimizing channel. What happens if the “ $A - B$ ” channel is still an AWGN channel but no longer matched to the source? For example the dimension of the source and the dimension of the channel, i.e. the number of parallel channels, may not be equal.

It is very difficult to solve this problem for the general case. See the work of Basar and Bansal. [BB] In general the optimal source-channel encoder and channel-source decoder will be nonlinear.

One can compute the optimal affine source-channel encoder and affine channel-decoder. This is done by Lee and Peterson. [LP] There they provide conditions on the eigenvalues of the source covariance and the channel noise covariance to insure that they are matched. In the case when they are not matched they show how to approximately match the source and channel.

5.6.2 Sequential Rate Distortion for Differing Information Patterns

We have shown that the SRD infimizing channel can realized over memoryless AWGN channels with noiseless feedback. Often times in practice this feedback is unavailable to the encoder. (Recall from chapter two the definitions of encoders with differing information patterns.)

We are interested in channels that can be realized over memoryless AWGN channels without feedback. We now define two variants of the sequential rate distortion problem with differing information patterns. Recall definitions 5.3.5 and 5.3.6 where we defined the SRD problem.

- (1) Memoryless Sequential Rate Distortion (MSRD)

In this case the source-channel encoder at time t is independent of X^{t-1} and Y^{t-1} .

- (2) Innovation Sequential Rate Distortion (ISRD)

In this case the source-channel encoder at time t transmits information only about the innovation Z_t . (Recall equation (5.5).)

In the MSRD we do not allow the source-channel encoder access to the past channel outputs or inputs. In the ISRD formulation the source-channel encoder is allowed to transmit information only about the innovation Z_t . The ISRD problem was first formulated by Borkar and Mitter in [BM].

We now characterize the solutions for these two variants of the sequential rate distortion problem.

Memoryless Sequential Rate Distortion (MSRD)

First we solve the scalar case. We then solve the vector case in the high rate regime. The source-channel encoder is independent of the past source values and past source reconstructions. Thus the best distortion we can hope to achieve over an AWGN of rate R_t at time step t is $D_t = \lambda_{X_t} 2^{-2R_t}$. Specifically the encoder looks at the marginal statistics of X_t at time t and then encodes that. In the N -dimensional vector case and high rate regime we get $\Lambda_t = \Lambda_{X_t} 2^{\frac{-2R_t}{N}}$. And thus $D_t = \text{trace}(\Lambda_t)$.

Note that if A is unstable then the covariance of the state Λ_{X_t} is growing exponentially. Thus we can never hope to achieve a finite distortion D_t at a finite rate for all time. In the case when A is stable and $\Lambda_{Z_t} = \Lambda_Z$, $\forall t$ we know Lyapunov's equation defines the steady state covariance:

$$\Lambda_X = A\Lambda_X A' + \Lambda_Z.$$

If we use the same rate at each time step then the steady state reconstruction error covariance satisfies

$$\Lambda = A\Lambda A' + \Lambda_Z 2^{-\frac{2R}{N}}.$$

In the scalar case we can compute explicitly the steady state covariance $\lambda_X = \frac{\lambda_Z}{1-a^2}$. In the case where $D_t = D$, $\forall t$ and $\lambda_{Z_t} = \lambda_Z$, $\forall t$ we have

$$\lim_{T \rightarrow \infty} R_T^{\text{MSRD}}(D, D, \dots, D) = \max \left\{ 0, \frac{1}{2} \log \left(\frac{\lambda_Z}{(1-a^2)D} \right) \right\}.$$

Innovation Sequential Rate Distortion (ISRD)

In this case the source-channel encoder computes the innovation Z_t and transmits that over an AWGN channel of rate R_t .

We first treat the scalar case. The innovation can be reconstructed over an AWGN channel of rate R_t at a distortion $\tilde{D}_t = \lambda_{Z_t} 2^{-2R_t}$. We use the tilde to represent the distortion on the innovation as opposed to the distortion on the source reconstruction. It should be clear that the optimal channel-source decoder computes the estimate of the state as follows $Y_{t+1} = aY_t + \tilde{Z}_t$. Thus the distortion follows the recursions $D_{t+1} = a^2 D_t + \lambda_{Z_t} 2^{-2R_t}$. For the vector case and high rate regime we get $\Lambda_{t+1} = A\Lambda_t A' + \Lambda_{Z_t} 2^{-\frac{2R_t}{N}}$. With distortion $D_t = \text{trace}(\Lambda_t)$.

Note that if A is unstable then the distortion D_t is growing exponentially. Thus we can never hope to achieve a finite distortion D_t at a finite rate for all time.

In the case when A is stable, $\Lambda_{Z_t} = \Lambda_Z \forall t$, and we use the same rate at each time step we know Lyapunov's equation define the steady state covariance:

$$\Lambda = A\Lambda A' + \Lambda_Z 2^{-\frac{2R}{N}}.$$

Note that this is exactly the same as the MSRD solution. Thus we have just shown that the ISRD and MSRD formulations are equivalent. This curious phenomena states that in the innovation scheme the rate gained by decorrelating the process into innovations is equal to the rate lost in trying to account for the accumulating errors. Thus there is no gain in coding the innovation as opposed to coding the state with respect to its marginal distribution.

A similar result for block-coding of Gauss-Markov sources is shown by Berger. See theorem 6.3.3 of [Berg]. Specifically he shows for the high rate regime that coding the source as a block and coding the innovations as a block lead to the same rate-distortion functions. We have just shown a sequential version of the Berger result continues to hold.

In summary the ISRD and MSRD formulations are equivalent. Furthermore for unstable sources there does not exist a finite rate, as there did in the SRD formulation, that insures finite distortion for all time. This suggests that for unstable sources one cannot achieve finite distortion at a finite rate for all time if there is no feedback to the source-channel encoder. See Sahai's thesis for more discussion along these line. [Sah]

Operational Equivalence

The curious phenomena of the ISRD and MSRD equivalence continues to hold for the scalar stable Gauss-Markov source and a noiseless digital channel.

Proposition 5.6.3 *The operational rates for the ISRD problem and the MSRD problem for scalar Gauss-Markov processes are the same.*

Proof: Our source is a scalar Gauss-Markov process $X_{t+1} = aX_t + Z_t$ where $|a| < 1$ and $Z_t \sim N(0, \lambda_Z)$. Assume we are in steady state. Then $X_t \sim N(0, \frac{\lambda_Z}{1-a^2})$.

Let Q_X be the optimal quantizer for X_t . That is it achieves the smallest entropy while constraining average distortion to be at most D . Let it's rate be R_{Q_X} . Given that we have the optimal quantizer for the random variable $X \sim N(0, \frac{\lambda_Z}{1-a^2})$ can we determine the optimal quantizer for another random variable $Z \sim N(0, \lambda_Z)$? Yes. For the innovation scheme we know the quantizer Q_Z that we use to quantize the innovation must be constrained to have distortion less than or equal to $D(1-a^2)$.

Let the optimal quantizer be $Q_X = \{\{R_i\}, \{q_i\}\}$ where R_i 's represent the quantizing regions and the q_i represent the centroids of the regions. Then $D_{Q_X} = \sum_i \int_{R_i} (x - q_i)^2 p(x) dx$. Let $Q_Z = \{\{\frac{1}{\sqrt{1-a^2}}R_i\}, \{\frac{1}{\sqrt{1-a^2}}q_i\}\} = \{\{S_i\}, \{v_i\}\}$. By symmetry this is the optimal quantizer with rate $R_{Q_Z} = R_{Q_X}$ and distortion $D_{Q_Z} = \sum_i \int_{S_i} (z - v_i)^2 p(z) dz$. Now in terms of distribution $X = \frac{1}{\sqrt{1-a^2}}Z$. After substitution we see $D_{Q_Z} = \sum_i \int_{S_i} (z - v_i)^2 p(z) dz = \sum_i \int_{R_i} (\sqrt{1-a^2}x - \sqrt{1-a^2}q_i)^2 \frac{p(x)}{\sqrt{1-a^2}} \sqrt{1-a^2} dx = \sum_i \int_{R_i} (1-a^2)(x - q_i)^2 p(x) dx = D_{Q_X}(1-a^2)$. Thus for the same rate quantizers we achieve the same distortion whether we use the innovation coding scheme or the memoryless coding scheme. \square

It is not clear, though, if the result continues to hold for the vector-valued Gauss-Markov case.

5.6.3 High Rate Asymptotics

We conclude this section by showing that we can achieve the sequential rate distortion rate over a digital noiseless channel in the limit of low distortion.

Uniform Sources

First we give a quick discussion of the SRD for uniform sources. Let $x_1 \in \Omega \subset \mathbb{R}^N$ where Ω is a bounded set. Let $\|z_t\| \leq L$. And let the dynamics follow $x_{t+1} = Ax_t + z_t$. We want

to achieve a distortion $\|x_t - y_t\|_2^2 \leq D_t$ for each t . This is nothing more than the problem of observability under bounded disturbances treated in chapter three. Specifically proposition 3.5.2. gives an upper bound on the rate required to insure a given distortion schedule.

High Rate Asymptotics for the Gauss-Markov Source

We now treat the problem of high rate asymptotics for the Gauss-Markov case. We solve the scalar case. The vector case is more difficult and is left to future work. The main ideas, though, are captured in the scalar case. The source is $X_{t+1} = aX_t + Z_t$ with $X_1 \sim \mathcal{N}(0, \lambda_{X_1})$ and $Z_t \sim \mathcal{N}(0, \lambda_Z)$. Assume without loss of generality that $a \geq 0$.

Fix a finite horizon T . We want to achieve a distortion D at each time step while maintaining a rate close to $R_1 = \frac{1}{2} \log \left(\frac{\lambda_{X_1}}{D} \right)$ and $R_t = \frac{1}{2} \log \left(a^2 + \frac{\lambda_Z}{D} \right)$ for all $t = 2, \dots, T$.

We will define two dynamic ranges $[-L_{X_1}, L_{X_1}]$ and $[-L_Z, L_Z]$. The idea is that if $X_1 \in [-L_{X_1}, L_{X_1}]$ and for all t the innovation $Z_t \in [-L_Z, L_Z]$ then we will reconstruct the source with a distortion bounded by $\frac{D}{2}$. We will show that the distortion accrued by falling outside these dynamic ranges can be bounded by $\frac{D}{2}$.

We want to choose R_1 so that $\left(\frac{L_{X_1}}{2^{R_1}} \right)^2 \leq \frac{D}{2}$. And we will choose R_t for $t = 2, \dots, T$ so that $\left(\frac{a\sqrt{\frac{D}{2}} + L_Z}{2^{R_t}} \right)^2 \leq \frac{D}{2}$. Thus let $R_1^s = \frac{1}{2} \log \frac{2L_{X_1}^2}{D}$ and $R_t^s = \log \left(a + \sqrt{\frac{2}{D}} L_Z \right)$. (Recall “s” represent “scheme.”)

