TO CODE OR NOT TO CODE

THÈSE N^{O} 2687 (2002)

PRÉSENTÉE À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS

SECTION DES SYSTÈMES DE COMMUNICATION

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Michael GASTPAR

ingénieur électricien diplômé EPF, M.Sc. University of Illinois, Urbana-Champaign, Etats Unis
de nationalité suisse et originaire de Zurich (ZH) et Lucerne (LU)

acceptée sur proposition du jury:

Prof. M. Vetterli, Prof. B. Rimoldi, directeurs de thèse
Prof. R. Gallager, rapporteur
Prof. A. Lapidoth, rapporteur
Prof. J. Massey, rapporteur
Prof. E. Telatar, rapporteur
Prof. S. Verdú, rapporteur

Lausanne, EPFL
2002

Acknowledgments

If a man will begin with certainties, he shall end in doubts;
but if he will be content to begin with doubts, he shall end in certainties.
– FRANCIS BACON, *The Advancement of Learning, Book One (1605)*[1]

Einer hat immer Unrecht: aber mit Zweien beginnt die Wahrheit.
Einer kann sich nicht beweisen: aber Zweie kann man bereits nicht widerlegen.
– FRIEDRICH NIETZSCHE, *Die fröhliche Wissenschaft (1882)*

It was a great pleasure to be the pupil and student of Bixio Rimoldi and Martin Vetterli. If I haven't learned my trade from them, I'm never going to learn it. They offered an invaluable combination of craftsmanship and art, of detail and the big picture, of precision and inspiration, of dogma and *Freidenkertum*. They taught me that research is not only the quest for truth, but just as well an everyday business, an *art pour l'art*, and a way to have fun.

As if this extraordinary constellation was not enough luck by itself, I also had the opportunity to interact with Emre Telatar, who influenced this work in many different ways, and with Rüdiger Urbanke.

An exciting summer internship at Bell Labs, Lucent Technologies, offered the opportunity to work with Piyush Gupta, and with Gerhard Kramer, who first taught me research five years ago at ETH Zürich, and who remained a great source of inspiration, a teacher of many intricate details, and a friend.

Another important acknowledgment goes to Jim Massey for being among the prime suspects for getting me into all of this, and for his detailed and much appreciated comments on the first version of this thesis.

[1] Francis Bacon, A critical edition of the major works (The Oxford Authors). Edited by B. Vickers. Oxford University Press, 1996. p. 147.

It was a great pleasure and honor for me to have, on my *jury de thèse*, Bob Gallager, Amos Lapidoth, Jim Massey, Sergio Verdú, and Emre Telatar, and I would like to thank them for accepting to read and review the first version of this thesis.

A team of outstanding researchers, latex specialists, and other experts contributed in various ways to this thesis, including Nicolae Chiurtu, Minh Do, Pier Luigi Dragotti, Peggy Gerardin, David Hasler, Raymond Knopp, Jérôme Lebrun, Olivier Lévêque, Ninoslav Marina, Pina Marziliano, Pietro Polotti, Andrea Ridolfi, Patrick Thiran, Daniela Tuninetti, Harald Viste, and Claudio Weidmann.

Un grand merci à Françoise Behn et à Jocelyne Plantefol: Your administration of the laboratories was both professional and enjoyable. And thanks to Denice Deatrich and Williams Devadson for a reliable computer network.

Doch ein Sandwich ohne Käse ist nichts als Brot. Dass diese Dissertation in drei glücklichen Jahren entstand, verdanke ich insbesondere Nicole, Fränä, Sabine, Fetsch, meinen Eltern Regine und Huldrych, und meinen Brüdern Andreas und Tobias.

Cette thèse est dédiée à Nicole qui y tient la partie la plus importante, bien au-delà de m'aider à réussir la preuve mathématique la plus difficile.

Abstract

It is well known and surprising that the uncoded transmission of an independent and identically distributed Gaussian source across an additive white Gaussian noise channel is optimal: No amount of sophistication in the coding strategy can ever perform better.

What makes uncoded transmission optimal? In this thesis, it is shown that the optimality of uncoded transmission can be understood as the perfect match of four involved measures: the probability distribution of the source, its distortion measure, the conditional probability distribution of the channel, and its input cost function.

More generally, what makes a source-channel communication system optimal? Inspired by, and in extension of, the results about uncoded transmission, this can again be understood as the perfect match, now of six quantities: the above, plus the encoding and the decoding functions. The matching condition derived in this thesis is explicit and closed-form. This fact is exploited in various ways, for example to analyze the optimality of source-channel coding systems of finite block length, and involving feedback.

In the shape of an intermezzo, the potential impact of our findings on the understanding of biological communication is outlined: owing to its simplicity, uncoded transmission must be an interesting strategy, e.g., for neural communication. The matching condition of this thesis shows that, apart from being simple, uncoded transmission may also be information-theoretically optimal.

Uncoded transmission is also a useful point of view in network information theory. In this thesis, it is used to determine network source-channel communication results, including a single-source broadcast scenario, to establish capacity results for Gaussian relay networks, and to give a new example of the fact that separate source and channel coding does not lead to optimal performance in general networks.

Kurzfassung

Es ist eine wohlbekannte und unverschämte Tatsache, dass die *unkodierte* Übertragung einer Gaussischen Quelle über einen Gaussischen Kanal optimal ist. Dramatischer ausgedrückt: in diesem Beispiel kann auch die durchdachteste und komplizierteste Übertragungstechnik nicht besser sein als die schlichte unkodierte.

Was macht unkodierte Übertragung in diesem Beispiel optimal? Die Nemesis des Glückfalls ist, nur unter äusserst günstigen Umständen einzutreffen. Was das genau heisst, wird in dieser Dissertation gezeigt. Vier Masse definieren ein Kommunikationsproblem, als da wären: die Wahrscheinlichkeitsverteilungen der Quelle und des Kanals, das Verzerrungsmass und die Kostenfunktion des Kanals. Wenn diese vier Masse bereits perfekt aufeinander abgestimmt sind, dann ist jegliche Kodierung überflüssig und nur durch Eitelkeit zu rechtfertigen.

Allgemeiner gefragt, was macht *kodierte* Übertragung optimal? Auch dies kann als optimale Abstimmung der betreffenden Masse verstanden werden: wiederum sind die Quellen- und Kanalparameter im Spiel, daneben aber auch die Kodierungs- und die Dekodierungsfunktion. In dieser Dissertation wird eine explizite Formel für die optimale Abstimmung hergeleitet. Dann wird diese Formel auf verschiedene Probleme angewandt, u.a. auf Kodierungssysteme mit endlicher Blocklänge und mit Feedback.

In Form eines Intermezzo zeigen wir auf, in welcher Form unsere Resultate dem Verständnis der Kommunikation in biologischen Systemen förderlich sein könnten. Am Beispiel der neuronalen Kommunikation illustrieren wir, dass die unkodierte Übertragung nicht nur wegen ihrer unübertrefflichen Einfachheit interessant ist, sondern dass sie auch optimal sein könnte, wenn nur die Masse abgestimmt wären. Und die Evolution hätte die Möglichkeit gehabt, eine solche Abstimmung herbeizuführen.

Schliesslich wird unkodierte Übertragung in Netzwerken diskutiert. Auch hier hält sie einige Überraschungen bereit, so zum Beispiel für ein einfaches Modell eines Broadcast-Netzwerkes, und für ein bestimmtes Netzwerk mit sehr vielen Helferknoten.

Contents

List of Figures

Introduction

Je planmässiger die Menschen vorgehen, desto wirkungsvoller vermag sie der Zufall zu treffen.

– FRIEDRICH DÜRRENMATT, *21 Punkte zu den Physikern*

Source-Channel Communication

In a telephone conversation or a television broadcast, a signal is to be transmitted across a noisy channel. The goal is by no means to transmit this signal *perfectly*, but simply to satisfy the human ear and eye.

The wireless sensors of a sensor network that monitors chemical concentrations can acquire large amounts of data, but their batteries are limited. Typically, it is impossible for the sensor to transmit *all* of the acquired information. Rather, the sensor must attempt to give the best approximation within the potential of its battery.

Unlike, say, the transmission of executable program files over the internet, the basic goal in the two scenarios outlined above is *neither* to find a non-redundant representation of the data *nor* to send bits across a channel at the smallest possible error rate. Rather, in both situations, the key task is to trade off the power of the transmission (or, more generally, its *cost*) against the quality of the data reconstruction.

The determination and characterization of the optimal trade-off between cost and distortion is one of the fundamental problems of information theory, sometimes called the *source-channel communication problem.* It can be addressed from various perspectives; among them, the themes of this thesis.

Themes

- **Separation.** The optimal trade-off between cost and distortion for a simple (ergodic) point-to-point link can be achieved by splitting the coding system into two parts: first, the source is represented in a non-redundant fashion at the required level of fidelity. Thereafter, this representation is transmitted across the channel, using a code with a negligible error probability. This is Shannon's separation theorem.

  The fact that the system can be separated (or modularized) in such a way without sacrificing optimality is conceptually as well as practically pleasing: For the source coding stage, no knowledge is required about the precise channel characteristics, and for the channel code design, the structure of the source is irrelevant, too. In Chapter 1 of this thesis, we review Shannon's separation theorem.

  The theme of modularization is revisited in the context of networks. Here, there is a double temptation of modularization: on the one hand, to modularize source and channel coding, like in the point-to-point case; on the other hand, to modularize the individual channels, i.e., to turn the network into a set of point-to-point channels. In Chapter 5, we illustrate that both kinds of modularization *do incur* a loss of optimality.

- **Optimal Uncoded Transmission.** While the separation theorem provides a very pleasing solution to the source-channel communication problem, it does not claim it to be unique. In fact, it provides a very expensive solution since infinite delay and complexity are needed in general. Less expensive solutions can be found by abandoning the property of modularization, that is, by considering *joint* source-channel codes. The extreme case of a joint source-channel code is *uncoded* (or straight-wire) transmission: the source output is directly fed into the channel input, and the channel output is the estimate of the source. It is well known (yet slightly puzzling) that for certain source/channel pairs, uncoded transmission *does* perform optimally: it achieves the same cost-distortion trade-off as the scheme composed of an optimal source coding stage, followed by an optimal channel code.

  This thesis determines the conditions for the optimality of uncoded transmission for the point-to-point case (Chapter 2) and for certain network communication scenarios (Chapter 5).

- **Measure-matching.** The point of view on source-channel communication that is suggested by the separation theorem could be called *rate-matching*: the source-channel communication problem is solved by matching the *rates* of the (otherwise independent) source and channel code designs: for the system to work, the rate of the source code cannot be larger than the capacity provided by the channel.

  In extension of our study of uncoded transmission, we suggest in Chapter 3 a different criterion for optimal systems; we call this *measure-matching*:

A source-channel communication system is optimal when/because certain measures are matched optimally to each other. These measures are: the source probability distribution, and the corresponding distortion measure; the channel conditional probability distribution, and the corresponding input cost function; and the encoding and the decoding functions.

Contributions

1. **When is uncoded transmission optimal?** Two examples are well known. In this thesis, we establish general conditions for the optimality of uncoded transmission in the point-to-point case (Chapter 2) and in certain network scenarios (Chapter 5).

2. **A different approach to source-channel communication.** The most common way to approach the source-channel communication problem is through Shannon's separation theorem. We call this approach *rate-matching* (as explained in the previous section). In this thesis, we propose a novel, alternative approach that we call *measure-matching*. This approach suggests new results, some of which can be found in Chapter 3.

3. **An autonomous theory of *joint* source-channel coding.** Joint source-channel coding is an area known for its ad-hoc concepts and lack of appropriate criteria (e.g., complexity); its theory is usually handled by the separation theorem. In this thesis, we present an autonomous theory of *joint* source-channel coding. In an operational sense, our theory is more general than the separation theorem: it provides explicit statements about arbitrary coding schemes, not only about separation-based designs.

4. **Capacity of large Gaussian relay networks.** In this thesis, we determine the capacity of certain large Gaussian relay networks. Our proof uses uncoded transmission as a tool (Chapter 5). This capacity result can also be extended to apply to certain common models of wireless networks.

5. **Optimal cost-distortion trade-offs in networks, and counterexamples to the separation paradigm.** The separation theorem does not extend to networks in general, and the optimal cost-distortion trade-offs in general networks are not known to date. In this thesis, we derive the optimal trade-off for two network topologies: for a single-source broadcast situation (Section 5.2), where only the Gaussian case has been known; and for a particular Gaussian sensor network situation (Section 5.4.5). Both cases lead to new examples of the fact that the separation paradigm does not extend to networks: in both cases, a separate source and channel code design leads to substantially suboptimal performance.

Outline

Chapter 1 discusses the source-channel communication problem using the separation theorem. This is motivated by the fact that in the standard literature

[7, 3, 4], the topic of source-channel communication tends to be treated as a corollary to the capacity and rate-distortion theorems. The goal of Chapter 1 is to invert this: the source-channel communication problem is treated as the fundamental problem, and one way to attack this fundamental problem is through the separation theorem, and hence through capacity and rate-distortion. Hence, most of Chapter 1 is a review. Towards the end, the chapter presents extensions of the separation theorem to networks, including certain new results.

Two well-known examples seem to suggest that the separation theorem does not explain all there is to know about the source-channel communication problem: When a binary uniform source is transmitted across a binary symmetric channel, uncoded transmission achieves just the same performance as the best and unboundedly complex coding system, designed e.g. according to the separation theorem. Similar behavior is observed for a certain scenario involving a Gaussian source and a Gaussian channel. Are these the only two "lucky" source-channel pairs? In **Chapter 2**, we derive conditions for the optimality of uncoded transmission, and we show that there is an unlimited number of examples where uncoded transmission is optimal, the condition being that the source and channel probability distributions, the distortion measure and the input cost function of the channel are favorably matched.

Chapter 3 discusses the source-channel communication problem using the *measure-matching* condition. This condition follows from the new results on uncoded transmission presented in Chapter 2; the match involves again the source and channel characteristics, but also the encoding and decoding functions. The measure-matching condition is used to establish a number of results, including the optimality of source-channel codes of finite block length, and a certain property of universality featured by source-channel coding systems. The measure-matching condition is also extended to communication systems with feedback.

The theory developed in Chapters 1 and 3 can be applied to a number of scenarios. **Chapter 4** is a brief intermezzo with the purpose of illustrating the interest in results about optimal uncoded transmission: we outline an application of our results to neural communication.

Chapter 5 investigates some of the potential of uncoded transmission in networks. First, in extension of the results developed in Chapters 2 and 3, we derive results about network source-channel communication, including a single-source broadcast and a multiple-description scenario. Thereafter, we show that uncoded transmission is also a useful tool beyond source-channel communication: we use it to prove capacity results for certain Gaussian relay networks, including a model of a wireless ad-hoc network. Using similar methods, we also analyze large Gaussian sensor networks. For sensor networks, the key question is the trade-off between the power used by the sensors and the quality at which the sensed signal can be reconstructed. The optimal trade-off is not known in general, but for a special situation, our results permit to determine it. This leads to a new example of the fact that the separation paradigm does not extend to networks.

Chapter 1

The Source-Channel Communication Problem, Part I

The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point.
– CLAUDE E. SHANNON, *A mathematical theory of communication [98]*

This chapter serves a double purpose: First, it sets the stage for subsequent chapters by quoting the relevant prior art. Second, many textbooks on information and coding theory do not explicitly treat the problem of joint source-channel coding, and if they do, it is simply appended at the end, as a corollary [3, 4]. As a case in point, the commemorative issue of the *IEEE Transactions on Information Theory* (published in October 1998) does not devote an article to this problem, either.

There are various ways to tackle the source-channel communication problem. The most successful and popular approach is Shannon's separation theorem [98, Thm. 21]. After quoting Shannon's definition of the communication problem (Section 1.1), the separation theorem for the point-to-point case is discussed in Section 1.2. Section 1.3 outlines briefly the limitations of the separation theorem. In particular, in the point-to-point case, if the communication system is non-ergodic, or if its delay and complexity are constrained, the separation theorem does not generally characterize the best possible performance. For topologies beyond point-to-point, the separation theorem does not generally apply, but it may in special cases. In Section 1.4, we discuss how the separation theorem may be extended. We make this precise for the case of feedback in Section 1.5, and in Section 1.6, we address genuine network topologies. It is

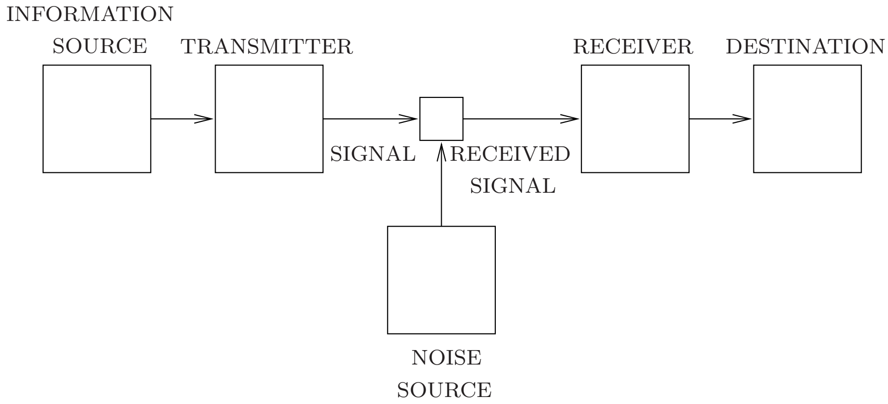

Figure 1.1: Shannon's "Schematic diagram of a general communication system."

first illustrated that the source-channel separation paradigm does *not* lead to optimal designs in general, i.e., no separation theorem applies. Thereafter, special topologies are presented for which a separation theorem *does* apply.

1.1 Shannon's Communication Problem

Shannon's communication system is best illustrated by the famous rendering in Figure 1.1 (taken from [98]). It consists of five parts: an *information source*, a *transmitter*, a *channel*, a *receiver*, and a *destination*. The crucial nature-given ingredients to Shannon's communication system are (informally):

- The information source, modeled by a random process. In Shannon's concept of communication, a deterministic source does not "produce" any information.
- The transmission channel, whose key ingredient is a random noise source. This models for example thermal noise in the electronic components, or fading effects due to (e.g.) multi-path propagation of the signal.

The task of the communication engineer is to design the boxes labeled *transmitter* and *receiver*, given the characteristics of the source and the channel.

1.1.1 Informal statement of the problem

The fundamental problem of communication is informally stated in the quote at the very beginning of this chapter. In this thesis, we are particularly interested in the second case: when the message is reproduced *approximately*; the exact reconstruction can be seen as a limiting case of the approximate reconstruction. Informally, the fundamental problem is captured by the following question: At fixed power budget, what is the highest quality at which the source can be represented at the destination?

More generally, "power" means a set of physical constraints on the channel, including bandwidth and battery charging properties, for example. Similarly, "quality" should be understood in a wide sense: If the source and/or the destination of information is human, then quality is determined for example by psycho-acoustic factors and the properties of the human visual system.

The goal is to find the best possible communication system for the given physical constraints. In this perspective, the complexity of computation as well as the delay due to such computation remain unconstrained; in other words, it is always assumed that an infinitely fast computer, with an infinite amount of memory and which runs at no power, is at hand. The goal is to obtain results that are valid "forever" within a given model of the physics of the communication problem, irrespective of the cleverness of the communications engineers and the advances in computer technology, for example.

1.1.2 Formal statement of the problem

A communication system is specified by six entities, grouped into three pairs: the source (p_S, d), consisting of a probability distribution p_S and a distortion function d; the channel $(p_{Y|X}, \rho)$, consisting of a conditional probability distribution $p_{Y|X}$ and a cost function ρ; and the code (F, G), consisting of the encoder F and the decoder G. The relationship between these six entities can be schematically rendered as in Figure 1.2. The following paragraphs serve to

Figure 1.2: The basic communication system.

make this brief summary more precise. For the purpose of this thesis, we are mostly concerned with discrete and finite alphabets.

Definition 1.1 (source) *A discrete-time memoryless source* (p_S, d) *is specified by a probability distribution* $p_S(s)$ *on an alphabet* $\mathcal{S}$ *and a nonnegative function*

$$d \quad : \quad \mathcal{S} \times \hat{\mathcal{S}} \to \mathbb{R}_+, \tag{1.1}$$

called the distortion measure. This implicitly specifies an alphabet $\hat{\mathcal{S}}$ *in which the source is reconstructed.*

When the alphabets are discrete*, we call this a* discrete *memoryless source, and the probability distribution becomes a probability mass function (pmf).*

Definition 1.2 (channel) *A discrete-time memoryless channel* $(p_{Y|X}, \rho)$ *is specified by a conditional probability distribution* $p_{Y|X}(y|x)$*, defined on two discrete alphabets* $\mathcal{X}$ *and* $\mathcal{Y}$ *and a nonnegative function*

$$\rho \quad : \quad \mathcal{X} \to \mathbb{R}_+, \tag{1.2}$$

called the channel input cost function.

When the alphabets are discrete*, we call this a* discrete *memoryless channel.*

Remark 1.1 (memory) *Many of the statements quoted in this chapter also apply or extend to sources and channels with memory, but for the purpose of this thesis, considerations are limited to memoryless source/channel pairs.*

Definition 1.3 (source-channel code of rate κ) *A source-channel code (F, G) of rate κ is specified by an encoding function*

$$F \quad : \quad \mathcal{S}^k \to \mathcal{X}^m, \tag{1.3}$$

and a decoding function

$$G \quad : \quad \mathcal{Y}^m \to \hat{\mathcal{S}}^k, \tag{1.4}$$

such that $k/m = \kappa$.

This means that even though the source and the channel are memoryless, the channel input sequence X^m is not a sequence of independent and identically distributed (iid) random variables in general, which implies that the channel output sequence Y^m and the source reconstruction sequence $\hat{S}^k$ are not generally iid sequences. This makes it necessary to define the distortion and the cost for sequences, rather than only for single letters as we have done above. This extension is straightforward: The distortion between sequences is simply the sample average of the component-wise distortions, and the cost of using a particular sequence is the sample average of the component-wise costs. If this component-wise decomposition does not apply, then we do not consider the source (p_S, d) or the channel $(p_{Y|X}, \rho)$, respectively, to be memoryless. In line with the standard literature [1, 3], we use the following notational conventions:

$$d(s^n, \hat{s}^n) \stackrel{def}{=} \frac{1}{n}\sum_{i=1}^{n} d(s_i, \hat{s}_i), \tag{1.5}$$

$$\rho(x^n) \stackrel{def}{=} \frac{1}{n}\sum_{i=1}^{n} \rho(x_i). \tag{1.6}$$

For a fixed source (p_S, d), a fixed channel $(p_{Y|X}, \rho)$ and a fixed code (F, G), we can then easily determine the average incurred distortion,

$$\Delta \stackrel{def}{=} Ed(S^k, \hat{S}^k), \tag{1.7}$$

and the average required cost,

$$\Gamma \stackrel{def}{=} E\rho(X^m), \tag{1.8}$$

and we define:

Definition 1.4 (cost-distortion pair) *For a fixed source (p_S, d), a fixed channel $(p_{Y|X}, \rho)$ and a fixed code (F, G), the cost-distortion pair (Γ, Δ) is given by (1.7) and (1.8), respectively.*

In this thesis, we will write (Γ, Δ) for the cost-distortion pair *achieved by a particular source-channel code* (F, G). A general cost-distortion pair (not necessarily achievable) will be denoted by (P, D).

A key question for the communications engineer is: Is a suggested source/channel code (F, G) optimal? Or would it be worth investing in the quest for better codes?

The first step to answering this is to agree on a definition of optimality. In our view, the most meaningful definition is the one requiring the cost Γ and the distortion Δ to be *simultaneously* optimal. In line with this, for the framework of this thesis, a communication system is considered optimal only if it lives up to the following definition.

Definition 1.5 *The performance* (Γ, Δ) *of a communication system is said to be optimal if*

1. Δ *could not be lowered without increasing the cost, and*

2. Γ *could not be lowered without increasing the distortion,*

irrespective of complexity and delay.[1]

The expression "irrespective of complexity and delay" looks somewhat imprecise, but it simply denotes the optimization over *all* source-channel codes of rate κ. Similar definitions are the OPTA[2] function [1, p. 156], and the LMTR[3] [4, p. 129, middle].

Remark 1.2 *We will see later on that the two conditions of Definition 1.5 imply one another in many cases of interest, but not always. The benchmark for optimality is always Definition 1.5.*

Remark 1.3 (optimality is a matter of marginals) *Optimality according to Definition 1.5 is a matter* only *of the marginal distributions of the channel inputs* $p(x_j)$, *for* $j = 1, \ldots, m$, *and of the (joint) marginal distributions of the source and corresponding reconstruction symbols,* $p(s_j, \hat{s}_j)$, *for* $j = 1, \ldots, k$. *In other words, any scheme that achieves the right marginals is optimal, irrespective of the correlation it introduces* among *the symbols.*

The communication problem according to Definition 1.5 is the trade-off between cost and distortion: The more cost one can use on the channel, the smaller

[1]Strictly speaking, according to this definition, many situations have no system with optimal performance. This was oberved by Prof. em. James L. Massey on the occasion of the *défense privée* of this thesis. While we feel it would be more appropriate to rewrite the thesis in that spirit, we choose to obey the delay constraints of thesis writing and continue to call an "optimal communication system" also the limit of the sequence of communication systems whose performance approaches the optimum (in line with a considerable part of the standard literature).

[2]optimum performance theoretically achievable

[3]limit of the minimum transmission ratio, see [4, p. 5]

the distortion one can achieve. For a given source (p_S, d) and a given channel $(p_{Y|X}, \rho)$, there is an entire set of optimal trade-offs.

One way to visualize this is to consider the *cost-distortion curve* for the source/channel pair. This curve gives, for every distortion Δ, the smallest possible cost Γ, and conversely, for every cost Γ, the smallest possible distortion Δ. A typical example of such a curve is shown in Figure 1.3. It can be shown that the cost-distortion curve is indeed convex, as suggested by the figure. The pair (Γ, Δ) in the figure represents a cost-distortion pair which is optimal by the standards of Definition 1.5; however, the pair (Γ', Δ') in the figure, even though on the curve, is not optimal: the cost could be reduced without changing the distortion. Hence, even though it carries the right intuition, the cost-distortion curve is not fully precise as a means of visualizing the optimal cost-distortion trade-off. We now clarify this issue.

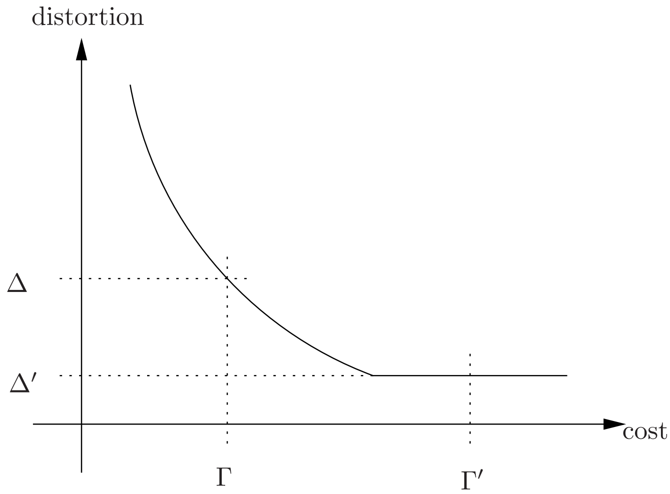

Figure 1.3: Schematic rendering of the cost-distortion trade-off.

The cost-distortion curve as drawn in Figure 1.3 is the union of the pairs (P, D) satisfying

$$D(P) = \lim_{n\to\infty} \min_{F_n, G_n : E\rho(X^n) \leq P} Ed(S^n, \hat{S}^n), \tag{1.9}$$

with the pairs (P, D) satisfying

$$P(D) = \lim_{n\to\infty} \min_{F_n, G_n : Ed(S^n, \hat{S}^n) \leq D} E\rho(X^n). \tag{1.10}$$

The pairs (P, D) satisfying (1.9) are precisely the pairs that satisfy the first condition of Definition 1.5. By analogy, the pairs (P, D) satisfying (1.10) are the pairs that satisfy the second condition of Definition 1.5. Hence, the optimal trade-offs are the pairs (P, D) in the *intersection* of the solutions to (1.9) and (1.10), rather than the union. This is why the pair (Γ', Δ') in Figure 1.3 does not represent an optimal communication system: it is only a solution to (1.9), not to (1.10).

The next step is to solve the optimization problems (1.9) and (1.10). The crux is that the optimization has to be carried out over all possible coding schemes; this could be a rather lengthy task. Once the optimization problems (1.9) and (1.10) are solved, the remaining task is to determine the *intersection* of their solution sets. The final result is a condition to check whether or not a given cost-distortion pair (Γ, Δ) represents an optimal trade-off.

The (very elegant) solution of this problem is due to Shannon [98]. His way of solving the problem is often called the *separation theorem*.

1.2 The Separation Theorem

The separation theorem owes much of its reputation to its operational relevance; here, however, we are interested in its ability to solve the optimization problems (1.9) and (1.10). The theorem is presented in two parts: Section 1.2.1 establishes a general bound on the performance of any (ergodic point-to-point) communication system; thereafter, it is shown in Section 1.2.2 that this bound is tight, i.e., that there exists always a scheme achieving (or at least approaching) the limit. The combination of these two parts yields the solution both to the optimization problem (1.9) and to the optimization problem (1.10).

1.2.1 Converse part

We start by giving the definitions of two well-known functions, the *rate-distortion function* and the *capacity-cost function*. They are defined to be the solutions of the following optimization problems involving the functional $I(\cdot;\cdot)$, called the *mutual information* and defined e.g. in [7, 3].

Definition 1.6 (rate-distortion function) *The rate-distortion function of the source* (p_S, d) *is defined as*

$$R(D) \quad = \quad \min_{p_{\hat{S}|S}: Ed(S,\hat{S}) \leq D} I(S;\hat{S}). \qquad (1.11)$$

Definition 1.7 (capacity-cost function) *The capacity-cost function of the channel* $(p_{Y|X}, \rho)$ *is defined as*

$$C(P) \quad = \quad \max_{p_X: E\rho(X) \leq P} I(X;Y). \qquad (1.12)$$

A special point of the capacity-cost function is the *unconstrained capacity* C_0:

Definition 1.8 (unconstrained capacity) *The* unconstrained capacity *of the channel* $(p_{Y|X}, \rho)$ *is the capacity of the channel disregarding input costs, that is*

$$C_0 \quad = \quad \max_{p_X} I(X;Y). \qquad (1.13)$$

Hence, C_0 is independent of the choice of ρ; it is solely a property of $p_{Y|X}$. When $\rho(x) < \infty, \forall x \in \mathcal{X}$, an equivalent definition is $C_0 = \lim_{P\to\infty} C(P)$.

It is important to note that the operational meaning of these two functions is not needed at this point; they are simply a shorthand for the corresponding optimization problems. In terms of these shorthands, we can express a general bound on the performance of any communication system in the following way.

Theorem 1.1 (separation theorem, converse part [98, Thm. 21]) *Any communication system for a discrete-time memoryless source and a discrete-time memoryless channel using a source-channel code of rate κ satisfies*

$$\kappa R(\Delta) \quad \leq \quad C(\Gamma), \tag{1.14}$$

where Δ is the distortion incurred and Γ the cost used by the communication system.

One reason why this theorem is called the *separation theorem* is because it permits one to *separate* the optimization problems (1.9) and (1.10) into two parts, the rate-distortion problem and the capacity-cost problem.

Proof. The proof given in [98, Thm. 21] is kept on a short and intuitive level. For completeness, we write it out for the case of discrete alphabets.

To make matters precise, we suppose that the code maps k source symbols onto m channel symbols, where, by assumption, $k/m = \kappa$. Then,

$$\begin{aligned}
kR(\Delta) &= kR(Ed(S^k, \hat{S}^k)) \\
&= kR\left(\frac{1}{k}\sum_{i=1}^{k} Ed(S_i, \hat{S}_i)\right) \\
&\overset{(a)}{\leq} k\sum_{i=1}^{k} \frac{1}{k} R(Ed(S_i, \hat{S}_i)) && (1.15) \\
&= \sum_{i=1}^{k} R(Ed(S_i, \hat{S}_i)) \\
&\leq \sum_{i=1}^{k} I(S_i; \hat{S}_i) && (1.16) \\
&= \sum_{i=1}^{k} H(S_i) - \sum_{i=1}^{k} H(S_i|\hat{S}_i) \\
&\overset{(b)}{\leq} \sum_{i=1}^{k} H(S_i) - \sum_{i=1}^{k} H(S_i|\hat{S}^n, S_{i-1}, \ldots, S_1) \\
&= H(S^k) - H(S^k|\hat{S}^k) && (1.17) \\
&= I(S^k; \hat{S}^k) && (1.18) \\
&\overset{(c)}{\leq} I(X^m; Y^m)
\end{aligned}$$

$$
\begin{aligned}
&= H(Y^m) - H(Y^m|X^m) \\
&= \sum_{i=1}^{m} H(Y_i|Y_{i-1}, \ldots, Y_1) - \sum_{i=1}^{m} H(Y_i|X^n, Y_{i-1}, \ldots, Y_1) \\
&= \sum_{i=1}^{m} H(Y_i|Y_{i-1}, \ldots, Y_1) - \sum_{i=1}^{m} H(Y_i|X_i) \qquad (1.19) \\
&\stackrel{(d)}{\leq} \sum_{i=1}^{m} H(Y_i) - \sum_{i=1}^{m} H(Y_i|X_i) \\
&= \sum_{i=1}^{m} I(X_i; Y_i) \qquad (1.20) \\
&\leq \sum_{i=1}^{m} C(E\rho(X_i)) \qquad (1.21) \\
&= m \sum_{i=1}^{m} \frac{1}{m} C(E\rho(X_i)) \\
&\stackrel{(e)}{\leq} mC\left(\frac{1}{m}\sum_{i=1}^{m} E\rho(X_i)\right) \qquad (1.22) \\
&= mC(E\rho(X^m)) \\
&= mC(\Gamma).
\end{aligned}
$$

The key arguments needed in the proof are the following:

(a) the convexity of the rate-distortion function,

(b) the fact that additional conditioning cannot increase entropy, see e.g. [7, Thm. 2.3.2] or [3, Thm. 2.6.5],

(c) the data processing inequality, see e.g. [7, Thm. 4.3.3] or [3, Thm. 2.8.1],

(d) again the fact that additional conditioning cannot increase entropy (or, more precisely, that removing conditioning cannot decrease entropy), and

(e) the concavity of the capacity-cost function.

Moreover, we have used the fact that the source is memoryless in (1.17), and the fact that the channel is memoryless in (1.19).

The remaining steps involve minor arguments, including for example the definition of the rate-distortion function in (1.16) and of the capacity-cost function in (1.21).

The proof of this theorem is implicit in the standard treatments on information theory, e.g., [7] and [3]. For continuous alphabets, there is a number of subtleties. For the rate-distortion side, see [1]. For the capacity side, the subtleties are explained in detail in [7, Thm.7.2.2]. □

The proof of the converse part of the separation theorem rests on two pillars. The first pillar is the reduction from a consideration of potentially infinitely long codes to single-letter quantities. This is enabled by the fact that additional conditioning can only decrease entropy, which is directly due to the properties of the logarithm. The second pillar is the data processing inequality, which is also directly due to the properties of the logarithm.

For this reason, it is not surprising that researchers have been tempted by the idea of replacing the logarithm by another function while retaining these two key features. We discuss one such approach below in Section 1.4.

1.2.2 Direct part

Theorem 1.1 naturally raises the question whether the inequality in (1.14) is tight: Is there always a communication system, potentially very complex, that achieves equality in (1.14)?

For a fixed source (p_S, d), a fixed channel $(p_{Y|X}, \rho)$, and a fixed *rate* κ of the source-channel code, we can illustrate (1.14) by plotting the rate-distortion function next to the capacity-cost function. Figure 1.4 illustrates the two different possible behaviors. The interesting class of source/channel pairs is illustrated by the solid lines. For these pairs, there exist cost-distortion pairs (P, D) such that $\kappa R(D) = C(P)$. In this section, we outline how Shannon argued that this also implies that there *exists* a sequence of encoder/decoder pairs that approaches equality in (1.14). The second possible behavior for a source/channel pair is illustrated by the dashed lines in Figure 1.4: For those source/channel pairs, there do not exist cost-distortion pairs (P, D) such that $\kappa R(D) = C(P)$. In other words, for those source/channel pairs, the capacity is always larger than the source entropy.[4]

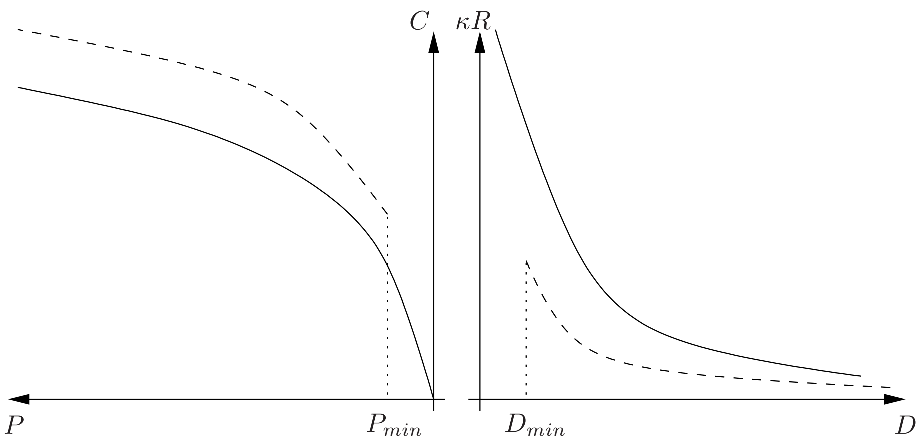

Figure 1.4: The capacity-cost function (left) and the rate-distortion function (right), schematic rendering.

This thesis focuses exclusively on source/channel pairs according to the solid

[4]Remark: To be more precise, it is not the entropy that is relevant, but the rate corresponding to the minimum distortion, $\sup_D R(D)$. For most interesting distortion measures, this is indeed the source entropy.

lines in Figure 1.4. Source/channel pairs according to the dashed lines are degenerate cases in the sense that the trade-off between cost and distortion is trivial: there is only one optimal operating point, namely $(P = P_{min}, D = D_{min})$.

Shannon's proof that there exist source-channel codes (F, G) that approach $\kappa R(\Delta) = C(\Gamma)$ as closely as desired (at least for the source/channel pairs according to the solid lines in Figure 1.4) is a direct consequence of the operational meaning of the rate-distortion function and of the capacity-cost function.

Somewhat informally, the operational rate-distortion problem can be stated as follows: For a given source (p_S, d), what is the least number of bits per source symbol needed to describe the source in such a way that these bits permit to reconstruct the source at an average distortion of at most D? It has been shown that this least number of bits is precisely the rate-distortion function $R(D)$. We quote the following result:

Theorem 1.2 (source coding) *It is possible to represent n outputs of the discrete-time memoryless source with pmf p_S by $2^{nR(D)}$ codewords of length n, $\{\hat{s}_1^n, \hat{s}_2^n, \ldots, \hat{s}_{2^{nR(D)}}^n\}$, such that* $\lim_{n\to\infty} E \min_i d(S^n, \hat{s}_i^n) \leq D$.

Proof. For a proof of this theorem, see for example [3, Thm. 3.2.1] or [7, Thm. 9.3.1].

By analogy, a somewhat informal statement of the capacity-cost problem is: For a given channel $(p_{Y|X}, \rho)$, what is the maximum number of bits that can be transmitted reliably when the permissible channel input cost is at most P? Note that the subtlety is in the word "reliable." It means that the error probability tends to zero as the coding complexity and delay become unconstrained. It has been shown that this maximum number of bits is precisely the capacity-cost function $C(P)$. We quote the following result:

Theorem 1.3 (channel coding) *For a discrete-time memoryless channel $(p_{Y|X}, \rho)$ with input X and output Y: for any $\epsilon > 0$ and any $R < C(P)$, there exist an n and 2^{nR} codewords of length n, $\{x_1^n, x_2^n, \ldots, x_{2^{nR}}^n\}$, where the components of the i-th codeword satisfy $\frac{1}{n}\sum_{j=1}^{n} \rho(x_{i,j}^n) \leq P$, for $i = 1, 2, \ldots, 2^{nR}$, along with a decoding function g, such that* $\max_i Pr(g(Y^n) \neq i | X^n = x_i^n) < \epsilon$.

Proof. For a proof of this theorem, see for example [7, Thms. 5.6.2, 7.2.1] or [3, Ch. 8].

The combination of these two theorems for the source/channel pairs represented by the solid lines in Figure 1.4 leads to the fact that there exist source-channel codes (F, G) that approach $\kappa R(\Delta) = C(\Gamma)$ as closely as desired: The source is described with $C(\Gamma)/\kappa - \epsilon$ bits per source symbol. The resulting bits can be transmitted without error across the channel (by Theorem 1.3); at the same time, they permit to reconstruct the source at distortion Δ (by Theorem 1.2). This leads to the following statement.

Theorem 1.4 (separation theorem, approachability) *For any discrete-time memoryless source/channel pair (p_S, d) and $(p_{Y|X}, \rho)$, and any $\epsilon > 0$,* *there exists a source-channel code (F, G) of rate κ using average cost Γ and incurring average distortion Δ such that either 1. or 2. is satisfied:*

1. $\kappa R(\Delta) = C(\Gamma) - \epsilon$
2. $\Delta = D_{min} + \epsilon$ *and* $\Gamma = P_{min} + \epsilon$.

Proof. A proof is given in [98, Thm. 21]; the theorem is an immediate consequence of the combination of Theorem 1.2 with Theorem 1.3. □

The combination of Theorems 1.1 and 1.4 permits one to solve both the optimization problem (1.9) and the optimization problem (1.10). This is discussed in the next section.

1.2.3 Optimal source-channel communication systems

Consider a source/channel pair (p_S, d) and $(p_{Y|X}, \rho)$ for which there exists values of P and D such that $\kappa R(D) = C(P)$. For this situation, Theorem 1.1 shows that any source-channel code (F, G) of rate κ must satisfy $\kappa R(\Delta) \leq C(\Gamma)$, and Theorem 1.4 shows that there exists a code (F, G) that satisfies $\kappa R(\Delta) = C(\Gamma)$. This suggests the condition that a communication system is optimal if and only if it satisfies $\kappa R(\Delta) = C(\Gamma)$.

Figure 1.5 illustrates that this is almost true, subject only to a minor technical condition: Both (Γ_1, Δ) and (Γ_2, Δ) satisfy $\kappa R(\Delta) = C(\Gamma)$, but it is clear that (Γ_1, Δ) is not an optimal cost-distortion trade-off: The cost can be reduced without changing the distortion. In summary, the following necessary and sufficient condition can be given.

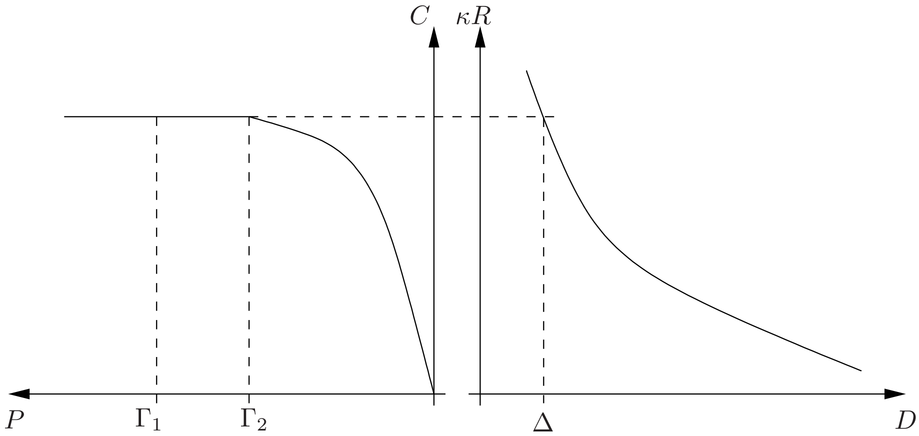

Figure 1.5: The capacity-cost function (left) and the rate-distortion function (right), schematic rendering.

Theorem 1.5 (separation theorem) *For a discrete-time memoryless source/channel pair (p_S, d) and $(p_{Y|X}, \rho)$ for which there exist values of P and*

D such that $\kappa R(D) = C(P)$, a cost-distortion trade-off (Γ, Δ) *is optimal if and only if*

(i) $\kappa R(\Delta) = C(\Gamma)$, *and*

(ii) Δ *cannot be lowered without increasing* $R(\Delta)$ *nor can* Γ *be lowered without decreasing* $C(\Gamma)$.

The following simpler (but less explicit) version of Theorem 1.5 will also be of interest later on. For this reason, we state it explicitly:

Corollary 1.6 *For a discrete-time memoryless source/channel pair* (p_S, d) *and* $(p_{Y|X}, \rho)$ *for which there exist values of P and D such that $\kappa R(D) = C(P)$, a cost-distortion trade-off* (Γ, Δ) *is achievable if and only if*

$$\kappa R(\Delta) \quad \leq \quad C(\Gamma). \tag{1.23}$$

Remark 1.4 *For source/channel pairs for which there do not exist values of P and D such that $\kappa R(D) = C(P)$, a code is optimal if and only if it satisfies Case 2 of Theorem 1.4. As pointed out earlier, these degenerate source/channel pairs will not be studied any further in this thesis.*

Remark 1.5 *Condition* (i) *of Theorem 1.5 characterizes the* union *of the solutions to the optimization problems (1.9) and (1.10). More precisely, by the aid of Theorem 1.5, the solution to the optimization problem (1.9) is found to be*

$$D \quad = \quad R^{-1}\left(\frac{C(P)}{\kappa}\right), \tag{1.24}$$

where $R^{-1}(\cdot)$ denotes the inverse of the rate-distortion function. By analogy, the solution to the optimization problem (1.10) is found to be

$$P \quad = \quad C^{-1}(\kappa R(D)), \tag{1.25}$$

where $C^{-1}(\cdot)$ denotes the inverse of the capacity-cost function.