In the following let $G_Z(L_Z) = 2 \int_{L_Z}^{\infty} z^2 p(z) dz$ where $p(z) \sim \mathcal{N}(0, \lambda_Z)$ and $G_{X_1}(L_{X_1}) = 2 \int_{L_{X_1}}^{\infty} x^2 p(x) dx$ where $p(x) \sim \mathcal{N}(0, \lambda_{X_1})$. Then

$$D_1^s \leq \left(\frac{L_{X_1}}{2^{R_1^s}} \right)^2 + G_{X_1}(L_{X_1}),$$

$$D_2^s \leq \left(\frac{a\sqrt{\frac{D}{2}} + L_Z}{2^{R_2^s}} \right)^2 + a^2 G_{X_1}(L_{X_1}) + G_Z(L_Z),$$

and in general

$$D_t^s \leq \left(\frac{a\sqrt{\frac{D}{2}} + L_Z}{2^{R_t^s}} \right)^2 + a^{2(t-1)} G_{X_1}(L_{X_1}) + \sum_{i=0}^{t-2} a^{2i} G_Z(L_Z).$$

By our choice of the R_t^s 's the first term on the right hand side of each inequality is equal to $\frac{D}{2}$ and represents the distortion accrued when the source has fallen into the dynamic range of the quantizer. The subsequent terms represent the distortion due to the quantizer overflow.

We now bound the distortion due to overflow. To that end we choose L_{X_1}, L_Z such that

$$\max_{t=1, \dots, T} a^{2(t-1)} G_{X_1}(L_{X_1}) + \sum_{i=0}^{t-2} a^{2i} G_Z(L_Z) \leq \frac{D}{2}.$$

From section 4.2.5 we know there exist L_Z, L_{X_1} large enough so that $G_Z(L_Z) \leq e^{-\frac{L_Z^2}{4\lambda_Z}}$ and $G_{X_1}(L_{X_1}) \leq e^{-\frac{L_{X_1}^2}{4\lambda_{X_1}}}$. Assume L_{X_1} and L_Z are large enough to satisfy these inequalities. Then choose L_{X_1} so that

$$\max_{t=1, \dots, T} a^{2(t-1)} G_{X_1}(L_{X_1}) \leq \frac{D}{4}.$$

Specifically if $a < 1$ then let $L_{X_1} = \sqrt{4\lambda_{X_1} \ln \frac{4}{D}}$ and if $a > 1$ then let $L_{X_1} = \sqrt{4\lambda_{X_1} \ln \frac{4}{a^{2(T-1)}D}}$ and if $a > 1$. Also choose L_Z so that

$$\max_{t=1, \dots, T} \sum_{i=0}^{t-2} a^{2i} G_Z(L_Z) \leq \frac{D}{2}.$$

Specifically let $L_Z = \sqrt{4\lambda_Z \ln \frac{4(a^{2(T-1)}-1)}{(a^2-1)D}}$.

Now upon substitution for L_{X_1} we get

$$R_1^s = \frac{1}{2} \log \frac{8\lambda_{X_1} \ln \frac{4}{D}}{D} \text{ or } R_1^s = \frac{1}{2} \log \frac{8\lambda_{X_1} \ln \frac{4}{a^{2(T-1)}D}}{D}$$

depending on whether $a \leq 1$ or $a > 1$ respectively. For $t = 2, \dots, T$ we get

$$R_t^s = \log \left(a + \sqrt{\frac{8\lambda_Z}{D} \ln \frac{4(a^{2(T-1)}-1)}{(a^2-1)D}} \right).$$

By construction we have achieved the distortion D at each time step. We need only check that the rate of our scheme is close to the SRD rates. Specifically we need to show that $\frac{R_t^s}{R_t}$ goes to one as D goes to zero. That $\lim_{D \rightarrow 0} \frac{R_1^s}{R_1} = 1$ follows from section 5.2.5. We check the case for $t = 2, \dots, T$. We have

$$\begin{aligned} \lim_{D \rightarrow 0} \frac{R_t^s}{R_t} &= \lim_{D \rightarrow 0} \frac{\log \left(a + \sqrt{\frac{8\lambda_Z}{D} \ln \frac{4(a^{2(T-1)}-1)}{(a^2-1)D}} \right)}{\frac{1}{2} \log \left(a^2 + \frac{\lambda_Z}{D} \right)} \\ &= 1 + \lim_{D \rightarrow 0} \frac{\log \left(\frac{\sqrt{\frac{8\lambda_Z}{D} \ln \frac{4(a^{2(T-1)}-1)}{(a^2-1)D}}}{\sqrt{\frac{\lambda_Z}{D}}} \right)}{\frac{1}{2} \log \left(\frac{\lambda_Z}{D} \right)} \\ &= 1 + \lim_{D \rightarrow 0} \frac{\log \left(8 \ln \frac{4(a^{2(T-1)}-1)}{(a^2-1)D} \right)}{\log \left(\frac{\lambda_Z}{D} \right)} \end{aligned}$$

The second term goes to zero as D goes to zero.

Thus we have shown in the high rate case that we can achieve the SRD rates over a noiseless digital channel.

5.7 Summary

In this chapter we formulated the sequential rate distortion problem. This problem is a generalization of the traditional rate distortion problem to processes over time. We showed that the sequential rate distortion provided a general framework for viewing rate distortion and successive refinement as special cases.

For Markov source we applied the tools of dynamic programming to characterize the conditional channel laws in the rate distortion problem as “policies.” These infimizing laws were called *matched* channels. We showed that in general the separation between source and channel coding does not hold for small delays. But when the source and channel are matched we can achieve the sequential rate distortion bounds.

We provided explicit solutions to the sequential rate distortion function for Gauss-Markov processes. We showed that unstable Gauss - Markov processes cannot be transmitted across a noisy channel at the rate determined by the rate distortion function unless there is feedback.

We reexamined the Schalkwijk-Kailath feedback coding scheme by representing it as a successive refinement problem with exponentially decaying distortion schedule. The successive refinement formulation was also used to compute bounds on the coding schemes presented in chapter three.

We provided high rate asymptotics for the sequential rate distortion function.

Chapter 6

Control of Stochastic Systems Under Communication Constraints

6.1 Introduction

In this chapter we examine the stochastic control problem when there is a communication channel connecting the sensor to the controller. See figure 6-1. This problem arises when the plant and the controller are geographically separated and there is a noisy or band-limited communication channel connecting them. One example of this occurs in remote control over wireless links. The analysis of this problem will require the tools developed in all of the preceding chapters.

We first formulate the control problem using the framework introduced in chapter two. The system consists of a plant, an encoder, a channel, a decoder, and a controller. The plant and the channel are given to us. We must design the encoder, decoder, and controller to satisfy some control objective. We look only at the case where there is a communication channel between the sensor and the controller. The link between the controller and the plant is assumed to be noiseless. The problem is already difficult with one communication link and thus we leave the case of two communication links to future work. Many of the insights, though, provided for the one communication link case carry over to the more general case.

In the previous chapters we have treated the problem of encoder/decoder design for the problem of channel coding with the objective of minimizing the probability of error and for the problem of joint source-channel coding with the objective of minimizing the end-to-end distortion. For the control problem we need to design the encoder/decoder pair as well as the controller. We assume centralized design. By the results in chapter two we know this problem can be solved, in principle, via dynamic programming. In general, though, the joint optimization of the encoder/decoder pair and the controller is hard to solve.

We specialize the control problem to the linear quadratic Gaussian (LQG) control problem. We show that if the encoder and decoder are equi-memory and the encoder has access to the controls then the “control” aspect and the “communication” aspect of the problem can be separated. See the dashed box in figure 6-1. The traditional separation theorem for the partially observed LQG problem states that the controller and the estimator can be separated. Specifically we can apply a certainty equivalent controller to the state estimate.

Here we will show that a more general notion of separation holds. Under the hypothesis mentioned the certainty equivalent controller is still optimal. The encoder and decoder are designed to provide the “best” state estimate. The controller design is independent of the channel, the encoder, and the decoder. It will turn out that the “communication” aspect of the problem can be reduced to a sequential rate distortion problem over a given noisy channel with a particular weight matrix tuned to the underlying Riccati equation. These ideas will be discussed in more detail in the sequel.

We examine two kinds of channels: the digital noiseless channel and the additive white Gaussian noise (AWGN) channel. For the digital channel we can use the operational sequential rate distortion function to bound the end-to-end distortion for a given channel rate. For the AWGN channel we can use the sequential rate distortion function to bound the end-to-end distortion for a given channel capacity. Furthermore under conditions of joint source-channel matching we can show this bound to be tight. The optimal quadratic cost then decomposes into two pieces: a full knowledge cost and a sequential rate distortion cost.

In summary the main contribution of this chapter is the application of our previous results to the problem of stochastic control under a communication constraint. We show for the LQG problem, under suitable conditions, that a generalized separation principle continues to hold.

In section 6.2 we formulate the stochastic control problem using the framework of chapter two. In section 6.3 we examine the tracking and LQG problems. In section 6.4 we conclude.

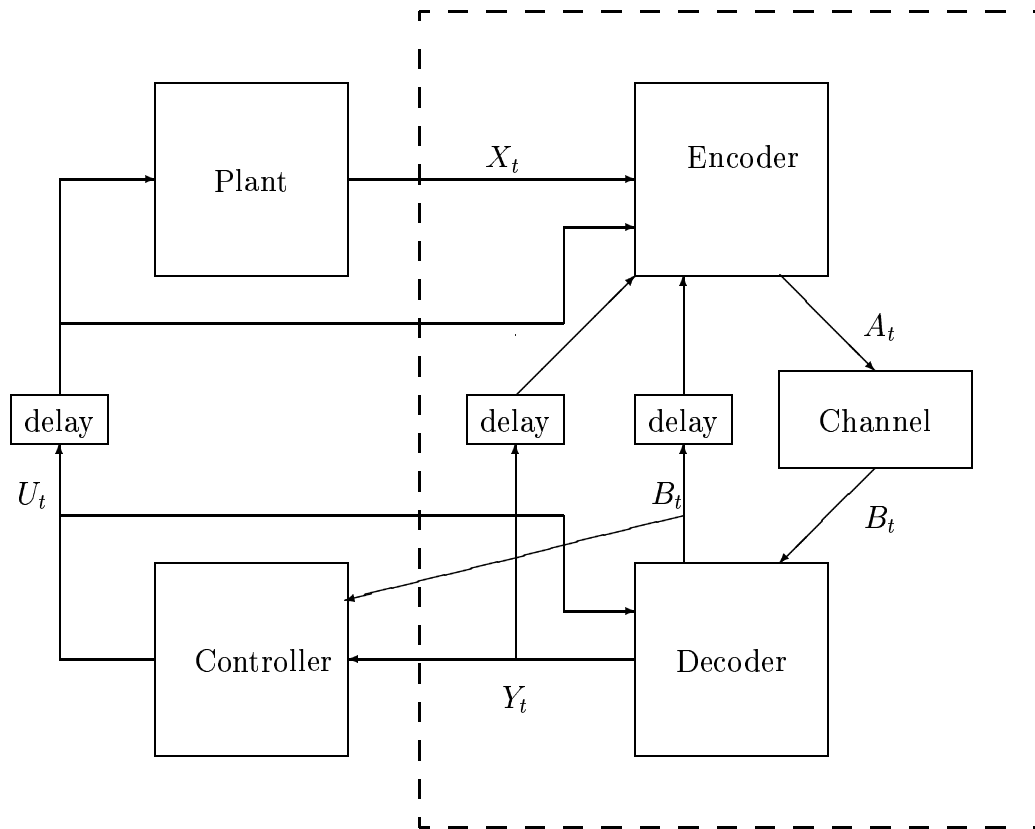


Figure 6-1: System

6.2 Stochastic Control Problem

The general setup consists of a plant, an encoder, a channel, a decoder, and a controller. There are five different kinds of signals: state X_t , channel input A_t , channel output B_t , decoder output Y_t , and control U_t . The time-ordering is

$$X_1, A_1, B_1, Y_1, U_1, \dots, X_T, A_T, B_T, Y_T, U_T.$$

Now we define where the signals live, the system specifications, and the information pattern. The state and decoder output X_t, Y_t are \mathbb{R}^d -valued processes. The control U_t is a \mathbb{R}^m -valued process. Let the channel input and channel output live in $A_t \in \mathcal{A}$ and $B_t \in \mathcal{B}$ respectively. We will treat two cases: \mathcal{A} is a finite set or $\mathcal{A} = \mathbb{R}^l$. The former will be used when we treat the noiseless digital channel and the latter will be used when we treat the AWGN channel. Furthermore we assume that the channel input space equals the channel output space: $\mathcal{A} = \mathcal{B}$. See figure 6-1.

Recall, from definition 2.2.1, that a model of a system is the set of all joint measures P that are consistent with the time-ordering, system specification, and information pattern. We describe the different system and decision kernels now.

Plant

To define the plant we need to specify the stochastic kernels

$$Q(dX_{t+1} | x^t, a^t, b^t, y^t, u^t), \quad t = 1, \dots, T$$

We are interested in the time-invariant, linear Gaussian plant. Thus we can equivalently write these kernels as

$$X_1, \quad X_{t+1} = Fx_t + Gu_t + W_t, \quad t = 1, \dots, T-1$$

where $\{W_t\}$ are IID Gaussian $\sim \mathcal{N}(0, K_W)$, the initial position $X_1 \sim \mathcal{N}(0, K_X)$, and F, G are system matrices of suitable dimensions.

Note that X_{t+1} is independent of $X^{t-1}, A^t, B^t, Y^t, U^{t-1}$ given X_t, U_t .

Channel

We treat two channels: the digital noiseless channel and the time-invariant, memoryless, AWGN channel. Abstractly each channel can be treated as a sequence of stochastic kernels of the form

$$Q(dB_t | x^t, a^t, b^{t-1}, y^{t-1}, u^{t-1}), \quad t = 1, \dots, T.$$

For the digital channel we can equivalently write these kernels as

$$Q(dB_t | x^t, a^t, b^{t-1}, y^{t-1}, u^{t-1}) = \delta_{(B_t=a_t)}.$$

And for the time-invariant, memoryless, AWGN channel we can equivalently write these kernels as

$$B_t = a_t + V_t$$

where the $\{V_t\}$ are IID and distributed as $V_t \sim \mathcal{N}(0, K_V)$. Furthermore there is a power constraint: $E_P(A_t' A_t) \leq L, \quad \forall t$. Where P represents the overall joint measure. (We have not completely specified it yet.)

Note that B_t independent of $X^t, A^{t-1}, B^{t-1}, Y^{t-1}, U^{t-1}$ when given A_t .

We have just defined the system specifications. The encoder, decoder, and controller are designed by the designer. Now we need to describe their information patterns.

Encoder

The encoder is specified by the designer

$$Q(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1}, u^{t-1}), \quad t = 1, \dots, T.$$

We make no restrictions on the information pattern here. The restrictions will come when we define equi-memory. Note that in general there are five kinds of feedback to the encoder: the state, the past channel inputs, the past channel outputs, the past decoder outputs, and the past controls. In the sequel we will discuss how these feedback paths may be realized. If the encoder is allowed to see the past controls then in the parlance of chapter three the encoder is in encoder class one. Similarly if it cannot observe the control then it is in encoder class two.

Decoder

The decoder is specified by the designer

$$Q(dY_t | x^t, a^t, b^t, y^{t-1}, u^{t-1}), \quad t = 1, \dots, T.$$

The information pattern here stipulates that Y_t be independent of X^t, A^t given B^t, Y^{t-1}, U^{t-1} .

Controller

The controller is specified by the designer

$$Q(dU_t | x^t, a^t, b^t, y^t, u^{t-1}), \quad t = 1, \dots, T.$$

The information pattern here stipulates that U_t be independent of X^t, A^t given B^t, Y^t, U^{t-1} .

In summary our model consists of all joint measures P such that

(1) Plant:

$$\begin{aligned} & P(dX_{t+1} | X^t = x^t, A^t = a^t, B^t = b^t, Y^t = y^t, U^t = u^t) \\ &= P(dX_{t+1} | X_t = x_t, U_t = u_t) \\ &= Q(dX_{t+1} | x_t, u_t) \quad P(X^t, A^t, B^t, Y^t, U^t) - a.s. \quad t = 1, \dots, T-1 \end{aligned}$$

(2) Channel:

$$\begin{aligned} & P(dB_t | X^t = x^t, A^t = a^t, B^{t-1} = b^{t-1}, Y^{t-1} = y^{t-1}, U^{t-1} = u^{t-1}) \\ &= P(dB_t | A_t = a_t) \\ &= Q(dB_t | a_t) \quad P(X^t, A^t, B^{t-1}, Y^{t-1}, U^{t-1}) - a.s. \quad t = 1, \dots, T \end{aligned}$$

(3) Encoder: there are no restrictions on the information pattern.

$$P(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1}, u^{t-1}) \quad t = 1, \dots, T$$

(4) Decoder: there are restrictions on the information pattern.

$$P(dY_t | x^t, a^t, b^t, y^{t-1}, u^{t-1}) = P(dY_t | b^t, y^{t-1}, u^{t-1}) \quad t = 1, \dots, T$$

(5) Controller: there are restriction on the information pattern.

$$P(dU_t | x^t, a^t, b^t, y^t, u^{t-1}) = P(dU_t | b^t, y^t, u^{t-1}) \quad t = 1, \dots, T$$

The resulting joint measure factors as

$$\begin{aligned} & P(dX^T, dA^T, dB^T, dY^T, dU^T) \\ = & \bigotimes_{t=1}^T \left\{ P(dX_{t+1} | x^t, a^t, b^t, y^t, u^t) \otimes P(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1}, u^{t-1}) \right. \\ & \quad \otimes P(dB_t | x^t, a^t, b^{t-1}, y^{t-1}, u^{t-1}) \otimes P(dY_t | x^t, a^t, b^t, y^{t-1}, u^{t-1}) \\ & \quad \left. \otimes P(dU_t | x^t, a^t, b^t, y^t, u^{t-1}) \right\} \\ = & \bigotimes_{t=1}^T \left\{ Q(dX_{t+1} | x_t, u_t) \otimes Q(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1}, u^{t-1}) \otimes Q(dB_t | a_t) \right. \\ & \quad \left. \otimes Q(dY_t | b^t, y^{t-1}, u^{t-1}) \otimes Q(dU_t | b^t, y^t, u^{t-1}) \right\} \end{aligned}$$

This completes the specification of the linear Gaussian model under communication constraints.

Centrally Designed Model

We assume that the encoder, decoder, and controller all have complete system and policy knowledge. Thus, by definition 2.2.5, the system is a centrally designed model.

Performance

A common objective is to minimize the average cost

$$\frac{1}{T} E_P \sum_{t=1}^T c(X_t, U_t)$$

for some integrable running cost $c(\cdot, \cdot)$. In principle we can solve this problem by dynamic programming. Unfortunately this is very difficult.

In the next subsection we examine the quadratic cost problem. Under suitable restrictions on the information patterns of the encoder and decoder we will show that a general separation principle holds.

6.3 Linear Quadratic Gaussian Problem

In chapter three we defined the notion of equi-memory. We extend that definition to the stochastic case. This equi-memory condition allows the encoder and decoder to be coordinated in their actions. We then define expectation predictive encoders and decoders. This allows both the encoder and the decoder the ability to remove the effect of the control signals U_t .

Definition 6.3.1 *An encoder and decoder are said to be equi-memory if there exists a sigma-field $\sigma_t \subset \sigma(B^{t-1}, Y^{t-1}, U^{t-1})$ such that the stochastic kernels defining the encoders*

$$Q(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1}, u^{t-1})$$

are $\sigma(X_t) \times \sigma_t$ -measurable. The stochastic kernels defining the decoders

$$Q(dY_t | b^t, y^{t-1}, u^{t-1})$$

are $\sigma(B_t) \times \sigma_t$ measurable.

In words this states that the encoder at time t is a measure on \mathcal{A} parameterized by x_t and some measurable function $f(b^{t-1}, y^{t-1}, u^{t-1})$ and the decoder is a measure on \mathcal{Y} parameterized by b_t and the same function $f(b^{t-1}, y^{t-1}, u^{t-1})$. Thus the encoder and decoder use the same information excepting that the encoder also observes X_t and the decoder also observes B_t .

Definition 6.3.2 *An equi-memory encoder and decoder with $\sigma_t = \sigma(Y^{t-1}, U^{t-1})$ are called an expectation predictive encoder and an expectation predictive decoder respectively if the encoder is of the form*

$$A_t = g \left(x_t - E(X_t | b^{t-1}, y^{t-1}, u^{t-1}) \right)$$

for some measurable function g and the decoder is of the form

$$Y_t = E(X_t | b^t, y^{t-1}, u^{t-1}).$$

Note that in this case both the encoder, $Q(dA_t | x^t, a^{t-1}, b^{t-1}, y^{t-1}, u^{t-1})$, and the decoder, $Q(dY_t | b^t, y^{t-1}, u^{t-1})$, are Dirac measures (i.e. they are functions.)

The encoder has access to the past y 's. This requires a dedicated link between the decoder output and the encoder. Now Y_t is a function of (b^t, y^{t-1}, u^{t-1}) . Thus we can send B_t to the encoder instead of sending Y_t . Upon receiving B_t the encoder can compute Y_t (recall we have assumed centralized design.) For the noiseless digital channel case we do not need a special link because the encoder can compute y^{t-1} locally.

In the noiseless digital channel case the function g is quantizer that is applied to the *innovation*. In the AWGN channel case the function g is a gain matrix magnifying the value of the innovation while maintaining an average power constraint.

Let the *estimation error* be denoted $\Delta_t \triangleq X_t - Y_t$.

Lemma 6.3.1 *For equi-memory expectation predictive encoders and decoders the error, Δ_t , is uncorrelated with U^{t-1} .*

Proof:

First note

$$\begin{aligned}\Delta_{t+1} &= FX_t + GU_t + W_t - E(FX_t + GU_t + W_t \mid B^{t+1}, Y^t, U^t) \\ &= F\Delta_t + W_t - E(F\Delta_t + W_t \mid B^{t+1}, Y^t, U^t)\end{aligned}$$

Thus $E(\Delta_{t+1}) = 0$. Now

$$\begin{aligned}E(\Delta_{t+1} \mid U^t) &= E[F\Delta_t + W_t - E(F\Delta_t + W_t \mid B^{t+1}, Y^t, U^t) \mid U^t] \\ &= E[F\Delta_t + W_t \mid U^t] - E[E(F\Delta_t + W_t \mid B^{t+1}, Y^t, U^t) \mid U^t] \\ &= E[F\Delta_t + W_t \mid U^t] - E[F\Delta_t + W_t \mid U^t] \\ &= 0\end{aligned}$$

Thus we see that Δ_{t+1} is uncorrelated with U^t . \square

6.3.1 Tracking

We treat the tracking problem first.