In the simplest case when the cost-distortion pair (P, D) is such that $(P, C(P))$ is a point in the strictly concave region of the capacity-cost function, and $(D, \kappa R(D))$ is in the strictly convex region of the rate-distortion function, both (1.24) and (1.25) boil down to the condition $\kappa R(D) = C(P)$.

In other words, conditions (1.24) and (1.25) are necessary *conditions for the optimality of the cost-distortion pair (P, D); they describe the* union *of the solutions to (1.9) and (1.10). Condition (ii) then filters out their intersection.*

Remark 1.6 (optimality and marginals) *In Remark 1.3, we pointed out that optimality is a matter only of the right marginals. This follows directly from the definition of optimality. The separation theorem permits one to strengthen this insight: in an optimal system, the marginals of the channel inputs $p(x_j)$, for $j = 1, \ldots, m$, must all achieve the* same *Γ_j, and the (joint) marginal distributions of the source and corresponding reconstruction symbols, $p(s_j, \hat{s}_j)$, for*

$j = 1, \ldots, k$, must all achieve the same Δ_j. *This follows directly from the condition for equality in (1.15) and (1.22), respectively. In many cases, this implies that the marginals themselves must all be the same.*

Proof. By Theorem 1.1, any source-channel code (F, G) of rate κ satisfies

$$\kappa R(\Delta) \quad \leq \quad C(\Gamma). \tag{1.26}$$

Suppose that a certain source-channel code (F, G) satisfies

$$\kappa R(\Delta) \quad = \quad C(\Gamma) - \epsilon. \tag{1.27}$$

But then, by Theorem 1.4, there is a better source-channel code (F', G'), namely one that satisfies $\kappa R(\Delta) = C(\Gamma) - \epsilon'$, with $\epsilon' < \epsilon$. Hence, the optimal cost-distortion pair (Γ, Δ) satisfies $\kappa R(\Delta) = C(\Gamma)$. However, this is not sufficient as it may occur that Δ can be reduced without increasing $R(\Delta)$. This is prevented by condition *(ii)*. The same comment applies to $C(\Gamma)$. Hence, for source/channel pair (p_S, d) and $(p_{Y|X}, \rho)$ that admit cost-distortion pairs (P, D) satisfying $\kappa R(D) = C(P)$, *(i)* and *(ii)* together are necessary and sufficient conditions for optimality. □

Figure 1.5 illustrates Condition *(ii)* of Theorem 1.5 schematically. The following is a concrete example of a source/channel pair where the issue of Condition *(ii)* occurs.

Example 1.1 ($\kappa R(\Delta) = C(\Gamma)$ is not always sufficient) *Let all involved alphabets be $\mathcal{S} = \mathcal{X} = \mathcal{Y} = \hat{\mathcal{S}} = \{0, 1, \ldots, L-1\}$, where L is an even integer. The channel conditional pmf is the noisy typewriter channel as in [3, p. 185], that is, $p_{Y|X}(k|k) = 1/2$ and $p_{Y|X}((k+1) \mod L|k) = 1/2$, for all k. The unconstrained capacity (Definition 1.8) of this channel is found to be $C_0 = \log_2 \frac{L}{2}$. Let the encoder and decoder be the identity function. For the source pmf, define $p_{odd}(s)$ to be the uniform pmf over the odd inputs, and $p_{even}(s)$ the uniform pmf over the even inputs. Let the source pmf be a convex combination of these two, i.e. $p_\lambda(s) = \lambda p_{odd}(s) + (1-\lambda) p_{even}(s)$, where $0 \leq \lambda \leq 1$. Notice that $p_\lambda(s)$ achieves capacity on the unconstrained noisy typewriter channel for any λ.*

Define the following distortion measure:

$$d(s, \hat{s}) \quad = \quad \begin{cases} 0, & \hat{s} = s \text{ or } \hat{s} = (s+1) \mod L \\ 1, & \text{otherwise.} \end{cases} \tag{1.28}$$

Certainly, $\Delta = Ed(S, \hat{S}) = 0$. Moreover, we find that for any λ,

$$R(\Delta = 0) \quad = \quad \log_2 \frac{L}{2}. \tag{1.29}$$

Let the input cost function be

$$\rho(x) \quad = \quad \begin{cases} 1, & x \text{ even,} \\ 0, & x \text{ odd.} \end{cases} \tag{1.30}$$

Suppose now that the source has $\lambda = 1/2$*. Is the overall communication system optimal in that case? For* $\lambda = 1/2$*, we compute* $\Gamma = \frac{L}{2}$*, and hence*

$$C(\Gamma) \quad = \quad C_0 = \log_2 \frac{L}{2}. \tag{1.31}$$

Evidently, the condition $R(\Delta) = C(\Gamma)$ *is satisfied. However, this is* not *an optimal communication system. Consider for example the source with parameter* $\lambda = 1$*. We compute* $\Gamma' = 0 < \Gamma$*, but clearly,* $C(\Gamma) = C(\Gamma')$*. Hence, Condition (ii) of Theorem 1.5 is violated: It is indeed possible in this case to lower* Γ *without changing* $C(\Gamma)$*. Practically, this means that for the source with parameter* $\lambda = 1/2$*, there exists a* coded *communication system that achieves the same distortion but requires lower cost.*

As a last remark, let us point out that the fact that the distortion and the cost Γ' *are zero is* not *crucial for this example.*

In summary, let us emphasize that the key contents of Theorem 1.5 is the condition $\kappa R(\Delta) = C(\Gamma)$. Condition *(ii)* concerns cases of limited interest for the purpose of this thesis. We will discuss that condition, and hence the problem illustrated by Example 1.1, in more detail in Section 2.2.2. Another case of limited interest is Case *2.* of Theorem 1.4. As pointed out in Remark 1.4, this issue will not be studied any further in this thesis.

1.2.4 The double role of the separation theorem

So far, we have discussed the optimal cost-distortion trade-off for a given source/channel pair. More precisely, the goal was to determine the intersection of the solutions to the optimization problem (1.9) and the optimization problem (1.10).

Shannon solved this by the separation theorem, which splits the problem into two subproblems, the rate-distortion and the capacity-cost problem. This is the first role of the separation theorem.

The second role of the separation theorem is its practical significance. Loosely speaking, it states that an optimal communication system can be implemented by cascading an optimal source coding system with an optimal channel coding system. These two component coding systems can be designed independently of one another: the design of the channel code does not involve any knowledge of the source properties, nor does the design of the source code require knowledge about the channel. Hence, from a practical point of view, the separation theorem is a very powerful tool: it permits one to build optimal communication systems in a modular fashion. Source coding techniques can be devised for sources of practical interest, irrespective of the channel to be used. Similarly, channel coding systems, once developed, serve equally well for the transmission of all kinds of sources.

Conceptually, the separation theorem leads thus to a double perspective onto the source-channel communication problem which can be described as follows: on the channel side, the goal is to eliminate all the randomness by the aid of suitable coding, and turn the channel into what is sometimes called an "error-free bit-pipe". On the source side, the set of possible source outputs is partitioned into disjoint regions in an optimal fashion. The source-channel communication problem is turned into the problem of *rate-matching*: the rate of the source code is matched to the rate of the channel code. In Chapter 3, we revisit the source-channel communication problem, investigating an alternative point of view.

1.3 The Separation Theorem, Revisited

Clearly, it is tempting to apply the strategy of separating source from channel coding to *any* communication system, even to those for which it has not been established that it leads to *optimal* performance. In the sequel, we will use the term *separation-based design* precisely to denote the design strategy of splitting the coding system into two steps in such a way that the first step (the source coding) handles all the distortion, while the second step (the channel coding) provides a noise-free channel. In contrast to this, the term *separation theorem* continues to mean exclusively Theorem 1.5 above: that in certain situations, the separation-based design *does* lead to optimal performance.

In point-to-point communication, there are two key challenges to the separation-based approach: First, as soon as delay and complexity are constrained, it does not lead to optimal designs anymore, and second, the separation theorem does not apply to non-ergodic sources and channels in general. A third key challenge concerns communication networks: even for quite simple network topologies, the separation-based approach may lead to a considerably suboptimal performance. This will be discussed in Section 1.6 below.

A more subtle limitation of the separation principle is that it invariably leads to a *deterministic* end-to-end mapping, i.e., the same realization s^n of a source output sequence leads (with high probability) to the same realization $\hat{s}^n$ at the destination. This is not required for optimality according to Definition 1.5.[5] As a matter of fact, Chapters 2 and 3 focus on systems that are optimal according to Definition 1.5, yet are not asymptotically deterministic in the above sense.

1.3.1 Delay and complexity

In the proof of the achievability, one has to allow for codes of arbitrary length. No practical system can be expected to be as generous as that. In other words, as soon as constraints on the coding complexity and on the delay are imposed, the separation theorem *strictu sensu* is no longer applicable. As a matter of fact, it is easy to devise catastrophic examples: when the channel code is of finite length, it cannot prevent all errors. However, depending on the source code, one single error on the channel may have a catastrophic effect on the source reconstruction. To prevent this, it is necessary to design the channel

[5] This is further elaborated in [87].

code *taking into account* the properties of the source, and vice versa. This is called a *joint source/channel code*.

A vast literature concerns somewhat ad-hoc, but often imaginative joint source-channel coding schemes. Recent research efforts concern unequal error protection, including [14]. Another effort investigates so-called index assignment methods, by which one means that the indices of the source codebook are cleverly assigned to the channel codewords. An early example of an index assignment method is the Gray code: codewords representing adjacent integers differ by only one binary digit. For recent overviews, we refer to [69]; for more elaborate source-channel scenarios, see e.g. [84]. It is, however, beyond the scope of this thesis to give a comprehensive guide to joint source-channel coding techniques.

On the theoretical side, the key challenge of joint source-channel coding is to address the trade-off between the performance of the system and the required coding complexity and delay. However, it seems very difficult to obtain meaningful answers.

1.3.2 Non-ergodic systems

The second important challenge to the general relevance of the separation principle arises in the context of non-ergodic systems, which is of great practical interest in some applications. Vembu, Verdú and Steinberg [105] constructed an insightful (yet slightly contrived) example of the fact that the separation principle does not extend to general non-ergodic situations. We give a simplified and shortened version of their example:

Example 1.2 (non-ergodic source/channel pair [105]) *The source switches between a Bernoulli*$(1/2)$ *source (state 1) and a deterministic source (state 2), and the channel between a perfect binary channel (state 1) and a binary symmetric channel of transition probability* $1/2$ *(state 2). The switching is* synchronous, *i.e., either both are in state 1, or both are in state 2. Clearly, irrespective of the switching schedule, the source can be transmitted perfectly across the channel without any further coding.*

The trick of [105] is to devise a switching schedule that makes the capacity of the channel smaller *than the entropy of the source; in other words, the separation principle would lead one to guess that error-free transmission is impossible.*

The particular switching schedule of [105] is deterministic*: switching occurs right before times* 2^i, $i = 1, 2, 3, \ldots$. *It is not hard to verify that the entropy rate of the source under this switching schedule becomes* $2/3$, *while the guaranteed capacity of the channel is only* $1/3$ *(see [108]).*

Note that this example also underlines the wide validity of the separation theorem: Counterexamples have to be constructed carefully (or so it seems). An extension of [105] can be found in [24].

1.4 Other Separation Theorems

In Shannon's original paper [98], the separation theorem is presented in the section entitled "The rate for a source relative to a fidelity criterion" (as Theorem 21), and the converse to the separation theorem is proved by the following argument:

"...follows immediately from the definition of R_1 and previous results."

R_1 is precisely the rate-distortion function, and the "immediately" refers to the fact that had it been possible to send the source through a channel of $C < R_1$ and still achieve a fidelity D, then this would contradict the definition of R_1.

While this may not be entirely rigorous, it immediately suggests the discovery of other separation theorems: wherever a rate-distortion theorem has been established, corresponding separation theorems can be formulated. Candidates are therefore the Wyner-Ziv rate-distortion function, and the multiple description problem (for the Gaussian source and mean-square error). These cases will be discussed in Section 1.6.

But Shannon's insight also works in the other direction: had it been possible to send the source of $R(D) > C$ through a channel of capacity C and still achieve a fidelity D, then this would, by the same token, contradict the definition of C. Hence, for any capacity result, corresponding separation theorems can be formulated. A prime example is the memoryless channel with feedback. This will be discussed in Section 1.5

Before stating these generalizations, we point to an extension of the separation theorem that is of a quite different flavor.

Other functionals satisfying a data processing inequality

The key ingredient that makes the separation theorem work is the data processing inequality. It has been noted by Ziv and Zakai in [117] that mutual information is not the only functional that satisfies such an inequality. As the block length of the involved codes tends to infinity, the separation theorem following from the mutual information functional is tight (it is both an upper bound and achievable). Hence, in that case, replacing mutual information by another functional cannot give better bounds. However, for codes of finite block length, the separation theorem is not tight. More precisely, it predicts an unachievable cost-distortion trade-off. In that case, replacing mutual information by a different functional may actually lead to better bounds on the achievable performance. In [117], Ziv and Zakai give an example of a source/channel pair and an alternative functional which results in a tighter bound on the achievable cost-distortion trade-off.

1.5 The Separation Theorem with Feedback

In this section, we consider the situation where the encoder has perfect (causal) feedback: at time n, it knows the received symbols $Y_{n-1}, Y_{n-2}, \ldots$. This is illustrated in Figure 1.6. To recall that the encoder has feedback available, we denote it by F°. The definition of the channel (Definition 1.2) is a somewhat

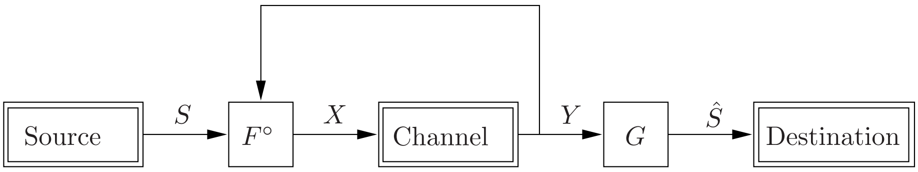

Figure 1.6: The basic communication system with feedback.

subtle task in the presence of feedback. For example, Definition 1.2 is an indefensible definition of a discrete-time *memoryless* channel when feedback is used. Rather, the definition has to be extended by the following condition (see [74]): a channel is called memoryless if it satisfies

$$p(y_i|x^i, y^{i-1}) \quad = \quad p_{Y|X}(y_i|x_i). \tag{1.32}$$

For the case of a discrete-time memoryless source/channel pair, the separation theorem applies.

Theorem 1.7 *For a discrete-time memoryless source* (p_S, d) *and a discrete-time memoryless channel* $(p_{Y|X}, \rho)$ *(where memoryless is taken in the sense of Equation (1.32)),* whether or not feedback is used*, a cost-distortion pair* (Γ, Δ) *is achievable for a source-channel coding rate* κ *if and only if it satisfies*

$$\kappa R(\Delta) \quad \leq \quad C(\Gamma). \tag{1.33}$$

Remark 1.7 *This separation theorem seems to extend also to systems with memory: For sources with memory, it can be established along the lines of the proof given below. For* channels *with memory, the argument is somewhat subtler because no simple characterization of their feedback capacity is known. The discussion in this thesis is limited to the memoryless case.*

Proof. The converse follows from a rather simple modification of the proof that feedback cannot increase the capacity of a discrete-time memoryless channel, see e.g. [7, p.520] or [3, p.214].

Another way to prove the converse part of this theorem is by the aid of the elegant concept of *directed information* [74]. We restrict our version of the proof to the case of discrete alphabets. Massey defines the directed information in [74] as follows:

$$I(X^m \to Y^m) \quad \stackrel{def}{=} \quad \sum_{i=1}^{m} I(X^i; Y_i | Y^{i-1}). \tag{1.34}$$

We refer the reader to [74, 66] for excellent treatments of directed information. For the purpose of our proof, we only need two properties of the functional

(1.34). The first property is the inequality[6]

$$I(X^m \to Y^m) \quad \leq \quad \sum_{i=1}^{m} I(X_i; Y_i), \tag{1.35}$$

with equality if and only if $Y_1, Y_2, \ldots, Y_m$ are independent [74]. The second property is that for the scenario of Figure 3.3,

$$I(S^k; Y^m) \quad \leq \quad I(X^m \to Y^m), \tag{1.36}$$

which is also proved in [74].

For any source-channel communication system using a block feedback encoder of rate κ,

$$\begin{aligned} kR(\Delta) &\overset{(a)}{\leq} I(S^k; \hat{S}^k) \\ &\overset{(b)}{\leq} I(S^k; Y^m) \\ &\overset{(c)}{\leq} I(X^m \to Y^m) \\ &\overset{(d)}{\leq} \sum_{i=1}^{m} I(X_i; Y_i) \\ &\overset{(e)}{\leq} mC(\Gamma), \end{aligned} \tag{1.37}$$

where (a) follows by the same arguments as Equation(1.18), (b) is the data processing inequality, see e.g. [7, Thm. 4.3.3] or [3, Thm. 2.8.1], (c) and (d) are the properties of directed information mentioned above and are proved in [74], and (e) holds by the same arguments as from Equation (1.20) onwards.

The achievability is immediate since feedback does not decrease the capacity: simply ignore the feedback and use Theorem 1.4. □

Remark 1.8 *There is one subtle difference between the separation theorems with and without feedback: in the latter case, the separation theorem says that for any system*

$$kR(\Delta) \leq I(S^k; \hat{S}^k) \leq I(X^m; Y^m) \leq mC(\Gamma), \tag{1.38}$$

and hence, an optimal system satisfies this with equality throughout (except for certain degenerate cases), i.e., a necessary condition for an optimal coding system is

$$I(S^k; \hat{S}^k) = I(X^m; Y^m). \tag{1.39}$$

With feedback, this is not *a necessary condition; it may happen that*

$$I(X^m; Y^m) > mC(\Gamma). \tag{1.40}$$

[6]See Remark 1.8 for an explanation that, in the feedback case, the inequality $I(X^m; Y^m) \leq \sum_{i=1}^{m} I(X_i; Y_i)$ does not hold.

To illustrate this by a very simple example, take the binary symmetric channel of capacity zero. Suppose that the feedback scheme is $X_i = Y_{i-1}$, for $i = 2, 3, \ldots$. Then, using, first, the chain rule of mutual information, and then, its definition,

$$I(X_1, X_2; Y_1, Y_2) = I(X_2; Y_1) + A = H(Y_1) + A, \qquad (1.41)$$

where $A \geq 0$ by the non-negativity of mutual information. This is certainly larger than zero.

This argument shows that $I(X^m; Y^m)$ is not a quantity of interest in feedback systems; a system may be optimal even though it satisfies

$$I(S^k; \hat{S}^k) < I(X^m; Y^m). \qquad (1.42)$$

These observations are of importance to the discussion in Section 3.5 below.

1.6 The Separation Theorem in Communication Networks

1.6.1 General network situations

The separation-based design can also be extended to networks. In this section, it is briefly discussed how this can be done, and then shown (by well-known counterexamples) that this does not generally lead to optimal communication systems. This is the third and probably most fundamental challenge to the paradigm of separation.

To illustrate this point, we prefer to consider the simple network topology shown in Figure 1.7, rather than to attack the most general case. There are two channel inputs, X_1 and X_2, with their input constraints $E\rho_1(X_1) \leq P_1$ and $E\rho_2(X_2) \leq P_2$, respectively. There are three sources and two destinations. Each destination wishes to reconstruct two of the sources, as shown in the figure. The first destination reconstructs S_1 at distortion $D_1 = Ed_1(S_1, \hat{S}_1)$ and S_3 at distortion $D_{31} = Ed_{31}(S_3, \hat{S}_{31})$. The second destination reconstructs S_2 at distortion $D_2 = Ed_2(S_2, \hat{S}_2)$ and S_3 at distortion $D_{32} = Ed_{32}(S_3, \hat{S}_{32})$. Finally, there are two encoders and two decoders available, with inputs and outputs as drawn in the figure. The basic question is again: What are optimal cost-distortion tuples $(P_1, P_2, D_1, D_2, D_{31}, D_{32})$?

By analogy to the point-to-point problem, a network performs optimally if and only if all elements of the cost-distortion tuple $(P_1, P_2, D_1, D_2, D_{31}, D_{32})$ are simultaneously minimal. Hence, the goal is to formulate optimization problems like (1.9) and (1.10), and again to study the intersection of their solution sets. The individual optimization problems in the spirit of (1.9) will be of the form

$$D_1(D_2, D_{31}, D_{32}, P_1, P_2) = \min_{F_{12}, F_3, G_1, G_{23}} Ed_1(S_1, \hat{S}_1), \qquad (1.43)$$

where the minimization is taken over all codes F_{12}, F_3, G_1, G_{23} such that the other distortions are at most D_2, D_{31}, D_{32} and the costs are at most P_1, P_2.

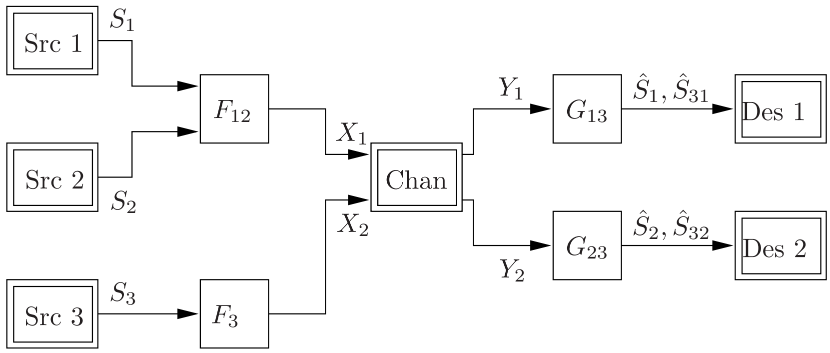

Figure 1.7: Simple example of a source-channel network.

Remark 1.9 (optimality and marginals) *In Remark 1.3, we pointed out that in the point-to-point case, optimality is a matter only of the right marginals. This follows directly from the definition of optimality. Clearly, for the same reason, this remark also applies to networks.— Remark 1.6 then showed that all the marginals have to be (essentially) identical for optimal performance. This followed from the separation theorem. Since there is no separation theorem for the network case, we cannot readily infer that the property of identical marginals also applies to networks.*

The notions of capacity and rate-distortion have been extended to the network case. For a channel network with K desired connections of rates $R_1, R_2, \ldots, R_K$, respectively, the set of achievable rate tuples is generally called the *capacity region* of the network [4, p. 271]. We denote it by

$$\mathcal{C}(P_1, P_2, \ldots, P_K), \tag{1.44}$$

where $P_1, P_2, \ldots, P_K$ denote the respective input constraints. Similarly, for a network of sources with M desired reconstructions, the set of achievable rate tuples is generally called the *rate(-distortion) region* of the source network. We denote it by

$$\mathcal{R}(D_1, D_2, \ldots, D_M), \tag{1.45}$$

where $D_1, D_2, \ldots, D_M$ denote the requested distortions for each of the reconstructions.

By complete analogy to the point-to-point case, the rate-distortion and the capacity results can be combined to yield the following statement:

Theorem 1.8 *If*

$$\mathcal{R}(D_1, D_2, \ldots, D_M) \cap \mathcal{C}(P_1, P_2, \ldots, P_K) \quad \neq \quad \emptyset, \tag{1.46}$$

then $(D_1, D_2, \ldots, D_M, P_1, P_2, \ldots, P_K)$ *is achievable.*

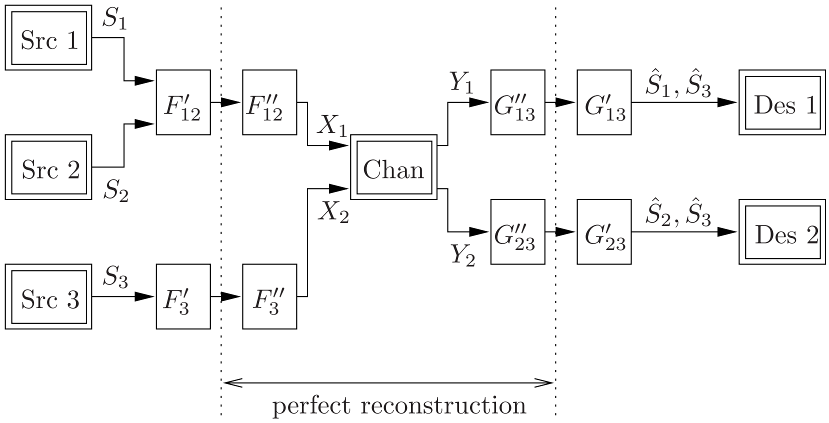

Figure 1.8: It is generally *suboptimal* to implement a network communication system in this fashion, i.e., by making the channel asymptotically error-free.

Proof. The proof follows directly from the operational meaning of the rate-distortion and capacity region. □

The converse to Theorem 1.8 is not true: even if rate-distortion region and capacity region do not intersect, a cost-distortion tuple may be achievable. In other words, it is *not* true that an optimal system can be implemented as suggested by Figure 1.8. We illustrate this by two examples. The first example concerns a multi-access situation and is taken from [3, p. 449].

Example 1.3 (multi-access with correlated sources [3]) *Two (correlated) binary sources X_1 and X_2 are to be transmitted over a multiple access channel. The channel takes binary inputs, and its output is the (real) sum of the inputs. It can be shown that its capacity region is contained in the triangle $0 \leq R_1 + R_2 \leq 1.5$. Let the sources have the following joint distribution: $(X_1 = 1, X_2 = 0)$ never occurs; the other three events are equally likely. To encode this source in a lossless way, a total rate of $R_1 + R_2 = \log_2 3 = 1.585$ bits is required at least. Assuming the separation principle applies, we conclude that it is not possible to transmit this source over the above channel without error. However, if we let the channel inputs be the two sources without further encoding, then their sum uniquely determines both X_1 and X_2.*

It is relatively easy to locate the problem with the separation paradigm in this example: The capacity region for the multi-access channel was computed supposing *independent* sources. Clearly, if the sources can be dependent, the capacity region is generally larger. This is precisely what happens in this example. If, in the definition of the capacity region, we allow for arbitrarily dependent channel inputs, then the separation theorem applies again, and the problem pointed out by the above example does not arise. This is not at all surprising: allowing for arbitrarily dependent sources turns the problem effectively into a point-to-point problem. Yet, the gist of this example is subtler than one might think at this point of the discussion: the two channel inputs *cannot* be arbitrarily dependent since they have to be separate. Slepian-Wolf coding turns the originally dependent sources into two new sources that are *independent*. For the source coding problem, this is optimal. However, for the joint source/channel coding problem, removing the dependence is not optimal: the dependence might be fit to the channel (as it is the case in this example), and once it is removed, it cannot be brought in again. A similar example, featuring a Gaussian multiple-access channel, is given below in Section 5.4.6.

While the above example is interesting, it may not convince the reader of the shortcomings of the separation paradigm in a network context. To this end, we consider a second well-known example involving the transmission of one Gaussian source to many users. As we will see, it is harder to locate the breakdown of the separation paradigm in this case.

Example 1.4 (Gaussian single-source broadcast) *An iid Gaussian source S of mean zero and variance P is to be transmitted across a Gaussian two-user broadcast channel with power constraint P and with noise variances $\sigma_1^2 < \sigma_2^2$. Denote the capacities of the two underlying point-to-point channels by $C_1 = 1/2\log_2(1 + P/\sigma_1^2)$ and $C_2 = 1/2\log_2(1 + P/\sigma_2^2)$. We want $D_2 = D_{\mathcal{N}}(C_2)$. What is the smallest achievable D_1?*

First, suppose we use a separation-based design. In order to satisfy the constraint on D_2, we have to use a capacity-achieving code of rate C_2 to send to user 2 (the worse user). Can we superimpose another code for user 1? The answer is no: any superposition would be noise for user 2, and hence compromise D_2. Since the broadcast channel at hand is degraded, user 1 can also decode the codewords destined for user 2, hence the smallest D_1 that can be achieved with this scheme is $D_1 = D_2 = D_{\mathcal{N}}(C_2)$.

However, it is quickly verified[7] for this example that simply sending S without further coding achieves the distortions $D_1^* = D_{\mathcal{N}}(C_1)$ and $D_2^* = D_{\mathcal{N}}(C_2)$. In other words, the cost-distortion tuple (P, D_1^*, D_2^*) does not satisfy the relationship (1.46), yet it is achievable. Further details for this setup are given in Example 5.1 of Chapter 5.

The failure of the separation theorem in this example is considerably more intricate than in Example 1.3. In Section 5.2 of Chapter 5, we provide an explanation in terms of the concepts that are introduced in Chapters 2 and 3 of this thesis.

While these two examples clearly show that the separation-based design does not lead to an optimal system in general network scenarios, the following sections concern special networks where separation *does* lead to optimal systems.

[7]See also Example 2.2 below.

1.6.2 Independent sources on the multi-access channel

The simplest extension of the separation theorem to a network context is the case of independent sources, transmitted across a set of independent point-to-point channels. In extension of this, we now consider independent sources, transmitted across a *multiple-access channel*.

The rate-distortion region for independent sources is known; it follows directly from the rate-distortion functions of the sources. Following Shannon's insight, this leads to potentially multiple separation theorems, including the setup of Figure 1.9. For this setup, the separation theorem says that an optimal

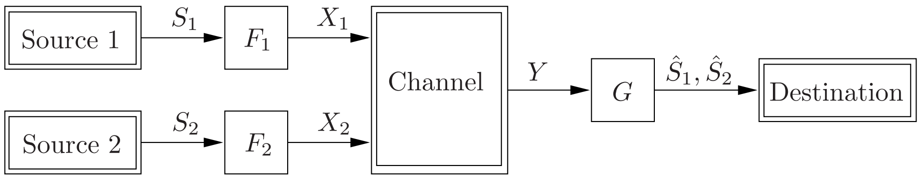

Figure 1.9: Independent sources on the multi-access channel.

communication strategy can be implemented by compressing each source separately and transmitting the resulted codewords across the multi-access channel, using a multi-access code that achieves (approaches) the boundary of the capacity region of the multi-access channel. This is certainly not an unexpected fact.

More precisely, for given costs P_1 and P_2 at the respective inputs of the multi-access channel, a capacity region

$$\mathcal{C}(P_1, P_2) \tag{1.47}$$

can be determined, see e.g. [3, Thm. 14.3.1]. Moreover, for given distortions D_1 and D_2, the rate region is given by

$$R(D_1, D_2) \quad = \quad \{(R_1, R_2) : R_1 \geq R_{S_1}(D_1), R_2 \geq R_{S_2}(D_2)\}, \tag{1.48}$$

where R_{S_i} denotes the rate-distortion function of source i. To emphasize that this rate region is essentially a single-terminal object (it is rectangular), we use the notation $R(D_1, D_2)$, reserving the script notation $\mathcal{R}(\cdot)$ for true multi-terminal regions.

Theorem 1.9 (independent sources on the MAC) *For independent discrete-time memoryless sources on the discrete-time memoryless MAC, a cost-distortion trade-off* $(P_1, P_2), (D_1, D_2)$ *is achievable if and only if*[8]

$$R(D_1, D_2) \cap \mathcal{C}(P_1, P_2) \neq \emptyset. \tag{1.49}$$

[8] Here, we also assume that the sources and the channel are such that there exist values of $(P_1, P_2), (D_1, D_2)$ that satisfy Condition (1.49).

Proof. The achievability follows immediately from the operational meaning of the rate-distortion function and of the capacity region. For the converse, suppose that a certain code (F_1, F_2, G) achieves $(P_1, P_2), (D_1, D_2)$. By the point-to-point separation theorem,

$$\begin{aligned} nR_{S_1}(D_1) &\overset{(a)}{\leq} I(S_1^n; Y^n) \\ &\overset{(b)}{\leq} I(S_1^n; Y^n | X_2^n) \\ &\overset{(c)}{\leq} I(X_1^n; Y^n | X_2^n), \end{aligned}$$

where (a) follows immediately from the proof of Theorem 1.1, (b) holds because S_1 and X_2 are independent and (c) is the data processing inequality. For the sum rate bound, note that for any source-channel code of block length n,

$$\begin{aligned} nR_{S_1}(D_1) + nR_{S_2}(D_2) &\overset{(a)}{\leq} I(S_1^n; \hat{S}_1^n) + I(S_2^n; \hat{S}_2^n) \\ &\overset{(b)}{\leq} I(S_1^n, S_2^n; \hat{S}_1^n, \hat{S}_2^n) \\ &\overset{(c)}{\leq} I(X_1^n, X_2^n; Y^n), \end{aligned}$$

where (a) follows from the definition of the rate-distortion function (and is analogous to the proof of Theorem 1.1), (b) holds because S_1 and S_2 are independent, and (c) is the data processing inequality. By assumption, X_1 and X_2 satisfy the cost constraints, and since they are generated independently from independent random variables S_1 and S_2, they can at most have a dependency of the time-sharing kind (see e.g. [3, p.397]). Hence, by the definition of the capacity region, the expressions $I(X_1^n; Y^n | X_2^n)$, $I(X_2^n; Y^n | X_1^n)$, and $I(X_1^n, X_2^n; Y^n)$ describe a point inside the capacity region $\mathcal{C}(P_1, P_2)$ of the multi-access channel. □

1.6.3 The Wyner-Ziv separation theorem

In this section, we study the communication scenario depicted in Figure 1.10. The double boxes mark the source, the channel and the destination. Filling the single boxes is the task of the communications engineer.

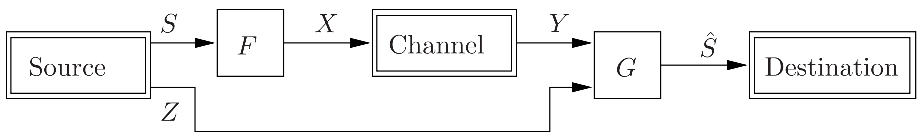

Figure 1.10: The Wyner-Ziv communication scenario.

In the spirit of the separation paradigm, we hope to implement an optimal system as follows: For the channel, we use a capacity-achieving code, thus turning it into an essentially error-free bit pipe of C bits per channel use. Then, we encode the source at C bits per source sample, knowing that the decoder will have side information Z available. This special rate-distortion problem has been studied by Wyner and Ziv, and the solution is known as the Wyner-Ziv rate-distortion function that we denote as $R^{WZ}_{S|Z}(D)$, see e.g. [3, Section 14.9].

Such a separation theorem can indeed by proved for the communication scenario of Figure 1.10. It is given in the following theorem.

Theorem 1.10 (Wyner-Ziv separation theorem) *For the Wyner-Ziv communication scenario with a discrete-time memoryless Wyner-Ziv source and a discrete-time memoryless channel, the cost-distortion pair* (P, D) *is achievable if and only if*[9]

$$R^{WZ}_{S|Z}(D) \quad \leq \quad C(P). \tag{1.50}$$

Proof. Achievability is a straight consequence of the operational meaning of the Wyner-Ziv rate-distortion function, see [3, p. 438]. For the converse, we can write out as follows:

$$\begin{aligned} nR^{WZ}_{S|Z}(D) &= n \min_{W:\exists g \text{ s.t. } Ed(S,g(W,Z))\leq D} (I(S;W) - I(Z;W)) \\ &\overset{(a)}{=} \min_{W^n:\exists g \text{ s.t. } Ed(S^n,g(W^n,Z^n))\leq D} (I(S^n;W^n) - I(Z^n;W^n)) \\ &\overset{(b)}{\leq} I(S^n;Y^n) - I(Z^n;Y^n) \\ &\leq I(S^n;Y^n) \\ &\overset{(c)}{\leq} I(X^n;Y^n) \\ &\leq nC(P), \end{aligned}$$

where (a) follows from the proof of the Wyner-Ziv rate-distortion function (see e.g. [3, Section 14.9]), (b) follows because by assumption, Y^n is such that there exists a decoding function g yielding $Ed(S^n, g(Y^n, Z^n)) \leq D$, hence it is in the set over which the minimization is performed, and (c) is the data processing theorem. □

The Wyner-Ziv rate-distortion function has been determined for discrete memoryless sources in [113], and extended to sources with continuous alphabets in [112]. Unlike in the scenario of Slepian and Wolf [100], there is usually a rate loss due to the fact that the side information is not known at the encoder [113, 115], but in the special case where the source and the side information are jointly Gaussian, there is none: the Wyner-Ziv rate-distortion function coincides with the conditional rate-distortion function (defined e.g. in [48]). In extension of this fact, the Wyner-Ziv rate-distortion function has also been determined for Gaussian vector sources [37, 81].

[9] Here, we also assume that the source and the channel are such that there exist values of (P, D) that satisfy Condition(1.50).

A more general version of the simple Wyner-Ziv separation theorem above was derived by Merhav and Shamai [75]: they consider the case of the Wyner-Ziv source and the Gel'fand-Pinsker channel.[10]

Another extension of the simple Wyner-Ziv separation theorem has been presented by Shamai, Verdú and Zamir in [96, 97]. Their scenario is illustrated Figure 1.11. There are two channels, channel A and channel D. Channel A is the "analog" channel, i.e., its input alphabet $\mathcal{S}'$ is the same as the source alphabet $\mathcal{S}$. Channel A has capacity C_A.

The practical motivation for the study of Figure 1.11 is as follows: Channel A, together with the encoder F', models an existing system (e.g., analog television). Channel D are additional resources, available to preferred customers. In particular, the encoder F' is fixed to be the identity map (i.e., uncoded transmission), and the question is: when is it possible to cater to the preferred customers at the same quality as if F' could be designed freely? Notice that the analog system by itself (the non-preferred customer's system) is *not* required to perform optimally. Denoting the conditional rate-distortion function (defined e.g. in [48]) by $R_{S|Z}(D)$, the answer can be stated as follows:

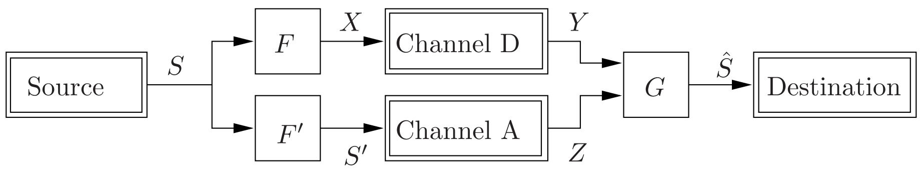

Figure 1.11: The communication scenario studied by Shamai, Verdú and Zamir in [97].

Theorem 1.11 (Shamai, Verdú, Zamir) *There exists an optimal code (F, G) in the communication scenario of Figure 1.11 that permits F' to be the identity mapping (i.e., uncoded transmission) if and only if the following three conditions are satisfied:*

(i) $I(S; Z) = C_A$

(ii) $R_{S|Z}(D) = R_{S|Z}^{WZ}(D)$

(iii) $R_S(D) = R_{S|Z}(D) + I(S; Z)$,

where all involved quantities are calculated for the case where F' in Figure 1.11 is the identity mapping.

For a proof, the reader is referred to the original paper [97]. Theorem 1.11 is a separation theorem in the sense that it determines the performance of the optimal communication system for the setup of Figure 1.11 in terms of rate-distortion and capacity-cost functions.

[10] After establishing their separation theorem, Merhav and Shamai discuss uncoded transmission for the same scenario, using an approach that was inspired by the results presented in Chapters 2 and 3 of this thesis.

1.6.4 The multiple-description separation theorem

Another rate-distortion problem which has been solved (at least in one special case) is the multiple-description problem. This leads to the consideration of the communication scenario in Figure 1.12. There is one source, S, and one encoder F, emitting two codewords. One is sent through Channel 1, the other through Channel 2. There are three destinations interested in reconstructing S. Destination 0 observes both channel outputs, destination 1 only the output of Channel 1 and destination 2 only the output of Channel 2. For given costs P_1 and P_2 on the two channels, respectively, what is the set of achievable distortions D_0, D_1 and D_2?

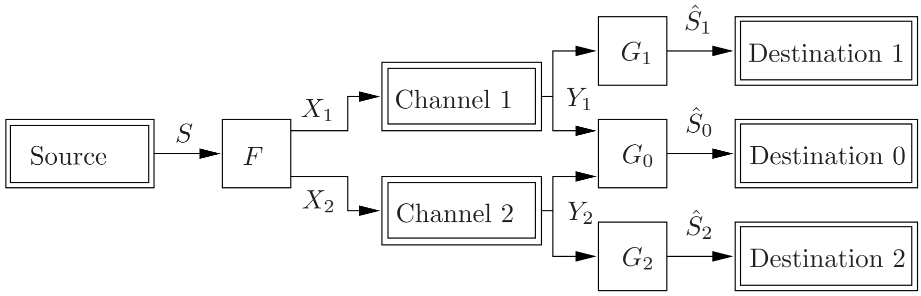

Figure 1.12: The multiple description communication scenario.

The corresponding source coding problem is the multiple-description source coding problem. The Gaussian case has been solved by Ozarow in [80].[11] The rate-distortion region $\mathcal{R}^{(MD)}(D_0, D_1, D_2)$ is given by [8]

$$\begin{aligned} R_1 &\geq R(D_1) & (1.51)\\ R_2 &\geq R(D_2) & (1.52)\\ R_1 + R_2 &\geq R(D_0) + \gamma_R(D_0, D_1, D_2), & (1.53) \end{aligned}$$

where $\gamma_R = 0$ for $D_1 + D_2 - D_0 \geq 1$, and

$$\gamma_R = \frac{1}{2}\log_2 \frac{(1-D_0)^2}{(1-D_0)^2 - \left(\sqrt{(1-D_1)(1-D_2)} - \sqrt{(D_1-D_0)(D_2-D_0)}\right)^2} \quad (1.54)$$

for $D_1 + D_2 - D_0 < 1$.

The capacity region for two independent parallel channel is trivial,

$$C(P_1, P_2) = \{(R_1, R_2) : 0 \leq R_1 \leq C_1(P_1), 0 \leq R_2 \leq C_2(P_2)\}. \quad (1.55)$$

Theorem 1.12 *For the multiple description communication system with an iid Gaussian source and mean-squared error distortion, and two independent*

[11]Recent advances in multiple description coding [106, 83] have provided new achievable rate points, but no converses.

discrete-time memoryless channels, the cost-distortion tuple $(P_1, P_2, D_0, D_1, D_2)$ *is achievable if and only if*

$$\mathcal{R}^{(MD)}(D_0, D_1, D_2) \cap C(P_1, P_2) \neq \emptyset. \quad (1.56)$$

Proof. We follow along the lines of Ozarow [80]. If the source-channel code (F, G_0, G_1, G_2) achieves the cost-distortion tuple $(P_1, P_2, D_0, D_1, D_2)$, then, by the point-to-point separation theorem (more precisely, by Theorem 1.1),

$$\begin{aligned} R(D_1) &\leq C_1(P_1) \\ R(D_2) &\leq C_2(P_2). \end{aligned}$$

Moreover,

$$\begin{aligned} n(R(D_0) + \gamma_R(D_0, D_1, D_2)) &\overset{(a)}{\leq} I(S^n; \hat{S}_0^n) \overset{(b)}{\leq} I(S^n; Y_1^n, Y_2^n) \\ &\overset{(c)}{\leq} I(S^n; Y_1^n) + I(S^n; Y_2^n) \\ &\overset{(b)}{\leq} I(X_1^n; Y_1^n) + I(X_2^n; Y_2^n) \\ &\overset{(d)}{\leq} C_1(P_1) + C_2(P_2). \end{aligned}$$

where (a) holds by the definition of the multiple-description rate-distortion region, (b) is the data processing inequality, (c) is due to the fact that Y_1 and Y_2 are conditionally independent given S, as follows:

$$I(S^n; Y_1^n, Y_2^n) = I(S^n; Y_1^n) + I(S^n; Y_2^n | Y_1^n).$$

The last term can be developed

$$\begin{aligned} I(S^n, Y_1^n; Y_2^n) &= I(S^n; Y_2^n) + I(Y_1^n; Y_2^n | S^n) \\ &= I(Y_1^n; Y_2^n) + I(S^n; Y_2^n | Y_1^n), \end{aligned}$$

and since $I(Y_1^n; Y_2^n | S^n) = 0$, it follows that

$$I(S^n; Y_2^n | Y_1^n) = I(S^n; Y_2^n) - I(Y_1^n; Y_2^n) \leq I(S^n; Y_2^n).$$

Inequality (d) holds by the definition of capacity. □

1.7 Problems

Problem 1.1 *Consider two binary sources* S_1 *and* S_2. S_1 *is simply a Bernoulli-1/2 process, and* S_2 *can be defined as* $S_2 = S_1 + E$, *where* E *is Bernoulli-p. The two sources should be compressed independently of one another.* S_1 *is compressed using* R_1 *bits,* S_2 *is compressed using* R_2 *bits.*

(i) The goal of a common decoder is to reconstruct S_1 *and* S_2 *without error, as illustrated in Figure 1.13. Determine the achievable rate region.*

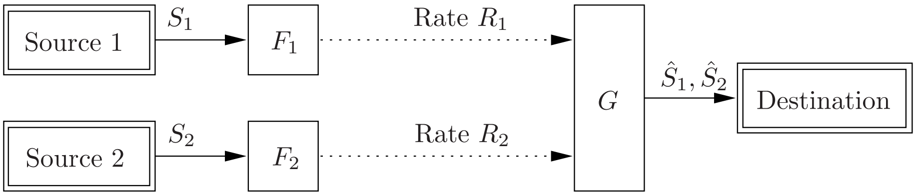

Figure 1.13: The problem considered in Part *(i)*.