Definition 6.3.3 *A linear Gaussian system with a communication channel is said to be trackable, independent of control, at distortion D , if there exists a channel encoder and decoder such that the squared state estimation error $E(\|X_t - Y_t\|^2) \leq D$ for all control sequences and times t .*

The following proposition gives a lower bound on the capacity of a channel in order to achieve trackability. This is essentially the converse theorem for sequential rate distortion.

Proposition 6.3.1 *Let $R_T^{Seq}(D, D, \dots, D)$ be the sequential rate distortion function for the $\{X_{t+1} = Fx_t + W_t\}_{t=1}^T$ source. A necessary condition on the capacity of the channel for the linear Gaussian system to be trackable, independent of control, at distortion D is $C_T \geq R_T^{Seq}(D, D, \dots, D)$.*

Proof: Note that if we use the zero controller, i.e. a controller that only outputs the zero control, then the problem of trackability reduces to a sequential rate distortion problem for the source $\{X_{t+1} = Fx_t + W_t\}_{t=1}^T$. By theorem 5.3.2 we know that a lower bound on the capacity of the channel that achieves a distortion D at each time step is $R_T^{Seq}(D, D, \dots, D)$.

Since this is a lower bound on the capacity when we apply the all zero control it must be a lower bound on the capacity for the system to be trackable, independent of control, at distortion D . \square

The previous proposition provides a lower bound on the capacity of the channel needed to achieve trackability independent of control at distortion D . Note that this necessary condition holds independently of whether we use an equi-memory predictive encoder and decoder or not.

The next proposition gives upper bounds on the capacity needed for trackability.

Proposition 6.3.2 *The following sufficient conditions on capacity hold*

- (1) *If the channel is a digital channel with rate $C_T \geq R_T^{\text{Seq}, 0}(D, \dots, D)$ then there exists an encoder and decoder such that the linear Gaussian system is trackable, independent of control, at distortion D . (Where here $R_T^{\text{Seq}, 0}$ is the operational sequential rate distortion function for the source $\{X_{t+1} = Fx_t + W_t\}_{t=1}^T$.)*
- (2) *If the channel is a time-invariant, memoryless, AWGN channel with capacity $C_T \geq R_T^{\text{Seq}}(D, \dots, D)$ and furthermore the channel is matched to the source $\{X_{t+1} = Fx_t + W_t\}_{t=1}^T$ at distortion D then there exists an encoder and decoder such that linear Gaussian system is trackable, independent of control, at distortion D .*

Proof: For both channels assume the encoder and decoder are equi-memory and expectation predictive. By lemma 6.3.1 the error, Δ_t , is uncorrelated with the past U^{t-1} . Thus the problem reduces to sequential rate distortion on the source: $\{X_{t+1} = Fx_t + W_t\}_{t=1}^T$. For the digital channel we know that $R_T^{\text{Seq}, 0}(D, \dots, D)$ is an achievable rate. For an the AWGN channel that is matched to the source we know that $R_T^{\text{Seq}}(D, \dots, D)$ is an achievable rate. \square

Note that in the AWGN channel case the decoder is required to feedback B_t to the encoder in order to realize the sequential rate distortion infimizing channel.

6.3.2 Quadratic Performance

Now we treat the LQG problem. Our goal is to minimize the long term average cost

$$\limsup_{T \rightarrow \infty} \frac{1}{T} E \left[\sum_{t=1}^T X_t' Q X_t + U_t' S U_t \right] \quad (6.1)$$

where Q is positive semidefinite and S is positive definite.

Assume that the pair (F, G) is controllable. Then under full state observation it is well known that the optimal steady state control law is a linear gain of the form $U_t = L X_t$ where

$$L = -(G' P G + S)^{-1} G' P F X \quad (6.2)$$

and P satisfies the Riccati equation

$$P = F' \left(P - PG (G'PG + S)^{-1} G'P \right) F + Q. \quad (6.3)$$

Furthermore the optimal cost is

$$E(W'PW) = \text{tr}(PK_W) \quad (6.4)$$

These standard results can be found in [Bert].

Proposition 6.3.3 *For expectation predictive encoders and decoders the optimal control law, for the quadratic cost, is the certainty equivalent control law. Specifically the optimal steady state control law is of the form $U_t = LY_t$ where L is given by (6.2).*

Proof: Bar-Shalom and Tse prove a general theorem that states that the certainty equivalent controller is optimal if and only if the state estimation error is uncorrelated with the past controls. [BT] We have shown in lemma 6.3.1 that the state estimation error, Δ_t , is uncorrelated with the past controls. \square

We will now convert the problem into a fully observed LQ problem with Y_t being the new state process. This approach follows Borkar and Mitter. [BM] Note that

$$\begin{aligned} Y_{t+1} &= X_{t+1} - \Delta_{t+1} \\ &= FY_t + GU_t + (F\Delta_t + W_t - \Delta_{t+1}) \\ &= FY_t + GU_t + \tilde{W}_t \end{aligned}$$

where $\tilde{W}_t = (F\Delta_t + W_t - \Delta_{t+1})$. Our new system has dynamics $Y_{t+1} = FY_t + GU_t + \tilde{W}_t$.

We need to show that $\{\tilde{W}_t\}$ are uncorrelated.

Lemma 6.3.2 *For expectation predictive encoders and decoders the random variables $\{\tilde{W}_t\}$ are uncorrelated.*

Proof: We need to show that $E(\tilde{W}_t \tilde{W}_s') = 0$ for all s, t . We will prove $E(\tilde{W}_t \tilde{W}_{t+1}') = 0$. The more general case will then follow.

Since we are using expectation predictive encoders and decoders we know \tilde{W}_t is uncorrelated with Δ_{t+1} . Furthermore \tilde{W}_t and Δ_{t+1} are uncorrelated with W_{t+1} . Thus \tilde{W}_t is uncorrelated with $F\Delta_{t+1} + W_{t+1}$. But $F\Delta_{t+1} + W_{t+1} = \Delta_{t+2} + \tilde{W}_{t+1}$. Thus \tilde{W}_t is uncorrelated with $\Delta_{t+2} + \tilde{W}_{t+1}$.

Once again since we are using expectation predictive encoder and decoders we know \tilde{W}_{t+1} is uncorrelated with Δ_{t+2} . Thus it must be the case that \tilde{W}_t is uncorrelated with \tilde{W}_{t+1} . \square

The running cost for the original problem can be decomposed as

$$E(X_t' Q X_t + U_t' S U_t) = E(Y_t' Q Y_t + U_t' S U_t) + E(\Delta_t' Q \Delta_t).$$

The optimal cost under the Y_t process is, by equation (6.4), equal to $\text{tr}(PK_{\tilde{W}})$.

Assume that $E(\Delta_t \Delta_t') = H$ for all t . By propositions 6.3.1 and 6.3.2 we have both necessary conditions and sufficient conditions on the channel for this to occur. Then the optimal cost for the original problem is

$$\begin{aligned} \limsup_{T \rightarrow \infty} \frac{1}{T} E\left[\sum_{t=1}^T X_t' Q X_t + U_t' S U_t\right] &= \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T E(Y_t' Q Y_t + U_t' S U_t) + E(\Delta_t' Q \Delta_t) \\ &= \text{tr}(PK_{\tilde{W}}) + \text{tr}(QH) \\ &= \text{tr}(PK_W) + \text{tr}((F' P F - P + Q)H) \end{aligned}$$

The optimal cost decomposes into two terms. The first term is the full state cost and the second term depends only on H the state estimation error covariance. Thus we have reduced the problem of computing the optimal cost to that of minimizing $\text{tr}((F' P F - P + Q)H)$ over a given channel. But this latter problem is nothing more than a sequential rate distortion problem with squared error distortion and weight matrix $(F' P F - P + Q)$. This problem was treated in chapter five.

In the scalar case we can compute this cost explicitly. Specifically for channels with capacity $R > \max\{0, \log |F|\}$ we have

$$\text{Ave. Cost} \geq PK_W + \frac{K_W (F^2 P - P + Q)}{2^{2R} - F^2}$$

Equality is achieved if our channel is a matched AWGN channel. For the digital channel we need to replace the second addend with the operational distortion at rate R . If $|F| > 1$ and $R < \log |F|$ then the cost equals infinity. In summary we have characterized the tradeoff between the channel capacity, R , and the LQG performance.

6.4 Summary

In this chapter we formulated a stochastic control problem with a communication link connecting the sensor to the controller. At this generality the problem is difficult to solve. We simplified the problem by finding conditions that allow us to separate the “control” part from the “communication” part. Specifically we examined the LQG problem with communication constraints. We showed that under expectation predictive encoders and decoders the certainty equivalent controller is optimal. Thus the “control” part is indeed separated from the “communication” part. The solution to the “communication” part was shown to be equivalent to the solution of a sequential rate distortion problem. Finally we showed that the optimal cost separates into two pieces: a full knowledge cost and a sequential rate distortion cost.

Chapter 7

Summary

This thesis explores the distributed control problem where there are communication channels between the different components of the system. In this setting, traditional information theory, which codifies the fundamental limitations of reliable communication over noisy channels, is not directly applicable. The reason is traditional information theory is asymptotic, has delays, and does not completely deal with feedback. Since feedback is an essential element of control in the presence of uncertainty and since delays have to be taken into account in control problems, especially for unstable systems, it is natural to look for a unification of information theory and stochastic control when components of the control system are interconnected through communication channels.

We provide such a unification in chapter two where we define our general model. This unified view of control and communication clarifies many of the conceptual issues underlying the distributed control problem. Our model of a system is defined to be the set of all probability measures that satisfy the three specifications:

- (1) A time ordering on the variables of interest.
- (2) A specification of the stochastic kernels representing the plants and channels in the system.
- (3) A specification of the information patterns for the different decision variables.

The control problem becomes one of choosing a sequence of controller stochastic kernels to interconnect the partially specified system kernels so as to satisfy some control objective.

This framework allows us to clarify the different kinds of knowledge: knowledge of the signals, knowledge of the system and policies, and knowledge of the objective or objectives. We show that within this framework and under the assumption of a centralized design we can formulate a broad class of distributed control problems as dynamic programming problems. Furthermore we can use this framework to understand the fundamental limits to performance in distributed systems when there are channel constraints.

In chapter three we examine the deterministic control problem with a communication channel from the sensor, measuring system variables at the plant, to the controller. We formulate the control problem and discuss the role of different information patterns. We provide a lower bound on the rate required to achieve different control objectives that is

independent of the information pattern used. We show that this rate can be achieved for a particular class of information patterns. We introduce the concept of covering number.

In chapter four we treat the feedback channel coding problem. We provide a rather general coding theorem. We show that the directed mutual information is the correct notion of mutual information when computing the capacity of a channel. It allows us to move from code-function distributions to channel input distributions. Next we show that for Markov channels the capacity optimization problem can be formulated as a dynamic programming problem. Lastly we provide a directed version of the data processing inequality.

In chapter five we examine the sequential rate distortion problem. We provide a coding theorem for noiseless digital channels. For Markov sources we show that the capacity optimization problem can be formulated as a dynamic programming problem. We also examine the problem of joint source-channel coding. We discuss the role of *matching*.

In chapter six we put all the pieces together to examine a stochastic control problem with a communication channel connecting the sensor, at the plant, to the controller. We show conditions on the information pattern that effectively separate the design of the controller from the encoder and decoder. Specifically for the LQG control problem we give conditions for the optimality of the certainty equivalent controller.

Appendix A

Background and Supporting Results

A.1 Chapter 2

Stochastic Kernels

A Polish space, \mathcal{X} , is a complete, separable, metric space. We endow it with the topology induced by the metric. Let X be a random variable defined on $(\Omega, \mathcal{F}, \mathcal{P})$ and taking values in the Polish space $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ where $\mathcal{B}(\mathcal{X})$ is the Borel σ -field. We use lowercase “ x ” to represent a particular realization of the random variable X . Define the induced probability measure on \mathcal{X} as $P_X(A)$ where $A \in \mathcal{B}(\mathcal{X})$. Specifically $P_X(A) \triangleq \mathcal{P}(\{\omega : X(\omega) \in A\})$.

We now provide the basic definition and theorems for stochastic kernels. For more details the reader should see section 7.4.3 of [BS].

Definition A.1.1 *Let \mathcal{X}, \mathcal{Y} be Polish spaces. A stochastic kernel on \mathcal{Y} given \mathcal{X} is a function $P_{Y|X}(\cdot|\cdot)$ such that*

- (a) $P_{Y|X}(\cdot|x)$ is a probability measure on \mathcal{Y} for each fixed $x \in \mathcal{X}$
- (b) $P_{Y|X}(B|\cdot)$ is a measurable function on \mathcal{X} for each fixed $B \in \mathcal{B}(\mathcal{Y})$.

Theorem A.1.1 *Let $\mathcal{X}_1, \mathcal{X}_2, \dots$ be a sequence of Polish spaces and, for $n = 1, 2, \dots$ define $\mathcal{Y}_n = \mathcal{X}_1 \times \dots \times \mathcal{X}_n$ and $\mathcal{Y} = \prod_{i \geq 1} \mathcal{X}_i$. Let ν be an arbitrary measure on \mathcal{X}_1 and for every $n = 1, 2, \dots$ let $P_n(dx_{n+1}|y_n)$ be a stochastic kernel on \mathcal{X}_{n+1} given y_n . Then there exists a unique probability measure P_ν on \mathcal{Y} such that, for every measurable rectangle $B_1 \times \dots \times B_n \in \mathcal{Y}_n$ we have*

$$P_\nu(B_1 \times \dots \times B_n) = \int_{B_1} \nu(dx_1) \int_{B_2} P_1(dx_2|x_1) \cdots \int_{B_n} P_{n-1}(dx_n|x_1, \dots, x_{n-1}).$$

Proof: See [BS] proposition 7.28. \square

Theorem A.1.2 *Assume the hypothesis of the last theorem. Given a joint measure $P_\nu(B_1 \times \dots \times B_n)$ we can disintegrate it as*

$$P_\nu(B_1 \times \dots \times B_n) = \int_{B_1} \nu(dx_1) \int_{B_2} P_1(dx_2|x_1) \cdots \int_{B_n} P_{n-1}(dx_n|x_1, \dots, x_{n-1})$$

where each stochastic kernel $P_i(dX_i|x_1, \dots, x_{i-1})$ is determined uniquely $P_\nu(dx_1, \dots, dx_{i-1})$ -almost surely. And $\nu(dX_1)$ is just the marginal of the joint measure.

Proof: See [BS] corollary 7.27.2. \square

Thus we can integrate stochastic kernels and disintegrate joint measures in a well defined and almost surely unique way. We will use the shorthand

$$P_\nu(dX_1, \dots, dX_n) = \nu(dX_1) \otimes P_{X_2|X_1}(dX_2|x_1) \otimes \cdots \otimes P_{X_n|X_1, \dots, X_{n-1}}(dX_n|x_1, \dots, x_{n-1})$$

(where each stochastic kernel is defined uniquely almost surely as described in theorem A.1.2) For a given stochastic kernel we may suppress the subscript and write $P_{Y|X}(dY|x)$ as $P(dY|x)$.

We are given a measure $P_{X,Y,Z}$. Define $P_{X|Y}P_{Z|Y}P_Y$ to be the measure

$$P_{X|Y}P_{Z|Y}P_Y(B_x, B_y, B_z) \triangleq \int_{B_y} P(B_x|y)P(B_z|y)P(dy)$$

where $P(dX|y)$ and $P(dZ|y)$ are the regular conditional probabilities of $P_{X,Y,Z}$.

Definition A.1.2 *The random variables X, Y, Z form a Markov chain, denoted $X - Y - Z$, if $P_{X,Y,Z} = P_{X|Y}P_{Z|Y}P_Y$.*

A.2 Chapter 3

Key Technical Lemma

Let $A \in \mathbb{R}^{d \times d}$. Then by theorem 3.4.1 A has a real Jordan canonical form.

Lemma 3.4.1 $H^t \Upsilon H^{-t} = \Upsilon$

Proof: $H^t \Upsilon H^{-t}$ is the product of three block diagonal matrices. Thus we need only check that it holds for each of the blocks. The blocks come in two types: those associated with real eigenvalues and those associated with complex conjugate eigenvalues. For the the real eigenvalue case H_j is identity. Thus clearly $I^t J_j I^{-t} = J_j$. Let us examine the complex conjugate eigenvalue case:

$$\begin{aligned}
 & H_j^t J_j H_j^{-t} \\
 = & \begin{bmatrix} r(\theta)^{-t} & & & \\ & r(\theta)^{-t} & & \\ & & \dots & \\ & & & r(\theta)^{-t} \end{bmatrix} \begin{bmatrix} \rho r(\theta) & I & & \\ & \rho r(\theta) & I & \\ & & \dots & \\ & & & \rho r(\theta) \end{bmatrix} \begin{bmatrix} r(\theta)^t & & & \\ & r(\theta)^t & & \\ & & \dots & \\ & & & r(\theta)^t \end{bmatrix} \\
 = & \begin{bmatrix} \rho r(\theta)^{-t+1} & r(\theta)^{-t} & & \\ & \rho r(\theta)^{-t+1} & r(\theta)^{-t} & \\ & & \dots & \\ & & & \rho r(\theta)^{-t+1} \end{bmatrix} \begin{bmatrix} r(\theta)^t & & & \\ & r(\theta)^t & & \\ & & \dots & \\ & & & r(\theta)^t \end{bmatrix} \\
 = & \begin{bmatrix} \rho r(\theta) & I & & \\ & \rho r(\theta) & I & \\ & & \dots & \\ & & & \rho r(\theta) \end{bmatrix} \\
 = & J_j.
 \end{aligned}$$

□

Lemma 3.4.2 *If for all i we have $R_i > \max\{0, \log |\lambda_i|\}$ then $\overline{\Upsilon} F_{\underline{R}}$ is stable. If there exists at least one i such that*

$$R_i < \max\{0, \log |\lambda_i|\}$$

then $\overline{\Upsilon} F_{\underline{R}}$ is unstable.

Proof: $\overline{\Upsilon} F_{\underline{R}}$ is block diagonal. Each block is upper triangular. If $R_i > \max\{0, \log |\lambda_i|\}$ then the diagonal of each block will consist of numbers in $[0, 1)$. These are the eigenvalues of that block. Therefore each block is stable. Thus $\overline{\Upsilon} F_{\underline{R}}$ is stable. (See page 39 of [HJo].)

If $R_i < \max\{0, \log |\lambda_i|\}$ for at least one i then there will be a block with a particular element on the diagonal with value greater than one. This is an eigenvalue of value greater than one. Therefore that block is unstable. Thus $\overline{\Upsilon} F_{\underline{R}}$ is unstable. □

Stability of certain matrix products

Lemma 3.5.1 *Let A be a stable matrix. Let B_t be a set of matrices such that $\|B_t\| \leq L$ and the limit $\lim_{t \rightarrow \infty} B_t \rightarrow 0$. Let $S_t = \sum_{i=0}^{t-1} A^{t-1-i} B_i$ then $\lim_{t \rightarrow \infty} S_t = 0$.*

Proof: Since A is stable there exists $c \geq 0$ and $0 \leq \lambda < 1$ such that $\|A^t\| \leq c\lambda^t$. For all $\epsilon > 0$ there exists a $T(\epsilon)$ such that $\|B_t\| \leq \epsilon \forall t \geq T(\epsilon)$. Let $t > T(\epsilon)$. Then

$$\begin{aligned}
 \left\| \sum_{j=0}^{t-1} A^{t-j-1} B_j \right\| &\leq \sum_{j=0}^{t-1} \|A^{t-j-1}\| \|B_j\| \\
 &\leq c \sum_{j=0}^{t-1} \lambda^{t-j-1} \|B_j\| \\
 &\leq c \left\{ \lambda^{t-T(\epsilon)-1} \sum_{j=0}^{T(\epsilon)} \lambda^{T(\epsilon)-j} L + \sum_{j=T(\epsilon)+1}^{t-1} \lambda^{t-j-1} \epsilon \right\} \\
 &\leq \frac{c}{1-\lambda} \left\{ \lambda^{t-T(\epsilon)-1} L + \epsilon \right\}
 \end{aligned}$$

Now we can choose ϵ small enough and a t large enough so that the sum is arbitrarily small. \square

A.3 Chapter 4

Divergence and Mutual Information

We use “log” to represent logarithm base 2 and “ln” to represent logarithm base e . Given two measures P and Q defined on the same space we say $P \ll Q$ if P is absolutely continuous with respect to Q . Denote the Radon-Nikodym derivative as $\frac{dP}{dQ}$. If P is not absolutely continuous with respect to Q then define $\frac{dP}{dQ} \triangleq \infty$.

Definition A.3.1 Let Π be a measurable partition of (Ω, \mathcal{F}) . Denote the elements of the partition by $\pi \in \Pi$. Define the divergence between two measures P, Q on (Ω, \mathcal{F}) to be

$$D(P \mid Q) = \sup_{\Pi} \sum_{\pi \in \Pi} P(\pi) \log \frac{P(\pi)}{Q(\pi)}.$$

where the supremum is over all measurable partitions.

Theorem A.3.1 The divergence can be characterized as:

$$D(P \mid Q) = \begin{cases} \int \log \frac{dP}{dQ} dP & \text{if } P \ll Q \\ +\infty & \text{else} \end{cases}$$

Proof: See theorem 2.4.2 of [Pin].

Given a joint measure $P_{X,Y}$ on $\mathcal{X} \times \mathcal{Y}$ we denote the marginal of $P_{X,Y}$ on \mathcal{X} by P_X . Where $P_X(dX) = \int_{\mathcal{Y}} P_{X,Y}(dX, dy)$.

Definition A.3.2 The mutual information between two random variables X and Y is defined as $I_{P_{X,Y}}(X; Y) \triangleq D(P_{X,Y} \mid P_X P_Y)$ where $P_X P_Y$ denotes the product of measures on $\mathcal{X} \times \mathcal{Y}$. (We will use $I(X; Y)$ when the underlying joint measure is obvious.)