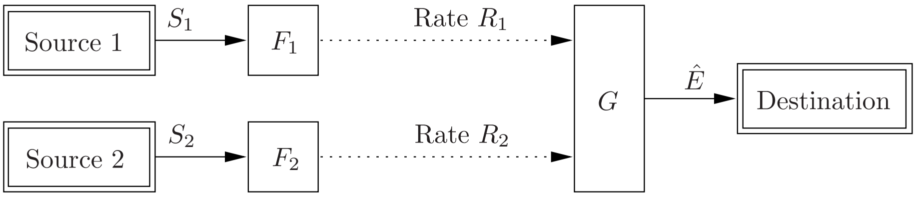

Figure 1.14: The problem considered in Part *(ii)*.

(ii) The goal of a common decoder is to reconstruct E without error, as illustrated in Figure 1.14. Show that the rates $R_1 = R_2 = h_b(p)$[12] *are sufficient. (This result is due to Körner and Marton [62].)*

(iii) Consider the transmission of the sources S_1 and S_2 across the additive modulo-2 multiple-access channel,

$$Y = X_1 \oplus X_2, \tag{1.57}$$

where $\oplus$ denotes modulo-2 addition. The goal of a common decoder is again to reconstruct E without error, like in Part (ii). The situation is illustrated in Figure 1.15.

(a) Suppose that a separation-based design is used: S_1 and S_2 are encoded using the optimal scheme of Part (ii). The resulting codewords are transmitted using a capacity-achieving code for the above multi-access channel. Show that it is not always possible for the decoder to reconstruct E.

(b) Show that the strategy in Part (a) is suboptimal, in line with Example 1.3. Hint: *Indicate a better strategy for "large" p, i.e., $0.11 < p < 0.5$. Recall the title of this thesis.*

Solution 1.1 *For (i), the Slepian-Wolf region is easily determined (see e.g. [3, Thm. 14.4.1]) in terms of $h_b(p)$:*

$$\begin{aligned} R_1 &\geq H(S_1|S_2) = h_b(p) \\ R_2 &\geq H(S_2|S_1) = h_b(p) \\ R_1 + R_2 &\geq H(S_1, S_2) = 1 + h_b(p), \end{aligned}$$

[12]Here, $h_b(\cdot)$ denotes the binary entropy function, i.e., $h_b(p) = -p \log_2 p - (1-p) \log_2 (1-p)$

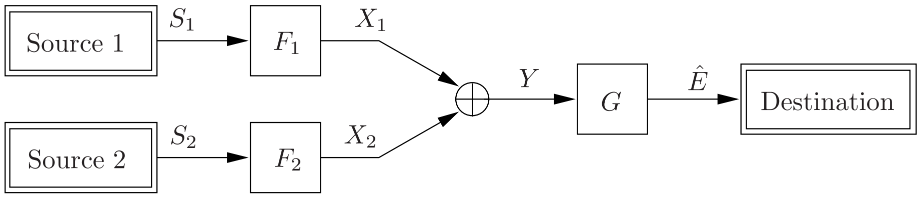

Figure 1.15: The problem considered in Part *(iii)*.

For (ii), the idea is for both S_1 *and* S_2 *to encode only the "residual" with respect to the other. More precisely, encoder* F_1 *uses the Slepian-Wolf coding technique as if the decoder knew* S_2*, see [100], and encoder* F_2 *does the same as if the decoder knew* S_1*. It is shown in [62] that the resulting two codewords enable the decoder to perfectly reconstruct* E*.*

For (iii), the maximum sum rate on the binary modulo-2 MAC is 1*. Coding according to (ii) requires a sum rate of* $2h_b(p)$*. For large* p*, this is larger than* 1*, hence the separation approach of Part (a) will not work. However, for Part (b), straight uncoded transmission obviously always works.*

Chapter 2

To Code, Or Not To Code?

The ergodic point-to-point source-channel communication problem can be solved by the separation theorem. As reviewed in Chapter 1, it gives the solution to the *theoretical* side of the source-channel communication problem: it permits one to determine the optimal cost-distortion trade-offs.

Moreover, the separation theorem also provides a solution to the *practical* side of the source-channel communication problem, at least in an asymptotic sense: as the encoding and the decoding are allowed to incur unbounded delay and complexity, optimal source-channel communication schemes can be implemented in a two-stage fashion: optimal source compression followed by optimal channel coding. This has split the research community into two camps, those who examine source compression and those who investigate channel coding.

But the separation theorem does not claim to be the *only* solution to the source-channel communication problem. In fact, it is a very expensive solution, owing to its disregard of delay and complexity issues. In sharp contrast to this, it is well known that in certain examples, *uncoded transmission* achieves optimal cost-distortion trade-offs. The behavior observed in these examples cannot be explained from the perspective of the separation theorem: uncoded transmission is a *joint* source-channel "coding" scheme.

In this chapter, we first review the two well-known examples of such behavior (Section 2.1). In Section 2.2, we develop a theory that explains the optimality of uncoded transmission and extends the well-known examples. In Section 2.3, we illustrate our findings by examples, but also by a number of lemmata whose proof is enabled or simplified by the results of Section 2.2.[1]

[1]The contents of this chapter is (up to editorial changes) identical with [43]. Sections 2.1 and 2.2.4 have been added, and Section 2.3 has been extended.

2.1 Two Inspiring Examples

The first example involves a binary source/channel pair. It appears e.g. in [9, Sec. 11.8], or [10, p. 117].

Example 2.1 (Binary/Hamming) *In this example, the source is defined by the binary uniform random variable S, and the distortion measure*

$$d(s,\hat{s}) = \begin{cases} 0, & \text{if } s=\hat{s}, \\ 1, & \text{otherwise.} \end{cases} \tag{2.1}$$

This is usually called Hamming *distortion. The channel is the binary symmetric channel, i.e., it has binary input X and binary output Y, and conditional probability mass function*

$$p(y|x) = \begin{cases} 1-\epsilon, & \text{if } y=x, \\ \epsilon, & \text{otherwise.} \end{cases} \tag{2.2}$$

We assume that $\epsilon < 1/2$. Suppose the channel is used once per source symbol. The goal is to determine the smallest achievable average distortion D_{min}. In this example, $\max_D R(D) = H(S) = 1 \geq C$, which implies that there are values of D such that $R(D) = C$ (the channel is unconstrained in this example, hence there is no cost P involved). This means that Theorem 1.5 applies. Hence, the minimum distortion D_{min} satisfies

$$R(D_{min}) = C. \tag{2.3}$$

For the binary symmetric channel, $C = 1 - h_b(\epsilon)$,[2] and for the binary uniform source with Hamming distortion, $R(D) = 1 - h_b(D)$. Hence,

$$D_{min} = \epsilon. \tag{2.4}$$

Suppose now that the source symbol S is directly used as the channel input X, and that the channel output Y is the source estimate $\hat{S}$. This "uncoded" transmission scheme achieves the following average distortion:

$$\begin{aligned} \Delta &= Ed(S,\hat{S}) \\ &= 1 \cdot Prob(S \neq \hat{S}) + 0 \cdot Prob(S = \hat{S}) \\ &= \epsilon. \end{aligned}$$

Hence, uncoded transmission is optimal in this example.

This example illustrates the fact that an optimal cost-distortion trade-off may be achievable by a very simple coding scheme. More precisely, note that the optimal scheme designed according to the separation theorem requires as a component a capacity-achieving code for the binary symmetric channel. Such a code requires infinite delay, which follows immediately from Theorem 1.4: For

[2] Here, $h_b(\cdot)$ denotes the binary entropy function, i.e., $h_b(p) = -p\log_2 p - (1-p)\log_2(1-p)$

the system to work, the error probability on the channel must be vanishingly small. For this to happen on the binary symmetric channel, infinite delay is necessary.

Apart from the delay (and complexity) issues, there are two other key differences between the two communication schemes: First, note that the uncoded communication system does *not* implement a deterministic end-to-end mapping (see also Section 1.3). Second, the uncoded communication system is less flexible than the one designed according to the separation principle (see also Section 1.2.4). This point is further discussed in [73].

The second well-known example where uncoded transmission is optimal involves the iid Gaussian source and the additive white Gaussian noise channel [47].

Example 2.2 (Gaussian/MSE) *Consider a discrete-time source whose outputs are independent and identically distributed (iid) Gaussian random variables S with zero mean and variance σ_S^2. The distortion measure is the mean-square error,*

$$d(s, \hat{s}) \quad = \quad (s - \hat{s})^2. \tag{2.5}$$

The channel is the standard additive white Gaussian noise (AWGN) channel with real-valued input X and real-valued output Y, i.e.,

$$Y = X + Z, \tag{2.6}$$

where Z is white Gaussian noise of variance σ^2. The channel input is power constrained: $EX^2 \leq P$. Suppose that the channel is used once per source symbol. The rate-distortion function for the Gaussian source and mean-square error is known to be

$$R(D) \quad = \quad \begin{cases} \frac{1}{2}\log_2 \frac{\sigma_S^2}{D}, & \text{if } D < \sigma_S^2, \\ 0, & \text{otherwise.} \end{cases} \tag{2.7}$$

The capacity of the AWGN channel is

$$C(P) \quad = \quad \frac{1}{2}\log_2\left(1 + \frac{P}{\sigma^2}\right). \tag{2.8}$$

Since $\max_D R(D) \geq C(P)$, there are values of P and D such that $R(D) = C(P)$, and hence, Theorem 1.5 applies: an optimal power-distortion pair (P, D) satisfies

$$R(D) \quad = \quad C(P). \tag{2.9}$$

Hence, the smallest achievable distortion is found to be

$$D_{min} \quad = \quad \frac{\sigma_S^2 \sigma^2}{P + \sigma^2}. \tag{2.10}$$

Next, we consider the performance of a particular joint source-channel coding strategy. The encoder is defined by the symbol-by-symbol mapping

$$X = \sqrt{\frac{P}{\sigma_S^2}} S, \tag{2.11}$$

and the decoder by the symbol-by-symbol mapping

$$\hat{S} = \sqrt{\frac{\sigma_S^2}{P}} \frac{P}{P+\sigma^2} Y. \tag{2.12}$$

The resulting system is illustrated in Figure 2.1. The distortion can be computed simply by evaluating the expectation:

$$\begin{aligned} \Delta &= E|S-\hat{S}|^2 \\ &= \left(1-\frac{P}{P+\sigma^2}\right)^2 E|S|^2 + \frac{\sigma_S^2}{P}\left(\frac{P}{P+\sigma^2}\right)^2 E|Z|^2 \\ &= \frac{\sigma_S^2\sigma^4}{(P+\sigma^2)^2} + \frac{\sigma_S^2\sigma^2 P}{(P+\sigma^2)^2} \\ &= \frac{\sigma_S^2\sigma^2}{P+\sigma^2}. \end{aligned}$$

That is, the nearly uncoded transmission strategy defined by Equations (2.11) and (2.12) performs optimally.

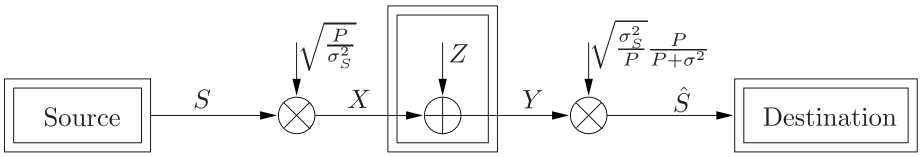

Figure 2.1: An optimal coding scheme for the transmission of an iid Gaussian source across an AWGN channel.

This Gaussian example appears in many places; it is often credited to Goblick [47]. It has been extended and reconsidered many times, including [32, 47, 76].

2.2 Single-letter Codes that Perform Optimally

It is well known that there are instances of source/channel pairs for which single-letter codes achieve the best possible performance. This result is particularly surprising since such codes are extremely easy to implement and operate at zero delay. In this section, we derive necessary and sufficient conditions under which single-letter codes are optimal. In line with Definition 1.3, we define:

Definition 2.1 (single-letter source-channel code) *A single-letter source-channel code (f, g) is specified by an encoding function $f(\cdot) : \mathcal{S} \to \mathcal{X}$ and a decoding function $g(\cdot) : \mathcal{Y} \to \hat{\mathcal{S}}$.*

Note that for single-letter codes, $\kappa = 1$. Theorem 1.5 contains two conditions that together are necessary and sufficient to establish the optimality of any communication system, including those that use single-letter codes. These conditions will now be examined in detail. In Section 2.2.1, we elaborate on the first condition, i.e., $R(\Delta) = C(\Gamma)$. The second condition is somewhat subtler; it will be discussed in Section 2.2.2. In Section 2.2.3, the results are combined to yield a general criterion for the optimality of single-letter codes.

Figure 2.2: The symbol-by-symbol communication system: $f : \mathcal{S} \to \mathcal{X}, g : \mathcal{Y} \to \hat{\mathcal{S}}$.

2.2.1 Condition (i) of Theorem 1.5

Optimal communication systems satisfy the conditions of the separation theorem, Theorem 1.5. In this section, we examine condition *(i)* of Theorem 1.5 for the special case when the code is a single-letter code. Since a single-letter code has $\kappa = 1$, that condition becomes $R(\Delta) = C(\Gamma)$. As a first step, we reformulate this condition more explicitly as follows:

Lemma 2.1 *For a discrete-time memoryless source/channel pair* (p_S, d) *and* $(p_{Y|X}, \rho)$ *and a single-letter code* (f, g) *as in Figure 2.2,* $R(\Delta) = C(\Gamma)$ *holds if and only if the following three conditions are simultaneously satisfied:*

(i) the distribution p_X *of* $X = f(S)$ *achieves capacity on the channel* $(p_{Y|X}, \rho)$ *at input cost* $\Gamma = E\rho(X)$*, i.e.,* $I(X;Y) = C(\Gamma)$*,*

(ii) the conditional distribution $p_{\hat{S}|S}$ *of* $\hat{S} = g(Y)$ *given* S *achieves the rate-distortion function of the source* (p_S, d) *at distortion* $\Delta = Ed(S, \hat{S})$*, i.e.* $I(S;\hat{S}) = R(\Delta)$*, and*

(iii) $f(\cdot)$ *and* $g(\cdot)$ *are such that* $I(S;\hat{S}) = I(X;Y)$*, i.e.,* $f(\cdot)$ *and* $g(\cdot)$ *are "information lossless."*

Proof. For any source-channel communication system that employs a single-letter code,

$$\begin{aligned} R(\Delta) &= \min_{q_{\hat{S}|S}: Ed(S,\hat{S}) \le \Delta} I(S;\hat{S}) \overset{(a)}{\le} I(S;\hat{S}) \overset{(b)}{\le} I(X;Y) \\ &\overset{(c)}{\le} \max_{q_X : E\rho(X) \le \Gamma} I(X;Y) = C(\Gamma), \end{aligned} \tag{2.13}$$

where (b) is the data processing inequality. Equality holds in (a) if and only if $p_{\hat{S}|S}$ achieves the rate-distortion function of the source, and in (c) if and only if p_X achieves the capacity-cost function of the channel. Thus, $R(\Delta) = C(\Gamma)$

is satisfied if and only if all three conditions in Lemma 2.1 are satisfied, which completes the proof. □

Remark 2.1 (test channel) *The conditional distribution achieving the rate-distortion function is sometimes called the "test channel", see e.g. [3, p. 345]. In this terminology, uncoded transmission is optimal if the actual channel is precisely the test channel for the source,*[3] *and the cost function for the channel is chosen appropriately.*

There are various ways to verify whether the requirements of Lemma 2.1 are satisfied. Some of them lead to problems that notoriously do not admit analytical solutions. For example, we could compute the capacity-cost function $C(\cdot)$ of the channel $(p_{Y|X}, \rho)$ and evaluate it at Γ. This is known to be a problem that does not have a closed-form solution for all but a small set of channels. Similarly, one could compute the rate-distortion function $R(\cdot)$ of the source (p_S, d) and evaluate it at Δ. Again, closed-form solutions are known only for a handful of special cases. Hence, following this approach, we have to resort to numerical methods via the Arimoto-Blahut algorithm.

More precisely, fix the channel conditional distribution to be $p_{Y|X}$. For a given cost function ρ, there is no general closed-form expression for the channel input distribution that achieves capacity. The key idea of the following lemma is to turn this game around: for any distribution q_X over the channel input alphabet $\mathcal{X}$, there exists a closed-form solution for the input cost function ρ such that the distribution q_X achieves capacity.

Lemma 2.2 *For fixed discrete source distribution p_S, single-letter encoder f and discrete channel conditional distribution $p_{Y|X}$ with unconstrained capacity C_0 (see Definition 1.8):*

(i) If $I(X;Y) < C_0$, the first condition of Lemma 2.1 is satisfied if and only if the input cost function satisfies,

$$\rho(x) \begin{cases} = & c_1 D(p_{Y|X}(\cdot|x)||p_Y(\cdot)) + \rho_0, \quad \text{if } p(x) > 0, \\ \geq & c_1 D(p_{Y|X}(\cdot|x)||p_Y(\cdot)) + \rho_0, \quad \text{otherwise,} \end{cases} \tag{2.14}$$

where $c_1 > 0$ and ρ_0 are constants, and $D(\cdot||\cdot)$ denotes the Kullback-Leibler distance between two distributions.

(ii) If $I(X;Y) = C_0$, the first condition of Lemma 2.1 is satisfied for any function $\rho(x)$.

Proof. See Appendix 2.A.

This lemma appears in a different shape as Problem 2 (without explicit proof) in [4, p. 147]; it can also be seen as an extension of Theorem 4.5.1

[3] Prof. em. James L. Massey remarked this on the occasion of the *défense privée* of this thesis and told the author that this insight had been part of the lectures of Claude E. Shannon.

of [7] to the case of constrained channel inputs. To gain insight, let q_X be the channel input distribution induced by some source distribution through the encoder f. For any cost function ρ, one finds an expected cost and a set of admissible input distributions leading to the same (or smaller) average cost. The input distribution q_X lies in that set, but it does not necessarily maximize mutual information. The key is now to find the cost function, and thus the set of admissible input distributions, in such a way that the input distribution q_X maximizes mutual information within the set. In the special case where the input distribution q_X achieves C_0, it clearly maximizes mutual information among distributions in *any* set, regardless of ρ. Hence, in that case, the choice of the cost function ρ is unrestricted.

Lemma 2.2 gives an explicit formula to select the input cost function ρ for given channel conditional and input distributions. By analogy, the next lemma gives a similar condition for the distortion measure.

Lemma 2.3 *For fixed discrete source distribution p_S, discrete channel conditional distribution $p_{Y|X}$ and single-letter code (f, g):*

(i) If $0 < I(S;\hat{S})$, the second condition of Lemma 2.1 is satisfied if and only if the distortion measure satisfies

$$d(s,\hat{s}) \quad = \quad -c_2 \log_2 p(s|\hat{s}) + d_0(s), \tag{2.15}$$

where $c_2 > 0$ and $d_0(\cdot)$ is an arbitrary function.

(ii) If $I(S;\hat{S}) = 0$, the second condition of Lemma 2.1 is satisfied for any function $d(s,\hat{s})$.

Proof. See Appendix 2.A.

This lemma should be understood by complete analogy to Lemma 2.2; it appears in a different shape as Problem 3 (without explicit proof) in [4, p. 147]. That is, let $q_{\hat{S}|S}$ be the conditional distribution induced by some channel conditional distribution through the encoder f and the decoder g. For any distortion measure d, an average distortion $\Delta = E_{q_{\hat{S}|S}} d(S,\hat{S})$ can be computed, which implies a set of alternative conditional distributions that also yield distortion Δ. The key is to find d in such a way that the chosen $q_{\hat{S}|S}$ minimizes $I(S;\hat{S})$ among all conditional distributions in the set.

Apparently, there is a slight asymmetry between Lemmata 2.2 and 2.3: In the former, when $p(x) = 0$, $\rho(x)$ satisfies a less stringent condition. In the latter, a similar behavior occurs: when $p(s,\hat{s}) = 0$, the condition can be relaxed to $d(s,\hat{s}) \geq -c_2 \log_2 p(s|\hat{s}) + d_0(s)$. However, since the right hand side is infinity in that case, requiring equality is equivalent.

In summary, our discussion of the requirement $R(\Delta) = C(\Gamma)$ produced a set of explicitly verifiable conditions that together ensure $R(\Delta) = C(\Gamma)$. However, to obtain an explicit criterion that can establish the optimality of a single-letter code, it still remains to scrutinize the condition *(ii)* of Theorem 1.5. This is the goal of the next section.

2.2.2 Condition (ii) of Theorem 1.5

Theorem 1.5 contains two simultaneous requirements to ensure the optimality of a communication system that employs single-letter codes. The first requirement, $R(\Delta) = C(\Gamma)$, was studied and developed in detail in Section 2.2.1; in this section, we examine the second condition, namely, when is it impossible to lower Δ without changing $R(\Delta)$ and when it is impossible to lower Γ without changing $C(\Gamma)$. This permits us to give a general criterion to establish the optimality of any communication system that uses single-letter codes.

The crux of the problem is illustrated in Figure 2.3. It shows simultaneously the capacity-cost function of the channel (left) and the rate-distortion function of the source (right). Problematic cases may occur only if either $R(\cdot)$ or $C(\cdot)$ are horizontal, i.e. when they have reached their asymptotic values $R(D \to \infty)$ and $C(P \to \infty)$. This happens only when the mutual information is zero or C_0 (the unconstrained capacity of the channel, see Definition 1.8). For example, both the cost-distortion pair (Γ_1, Δ) and the cost-distortion pair (Γ_2, Δ) satisfy the condition $R(\Delta) = C(\Gamma)$; however, only the pair (Γ_2, Δ) corresponds to an optimal transmission strategy. By analogy, an example can be given involving two different distortions. A concrete example of a system where the condition $R(\Delta) = C(\Gamma)$ is not sufficient is given as Example 1.1 above.

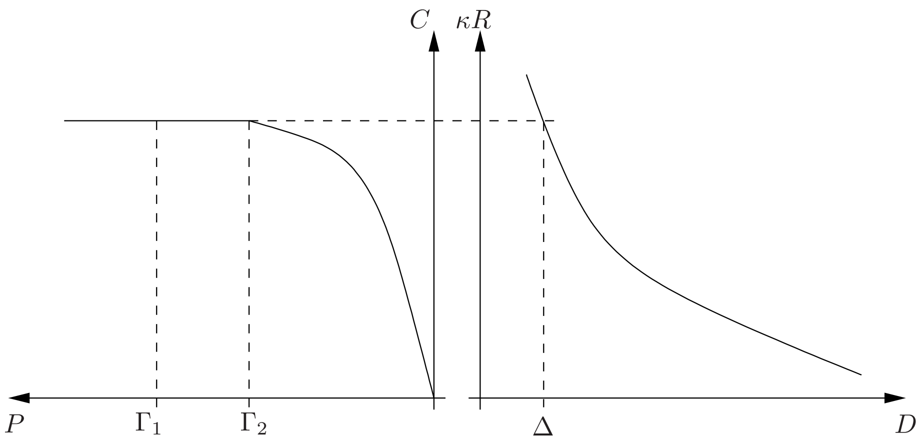

Figure 2.3: When $R(\Delta) = C(\Gamma)$ is not sufficient to guarantee optimality..

Continuing in this line of thought, we obtain the following proposition.

Proposition 2.4 *Suppose that the transmission of the discrete-time memoryless source* (p_S, d) *across the discrete-time memoryless channel* $(p_{Y|X}, \rho)$ *using the single-letter code* (f, g) *satisfies* $R(\Delta) = C(\Gamma)$. *Then,*

(i) Γ *cannot be lowered without changing* $C(\Gamma)$ *if and only if one of the following two conditions is satisfied:*

(a) $I(X;Y) < C_0$, *or*

(b) $I(X;Y) = C_0$ *and among the distributions that achieve* C_0, p_X *is one with lowest cost. In particular, the last condition is trivially satisfied whenever* p_X *is the unique channel input distribution achieving* C_0.

(ii) Δ *cannot be lowered without changing* $R(\Delta)$ *if and only if one of the following two conditions is satisfied:*

(a) $I(S;\hat{S}) > 0$, *or*

(b) $I(S;\hat{S}) = 0$ *and among the conditional distributions for which* $I(S;\hat{S}) = 0$, $p_{\hat{S}|S}$ *is one with lowest distortion. In particular, the last condition is trivially satisfied if* $p_{\hat{S}|S}$ *is the unique conditional distribution achieving* $I(S;\hat{S}) = 0$.

Proof. Part *(i)*: To see that condition *(a)* is sufficient, define $\Gamma_{max} = \min\{P : C(P) = C_0\}$. For every $\Gamma < \Gamma_{max}$, the value $C(\Gamma)$ uniquely specifies Γ. This follows from the fact that $C(\cdot)$ is convex and nondecreasing. From Lemma 2.1, $R(\Delta) = C(\Gamma)$ implies $C(\Gamma) = I(X;Y)$. Hence, $I(X;Y) < C_0$ implies $C(\Gamma) < C_0$, which in turn implies that it is not possible to change Γ without changing $C(\Gamma)$. To see that condition *(b)* is sufficient, note that if among the achievers of C_0, p_X belongs to the ones with lowest cost, then it is indeed impossible to lower Γ without changing $C(\Gamma)$. In particular, if p_X is the only achiever of C_0, then there cannot be another p_X that achieves the same rate, namely C_0, but with smaller cost, simply because there is no other p_X that achieves C_0.
It remains to show that if neither *(a)* nor *(b)* is satisfied, then Γ can indeed be lowered. In that case, $I(X;Y) = C_0$ (it cannot be larger than C_0). Moreover, there must be multiple achievers of C_0, and p_X is not one minimizing Γ. In other words, Γ can indeed be lowered without changing $C(\Gamma) = C_0$.
The proof of part *(ii)* of the proposition goes along the same lines. To see that condition *(a)* is sufficient, define $\Delta_{max} = \min\{D : R(D) = 0\}$. For every $\Delta < \Delta_{max}$, the value $R(\Delta)$ uniquely specifies Δ. This follows from the fact that $R(\cdot)$ is convex and non-increasing. From Lemma 2.1, $R(\Delta) = C(\Gamma)$ implies $R(\Delta) = I(S;\hat{S})$. Hence, $0 < I(S;\hat{S})$ implies $0 < R(\Delta)$, which in turn implies that it is not possible to change Δ without changing $R(\Delta)$. For condition *(b)*, note that if among the achievers of zero mutual information, $p_{\hat{S}|S}$ is one with lowest distortion, then it is indeed impossible to lower Δ without changing $R(\Delta)$. In particular, if $p_{\hat{S}|S}$ is the unique conditional distribution achieving zero mutual information, then there cannot be another conditional distribution achieving the same rate (zero) but with smaller distortion, simply because by assumption, there is no other conditional distribution achieving zero mutual information.
It remains to show that if neither *(a)* nor *(b)* is satisfied, then Δ can indeed be lowered. In that case, $I(S;\hat{S}) = 0$ (it cannot be smaller than 0). Moreover, there must be multiple achievers of zero mutual information, and $p_{\hat{S}|S}$ does not minimize the distortion among them. In other words, Δ can indeed be lowered without changing $R(\Delta) = 0$. □

Remark 2.2 *In the most general case of Proposition 2.4, it is necessary to specify the cost function and the distortion measure before the conditions can be verified. Let us point out, however, that, in many cases of practical interest, this is* not *necessary. In particular, if* $I(X;Y) < C_0$*, or if* $I(X;Y) = C_0$ *but* p_X *is the unique distribution that achieves* C_0*, then Part* (i) *is satisfied irrespective of the choice of the cost function. By analogy, if* $0 < I(S;\hat{S})$*, or if* $I(S;\hat{S}) = 0$ *but* $p_{\hat{S}|S}$ *is the unique conditional distribution for which* $I(S;\hat{S}) = 0$*, then Part* (ii) *is satisfied irrespective of the choice of the distortion measure.*

In summary, our discussion of Condition *(ii)* of Theorem 1.5 supplied a set of explicitly verifiable criteria. The main result of this chapter is obtained by combining this with the results of Section 2.2.1.

2.2.3 To code, or not to code?

The main result of this chapter is a simple criterion to check whether a given single-letter code performs optimally for a given source/channel pair. Theorem 1.5 showed that, on the one hand, the system has to satisfy $R(\Delta) = C(\Gamma)$. The choice of the cost function ρ as in Lemma 2.2 ensures that the channel input distribution achieves capacity. Similarly, the choice of the distortion measure according to Lemma 2.3 ensures that the conditional distribution of $\hat{S}$ given S achieves the rate-distortion function of the source. Together with the condition that $I(S;\hat{S}) = I(X;Y)$, this ensures that $R(\Delta) = C(\Gamma)$. But Theorem 1.5 required on the other hand that Γ cannot be lowered without changing $C(\Gamma)$, and that Δ cannot be lowered without changing $R(\Delta)$. Recall that this is *not* ensured by Lemmata 2.2 and 2.3. Rather, it was discussed in Section 2.2.2 and led to Proposition 2.4. It is now a simple matter to combine the insight gained in the latter proposition with the statements from Lemmata 2.2 and 2.3. This leads to a quite simple criterion to establish the optimality of a large class of communication systems that employ single-letter codes:

Theorem 2.5 *Consider a discrete memoryless source* (p_S, d) *and a discrete memoryless channel* $(p_{Y|X}, \rho)$ *for which there exists values of* P *and* D *such that* $R(D) = C(P)$*. For the transmission using a single-letter code* (f, g)*, the following statements hold:*

(o) If $I(S;\hat{S}) \neq I(X;Y)$*, then the system does not perform optimally.*

(i) If $0 < I(S;\hat{S}) = I(X;Y) < C_0$*, the system is optimal if and only if* $\rho(x)$ *and* $d(s,\hat{s})$ *satisfy Lemmata 2.2 and 2.3, respectively.*

(ii) If $0 < I(S;\hat{S}) = I(X;Y) = C_0$*, the system is optimal if and only if* $d(s,\hat{s})$ *satisfies Lemma 2.3, and* $\rho(x)$ *is such that* $E\rho(X) \leq E_{\tilde{p}_X}\rho(X)$ *for all other achievers* $\tilde{p}_X$ *of* C_0*. In particular, the last condition is trivially satisfied if* p_X *is the unique channel input distribution achieving* C_0*.*

(iii) If $0 = I(S;\hat{S}) = I(X;Y) < C_0$*, the system is optimal if and only if* $\rho(x)$ *satisfies Lemma 2.2, and* $d(s,\hat{s})$ *is such that* $Ed(S,\hat{S}) \leq E_{\tilde{p}_{\hat{S}|S}} d(S,\hat{S})$ *for all other achievers* $\tilde{p}_{\hat{S}|S}$ *of* $I(S;\hat{S}) = 0$*. In particular, the last condition is trivially satisfied if* $p_{\hat{S}|S}$ *is the unique conditional distribution for which* $I(S;\hat{S}) = 0$.

(iv) If $C_0 = 0$*, then the system is optimal if and only if* $E\rho(X) \leq E_{\tilde{p}_X}\rho(X)$ *for all channel input distributions* $\tilde{p}_X$*, and* $Ed(S,\hat{S}) \leq E_{\tilde{p}_{\hat{S}|S}} d(S,\hat{S})$ *for all conditional distributions* $\tilde{p}_{\hat{S}|S}$.

Proof. Part *(o)*. From the data processing inequality (see e.g. [7, Thm. 4.3.3] or [3, Thm. 2.8.1]), $I(S;\hat{S}) \neq I(X;Y)$ implies $I(S;\hat{S}) < I(X;Y)$. Moreover, $I(S;\hat{S}) < I(X;Y)$ implies $R(\Delta) < C(\Gamma)$ (see also the proof of Lemma 2.1). But then, by Theorem 1.5, the system does not perform optimally.
Part *(i)*. If $0 < I(S;\hat{S})$ and $I(X;Y) < C_0$, the system is optimal *if and only if* $R(\Delta) = C(\Gamma)$ (Theorem 1.5 with Proposition 2.4). We have shown that this is equivalent to requiring the three conditions of Lemma 2.1 to be satisfied. The third of these conditions, $I(S;\hat{S}) = I(X;Y)$, is satisfied by assumption. As long as $0 < I(S;\hat{S})$ and $I(X;Y) < C_0$, Lemmata 2.2 and 2.3 establish that the first two are satisfied *if and only if* ρ and d are chosen according to Formulae (2.14) and (2.15), respectively.
Part *(ii)*. If $I(X;Y) = C_0$, the system is optimal *if and only if* $R(\Delta) = C(\Gamma)$ and among the achievers of C_0, p_X belongs to the ones with lowest cost (Theorem 1.5 with Proposition 2.4). The condition $R(\Delta) = C(\Gamma)$ is satisfied *if and only if* the three conditions of Lemma 2.1 are satisfied. The third of these conditions, $I(S;\hat{S}) = I(X;Y)$, is satisfied by assumption. When $0 < I(S;\hat{S})$ but $I(X;Y) = C_0$, Lemmata 2.2 and 2.3 establish that the first two are satisfied *if and only if* d is chosen according to Formula (2.15).
Part *(iii)*. If $0 = I(S;\hat{S})$, the system optimal *if and only if* $R(\Delta) = C(\Gamma)$ and among the conditional distributions for which $I(S;\hat{S}) = 0$, $p_{\hat{S}|S}$ belongs to the ones with lowest distortion (Theorem 1.5 with Proposition 2.4). The condition $R(\Delta) = C(\Gamma)$ is satisfied *if and only if* the three conditions of Lemma 2.1 are satisfied. The third of these conditions, $I(S;\hat{S}) = I(X;Y)$, is satisfied by assumption. When $I(X;Y) < C_0$ but $I(S;\hat{S}) = 0$, Lemmata 2.2 and 2.3 establish that the first two are satisfied *if and only if* ρ is chosen according to Formula (2.14).
Part *(iv)* has been added for completeness only. It should be clear that if $C_0 = 0$, then automatically all the mutual information conditions are satisfied since all mutual informations must be zero. All that has to be checked is that the cost and the distortion are minimal. Obviously, this case is of limited interest. □

2.2.4 Extension to continuous alphabets

While Theorem 2.5 was proved for discrete alphabets only, some results of this chapter apply also to continuous alphabets. In this section, we establish one consequence of this fact.

More precisely, we now consider *discrete-time memoryless* sources and channels as defined in Definitions 1.1 and 1.2, respectively For those, we can establish the following theorem:

Theorem 2.6 *Consider a discrete-time memoryless source* (p_S, d) *and a discrete-time memoryless channel* $(p_{Y|X}, \rho)$ *for which there exist values of* P *and* D *such that* $R(D) = C(P)$*. Suppose that the single-letter source-channel code* (f, g) *is such that* $0 < I(S;\hat{S}) = I(X;Y) < C_0$*. If the source, the channel and the code satisfy*

$$\rho(x) \quad \begin{cases} = & c_1 D(p_{Y|X}(\cdot|x)||p_Y(\cdot)) + \rho_0 \quad \text{if } p(x) > 0 \\ \geq & c_1 D(p_{Y|X}(\cdot|x)||p_Y(\cdot)) + \rho_0 \quad \text{otherwise.} \end{cases} \tag{2.16}$$

$$d(s, \hat{s}) \quad = \quad -c_2 \log_2 p(s|\hat{s}) + d_0(s), \tag{2.17}$$

then they constitute an optimal communication system.

Proof. The proof of Lemma 2.2 shows that even for continuous alphabets, if $\rho(x)$ is chosen according to (2.16), then the underlying channel input density $p(x)$ achieves capacity. Moreover, the proof of Lemma 2.3 shows that even for continuous alphabets, if $d(s, \hat{s})$ is chosen according to (2.17), then the underlying distribution $p(\hat{s}|s)$ achieves the rate-distortion function. Furthermore, by assumption, the code (f, g) is such that $I(S;\hat{S}) = I(X;Y)$. Hence, by Lemma 2.1, the communication system satisfies $R(\Delta) = C(\Gamma)$. Hence, Condition (i) of Theorem 1.5, i.e., $R(\Delta) = C(\Gamma)$, is satisfied. Since by assumption, $0 < I(S;\hat{S}) = I(X;Y) < C_0$, Proposition 2.4 ensures that Condition (ii) is also satisfied, hence the communication system performs optimally. □

Note that in contrast to Theorem 2.5, Equations (2.16) and (2.17) of Theorem 2.6 are *not* necessary conditions — there may generally be other choices of ρ and d for which the system also performs optimally. A stronger statement could be made by constraining the considered class of probability distributions, and further technical conditions. Alternatively, the methods described in [28] may be helpful. This, however, is beyond the scope of this thesis.

2.3 Illustrations of Theorems 2.5 and 2.6

To illustrate Theorem 2.5, pick any probability measures for the source and the channel, and determine the cost function and distortion measure according to Lemmata 2.2 and 2.3, respectively. For the well-known example of a binary uniform source across a binary symmetric channel, this is done as follows.

Example 2.3 (binary) *Consider again Example 2.1. For that example, it turned out that uncoded transmission is an optimal source-channel communication strategy. Let us establish the same fact using Theorem 2.5. Since $I(X;Y) = C_0$, this example falls in case* (ii). *The distortion measure requirement is evaluated to be*

$$p(s|\hat{s}) = \frac{p_{Y|X}(\hat{s}|s)p_S(s)}{p_Y(\hat{s})} = p_{Y|X}(\hat{s}|s) = \begin{cases} 1-\epsilon, & \text{if } \hat{s} = s, \\ \epsilon, & \text{otherwise.} \end{cases} \quad (2.18)$$

Taking $d_0(s) = \frac{\log_2(1-\epsilon)}{\log_2 \frac{1-\epsilon}{\epsilon}}$ and $c_2 = \frac{1}{\log_2 \frac{1-\epsilon}{\epsilon}}$ in Lemma 2.3 reveals that one of the distortion measures that satisfy the requirement in Theorem 2.5 is the Hamming distance, confirming again that for the setup of Example 2.1, uncoded transmission is an optimal strategy.

As shown in this example, Theorem 2.5 can be applied directly by fixing the probability measures and the single-letter code, and determining the cost function and distortion measures according to the formulae. But the conditions of Theorem 2.5 are also useful if, e.g., the channel conditional probability distribution and the cost function are specified, and the source probability distribution and the distortion measure have to be determined accordingly, as illustrated by the following example [87]:

Example 2.4 *In this example, the alphabets are binary sequences of length n, denoted by bold symbols $\mathbf{x}$. Let the channel conditional distribution be any permutation of the length-n sequence, i.e.,*

$$p(\mathbf{y}|\mathbf{x}) = \begin{cases} \binom{n}{w(\mathbf{x})}^{-1}, & \text{if } w(\mathbf{x}) = w(\mathbf{y}), \\ 0, & \text{otherwise,} \end{cases} \quad (2.19)$$

where $w(\mathbf{x})$ denotes the Hamming weight (number of 1s) in the sequence $\mathbf{x}$. Moreover, let the cost function be

$$\rho(x) = a_1 w(\mathbf{x}) + a_0. \quad (2.20)$$

This can be seen as a simple model of neural communication [87]. By Lemma 2.2, the capacity-achieving input distribution satisfies

$$a_1 w(\mathbf{x}) + a_0 = c_1 D\left(p_{\mathbf{Y}|\mathbf{X}}(\cdot|\mathbf{x})||p_{\mathbf{Y}}(\mathbf{y})\right). \quad (2.21)$$

In [87], this condition is used to determine the capacity-achieving input distribution. The probability that $w(\mathbf{x}) = k$ is found to be

$$q(k) = \frac{b^n}{b^{n+1}-1}(b-1)b^{-k}, \quad (2.22)$$

for $k \in \{0, 1, \ldots, n\}$. In [87], the distortion measure according to Lemma 2.3 is also determined.

Example 2.4 considered a simple model of neural communication. This also illustrates the point that in certain applications, the source and the channel *can* be selected in a favorable fashion: For the case of neural communication, evolution had the opportunity to do so. This point is discussed in Chapter 4.

Beyond such a direct application, Theorem 2.5 is also useful in certain proofs. Example 2.3 suggests the question of the *uniqueness* of the solution: Suppose that all involved alphabets are binary, the distortion measure is Hamming, and the channel input cost function a constant. Then, is Example 2.3 the unique instance of optimal uncoded transmission? Using Theorem 2.5, one can establish the following Lemma:

Lemma 2.7 (binary) *Let $\mathcal{S} = \mathcal{X} = \mathcal{Y} = \hat{\mathcal{S}} = \{0,1\}$, $\rho(x) = const.$, and $d(s,\hat{s}) = 1$ if $s \neq \hat{s}$, and $d(s,\hat{s}) = 0$ otherwise (Hamming distortion). Suppose that the channel has nonzero capacity. Then, there exists a single-letter code with optimal performance if and only if the source probability mass function (pmf) p_S is uniform and the channel conditional pmf $p_{Y|X}$ is symmetric.*

Proof. The proof is given in Appendix 2.A.

If the alphabets are not binary, the following similar result can be established:

Lemma 2.8 (L-ary uniform) *Let $\mathcal{S}, \mathcal{X}, \mathcal{Y}$ and $\hat{\mathcal{S}}$ be L-ary, $\rho(x) = const.$, for all x, $d(s,\hat{s}) = 1$ if $s \neq \hat{s}$, and $d(s,\hat{s}) = 0$ otherwise (Hamming distortion), and p_S be uniform. Moreover, let the channel have nonzero capacity C_0. Then, there exists a single-letter code with optimal performance if and only if the channel conditional pmf is $p_{Y|X}(y|x) = const.$, for $y \neq x$ (or a permutation thereof).*

Proof. The proof is given in Appendix 2.A.

There is a nice intuition going along with the last result: Suppose that the channel is symmetric ([3, p. 190]) and that the probabilities of erroneous transition are $\{\epsilon_1, \ldots, \epsilon_{L-1}\}$ for every channel input. The distortion achieved by uncoded transmission is simply the sum of these probabilities. However, the distortion achieved by coded transmission depends on the capacity of the channel. Therefore, if uncoded transmission should have a chance to be optimal, we have to minimize the capacity of the channel subject to a fixed sum $\sum_{i=1}^{L-1} \epsilon_i$. But this is equivalent to maximizing the entropy of the "noise" $Z = Y - X$ subject to a fixed probability $p_Z(z = 0)$. Clearly, this maximum occurs when all the ϵ_i are equal.

As shown in Theorem 2.6, our results also apply to continuous alphabets, albeit in a weaker form. We illustrate this by the well-known Gaussian example, followed by a Laplacian example that was constructed from the theory developed in this chapter.

Example 2.5 (Gaussian) *Consider again Example 2.2. For that example, it is well known that (almost) uncoded transmission as described by Equations (2.11) and (2.12) is an optimal communication strategy. Let us establish the* *same fact using Theorem 2.6. It is clear that* $I(X;Y) = I(S;\hat{S})$ *since both the encoder (2.11) and the decoder (2.12) are bijective maps. Moreover, since* C_0 *is infinite for the AWGN channel, Theorem 2.6 can be applied. For the cost function, we first determine*

$$\begin{aligned} & D(p_{Y|X}(\cdot|x)||p_Y(\cdot)) \\ = & \; D\left(p_Z\left(\cdot - x\right) \;||\; p_Y(\cdot)\right) \\ = & \; -h\left(p_Z\left(\cdot - x\right)\right) - \int p_Z\left(y-x\right)\log_2 p_Y(y) dy. \end{aligned} \tag{2.23}$$

Since the entropy of a Gaussian is independent of its mean, the first term is a constant, say a. *Hence,*

$$\begin{aligned} & D(p_{Y|X}(\cdot|x)||p_Y(\cdot)) \\ & = a - \int \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(y-x)^2}{2\sigma^2}} \left(\log_2 \frac{1}{\sqrt{2\pi(\alpha^2\sigma_S^2 + \sigma^2)}} - \frac{y^2}{2(\alpha^2\sigma_S^2+\sigma^2)} \right) dy \\ & = a_1 + a_2 \int \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(y-x)^2}{2\sigma^2}} y^2 dy = a_1 + a_2(\sigma^2 + x^2) = b_1 x^2 + b_2, \end{aligned} \tag{2.24}$$

where the a_i *and* b_i *are appropriate constants. Since the formula of Theorem 2.6 only specifies* $\rho(x)$ *up to an affine transform, their precise value is irrelevant. For example, by choosing (in Theorem 2.6)* $c_1 = 1/b_1$ *and* $\rho_0 = -b_2/b_1$, *Eqn. (2.16) reads* $\rho(x) = x^2$.