By definition A.3.1 we have

$$I(X; Y) = \sup_{\Pi} \sum_{\pi \in \Pi} P_{X,Y}(\pi) \log \frac{P_{X,Y}(\pi)}{P_X P_Y(\pi)} \quad (\text{A.1})$$

where the supremization is over all partitions Π measurable with respect to $(\mathcal{X} \times \mathcal{Y}, \mathcal{B}(\mathcal{X} \times \mathcal{Y}))$. Furthermore, as the following theorem states, we can restrict the partitions to be “product” partitions.

Theorem A.3.2 In equation (A.1) there is no loss in generality by restricting to partitions Π whose elements π are of the form $\pi = E \times F$ where $E \in \mathcal{B}(\mathcal{X})$ and $F \in \mathcal{B}(\mathcal{Y})$.

Proof: See theorem 2.1.1 of [Pin]. \square

Recall the notation $P_{X|Y} P_{Z|Y} P_Y$ in definition A.1.2.

Definition A.3.3 The conditional mutual information between X and Z given Y is defined as

$$I_{P_{X,Y,Z}}(X; Z|Y) \triangleq D(P_{X,Y,Z} \mid P_{X|Y} P_{Z|Y} P_Y).$$

This is a measure of the distance between $P_{X,Y,Z}$ and its “Markov” form $X - Y - Z$. By definition A.3.1 we have

$$I(X; Z|Y) = \sup_{\Pi} \sum_{\pi \in \Pi} P_{X,Y,Z}(\pi) \log \frac{P_{X,Y,Z}(\pi)}{P_{X|Y}P_{Z|Y}P_Y(\pi)} \quad (\text{A.2})$$

where the supremization is over all partitions Π measurable with respect to $(\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}, \mathcal{B}(\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}))$. Furthermore, as in theorem A.3.2, we can restrict the partitions to be “product” partitions.

Theorem A.3.3 *In equation (A.2) there is no loss in generality by restricting to partitions Π whose elements π are of the form $\pi = E \times F \times G$ where $E \in \mathcal{B}(\mathcal{X})$, $F \in \mathcal{B}(\mathcal{Y})$, and $G \in \mathcal{B}(\mathcal{Z})$.*

Proof: See theorem 3.5.1 of [Pin]. \square

Note that in general $I(X; Z|Y) \neq \sup_{\Pi} \sum_{\pi \in \Pi} P_{X,Y,Z}(\pi) \log \frac{P_{X,Z|Y}(\pi)}{P_{X|Y}P_{Z|Y}(\pi)}$. See exercise 2.27 in Gallager’s text for a counterexample. [Gal]

Definition A.3.4 *If X is a random variable taking on a countable number of values, x_i , then its entropy is defined as $H(X) \triangleq -\sum_i P(X = x_i) \log P(X = x_i)$.*

One can easily show for this case that $H(X) = I(X; X)$.

Definition A.3.5 *If X is a random variable admitting a density, p_X , then its differential entropy is defined as $h(X) \triangleq -\int p_X(x) \log p_X(x) dx$.*

Theorem A.3.4 *Given three random variable X, Y, Z .*

- (a) $P_{X,Z} \ll P_X P_Z$ and $P_{X,Y,Z} \ll P_{Y|X} P_{Z|X} P_X$ if and only if $P_{X,Y,Z} \ll P_{X,Y} P_Z$.
- (b) If $P_{X,Y,Z} \ll P_{X,Y} P_Z$ then

$$\frac{dP_{X,Y,Z}}{dP_{X,Y} P_Z} = \frac{dP_{X,Y,Z}}{dP_{Y|X} P_{Z|X} P_X} \times \frac{dP_{X,Z}}{dP_X P_Z}$$

almost surely $P_{X,Y,Z}$.

- (c) $I(X, Y; Z) = I(Y; Z | X) + I(X; Z)$

Proof: See theorem 3.6.1 of [Pin]. \square

The following proposition lists some further useful properties we will need.

Proposition A.3.1 *The information measures defined above are well defined and*

- (a) $D(P | Q) \geq 0$ with equality if and only if $P = Q$.
- (b) $I(X_1; X_2) = I(X_2; X_1)$.

- (c) *Mutual information is invariant under injective transformations. Let f be an injective measurable function from \mathcal{X} to \mathcal{X} . Then $I(X_1; X_2) = I(f(X_1); f(X_2))$.*
- (d) *$P_{X,Y,Z} = P_{X|Y}P_{Z|Y}P_Y$, i.e. $X - Y - Z$ form a Markov chain, if and only if $I(X; Z|Y) = 0$.*

Proof: Proof of these results can be found in Pinsker's text [Pin]. \square

The following proposition is called the data-processing inequality.

Proposition A.3.2 *Let $X - A - B - Y$ be Markov chain. Then $I(X; Y) \leq I(A; B)$.*

Proof: We know both $I(X, B | A) = 0$ and $I(A; Y | B) = 0$. Thus $I(X; B) = I(A; B) + I(X; B | A) - I(A; B | X) \leq I(A; B)$. A similar calculation shows $I(X; Y) \leq I(X; B)$. \square

Information Stability

In order to prove lemma 4.3.2 we need the following three lemmas. Combined they state that the mass of $\vec{i}(A^T; B^T)$ at the tails is small. Recall that \mathcal{A} and \mathcal{B} are finite spaces.

Lemma A.3.1 *Let $G > \log |\mathcal{A}|$. For any sequence of measures $\{P_{A^T}\}_{t=1}^T$ we have*

$$\lim_{T \rightarrow \infty} E \left[\frac{1}{T} \log \frac{1}{P(A^T)} 1_{\left\{ \frac{1}{T} \log \frac{1}{P(A^T)} \geq G \right\}} \right] = 0.$$

Proof: This proof is adapted from lemma A1 of [HV]. Let

$$\Omega = \{a^T : P(a^T) \leq 2^{-TG}\}.$$

Now

$$\begin{aligned} E \left[\frac{1}{T} \log \frac{1}{P(A^T)} 1_{\left\{ \frac{1}{T} \log \frac{1}{P(A^T)} \geq G \right\}} \right] &= \frac{1}{T} \sum_{a^T \in \Omega} P(a^T) \log \frac{1}{P(a^T)} \\ &= \frac{1}{T} P(\Omega) \sum_{a^T \in \Omega} \frac{P(a^T)}{P(\Omega)} \log \frac{1}{\frac{P(a^T)}{P(\Omega)}} - \frac{1}{T} P(\Omega) \log P(\Omega) \\ &\leq \frac{1}{T} P(\Omega) \log |\mathcal{A}^T| - \frac{1}{T} P(\Omega) \log P(\Omega) \\ &\leq \frac{1}{T} P(\Omega) \log |\mathcal{A}^T| + \frac{1}{T} \end{aligned}$$

where the first inequality follows because entropy is maximized by the uniform distribution and the second inequality follows because $-x \log x \leq 1$, $0 \leq x \leq 1$.

Now $P(\Omega) \leq |\Omega|2^{-TG} \leq |\mathcal{A}^T|2^{-TG}$. Thus

$$E \left[\frac{1}{T} \log \frac{1}{P(A^T)} \mathbf{1}_{\left\{ \frac{1}{T} \log \frac{1}{P(A^T)} \geq G \right\}} \right] \leq \log |\mathcal{A}| 2^{-T(G - \log |\mathcal{A}|)} + \frac{1}{T}.$$

This upper bound goes to zero as T goes to infinity. \square

Lemma A.3.2 *Let $G > \log |\mathcal{A}|$. For any sequence of joint measures $\{P_{A^T, B^T}\}_{T=1}^\infty$ we have*

$$\lim_{T \rightarrow \infty} E \left[\frac{1}{T} \vec{i}(A^T; B^T) \mathbf{1}_{\left\{ \frac{1}{T} \vec{i}(A^T; B^T) \leq -G \right\}} \right] = 0$$

Proof: Let

$$\Omega_{b^T} = \{a^T : P(a^T | b^T) \leq 2^{-TG}\}.$$

Note that for $P(b^T) > 0$ we have

$$P(a^T | b^T) = \frac{\vec{P}(a^T | b^T) \vec{P}(b^T | a^T)}{P(b^T)} \leq \frac{\vec{P}(b^T | a^T)}{P(b^T)}.$$

Now

$$\begin{aligned} & E \left[\frac{1}{T} \vec{i}(A^T; B^T) \mathbf{1}_{\left\{ \frac{1}{T} \vec{i}(A^T; B^T) \leq -G \right\}} \right] \\ &= E \left[\frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \mathbf{1}_{\left\{ \frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \leq -G \right\}} \right] \\ &\geq E \left[\frac{1}{T} \log P(A^T | B^T) \mathbf{1}_{\left\{ \frac{1}{T} \log P(A^T | B^T) \leq -G \right\}} \right] \\ &= - \sum_{b^T} P(b^T) E \left[\frac{1}{T} \log \frac{1}{P(A^T | b^T)} \mathbf{1}_{\{A^T \in \Omega_{b^T}\}} \mid b^T \right] \\ &\geq - \sum_{b^T} P(b^T) \left(\log |\mathcal{A}| 2^{-T(G - \log |\mathcal{A}|)} + \frac{1}{T} \right) \\ &= \left(\log |\mathcal{A}| 2^{-T(G - \log |\mathcal{A}|)} + \frac{1}{T} \right) \end{aligned}$$

Where the last inequality follows from lemma A.3.1. This lower bound goes to zero as T goes to infinity. \square

Lemma A.3.3 *Let $G > \log |\mathcal{A}|$. For any sequence of joint measures $\{P_{A^T, B^T}\}_{T=1}^\infty$ we have*

$$\lim_{T \rightarrow \infty} E \left[\frac{1}{T} \vec{i}(A^T; B^T) \mathbf{1}_{\left\{ \frac{1}{T} \vec{i}(A^T; B^T) \geq G \right\}} \right] = 0.$$

Proof: Let

$$\vec{\Omega}_{b^T} = \{a^T : \vec{P}(a^T | b^T) \leq 2^{-TG}\}.$$

Note that for $P(b^T) > 0$ we have

$$\frac{1}{\vec{P}(a^T | b^T)} = \frac{\vec{P}(b^T | a^T)}{P(a^T | b^T)P(b^T)} \geq \frac{\vec{P}(b^T | a^T)}{P(b^T)}.$$

Now

$$\begin{aligned} & E \left[\frac{1}{T} \vec{i}(A^T; B^T) \mathbf{1}_{\{\frac{1}{T} \vec{i}(A^T; B^T) \geq G\}} \right] \\ &= E \left[\frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \mathbf{1}_{\{\frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \geq G\}} \right] \\ &\leq E \left[\frac{1}{T} \log \frac{1}{\vec{P}(A^T | B^T)} \mathbf{1}_{\{\frac{1}{T} \log \frac{1}{\vec{P}(A^T | B^T)} \geq G\}} \right] \\ &= \sum_{b^T} P(b^T) E \left[\frac{1}{T} \log \frac{1}{\vec{P}(A^T | b^T)} \mathbf{1}_{\{A^T \in \vec{\Omega}_{b^T}\}} \middle| b^T \right] \\ &\leq \sum_{b^T} P(b^T) \left(\log |\mathcal{A}| 2^{-T(G - \log |\mathcal{A}|)} + \frac{1}{T} \right) \\ &= \left(\log |\mathcal{A}| 2^{-T(G - \log |\mathcal{A}|)} + \frac{1}{T} \right) \end{aligned}$$

Where the last inequality follows from lemma A.3.1. This upper bound goes to zero as T goes to infinity. \square

Now we can prove lemma 4.3.2.