For the distortion measure, we have to determine $p(s|\hat{s})$. *For notational convenience, denote the encoder given in Equation (2.11) by* $x = \alpha s$, *and the decoder given in Equation (2.12) by* $\hat{s} = \beta y$. *With this,* $p(\hat{s}|s)$ *can be expressed as*

$$p(\hat{s}|s) \;=\; \frac{1}{\beta} p_{Y|X}\left(\hat{s}/\beta \,|\, \alpha s\right) \;=\; \frac{1}{\sqrt{2\pi}\sigma\beta} e^{-\frac{1}{2\sigma^2\beta^2}(\hat{s}-\alpha\beta s)^2}. \tag{2.25}$$

The marginal of $\hat{S}$ *can be determined by recalling that* Y *is Gaussian with variance* $P + \sigma^2$. *Hence,* $\hat{S}$ *is Gaussian with variance* $\beta^2(P+\sigma^2)$. *Plugging in, we find*

$$\log_2 \frac{p(\hat{s}|s)}{p(\hat{s})} \;=\; \log_2 \frac{\beta\sqrt{P+\sigma^2}}{\sigma} e^{-\frac{1}{2\sigma^2\beta^2}(\hat{s}-\alpha\beta s)^2 + \frac{1}{2\beta^2(P+\sigma^2)}\hat{s}^2}, \tag{2.26}$$

which gives (by defining c_2 *and* $d_0(s)$ *in Theorem 2.6 appropriately)*

$$d(s,\hat{s}) \;=\; c_2 \left(\hat{s} - \frac{\alpha\beta(P+\sigma^2)}{P} s \right)^2 + d_0(s). \tag{2.27}$$

It is quickly verified that inserting the definitions of α *and* β *yields the standard mean-square error distortion. Hence, Theorem 2.6 allows to conclude that the suggested communications scheme performs optimally.*

As a side note, suppose that the coefficients α *and* β *are chosen differently, which means that the single-letter code is "mismatched." Then, the above derivation shows that the code performs optimally with respect to a "weighted" MSE distortion, with weighting as given by the last equation.*

By analogy to Lemma 2.7, it is tempting to claim that whenever the alphabets are continuous, the cost function is the square, and the distortion is the mean-square error, then the source and channel distributions must be Gaussian. It is beyond the scope of this thesis to establish this result, but note that certain other results point in this direction. For example, in [88], it is shown that for mean-square error distortion, the rate-distortion-achieving distribution is a probability *density* function if and only if the source is a mixture of Gaussians. Another clue comes from the Darmois-Skitovič Theorem (see e.g. [5]).

The last example of this section serves to illustrate that our results are not limited to binary and Gaussian examples.

Example 2.6 (Laplacian) *This example studies the transmission of a Laplacian source across an additive white Laplacian noise (AWLN) channel, defined as follows:*

$$p_S(s) = \frac{\alpha_0}{2} e^{-\alpha_0 |s|} \tag{2.28}$$

$$p_{Y|X}(y|x) = \frac{\alpha}{2} e^{-\alpha |y-x|}. \tag{2.29}$$

We also use $Z = Y - X$ *to denote the additive noise. Hence,* Z *is Laplacian with parameter* α*. Assume that* $\alpha_0 < \alpha$ *(which implies* $ES^2 > EZ^2$*). Note that with trivial changes, the derivations can be altered for the case* $\alpha_0 \geq \alpha$*. Moreover, let the encoding and the decoding function be simply the identity (in other words, consider uncoded transmission). The corresponding output distribution* $p_Y(y)$ *is found to be*

$$p_Y(y) = \frac{\alpha_0 \alpha}{2} \frac{\alpha e^{-\alpha_0 |y|} - \alpha_0 e^{-\alpha |y|}}{\alpha^2 - \alpha_0^2}. \tag{2.30}$$

Since the channel is an independent additive noise channel, the formula in Theorem 2.2 can be rewritten as

$$\rho(x) = -\int_z p_Z(z) \log_2 p_Y(x+z) dz. \tag{2.31}$$

A numerical approximation to this is illustrated in Fig. 2.4 for a particular choice of the parameters: $\alpha_0 = 3$ *and* $\alpha = 9$*, hence the signal-to-noise ratio in the example is* $\alpha^2/\alpha_0^2 = 9$*. Note that* $\rho(x)$ *as in Eqn. (2.31) is negative for some values of* x*. For the figure, we have added a suitable constant. The figure reveals that* $\rho(x)$ *is similar to the magnitude function (at least for our choice of the parameters). The next step is to compute the distortion measure that makes the system optimal. According to Theorem 2.3, we need to determine*

$$\begin{aligned} -\log_2 p(s|\hat{s}) &= -\log_2 \frac{p_{Y|X}(\hat{s}|s) p_S(s)}{p_Y(\hat{s})} \\ &= -\log_2 \frac{\alpha^2 - \alpha_0^2}{2} \frac{e^{-\alpha|\hat{s}-s| - \alpha_0 |s|}}{\alpha e^{-\alpha_0 |\hat{s}|} - \alpha_0 e^{-\alpha |\hat{s}|}}. \end{aligned} \tag{2.32}$$

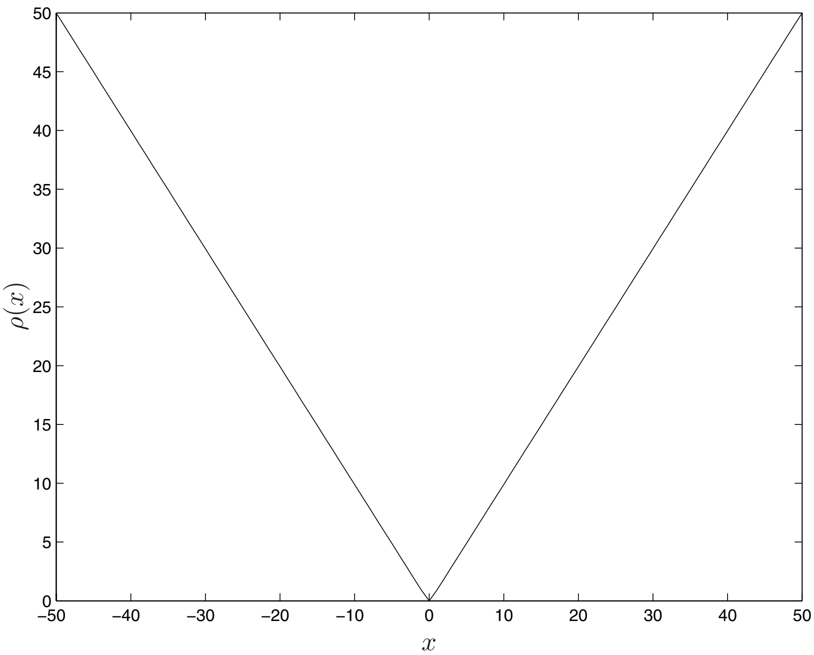

Figure 2.4: Channel input cost function $\rho(x)$ according to Eqn. (2.31).

However, this function is negative for some $(s, \hat{s})$*. To make it nonnegative, we add, for each* s*, an appropriate constant, namely the* $\log_2$ *of*

$$\begin{aligned}\max_{\hat{s}} p(s|\hat{s}) &= \frac{\alpha^2 - \alpha_0^2}{2} e^{-\alpha_0|s|} \max_{\hat{s}} \frac{e^{-\alpha|\hat{s}-s|}}{\alpha e^{-\alpha_0|\hat{s}|} - \alpha_0 e^{-\alpha|\hat{s}|}} \\ &= \frac{\alpha^2 - \alpha_0^2}{2} \frac{1}{\alpha - \alpha_0 e^{-(\alpha-\alpha_0)|s|}}.\end{aligned}$$

Substituting, we obtain

$$d(s, \hat{s}) = |\hat{s} - s| + \frac{1}{\alpha} \log_2 \frac{\alpha e^{-\alpha_0|\hat{s}|} - \alpha_0 e^{-\alpha|\hat{s}|}}{\alpha e^{-\alpha_0|s|} - \alpha_0 e^{-\alpha|s|}}. \tag{2.33}$$

This is illustrated in Fig. 2.5 for the above choice of the parameters ($\alpha_0 = 3$ *and* $\alpha = 9$*).*[4] *To conclude this example, let us point out that there is no straightforward answer to the question whether this distortion measure is practically meaningful. To judge that, physical objectives have to be taken into consideration.*

2.4 Summary and Conclusions

To code, or not to code: that is the question. Undoubtedly, “not to code” is very appealing when it leads to an optimal cost-distortion trade-off, since it involves

[4]Note that the figure does not exactly depict Eqn. (2.33); rather, additive and multiplicative constants have been selected to get a clearer picture.

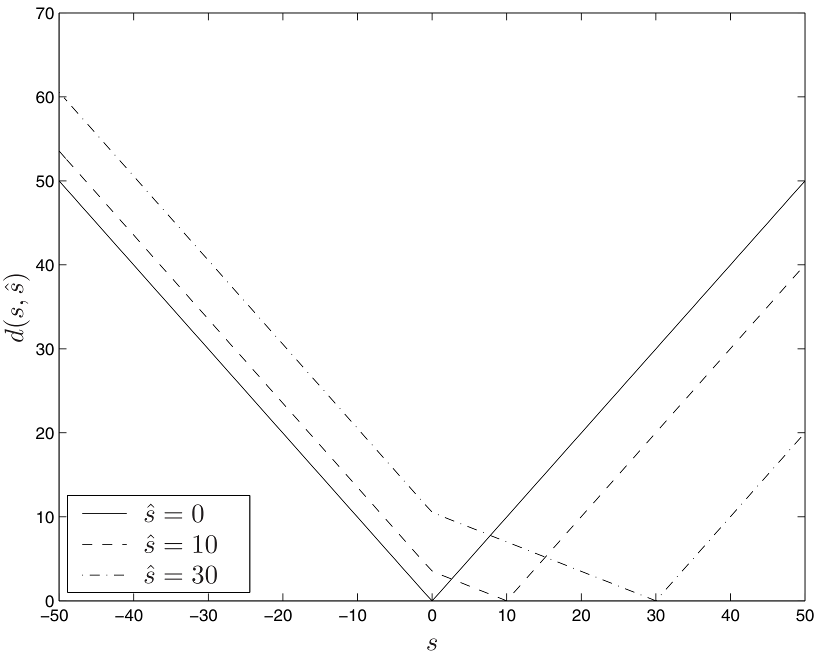

Figure 2.5: Distortion measure $d(s, \hat{s})$ according to Eqn. (2.33) for fixed $\hat{s}$.

the smallest possible delay and complexity. For uncoded transmission, optimality was shown to be a matter of matching four quantities: the source probability distribution and the corresponding distortion measure, and the channel conditional probability distribution and the corresponding cost function. Since these are all measures of some kind, we call our main result, Theorem 2.5, the *measure-matching* conditions.

One attractive feature of the conditions given in Theorem 2.5 is that they are explicit for fixed source and channel (conditional) distributions: We provide closed-form formulae for the channel input cost function ρ and the distortion measure d such that no loss in transmission quality is incurred with respect to the best possible communication scheme, irrespective of delay and complexity. More explicitly, it achieves the same cost-distortion trade-off as the best source compression followed by the best channel code. We also show that it is always possible to select ρ and d in such a way. In other words, there are infinitely many examples where uncoded transmission is an optimal communication strategy, just like in the examples of Section 2.1. We showed that in many cases, these formulae are also necessary conditions in the sense that if ρ and d are not chosen according to them, then the overall system performs suboptimally.

In the next chapter, we extend the measure-matching conditions of Theorem 2.5 to arbitrary source-channel codes. This leads to a different perspective on the general source-channel communication problem.

Appendix 2.A Proofs

Proof of Lemma 2.2. This lemma appears as Problem 2 in [4, p. 147], and its proof is a consequence of [4, Thm. 3.4]. In the following, we prove the sufficiency of the formula for $\rho(x)$ using a slightly different approach. Note that our proof also applies to continuous alphabets.

Let $p_{Y|X}$ be fixed. For any distribution p_X on $\mathcal{X}$, define

$$I'_{p_X}(x) \quad = \quad D(p_{Y|X}(\cdot|x)||p_Y), \tag{2.34}$$

where $p_Y(y) = Ep_{Y|X}(y|X)$ is the marginal distribution of Y when X is distributed according to p_X.

It is quickly verified that with this definition, $I_{p_X}(X;Y) = \langle p_X, I'_{p_X}\rangle$, where $\langle f,g\rangle$ denotes the standard inner product, i.e. for discrete alphabets, $\langle f,g\rangle = \sum_x f(x)g(x)$ and for continuous alphabets, $\langle f,g\rangle = \int f(x)g(x)dx$. With this notation, we may write $D(p_{Y|X}(\cdot|x)||p_Y) = \langle p_{Y|X}, \log_2 \frac{p_{Y|X}}{p_Y}\rangle_y$, where the subscript emphasizes that the inner product is taken in the variable y. The following auxiliary lemma is crucial for the proof:

Lemma: For any p_X and $\tilde{p}_X$, $I_{\tilde{p}_X}(X;Y) - I_{p_X}(X;Y) \ \le\ \langle \tilde{p}_X - p_X, I'_{p_X}\rangle$.

To see this, note first that since $I_{p_X}(X;Y) = \langle p_X, I'_{p_X}\rangle$, we equivalently prove the inequality $\langle \tilde{p}_X, I'_{p_X}\rangle - I_{\tilde{p}_X}(X;Y) \ \ge\ 0$, for any $p_X, \tilde{p}_X$.

$$\begin{aligned}
\langle \tilde{p}_X, I'_{p_X}\rangle - I_{\tilde{p}_X}(X;Y) &= \langle \tilde{p}_X, I'_{p_X}\rangle - \langle \tilde{p}_X, I'_{\tilde{p}_X}\rangle \\
&= \langle \tilde{p}_X, I'_{p_X} - I'_{\tilde{p}_X}\rangle \\
&= \langle \tilde{p}_X, D(p_{Y|X}||p_Y) - D(p_{Y|X}||\tilde{p}_Y)\rangle \\
&= \langle \tilde{p}_X, \langle p_{Y|X}, \log_2 \frac{\tilde{p}_Y}{p_Y}\rangle_y\rangle_x \\
&\overset{(a)}{=} \langle\langle \tilde{p}_X, p_{Y|X}\rangle_x, \log_2 \frac{\tilde{p}_Y}{p_Y}\rangle_y \\
&= \langle \tilde{p}_Y, \log_2 \frac{\tilde{p}_Y}{p_Y}\rangle_y \\
&= D(\tilde{p}_Y||p_Y) \ \ge\ 0,
\end{aligned} \tag{2.35}$$

where (a) is a change of summation (or integration) order and the inequality follows from the fact that the Kullback-Leibler distance is nonnegative. Lemma 2.2 can then be proved as follows.

($\Leftarrow$.) (Sufficiency of the formula.) Fix a distribution p_X over the channel input alphabet. Let ρ be arbitrary and let $\tilde{p}_X$ be any channel input distribution such that

$$E_{\tilde{p}_X}\rho(X) \quad \le \quad E_{p_X}\rho(X). \tag{2.36}$$

For any $\lambda \ge 0$,

$$\begin{aligned}
I_{p_X}(X;Y) - I_{\tilde{p}_X}(X;Y) &\ge \langle p_X - \tilde{p}_X, I'_{p_X}\rangle \\
&\ge \langle p_X - \tilde{p}_X, I'_{p_X} - \lambda\,\rho\rangle,
\end{aligned} \tag{2.37}$$

where the first inequality is the last lemma, and the second follows by assumption on $\tilde{p}_X$. If $\lambda\,\rho(x) = I'_{p_X}(x)+c$ for all x with $p(x) > 0$, then the last expression is zero, proving that $I_{p_X}(X;Y)$ indeed maximizes mutual information.

When $I_{p_X}(X;Y) = C_0$, then the input distribution p_X maximizes $I(X;Y)$ regardless of $\rho(x)$ and trivially fulfills the expected cost constraint. □

Proof of Lemma 2.3. This lemma appears as Problem 3 in [4, p. 147], and its proof is a consequence of [4, Thm. 3.7]. In the following, we prove the sufficiency of the formula for $d(s,\hat{s})$ using a slightly different approach. Note that our proof also applies to continuous alphabets.

To simplify the notation, we will use the symbol W in place of $p_{\hat{S}|S}$ in the proof. Define

$$I'_W(s,\hat{s}) \quad = \quad \log_2 \frac{W(\hat{s}|s)}{p_{\hat{S}}(\hat{s})}, \tag{2.38}$$

where $p_{\hat{S}}$ is the marginal distribution of $\hat{S}$.

In particular, note that with this definition, $I_W(S;\hat{S}) = \langle p_S W, I'_W\rangle$, where with slight abuse of notation, we have used $\langle p_S W, I'_W\rangle$ to mean $\int\int p_S(s)W(\hat{s}|s)I'_W(s,\hat{s})dsd\hat{s}$. In the proof, we use the following auxiliary lemma:

Lemma: For any W and $\tilde{W}$, $I_{\tilde{W}}(S;\hat{S}) - I_W(S;\hat{S}) \geq \langle p_S\tilde{W} - p_S W, I'_W\rangle$.

Using the fact that $I_W(S;\hat{S}) = \langle p_S W, I'_W\rangle$, we consider

$$\begin{aligned} I_{\tilde{W}}(S;\hat{S}) - \langle p_S\tilde{W}, I'_W\rangle &= \langle p_S\tilde{W}, \log_2 \frac{\tilde{W}}{\tilde{p}_{\hat{S}}}\rangle - \langle p_S\tilde{W}, \log_2 \frac{W}{p_{\hat{S}}}\rangle \\ &= \langle p_S\tilde{W}, \log_2 \frac{\tilde{V}}{p_S}\rangle - \langle p_S\tilde{W}, \log_2 \frac{V}{p_S}\rangle \\ &= \langle p_{\hat{S}}\tilde{V}, \log_2 \frac{\tilde{V}}{V}\rangle \;=\; \langle p_{\hat{S}}, D(\tilde{V}\|V)\rangle \\ &\geq 0, \end{aligned} \tag{2.39}$$

where we have used V to denote the conditional distribution of S given $\hat{S}$ under W, i.e. $V(s|\hat{s}) = W(\hat{s}|s)p(s)/p(\hat{s})$, and correspondingly $\tilde{V}$ to denote the same distribution, but under $\tilde{W}$, i.e. $\tilde{V}(s|\hat{s}) = \tilde{W}(\hat{s}|s)p(s)/\tilde{p}(\hat{s})$. $D(\tilde{V}\|V)$ denotes the Kullback-Leibler distance between $\tilde{V}$ and V in the variable s, hence it is a function of $\hat{s}$. The last inner product is thus one-dimensional in the variable $\hat{s}$. The inequality follows from the fact that the Kullback-Leibler distance is nonnegative.

With this, we are ready to prove Lemma 2.3.

($\Leftarrow$.) (Sufficiency of the formula.) Let d be arbitrary, let $\tilde{W}$ be an arbitrary conditional distribution such that

$$E_{p_S\tilde{W}}d(S,\hat{S}) \quad \leq \quad E_{p_S W}d(S,\hat{S}). \tag{2.40}$$

For any $\lambda > 0$,

$$\begin{aligned} I_{\tilde{W}}(S;\hat{S}) - I_W(S;\hat{S}) &\geq \langle p_S\tilde{W} - p_S W, I'_W(s,\hat{s})\rangle \\ &\geq \langle p_S\tilde{W} - p_S W, I'_W + \lambda d\rangle, \end{aligned} \tag{2.41}$$

where the first inequality is the last lemma, and the second follows by assumption on $\tilde{W}$. If $\lambda d(s,\hat{s}) = -I'_W(s,\hat{s}) + \tilde{d}_0(s)$ for all pairs $(s,\hat{s})$ with $p(s,\hat{s}) > 0$, then the last expression is zero, proving that $I_W(S;\hat{S})$ indeed minimizes mutual information. Setting $\tilde{d}_0(s) = -\log_2 p(s) + \lambda d_0(s)$ gives the claimed formula (2.15).

When $I_W(S;\hat{S}) = 0$, then trivially W achieves the minimum mutual information $I(S;\hat{S})$ over all $\tilde{W}$ that satisfy $E_{\tilde{W}} d(S,\hat{S}) \leq E_W d(S,\hat{S})$, regardless of d. □

Proof of Lemma 2.7. Assume that $X = S$ and $\hat{S} = Y$. This is without loss of generality, since the only two alternatives are *(i)* that the encoder permutes the source symbols, which is equivalent to swapping the channel transition probabilities (by the symmetry of the problem), and *(ii)* that the encoder maps both source symbols onto one channel input symbol, which is always suboptimal except when the channel has capacity zero. We will use the following notation: $\epsilon = p_{Y|X}(1|0)$, $\delta = p_{Y|X}(0|1)$, $p_X(x=0) = \bar{\pi}$ and $p_X(x=1) = \pi$. For the system to be optimal, since the channel is left unconstrained, it is necessary that $I(X;Y) = C_0$. Therefore, Case *(ii)* of Theorem 2.5 applies. Hence, it is *necessary* that $d(s,\hat{s})$ be chosen according to Eqn. (2.15); i.e., we require that $-\log_2 p(s|\hat{s}) = -\log_2 p(x|y)$ be equivalent to the Hamming distortion. This is the same as requiring that $p_{X|Y}(0|1) = p_{X|Y}(1|0)$. Expressing $p(x|y)$ as a function of $\epsilon, \delta, \bar{\pi}$ and π, the latter implies that $\pi = \sqrt{(\epsilon(1-\epsilon))/(\delta(1-\delta))}\bar{\pi}$. Since moreover, $\pi + \bar{\pi} = 1$, we find

$$\pi = \frac{1}{1 + \sqrt{(\delta(1-\delta))/(\epsilon(1-\epsilon))}}. \tag{2.42}$$

We show that for channel of nonzero capacity, this is the capacity-achieving distribution if and only if $\epsilon = \delta$, which completes the proof. The derivative

$$\frac{d}{d\pi} I(X;Y) = (\epsilon + \delta - 1)\log_2 \frac{1 - ((1-\pi)(1-\epsilon) + \pi\delta)}{(1-\pi)(1-\epsilon) + \pi\delta} + H_b(\epsilon) - H_b(\delta)$$

vanishes at the capacity-achieving input distribution. Plugging in π from above yields

$$2^{\frac{H_b(\delta) - H_b(\epsilon)}{1-\delta-\epsilon}} = \frac{(1-\epsilon)\sqrt{\delta(1-\delta)} + \delta\sqrt{\epsilon(1-\epsilon)}}{\epsilon\sqrt{\delta(1-\delta)} + (1-\delta)\sqrt{\epsilon(1-\epsilon)}}. \tag{2.43}$$

Clearly, equality holds if $\epsilon = \delta$ (and thus $\bar{\pi} = \pi$), but also if $\epsilon = 1 - \delta$. In the latter case, the channel has zero capacity. To see that there are no more values of ϵ and δ for which equality holds, fix (for instance) δ and consider the curves defined by the right side and the left side of Eqn. (2.43), respectively. The left side is convex and decreasing in ϵ. For $0 \leq \epsilon \leq 1 - \delta$, the right side is also convex and decreasing. Hence, at most 2 intersections can occur in this interval, and we already know them both. By continuing in this fashion, or by upper and lower bounds, one can establish that there are no more intersections. □

Proof of Lemma 2.8. Pick an arbitrary channel conditional distribution $p_{Y|X}$ for which there exists a single-letter code (f, g) that makes the overall system optimal. From Lemma 2.1, this implies that $I(X;Y) = C(\Gamma)$. Since the channel is unconstrained here, $C(\Gamma) = C_0$. Therefore, Case *(ii)* of Theorem 2.5 applies. That is, to perform optimally, the distortion measure must be chosen as a scaled and shifted version of $-\log_2 p(s|\hat{s})$. But since by assumption, the distortion measure must be the Hamming distance, we must have that $-\log_2 p(s|\hat{s}) = c_2(1 - \delta(s - \hat{s})) + d_0(s)$, where $\delta(\cdot)$ denotes the Kronecker delta function (i.e. it is one if the argument is zero, and zero otherwise). Equivalently, $p(s|\hat{s})$ must satisfy

$$p(s|\hat{s}) = \begin{cases} 2^{-d_0(s)}, & s = \hat{s}, \\ 2^{-c_2 - d_0(s)}, & s \neq \hat{s}. \end{cases} \tag{2.44}$$

The L simultaneous equations $\sum_s p(s|\hat{s}) = 1$ imply a full-rank linear system of equations in the variables $2^{-d_0(s)}$, from which it immediately follows that $d_0(s) = const$. But this means that $p(s|\hat{s})$ must satisfy

$$p(s|\hat{s}) = \begin{cases} \alpha, & s = \hat{s}, \\ \frac{1-\alpha}{L-1}, & s \neq \hat{s}. \end{cases} \tag{2.45}$$

By assumption, $p(s)$ is uniform, which implies that $p(\hat{s})$ is also uniform. But since all alphabets are of the same size, the condition that $I(S;\hat{S}) = I(X;Y)$ implies that $p(x)$ and $p(y)$ are also uniform, and that $p(x|y)$ is a permutation of

$$p(x|y) = \begin{cases} \alpha, & y = x, \\ \frac{1-\alpha}{L-1}, & y \neq x. \end{cases} \tag{2.46}$$

But this implies that the channel $p(y|x)$ has to be symmetric with $p(y|x) = \alpha$ for $y = x$, and $p(y|x) = (1 - \alpha)/(L - 1)$ for $y \neq x$, or a permutation thereof. □

Chapter 3

The Source-Channel Communication Problem, Part II

Source-channel communication is the problem of matching a source to a channel in an optimal fashion. One way to achieve an optimal match is described by the separation theorem (Theorem 1.5): If the rate of the source code is *matched* to the rate of the channel code, then the source-channel communication system performs optimally. This could be termed the *rate-matching* condition.

The purpose of the present chapter is to revisit optimal source-channel communication from the perspective of *matched measures*: In extension of the results of Chapter 2, we argue in Section 3.1 that an optimal source-channel communication system can also be characterized by the fact that the source and the channel probability distributions, the channel input cost function, and the distortion measure are matched in the optimal fashion. Since these are all measures of some kind, we call the resulting criterion for optimality the *measure-matching* condition.

While rate-matching makes operational sense only for systems with separate source and channel coding, the measure-matching condition makes sense for *arbitrary* "transducers" between the source output and the channel input, and between the channel output and the destination input, respectively.

The rest of the chapter illustrates and exploits the measure-matching conditions. In Section 3.2, we establish the fact that for a class of discrete memoryless source/channel pairs, the optimal match is obtained by a source-channel code either of block length one, or of block length infinity. Section 3.3 makes the connection between the separation theorem and measure-matching explicit. In Section 3.4, we introduce a notion of universality for source-channel codes, and we prove elementary facts about it. This notion is also of interest in our study of network source-channel communication (in particular, in Section 5.2). In Section 3.5, we establish a statement about optimal source-channel communication with feedback, and we give a new example of an optimal system that uses a very simple feedback code. Section 3.6 gives further connections to other results, most notably "bits through queues."[1]

3.1 Measure-matching

3.1.1 Block source-channel codes

In Section 2.2, we developed results for single-letter codes. It is clear that any source-channel code can be seen as a single-letter code in appropriately extended alphabets, at least as long as all alphabets are assumed to be discrete. For continuous alphabets, a similar extension is possible. Hence, the results of Section 2.2 can be applied directly to arbitrary source-channel codes, leading to a general criterion to establish the optimality of any source-channel communication system.

More precisely, suppose that a source-channel code (F, G) is used, with $F : \mathcal{S}^k \to \mathcal{X}^m$ and $G : \mathcal{Y}^m \to \hat{\mathcal{S}}^k$. To address this situation, k source symbols are merged into one new symbol, or symbol vector, denoted by $s^k \in \mathcal{S}^k$. The source distribution is given by

$$p(s^k) \quad = \quad p(s_1, s_2, \ldots, s_k), \tag{3.1}$$

and the source is reconstructed in the alphabet $\hat{\mathcal{S}}^k$, with respect to a distortion measure

$$d^{(k)}(s^k, \hat{s}^k) \quad = \quad d^{(k)}((s_1, s_2, \ldots, s_k), (\hat{s}_1, \hat{s}_2, \ldots \hat{s}_k)). \tag{3.2}$$

Hence, the source may have block memory of length k, but this memory is synchronized with the code (F, G): subsequent input blocks to the encoder are independent and identically distributed.[2] An interesting special case of such a source is the memoryless source. In that case, the source distribution and the distortion measure simplify to

$$p(s^k) \quad = \quad \prod_{j=1}^{k} p(s_j), \tag{3.3}$$

$$d^{(k)}(s^k, \hat{s}^k) \quad = \quad \sum_{i=1}^{k} d(s_i, \hat{s}_i). \tag{3.4}$$

The channel can also be seen in a vector perspective: It takes as inputs blocks of length m, namely the output blocks of the encoder, and its outputs are again blocks of length m. Hence, the conditional distribution is generally

$$p(y^m|x^m) \quad = \quad p(y_1, y_2, \ldots, y_m|x_1, x_2, \ldots, x_m), \tag{3.5}$$

[1]Sections 3.1, 3.2 and 3.4 appear in [43].

[2]This is less restrictive than it seems: suppose e.g. that the source has block memory of length k'; then, the theory can be applied to blocks of length kk'.

and the cost at the channel input is also defined on blocks of length m, i.e.,

$$\rho^{(m)}(x^m) = \rho^{(m)}(x_1, x_2, \ldots, x_m). \tag{3.6}$$

Hence, this model again naturally accommodates block memory, as long as it is synchronized with the encoder (see also the comment above). An interesting special case of such a channel is the memoryless channel. In that case, the channel conditional distribution and the input cost function simplify to

$$p(y^m|x^m) = \prod_{j=1}^{m} p(y_j|x_j), \tag{3.7}$$

$$\rho^{(m)}(x^m) = \sum_{i=1}^{m} \rho(x_i). \tag{3.8}$$

Clearly, the system considered in the extended alphabet (or vector) perspective is again a single-letter code. Lemma 2.1 applies unchanged; we quote it for future reference.

Lemma 3.1 *For a discrete-time memoryless* vector *source/channel pair* $(p_{S^k}, d^{(k)})$ *and* $(p_{Y^m|X^m}, \rho^{(m)})$ *and a single-letter (i.e., single-vector) code* (F, G), $R^{(k)}(\Delta) = C^{(m)}(\Gamma)$ *holds if and only if the following three conditions are simultaneously satisfied:*

(i) the distribution p_{X^m} *of* $X^m = F(S^k)$ *achieves capacity on the channel* $(p_{Y^k|X^k}, \rho^{(k)})$ *at input cost* $\Gamma = E\rho^{(m)}(X^m)$, *i.e.,* $I(X^m; Y^m) = C^{(m)}(\Gamma)$,

(ii) the conditional distribution $p_{\hat{S}^k|S^k}$ *of* $\hat{S}^k = G(Y^m)$ *given* S^k *achieves the rate-distortion function of the source* $(p_{S^k}, d^{(k)})$ *at distortion* $\Delta = Ed^{(k)}(S^k, \hat{S}^k)$, *i.e.* $I(S^k; \hat{S}^k) = R^{(k)}(\Delta)$, *and*

(iii) F *and* G *are such that* $I(S^k; \hat{S}^k) = I(X^m; Y^m)$, *i.e., they are "information lossless."*

Remark 3.1 *Note that in the above lemma,* $R^{(k)}(D)$ *denotes the rate-distortion function of the source that emits blocks of length* k. *Similarly,* $C^{(m)}(P)$ *denotes the capacity-cost function of the channel that operates with blocks of length* m, *and* $C_0^{(m)}$ *denotes the unconstrained capacity of that channel.*

Proof. The proof is the same as that of Lemma 2.1, except that all involved quantities are reformulated for the vector case. □

Remark 3.2 (memoryless source) *If the source is memoryless as in (3.3) and (3.4), then, condition* (ii) *of Lemma 3.1 can be expressed by the following two conditions:*

(a) $p(s^k|\hat{s}^k) = \prod_{i=1}^{k} p(s_i|\hat{s}_i)$, *and*

(b) the distributions $p(s_i|\hat{s}_i)$ all achieve the rate-distortion function of the source (p_S, d) at the same distortion $\Delta = Ed(S, \hat{S})$, i.e., $I(S; \hat{S}) = R(\Delta)$. Here, $R(\Delta)$ denotes the the rate-distortion function of the source (p_S, d).

This follows immediately from the proof of Theorem 1.1, in particular inequalities (a) *and* (b) *of the proof.*

Remark 3.3 (memoryless channel) *If the channel $(p_{Y^m|X^m}, \rho^{(m)})$ is memoryless as in (3.7) and (3.8), then, condition* (i) *of Lemma 3.1 can be expressed by the following two conditions:*

(a) $Y_1, Y_2, \ldots$ are independent random variables, and

(b) the marginals of $X_1, X_2, \ldots$ all achieve capacity on the channel $(p_{Y|X}, \rho)$ at the same input cost $\Gamma = E\rho(X)$, i.e., $I(X; Y) = C(\Gamma)$. Here, $C(\Gamma)$ denotes the capacity-cost function of the channel $(p_{Y|X}, \rho)$

This follows immediately from the proof of Theorem 1.1, in particular inequalities (d) *and* (e) *of the proof, and is also a special case of Lemma 3.8, to be proved below.*

3.1.2 Discrete alphabets

By analogy to Chapter 2, the cases of discrete and continuous alphabets are again treated separately. For discrete alphabets, the situation is particularly simple since the extension source and channel alphabets are still discrete alphabets. Hence, Theorem 2.5 can be applied directly to the extension alphabet, which leads to the following statement.

Corollary 3.2 *For a discrete memoryless source $(p_{S^k}, d^{(k)})$, and a discrete memoryless channel $(p_{Y^m|X^m}, \rho^{(m)})$, suppose that there exist values of P and D such that $R(D) = C(P)$. Consider transmission using a single-letter source-channel code (F, G) with $F : \mathcal{S}^k \to \mathcal{X}^m$ and $G : \mathcal{Y}^m \to \hat{\mathcal{S}}^k$, and suppose that $0 < I(S^k; \hat{S}^k) = I(X^m; Y^m) < C_0$. This is optimal if and only if*

$$\rho^{(m)}(x^m) \quad \begin{cases} = & c_1 D(p_{Y^m|X^m}(\cdot|x^m)||p_{Y^m}(\cdot)) + \rho_0 \quad \text{if } p(x^m) > 0 \\ \geq & c_1 D(p_{Y^m|X^m}(\cdot|x)||p_{Y^m}(\cdot)) + \rho_0 \quad \text{otherwise.} \end{cases} \tag{3.9}$$

$$d^{(k)}(s^k, \hat{s}^k) \quad = \quad -c_2 \log_2 p(s^k|\hat{s}^k) + d_0(s^k). \tag{3.10}$$

Proof. This corollary is Theorem 2.5, Part *(i)*, applied to suitably extended alphabets. □

Corollary 3.2 makes the concept of *measure-matching* precise: Optimality is a matter of matching up six functions, namely the encoding and the decoding function with the source and channel parameters. The latter are, in the case of the source, the probability distribution and the distortion measure, and in the case of the channel, the conditional probability distribution and the channel input cost function. Since these are all measures of some kind, and since this term seems to be their "greatest common divisor" (or their least common multiple?), we call the resulting concept *measure-matching*. In one direction, the measure-matching conditions stated in Corollary 3.2 are explicit: For fixed source and channel distributions and encoding/decoding functions, the condition specifies in closed form the shape of the channel input cost function and the distortion measure. We will exploit this feature of Corollary 3.2 in the sequel in various ways.

It is clear that longer codes generally permit one to better match the source and the channel. Corollary 3.2 can therefore also be interpreted as follows: Suppose a certain finite complexity is available to implement a source-channel communication system. Following Theorem 1.5, one would design (suboptimal) source and channel codes independently. The advantage of additional complexity appears as a lower error probability on the channel and a smaller size of the quantization cells for the source. In contrast to this, Corollary 3.2 suggests a very different perspective: additional coding complexity (in the shape of longer codes) is used to better match $\rho^{(m)}$ and $d^{(k)}$ to the desired cost and distortion measures.

3.1.3 Continuous alphabets

For continuous alphabets, a statement similar to Corollary 3.2 can be obtained. However, in line with Theorem 2.6, the formulae for ρ and d are merely sufficient, rather than also necessary, conditions. The extension of Theorem 2.6 to vector sources and channels can be phrased as follows.

Theorem 3.3 *Consider a discrete-time memoryless* vector *source/channel pair* $(p_{S^k}, d^{(k)})$ *and* $(p_{Y^m|X^m}, \rho^{(m)})$ *for which* $R^{(k)}(D) = C^{(m)}(P)$ *is feasible. Suppose that the single-vector source-channel code* (F, G) *is such that* $0 < I(S^k; \hat{S}^k) = I(X^m; Y^m) < C_0^{(m)}$. *If the source, the channel and the code satisfy*

$$\rho^{(m)}(x^m) \quad \begin{cases} = & c_1 D(p_{Y^m|X^m}(\cdot|x^m)||p_{Y^m}(\cdot)) + \rho_0 \quad \text{if } p(x^m) > 0 \\ \geq & c_1 D(p_{Y^m|X^m}(\cdot|x)||p_{Y^m}(\cdot)) + \rho_0 \quad \text{otherwise.} \end{cases} \tag{3.11}$$

$$d^{(k)}(s^k, \hat{s}^k) \quad = \quad -c_2 \log_2 p(s^k|\hat{s}^k) + d_0(s^k), \tag{3.12}$$

then they constitute an optimal communication system.

Proof. The proof of Lemma 2.2 can be extended to show that if $\rho^{(k)}$ is chosen according to Equation (3.11), then the corresponding channel input distribution $p(x^k)$ achieves capacity. In particular, the chain of equalities (2.35) applies unchanged if the notation $\langle\cdot\rangle$ is taken to mean

$$\langle f, g\rangle \quad = \quad \int_{x_1} \cdots \int_{x_k} f(x_1, \ldots, x_k) g(x_1, \ldots, x_k) dx_1 \cdots dx_k. \tag{3.13}$$

Lemma 2.3 can be extended by the same token. By Lemma 3.1, these two facts, together with the hypothesis that $I(S^k;\hat{S}^k) = I(X^m;Y^m)$, establish that the system satisfies $\kappa R(\Delta) = C(\Gamma)$. To conclude that the system performs optimally requires the appropriate extension of Proposition 2.4. For the case $0 < I(S^k;\hat{S}^k)$ and $I(X^m;Y^m) < C_0$, this is straightforward. Details are omitted. □

While this short argument shows how to handle simple instances of the problem of continuous sources and channels with memory, a number of subtleties have been swept under the rug. Most importantly, as the length of the blocks under consideration tends to infinity (e.g., for a simple first-order Markov source), issues of convergence have to be resolved in detail. This is left as future work. For the particular case where the source is Gaussian, the channel is also Gaussian (both with memory), and uncoded transmission is used, a condition that is both necessary and sufficient (and in this sense stronger than Theorem 3.3) has been recently presented in [56].

Remark 3.4 (MIMO systems) *The setup of Theorem 3.3 is also a model for multi-input multi-output (MIMO) source-channel communication. In particular, Theorem 3.3 considers a vector source (of length k) that is transmitted across a vector channel, also of length k, using a code (F, G) that maps each source output vector to a channel input vector, and each channel output vector to a source reconstruction vector. The theorem says that if ρ and d satisfy the stated conditions, then the MIMO source-channel communication system performs optimally.*

3.2 Source-Channel Codes of Finite Block Length

While longer codes permit us to *better* match the source (p_S, d) to the channel $(p_{Y|X}, \rho)$, we would also like to know what code length is necessary to obtain the *optimal* match. We again restrict our attention to discrete memoryless sources and channels as defined in Definitions 1.1 and 1.2, but the code is now an arbitrary source-channel code of (finite) length M. For simplicity, we consider only codes of rate $\kappa = 1$. Corollary 3.2 gives the cost function and the distortion measure on length-M blocks that are necessary for optimal performance. However, the underlying source and channel are *memoryless*. Therefore, by definition, it must be possible to express the cost function on length-M blocks as a sum of M individual terms, and the same must be true for the distortion measure. This excludes certain M-letter codes. Our conjecture is that a finite-length code with optimal performance exists if and only if there exists also a single-letter code with optimal performance for the same source/channel pair. We can prove this conjecture under some additional assumptions:

Theorem 3.4 *Let (p_S, d) and $(p_{Y|X}, \rho)$ be a discrete memoryless source and a discrete memoryless channel, respectively. Suppose that all alphabets are of* *the same size, that* $p(s) > 0$ *for all* $s \in \mathcal{S}$*, that the distortion measure has the property that the matrix* $\{2^{-d(s,\hat{s})}\}_{s,\hat{s}}$ *is invertible and that the channel transition probability matrix is invertible. Then, there exists a source-channel code of finite block length that performs optimally if and only if, for the same source/channel pair, there exists also a single-letter source-channel code that performs optimally.*

Proof. The proof of this theorem is given in Appendix 3.A.

Among the restrictions imposed by the last theorem, the one on the distortion measure may seem somewhat unusual. Note however that the standard distortion measures like the Hamming distance and the squared-error distortion satisfy that restriction. In fact, any distortion measure under which the mapping $T(s) = \arg\min_{\hat{s}} d(s, \hat{s})$ is one-to-one satisfies the requirement.

3.3 Successive Measure-matching

The separation theorem furnishes one way of satisfying the measure-matching condition (Corollary 3.2), at least asymptotically in the block length of the code. It does so in two steps, source coding and channel coding. The source coding maps the source outputs into an intermediate space, and the channel coding maps from this intermediate space into the optimal channel inputs. The separation theorem shows that this intermediate space is the *same* for *all* source/channel pairs, namely bits. We will refer to this intermediate space more precisely as the *(asymptotically) perfect reconstruction space*.[3] The source output sequence is described by one out of M different messages, and the channel code ensures that the index of the selected message can be communicated to the decoder reliably. This implies one of the conceptually and practically most attractive features of the separation theorem: It provides a modularization. For example, on the source side, an optimal source code of rate R can be optimally used for transmission across *any* channel of capacity R, irrespective of the precise channel structure. The disadvantage of this modularization is the requirement for the channel to transmit *reliably*; this implies infinite delays in many cases.

In contrast to this, the source-channel codes considered in Section 3.1 are optimal *for one special source and one special channel*, but they are of unbeatably low complexity.

Clearly, it would be interesting to identify solutions that lie between these two extremes: to find intermediate spaces

- that provide a limited modularization in the sense that only sources and channels out of a suitably defined class are admissible,
- but that, in turn, do not require infinite delay.

This motivates to study more closely coding systems whose encoder and decoder are split into two steps. We call this *successive measure-matching*.

[3] or: asymptotically *deterministic* reconstruction space, see also the discussion in Section 1.3.

3.3.1 Successive measure-matching

Consider the system shown in Figure 3.1. A source (p_S, d) is first encoded into an auxiliary random variable $S' = F'(S)$. Thereafter, S' is mapped into the channel input alphabet by $X = F''(S')$. Clearly, we can put the original source and the first encoder F' into a black box which serves as our new source S', and we can discuss the optimality of the communication system whose source is S'. By the same token, we can put the original channel together with the encoder F'' into a black box which serves as our new channel $p_{Y|X'}$. This explains why in Figure 3.1, we denote the same signal by S' (in the source perspective) and by X' (in the channel perspective).

The goal of successive measure-matching is to associate the optimality of subsystems, obtained by applying black boxes as discussed above, with the optimality of the overall system.

Figure 3.1: Successive measure-matching.

For the following theorem, recall that "optimal communication system" denotes a communication system according to Definition 1.5. Moreover, the symbol I denotes the identity function, i.e., straight uncoded transmission.

Theorem 3.5 (successive measure-matching) *In Figure 3.1, suppose that (p_S, d) and $(p_{S'}, d')$ are discrete-time memoryless sources, and $(p_{Y'|X'}, \rho')$ and $(p_{Y|X}, \rho)$ are discrete-time memoryless channels. Suppose that for each of the following communication systems, there exist values of P and D such that $\kappa' R(D) = C(P)$, for the corresponding values of κ'. Then, the following holds:*

The two communication systems

$$(i) \qquad (p_S, d), (p_{Y'|X'}, \rho'), (F', G') \tag{3.14}$$

$$(ii) \qquad (p_{S'}, d'), (p_{Y|X}, \rho), (F'', G'') \tag{3.15}$$

are both optimal if and only if the two communication systems

$$(i) \qquad (p_S, d), (p_{Y|X}, \rho), (F'' \circ F', G' \circ G'') \tag{3.16}$$

$$(ii) \qquad (p_{S'}, d'), (p_{Y'|X'}, \rho'), (I, I) \tag{3.17}$$

are also both optimal.

Proof. A communication system for which there exist values of P and D such that $\kappa' R(D) = C(P)$, for the corresponding values of κ', is optimal if and only if it satisfies Theorem 1.5.

($\Rightarrow$.) The optimality of (3.14) implies that $p_{\hat{S}|S}$ achieves the rate-distortion function, and that Δ cannot be lowered without changing $R(\Delta)$. The optimality of (3.15) implies that p_X achieves the capacity-cost function, and that Γ

cannot be lowered without changing $C(\Gamma)$. The optimality of (3.14) and (3.15) together implies that $I(S^k; \hat{S}^k) = I(X^m; Y^m)$. Using Lemma 3.1, we find that the communication system (3.16) satisfies $\kappa R(\Delta) = C(\Gamma)$, and using Theorem 1.5, we conclude that (3.16) is optimal. The optimality of (3.17) is established along the same lines.

($\Leftarrow$.) The proof proceeds along the same lines as the proof of ($\Rightarrow$). □

This theorem permits to clarify the term "intermediate space" that was used in the discussion above: The intermediate space is characterized by the intermediate communication system $(p_{S'}, d'), (p_{Y'|X'}, \rho'), (I, I)$. More precisely, to specify an intermediate space, the following have to be determined:

(i) The alphabet $\mathcal{S}' = \mathcal{X}'$, and the alphabet $\mathcal{Y}' = \hat{\mathcal{S}}'$.

(ii) The source $(p_{S'}, d')$.

(iii) The channel $(p_{Y'|X'}, \rho')$.

Our next goal is to illustrate how Theorem 3.5 relates to the separation theorem. The key step is the specification of the intermediate space, determined by $(p_{S'}, d')$ and $(p_{Y'|X'}, \rho')$, for the case of the separation theorem.