Lemma 4.3.2 *For any sequence of joint measures $\{P_{A^T, B^T}\}_{T=1}^{\infty}$ we have*

$$\underline{I}(A \rightarrow B) \leq \liminf_{T \rightarrow \infty} \frac{1}{T} I(A^T \rightarrow B^T) \leq \limsup_{T \rightarrow \infty} \frac{1}{T} I(A^T \rightarrow B^T) \leq \bar{I}(A \rightarrow B)$$

Proof: We first prove the leftmost inequality. For any $\epsilon > 0$ we have

$$\begin{aligned} \frac{1}{T} I(A^T \rightarrow B^T) &\geq E \left[\frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \mathbf{1}_{\{\frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \leq -G\}} \right] \\ &\quad - GP \left[-G \leq \frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \leq \underline{I}(A \rightarrow B) - \epsilon \right] \\ &\quad + (\underline{I}(A \rightarrow B) - \epsilon) P \left[\frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \geq \underline{I}(A \rightarrow B) - \epsilon \right] \end{aligned}$$

The first addend goes to zero by lemma A.3.2, the second addend goes to zero by definition of \underline{I} , and the probability in the last addend goes to 1. Thus for T large enough we have $\frac{1}{T}I(A^T \rightarrow B^T) \geq \underline{I} - 2\epsilon$. Since ϵ is arbitrary we see that $\underline{I}(A \rightarrow B) \leq \liminf_{T \rightarrow \infty} \frac{1}{T}I(A^T \rightarrow B^T)$.

The second inequality, $\liminf_{T \rightarrow \infty} \frac{1}{T}I(A^T \rightarrow B^T) \leq \limsup_{T \rightarrow \infty} \frac{1}{T}I(A^T \rightarrow B^T)$, is obvious.

Now we treat the rightmost inequality. For any $\epsilon > 0$ we have

$$\frac{1}{T}I(A^T \rightarrow B^T) \leq E \left[\frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} 1_{\left\{ \frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \geq G \right\}} \right] \quad (\text{A.1})$$

$$+ GP \left[G \geq \frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{P(B^T)} \geq \bar{I}(A \rightarrow B) + \epsilon \right] \quad (\text{A.2})$$

$$+ (\bar{I}(A \rightarrow B) + \epsilon) P \left[\frac{1}{T} \log \frac{\vec{P}(B^T | A^T)}{q(B^T)} \leq \bar{I}(A \rightarrow B) + \epsilon \right] \quad (\text{A.3})$$

The first addend goes to zero by lemma A.3.3, the second addend goes to zero by definition of \bar{I} , and the probability in the last addend goes to 1. Thus for T large enough we have $\frac{1}{T}I(A^T \rightarrow B^T) \leq \bar{I} + 2\epsilon$. Since ϵ is arbitrary we see that $\limsup_{T \rightarrow \infty} \frac{1}{T}I(A^T \rightarrow B^T) \leq \bar{I}(A \rightarrow B)$. \square

A.4 Chapter 5

Lemma A.4.1 *If $X \sim \mathcal{N}(0, \lambda)$ then for $L > 0$*

$$\int_L^\infty x^2 p(x) dx = \sqrt{\frac{\lambda}{2\pi}} L e^{-\frac{L^2}{2\lambda}} + \lambda \Pr(X \geq L).$$

Proof: We prove this by integration by parts.

$$\begin{aligned} \int_L^\infty \frac{1}{\sqrt{2\pi\lambda}} x^2 e^{-\frac{x^2}{2\lambda}} dx &= \frac{-\lambda x}{\sqrt{2\pi\lambda}} e^{-\frac{x^2}{2\lambda}} \Big|_L^\infty + \int_L^\infty \frac{\lambda}{\sqrt{2\pi\lambda}} e^{-\frac{x^2}{2\lambda}} dx \\ &= \sqrt{\frac{\lambda}{2\pi}} L e^{-\frac{L^2}{2\lambda}} + \lambda \Pr(X \geq L). \end{aligned}$$

□

The proof of the following lemma follows from [OR].

Lemma A.4.2 (a) $P(a^N \in \Omega_A^{\delta, N}) \geq 1 - \epsilon_A^{\delta, N}$.

(b) For all $(a^N, b^N) \in \Omega_{A,B}^{\delta, N}$ we have $2^{-N(1+\delta)H(B|A)} \leq P(b^N | a^N) \leq 2^{-N(1-\delta)H(B|A)}$.

(c) Let $a^N \in \Omega_A^{\delta_1, N}$. Let b^N be drawn from $P(B^N | a^N)$. Then

$$P(b^N \in \Omega_{A,B}^{\delta_2, N}(a^N) | A^N = a^N) \geq 1 - \epsilon_{A,B}^{\delta_1, \delta_2, N}.$$

(d) For every $a^N \in \Omega_A^{\delta_1, N}$ we have $|\Omega_{A,B}^{\delta_2, N}(a^N)| \geq (1 - \epsilon_{A,B}^{\delta_1, \delta_2, N}) 2^{N(1-\delta_2)H(B|A)}$.

Proof:

(a) By Chernoff's bound we have $P(|\nu_{A^N}(a) - P(a)| \geq \delta P(a)) \leq e^{-N \frac{\delta^2 P(a)}{3}}$. Use this and the union bound to get result (a).

(b) Since $(a^N, b^N) \in \Omega_{A,B}^{\delta, N}$ we know

$$\begin{aligned} P(b^N | a^N) &= 2^{-N \left(\sum_{a,b} \nu_{(a^N, b^N)}(a,b) \log \frac{1}{P(b|a)} \right)} \\ &= 2^{-N \left(H(B|A) + \sum_{a,b} (\nu_{(a^N, b^N)}(a,b) - P(a,b)) \log \frac{1}{P(b|a)} \right)} \end{aligned}$$

Now $\left| \sum_{a,b} (\nu_{(a^N, b^N)}(a,b) - P(a,b)) \log \frac{1}{P(b|a)} \right| \leq \delta \sum_{a,b} P(a,b) \log \frac{1}{P(b|a)} = \delta H(B|A)$.

(c) Note that $E\left(\nu_{(a^N, B^N)}(a, b) | a^N\right) = \nu_{a^N}(a)P(b|a)$

$$\begin{aligned} & P\left(|\nu_{(a^N, B^N)}(a, b) - P(a, b)| \geq \delta_2 P(a, b) \mid a^N\right) \\ = & P\left(|\nu_{(a^N, B^N)}(a, b) - P(a, b)| \geq \delta_2 \frac{P(a)}{\nu_{a^N}(a)} \nu_{a^N}(a) P(b|a) \mid a^N\right) \\ \leq & e^{-N \frac{(\delta_2 - \delta_1)^2 p(a, b)}{3(1 + \delta_1)}}. \end{aligned}$$

where the third line follows from the Chernoff bound and a rearranging of terms. Now use this and the union bound to get (c).

(d)

$$\begin{aligned} 1 - \epsilon_{A, B}^{\delta_1, \delta_2, N} & \leq P\left(b^N \in \Omega_{A, B}^{\delta_2, N}(a^N) \mid A^N = a^N\right) \\ & = \sum_{b^N \in \Omega_{A, B}^{\delta_2, N}(a^N)} P(b^N | a^N) \\ & \leq |\Omega_{A, B}^{\delta_2, N}(a^N)| 2^{-N(1 - \delta_2)H(B|A)}. \end{aligned}$$

where the first line follows from part (c) and the last line follows from part (b).

□

The following is the key technical lemma in the proof of the coding theorem.

Lemma A.4.3 *Let $(x^{t, N}, y^{t-1, N}) \in \Omega_{X^t, Y^{t-1}}^{\delta_1, N}$. Let y_t^N be drawn from $P(Y_t^N | y^{t-1, N})$. Then*

$$\begin{aligned} & P\left(y_t^N \in \Omega_{X^t, Y^{t-1}}^{\delta_2, N}(x^{t, N}, y^{t-1, N}) \mid x^{t, N}, y^{t-1, N}\right) \\ \geq & (1 - \epsilon_{X^t, Y^t}^{\delta_1, \delta_2, N}) 2^{-N(I(X^t; Y_t | Y^{t-1}) + 2\delta_2 H(Y_t | Y^{t-1}))}. \end{aligned}$$

Proof:

$$\begin{aligned} & P\left(y_t^N \in \Omega_{X^t, Y^{t-1}}^{\delta_2, N}(x^{t, N}, y^{t-1, N}) \mid x^{t, N}, y^{t-1, N}\right) \\ = & \sum_{y_t^N \in \Omega_{X^t, Y^{t-1}}^{\delta_2, N}(x^{t, N}, y^{t-1, N})} P(y_t^N | y^{t-1, N}) \\ \geq & \sum_{y_t^N \in \Omega_{X^t, Y^{t-1}}^{\delta_2, N}(x^{t, N}, y^{t-1, N})} 2^{-N(1 + \delta_2)H(Y_t | Y^{t-1})} \\ = & |\Omega_{X^t, Y^{t-1}}^{\delta_2, N}(x^{t, N}, y^{t-1, N})| 2^{-N(1 + \delta_2)H(Y_t | Y^{t-1})} \\ \geq & (1 - \epsilon_{X^t, Y^t}^{\delta_1, \delta_2, N}) 2^{N(1 - \delta_2)H(Y_t | X^t, Y^{t-1})} 2^{-N(1 + \delta_2)H(Y_t | Y^{t-1})} \\ \geq & (1 - \epsilon_{X^t, Y^t}^{\delta_1, \delta_2, N}) 2^{-N(I(X^t; Y_t | Y^{t-1}) + 2\delta_2 H(Y_t | Y^{t-1}))}. \end{aligned}$$

Where the third line follows from lemma A4.2(b) and the fifth line from lemma A4.2(d).
The last line comes from noting

$$\begin{aligned}
& (1 + \delta_2)H(Y_t|Y^{t-1}) - (1 - \delta_2)H(Y_t|X^t, Y^{t-1}) \\
= & H(Y_t|Y^{t-1}) - H(Y_t|X^t, Y^{t-1}) + \delta_2 \left(H(Y_t|Y^{t-1}) + H(Y_t|X^t, Y^{t-1}) \right) \\
= & I(X^t; Y_t|Y^{t-1}) + \delta_2 \left(H(Y_t|Y^{t-1}) + H(Y_t|X^t, Y^{t-1}) \right) \\
\leq & I(X^t; Y_t|Y^{t-1}) + 2\delta_2 H(Y_t|Y^{t-1})
\end{aligned}$$

□

Lemma A.4.4 *If $(x_t^N, y_t^N) \in \Omega_{X_t, Y_t}^{\delta, N}$ then $d_N(x_t^N, y_t^N) \leq (1 + \delta)E(d(X_t, Y_t))$.*