3.3.2 Connection to the separation theorem

The measure-matching conditions (as stated, e.g., in Corollary 3.2) were derived essentially from the separation theorem. A natural question to ask is therefore whether the separation theorem, in turn, can be derived from the measure-matching conditions. In this section, we provide a part of this derivation.

In particular, the separation theorem contains a statement of the form of Theorem 3.5, for a special intermediate space $(p_{S'}, d')$ and $(p_{Y'|X'}, \rho')$. In this section, we determine these four entities. As explained briefly at the beginning of this section, these must be expected to be quite singular: the intermediate space in the case of the separation theorem is the perfect reconstruction space: the error probability must tend to *zero* asymptotically. This implies in particular that the intermediate distortion measures d' and the intermediate conditional distribution $p_{Y'|X'}$, are singular objects. We define the following special source and special channel:

Definition 3.1 (M-ary asymptotically perfect source) *The M-ary asymptotically perfect source has an M-ary source alphabet and is reconstructed in the same alphabet. Its probability distribution is the M-ary uniform distribution, denoted by*

$$u^{(M)}. \tag{3.18}$$

The distortion measure is defined on blocks of n source symbols, denoted by s^n, and blocks of n reconstruction symbols, denoted by $\hat{s}^n$. It has the following form:

$$d_u(s^n, \hat{s}^n) \quad = \quad \begin{cases} 0, & \text{if } s^n = \hat{s}^n \\ d(n), & \text{otherwise.} \end{cases} \tag{3.19}$$

where $d(n)$ *is an increasing positive function with* $\lim_{n\to\infty} d(n) = \infty$.

Definition 3.2 (M'-ary asymptotically perfect channel) *The M'-ary asymptotically perfect channel has an M'-ary input alphabet and the same M'-ary output alphabet. Its conditional probability distribution is defined on blocks of n input symbols, and blocks of n output symbols, as follows:*

$$w^{(M')}(y^n|x^n) = \begin{cases} 1 - e(n), & \text{if } y^n = x^n \\ \frac{e(n)}{n^{M'}-1}, & \text{otherwise,} \end{cases} \tag{3.20}$$

where $e(n) < 1$ *is a decreasing function with* $\lim_{n\to\infty} e(n) = 0$. *Its cost function is uniform, i.e.,*

$$\rho_w(x) = 0, \forall x \in \mathcal{X}. \tag{3.21}$$

Using this in Theorem 3.5, we obtain the following immediate corollary:

Corollary 3.6 (separation theorem, achievability) *In Figure 3.1, suppose that (p_S, d) and $(p_{S'}, d')$ are discrete-time memoryless sources, and $(p_{Y'|X'}, \rho')$ and $(p_{Y|X}, \rho)$ are discrete-time memoryless channels. Suppose that for each of the following communication systems, there exist values of P and D such that $\kappa' R(D) = C(P)$, for the corresponding values of κ'. Then, the following holds:*
The two communication systems

$$(i) \qquad (p_S, d), (w^{(M)}, \rho_w), (F', G'), \tag{3.22}$$

$$(ii) \qquad (u^{(M)}, d_u), (p_{Y|X}, \rho), (F'', G''), \tag{3.23}$$

are both optimal if and only if the two communication systems

$$(i) \qquad (p_S, d), (p_{Y|X}, \rho), (F'' \circ F', G' \circ G'') \tag{3.24}$$

$$(ii) \qquad (u^{(M)}, d_u), (w^{(M)}, \rho_w), (I, I) \tag{3.25}$$

are also both optimal.

This is illustrated in Figure 3.2.

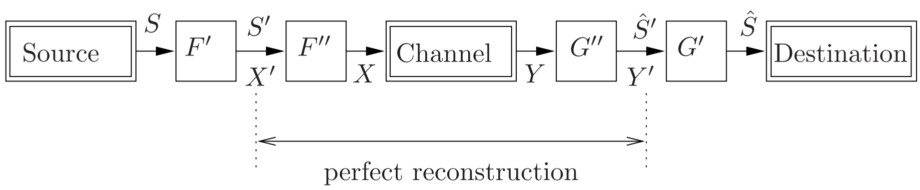

Figure 3.2: Illustration of the separation theorem as successive measure-matching

The separation theorem establishes moreover that for an arbitrary source/channel pair (p_S, d) and $(p_{Y|X}, \rho)$ and an arbitrary M (within certain bounds), there exists always a pair of codes, (F', G') and (F'', G''), that satisfies the conditions of Corollary 3.6. It would be interesting to prove this stronger statement using the arguments developed in Chapters 2 and 3. This is left as future work.

In summary, from the perspective of successive measure-matching, the separation theorem is the statement that *any* match can be achieved (asymptotically) by a two-step matching operation that passes through the perfect reconstruction space.

3.3.3 Weak separation theorems

As explained at the beginning of Section 3.3, the challenge of successive measure-matching is to find other "separation theorems" that do not pass through perfect reconstruction. They may not provide true universality in the sense that they do not apply to *all* sources and channels, but rather to specific *classes*. Such separation theorem may be called "weak" separation theorems.

While a comprehensive theory of weak separation theorems is left for future research, we can provide at this point one very simple example that illustrates the point.

Example 3.1 (Gaussian separation theorem) *Using the notation of Figure 3.1, consider the intermediate space characterized by the source*

$$(\mathcal{N}_1, (s' - \hat{s}')^2) \tag{3.26}$$

where $\mathcal{N}_1$ *is the iid Gaussian source of unit variance, and the channel*

$$(\mathcal{N}_{x',1}, {x'}^2), \tag{3.27}$$

where $\mathcal{N}_{x',1}$ *denotes the additive white Gaussian noise channel of unit noise variance.*

Then, we can give the following weak separation theorem: Suppose that any iid Gaussian source (p_S, d) *is to be transmitted across any AWGN channel* $(p_{Y|X}, \rho)$*. An optimal coding strategy can be decomposed into two stages: The first (outer) stage is given by*

$$S' = F'(S) \quad = \quad \frac{1}{\sqrt{\sigma_S^2}} S, \tag{3.28}$$

and

$$\hat{S} = G'(\hat{S}') \quad = \quad \sqrt{\sigma_S^2}\hat{S}'. \tag{3.29}$$

It is quickly verified that (e.g., using Example 2.2)

$$(p_S, d), (\mathcal{N}(x', 1), {x'}^2), (F', G') \tag{3.30}$$

is an optimal communication system.

The second (inner) stage is given by

$$X = F''(X') \quad = \quad \sqrt{P}X', \tag{3.31}$$

and

$$Y' = G''(Y) \quad = \quad \frac{1}{\sqrt{P}} \frac{P}{P+\sigma^2} Y. \tag{3.32}$$

It is quickly verified that (e.g., using Example 2.2)

$$(\mathcal{N}_1, (s' - \hat{s}')^2), (p_{Y|X}, \rho), (F'', G'') \tag{3.33}$$

is an optimal communication system.

Hence, by Theorem 3.5, we conclude that

$$(p_S, d), (p_{Y|X}, \rho), (F'' \circ F', G' \circ G'') \tag{3.34}$$

is an optimal communication system, confirming again Example 2.2.

In this simple example, a (weak) separation theorem applies, but the intermediate space is not the perfect reconstruction space, and infinite delays are not necessary. We hope in the future to extend this example using the results of Chapters 2 and 3.

3.4 The Universality of Source-Channel Codes

Optimal communication systems designed according to the separation principle may be quite sensitive to parameter mismatch. Suppose e.g. that the capacity of the channel turns out to be smaller than the rate of the channel code that is used. The effect of this parameter mismatch on the final reconstruction of the data may be catastrophic.

Source-channel codes may feature a graceful degradation as a function of mismatched parameters. In fact, in some cases, one and the same source-channel code achieves *optimal* performance for *multiple* source/channel pairs. In this sense, certain source-channel codes have a universality property. The following example illustrates this.

Example 3.2 (fading) *Let the source be the binary uniform source as in Example 2.3. The channel is slightly different from Example 2.3: the transition probability ϵ varies in a memoryless fashion during transmission. Take the encoder and the decoder to be identity mappings (i.e., uncoded transmission). From Example 2.3, it is clear that this code performs optimally irrespective of the value of ϵ.*

In this example, the suggested code is universal for the transmission of a binary uniform source across any one out of an entire class of channels. In the spirit of this example, we introduce the following definition:

Definition 3.3 (universality) *The source-channel code (F, G) is called universal for the source (p_S, d) and the class of channels given by $\mathcal{W} = \{(p^{(0)}_{Y|X}, \rho^{(0)}), (p^{(1)}_{Y|X}, \rho^{(1)}), \ldots\}$ if, for all i, the transmission of the source (p_S, d) across the channel $(p^{(i)}_{Y|X}, \rho^{(i)})$ using the code (F, G) is optimal.*

Note that by complete analogy, one can define the universality of a code with respect to a *class* of sources and a class of channels. In order to keep notation simple, we leave this as an exercise to the reader. Instances of universality can be characterized by direct application of Theorem 2.5 to the present scenario. For example, for single-letter codes, Theorem 2.5, Part *(i)*, provides the following corollary:

Corollary 3.7 *Consider a source* (p_S, d) *and a class of channels* $\mathcal{W}$ *such that for every channel in* $\mathcal{W}$*, there exist values of* P_i *and* D_i *such that* $R(D_i) = C_i(P_i)$*. Suppose that for the single-letter code* (f, g)*, it is true that* $0 < I(S; \hat{S}^{(i)}) = I(X; Y^{(i)}) < C_0^{(i)}$ *for all* i*. The single-letter code* (f, g) *is universal if and only if for all* i*,*

$$\rho^{(i)}(x) \quad \begin{cases} = & c_1^{(i)} D(p_{Y|X}^{(i)}(\cdot|x)||p_Y(\cdot)) + \rho_0^{(i)} \quad \text{if } p(x) > 0 \\ \geq & c_1^{(i)} D(p_{Y|X}^{(i)}(\cdot|x)||p_Y(\cdot)) + \rho_0^{(i)} \quad \text{otherwise.} \end{cases} \tag{3.35}$$

$$d(s, \hat{s}) \quad = \quad -c_2^{(i)} \log_2 p^{(i)}(s|\hat{s}) + d_0^{(i)}(s), \tag{3.36}$$

where $c_1^{(i)} > 0$*,* $c_2^{(i)} > 0$ *and* $\rho_0^{(i)}$ *are constants, and* $d_0^{(i)}(s)$ *is an arbitrary function.*

Proof. Follows directly from Theorem 2.5. □

By analogy, one can again include all the special cases of Theorem 2.5. This is left to the reader. The main reason for studying this particular property of source-channel codes lies in its practical implications. One implication is to time-varying (fading) channels, as illustrated by the above example: The channel varies over time, but it always remains inside the class $\mathcal{W}$. For that case, it is immediate that a universal source-channel code achieves the performance of the best source compression followed by the best channel code. However, the significance of universal source-channel codes extends beyond the validity of the separation theorem. To end this short discussion, we mention two key scenarios under which source-channel codes outperform any code designed according to the separation paradigm.

The first scenario concerns communication under channel uncertainty.

Implication 1 (non-ergodic channels). Let the source-channel code (F, G) be universal for the source (p_S, d) and the class of channels $\mathcal{W}$. Let the channel be in $\mathcal{W}$, but not determined at the time of code design. Then, transmission using the source-channel code (F, G) achieves optimal performance, regardless of which particular channel is selected.

Capacity issues under channel uncertainty have been studied in detail, see e.g. [71]. As the above shows, joint source-channel coding may be another valuable perspective in the analysis of channel uncertainty. This is an object of future research.

The second scenario concerns a very simple network topology.

Implication 2 (single-source broadcast). Let the source-channel code (F, G) be universal for the source (p_S, d) and the class of channels $\mathcal{W}$. In the particular broadcast scenario where the single source (p_S, d) is transmitted across multiple channels $(p^{(i)}_{Y|X}, \rho^{(i)}) \in \mathcal{W}$, transmission using the source-channel code (F, G) achieves optimal performance on each channel individually.

This will be discussed and illustrated again in the context of the significance of source-channel codes for networks, in particular in Section 5.2.

3.5 Measure-matching through Feedback

The well-known Gaussian example has been extended to the situation with noiseless feedback in 1967 simultaneously by several authors [27, 57, 91]. Their work was an extension of the work of Schalkwijk and Kailath [90, 92] and Omura [77].

Example 3.3 (the simplest case of [27, 57, 91]) *Let the source be the iid Gaussian source with variance σ_S^2, and the channel the additive white Gaussian noise channel with power constraint P and additive noise variance σ^2, exactly as in Example 2.2. The difference is that* two *channel uses are available per source sample, i.e., we consider source-channel communication strategies with $\kappa = 1/2$, see Definition 1.3. From Theorem 1.7, the minimum achievable distortion is given by*

$$D_{min} = D_{\mathcal{N}}(R = 2C) = \sigma_S^2 2^{-4C} \tag{3.37}$$

$$= \sigma_S^2 \left(\frac{\sigma^2}{P + \sigma^2} \right)^2 . \tag{3.38}$$

Since feedback does not increase the capacity of a memoryless channel, this result holds whether or not feedback is available.

It is quickly verified that simply sending every source letter S twice across the channel without further coding does not achieve optimal performance; the resulting distortion is larger than D_{min}. Instead, suppose now that noiseless feedback is available, and consider the following scheme:

$$X_{2n} = \alpha_1 S_n, \tag{3.39}$$

$$X_{2n+1} = \alpha_2 (S_n - \beta Y_{2n}). \tag{3.40}$$

This means that the decoder has the two following observations available:

$$Y_{2n} = \alpha_1 S_n + W_{2n}, \tag{3.41}$$

$$Y_{2n+1} = \alpha_2 (S_n - \beta Y_{2n}) + W_{2n+1}, \tag{3.42}$$

leading to an estimate $\hat{S}_n = \gamma_1 Y_{2n} + \gamma_2 Y_{2n+1}$. If $\alpha_1, \alpha_2, \beta, \gamma_1$ and γ_2 are chosen optimally, the scheme achieves optimal performance.

To furnish a concrete example, suppose that all variances are of the same value, $\sigma_S^2 = \sigma^2 = P$. Then, $C = 1/2$ and hence, $D_{min} = P/4$. The "uncoded" *strategy described above uses the following parameters:*

$$\alpha_1 = 1, \qquad \alpha_2 = \sqrt{2}, \qquad \beta = \frac{1}{2}, \tag{3.43}$$

which satisfies the power constraint. The optimum estimators are found from standard arguments (e.g., derivatives) to be

$$\gamma_1 = \frac{1}{2}, \qquad \gamma_2 = \frac{1}{2\sqrt{2}}. \tag{3.44}$$

For this choice of the parameters, the distortion is found to be

$$\begin{aligned}
&E|S_n - \hat{S}_n|^2 \\
&\quad = E|S_n - \gamma_1 Y_{2n} - \gamma_2 Y_{2n+1}|^2 \\
&\quad = E\left|S_n - \frac{1}{2}(S_n + W_{2n}) - \frac{1}{2\sqrt{2}}\left(\sqrt{2}\left(S_n - \frac{1}{\sqrt{2}}Y_{2n}\right) + W_{2n+1}\right)\right|^2 \\
&\quad = E\left|-\frac{1}{2}W_{2n} + \frac{1}{4}Y_{2n} - \frac{1}{2\sqrt{2}}W_{2n+1}\right|^2 \\
&\quad = E\left|-\frac{1}{2}W_{2n} + \frac{1}{4}(S_n + W_{2n}) - \frac{1}{2\sqrt{2}}W_{2n+1}\right|^2 \\
&\quad = E\left|\frac{1}{4}S_n - \frac{1}{4}W_{2n} - \frac{1}{2\sqrt{2}}W_{2n+1}\right|^2 \\
&\quad = \frac{1}{16}P + \frac{1}{16}P + \frac{1}{8}P = \frac{P}{4},
\end{aligned} \tag{3.45}$$

confirming that it was an optimal choice of the parameters: $P/4$ is the smallest possible distortion. Explanations and intuition for this result are given below in Example 3.4.

Beyond the Gaussian case, no results seem to appear in the literature. Using the results of Chapters 2 and 3, this can be extended to more general cases. In order to do so, recall the line of thought leading to our main result, formulated as Corollary 3.2: The first step was to rewrite the separation theorem (Theorem 1.5) in the shape of Lemma 2.1. This led to three simultaneous conditions. The second step was to develop these three conditions individually (Lemmata 2.2 and 2.3), which eventually led to Corollary 3.2.

The first step, Theorem 1.5, applies unchanged to the feedback case (in the shape of Theorem 1.7). The next step is to rewrite the condition $\kappa R(\Delta) = C(\Gamma)$ in a more explicit form. The equivalent of Lemma 3.1 has to be determined for the case of discrete-time memoryless systems with feedback. This is the next issue in our discussion. The feedback situation is illustrated in Figure 3.3. For the purpose of this section, and in line with Example 3.3, we consider the following limited kind of feedback encoders:

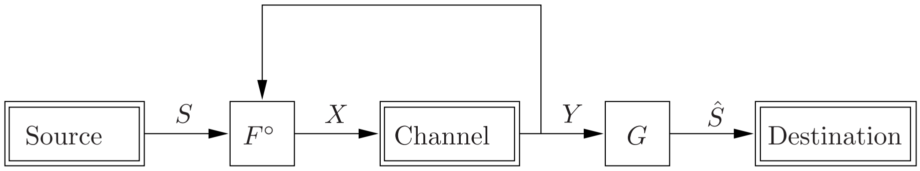

Figure 3.3: The basic communication system with feedback.

Definition 3.4 (block feedback encoder) *A block feedback encoder is defined by a function*

$$F^\circ \quad : \quad \mathcal{S}^k \times \mathcal{Y}^m \to \mathcal{X}^m, \tag{3.46}$$

with the restriction that the dependence on Y^m must be causal, i.e., the jth output of the encoder is

$$X_j \quad = \quad F_j^\circ(S^k, Y_{j-1}, Y_{j-2}, \ldots, Y_1). \tag{3.47}$$

Note that the definition of a block feedback encoder incorporates only a reduced feedback, namely only *within* one block of k source or m channel symbols, respectively. Under this restriction, subsequent blocks are independent.

This leads to the following generalization of Lemma 2.1:

Lemma 3.8 (generalization of Lemma 2.1) *For a discrete-time memoryless source (p_S, d) and a discrete-time memoryless channel $(p_{Y|X}, \rho)$ (where memoryless is taken in the sense of Equation (1.32)), and a block feedback code (F°, G) as in Figure 3.3, $\kappa R(\Delta) = C(\Gamma)$ holds if and only if the following three conditions are simultaneously satisfied:*

(i) (a) $Y_1, Y_2, \ldots$ are independent random variables, and

(b) the marginals of $X_1, X_2, \ldots$ all achieve capacity on the channel $(p_{Y|X}, \rho)$ at the same input cost $\Gamma = E\rho(X)$, i.e., $I(X;Y) = C(\Gamma)$,

(ii) the conditional distribution $p_{\hat{S}^k|S^k}$ of $\hat{S}^k = G(Y^m)$ given S^k achieves the rate-distortion function of the source (p_S, d) at distortion $\Delta = Ed(S^k, \hat{S}^k)$, i.e. $I(S^k; \hat{S}^k) = kR(\Delta)$, and

(iii) F° and G are such that $I(S^k; \hat{S}^k) = \sum_{i=1}^m I(X_i; Y_i)$, i.e., they are "information lossless."

Proof. We refer to the proof of Theorem 1.7. The conditions for equality in the inequalities (1.37) are as follows: Equality holds in (a) if and only if $p_{\hat{S}^k|S^k}$ achieves the rate-distortion function of the source, and in (d) if and only if $Y_1, \ldots, Y_m$ are independent random variables (which is proved in [74], see also the proof of Theorem 1.7), and (e) if and only if p_{X^n} achieves the maximum directed information, which is equal to $mC(\Gamma)$. Thus, $\kappa R(\Delta) = C(\Gamma)$ is satisfied if and only if all three conditions in Lemma 3.8 are satisfied, which completes the proof. □

Finally, we put things together to obtain the equivalent of Theorem 2.5, or more precisely, of Theorem 3.3:

Theorem 3.9 *For a discrete-time memoryless source* (p_S, d)*, and a discrete-time memoryless channel* $(p_{Y|X}, \rho)$ *(where memoryless is taken in the sense of Equation (1.32)), suppose that there exist values of* P *and* D *such that* $R(D) = C(P)$*. Consider the transmission using a block feedback source-channel code* (F°, G) *with* $F^\circ : \mathcal{S}^k \times \mathcal{Y}^m \to \mathcal{X}^m$ *and* $G : \mathcal{Y}^m \to \hat{\mathcal{S}}^k$*, and suppose that* $0 < I(S^k; \hat{S}^k) = \sum_{i=1}^m I(X_i; Y_i) < mC_0$*. If*

$$p(y_1, \ldots, y_m) \quad = \quad p_Y(y_1) \cdots p_Y(y_m), \tag{3.48}$$

for some $p_Y(\cdot)$*, and*

$$\rho(x) \quad \begin{cases} = & c_1 D(p_{Y|X}(\cdot|x) || p_Y(\cdot)) + \rho_0 & \text{if } p(x) > 0 \\ \geq & c_1 D(p_{Y|X}(\cdot|x) || p_Y(\cdot)) + \rho_0 & \text{otherwise,} \end{cases} \tag{3.49}$$

$$d(s, \hat{s}) \quad = \quad -c_2 \log_2 p_{S_j|\hat{S}_j}(s|\hat{s}) + d_0(s), \text{ for } j = 1, \ldots, k, \tag{3.50}$$

then the system performs optimally.

Outline of the proof. By assumption, Y_i are independent random variables. If ρ is chosen according to the formula, then the channel input achieves capacity. If d is chosen according to the formula, then the overall conditional achieves the rate-distortion function. Moreover, by assumption, $I(S^k; \hat{S}^k) = \sum_{i=1}^m I(X_i; Y_i)$. All of this together implies that the system satisfies $\kappa R(\Delta) = C(\Gamma)$ (by Lemma 3.8). To ensure the optimality of the system, one has to verify that Δ cannot be reduced without reducing $R(\Delta)$, and that Γ cannot be reduced without reducing $C(\Gamma)$. By the arguments of Proposition 2.4, the condition $0 < I(S^k; \hat{S}^k) = \sum_{i=1}^m I(X_i; Y_i) < mC_0$ is sufficient for this. □

Example 3.4 (continuation of Example 3.3) *Both the receiver and (through the feedback) the transmitter can form the minimum mean-squared error (MMSE) estimate* $\hat{S}'_n$ *based on* Y_{2n}*. By the orthogonality principle (see e.g. [11, Proposition V.C.2]), the error associated with* $\hat{S}'_n$ *is orthogonal to the observations that led to* $\hat{S}'_n$*, i.e.,* $E[(\hat{S}'_n - S_n)Y_{2n}] = 0$*. In the Gaussian case, this implies that* $(\hat{S}'_n - S_n)$ *and* Y_n *are independent. Hence, if the transmitter sends* $X_{2n+1} = \alpha_2(\hat{S}'_n - S_n)$*, then* Y_{2n+1} *is independent of* Y_{2n}*. It is quickly verified that with the suggested choice of* α_2 *and* β*,* Y_1 *and* Y_2 *are independent and identically distributed. More precisely, they are normal with mean zero and variance* $2P$*. To apply Theorem 3.9, we have to determine* $I(S; \hat{S})$*. From the development (3.45), we know that*

$$\hat{S}_n \quad = \quad \frac{3}{4} S_n + \frac{1}{4} W_{2n} + \frac{1}{2\sqrt{2}} W_{2n+1}, \tag{3.51}$$

from which we find immediately

$$I(S_n; \hat{S}_n) = \frac{1}{2}\log_2\left(1 + \frac{\frac{9}{16}P}{\frac{1}{16}P + \frac{1}{8}P}\right) = 1. \quad (3.52)$$

Moreover, $\sum_{i=0}^{1} I(X_{2n+i}; Y_{2n+i})$ *can be found easily by noting that the marginals of* X_{2n} *and* X_{2n+1} *are both capacity achieving, hence*

$$\sum_{i=0}^{1} I(X_{2n+i}; Y_{2n+i}) = 2\frac{1}{2}\log_2\left(1 + \frac{P}{P}\right) = 1. \quad (3.53)$$

Hence, in this example, the condition $I(S_n; \hat{S}_n) = \sum_{i=0}^{1} I(X_{2n+i}; Y_{2n+i})$ *is satisfied.*

The next step is to determine $\rho(x)$ *as in Theorem 3.9. The exact same calculation as in Example 2.5 reveals that* $\rho(x) = x^2$*. For the distortion measure, one has to determine the conditional density of* S *given* $\hat{S}$*. Again, from the development (3.45),*

$$S_n = \frac{4}{3}\hat{S}_n - \frac{1}{3}W_{2n} - \frac{\sqrt{2}}{3}W_{2n+1}, \quad (3.54)$$

The exact same calculation as in Example 2.5 reveals that with appropriate scaling, $d(s, \hat{s}) = (s - \hat{s})^2$*.*

We conclude this short discussion of source-channel communication with feedback with a simple discrete example.

Example 3.5 (quaternary source and binary channel) [4] *In this example, we consider a quaternary source whose symbols are denoted by 0, 1, 2, and 3, with distribution*

$$p(s) = \begin{cases} \frac{1}{3}, & \text{if } s = 0 \text{ or } s = 3, \\ \frac{1}{6}, & \text{otherwise.} \end{cases} \quad (3.55)$$

The channel is the binary symmetric channel with transition probability $\epsilon = \frac{1}{4}$*.*

Suppose there are two channel uses available per source sample. We consider the following feedback encoder, mapping one source symbol, s*, onto two channel inputs,* x_1 *and* x_2*, as follows:*

$$x_1 = \begin{cases} 0, & \text{if } s = 0, 1 \\ 1, & \text{otherwise.} \end{cases} \quad (3.56)$$

$$x_2 = \begin{cases} 0, & \text{if } s = 0, \quad \text{and} \quad y_1 = 0, \quad \text{or} \\ & \text{if } s = 0, 1, 2, \quad \text{and} \quad y_1 = 1, \\ 1, & \text{otherwise.} \end{cases} \quad (3.57)$$

The conditional joint distribution of Y_1 *and* Y_2 *given* $S = 0$ *is*

[4] This example appears to be new here; it is inspired by the capacity considerations of Horstein in [55].

	$Y_1 = 0$	$Y_1 = 1$
$Y_2 = 0$	$(1-\epsilon)^2 = \frac{9}{16}$	$\epsilon(1-\epsilon) = \frac{3}{16}$
$Y_2 = 1$	$(1-\epsilon)\epsilon = \frac{3}{16}$	$\epsilon^2 = \frac{1}{16}$

The table for $S = 3$ is obtained by swapping the roles of $Y_1 = 0$ and $Y_1 = 1$ as well as the roles of $Y_2 = 0$ and $Y_2 = 1$. The conditional joint distribution of Y_1 and Y_2 given $S = 1$ is

	$Y_1 = 0$	$Y_1 = 1$
$Y_2 = 0$	$(1-\epsilon)\epsilon = \frac{3}{16}$	$\epsilon(1-\epsilon) = \frac{3}{16}$
$Y_2 = 1$	$(1-\epsilon)^2 = \frac{9}{16}$	$\epsilon^2 = \frac{1}{16}$

The table for $S = 2$ is obtained by swapping the roles of $Y_1 = 0$ and $Y_1 = 1$ as well as the roles of $Y_2 = 0$ and $Y_2 = 1$. By adding the tables,

$$p(y_1, y_2) \quad = \quad \sum_{s=0}^{3} p(y_1, y_2|s)p(s), \tag{3.58}$$

it is quickly verified that Y_1 and Y_2 are independent and uniform binary random variables. From this, it follows immediately that the scheme achieves the unconstrained capacity C_0 of the channel, hence any cost function will do.

Furthermore, define the decoding operation as follows:

$$\hat{s} \quad = \quad \begin{cases} 0, & \text{if } y_1 = 0 \quad \text{and} \quad y_2 = 0, \\ 1, & \text{if } y_1 = 0 \quad \text{and} \quad y_2 = 1, \\ 2, & \text{if } y_1 = 1 \quad \text{and} \quad y_2 = 0, \\ 3, & \text{if } y_1 = 1 \quad \text{and} \quad y_2 = 1. \end{cases} \tag{3.59}$$

To verify the condition $I(S;\hat{S}) = I(X_1;Y_1) + I(X_2;Y_2)$, write out the discrete memoryless channel between S and $\hat{S}$, and take the input distribution of Equation (3.55). This reveals that $I(S;\hat{S}) = 2 - 2h_b(\epsilon)$,[5] which is equal to $I(X_1;Y_1) + I(X_2;Y_2)$, since Y_1 and Y_2 were shown to be independent uniform binary random variables. According to Theorem 3.9, a distortion measure under which this system is optimal is

$$d(s, \hat{s}) \quad = \quad -c_2 \log_2 p(s|\hat{s}) + d_0(s). \tag{3.60}$$

Using the above tables, this can be evaluated easily:

$$\begin{aligned} p(s = i|\hat{s} = j) &= p(s = i|y_1 = k_1(j), y_2 = k_2(j)) \\ &= \frac{p(y_1 = k_1(j), y_2 = k_2(j)|s = i)p(s = i)}{p(y_1 = k_1(j), y_2 = k_2(j))} \\ &= p(y_1 = k_1(j), y_2 = k_2(j)|s = i)\frac{p(s = i)}{1/4}, \end{aligned}$$

and the expression $p(y_1 = k_1(j), y_2 = k_2(j)|s = i)$ only assumes one out of three different values, namely $\epsilon^2, \epsilon(1-\epsilon)$ or $(1-\epsilon)^2$. To normalize this distortion measure, we pick c_2 and $d_0(s)$ in Theorem 3.9 as

$$c_2 = \frac{1}{\log_2 3} \quad \text{and} \quad d_0(s) = -\frac{1}{\log_2 3}\log_2(4(1-\epsilon)^2 p(s)).$$

[5] Here, $h_b(\cdot)$ denotes the binary entropy function, i.e., $h_b(p) = -p\log_2 p - (1-p)\log_2(1-p)$

In other words, an optimal distortion measure is the following:

$$d(s,\hat{s}) = \frac{1}{\log_2 3}\left(\log_2 p(s|\hat{s}) - \log_2(4(1-\epsilon)^2 p(s))\right). \tag{3.61}$$

This evaluates, for $j = 0, 3$, to

$$d(s = j, \hat{s} = i) = \begin{cases} 0, & \text{if } j = i, \\ 1, & \text{if } |j-i| = 1, \text{ or } |j-i| = 2 \\ 2, & \text{if } |j-i| = 3, \end{cases} \tag{3.62}$$

which is thus the Hamming distance between the binary representations of the integers $0, 1, 2,$ *and* 3*. For* $j = 1, 2$,

$$d(s = j, \hat{s} = i) = |j - i|. \tag{3.63}$$

It should be clear that these two cases cannot easily be expressed in terms of one and the same distance measure, simply because the distance between integers, and the Hamming distance between their binary representations impose a different structure. The distortion can still be expressed in one formula, as follows:

$$d(s,\hat{s}) = \begin{cases} 0, & \text{if } s = \hat{s}, \\ 2, & \text{if } |s-\hat{s}| = \max_t |s-t| \\ 1, & \text{otherwise.} \end{cases} \tag{3.64}$$

The main goal of this example is to illustrate that measure-matching through feedback is not limited to the Gaussian case: rather, it reveals a new class of source/channel pairs for which very simple source-channel codes, now with feedback, achieve optimal performance.

3.6 Connections to Other Results

The results discussed in this chapter provide different perspectives on other results. For example, they have been used in [82] to analyze the duality between source and channel coding, with and without side information, and in [75] to analyze uncoded transmission for the Wyner-Ziv source and the Gel'fand-Pinsker channel. In this section, we consider the queuing channel studied in [15] from the point of view of our results.

3.6.1 Bits through queues

Anantharam and Verdú derive in [15] a capacity result for a certain queuing channel. The channel is a Poisson queue of service rate μ, and the channel inputs are impulses. The contents of these impulses is of no relevance to the problem studied in [15]. They may be information packets, for example. The main insight is that the *timing* of these impulses can be used to transmit information, similar to pulse position modulation. Each pulse is delayed by a random amount of time, but this random amount depends on the delay of the previous pulses. It is found that the capacity-achieving input distribution is a Poisson process.[6]

The problem can also be considered as a joint source-channel communication problem, as follows: A Poisson process of rate λ is transmitted without further coding across a Poisson queue of service rate μ. This can be plugged into Theorem 2.5 (or, more particularly, Theorem 3.3) to yield the cost and distortion functions that make this an optimal source-channel communication system.

The input to the queue is a sequence of impulses with random intervals. Denote the ith interval by A_i. In our case, the input process is Poisson, and hence, A_i is exponentially distributed. A vector of n such intervals has joint distribution

$$p(a_1, \ldots, a_n) = \prod_{i=1}^{n} \mathrm{e}_\lambda(a_i), \tag{3.65}$$

where $\mathrm{e}_\lambda(\cdot)$ denoted the exponential distribution with mean $1/\lambda$. For a Poisson queue of service rate $\mu > \lambda$, it can be shown [15, Eqn. (2.23)] (see also [107]) that the output is also a Poisson process. Denote the intervals between the output impulses by D_i. The joint distribution of n such intervals can be given as

$$p(d_1, \ldots, d_n) = \prod_{i=1}^{n} \mathrm{e}_\lambda(d_i). \tag{3.66}$$

The conditional density is given by the queue specification. As shown in [15], it can be expressed as

$$D_i = W_i + S_i, \tag{3.67}$$

where S_i is the service time and is independent of the entire sequence $\{A_j\}$ and of the past $\{D_j\}_{j=1}^{i-1}$, and W_i are the idling times,

$$W_i = \max\left\{0, \sum_{j=1}^{i} A_j - \sum_{j=1}^{i-1} D_j\right\}, \tag{3.68}$$

i.e., W_i is a deterministic function of $\{A_j\}_{j=1}^{i}$ and $\{D_j\}_{j=1}^{i}$. Hence, we can write

$$p(d_1, \ldots, d_n | a_1, \ldots, a_n) = p_{S_1,\ldots,S_n}(d_1 - w_1, \ldots, d_n - w_n), \tag{3.69}$$

where $p_{S_1,\ldots,S_n}$ is the distribution of the service times of the queue. For the Poisson queue with service parameter μ,

$$p(d_1, \ldots, d_n | a_1, \ldots, a_n) = \mathrm{e}_\mu(d_1 - w_1) \cdots \mathrm{e}_\mu(d_n - w_n). \tag{3.70}$$

[6] In a discussion on the matters of Chapter 2 during the 2001 DSC Summer Research Institute at EPFL, Prof. Sergio Verdú suggested to us to study this “Poisson-over-Poisson” example.

Theorem 3.3 tells us that an optimal choice for the cost function is given by

$$\begin{aligned}\rho(a_1,\ldots,a_n) &= D(p(d_1,\ldots,d_n|a_1,\ldots,a_n)||p(d_1,\ldots,d_n)) \\ &= \int_{d_1=w_1}^{\infty}\cdots\int_{d_n=w_n}^{\infty}\left(\prod_{i=1}^{n}\mu e^{-\mu(d_i-w_i)}\right) \\ &\qquad \left(-\mu\sum_{j=1}^{n}(d_j-w_j)+\lambda\sum_{j=1}^{n}d_j\right)dd_1\cdots dd_n,\end{aligned}$$

which is well defined: w_i can be easily computed for given $a_1,\ldots,a_i$ and given $d_1,\ldots,d_{i-1}$. We did not, however, find a closed-form solution for this integral.

In order to still gain some insight, we change the integration variables. Introduce $s_i = d_i - w_i$. The integral can be expressed as

$$\int_{s_1=0}^{\infty}\cdots\int_{s_n=0}^{\infty}\left(\prod_{i=1}^{n}\mu e^{-\mu s_i}\right)\left(-\mu\sum_{j=1}^{n}s_j+\lambda\sum_{j=1}^{n}(s_j+w_j)\right)ds_1\cdots ds_n.$$

By the construction of w_i, the variable s_i does not depend on the values of $a_1,\ldots,a_n$, i.e., the dependence on $a_1,\ldots,a_n$ resides in $w_1,\ldots,w_n$ exclusively. In other words, we can write

$$\rho(a_1,\ldots,a_n) = \int_{s_1=0}^{\infty}\cdots\int_{s_n=0}^{\infty}\left(\prod_{i=1}^{n}\mu e^{-\mu s_i}\right)\lambda\sum_{j=1}^{n}w_j ds_1\cdots ds_n + const.,$$

where w_i is determined by $a_1,\ldots,a_i$ and on $s_1,\ldots,s_{i-1}$ through (3.67) and (3.68):

$$w_i = \max\left\{0, \sum_{j=1}^{i}a_j - \sum_{j=1}^{i-1}(w_j - s_j)\right\}. \qquad (3.71)$$

To interpret the expression for $\rho(a_1,\ldots,a_n)$, note that for given $a_1,\ldots,a_n$, all the randomness of W_i is due to $S_1,\ldots,S_{i-1}$, hence the above integral is the expectation of $\sum_{i=1}^{n} W_i$ for given $a_1,\ldots,a_n$. In words, the cost of using a certain input sequence $a_1,\ldots,a_n$ is the resulting expected sum of the idling times.

By the same token, one can determine the distortion measure according to the formula of Theorem 3.3. For the Poisson source, rate-distortion considerations have appeared in [89] and [107] for two different distortion measures. Only the latter led to closed-form results. The distortion measure according to

Theorem 3.3 does not seem to coincide with either. More precisely, we find

$$
\begin{aligned}
& d((a_1, \ldots, a_n), (d_1, \ldots, d_n)) \\
& \quad = -\log_2 p(a_1, \ldots, a_n | d_1, \ldots, d_n) \\
& \quad = -\log_2 \frac{p(d_1, \ldots, d_n | a_1, \ldots, a_n) p(a_1, \ldots, a_n)}{p(d_1, \ldots, d_n)} \\
& \quad = -\log_2 p(d_1, \ldots, d_n | a_1, \ldots, a_n) - \log_2 \frac{p(a_1, \ldots, a_n)}{p(d_1, \ldots, d_n)} \\
& \quad = \mu \sum_{j=1}^{n} (d_j - w_j) + \lambda \sum_{j=1}^{n} (a_j - d_j) + const. \qquad (3.72)
\end{aligned}
$$

By Theorem 3.3, we can add or subtract an arbitrary function of $a_1, \ldots, a_n$, and hence, an equivalent distortion measure is

$$
d((a_1, \ldots, a_n), (d_1, \ldots, d_n)) \quad = \quad (\mu - \lambda) \sum_{j=1}^{n} d_j - \mu \sum_{j=1}^{n} w_j. \qquad (3.73)
$$

The result of this consideration is a new example of optimal uncoded transmission, and can be summarized as follows: When a Poisson process is to be transmitted across a Poisson queue channel, and the cost is the expected sum of the idling times of the queue, then uncoded transmission is optimal.

3.7 Summary and Conclusions

The (ergodic, point-to-point) source-channel communication problem can be solved by the separation theorem. This was discussed in Chapter 1. The resulting perspective can be called *rate-matching*: The solution to the source-channel communication problem is found by matching the rate of the source (its rate-distortion function) to the rate of the channel (its capacity).

In this chapter, we developed a different point of view: optimal solutions to the source-channel communication problem can also be characterized by a particular relationship between the source and channel parameters and the encoding and the decoding functions. In particular, the source parameters consist of the source probability distribution and the distortion measure. The channel parameters consist of the channel conditional probability distribution and the channel input cost function. Since these can all be seen as measures of some kind, we call the resulting perspective *measure-matching*, and the conditions expressed in Corollary 3.2 the *measure-matching conditions.*

One special feature of the measure-matching conditions given in Corollary 3.2 is the fact that they are explicit, at least in one direction: for fixed probability distributions and encoding/decoding functions, the optimal cost function and the distortion measure are derived in closed form. This fact was exploited in various ways in this chapter. For example, it was derived that for a relevant class of source/channel pairs, the optimal match can be achieved either by a code of block length one, or else by no code of finite block length. Measure-matching also serves to define a property of universality featured by certain source-channel codes, and to analyze source-channel communication systems with feedback.

The search for further instances where the measure-matching conditions can be successfully applied is an object of future research.

Appendix 3.A Proof of Theorem 3.4

Proof of Theorem 3.4. ($\Leftarrow$.) If there is a single-letter code with optimal performance, then trivially there is also a code of length M with optimal performance. ($\Rightarrow$.) Under the stated assumptions, the existence of a code of length M with optimal performance implies the existence of a single-letter code with optimal performance for the same source and channel. To prove this, we consider single-letter codes for the length-M extension source and channel.

Notation: We use the notation introduced in Section 3.1, and reprinted here for convenience. Let $s^M = (s_1, \ldots, s_M)$ be the vector of M consecutive source symbols, and define $\hat{s}^M$ accordingly. By assumption, all alphabets are of the same size. Without (further) loss of generality, we use the generic alphabet $\{1, 2, \ldots, K\}$. The length-M extension source is $p(s^M) = \prod_{m=1}^M p_S(s_m)$ with $d^{(M)}(s^M, \hat{s}^M) = \sum_{m=1}^M d(s_m, \hat{s}_m)$. For some of the considerations below, it will be more convenient to map s^M into an *extension alphabet* of size K^M. We use bold face characters to denote symbols of the extension alphabets, and the mapping is carried out according to $\mathbf{s} = \sum_{i=1}^M K^i s_i$. Both representations will be used interchangeably. Similarly, the extension channel is $p(y^M|x^M) = \prod_{m=1}^M p_{Y|X}(y_m|x_m)$ with $\rho^{(M)}(x^M) = \sum_{m=1}^M \rho(x_m)$. In the proof, it will also be handy to use matrix notation. We will use $P_{Y|X}$ for the matrix of channel transition probabilities, where y indexes the rows and x the columns. Note that in the extension alphabet, $P_{\mathbf{Y}|\mathbf{X}} = P_{Y|X} \otimes \ldots \otimes P_{Y|X}$ (M terms), where $\otimes$ denotes the Kronecker product (tensor product).

Outline: The single-letter code for the length-M extension will be denoted (F, G). Obviously, this is an M-letter code for the original source and channel. We will now apply the theory developed in Chapters 2 and 3 to the extension source and channel, and their single-letter code (F, G). Plugging $p_{\mathbf{S}}$, $p_{\mathbf{Y}|\mathbf{X}}$ and the code (F, G) into Formulae (2.14) and (2.15) of Theorems 2.2 and 2.3, we obtain the $\rho^{(M)}$ and $d^{(M)}$ that are necessary and sufficient for optimal performance.[7] However, by assumption, they have to be averaging (or *single-letter*) measures, that is, $\rho^{(M)}(x^M) = \sum_{m=1}^M \rho(x_m)$ for some cost function $\rho(\cdot)$, and $d^{(M)}(s^M, \hat{s}^M) = \sum_{m=1}^M d(s_m, \hat{s}_m)$ for some distortion measure $d(\cdot, \cdot)$. This excludes many of the possible M-letter codes (F, G).

From Theorem 2.3, the distortion measure has to be chosen as

$$d^{(M)}(s^M, \hat{s}^M) \quad = \quad -\log_2 p(s^M|\hat{s}^M). \tag{3.74}$$

Clearly, for this to split additively into equal functions each of which depends only on one of the pairs $(s_i, \hat{s}_i)$, it is necessary that $p(s^M|\hat{s}^M) = p_{S|\hat{S}}(s_1|\hat{s}_1) \cdot \ldots \cdot p_{S|\hat{S}}(s_M|\hat{s}_M)$. This can also be inferred from the argument given in Remark 3.2. In terms of transition probability matrices, this can be expressed as

$$P_{\mathbf{S}|\hat{\mathbf{S}}} \quad = \quad P_{S|\hat{S}} \otimes \ldots \otimes P_{S|\hat{S}}. \tag{3.75}$$

[7]Suppose there exists a code (F, G) such that $I(X^M; Y^M) = MC_0$. When all alphabets are of the same cardinality and $p(s) > 0$ for all s, it is a simple matter to prove that there exists also a single-letter code that achieves $I(X; Y) = C_0$. For this reason, the interesting case is when $I(X^M; Y^M) < MC_0$, in which case the formula for ρ is indeed a necessary condition. A similar comment applies to the case $I(S; \hat{S}) = 0$.

By symmetry, the second key insight follows from the fact that the cost function has to split additively. From Remark 3.3, we know that Y are iid. However, the derivation is somewhat more technical. Therefore, we state the result in the shape of the following lemma, to be proved below:

Lemma 3.10 *If $\rho^{(M)}$ is averaging and $P_{Y|X}$ invertible, then X and Y are iid.*

Note that the fact that Y are iid follows by the argument given in Remark 3.3. Hence, it remains to prove the assertion for X.

The third insight is that under the additional assumptions on the alphabet sizes and $p(s)$, the encoder and decoder have to be bijective. It is given by the following lemma (to be proved below):

Lemma 3.11 *If all alphabets are of the same cardinality, $p(s) > 0$ for all s, $P_{S|\hat{S}}$ and $P_{Y|X}$ are invertible and $d^{(M)}$ is averaging, then encoder F and decoder G are permutation matrices.*

To complete the proof, consider first the encoder. Suppose that for fixed distribution of S and X, there exists indeed a bijective encoder F that maps $\mathbf{S}$ to $\mathbf{X}$. Equivalently, this means that there exists a permutation matrix F such that $p_{\mathbf{X}} = F p_{\mathbf{S}}$, where $p_{\mathbf{X}}$ here means a vector containing the probabilities $p_{\mathbf{X}}(\mathbf{x})$, and $p_{\mathbf{S}}$ the corresponding for the random variable $\mathbf{S}$. By Lemma 3.10, X is iid, hence we can write

$$p_X \otimes \ldots \otimes p_X \quad = \quad F(p_S \otimes \ldots \otimes p_S). \tag{3.76}$$

But this can only be true if there exists also a permutation matrix f such that

$$p_X \quad = \quad f p_S. \tag{3.77}$$

In other words, there exists also a single-letter encoder f that maps S to X.

This argument can be applied to the matrix $P_{\mathbf{S}|\hat{\mathbf{S}}}$ to conclude that the decoder can also be implemented by a single-letter mapping. First, recall that $P_{\mathbf{S}|\hat{\mathbf{S}}} = P_{\mathbf{S}|\mathbf{X}} P_{\mathbf{X}|\mathbf{Y}} P_{\mathbf{Y}|\hat{\mathbf{S}}}$. On the right hand side, $P_{\mathbf{X}|\mathbf{Y}}$ can be written as an M-fold Kronecker product because the channel is memoryless and X and Y are iid. Moreover, we have just shown that the encoder is a permutation matrix, and that it can be written also as an M-fold Kronecker product. Using Eqn. (3.75), we find

$$\begin{aligned} P_{S|\hat{S}} \otimes \ldots \otimes P_{S|\hat{S}} &= (P_{S|X} \otimes \ldots \otimes P_{S|X})(P_{X|Y} \otimes \ldots \otimes P_{X|Y}) P_{\mathbf{Y}|\hat{\mathbf{S}}} \\ &= (A \otimes \ldots \otimes A) P_{\mathbf{Y}|\hat{\mathbf{S}}} \end{aligned} \tag{3.78}$$

for some matrix A. But if there does indeed exist a permutation matrix $P_{\mathbf{Y}|\hat{\mathbf{S}}}$ (namely, the decoder G) that satisfies the above equation, then there exists also a permutation matrix $P_{Y|\hat{S}}$ that satisfies $P_{S|\hat{S}} = A P_{Y|\hat{S}}$, which implies the existence of a single-letter decoder g. □

Remark 3.5 *Let us explain at this point why the additional assumptions in Theorem 3.4 are necessary: To ensure that the encoder and the decoder are permutation matrices. Using this, the step from Eqn. (3.76) to Eqn. (3.77) is simple. However, if this is not ensured, then the step from Eqn. (3.76) to Eqn. (3.77) seems to become surprisingly tricky. Relatively little seems to be known about Kronecker products.*

Proof of Lemma 3.10.

Part 1: *Under the stated conditions, Y_i are iid random variables.* This can be inferred from the argument in Remark 3.3. Here, we provide an alternative direct proof, using matrix arguments.

From Theorem 2.2, the cost function $\rho(x^M)$ has to be chosen as

$$\rho(x^M) \quad = \quad D\left(p_{Y^M|X^M}(\cdot|x^M)||p_{Y^M}(\cdot)\right) = \sum_{y^M} p(y^M|x^M)\log_2 \frac{p(y^M|x^M)}{p(y^M)} \tag{3.79}$$

By definition, the cost function of a memoryless channel has to split additively into M equal functions, each depending only on one of the x_i. It is now shown that this implies that $p(y_1, \ldots, y_M) = p_Y(y_1) \cdot \ldots \cdot p_Y(y_M)$. For the case $M = 2$,

$$\begin{aligned} \rho(x_1, x_2) \quad &= \quad H(Y|X = x_1) + H(Y|X = x_2) \\ &\quad - \sum_{y_1, y_2} p(y_1|x_1)p(y_2|x_2)\log_2 p(y_1, y_2). \end{aligned} \tag{3.80}$$

The last double sum has to split additively into two parts, one depending only on x_1, the other only on x_2. As a first step, we now show that this implies that Y_1 and Y_2 are independent random variables. Equivalently, we show that the matrix $P_{Y_1Y_2}$ containing the joint pmf of Y_1 and Y_2 has rank at most 1.

To see why this holds, let us introduce the following shorthand: $z_i^j = p(y = j|x_i)$, where $1 \leq i \leq K$ and $1 \leq j \leq K$. Moreover, in this paragraph, we use $p(\cdot, \cdot)$ in place of $p_{Y_1Y_2}(\cdot, \cdot)$ to make the formulae more readable. With this, we can rewrite the double sum on the RHS of Eqn. (3.80) as

$$\begin{aligned} & z_1 z_2 \log_2 p(1,1) + z_1 z_2^2 \log_2 p(1,2) + \ldots + z_1 z_2^K \log_2 p(1,K) \\ + \quad & z_1^2 z_2 \log_2 p(2,1) + z_1^2 z_2^2 \log_2 p(2,2) + \ldots + z_1^2 z_2^K \log_2 p(1,K) \\ + \quad & \vdots \\ + \quad & z_1^K z_2 \log_2 p(K,1) + z_1^K z_2^2 \log_2 p(K,2) + \ldots + z_1^K z_2^K \log_2 p(K,K), \end{aligned} \tag{3.81}$$

with the constraint

$$z_i + z_i^2 + \ldots + z_i^K \quad = \quad 1, \text{ for all } i. \tag{3.82}$$

To split the sum additively into terms that depend only on one of the x_i (or, equivalently, of the z_i), it is necessary that the coefficients of all terms that involve more than one of the variables z_i are zero. Substitute for instance

$z_1^2 = 1 - z_1 - z_2^3 - \ldots - z_2^K$ and $z_2^2 = 1 - z_2 - z_2^3 - \ldots - z_2^K$. Then it is quickly verified that the coefficient of $z_1 z_2$ is

$$\log_2 p(1,1) + \log_2 p(2,2) - \log_2 p(1,2) - \log_2 p(2,1). \tag{3.83}$$

But this is precisely the determinant of a 2×2 submatrix of $P_{Y_1 Y_2}$. In a similar fashion, we find that the determinants of all 2×2 submatrices of $P_{Y_1 Y_2}$ have to be zero. But this implies that $\operatorname{rank} P_{Y_1 Y_2} \leq 1$ (a well-known fact for which we did not find a reference, but which has a short proof; therefore it is given below as Lemma 3.12), which implies that Y_1 and Y_2 must be independent random variables.

For $M > 2$, define two sets of indices, $\mathcal{I}$ and $\mathcal{J}$, such that $\mathcal{I} \cap \mathcal{J} = \emptyset$. Let $Y^{(I)} = \{Y_i : i \in \mathcal{I}\}$ and $Y^{(J)} = \{Y_j : j \in \mathcal{J}\}$. But since Y are discrete random variables, $Y^{(I)}$ and $Y^{(J)}$ can be interpreted as two discrete random variables over larger alphabets. Denote the joint pmf matrix of $Y^{(I)}$ and $Y^{(J)}$ by P_{IJ}. For this matrix, it can again be shown that all 2×2 submatrices have zero determinant, and from Lemma 3.12, that P_{IJ} has rank one. Hence, the joint distribution matrix is $P_{IJ} = p_{Y^{(I)}} p'_{Y^{(J)}}$. Since this holds for any two index sets, it follows that the Y_i are independent random variables.

Up to now, we have established that $Y_1, \ldots, Y_M$ have to be independent random variables, thus we can write

$$\begin{aligned}\rho(x_1, \ldots, x_M) &= H(Y|X = x_1) - \sum_{y_1} p(y_1|x_1) \log_2 p(y_1) \\ &\quad + \ldots + H(Y|X = x_M) - \sum_{y_2} p(y_2|x_M) \log_2 p(y_2),\end{aligned} \tag{3.84}$$

which indeed splits additively into M functions, each of which depends only on one of the x_i. Moreover, it has to split into *equal* functions. That is, whenever $x_i = x_j$, we must have that

$$\sum_{y_i} p(y_i|x_i) \log_2 p(y_i) = \sum_{y_j} p(y_j|x_j) \log_2 p(y_j), \tag{3.85}$$

which can be rewritten (by letting $x = x_i = x_j$)

$$\sum_{y} p(y|x)(\log_2 p_{Y_i}(y) - \log_2 p_{Y_j}(y)) = 0. \tag{3.86}$$

This must hold for every choice of x. In other words, the vector $\{\log_2 p_{Y_i}(y) - \log_2 p_{Y_j}(y)\}_y$ must be orthogonal to all of the K vectors $\{p(y|x)\}_y$. Hence, if those K vectors span the entire K-dimensional space, then $p_{Y_i} = p_{Y_j}$, and thus Y_i and Y_j are identically distributed random variables. Thus, if the channel transition probability matrix $P_{Y|X}$ admits a right inverse, then the channel outputs $Y_1, \ldots, Y_M$ must be iid random variables.

Part 2: *Under the stated conditions, X_i are iid random variables.*

Under certain circumstances, the fact that $Y_1, \ldots, Y_M$ are iid implies that $X_1, \ldots, X_M$ are also iid. A sufficient (but not necessary) condition for this is that the channel transition probability matrix $P_{Y|X}$ admit a left inverse. For codes of length $M = 2$, this can be shown as follows. Construct matrices $P_{Y_1Y_2} = \{p(y_1, y_2)\}_{y_1,y_2}$ and $P_{X_1X_2} = \{p(x_1, x_2)\}_{x_1,x_2}$. Then, we can write

$$P_{Y_1Y_2} = P_{Y|X} P_{X_1X_2} P_{Y|X}^T. \quad (3.87)$$

Denote the left inverse of $P_{Y|X}$ by $P_{Y|X}^L$. Then,

$$P_{Y|X}^L P_{Y_1Y_2} P_{Y|X}^{LT} = P_{X_1X_2}. \quad (3.88)$$

However, since Y_1 and Y_2 are iid, $P_{Y_1Y_2} = pp^T$ for some vector p, and thus

$$\operatorname{rank} P_{X_1X_2} = \operatorname{rank}(P_{Y|X}^L P_{Y_1Y_2} P_{Y|X}^{LT}) \leq \operatorname{rank} P_{Y_1Y_2} = 1. \quad (3.89)$$

Since moreover, $P_{X_1X_2} = P_{X_1X_2}^T$, there must exist a vector q such that $P_{X_1X_2} = qq^T$.

To extend this argument to $M > 2$, we use again the sets $\mathcal{I}$ and $\mathcal{J}$ as defined above. The joint distribution of $Y^{(I)}$ and $Y^{(J)}$ can thus be written in matrix form as $P_{IJ} = p_{Y^{(I)}} p'_{Y^{(J)}}$. This is a rectangular matrix of dimension $K^{|\mathcal{I}|} \times K^{|\mathcal{J}|}$. By construction, it has only one non-zero singular value. The transition probability matrices are Kronecker products of multiple copies of $P_{Y|X}$ and are therefore also left invertible. This implies (by analogy to the argument for $M = 2$) that the joint pmf matrix of $X^{(I)}$ and $X^{(J)}$ has also only one non-zero singular value, which means that it must be the outer product of two vectors, hence $X^{(I)}$ and $X^{(J)}$ are independent. But since this holds for arbitrary sets $\mathcal{I}$ and $\mathcal{J}$, we have that $X_1, \ldots, X_M$ must be independent. The fact that they are also identically distributed can then be derived by considering X_i and X_j for all $i \neq j$, and using the same argument as in the case $M = 2$. □

Proof of Lemma 3.11. Consider the matrix $P_{\mathbf{S}|\hat{\mathbf{S}}}$. It may be expressed as $P_{\mathbf{S}|\hat{\mathbf{S}}} = P_{\mathbf{S}|\mathbf{X}} P_{\mathbf{X}|\mathbf{Y}} P_{\mathbf{Y}|\hat{\mathbf{S}}}$. The distortion measure has to be averaging, which, by Eqn. (3.75), implies that $P_{\mathbf{S}|\hat{\mathbf{S}}} = P_{S|\hat{S}} \otimes \ldots \otimes P_{S|\hat{S}}$ (M terms). By assumption, $P_{S|\hat{S}}$ is nonsingular. This is true if and only if $P_{\mathbf{S}|\hat{\mathbf{S}}}$ is also nonsingular. Hence, $P_{\mathbf{S}|\mathbf{X}}$ and $P_{\mathbf{Y}|\hat{\mathbf{S}}}$ must be full-rank matrices.

Moreover, using the requirement that $I(\mathbf{X}; \mathbf{Y}) = I(\mathbf{S}; \hat{\mathbf{S}})$, we now infer that $P_{\mathbf{S}|\mathbf{X}}$ and $P_{\mathbf{Y}|\hat{\mathbf{S}}}$ have to be permutation matrices. Consider the mutual information

$$\begin{aligned} I(\mathbf{S}, \mathbf{X}; \mathbf{Y}) &= I(\mathbf{S}; \mathbf{Y}) + I(\mathbf{X}; \mathbf{Y}|\mathbf{S}) \\ &= I(\mathbf{X}; \mathbf{Y}) + I(\mathbf{S}; \mathbf{Y}|\mathbf{X}), \end{aligned} \quad (3.90)$$

where $I(\mathbf{S}; \mathbf{Y}|\mathbf{X}) = 0$ since $\mathbf{S} \to \mathbf{X} \to \mathbf{Y}$ is a Markov chain, and hence $I(\mathbf{X}; \mathbf{Y}) = I(\mathbf{S}; \mathbf{Y}) + I(\mathbf{X}; \mathbf{Y}|\mathbf{S})$. To satisfy $I(\mathbf{X}; \mathbf{Y}) = I(\mathbf{S}; \mathbf{Y})$, it is therefore necessary that $I(\mathbf{X}; \mathbf{Y}|\mathbf{S}) = H(\mathbf{X}|\mathbf{S}) - H(\mathbf{X}|\mathbf{Y}, \mathbf{S}) = 0$. This is true if and only if $\mathbf{X}$ and $\mathbf{Y}$ are independent given $\mathbf{S}$. Hence consider the joint distribution matrix $P_{\mathbf{Y},\mathbf{X}|\mathbf{S}=\mathbf{s}}$.

Denoting by $P_{\mathbf{X}|\mathbf{S}=\mathbf{s}}$ a diagonal matrix with entries $p(\mathbf{x}|\mathbf{s})$ along the diagonal, we can write

$$P_{\mathbf{Y},\mathbf{X}|\mathbf{S}=\mathbf{s}} = P_{\mathbf{Y}|\mathbf{X}} P_{\mathbf{X}|\mathbf{S}=\mathbf{s}}. \tag{3.91}$$

For $\mathbf{X}$ and $\mathbf{Y}$ to be independent given $\mathbf{S}$, the matrix $P_{\mathbf{Y},\mathbf{X}|\mathbf{S}=\mathbf{s}}$ has to have rank 1 for all $\mathbf{s}$. However, since by assumption, any set of two columns of $P_{\mathbf{Y}|\mathbf{X}}$ are linearly independent, the diagonal matrix $P_{\mathbf{X}|\mathbf{S}=\mathbf{s}}$ can have at most one non-zero entry, hence $p(\mathbf{x}|\mathbf{s}) = 1$ for exactly one of the $\mathbf{x}$. Hence, the matrix $P_{\mathbf{X}|\mathbf{S}}$ has only ones and zeros as entries. Moreover, $p(\mathbf{s}) > \mathbf{0}$ for all $\mathbf{s}$ and $P_{\mathbf{S}|\mathbf{X}}$ is invertible, which implies that $P_{\mathbf{X}|\mathbf{S}}$ is also invertible. Hence $P_{\mathbf{X}|\mathbf{S}}$ is a permutation matrix (and so is $P_{\mathbf{S}|\mathbf{X}}$).

By analogy, consider

$$\begin{aligned} I(\mathbf{S};\mathbf{Y},\hat{\mathbf{S}}) &= I(\mathbf{S};\mathbf{Y}) + I(\mathbf{S};\hat{\mathbf{S}}|\mathbf{Y}) \\ &= I(\mathbf{S};\hat{\mathbf{S}}) + I(\mathbf{S};\mathbf{Y}|\hat{\mathbf{S}}). \end{aligned} \tag{3.92}$$

To satisfy $I(\mathbf{S};\mathbf{Y}) = I(\mathbf{S};\hat{\mathbf{S}})$, we need that $I(\mathbf{S};\mathbf{Y}|\hat{\mathbf{S}}) = H(\mathbf{Y}|\hat{\mathbf{S}}) - H(\mathbf{Y}|\hat{\mathbf{S}},\mathbf{S}) = 0$. This is true if and only if $\mathbf{S}$ and $\mathbf{Y}$ are independent given $\hat{\mathbf{S}}$. By analogy to the first half of the proof, this implies that $P_{\mathbf{Y}|\hat{\mathbf{S}}}$ can have only zero or one as entries. Since moreover, it is invertible, it follows that $P_{\mathbf{Y}|\hat{\mathbf{S}}}$ is a permutation matrix.

To conclude the proof, note that permutation matrices represent bijective mappings. □

Lemma 3.12 *For any matrix* $A \in M(n \times m)$, $\operatorname{rank}(A) \leq 1$ *if and only if*

$$A_{ij}A_{kl} - A_{kj}A_{il} = 0, \tag{3.93}$$

for all $1 \leq i,k \leq n$ *and* $1 \leq j,l \leq m$.

Proof. $rank(A) \leq 1 \Leftrightarrow A = xy^T$ for some vectors x and y. But then, the forward part is immediate.

For the reverse, we show that any two rows of A are dependent. Pick row i and row k, and form the $2 \times m$ submatrix A'. The rows of A' are independent if and only if we can find two columns j and l that are linearly independent. This happens if and only if the 2×2 submatrix A'' containing only columns j and l of A' has full rank. However, by assumption, this matrix has determinant zero. Hence any two rows of A are dependent, and the rank cannot be larger than 1. □

Chapter 4

An Intermezzo — Uncoded Transmission and Biological Systems

This chapter illustrates an area where some of the theory developed in Chapters 2 and 3 may be of particular interest: communication in biological systems. Uncoded transmission is simple and operates at no delay. Neural communication, for example, may be very interested in such transmission schemes. Rather than aiming at an exhaustive treatment, this chapter merely outlines one line of thought, and it does so by reinterpreting a known example from neural communication, taken from [29].

We first review this example in its original formulation, taking a point of view that we call the *capacity perspective*. Here, the measures of interest are rate and mutual information. Thereafter, we suggest an alternative point of view that we call the *energy vs. accuracy perspective*. This is based on the results presented in Chapters 2 and 3, and its measures of interest are the *accuracy* of the representation of the stimulus, and the *cost* of having this accuracy.[1]

To provide a precise and complete historic treatment would be beyond the framework of the present thesis. References can be found e.g. in [12].

4.1 The Capacity Perspective

One attractive way of applying information theory to sensory processing is to measure the statistics of the involved signals, and to plug these measurements into Shannon's capacity formula for the Gaussian channel, see for example the figure in [12, p. 157]. This rationale is motivated by the desire to "establish an absolute scale for neural performance" [12, p. 267, second paragraph].

Such a computation can be performed for spiking neurons, as in [12]. We prefer here to study in more detail another neural system, the so-called graded-

[1]The ideas discussed in this chapter were presented, in part, in [40] and [41].

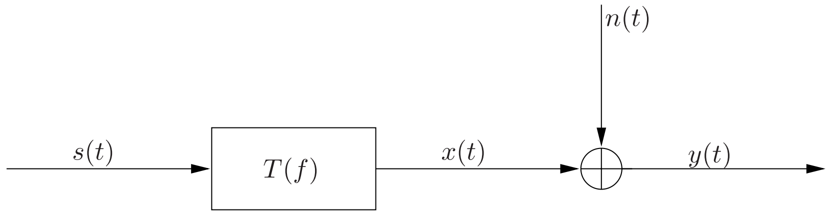

Figure 4.1: The system considered in [29].

potential synapse. A capacity calculation of the type of [12] was presented in [29]. The following example is a simplified version of that publication, concerning the early stages of the blowfly visual system.

Example 4.1 (graded potential synapse) *Consider the following line of thought, presented in [29]:*

1. *The system* **model** *is illustrated in Figure 4.1. The stimulus is first processed through a transfer function $T(f)$. Thereafter, (colored) Gaussian noise is added. Hence, the power spectral density $S_Y(f)$ of the output signal $y(t)$ is*

$$S_Y(f) \quad = \quad |T(f)|^2 S_S(f) + S_N(f), \tag{4.1}$$

where $S_S(f)$ is the power spectral density of the stimulus, $T(f)$ is the transfer function, and $S_N(f)$ is the power spectral density of the noise $n(t)$.

2. *The system* **parameters** *are determined from measurements. In particular, we need the transfer function and the power spectral density of the additive noise.*[2] *To achieve this goal, the* same *stimulus $s(t)$ is presented many times. The* average *response is taken to have power spectral density $|T(f)|^2 S_S(f)$. Since $S_S(f)$ is known, the average response permits one to determine*

$$T(f). \tag{4.2}$$

The deviations *from this average response are taken to be the noise, yielding the noise power spectral density*

$$S_N(f). \tag{4.3}$$

3. *Shannon's capacity formula for the system of Figure 4.1 is*

$$C \quad = \quad \max_{S_S(f)} \int_0^\infty df \log_2 \left(1 + \frac{|T(f)|^2 S_S(f)}{S_N(f)} \right). \tag{4.4}$$

[2] In [29], it is argued that the Gaussianity of the signals also follows from the measurements. However, [29] does not provide sufficient detail to be sure that this Gaussianity is not merely a consequence of the central limit theorem. Therefore, we here include the Gaussianity with the model parameters.

4. *The maximization is taken over all $S_S(f)$ that satisfy*

$$\int_0^\infty df S_S(f) \quad \leq \quad 0.1. \tag{4.5}$$

Hence, the maximizing $S_C(f)$ is found by water-filling on the power spectral density

$$\frac{S_N(f)}{|T(f)|^2}. \tag{4.6}$$

The result of the maximization is $C = 1650$ bits per second.

5. *It is argued: "The graded responses of these non-spiking cells transmit as much as 1650 bits per second."*

As this example demonstrates, the capacity perspective does provide an absolute scale for neural performance in the sense of Point 5. of Example 4.1, i.e., the result $C = 1650$ bits per second. A similar capacity result for the bullfrog auditory system can be found in [12, p. 185]. The remaining question is the significance of this absolute scale. We discuss some of this below in Section 4.1.1.

Another capacity perspective has been investigated by Berger [19]. His approach is almost exclusively based on models rather than on experimental data. More precisely, he considers "diffusive communication" across a synaptic gap: Initially, the gap is filled with a medium containing different types of molecules. Each type of molecule is present with its characteristic concentration. In other words, there is a *concentration distribution* over the types of molecules. When a spike arrives at one edge of the gap, the corresponding nerve cell injects molecules into the gap, altering the concentration distribution. By diffusion, this new concentration distribution can also be sensed at the far end of the gap. Hence, such a mechanism permits the spike to cross the synaptic gap.

In more detail, Berger's goal in [19] is to determine the capacity of this diffusion channel. The synaptic gap medium contains a total of b "chemical symbols" (e.g., molecules). There are M different types of chemical symbols. The simplest model for diffusion is that in each time slot, the cell at the far end of the gap grabs uniformly at random one of the b chemical symbols in the medium, and the cell at the near end of the gap injects a new chemical symbol into the gap. Hence, there are always b chemical symbols in the gap. A capacity result was given in [19], and referenced to yet unpublished work by Zhang and Berger. Asymptotically in b, the capacity behaves at least like $b^{-2/3}$. Obviously, the largest capacity is achieved by setting $b = 1$; this is simply a lossless channel. So will nature use that channel? In [19], Berger argues that such a channel is "expensive to build and run. Unless you are in a hurry (and sometimes you are), bits per energy matters more than bits per second." In a sense, Section 4.2 below formalizes and extends this last point. However, we trade off *accuracy* versus energy, rather than *bits*.

4.1.1 Critical review

While the capacity perspective is both elegant and seducing, it suffers from certain limitations. In particular, we note the following points:

1. The fact that from the measurements, one evaluates

$$C = \max_{S_C(f)} \int_0^\infty df \log_2 \left(1 + \frac{|T(f)|^2 S_C(f)}{N(f)}\right), \tag{4.7}$$

   does not mean that the system operates at C bits per second.

   Rather, it means that C bits per second are achievable by means of a generally very complex encoder/decoder. More explicitly, a system with a higher C does not necessarily provide a better stimulus reconstruction; this implication holds only if capacity-achieving encoding and decoding is used.

   It seems unlikely that a biological system has the kind of encoder/decoder pair that makes it possible to transmit C bits per second.

   Besides the fact that inspection does not seem to support the presence of such an encoder/decoder pair, it is also reasonable to assume that biological systems cannot afford the delay that communication very close to Shannon capacity implies.

2. More fundamentally, the Shannon capacity is the largest number of bits per second at which it is possible to make the *error probability* P_e *go to zero* (as the delay goes to infinity). It seems probable that neural communication (as opposed to communication via DNA) is not designed to make the error probability go to zero at such a high cost in terms of delay.

4.2 The Energy vs. Accuracy Perspective

Instead of maximizing the *number of bits* per second under the constraint that the error probability P_e must go to zero at some fixed cost (as in the capacity perspective), a more reasonable goal is to maximize the *accuracy* of the reconstruction at the same fixed cost.

As discussed in Chapters 1, 2 and 3 of this thesis, information theory also provides tools to attack this problem. The theory discussed in this thesis applies to discrete-time systems. In contrast, by nature, biological systems must be expected to operate in continuous time, which is also reflected by Example 4.1. For neural communication, there is a finite speed of reaction, and a finite transition time for electro-chemical processes. For these reasons, we assume in the sequel that one can model neural communication systems as *band-limited* systems, or that such a model is close enough to reality. Under this assumption, the continuous-time signals can be represented exactly by a sequence of samples, i.e., by a discrete-time signal. Hence, under the assumption that the involved systems are band-limited, a discrete-time consideration is sufficient. Our general framework can be outlined as follows:[3]

The source S with distribution $p(s)$ is communicated to a destination via a channel with conditional distribution $p(y|x)$. The encoder maps k source symbols onto m channel input symbols, and the decoder maps m channel output symbols onto k source reconstruction symbols. This is illustrated in Figure 4.2. The input X to the channel is any symbol x of the alphabet $\mathcal{X}$, and the cost

Figure 4.2: Conceptual framework of the energy vs. accuracy perspective.

(the "energy") of using the particular symbol x is denoted by $\rho(x)$. The average expected cost per channel use is

$$\Gamma = \frac{1}{m}\sum_{j=1}^{m} E\rho(X_j). \tag{4.8}$$

The channel output sequence Y^m is processed into an estimate $\hat{S}^k$ of the source sequence S^k. A distortion function $d(s, \hat{s})$ specifies how good or bad it is if the source letter s is reproduced by the reconstruction letter $\hat{s}$. The average distortion per source symbol is

$$\Delta = \frac{1}{k}\sum_{j=1}^{k} Ed(S_j, \hat{S}_j). \tag{4.9}$$

The source-channel communication problem is the trade-off between cost Γ and distortion Δ.

This does not only seem to be a more reasonable goal for a neural communication system; it is also well known that an optimal cost-distortion trade-off[4] can be achieved at low complexity in certain cases. This was discussed in detail in Chapters 2 and 3.[5]

More precisely, from Theorem 2.5 (and its extensions presented in Chapter 3), the optimum cost-distortion trade-off can be achieved at low complexity and delay, provided that the source and the channel are favorably matched. This may be too much to ask in a man-made communication system where the source and the channel are picked independently of each other. However, evolution

[3]In an attempt to make this chapter self-contained at the level of concepts, we briefly review the main ideas discussed in Chapter 1.

[4]Recall that *optimal* means: the best performance irrespective of delay and complexity. See Definition 1.5.

[5]When we presented the main results of Chapter 2 at the *International Symposium on Information Theory* (ISIT) in Sorrento in June 2000 [42], Prof. T. Berger suggested to us to apply our results to neural communication. He generously provided us with the models of neural communication he was studying [18], and later with his estimation and capacity considerations, mentioned above in Section 4.1 [19]. This was followed by a series of mutually inspiring discussions at the 2001 DSC Research Day at EPFL, and at ISIT 2001.

could have "designed" the source and the channel in a matching fashion. We are excited about the possibility that nature may indeed have followed such a path. It is in this spirit that we propose the following thoughts:

Example 4.2 *Starting from the channel model proposed in [29], illustrated in Figure 4.1, we select the channel to be an additive Gaussian noise channel with power spectral density given by* $S_N(f)$ *as in (4.3). We select the transfer function* $T(f)$ *of Figure 4.1 to be the encoder* F*, and we take the decoder* G *to be the identity function. This defines a system of the kind shown in Figure 4.2, and permits to apply the techniques developed in Chapters 2 and 3 of this thesis.*

For our arguments, we simplify Example 4.1 slightly by assuming that all processes are band-limited, and that therefore, a discrete-time consideration is sufficient. To make matters precise, we simplify this further by considering finite blocks of length k*: the source emits blocks of length* k*. The blocks are independent of one another, and every block has the same distribution. Similarly and synchronously, the channel noise emits blocks of length* k*, which are again assumed independent and identically distributed. Moreover, in the simplified model, the transfer function* $T(f)$ *is a* $k \times k$ *matrix, denoted by* T*. This is illustrated in Figure 4.3. We denote the* k*-dimensional Gaussian distribution in*

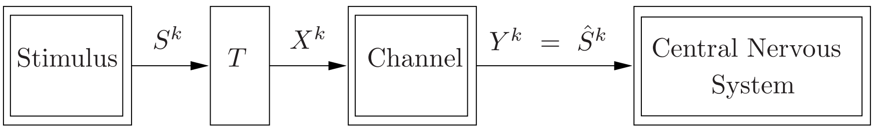

Figure 4.3: The simplified version of the model of Figure 4.1, as used for Example 4.2.

the (vector-valued) variable x^k *whose mean is given by the vector* m *(of length* k*) and whose covariance matrix is* Σ_X*, by*

$$\mathcal{N}(m, \Sigma_X)(x^k) = \frac{1}{(\sqrt{2\pi})^k \sqrt{\det \Sigma_X}} \exp(-\frac{1}{2}(x^k - m)^T \Sigma_X^{-1}(x^k - m). \quad (4.10)$$

We will also sometimes use the shorter notation $\mathcal{N}(m, \Sigma_X)$ *when the variable is not needed explicitly. The covariance matrices of the source, the channel noise, and the channel output will be denoted by* $\Sigma_S, \Sigma_N,$ *and* Σ_Y*, respectively. Theorem 3.3 says that if*

$$\begin{aligned} \rho^{(k)}(x^k) &= D(p(y^k|x^k)||p(y^k)) \\ &= D(\mathcal{N}(x^k, \Sigma_N)||\mathcal{N}(0, \Sigma_Y)), \end{aligned} \quad (4.11)$$

and, recalling that $p(s^k|\hat{s}^k)$ *is Gaussian with mean* $E[S^k|\hat{S}^k = \hat{s}^k]$ *and variance* $Var(S^k|\hat{S}^k = \hat{s}^k)$*,*

$$\begin{aligned} d^{(k)}(s^k, \hat{s}^k) &= -\log_2 p(s^k|\hat{s}^k) \\ &= -\log_2 \mathcal{N}(T\Sigma_S\Sigma_Y^{-1}\hat{s}, \Sigma_S - \Sigma_S T^T \Sigma_Y^{-1} T \Sigma_S)(s^k) \\ &= (s - P\hat{s})^T \Sigma^{-1}(s - P\hat{s}) + const., \end{aligned} \quad (4.12)$$

where $P = T\Sigma_S\Sigma_Y^{-1}$ *and* $\Sigma = \Sigma_S - \Sigma_S T^T \Sigma_Y^{-1} T \Sigma_S$, *then the overall system is optimal.*[6]

In order to gain insight into the structure of the cost and distortion functions according to Equations (4.11) and (4.12), respectively, the next goal is to rewrite ρ *and* d *in terms of well-known quantities, such as the power spectral densities of the involved signals. For simplicity, we assume that the covariance matrices* Σ_S, Σ_N *and* Σ_Y *as well as the encoder matrix* T *have identical eigenvectors. This is a reasonable assumption: If the source and the noise process are stationary (and here we mean stationary* inside *the blocks of length* k*), then their covariance matrices are Toeplitz. Moreover, if the encoder* T *can indeed be described by a transfer function (as is assumed in [29]), then the corresponding matrix is also Toeplitz. But in the appropriate sense, as* $k \to \infty$, *the eigenspace of a Toeplitz matrix is the Fourier space. This can be made precise, see e.g. [49].*

Under the assumption of equal eigenvectors of Σ_N *and* Σ_Y, *we can evaluate the above formula for* ρ. *Using the auxiliary Lemma 4.1 (see Appendix 4.A), we can rewrite (4.11) as*

$$\rho^{(k)}(x^k) = \sum_{j=1}^{k} D(\mathcal{N}(X_j, \lambda_{N,j}) || \mathcal{N}(0, \lambda_{Y,j})), \tag{4.13}$$

with $X^k = Qx^k$, *where* Q *is the matrix of eigenvectors of the covariance matrix* Σ_N, *and hence also of* Σ_Y. *Each term in the sum can be evaluated like in Example 2.5, yielding*

$$\rho^{(k)}(x^k) = \sum_{j=1}^{k} \alpha_j X_j^2, \tag{4.14}$$

where we have subtracted an appropriate constant. To interpret this formula, note that as $k \to \infty$, *the sequence* X^k *converges (in an appropriate sense) to the spectrum of the sequence* x^k, *see e.g. [49]. This suggests that in the formalism of Example 4.1, the cost of the signal* $x(t)$ *can be expressed in the spectral domain as*

$$\rho(X(f)) = \int df \alpha(f) |X(f)|^2 \tag{4.15}$$

for a function $\alpha(f)$ *that is essentially the spectrum of the noise.*

Similarly, under the assumption of equal eigenvectors of Σ_S, Σ_N, Σ_Y, *and* T, *the distortion function can be simplified along the following lines:*

$$\begin{aligned} d^{(k)}(s^k, \hat{s}^k) &= (s - P\hat{s})^H \Sigma^{-1} (s - P\hat{s}) \\ &= \sum_{j=1}^{k} \beta_j (S_j - \gamma_j \hat{S}_j)^2. \end{aligned} \tag{4.16}$$

with $S^k = Qs^k$ *and* $\hat{S}^k = Q\hat{s}^k$, *where* Q *is the matrix of eigenvectors of the covariance matrix* Σ_S, *and hence also of* Σ_N, Σ_Y *and* T. *Note that this implies*

[6] Recall that Theorem 3.3 does *not* establish that Equations (4.11) and (4.12) are the only optimal choices.

that Q is also the matrix of eigenvectors of the matrices P and Σ. This formula can be interpreted along the same lines as the formula for $\rho(x)$. In particular, in the appropriate sense, the sequences S^k and $\hat{S}_k$ converge again to the spectra of s^k and $\hat{s}^k$, respectively. This suggests that in the formalism of Example 4.1, we can express the distortion between a certain input signal $s(t)$ and its corresponding reconstruction $\hat{s}(t) = y(t)$ (recall that the decoder is taken to be the identity function) in the spectral domain as

$$d(S(f), Y(f)) = \int df \beta(f) |S(f) - \gamma(f) Y(f)|^2 \tag{4.17}$$

for a function $\gamma(f)$ that is essentially the spectrum of the encoder, and a function $\beta(f)$ that is the product of the square of the spectrum of the encoder with the spectrum of the noise.[7]

In conclusion, if we take for granted the model of a biological system given in [29], then the cost function $\rho(x)$ and the distortion measure $d(s, y)$ given in (4.14) and (4.16) are the ones for which the system is optimal. In words, no processing at the input or at the output can lead to a better average cost versus average distortion trade-off, regardless of complexity.

4.2.1 Critical review

As illustrated by Example 4.2, our approach is also based on statistical measurements, just like the capacity perspective of Section 4.1. But while the latter yields an "absolute" result, e.g., 1650 bits per second, our perspective yields a channel input cost function and a distortion measure. This is one step short of a final result: it remains to be argued that these functions make (physiological) sense. For instance, in Example 4.2, the last step is to argue that on the graded-potential synapses connecting the photo-receptor to the large monopolar cells in the blowfly visual system [29], it makes sense that the cost function is a weighted energy (with weighting in the frequency domain). Such arguments are beyond the scope of this thesis.

4.3 Further Perspectives

4.3.1 The redundancy perspective

Another application of information theory to sensory processing takes into account coding. The flavor of this perspective is quite different: The goal is to devise a sensible design principle for neural coding. This design principle is the *redundancy* of the stimulus representation, e.g., the redundancy in the neural signals that represent an image. The considered notion of redundancy is inspired by information theory:

$$\theta = 1 - H(M)/R, \tag{4.18}$$

[7] The precise analysis of the continuous-time problem is not pursued further.

where R is the rate[8] of the representation, and $H(M)$ is the entropy of the stimulus.

The problem of determining codes that minimize redundancy is known to be difficult. Therefore, simpler functionals were studied. One approach can be outlined as follows: Consider an input layer of n neurons whose activities are described by $L_1, \ldots, L_n$. The corresponding activities in the output layer are described by $O_1, \ldots, O_l$, and the response of the output neuron is assumed to be some general function,

$$O_i = K_i(L_1, \ldots, L_n).$$

Using this, a simplified redundancy function that retains some of the basic features is defined as

$$E[K_i] = \sum_{i=1}^{l} H(O_i) - 2\rho(H(O_1, \ldots, O_l) - H(L_1, \ldots, L_n)).$$

The goal is to optimize this over all possible recoding functions K_i. A further simplification is to consider only simple coding functions K_i, such as linear mappings.[9]

The line of thought furnished by this approach can be outlined as follows: Starting from statistical measurements, one determines the optimal coding. Then, it is verified (by "inspection", by experimentation, etc) that the actual coding is "close" to the optimum. The goal of this approach is to argue "that much of the processing in the early levels of sensory pathways might be geared towards building efficient representations of sensory stimuli in an animal's environment" [16]. A more detailed treatment of these ideas along with the relevant references is given in the review article [16].

Critical review

One key advantage of the redundancy perspective has been described as follows [12, p. 268, last paragraph]: "We can generalize this definition to include whole arrays of neurons, not just two." Recall that while the true information theoretic perspectives (as in Sections 4.1 and 4.2) are somewhat less ad-hoc, they are difficult to extend to the network case: only very few results in network information theory are known to date.

However, redundancy is an ad-hoc concept: There is no strong evidence that biological systems are interested in removing redundancy in general, or, as Atick puts it in [16, p. 220, second paragraph]: "In higher animals we feel it is more likely that cognitive benefits are the driving force towards efficient representation".

[8] In [16], Atick calls this the "capacity" of the representation.

[9] However, in contrast to Section 4.2, those simple codes are *not* argued to be information-theoretically optimal. Rather, *within* the class of simple codes, the one minimizing the redundancy R, or more precisely, the simplified functional $E[K_i]$, is determined.

4.4 Summary and Conclusions

In Chapter 3, we argued that optimal communication can be characterized by *matched measures*. Certain source/channel pairs are already favorably matched in such a way that an optimal match of the measures can be achieved by little coding, if any. The hypothesis of this chapter is that such a favorable match occurs in biological systems. However, an experimental "proof" of the hypothesis is beyond the scope of this thesis.

In particular, we illustrated how the results of Chapter 3 can be used to understand, from a different point of view, certain experiments that were published in the literature: rather than plugging the measurement data into Shannon's capacity function, and claim that the result is an "absolute scale for neural performance", we determine the framework in which the neural communication system achieves an information-theoretically optimal trade-off between the cost of transmission and the quality of reconstruction. This framework consists of the measured statistics, the coding functions, and the input cost function of the channel and the distortion measure.

This raises the future challenge of understanding the resulting cost and distortion measures from a physiological perspective — of associating these measures with observable physiological processes.

Appendix 4.A Auxiliary Lemma for Gaussian Divergence

Let the k-dimensional Gaussian distribution with mean m and covariance matrix Σ_X be denoted by $\mathcal{N}(m, \Sigma_X)$. This is a slightly simplified notation of (4.10). Then, the following is true:[10]

Lemma 4.1 *Suppose that Σ_Z and Σ_Y are covariance matrices of dimensions N and with the* same *eigenvectors $q_1, \ldots, q_N$, and that x is a constant vector of length N. Then,*

$$D(\mathcal{N}(x, \Sigma_Z)||\mathcal{N}(0, \Sigma_Y)) \quad = \quad \sum_{k=1}^{N} D(\mathcal{N}(X_k, \lambda_{Z,k})||\mathcal{N}(0, \lambda_{Y,k})), \qquad (4.19)$$

where $\lambda_{Z,k}$ and $\lambda_{Y,k}$ are the eigenvalues of Σ_Z and Σ_Y, respectively, corresponding to eigenvector q_k, and

$$X \quad = \quad Qx, \qquad (4.20)$$

where Q is the matrix whose columns are the eigenvectors $q_1, \ldots, q_N$.

Proof. The lemma could be established from the chain rule for divergence, see e.g. [3, Thm. 2.5.3]. We here prefer to simply write out the definition,

$$\begin{aligned} &D(\mathcal{N}(x, \Sigma_Z)||\mathcal{N}(0, \Sigma_Y)) \\ &= \frac{1}{(\sqrt{2\pi})^N \sqrt{\det \Sigma_X}} \int \cdots \int e^{-\frac{1}{2}(y-x)^T \Sigma_Z^{-1}(y-x)} \\ &\qquad \log_2 \frac{\det \Sigma_Y}{\det \Sigma_Z} e^{-\frac{1}{2}(y-x)^T \Sigma_Z^{-1}(y-x) + \frac{1}{2} y^T \Sigma_Y^{-1} y} dy_1 \cdots dy_N. \end{aligned}$$

Introduce the new variables

$$Y \quad = \quad Qy.$$

Since Q is a unitary transform, the determinant of the Jacobian is one. For convenience, we moreover replace

$$X \quad = \quad Qx.$$

Using this, the integral can be written in the new coordinates as

$$\begin{aligned} &\int \cdots \int e^{-\frac{1}{2}(Q^T Y - x)^T \Sigma_Z^{-1}(Q^T Y - x)} \\ &\qquad \log_2 \frac{\det \Sigma_Y}{\det \Sigma_Z} e^{-\frac{1}{2}(Q^T Y - x)^T \Sigma_Z^{-1}(Q^T Y - x) + \frac{1}{2} Y^T Q \Sigma_Y^{-1} Q^T Y} dY_1 \cdots dY_N \\ &= \int \cdots \int e^{-\frac{1}{2}(Y - X)^T Q \Sigma_Z^{-1} Q^T (Y - X)} \\ &\qquad \log_2 \frac{\det \Sigma_Y}{\det \Sigma_Z} e^{-\frac{1}{2}(Y - X)^T Q \Sigma_Z^{-1} Q^T (Y - X) + \frac{1}{2} Y^T Q \Sigma_Y^{-1} Q^T Y} dY_1 \cdots dY_N. \end{aligned}$$

[10] While we feel that this fact must be noted somewhere, we failed to discover an appropriate reference. Since the proof is rather straightforward, we include it here.

By assumption on the matrix Q, both the matrix $Q\Sigma_Z^{-1}Q^T$ and the matrix $Q\Sigma_Y^{-1}Q^T$ are diagonal matrices whose entries are the reciprocals of the corresponding eigenvalues. Hence,

$$\begin{aligned}
&\int \cdots \int e^{-\frac{1}{2}(Y-X)^T Q\Sigma_Z^{-1}Q^T(Y-X)} \\
&\qquad \log_2 \frac{\det \Sigma_Y}{\det \Sigma_Z} e^{-\frac{1}{2}(Y-X)^T Q\Sigma_Z^{-1}Q^T(Y-X)+\frac{1}{2}Y^T Q\Sigma_Y^{-1}Q^T Y} dY_1 \cdots dY_N \\
&= \int \cdots \int e^{-\frac{1}{2}\sum_{k=1}^N \frac{(Y_k-X_k)^2}{\lambda_{Z,k}}} \\
&\qquad \left(\sum_{k=1}^N \left(\log_2 \frac{\lambda_{Y,k}}{\lambda_{Z,k}} \right) \left(-\frac{(Y_k - X_k)^2}{\lambda_{Z,k}} + \frac{Y_k^2}{\lambda_{Y,k}} \right) \right) dY_1 \cdots dY_N \\
&= \sum_{k=1}^N \left(\log_2 \frac{\lambda_{Y,k}}{\lambda_{Z,k}} \right) \int \cdots \int e^{-\frac{1}{2}\sum_{k=1}^N \frac{(Y_k-X_k)^2}{\lambda_{Z,k}}} \\
&\qquad \left(-\frac{(Y_k - X_k)^2}{\lambda_{Z,k}} + \frac{Y_k^2}{\lambda_{Y,k}} \right) dY_1 \cdots dY_N. \qquad (4.21)
\end{aligned}$$

At this point, we bring the determinant term back in, noting that it can be split according to

$$\frac{1}{(\sqrt{2\pi})^N \sqrt{\det \Sigma_Z}} = \prod_{k=1}^N \frac{1}{\sqrt{2\pi}\sqrt{\lambda_{Z,k}}}.$$

With this, in each summand of (4.21), $N-1$ of the integrals yield one (since they are simply over a Gaussian distribution), and only the integral over Y_k remains, yielding

$$\begin{aligned}
&D(\mathcal{N}(x, \Sigma_Z)||\mathcal{N}(0, \Sigma_Y)) \\
&= \sum_{k=1}^N \frac{1}{\sqrt{2\pi}\sqrt{\lambda_{Z,k}}} \int e^{-\sum_{k=1}^N \frac{(Y_k-X_k)^2}{\lambda_{Z,k}}} \\
&\qquad \left(-\left(\log_2 \frac{\lambda_{Y,k}}{\lambda_{Z,k}} \right) \left(\frac{(Y_k - X_k)^2}{\lambda_{Z,k}} - \frac{Y_k^2}{\lambda_{Y,k}} \right) \right) dY_k \\
&= \sum_{k=1}^N D(\mathcal{N}(X_k, \lambda_{Z,k})||\mathcal{N}(0, \lambda_{Y,k})),
\end{aligned}$$

which concludes the proof. □

Chapter 5

Uncoded Transmission In Networks

This chapter explores some of the significance of uncoded transmission in a network context. Since the separation paradigm does not lead to optimal design in general, uncoded transmission can "beat capacity" in certain situations; i.e. it can achieve a better performance than source coding followed by capacity-approaching codes. Section 5.1 gives a brief introduction to network information theory, which serves to quote the relevant prior art and to relate the contributions of this chapter to the known results. Then, in Section 5.2, the results of Chapter 3 are applied to a single-source broadcast situation where one can indeed beat capacity in the above sense. New instances of such behavior are determined.[1] Section 5.3 gives an example of a multiple description network where uncoded transmission performs optimally.