Proof:

$$\begin{aligned}
d_N(x_t^N, y_t^N) &= \frac{1}{N} \sum_{n=1}^N d(x_{t,n}, y_{t,n}) \\
&= \sum_{x,y} \nu_{(x_{t,n}, y_{t,n})}(x, y) d(x, y) \\
&\leq \sum_{x,y} (1 + \delta) p(x, y) d(x, y) \\
&= (1 + \delta) E(d(X_t, Y_t)).
\end{aligned}$$

□

References

- [Ast] Karl Astrom, “Riemann and Lebesgue Sampling.” Advances in System Theory: A Symposium in Honor of Sanjoy K. Mitter, October 15-16, 1999, Cambridge, MA.
- [BB] R. Bansal and T. Basar, “Simultaneous Design of Measurement and Control Strategies for Stochastic Systems with Feedback.” *Automatica*, Volume 25, No. 5, pp 679-694, 1989.
- [BT] Y. Bar-Shalom and E. Tse, “Dual effect, Certainty Equivalence, and Separation in Stochastic Control.” *IEEE Trans. on Automatic Control*, Vol. AC-19, No. 5, October 1974, pp. 494-500.
- [Ben] W. Bennett, “Spectra of Quantized Signals.” *Bell Systems Technical Journal*, Vol. 27, July 1948, pp. 446-472.
- [Berg] T. Berger, *Rate Distortion Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1971.
- [Bert] D. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific, 1995.
- [BS] D. Bertsekas and S. Shreve, *Stochastic Optimal Control: the Discrete Time Case*. New York, Academic Press: 1978.
- [Bor] V. Borkar, “Average Cost Dynamic Programming Equations for Controller Markov Chains with Partial Observations.” To appear in the *SIAM Journal on Control and Optimization*.
- [BM] V. Borkar and S. Mitter, “LQG Control with Communication Constraints.” In *Communications, Computation, Control, and Signal Processing: a Tribute to Thomas Kailath*. Norwell, MA, Kluwer Academic Publishers: 1997.
- [CS] G. Caire and S. Shamai, “On the Capacity of Some Channels with Channel State Information.” *IEEE Trans. on Information Theory*, Vol. IT-45, No. 6, September 1999, pp. 2007-2019.
- [CP] T. Cover and S. Pombra, “Gaussian Feedback Capacity.” *IEEE Trans. on Information Theory*, Vol. IT-35, No. 1, January 1989, pp. 37-43.
- [CT] T. Cover and J. Thomas, *Elements of Information Theory*, John Wiley, New York, 1991.

- [Dob1] R. Dobrushin, "Transmission of Information in Channels with Feedback." *Theor. Prob. Appl.*, 3, N4, pp. 395-412, 1958.
- [Dob2] R. Dobrushin, "A General Formulation of the Basic Shannon Theorem in Information Theory," *Uspekhi Math. Nauk*, 14, N6, pp. 3-103, 1959.
- [Dud] R. Dudley, *Uniform Central Limit Theorems*, Cambridge University Press, New York, 1999.
- [Eli] N. Elia, personal communication, October 1999.
- [EM] N. Elia and S. Mitter, *Quantized Linear Systems*, LIDS report P 2442, September 1999.
- [EC] W. Equitz and T. Cover, "Successive Refinement of Information." *IEEE Trans. on Information Theory*, Vol. 37, March 1991, pp. 269-274.
- [Fei] A. Feinstein, "A New Basic Theorem of Information Theory." *IRE Trans. on Information Theory*, Vol 4, pp. 2-22, 1954
- [Gal] R. Gallager, *Information Theory and Reliable Communication*. New York: Wiley, 1968.
- [GV1] A. Goldsmith and P. Varaiya, "Capacity, mutual information, and coding for finite-state Markov channels," *IEEE Trans. Info. Theory*, vol. IT-42, pp. 868 - 886, May 1996.
- [GV2] A. Goldsmith and P. Varaiya, "Capacity of fading channels with channel side information." *IEEE Trans. Info. Theory*, vol. IT-43, pp. 1986 - 1992, November 1997.
- [GP] A. Gorbunov and M. Pinsker, "Prognostic Epsilon Entropy of a Gaussian Message and a Gaussian Source." Translation from *Problemy Peredachi Informatsii*, Vol. 10, No. 2, pp. 5-25, April-June 1973.
- [Han] T. Han, "An Information-Spectrum Approach to Source Coding Theorems with a Fidelity Criterion." *IEEE Trans. on Information Theory*, Vol. 43, July 1997, pp. 1145-1164.
- [HV] T. Han and S. Verdu, "Approximation Theory of Output Statistics." *IEEE Trans. on Information Theory*, Vol. 39, May 1993, pp. 752-772.
- [HLL] O. Hernandez-Lerma and J. Lasserre, *Discrete-Time Markov Control Processes*. New York, NY: Springer-Verlag, 1996.
- [HS] M. Hirsch and S. Smale, *Differential Equations, Dynamical Systems, and Linear Algebra*. San Diego, CA: Academic Press, Inc., 1974.
- [HJ] B. Hochwald and P. Jelenkovic, "State Learning and Mixing in Entropy of Hidden Markov Processes and the Gilbert-Elliott Channel." *IEEE Trans. Info. Theory*, vol. IT-45, pp. 128 - 138, January 1999.

- [HJo] R. Horn and C. Johnson, *Matrix Analysis*. New York, NY: Cambridge University Press, 1985.
- [Kle] T. Klein, "Lower Bound to the Feedback Capacity of the Colored Gaussian Noise Channel." Proceedings of the ISIT 2000 Conference.
- [Kra] G. Kramer, "Directed Information for Channels with Feedback." Ph.D. thesis, Swiss Federal Institute of Technology, Zurich, 1998.
- [Lap] A. Lapidoth, "On the role of mismatch in rate distortion theory." IEEE Trans. Info. Theory, vol. IT-43, pp. 38 - 47, January 1997.
- [LP] K.H. Lee and D.P. Petersen, "Optimal Linear Coding for Vector Channels." IEEE Transactions on Communications, Volume 24, No. 12, pp. 1283-1290, December 1976
- [LB] D. Liberzon and R. Brockett, "Quantized Feedback Stabilization of Linear Systems." IEEE Transactions on Automatic Control, Volume 45, no. 7, pp. 1279-1289, July 2000.
- [Mar] H. Marko, "The bidirectional Communication Theory - A Generalization of Information Theory." IEEE Trans. Comm., vol. COM-21, pp. 1345-1351, December 1973.
- [Mas] J. Massey, "Causality, Feedback, and Directed Information." Proceedings of the 1990 International Symposium on Information Theory and its Applications (ISITA-90), pp. 303-305.
- [Mit] S. Mitter, "Real-Time Information Theory." Allerton Workshop on Future Directions in System and Control. September 1997.
- [NE1] G. Nair and R. Evans, "State estimation under bit-rate constraints." Proceedings of the 37th IEEE Conference on Decision and Control, 1998, pp. 251-256 vol. 1.
- [NE2] G. Nair and R. Evans, "State estimation via a capacity-limited communication channel." Proceedings of the 36th IEEE Conference on Decision and Control, 1997, pp. 866-871 vol. 1.
- [OR] A. Orlitsky and J. Roche, "Coding for Computing." Proceedings of the 36th Annual Symposium on the Foundations of Computer Science, 1995, pp. 502 -511.
- [Pil] R. Pilc, "The Transmission Distortion of a Source as a Function of the Encoding Block Length." Bell Systems Technical Journal, Vol. 47, pp. 885, 1968.
- [Pea] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. Morgan-Kaufmann, 1988
- [Pin] M. S. Pinsker, *Information and Information Stability of Random Variables and Processes*. Translated by Amiel Feinstein. San Francisco, CA: Holden-Day, 1964.
- [Rad] R. Radner, "Team decision problems," Annals of Mathematical Statistics," Vol. 33, No. 3, pp. 857-881, 1962.

- [Rim] B. Rimoldi, Successive Refinement of Information: Characterization of the Achievable Rates." IEEE Trans. on Information Theory, Vol. 40, No. 1, pp. 253-259, January 1994.
- [Sah] A. Sahai, MIT Ph.D. thesis. In preparation, 2000.
- [Sau] F. Saussure, *Course in General Linguistics*. Translation of Cours de Linguistique Generale. London. Duckworth, 1983.
- [SK] Schalkwijk and Kailath, "A coding scheme for additive noise channels with feedback—I: No bandwidth constraint." IEEE Trans. Info. Theory, vol. IT-12, pp. 172 - 182, April 1966.
- [Sha1] C. Shannon, "A Mathematical Theory of Communication," Bell System Technical Journal, Vol. 27 (July and October 1948), pp. 379-423 and 623-656.
- [Sha2] C. Shannon, "The Zero Error Capacity of a Noisy Channel," Institute of Radio Engineers, Transactions on Information Theory, Vol. IT-2 (September, 1956), pp. S8-S19.
- [Sha3] C. Shannon, "Channels With Side Information at the Transmitter," IBM Journal Research and Development, Vol. 2 (1958), pp. 289-293. (Received Sept. 15, 1958.)
- [Son] E. Sontag, *Mathematical Control Theory, 2nd Edition*. New York. Springer-Verlag, 1998.
- [Str] C. Striebel, *Optimal Control of Discrete Time Stochastic Systems*. Berlin, Springer-Verlag, 1975.
- [SV] Y. Steinberg and S. Verdu, "Simulation of Random Processes and Rate-Distortion Theory." IEEE Trans. on Information Theory, Vol. 42, Jan 1996, pp. 63-86.
- [VVS] S. Vembu, S. Verdu, and Y. Steinberg, "The source-channel separation theorem revisited," IEEE Trans. Info. Theory, vol. IT-41, pp. 44 - 54, January 1995.
- [VH] S. Verdu, T. Han, "A general formula for channel capacity." IEEE Trans. Info. Theory, vol. IT-40, pp. 1147 - 1157, July 1994.
- [Vid] M. Vidyasagar, "On the Stabilization of Nonlinear Systems Using State Detection." IEEE Trans. on Automatic Control, Vol. 25, No. 3, pp. 504-509, June 1980.
- [Vis] H. Viswanathan, "Capacity of Markov Channels with Receiver CSI and Delayed Feedback." IEEE Trans. Info. Theory, vol. IT-45, No. 2, pp. 761-771, March 1999..
- [Wit] H. Witsenhausen, "Separation of Estimation and Control for Discrete Time Systems." Proceedings of the IEEE, Volume 59, No. 11, November 1971.
- [Wit2] H. Witsenhausen, "On Policy Independence of Conditional Expectations." Information and Control, Volume 28, pp. 65-75, 1975.

- [Wit3] H. Witsenhausen, "A Counterexample in Stochastic Optimum Control." SIAM J. Control, volume 6, pp. 131-147, 1968.
- [Wol] J. Wolfowitz, *Coding Theorems of Information Theory*. Berlin, Prentice-Hall, 1961.
- [WB1] W. Wong and R. Brockett, "Systems with Finite Communication Bandwidth Constraints I: State Estimation Problems," IEEE Trans. on Automatic Control, Vol. 42(9), pp. 1294-1299, Sept. 1997.
- [WB2] W. Wong and R. Brockett, "Systems with Finite Communication Bandwidth Constraints II: Stabilization with Limited Information Feedback," IEEE Trans. on Automatic Control, Vol. 44 (5), pp. 1049 -1053, May 1999.