The main part of the chapter, Section 5.4, concerns large Gaussian relay networks. We determine the capacity of a particular large Gaussian relay network in the limit as the number of relays tends to infinity. The upper bounds follow from a cut-set argument, and the lower bound follows from an argument involving uncoded transmission. We prove that in many cases of interest, upper and lower bounds coincide in the limit as the number of relays tends to infinity. Hence, this section gives one more example where the cut-set bound is achievable, and one more example where uncoded transmission achieves optimal performance.

To illustrate our findings, we first apply them to a sensor network situation. The comparison of our results with the CEO problem leads to a new instance of the fact that the source-channel separation paradigm does *not* extend to networks in general. Then, we show how to extend our approach to include certain ad-hoc wireless networks, which leads to a capacity result: When all nodes act purely as relays for a single source-destination pair, capacity grows with the logarithm of the number of nodes.[2]

[1]Section 5.2 is taken from [43].
[2]Section 5.4 is (up to editorial changes) identical to [45].

5.1 Network Performance Analysis

By means of well-known counterexamples, we pointed out in Chapter 1 that the separation paradigm does not extend to networks: Systems designed according to Theorem 1.8 cannot be argued to be optimal in general. This leads to a double perspective on network performance analysis: On the one hand, one can focus on the problem of achieving the best possible cost-distortion trade-off, and hence study the joint source-channel communication problem. On the other hand, one can focus on systems designed according to the separation theorem. Those will not generally achieve the best possible cost-distortion trade-off, but they are of independent interest precisely for their modular structure.

In this sense, both the source-channel *and* the rate-distortion/capacity perspectives are interesting points of view. As we will see, uncoded transmission is helpful in both cases. In the following subsections, we outline and quote a part of the prior art. This is needed both to establish the new results presented in the sequel, as well as to relate them to known results.

In particular, we first outline some known results on network joint source-channel coding; this is needed for the new results of Sections 5.2 and 5.4.6. Then, we quote certain known results on the capacity of channel networks, necessary for Section 5.4. Finally, we outline results on the rate-distortion behavior of source networks; such a result is needed in Section 5.4.6.

5.1.1 Network joint source-channel coding

As pointed out earlier (Remark 1.9), the network source-channel communication problem, just like in the point-to-point case, is a matter only of achieving the right marginal distributions, and the key difficulty is to identify the set of achievable marginals. In the point-to-point case, this set is characterized by Theorem 1.4. For the general network situation, the answer is not known.

In the absence of a general theory, and since the combination of rate-distortion and capacity does not lead to optimal cost-distortion trade-offs, an intermediate quest is for interesting examples of network joint source-channel codes; in particular, examples that outperform the separation-based approach.

A well-known case, described in Example 1.3, involves a simple binary multiple-access scenario with dependent binary sources. Using our results for the relay channel derived in Section 5.4, we provide in Section 5.4.6 an example that is similar in spirit to Example 1.3, but involves Gaussian sources and channels, and mean-squared error distortion.

Another well-known example of this approach is Example 1.4, i.e., broadcasting a single Gaussian source to multiple users across AWGN channels. In the spirit of Chapter 2, this example is extended to sources and channels other than the Gaussian (Section 5.2). Another extension of Example 1.4 was given in [85].

The theme of uncoded transmission in networks does not only appear in these two well-known examples; it has also been investigated in other contexts, including Gaussian sources in a particular linear network of Gaussian channels [32], and binary sources in a simple broadcast network [36].

5.1.2 Capacity of channel networks

Water networks

Consider a network of water pipes. There is one source of water, and one sink where all the water has to go. Given the structure and the parameters of the network, what is the maximum source output that can be accommodated by the network?

This question was studied by Ford and Fulkerson [6]. Their basic example is illustrated in Figure 5.1. The node labeled s is the source, the node labeled

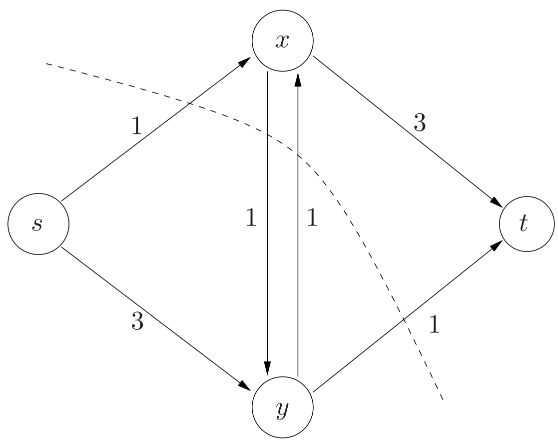

Figure 5.1: An example of the *max-flow min-cut* theorem, quoted from [6].

t is the sink, and the nodes labeled x and y are intermediate nodes. The labels on the edges are the corresponding flow capacities. A *cut* separates the network into two parts, one containing the source, the other containing the sink. The *flow* across a cut is just the sum of the capacities of the links that the cut cuts. One such cut is the dashed line in Figure 5.1. Its flow is 3, since the sum of the capacities of the edges crossing the cut from the source to the sink is 3. It is shown in [6] that the maximum flow from the source to the sink (the "max-flow") is equal to the minimum flow, minimized over all cuts (the "min-cut"),

$$\text{max-flow} \quad = \quad \text{min-cut},$$

often referred to as the *max-flow min-cut* theorem. The dashed cut in Figure 5.1 is a min-cut, as can be verified quickly by inspection, meaning that the capacity in the example is 3.

Suppose now that there are multiple sources and channels. This introduces a subtle distinction between the water network and, say, a postal delivery network: In the former, the goal is simply for the influx of water to disappear through the sinks. In the latter, it is not enough that a packet from source k arrives at some sink. Rather, it has to arrive at a specified destination. Hence, the latter is not just a trivial extension of max-flow min-cut: It is not sufficient that the pipes are wide enough; they also have to be in the right place. This problem was introduced and partially solved in [13]. An algebraic solution to the same problem has recently been proposed by Koetter and Médard [60, 61]. "Algebraic" here means that the authors define a number of multivariate polynomials whose coefficients follow from the network topology, and they show that the zeros of these polynomials constitute the answer to certain capacity questions. This has been extended in [54].

Channel networks under simplifying assumptions

The true behavior of information in a network, along with its capacity, is unknown to date. It has already been suggested that information in a network does *not* behave like water (see e.g. [13]). If we nevertheless start from the hydraulic perspective, then so-called *circuit-switched* networks (see e.g. [2]) result: for every source-channel pair, a (set of) pipes is created by putting the knobs in the right position.

As a next step, information can be constrained to behave like packets in a postal packet network. This perspective leads to so-called *packet-switched* networks (see e.g. [2]). As mentioned above, certain capacity results are known, see e.g. [60, 61].

With the advent of mobile communication, it was noticed that neither of these two perspectives is well fit to the wireless scenario. One of the main particularities of the wireless case is that every node receives the signals from every other node. This fact can be seen as a source of interference exclusively, a point of view taken e.g. in [51]. Under this simplifying assumption, the remaining problem is to schedule the transmissions in the best possible way. The basic trade-off is between low-power transmissions, which cause little interference but require the participation of a large number of intermediate nodes (so-called *hops*) on the one hand, and long-range transmission causing large interference but requiring few hops on the other hand. For the particular scenario considered in [51], it was found that as the number of nodes tends to infinity, the low-power, many-hops extreme of the above trade-off is optimal. For the same physical setup, but a different traffic scenario and different simplifying assumptions, we derive a result in Section 5.4.7.

The double temptation of the separation paradigm

It is conceptually and practically appealing to modularize a communication system. The point-to-point scenario can be modularized in the sense that source and channel codes are designed separately. There, the separation theorem guaranteed that this modularization comes without loss of optimality.

In the network context, there are two places where such a modularization could occur: On the one hand, source and channel codes could be designed separately. It has been pointed out above that this leads to suboptimal performance in general.

On the other hand, the capacity considerations described above have a basic feature in common: at an intermediate node, the incoming signal is completely decoded (i.e., the intermediate node gets to know the actual message), and subsequently re-encoded. This is reminiscent of the separation-based design in the sense that an intermediate node only needs local knowledge to operate. More precisely, an intermediate node needs to know the code used by his immediate neighbor, and the capacity of the link over which he forwards the message, but may ignore the rest of the network. In the following, we will refer to this loosely as the *modularization principle*. It is known that this does not lead to optimal network designs in general. We discuss this briefly in the next paragraph.

Capacity of channel networks

In extension of Shannon's definition of capacity in the point-to-point case, the capacity region of a channel network must be determined by optimizing over all possible coding schemes, irrespective of complexity and delay. This seems to be a hard problem. A coding scheme should really be thought of as a code *tree* for each terminal. The path through the tree is given by the information that the terminal wants to transmit *as well as* by the signals that are received by the terminal. A recent study along these lines has been presented by Kramer in [67]. There are no simple expressions for the result of this optimization in general, but for certain special cases, answers have been found, and these answers demonstrate that modularization as outlined above often leads to suboptimal performance. In the following, we discuss three reasons for this.

- **Traffic matrices that are not point-to-point.** The desired communication may not be point-to-point, but (e.g.) broadcast or multiple-access. As an example, consider the broadcast case: One transmitter sends independent information to two receivers, respectively. In line with the standard literature [3], suppose the two receivers cannot communicate with each other. Restricting to point-to-point transmission in this case means to serve only one of the destination nodes in a given frequency band and time slot. It is explained in [25] that this is suboptimal. Rather, superposition coding [25] can achieve higher rates (and was shown to be optimal in the so-called degraded case [22, 35]).

- **Relay nodes available.** Suppose that the desired communication is point-to-point, and that there is an otherwise idle node which serves as a relay for this point-to-point communication. Can this problem be modularized into two point-to-point transmissions, namely source to relay, relay to destination, without loss of optimality? It is well known that this is not the case [26]. Rather, in general, the source node's signal is received both by the relay and by the destination; thereafter, the relay helps to improve the destination's first reception. To make a subtler point, it is implicit in [26] and was shown explicitly in [38] for a simple Gaussian scenario that it is generally *suboptimal* for the relay to know the message at the end of the day, even if the channel from the source to the relay is better than the channel from the source to the destination. This provides strong evidence against the existence of a modularization principle for general networks. The main part of this chapter, Section 5.4, concerns large Gaussian relay networks. We show that as the number of relays tends to infinity, an optimal strategy is for the relays to apply uncoded forwarding. In that case, the relays cannot decode the message.

- **Collaboration between the nodes.** The wireless network situation suggests an even more optimistic perspective: the fact that every node automatically receives signals from every other node can be interpreted as an opportunity for the nodes to *collaborate* in achieving their respective goals. They can be, at the same time, relays for each other and sources/destinations of communication links. To illustrate this, consider the Gaussian multiple-access channel (MAC) (see e.g. [3, p. 378]). If the two terminals at the inputs of the MAC cannot cooperate, the sum of their respective rates is no larger than $\frac{1}{2}\log_2(2P/N)$. However, if they cooperate, they can achieve up to $\frac{1}{2}\log_2(4P/N)$, which is sometimes called the cooperative capacity of the MAC, see [3, p. 452]. A more precise (and very elegant) result has been found by Willems [111] for the network illustrated in Figure 5.2. The two transmitting terminals want to send the

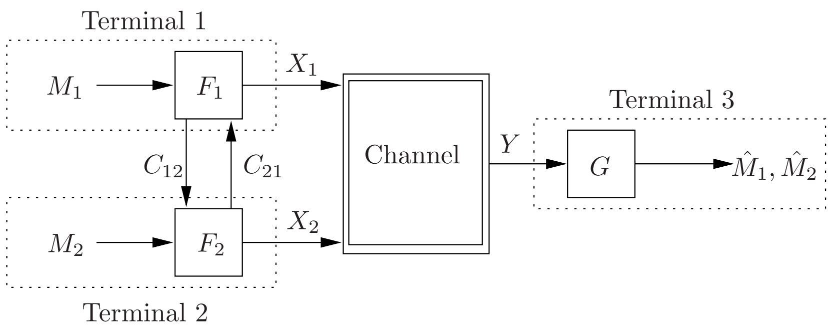

Figure 5.2: The three-terminal network studied by Willems in [111].

independent messages M_1 and M_2, respectively, to terminal 3. Moreover, terminals 1 and 2 can communicate over links of capacities C_{12} and C_{21}, respectively. The capacity region for this network is given in [111]. In the wireless case, the channels C_{12} and C_{21} in Figure 5.2 interfere with each other as well as with the multi-access channel. The capacity region is not known for this case. Achievable rates have been found in [23]. Recently, in view of wireless communications, such networks have received renewed interest, in particular scenarios that involve fading, see e.g. [70, 94, 95].

A subtler form of such collaboration is available when fading is present. To keep the argument simple, consider a multi-access situation where the channel gains change randomly over time but are known throughout the network. Then, in each time slot, only the best user transmits. This idea was first presented by Knopp and Humblet in [59], and further extended in a wealth of papers, most notably by Hanly, Tse *et al.* in e.g. [102, 53, 109].

Another form of collaboration occurs in networks with mobile nodes. Here, a part of the communication distance can be provided by the mobility of the nodes. However, the probability that one node meets its desired communication partner itself is small, at least in uniformly random models of mobility. This can be helped if the node uses many other nodes as relays: the chance that at least one of those relay nodes meets the desired destination can be shown to be large for a class of mobility patterns. This was shown in [50].

The cut-set upper bound on channel network capacity

We outlined certain capacity results for special cases of networks (simple network topologies, additional constraints on the coding scheme). For general networks and arbitrary coding schemes, capacity is not known.[3] A general upper bound on the capacity was given in [3, Thm.14.10.1]. We state it as follows: Suppose there are n nodes in a network. Node k receives Y_k and transmits X_k. Denote the rate at which node i sends to node j by R^{ij}. The nodes are split into two subsets, denoted by S and S^c. Using the shorthand $X^{(S)} \stackrel{def}{=} \{X_k\}_{k\in S}$, Theorem 14.10.1 of [3] can be stated as follows:

Theorem 5.1 (cut-set bound [3]) *If the rates $\{R^{ij}\}$ are achievable, then there exists a joint density $p(x_1, \ldots, x_n)$ (satisfying the channel network input constraints) such that*

$$\sum_{i\in S, j\in S^c} R^{ij} \quad \leq \quad I\left(X^{(S)}; Y^{(S^c)} \,\middle|\, X^{(S^c)}\right) \tag{5.1}$$

for all subsets S of the network.

Incidentally, this upper bound takes a form that is very much reminiscent of "max-flow min-cut": the set S characterizes the cut; the bound then says that the sum of the rates across the cut cannot be larger than the capacity of the cut. However, in contrast to the max-flow min-cut theorem of [6] (see the discussion at the beginning of this section), the bound of Theorem 5.1 it is *not* generally achievable; on the contrary, it is quite loose in many cases.

In order to exploit this theorem to its limit, one has to fix a joint distribution $p(x_1, \ldots, x_n)$ and calculate the mutual information for each cut through the network. This exercise must be repeated for every choice of $p(x_1, \ldots, x_n)$. Clearly, it is not generally simple to find the best bound.

Instead, it is possible to weaken the cut-set bound. This yields an upper bound which is easier to compute (and, in fact, also more easily proved).

[3] See also the explanations in Section I of [67].

Corollary 5.2 (weak cut-set bound) *For any subset S of the network,*

$$\sum_{i \in S, j \in S^c} R^{ij} \quad \leq \quad \max_{p_{X^{(S)}}} I\left(X^{(S)}; Y^{(S^c)} \,\middle|\, X^{(S^c)} \right). \tag{5.2}$$

Proof. While this corollary follows directly from Theorem 5.1, there is also a very simple way of proving it directly, namely by a "genie argument." To see this, suppose that all the terminals on the left of the cut can cooperate arbitrarily, and all the terminals on the right can cooperate arbitrarily. The resulting system cannot have smaller capacity than the original one simply because the original one is *one* way to implement the arbitrary cooperation. Hence, if we maximize the rate in the resulting system, this must lead to an upper bound on the rates in the original system. However, the resulting system is simply a point-to-point (vector) channel whose capacity is given by

$$\max_{p_{X^{(S)}}} I\left(X^{(S)}; Y^{(S^c)} \,\middle|\, X^{(S^c)} \right). \tag{5.3}$$

□

Remark 5.1 (feedback) *The bound of Corollary 5.2 applies also to the case where feedback is available, at least as long as conditional distribution of $Y^{(S^c)}$ given $X^{(S)}$ is memoryless. This follows directly from the proof: the network capacity is upper bounded by the capacity of a point-to-point channel, and feedback cannot increase the capacity of a memoryless point-to-point channel ([99]; see also e.g. [3, p. 213]).— While this is true for the upper bound given in Corollary 5.2, recall that feedback* can *increase the capacity of even a memoryless network, which was shown for the multi-access channel in [34], and for the broadcast channel in [30].*

This last formulation gives a simple means of finding upper bounds on the capacity of a communication network: Choose a subset S of the network, and maximize the mutual information across the cut; the sum of the rates across the cut must be smaller. Note that this is weaker than the cut-set bound (Theorem 5.1) because we maximize for each cut separately, rather than allowing only for one joint distribution $p(x_1, \ldots, x_n)$ for all cuts. However, if the set S is cleverly chosen, this may already give an interesting upper bound.

Note that the general bound, Theorem 5.1, is found by superimposing bounds of the type of Corollary 5.2. While the proof is technically somewhat more involved, even Theorem 5.1 remains in essence a two-terminal bound, and consequently fails to capture the true multi-terminal nature of the underlying system. For this reason, the general bound must be expected to be loose in general.

5.1.3 Rate-distortion behavior of source networks

The paradigmatic and fatal first problem involving a source network is illustrated in Figure 5.3. The two sources are to be represented by codewords of

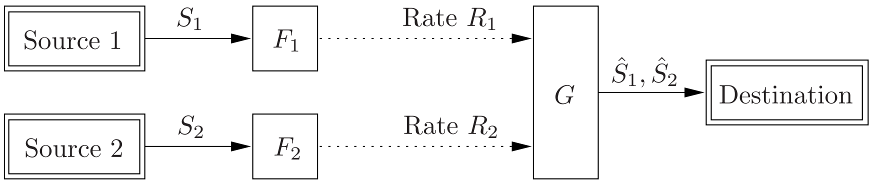

Figure 5.3: A basic source network problem.

rates R_1 and R_2 bits per source sample, respectively. The goal is to determine the set of rate pairs (R_1, R_2) such that the decoder can reconstruct source 1 at distortion D_1 and source 2 at distortion D_2. The answer to this problem is not known in general. The best known outer and inner bounds can be found e.g. in [17], and they do not coincide in general. Nevertheless, the answer to the problem of Figure 5.3 is known for a number of special cases: The case of perfect reconstruction has been solved by Slepian and Wolf in [100]. Other special cases include the situation when $R_1 \to \infty$ (or $R_2 \to \infty$), which is called the Wyner-Ziv problem [113] and was discussed above in Section 1.6.3; the case when $D_1 = 0$ (or $D_2 = 0$), see [58]; and the case when one random variable has to be reconstructed perfectly [20]. The Gaussian case where only one source is reconstructed (but both have to be encoded, using rates R_1 and R_2, respectively) was treated in [78].

In Section 5.4.6, we will encounter a close relative of the problem illustrated in Figure 5.3, the so-called CEO problem. For that case, asymptotic results (as the number of sources tends to infinity) are known.

5.2 Single-source Broadcast Networks

The Gaussian single-source broadcast example was briefly discussed in Section 1.6 (Example 1.4). It is a paradigmatic illustration of the failure of the separation-based design in networks. In this section, we revisit this example in detail and extend it using the results on uncoded transmission that were presented in Chapter 2.

Example 5.1 (single-source Gaussian broadcast) *As in Example 1.4, let the source S be iid Gaussian of zero mean and variance P, and consider a Gaussian broadcast channel with two users, as introduced in [25] and illustrated in Fig. 5.4. The noises Z_1 and Z_2 are zero-mean and Gaussian of variances σ_1^2 and σ_2^2, respectively, and we assume that $\sigma_1^2 < \sigma_2^2$.*

Consider first the performance of a separation-based communication system design that uses capacity-achieving codes on the broadcast channel. The broadcast channel is degraded according to [22], and its capacity region is known. Moreover, it is known that in this special case, user 1 (at the end of the better channel) can also decode the codeword destined to user 2. As a consequence of this fact, the problem of common information *[3, p. 421], i.e., information*

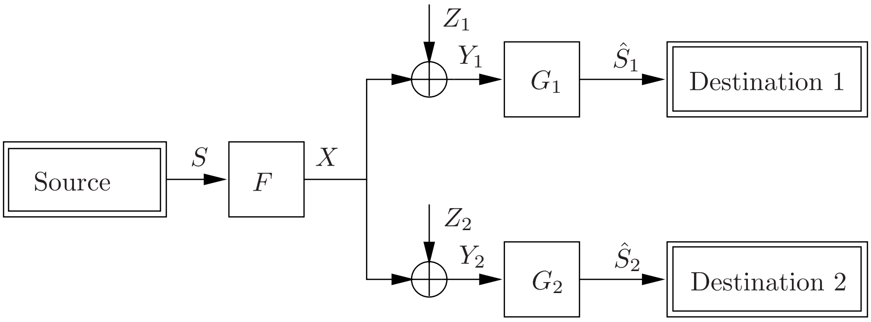

Figure 5.4: Single-source Gaussian broadcast.

destined to both users, is solved in the degraded case: it is simply the codeword destined to user 2. Hence, the largest possible common rate to both users is R_2*, on top of which user 1 gets a rate* R_1*, where* R_1 *and* R_2 *are in the capacity region of the considered Gaussian broadcast channel, found by Bergmans [22] and Gallager [35], see e.g. [3, p. 380]:* (R_1, R_2) *is in the capacity region if it satisfies*

$$R_1 \leq \frac{1}{2}\log_2\left(1+\frac{\alpha P}{\sigma_1^2}\right) \tag{5.4}$$

$$R_2 \leq \frac{1}{2}\log_2\left(1+\frac{(1-\alpha)P}{\alpha P+\sigma_2^2}\right), \tag{5.5}$$

for some $0 \leq \alpha \leq 1$.

Hence, the best possible distortions we can hope for are $D_1 = D_{\mathcal{N}}(R_1 + R_2)$ *for user 1 and* $D_2 = D_{\mathcal{N}}(R_2)$ *for user 2, where* $D_{\mathcal{N}}(\cdot)$ *denotes the distortion-rate function of the Gaussian source with respect to mean-squared error, see e.g. [3, p. 346].*

This performance is indeed achievable by successive refinement source coding [63, 64, 65, 33, 86]. Here, a source is described (coarsely) at rate R_2*, and refined at rate* R_1*. In general, it is not possible to choose the source code such that both the coarse description as well as the its combination with the refinement are rate-distortion optimal. However, for the Gaussian source and mean-squared error, this is possible (and it has recently been established that this is "nearly" true for all sources [72]). Hence, the separation-based communication system design can achieve* $D_1 = D_{\mathcal{N}}(R_1 + R_2)$ *for user 1 and* $D_2 = D_{\mathcal{N}}(R_2)$ *for user 2, where* R_1 *and* R_2 *satisfy (5.4) and (5.4), respectively, for some* $0 \leq \alpha \leq 1$*. The region of corresponding distortion points is illustrated by the shaded area in Figure 5.5.*

On the other hand, it is immediately clear from Example 2.2 that uncoded transmission (or more precisely, a single-letter code) achieves optimal performance. More precisely, consider the system illustrated by Figure 5.6, and take the single-letter "decoders" to be $\beta_1 = P/(P + \sigma_1^2)$ *and* $\beta_2 = P/P + \sigma_2^2)$*, respectively. It is clear that this achieves optimal single-user performance on each channel individually as if the other was not there, i.e.,* $\Delta_{u,1} = P\sigma_1^2/(P + \sigma_1^2)$

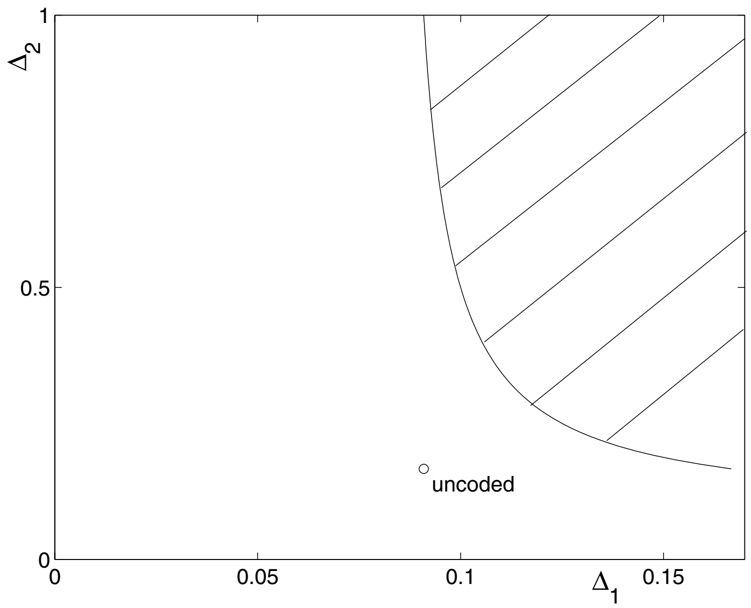

Figure 5.5: The distortion achievable by uncoded transmission (circle) versus the distortion region achievable by a transmission scheme based on the separation principle for Example 5.1. Parameters are $P = 1$, $\sigma_1^2 = 0.1$ and $\sigma_2^2 = 0.2$.

and $\Delta_{u,2} = P\sigma_2^2/(P + \sigma_2^2)$*. This is illustrated in Fig. 5.5 for a particular choice*

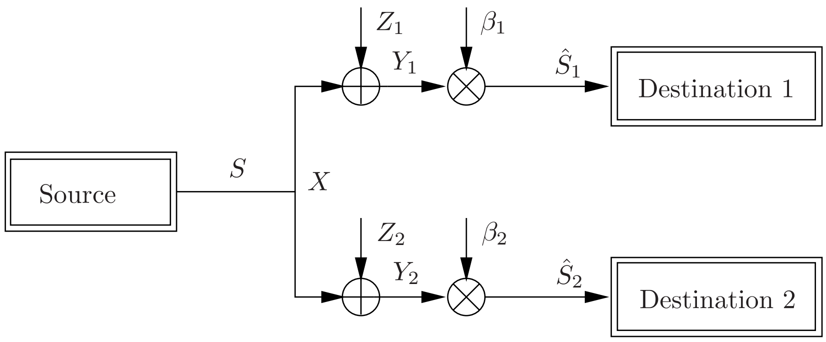

Figure 5.6: Uncoded transmission for the single-source Gaussian broadcast problem.

of the parameters. We observe that the distortion pair achieved by uncoded transmission lies strictly outside the distortion region for the separation-based approach that was described above.

In extension of the results on uncoded transmission of Chapter 2, we defined a certain universality in Section 3.4. One explanation of the superior performance of uncoded transmission in the Gaussian single-source broadcast scenario is precisely this universality: Uncoded transmission is optimal no matter what the variances of the branches of the Gaussian broadcast channel (as in Figure 5.4) are. Hence, we argue in this section that the superior performance of uncoded transmission (as illustrated by Figure 5.5) is not primarily due to the fact that all involved statistics are Gaussian, but due to the fact that all involved measures are matched in the optimal way.

To make this point more explicit, we now formulate a slightly more general result on single-source broadcast. Consider the single-source broadcast network depicted in Figure 5.7: one source is transmitted across a broadcast channel to various destinations. The broadcast channel is specified by a conditional distribution satisfying

$$p(y_1, y_2, \ldots, y_K|x) \quad = \quad p(y_1|x)p(y_2|x)\cdots p(y_K|x), \tag{5.6}$$

and K potentially different input cost functions $\rho_1(x), \rho_2(x), \ldots, \rho_K(x)$. The goal of destination k is to determine an estimate $\hat{S}_k$ of S in such a way as to minimize $Ed_k(S, \hat{S}_k)$. Note that our problem allows potentially for a different distortion measure for each destination.

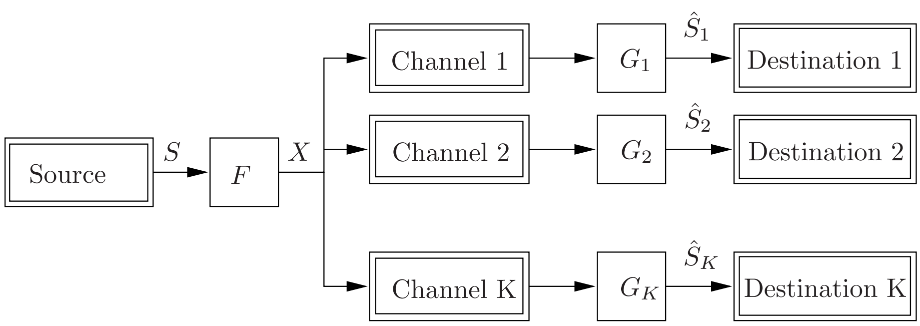

Figure 5.7: The considered single-source broadcast network.

With this, our problem can be stated as follows: For a single-source broadcast problem specified by a broadcast source $(p_S, d_1, \ldots, d_K)$ and a broadcast channel $(p_{Y_1,\ldots,Y_K|X}, \rho)$, determine the set of optimal cost-distortion tuples $(\Gamma, \Delta_1, \ldots, \Delta_K)$. This is yet an unsolved problem. However, a partial answer can be given easily in the following sense.

Lemma 5.3 *If for each $k \in \{1, \ldots, K\}$, (Γ_k, Δ_k) is an optimal cost-distortion tradeoff for the source (p_S, d_k) and the channel $(p_{Y_k|X}, \rho_k)$, then the cost-distortion tuple $(\Gamma_1, \ldots, \Gamma_K, \Delta_1, \ldots, \Delta_K)$ is optimal for the single-source broadcast problem.*

Proof. If it was possible to outperform this scheme, then (Γ, Δ_k) could not be an optimal cost-distortion trade-off for the corresponding point-to-point problem. Hence, $(\Gamma_1, \ldots, \Gamma_K, \Delta_1, \ldots, \Delta_K)$ is optimal. □

Remark 5.2 *The reason why we call Lemma 5.3 a partial answer is that in many single-source broadcast scenarios, cost-distortion tuples as good as the ones described by Lemma 5.3 are not achievable. For those cases, the best cost-distortion tuples are unknown to date.*

The separation theorem does not directly permit one to find instances that satisfy Lemma 5.3. This is due to the fact that there is only one encoder in the scenario, rather than one *per channel.* This cannot easily be taken into account by the separation theorem.

In contrast to this, the results about uncoded transmission for the point-to-point case established in Chapter 2 can be used to identify examples of single-source broadcast problems that satisfy the conditions in Lemma 5.3.

Theorem 5.4 *For the single-source broadcast network shown in Figure 5.7, the condition of Lemma 5.3 is satisfied if, for* $k = 1, \dots, K$*,*

$$\rho_k(x) = D(p_{Y_k|X}(\cdot|x)||p_{Y_k}(\cdot)) \tag{5.7}$$

$$d_k(s, \hat{s}_k) = -\log_2 p(s|\hat{s}_k). \tag{5.8}$$

Proof. The result follows directly from Theorem 2.5. □

Interesting questions in extension of Theorem 5.4 include: Under what conditions are the functions ρ_k the same, for $k = 1, \dots, K$? Similarly, under what conditions are the functions d_k the same, for $k = 1, \dots, K$? These questions are left as future work, but note that in the Gaussian example (Example 5.1), they are indeed all the same.

5.3 Multiple Description Networks

The multiple-description source coding problem has been solved only for the Gaussian case to date. In extension of this solution, a separation theorem was established in Theorem 1.12 for the communication scenario of Figure 1.12, hence for that particular scenario, the best possible cost-distortion trade-offs are known.

In the spirit of Chapter 2 we can ask: Are there instances where a simple *joint* source-channel code achieves those optimal cost-distortion trade-offs? For the special scenario where the two channels in Figure 1.12 are additive white Gaussian noise channels, it turns out that the answer is yes:

Example 5.2 (Gaussian multiple description coding) *An iid Gaussian source of variance* σ_S^2 *is to be transmitted across two AWGN noise variances* σ_1^2 *and* σ_2^2*, respectively, as illustrated in Figure 5.8. To keep notation simple, we assume that both channels have power constraint* σ_S^2*, equal to the source variance. The example can, however, easily be extended to the case where the two channels have different power constraints* P_1 *and* P_2*.*

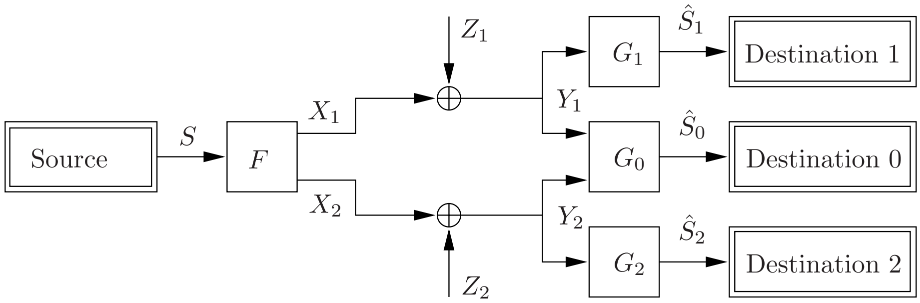

Figure 5.8: The multiple description communication scenario.

One strategy is to encode the source into two streams using multiple description source coding [80], and to transmit these two descriptions (without further losses) across the two AWGN channels. The achievable distortions have been determined in [80]. We have argued in Theorem 1.12 that this leads to optimal performance.

Let us now verify that uncoded transmission achieves a point on the boundary of the distortion region. The two distortions achieved by uncoded transmission on the two channels separately have been determined earlier in this thesis:

$$\Delta_1 = \frac{\sigma_S^2 \sigma_1^2}{\sigma_S^2 + \sigma_1^2} \quad \textit{and} \quad \Delta_2 = \frac{\sigma_S^2 \sigma_2^2}{\sigma_S^2 + \sigma_2^2}.$$

In order to obtain an estimate based on the outputs of both *channels, let us allow for optimal decoding of block length* one*. Then, we find*

$$\Delta_0 = \frac{\sigma_S^2 \sigma_1^2 \sigma_2^2}{\sigma_S^2 \sigma_1^2 + \sigma_S^2 \sigma_2^2 + \sigma_1^2 \sigma_2^2}.$$

The point $(\Delta_0, \Delta_1, \Delta_2)$ *lies indeed on the boundary of the distortion region corresponding to the rates* $R_1 = \frac{1}{2}\log_2(1+\sigma_S^2/\sigma_1^2)$ *and* $R_2 = \frac{1}{2}\log_2(1+\sigma_S^2/\sigma_2^2)$*, as a comparison with [80, Thm. 1] reveals.*

It would be interesting to extend this example to a joint source-channel coding theory of multiple description, just like the examples of Section 2.1 were extended in Chapters 2 and 3. However, since the multiple-description source coding problem is yet unsolved for all cases but the Gaussian, we do not know at present how to generalize Example 5.2.

5.4 Gaussian Relay Networks

The relay channel (both with a single relay and with multiple relays) was introduced by van der Meulen in his Ph.D. thesis [103] and in [104]. Key results for the single-relay channel have been found by Cover and El Gamal [26]. Their capacity results are restricted to the so-called *degraded* relay channel.[4] This restriction is considerably stronger than the common notion of degradedness in the case of broadcast channels introduced in [22]. For example, it is considered to be a weak model for the wireless relay channel. In extension of the single-relay channel, various relay network models have been studied in the literature, some of the most recent examples being [93, 52, 39, 68]. However, capacity results are rare. The simple Gaussian relay network studied in this section is illustrated in Figure 5.9. The corresponding single-relay channel, i.e., the system of Figure 5.9 with $M = 1$, is not a degraded relay channel according to [26]. Its capacity is unknown to date. In this section, we derive upper and lower bounds to the

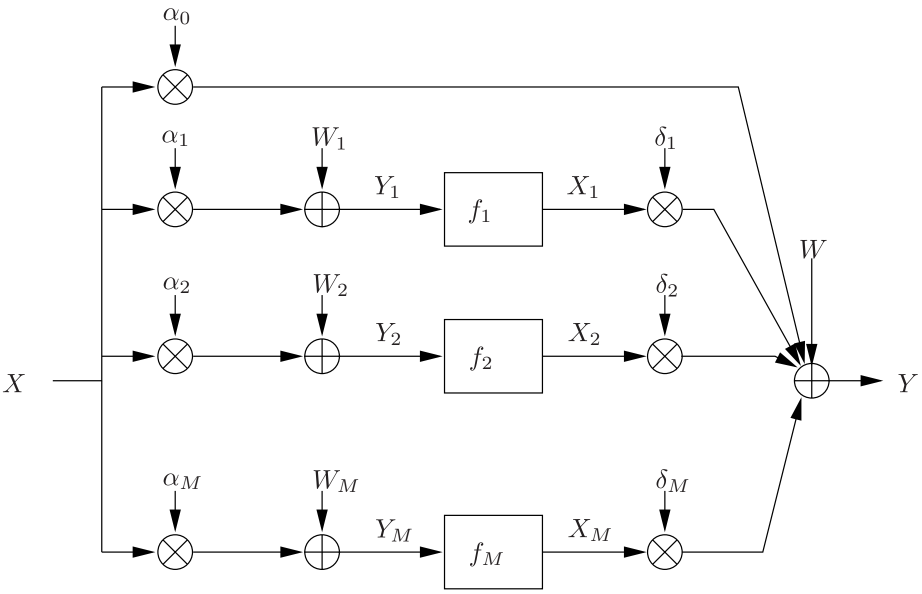

Figure 5.9: The considered Gaussian relay network.

capacity of the system depicted in Figure 5.9. While capacity is not known for any finite M, we show in this section that for many cases of interest, our upper and lower bounds coincide in the limit as the number of relays M tends to infinity, yielding an asymptotic capacity result.

After precisely defining the considered relay network in Section 5.4.1, we derive two upper bounds to the capacity in Section 5.4.2. These bounds can be seen as special cases of the cut-set bound [3, Thm. 14.10.1]. They can also be understood as the capacities corresponding to two idealizations of the considered system: the first upper bound is the capacity of a multi-antenna system with one transmit and $M + 1$ receive antennas, and the second upper bound is the capacity of a multi-antenna system with $M + 1$ transmit and one receive antenna.

[4] To our knowledge, there are only two other non-trivial capacity results: for the semi-deterministic relay channel [31], and for the case where the link from the relay to the destination is a separate point-to-point channel of fixed capacity C [116].

In Section 5.4.3, we determine a lower bound to the capacity of the relay network of Figure 5.9. More precisely, we analyze a particular communication strategy for the relay channel in which the relays use uncoded forwarding as a strategy. This should be expected to be suboptimal in general.

Then, in Section 5.4.4, we show that upper and lower bounds coincide as the number of relays M tends to infinity, at least under certain conditions. We illustrate these conditions by a number of examples.

In Section 5.4.5, the results of the previous sections are applied to sensor networks. In such a network, the relevant trade-off is between the power used by the sensors, and the fidelity (or distortion) at which the interested party can reconstruct the object of interest. While in general, the optimal such trade-off is not known to date, we show that for one particular sensor network situation, the arguments developed in this section lead to a definite result.

Section 5.4.6 provides a new example of the fact that the source-channel separation paradigm does *not* extend to networks. This example follows from the comparison of the result of Section 5.4.5 with the CEO problem [21, 110].

In Section 5.4.7, the results of the previous sections are extended to wireless ad-hoc networks, operated in relay mode. To this end, the simple relay network model considered in this earlier sections is slightly extended, and it is shown that our results permit to bound capacity to within a factor of two. A slightly different interpretation of this result was discussed in [46].

5.4.1 Definitions and notations

There is one input to our network, denoted by X. This input is complex-valued and has to satisfy the power constraint $E|X|^2 \leq P$. The M relays are the square boxes in Figure 5.9. At time n, relay k observes a noisy version of the input $X[n]$ at time n,

$$Y_k[n] \quad = \quad \alpha_k X[n] + W_k[n], \tag{5.9}$$

where $\{W_k[n]\}$ is a sequence (in n, for $n = 1, 2, 3, \ldots$) of independent and identically distributed (iid) circularly symmetric complex Gaussian random variables of mean zero and variance N. Moreover, we also assume that W_k and W_l are independent for all $k \neq l$. The assumption that all noise processes are of the same variance is made without loss of generality: different noise variances can be taken care of by appropriately adjusting the coefficients α_k. Using the sequence of observations $\{Y_k[n]\}$, the relay k produces suitable outputs[5] $\{X_k[n]\}$ that must satisfy two constraints: First, they must be causal, that is

$$X_k[n] \quad = \quad f_k(Y_k[n-1], Y_k[n-2], \ldots, Y_k[1]). \tag{5.10}$$

[5] In [32], Elias studies a similar Gaussian network topology. The difference lies precisely in the fact that we allow for *arbitrary* recoding functions, while the analysis in [32] is restricted to uncoded forwarding.

Second, they must satisfy a power constraint. We consider power constraints of the form

$$\sum_{k=1}^{M} E|X_k|^2 \quad \leq \quad c(M). \tag{5.11}$$

This means that our model allows power allocation among the relays. The destination observes the sum of the signals transmitted by the source and the relays, and additive white noise,

$$Y[n] \quad = \quad \alpha_0 X[n] + \sum_{k=1}^{M} \delta_k X_k[n] + W[n], \tag{5.12}$$

where $W[n]$ is a sequence of iid circularly symmetric complex Gaussian random variables of mean zero and variance N.

The significance and values of the coefficients α_k and δ_k is left open at present. We assume these coefficients to be known throughout the network for the scope of the present chapter. We also assume that α_0 is real, which is made without loss of generality: if it is not real, we can simply change the phase of X, which is equivalent to changing the phases of the transmitted signals X_k, for all k, by the same amount. The coefficients α_k and δ_k may represent the path loss of the signal and hence be related to the geometry of the network as

$$\alpha_k = \frac{1}{d_{0k}^r}, \quad \text{and} \quad \delta_k = \frac{1}{d_{kd}^r}, \tag{5.13}$$

for $k = 1, \ldots, M$, where d_{0k} is the distance from the source to relay k and d_{kd} is the distance from relay k to the destination. They can as well represent fading effects and hence be random variables.

For notational convenience, we define the following functions:

$$a(M) \quad = \quad \sum_{k=0}^{M} |\alpha_k|^2 \tag{5.14}$$

$$d(M) \quad = \quad \alpha_0^2 + \sum_{k=1}^{M} |\delta_k|^2 \tag{5.15}$$

$$b(M) \quad = \quad \sum_{k=1}^{M} |\alpha_k|^2 \frac{|\alpha_k|^2 P + N}{|\delta_k|^2}. \tag{5.16}$$

Recall that the function $c(M)$ denotes the total available relay power. All of our results can be stated in terms of these auxiliary functions.

5.4.2 Upper bounds to capacity

An upper bound to the capacity of the network of Figure 5.9 can be found from the weak cut-set bound, Corollary 5.2. From Figure 5.9, it is clear that our relay network model consists of a broadcast section and a multi-access section.

Hence, there are two natural cuts to consider first, the "broadcast cut" and the "multi-access cut." One way to visualize this is suggested in Figure 5.10: The black dots represent the terminals, i.e., the source, the destination, and the M relay nodes. The dotted lines suggest the connections that determine the bound resulting from the broadcast cut.

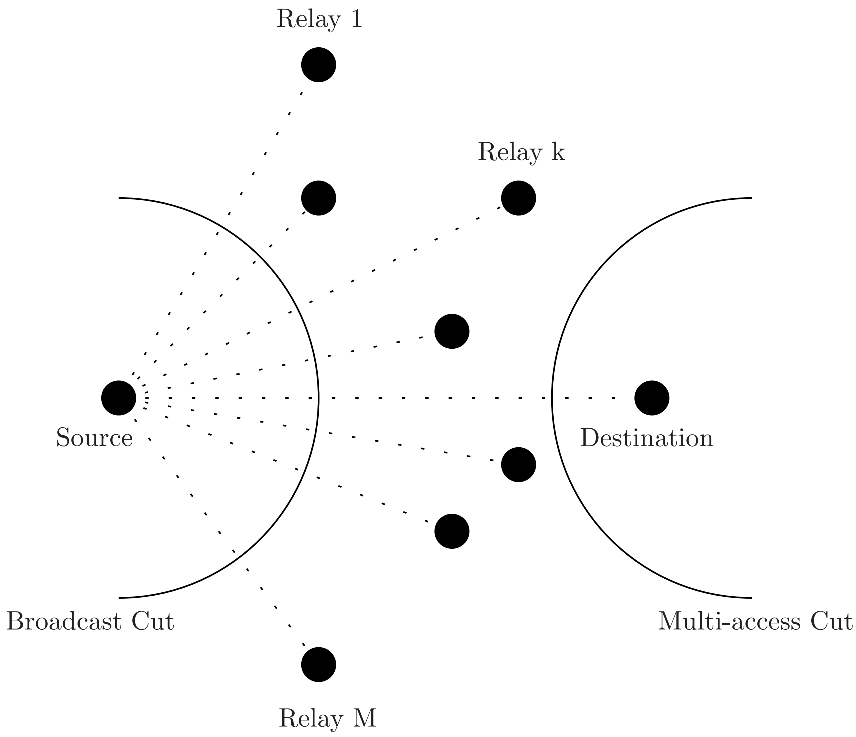

Figure 5.10: Two natural cuts through the considered Gaussian relay network.

For the broadcast cut, the mutual information expression to be maximized is

$$I(X;Y,Y_1,\ldots Y_M|X_1,\ldots X_M), \tag{5.17}$$

subject to the constraints

$$E|X|^2 \leq P \quad \text{and} \quad \sum_{k=1}^{M} E|X_k|^2 \leq c(M). \tag{5.18}$$

Under the assumption that the receiver knows the coefficients δ_k, we find

$$I(X;Y,Y_1,\ldots Y_M|X_1,\ldots X_M) \;=\; I(X;\tilde{Y},Y_1,\ldots Y_M), \tag{5.19}$$

where

$$\tilde{Y} \;=\; Y - \sum_{k=1}^{M} \delta_k X_k. \tag{5.20}$$

Hence, whether or not the receiver knows the coefficients δ_k, the capacity C of the relay network can be bounded by

$$C \leq C_{BC} \stackrel{def}{=} \max_{p_X : EX^2 \leq P} I(X; \tilde{Y}, Y_1, \ldots Y_M). \tag{5.21}$$

The right hand side can be evaluated using the results about Gaussian vector channels (see e.g. [101]), i.e., channels of the form $Y = AX + W$, where Y, X and W are vectors and A is generally a matrix. For the broadcast cut, the matrix A is simply the vector $(\alpha_0, \alpha_1, \ldots \alpha_M)$. It has only one singular value, namely $a(M)$. Hence, C_{BC} is found to be

$$C_{BC} = \log_2 \left(1 + \frac{P}{N} a(M)\right). \tag{5.22}$$

We comment on the tightness of this bound, which was essentially found under the additional hypothesis that the relays do not send anything, but are just the multiple antennae of the destination. However, under the hypothesis that the relays do not send anything, it is easy to find the *true* capacity from the source to the destination: $\log_2(1 + \alpha_0^2 P/N)$. Hence, the bound should not be expected to be tight.

The second cut for which we evaluate Corollary 5.2 is the multi-access cut, as drawn in Figure 5.10. This gives the bound

$$C \leq C_{MAC} \stackrel{def}{=} \max I(X, X_1, \ldots X_M; Y). \tag{5.23}$$

where the maximization is over all $p(x, x_1, \ldots x_M)$ that satisfy

$$E|X|^2 \leq P \quad \text{and} \quad \sum_{k=1}^{M} E|X_k|^2 \leq c(M). \tag{5.24}$$

To obtain a simple expression for the bound, we relax the power constraint to be $E|X|^2 + \sum_{k=1}^{M} E|X_k|^2 \leq P + c(M)$, which can only make the bound looser. Then, we can again use the results from [101] about Gaussian vector channels. Now, the matrix A is the vector $(\alpha_0, \delta_1, \ldots, \delta_M)$, with one singular value, namely $d(M)$. With this, we find

$$C_{MAC} \leq \log_2 \left(1 + \frac{P + c(M)}{N} d(M)\right). \tag{5.25}$$

The capacity must be smaller than either one of these two bounds. Hence, we have proved the following proposition.

Proposition 5.5 *The capacity of the Gaussian relay network of Figure 5.9 is upper bounded by*

$$C \leq \log_2 \left(1 + \frac{\min\{Pa(M), (P + c(M))d(M)\}}{N}\right). \tag{5.26}$$

Tighter bounds can be obtained by evaluating [3, Thm. 14.10.1]. For the scope of this chapter, however, we will use only the upper bound given by Proposition 5.5.

5.4.3 Lower bound to capacity

Lower bounds to capacity are found by analyzing transmission strategies. Such transmission strategies are usually set up in the following way: A code of 2^{nR} codewords of length n is proposed, and it is shown that the error probability goes to zero as $n \to \infty$. This means that the capacity of the considered channel is at least $C \geq R$. The art consists in cleverly choosing the 2^{nR} codewords.

However, whenever the separation theorem applies (e.g., for ergodic point-to-point channels), there is an alternative approach that also leads to lower bounds on capacity. We select any source (with any distortion measure). Then, we propose a transmission strategy for that source across the channel at hand (i.e., a joint source-channel code). Finally we evaluate the performance of that communication strategy; that is, we compute the achieved distortion Δ. By the separation theorem [98, Thm. 21],

$$R(\Delta) \quad \leq \quad C, \tag{5.27}$$

where $R(\cdot)$ denotes the rate-distortion function of the source with respect to the selected distortion measure (see e.g. [3, Thm. 13.2.1]). In other words, the capacity of the channel cannot be smaller than the rate necessary to encode the source at distortion Δ. Clearly, the art consists in selecting the right source with the right distortion measure to get the best lower bound.[6] The drawback of this approach is that we have to know the rate-distortion function of that source with respect to the selected distortion measure, or at least a lower bound to this function.

In this section, we use this approach to determine a lower bound to the capacity of the relay network. We start by fixing the functions according to which the relays operate. These functions must be chosen to satisfy the power and causality constraints. Once they are fixed, the relay network is turned into a point-to-point channel. Clearly, the capacity C' of this point-to-point channel cannot be larger than the capacity C of the relay network. The goal is to determine the capacity C' of this point-to-point channel.

In particular, for the Gaussian relay network of Figure 5.9, we propose the following joint source-channel coding problem: Suppose that an iid Gaussian source of variance P is transmitted without coding on the broadcast section of the relay channel. The relays simply delay the input by one time unit to satisfy causality, and scale it (up or down) to their power level. This coding technique is certainly suboptimal, but its complexity is the absolute minimum. Moreover, we will see later that this coding technique is sufficient to achieve the right scaling behavior in the number of relays M for a large class of Gaussian relay networks of the type depicted in Figure 5.9.

The goal is therefore to determine the distortion D_1 achieved by the suggested coding scheme when the source is zero-mean iid Gaussian of variance P.

[6]Note that with this insight, it becomes very simple to prove a lower bound to the capacity of the standard additive white Gaussian noise channel, i.e., Figure 5.9 with $M = 0$. See also Example 2.2.

The input of relay k at time n is

$$Y_k[n] \quad = \quad \alpha_k X[n] + W_k[n]. \tag{5.28}$$

The strategy of the relay is simply to scale this received value to meet its own power constraint P_k, and to transmit this result onwards. Hence, the output of relay k at time $n+1$ is

$$X_k[n+1] \quad = \quad e^{i\theta_k}\sqrt{\frac{P_k}{|\alpha_k|^2 P + N}} Y_k[n], \tag{5.29}$$

where θ_k is an appropriately chosen phase. There is a power allocation P_k that results in the following distortion:

Proposition 5.6 *The achievable distortion for the transmission of a Gaussian source across the Gaussian relay network of Figure 5.9 with M relays and power constraints $E|X|^2 \leq P$ and $\sum_{k=1}^{M} E|X_k|^2 \leq c(M)$ is not larger than D_1, where D_1 is given by*

$$\begin{aligned} D_1 \quad = \quad P & \frac{Nc(M)a(M) -}{Pc(M)a^2(M) + (N - 2\alpha_0^2 P)c(M)a(M) +} \\ & \frac{-\alpha_0^2 Nc(M) +}{+2\alpha_0 P\sqrt{c(M)}a(M)\sqrt{b(M)} + (\alpha_0^3 P - \alpha_0^2 N)c(M) -} \\ & \frac{+\alpha_0^2 P + N}{-2\alpha_0^4 P\sqrt{c(M)b(M)} + PN(\alpha_0^2 P + N)b(M)}. \end{aligned} \tag{5.30}$$

Proof. The proof is given in Appendix 5.A.

As discussed above, this implies a lower bound to capacity through the separation theorem. For the case at hand, the following statement can be made:

Corollary 5.7 *The capacity of the Gaussian relay network of Figure 5.9 with M relays and power constraints $E|X|^2 \leq P$ and $\sum_{k=1}^{M} E|X_k|^2 \leq c(M)$ is at least*

$$C \quad \geq \quad R_1 \stackrel{def}{=} \log_2 \frac{P}{D_1}, \tag{5.31}$$

with D_1 given by Equation (5.30).

Proof. By the separation theorem [98, Thm. 21], any ergodic point-to-point communication system satisfies $R(\Delta) \leq C$, where Δ is the distortion achieved by that communication system. □

The above transmission strategy is very simple, yet it is a genuine network coding strategy in a sense that we now clarify. In the first step (the broadcasting from the source node to the relays), a "code" is used that permits every relay to decode at its particular level of fidelity, or more precisely, the relays do not decode at all. This is clearly related to the fact that when one Gaussian source is sent across a Gaussian broadcast channel to multiple destinations, then uncoded transmission is an optimal strategy and actually *outperforms* any approach based on capacity-achieving codes. Extensions of this behavior were presented in [44]. In the second step (the multi-accessing from the relays to the destination), cooperative transmission is used to boost transmit power: the signals transmitted by the relays are all correlated.

5.4.4 Scaling behavior and asymptotic capacity

In this section, we compare the upper bounds derived in Section 5.4.2 to the lower bound found in Section 5.4.3. We split the main result of this section into two parts. In Theorem 5.8 below, we characterize the asymptotic difference between the lower bound R_1 and the upper bound that follows from the broadcast cut, C_{BC}. Thereafter, in Theorem 5.10, we characterize the asymptotic difference between R_1 and the upper bound that follows from the multi-access cut, C_{MAC}.

Theorem 5.8 (broadcast cut) *The capacity C of the Gaussian relay network of Figure 5.9 subject to the power constraints $E|X|^2 \leq P$ and $\sum_{k=1}^{M} E|X_k|^2 \leq c(M)$ is bounded between $C_{BC} \geq C \geq R_1$. Since $a(M)$ is a nondecreasing function of M, $\lim_{M\to\infty} 1/a(M) = \theta_a < \infty$. If moreover*

$$\lim_{M\to\infty} \frac{b(M)}{a(M)c(M)} = \tau < \infty, \tag{5.32}$$

then

$$\lim_{M\to\infty} (C_{BC} - R_1) = \gamma_{BC}. \tag{5.33}$$

The constant γ_{BC} takes the following values:

(i) If $\theta_a = 0$, then

$$\gamma_{BC} = \log_2\left(1 + \frac{\alpha_0^2 P + N}{N}\tau\right), \tag{5.34}$$

which means that when $\tau = 0$, then $\gamma_{BC} = 0$.

(ii) In general,

$$\gamma_{BC} = \log_2\left(\frac{1 + \frac{\alpha_0^2 P+N}{N}\tau +}{1 + \frac{N-2\alpha_0^2 P}{P}\theta_a + 2\alpha_0\sqrt{\theta_a\tau} + \left(\alpha_0^3 - \alpha_0^2\frac{N}{P}\right)\theta_a^2 -} \right.$$
$$\left. \frac{+\frac{N}{P}\theta_a - \alpha_0^2\frac{N}{P}\theta_a^2 + \frac{\alpha_0^2 P+N}{P}\theta_a\tau}{-2\alpha_0^4\theta_a\sqrt{\theta_a\tau} + \frac{\alpha_0^2 P+N}{P}\theta_a\tau}\right). \tag{5.35}$$

Proof. The proof is given in Appendix 5.A.

Theorem 5.8 gives general conditions for the convergence of C_{BC} and R_1. The following is a simple concrete illustration.

Example 5.3 (no attenuation, increasing total relay power) *To illustrate the conditions of Theorem 5.8, we now study the concrete example where* $|\alpha_k| = |\delta_k| = 1$ *for all* k*, and the power constraint on the relays is* $c(M) = M^u Q$ *for some constant* Q *and some* $u > 0$*. Hence,* $a(M) = M + 1$ *and*

$$b(M) = \sum_{k=1}^{M} |\alpha_k|^2 \frac{|\alpha_k|^2 P + N}{|\delta_k|^2} = M(P+N). \tag{5.36}$$

Substituting, we find (using $a(M) \geq M$*)*

$$\lim_{M\to\infty} \frac{b(M)}{a(M)c(M)} \leq \lim_{M\to\infty} \frac{M(P+N)}{M^{u+1}Q} = \lim_{M\to\infty} \frac{P+N}{M^u Q} = 0, \tag{5.37}$$

since $u > 0$ *by assumption. Hence, in Theorem 5.8, we have* $\theta = \tau = 0$*, which yields*

$$\lim_{M\to\infty} (C_{BC} - R_1) = 0, \tag{5.38}$$

which means that the capacity of this network behaves asymptotically like

$$C_{BC} = \log_2 \left(1 + \frac{(M+1)P}{N}\right). \tag{5.39}$$

In this example, Theorem 5.8 is asymptotically tight, i.e., it leads to a capacity result.

Example 5.3 is the simplest possible (non-trivial) case of a network according to Figure 5.9. In spite of this fact, the example generalizes in a straightforward manner to cover a large class of interesting cases: Whenever the total available power increases and the fading coefficients are lower and upper bounded (strictly larger than zero, strictly smaller than infinity), then Theorem 5.8 yields a capacity result. We formulate this in the shape of the following corollary:

Corollary 5.9 (bounded fading coefficients) *In the setup of Theorem 5.8, suppose that the fading coefficients are strictly bounded,* $0 < |\alpha_k|, |\delta_k| < \infty$*, for all* k*, and that the power constraint is* $c(M) = M^u Q$*, with* $u > 0$ *and* Q *some constant. Then,*

$$\lim_{M\to\infty} (C_{BC} - R_1) = 0, \tag{5.40}$$

i.e., in this case, the capacity behaves asymptotically like

$$C_{BC} = \log_2 \left(1 + \frac{a(M)P}{N}\right). \tag{5.41}$$

Proof. Since the fading coefficients are strictly bounded, we can upper bound $b(M) \leq M b_{max}^2$ and lower bound $a(M) \geq M a_{min}^2$, and hence,

$$\frac{b(M)}{a(M)c(M)} \leq \frac{M b_{max}^2}{M a_{min}^2 M^u Q} = \frac{b_{max}^2}{a_{min}^2 Q} \frac{1}{M^u}. \tag{5.42}$$

Since $u > 0$ by assumption, this implies that in Theorem 5.8, $\tau = 0$. By the same token,

$$\frac{1}{a(M)} \leq \frac{1}{Ma_{min}^2}, \tag{5.43}$$

and hence, $\theta_a = 0$. But then, Theorem 5.8 asserts the claim. □

This corollary also shows that there is a number of very natural scenarios for which the strategy of uncoded forwarding used to obtain Proposition 5.6 is an optimal strategy (asymptotically as M tends to infinity). To make a stronger point, suppose now that only point-to-point coding is used, as e.g. in [51]. Then, it is easy to see that the achievable rate remains constant, independent of M: The bottleneck is the source node; under the point-to-point constraint, it can only transmit to one relay node at a time. Hence, for the Gaussian relay network as considered here, network coding significantly changes the asymptotic behavior. This conclusion is certainly of interest in the interpretation of the result of [51]: it suggests the possibility that the asymptotic behavior of capacity *does change* when network coding rather than only point-to-point coding is allowed.

We now proceed to the examination of the asymptotic difference between the upper bound stemming from the multi-access cut (Equation (5.25)) and the achievable rate R_1. The main result is the following theorem, which is the analog of Theorem 5.8:

Theorem 5.10 (multi-access cut) *The capacity C of the Gaussian relay network of Figure 5.9 subject to the power constraints $E|X|^2 \leq P$ and $\sum_{k=1}^{M} E|X_k|^2 \leq c(M)$ is bounded between $C_{MAC} \geq C \geq R_1$. Since $a(M), c(M)$ and $d(M)$ are nondecreasing functions of M, $\lim_{M\to\infty} 1/a(M) = \theta_a < \infty$, $\lim_{M\to\infty} 1/c(M) = \theta_c < \infty$, and $\lim_{M\to\infty} 1/d(M) = \theta_d < \infty$ If moreover*

$$\lim_{M\to\infty} \frac{c(M)d(M)}{a(M)} = \tau_1 < \infty, \quad \text{and} \quad \lim_{M\to\infty} \frac{b(M)d(M)}{a^2(M)} = \tau_2 < \infty, \tag{5.44}$$

then

$$\lim_{M\to\infty} (C_{MAC} - R_1) = \gamma_{MAC}. \tag{5.45}$$

The constant γ_{MAC} takes the following values:

(i) If $\theta_a = \theta_d = 0$, then

$$\gamma_{MAC} = \log_2 \left(\tau_1\theta_c + \frac{\tau_1}{P} + \frac{(\alpha_0^2 P + N)(1 + P\theta_c)}{PN}\tau_2 \right). \tag{5.46}$$

Note that the bound is tight if the argument of the logarithm is 1.

(ii) In general,

$$\gamma_{MAC} =$$
$$\log_2\left(\frac{N\tau_1 - N\alpha_0^2\tau_1\theta_a + PN\tau_1\theta_c - PN\alpha_0^2\tau_1\theta_c\theta_a +}{PN - 2PN\alpha_0^2\theta_a + 2PN\alpha_0^4\theta_a^2 + 2\alpha_0 PN\sqrt{\tau_2\theta_c\theta_d} -}\right.$$
$$\left.\frac{+N^2\theta_a - N^2\alpha_0^2\theta_a^2 + (\alpha_0^2 P + N)(\tau_2 + P\tau_2\theta_c + N\tau_2\theta_c\theta_d)}{-2\alpha_0^3 N\sqrt{\tau_2\theta_c\theta_d}\theta_a + N^2\theta_a - \alpha_0^2 N^2\theta_a^2 + (\alpha_0^2 PN + N^2)\tau_2\theta_c\theta_d}\right). \tag{5.47}$$

Proof. The proof is given in Appendix 5.A.

Example 5.4 *Consider again the setup of Example 1. Theorem 5.10 is of no value here: bounding $a(M) \le 2M$, we find*

$$\frac{c(M)d(M)}{a(M)} \ \ge \ \frac{M^{u+1}Q}{2M} = M^u\frac{Q}{2}, \tag{5.48}$$

which does not converge as M tends to infinity. In other words, the multi-access cut leads to a very loose upper bound to capacity in this example.

Example 5.5 (no attenuation, constant total relay power) *As in Example 5.3, suppose that $|\alpha_k| = |\delta_k| = 1$. However, let the power constraint for the present example be $c(M) = Q$, where Q is some constant. Like in Example 5.3, $a(M) = M + 1$ and*

$$b(M) \ = \ \sum_{k=1}^{M} |\alpha_k|^2 \frac{|\alpha_k|^2 P + N}{|\delta_k|^2} = M(P+N). \tag{5.49}$$

First, consider Theorem 5.8. Substituting, we find

$$\lim_{M\to\infty} \frac{b(M)}{a(M)c(M)} \ = \ \lim_{M\to\infty} \frac{M(P+N)}{(M+1)Q} = \frac{P+N}{Q}. \tag{5.50}$$

Hence, in Theorem 5.8, we have $\theta = 0$, but $\tau = (P+N)/Q$, which yields

$$\lim_{M\to\infty}(C_{BC} - R_1) \ = \ \log_2\left(1 + \frac{(P+N)^2}{NQ}\right). \tag{5.51}$$

Hence, R_1 is asymptotically only a constant additive term away from C_{BC}; and thus, it also grows like $\log M$. *Since the capacity lies between $C_{BC} \ge C \ge R_1$, we conclude that C grows like* $\log M$ *as well. We also briefly discuss the difference $C_{BC} - R_1$. It is seen that this difference decreases with Q, but increases with P. This is due to the fact that for our decoding scheme, the original signal of power P is an* interferer *at the destination.*

In order to apply Theorem 5.10 to this example, we first have to determine τ_1 and τ_2 in (5.44), as follows:

$$\tau_1 \ = \ \lim_{M\to\infty} \frac{c(M)d(M)}{a(M)} = \lim_{M\to\infty} \frac{QM}{M+1} = Q, \tag{5.52}$$

and

$$\tau_2 = \lim_{M\to\infty} \frac{b(M)d(M)}{a^2(M)} = \lim_{M\to\infty} \frac{(P+N)M^2}{(M+1)^2} = P+N. \quad (5.53)$$

Since they are both finite, Theorem 5.10 does apply to this example. We evaluate moreover $\theta_a = 0$ and $\theta_d = 0$, and hence, the value of the bound supplied by Theorem 5.10 is determined by (5.46). With $\theta_c = 1/Q$, we find

$$\lim_{M\to\infty} (C_{MAC} - R_1) = \log_2\left(1 + \frac{Q}{P} + \frac{(\alpha_0^2 P + N)(P+N)(1+P/Q)}{PN}\right). \quad (5.54)$$

This bound is always weaker than (5.51). To verify this, note that in the present example, the broadcast bound C_{BC} is asymptotically always smaller than the multi-access bound C_{MAC}, which follows immediately from a comparison of Equations (5.21) and (5.25),

$$\lim_{M\to\infty} 2^{C_{MAC}-C_{BC}} = 1 + \frac{Q}{P}. \quad (5.55)$$

Example 5.6 *In this example, we study the scenario where $|\alpha_k|^2 = k$ and $|\delta_k| = 1$. We can bound $a_0M^2 \le a(M) \le a_1M^2$, $b_0M^3 \le b(M) \le b_1M^3$, and $d(M) = M$. Moreover, suppose a constant total power $c(M) = Q$. Then,*

$$\frac{b(M)}{a(M)c(M)} \ge \frac{b_0M^3}{a_1M^2Q}, \quad (5.56)$$

which diverges, and hence, Theorem 5.8 does not apply.

As for Theorem 5.10, we find

$$\frac{c(M)d(M)}{a(M)} \le \frac{QM}{a_0M^2}, \quad (5.57)$$

which tends to zero (and hence, $\tau_1 = 0$), and

$$\frac{b(M)d(M)}{a^2(M)} \le \frac{b_1M^3M}{a_0^2M^4} = \frac{b_1}{a_0}, \quad (5.58)$$

which converges. This means that in this example, Theorem 5.10 yields a tighter bound than Theorem 5.8. In fact, plugging into Equation (5.46), we find

$$\lim_{M\to\infty} (C_{MAC} - R_1) \le \log_2\left(\frac{(\alpha_0^2 P + N)(1+P/Q)}{PN}\tau_2\right). \quad (5.59)$$

It is also clear from this expression that a capacity result is obtained if it is possible to slightly alter α_k and δ_k in such a way as to make $\tau_2 = \frac{PN}{(\alpha_0^2P+N)(1+P/Q)}$ (while keeping all the limits fixed as in this example).

5.4.5 Application: Gaussian sensor network

The topology of our network model, Figure 5.9, also resembles a particular sensor network situation: in that case, X is the physical phenomenon to be measured, W_k are due to the fact that the phenomenon cannot be measured directly as well as due to measurement noise, and the relays are the sensors themselves. For the sensor network situation, we take $\alpha_0 = 0$. The goal is for a central unit to learn about the physical phenomenon X. The sensors communicate to the central unit over a common wireless channel, and their transmitted powers have to satisfy a sum power constraint. This is clearly a source-channel communication problem; the relevant trade-off is between the total power of the sensors and the fidelity at which the central unit can reconstruct X.

The optimal such trade-off is not known to date. Note that the separation paradigm does not extend to this case: compressing the sources (using the concepts of Slepian and Wolf [100], and their extension to the case of lossy compression [17]) and transmitting the source codewords using a capacity-achieving code on the multi-access channel is a suboptimal strategy. A simple illustration of this can be found in [3, p. 448]. We address this problem below in Section 5.4.6.

The results of this section permit one to determine the optimal trade-off for one special case, namely when X as well as the noises W_k are Gaussian (and in the limit as $M \to \infty$). We report our result in the following corollary:

Corollary 5.11 *Consider a physical phenomenon characterized by the sequence of complex-valued random variables $\{X[n]\}$. Suppose that $X[n]$ are iid circular complex Gaussian random variables of variance P. Sensor k measures $Y_k[n] = \alpha_k X[n] + W_k[n]$, where $W_k[n]$ is iid circular complex Gaussian noise of variance N. Sensor k is allowed to get to know the entire sequence $\{Y_k[n]\}$ before transmitting a sequence $\{X_k[n]\} = f_k(\{Y_k[n]\})$ at power $E|X_k|^2 \leq P_k$. There are M sensors, and their total power is constrained to be*

$$\sum_{k=1}^{M} P_k \quad \leq \quad MQ. \tag{5.60}$$

The final destination receives

$$Y[n] \quad = \quad \sum_{k=1}^{M} \delta_k X_k[n] + W[n], \tag{5.61}$$

where $W[n]$ is iid circular complex Gaussian noise of variance N. The final destination is allowed to get to know the entire sequence $\{Y[n]\}$ before producing the sequence of estimates $\{\hat{X}[n]\} = g(\{Y[n]\})$. If the conditions of Theorem 5.8 are satisfied with $\theta_a = 0$ and $\tau = 0$, then (for each n)

$$\lim_{M\to\infty} \min_{f_1,\dots,f_M,g} E\left|X[n] - \hat{X}[n]\right|^2 \quad = \quad \frac{PN}{a(M)P + N}, \tag{5.62}$$

and the minimum is achieved when the sensors use a simple scaling, $X_k[n] = \gamma_k Y_k[n]$, and the final destination uses $\hat{X}[n] = \gamma Y[n]$.

Remark 5.3 *The encoding functions f_k are much more general here than in the capacity consideration for the relay channel: there is no causality constraint like (5.10) in the considered sensor network model. This means that there are more degrees of freedom in choosing the encoding functions f_k in the present scenario. Nevertheless, our result says that the optimum (asymptotically as $M \to \infty$) can be achieved, without exploiting these additional degrees of freedom, by simple* causal *encoding functions respecting the constraint (5.10). More explicitly, in this example, the global optimum can be achieved by* real-time processing.

Proof. Suppose that the physical phenomenon $X[n]$ itself uses optimal coding. This is clearly an idealization, and hence leads to a lower bound to the distortion. For a given source, the achievable end-to-end distortion certainly cannot be smaller than the rate-distortion function of the source, evaluated at the capacity upper bound C_{BC}. This is immediate since the multi-antenna idealization of the multiple-relay channel is a simple ergodic point-to-point channel, hence the separation theorem applies. The distortion for this idealized system can be calculated as

$$D \geq D_{BC} \stackrel{def}{=} D_{\mathcal{N}}(C_{BC}) = \frac{PN}{a(M)P+N}, \tag{5.63}$$

where $D_{\mathcal{N}}(\cdot)$ denotes the distortion-rate function of the iid circularly complex Gaussian source. An achievable distortion D_1 has been found in Proposition 5.6. Now consider the quotient D_1/D_{BC}, for which (by substitution)

$$\frac{D_1}{D_{BC}} = 2^{C_{BC}-R_1}. \tag{5.64}$$

The convergence of the latter is established in the proof of Theorem 5.8. Hence, as M tends to infinity, the smallest achievable distortion behaves like the D_{BC}. □

5.4.6 Application: The CEO problem

The problem studied in Section 5.4.5 resembles the CEO problem, proposed and solved in [21]. The difference is that the problem of Section 5.4.5 is a joint source-channel coding problem, while the CEO problem is purely a source coding problem. For the problem considered in Section 5.4.5, we can compare the performance of two different modes of operation for the sensors: The first scheme is when the sensors perform simply uncoded transmission, leading to Corollary 5.11. In the second scheme, the sensors apply the best possible distributed compression, and send their respective codewords across the multi-access channel using capacity-achieving codes. That is, they apply the source-channel separation paradigm.

The performance of the first scheme, employing uncoded transmission, has been analyzed in Corollary 5.11. The coefficients are $\alpha_0 = 0$, and $\alpha_k = \delta_k = 1$, for $k = 1, \ldots, M$. Hence, $a(M) = M$ and $b(M) = M(P+N)$. From (5.60),

$c(M) = MQ$, and hence, $\theta_a = 0$ and $\tau = 0$. Therefore, the asymptotic behavior (as $M \to \infty$) of the distortion achieved by the uncoded strategy is found from Corollary 5.11 as

$$D = \frac{PN}{MP + N}. \tag{5.65}$$

Recall that this scheme satisfies an additional property of causality and real-time processing, as described in Corollary 5.11.

To analyze the performance of the second scheme, we use a result from [110]. Let R be the total rate available to encode M correlated sources $Y_1, Y_2, \ldots, Y_M$. The compression must be performed in a distributed fashion: Every source Y_k is encoded separately. A central decoder receives all M codewords, and produces an estimate $\hat{X}(R)$. In [110], it is shown that the smallest achievable distortion behaves like $E|X - \hat{X}(R)|^2 \sim R^{-1}$, in the limit as $R \to \infty$. The full rate-distortion function has subsequently been determined in [79], confirming the R^{-1} behavior for large R. To complete the example, the available rate R has to be calculated. It is determined by the channel characteristics: The maximum sum rate R on the additive white Gaussian multi-access channel specified by (5.60) and (5.61), even if cooperation between the terminals is allowed, is upper bounded by

$$R \leq \log_2\left(1 + \frac{M^2 Q}{N}\right), \tag{5.66}$$

which can be inferred from Equation (5.25). Hence, the distortion achieved by the second scheme behaves at best like

$$D \sim \frac{1}{\log_2\left(1 + \frac{M^2 Q}{N}\right)}. \tag{5.67}$$

The comparison of (5.65) with (5.67) reveals that the joint source-channel coding scheme clearly outperforms the scheme using separate source and channel codes, illustrating again that the separation paradigm does not extend to network situations in general, see also [3, p. 448].

5.4.7 Extension: Wireless networks

It is clear that the network of Figure 5.9 does not well model a wireless situation. Rather, we now consider the following extended network model: Relay k does not only receive the transmission from the source node, but also the transmissions from all other relay nodes. We replace (5.9) by

$$Y_k[n] = \alpha_k X[n] + \sum_{j=1, j\neq k}^{M} \alpha_{jk} X_k[n] + W_k[n]. \tag{5.68}$$

The network model considered by Gupta and Kumar in [51] is of this kind. Their network is discussed in more detail below in Section 5.4.7. The upper bound formulated in Proposition 5.5 applies unchanged: the cut-set idealization still yields upper bounds to capacity. The lower bound, however, is changed since there are many more interfering terms at the input of the relays, as a comparison of (5.9) with (5.68) reveals. One way to obtain a simple lower bound is to apply the same strategy, but to split it into two time slots: in the first time slot, the source sends to the relays, in the second time slot, the relays send to the destination. That is, a factor of two in rate is lost due to this time-multiplexing. Apart from that, the same rate is achievable as for the network specified by (5.9).[7] In the notation of Corollary 5.7, a rate of $R_1' \stackrel{def}{=} R_1/2$ is achievable. Hence, Theorems 5.8 and 5.10 apply almost unchanged, but they now give the asymptotic value of the difference $C_{BC} - 2R_1'$ and $C_{MAC} - 2R_1'$, respectively. In other words, capacity is determined up to a factor of two.

Corollary 5.12 *For the relay network model defined by Equation (5.68) and with the power constraints $E|X|^2 \leq P$ and $\sum_{k=1}^{M} E|X_k|^2 \leq c(M)$, suppose that the conditions of Theorem 5.8 are satisfied with $\theta_a = 0$ and $\tau = 0$. Then,*

$$\lim_{M\to\infty} C_{BC} - 2R_1' \quad = \quad 0, \tag{5.69}$$

i.e., the capacity C is bounded by $C_{BC}/2 \leq C \leq C_{BC}$ as $M \to \infty$.

Proof. The proof follows from the proof of Proposition 5.6. The exact same strategy is used, except that the input to the network is nonzero *only* during the even time slots. With this, the received value at the destination, during the odd time slots, is given by Equation (5.78), but without the first term, $\alpha_0 X[n+1]$. This additional interference is eliminated because the network is used only in the even time slots. Hence, the achieved rate in two time slots cannot be smaller than R_1, and the achieved rate per time slot is at least $R_1' = R_1/2$. The corollary now follows from Theorem 5.8. □

Clearly, Corollary 5.12 can be extended by taking into account the remaining cases covered by Theorem 5.8 as well as the ones covered by Theorem 5.10. We do not discuss this explicitly here.

Instead, the result of Corollary 5.12 is illustrated for two special classes of networks: For networks similar to the ones studied in [51], where the nodes are randomly placed *except* for certain dead zones, and for networks whose source node transmits only half of the time.[8]

Random node placement with dead zones

Consider the following network, illustrated in Figure 5.11: $M + 2$ nodes are placed arbitrarily in a disk of unit area, and the coefficients α_{jk} characterize the

[7] In fact, a slightly higher rate can be achieved since there is no interference between subsequent source symbols.

[8] The capacity result for the case of dead zones and the half-time source was presented at the 2001 IMA "Hot Topics" Workshop on Wireless Networks, August 8-10, 2001, Minneapolis, MN, and appeared in [46].

path losses, i.e.,

$$\alpha_{jk} = \frac{1}{d_{jk}^r}, \tag{5.70}$$

where d_{jk} is the Euclidean distance between nodes j and k, and r is the path loss exponent.[9] These networks are sometimes called *dense* networks since, as

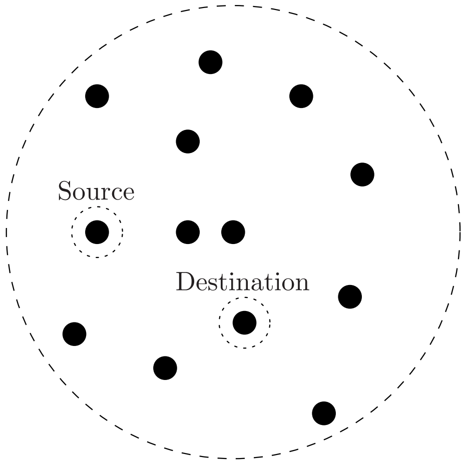

Figure 5.11: A wireless relay network model, based on [51].

the number of nodes n tends to infinity, the distance between adjacent nodes vanishes.

In [51], the nodes are partitioned into pairs (uniformly at random), each pair consisting of a source and a destination, and the goal of the analysis is to characterize the maximum throughput per pair. The analysis presented in [51] is limited to the case where all transmissions are carried out in a point-to-point fashion, considering simultaneous transmissions purely as noise.[10] In contrast to this, our argument considers the traffic scenario where two special nodes are selected, namely a source and a destination, and all other nodes serve purely as relays (see Figure 5.11), and *all possible* coding schemes are allowed, not only those satisfying the point-to-point coding hypothesis mentioned above.

To apply Corollary 5.12 to the network of Figure 5.11, we have to compute the values of θ_a and τ. To make a simple, but precise statement, suppose that there are "dead zones" around the source and the destination nodes: inside such a dead zone of radius $\epsilon > 0$, no other node can be placed.[11] This is suggested in Figure 5.11 by the dotted circles. Notice that the remaining area of the network is unconstrained — the relay nodes can be placed as close to each other as

[9] In [51], this exponent is denoted by α.

[10] More general results beyond this restriction on the coding scheme were recently presented in [114]. However, those results do not seem to apply to dense networks; rather, they impose a lower bound on the distance between any two nodes.

[11] A more restrictive dead zone assumption, not leading to dense networks, has recently been studied in [114].

desired. Using the dead zone assumption, the distance from the source to relay k cannot be smaller than ϵ, and since the network is on a disk of unit area, it cannot be larger than its diameter $\frac{2}{\sqrt{\pi}}$. Hence, α_k can be bounded as

$$\frac{1}{\left(\frac{2}{\sqrt{\pi}}\right)^r} \leq \alpha_k \leq \frac{1}{\epsilon^r}. \tag{5.71}$$

By analogy, the same bounds apply to δ_k. But then, it follows from the arguments leading to Corollary 5.9 that $\theta_a = 0$ and $\tau = 0$, and hence, Corollary 5.12 applies. This argument is true for all node placements that respect the dead zone assumption.

Suppose for example that the node placement is random according to some law that respects the dead zones. The expected asymptotic capacity is determined by $E[C_{BC}]$ in the sense of Corollary 5.12, i.e., up to a factor of two, since

$$E[C_{BC} - 2R_1'] \quad = \quad 0, \tag{5.72}$$

where the expectation is over all node placements. More particularly, the law for the node placement may be uniform like in [51] (except that it respects the dead zones). $E[C_{BC}]$ only depends on $\alpha_1, \ldots, \alpha_M$, and since the nodes are placed independently of one another, it is given by

$$E[C_{BC}] \quad = \quad \int_{\alpha_1} \cdots \int_{\alpha_M} \left(\prod_{k=1}^{M} p(\alpha_k) \right) \log_2 \left(1 + \frac{P}{N} \sum_{k=1}^{M} \alpha_k^2 \right) d\alpha_1 \cdots d\alpha_M.$$

Without calculating this explicitly, we can use (5.71) to bound the expected capacity as

$$\log_2 \left(1 + \frac{P}{N} \left(\frac{\pi}{4} \right)^r M \right) \leq \log_2 \left(1 + \frac{P}{N} \sum_{k=1}^{M} \alpha_k^2 \right) \leq \log_2 \left(1 + \frac{P}{N} \frac{1}{\epsilon^{2r}} M \right), \tag{5.73}$$

showing that the scaling behavior of capacity as a function of M is $\log M$ (at large M).

Half-time source

Corollary 5.12 determines the capacity up to a factor of two. For a special class of networks, Corollary 5.12 can be modified to avoid the factor of two, and hence to determine the asymptotic capacity precisely. In particular, consider the class of networks that observe the following additional constraint: *the source node sends only half of the time.* Under this hypothesis, the broadcast bound can be strengthened; it is easily verified that $C_{BC}' = C_{BC}/2$ is an upper bound to the capacity of the networks in this special class. The scheme used to prove the achievability in Corollary 5.12 satisfies the additional constraint, hence the rate R_1' is still achievable. Therefore, for the considered special class of networks, Corollary 5.12 can be strengthened to yield

$$\lim_{M \to \infty} C_{BC}' - R_1' \quad = \quad 0, \tag{5.74}$$

i.e., capacity behaves asymptotically like

$$C'_{BC} \quad = \quad \frac{1}{2}\log_2\left(1+\frac{P}{N}a(M)\right). \qquad (5.75)$$

5.5 Summary and Conclusions

This chapter discussed some of the potential of uncoded transmission and, more generally, joint source-channel codes in networks. After a brief review of some of the fundamentals of information-theoretic network performance analysis, it was shown that the results of Chapter 2 can be used to determine the optimal cost-distortion trade-off for certain single-source broadcast situations. This optimal trade-off cannot be determined using extensions of the separation theorem.

A short section then presented a certain Gaussian multiple description network for which it turned out that uncoded transmission is an optimal source-channel coding strategy.

Thereafter, the Gaussian relay network of Figure 5.9 was studied. Capacity is not known to date, not even for $M = 1$. In this chapter, we determined the asymptotic capacity in the limit as the number of relays M tends to infinity. For many interesting relay networks, we determined an *exact* asymptotic capacity result, most notably for all cases where the fading coefficients are strictly larger than zero and strictly smaller than infinity.

Beyond the exact capacity results, we show for a larger class of Gaussian relay networks that the typical *scaling behavior* of capacity is $\log M$, where M is the number of relay nodes. This is demonstrated even for network models beyond Figure 5.9, including certain wireless scenarios. In contrast to this, the point-to-point coding hypothesis of [51] leads only to a constant rate, independent of the number of relays. This shows that at least in certain situations, genuine network coding *can* alter the scaling behavior of the capacity of wireless networks.

Finally, we also demonstrate how our results can be applied to sensor networks. There, the trade-off is between sensor power and reconstruction fidelity, and is generally unknown to date. For a particular sensor network situation, we determine the optimal trade-off using the arguments developed in this chapter. We also demonstrate that this optimal trade-off cannot be achieved by separate source and channel code design, illustrating the fact that the separation paradigm does not extend to such sensor networks.

One potential extension of our work is to fading channels, and general situations with limited knowledge of the parameters of the network.

A more fundamental extension of our work is to cases beyond the Gaussian: Is it possible to again give a matching condition, in the spirit of Theorem 3.3? As we explained in this chapter, if the relays apply a separation-based strategy, suboptimal performance results. But there may be a measure-matching condition providing an insightful operating criterion for the relays such that optimal performance in achieved.

Appendix 5.A Proofs

Proof of Proposition 5.6. The relay receives

$$Y_k[n] \quad = \quad \alpha_k X[n] + W_k[n], \tag{5.76}$$

and transmits in the next time slot

$$X_k[n+1] \quad = \quad e^{i\theta_k}\sqrt{\frac{P_k}{|\alpha_k|^2 P + N}}Y_k[n], \tag{5.77}$$

where i is the imaginary unit (the square root of -1), $\theta_k \in [0, 2\pi)$ an appropriately chosen phase, and P_k an appropriately chosen nonnegative real constant. Note that this recoding coefficient makes the expected power of relay k equal to P_k. The received random variable at time $n+1$ is therefore

$$Y[n+1] \quad = \quad \alpha_0 X[n+1] + \sum_{k=1}^{M} \beta_k \left(\alpha_k X[n] + W_k[n]\right) + W[n+1], \tag{5.78}$$

where β_k is defined as

$$\beta_k \quad = \quad \delta_k e^{i\theta_k}\sqrt{\frac{P_k}{|\alpha_k|^2 P + N}}. \tag{5.79}$$

Suppose that we use $\gamma Y[n+1]$ as the estimate of $X[n]$. Then,

$$E\left|X[n] - \gamma Y[n+1]\right|^2 =$$
$$E\left|X[n]\left(1 - \gamma\sum_{k=1}^{M}\beta_k\alpha_k\right) - \gamma\alpha_0 X[n+1] - \gamma\sum_{k=1}^{M}\beta_k W_k[n] - \gamma W[n+1]\right|^2$$

which can be evaluated to yield

$$E\left|X[n] - \gamma Y[n+1]\right|^2 =$$
$$P\left|1 - \gamma\sum_{k=1}^{M}\beta_k\alpha_k\right|^2 + |\gamma|^2|\alpha_0|^2 P + |\gamma|^2 N\left(1 + \sum_{k=1}^{M}|\beta_k|^2\right). \tag{5.80}$$

Recall that α_0 is real by assumption, hence $|\alpha_0|^2 = \alpha_0^2$. The optimal single-letter decoding function γ can be found by taking the derivative of Expression (5.80) with respect to γ and setting this derivative equal to zero. This yields

$$\gamma_{opt} \quad = \quad \frac{P\overline{\sum_{k=1}^{M}\alpha_k\beta_k}}{P|\sum_{k=1}^{M}\alpha_k\beta_k|^2 + \alpha_0^2 P + N(1 + \sum_{k=1}^{M}|\beta_k|^2)}, \tag{5.81}$$

where $\overline{x}$ denotes the complex conjugate of x. The minimum achievable distortion becomes

$$\begin{aligned} D_{opt} &= E|X[n] - \gamma_{opt} Y[n+1]|^2 \\ &= P\left|1 - \gamma_{opt}\sum_{k=1}^{M}\beta_k\alpha_k\right|^2 + |\gamma_{opt}|^2\left(\alpha_0^2 P + N\left(1 + \sum_{k=1}^{M}|\beta_k|^2\right)\right) \\ &= \frac{P\left(\alpha_0^2 P + N\left(1 + \sum_{k=1}^{M}|\beta_k|^2\right)\right)}{P\left|\sum_{k=1}^{M}\alpha_k\beta_k\right|^2 + \alpha_0^2 P + N\left(1 + \sum_{k=1}^{M}|\beta_k|^2\right)}. \end{aligned} \tag{5.82}$$

The next goal is to find suitable β_k's, i.e., to find a good power allocation between the relays. To this end, we may rewrite D_{opt} using vector notation. Define the two vectors $\alpha = (\alpha_0, \alpha_1, \ldots, \alpha_M)$ and $\beta = (1, \beta_1, \ldots, \beta_M)$. Then,

$$D_{opt} = P\frac{\alpha_0^2 P + \langle \beta, \beta \rangle N}{|\langle \alpha, \overline{\beta} \rangle|^2 P + \langle \beta, \beta \rangle N} = P\frac{\frac{\alpha_0^2}{||\beta||^2} P + N}{\frac{|\langle \alpha, \overline{\beta} \rangle|^2}{||\beta||^2} P + N}. \tag{5.83}$$

Minimizing this expression over all vectors β under the given power constraint does not seem to have a simple solution. However, a good (but generally sub-optimal) solution is found by recalling that

$$\max_{\beta:||\beta||^2=const} \frac{\langle \alpha, \overline{\beta} \rangle^2}{||\beta||^2} = \langle \alpha, \alpha \rangle = ||\alpha||^2 \tag{5.84}$$

and that the maximum is achieved when $\overline{\beta} = B\alpha$, for a scalar B.

For the purposes of this section, this relay coding strategy is sufficient. That is, we pick the recoding function given by Equation (5.77) such that $\beta_k = B\overline{\alpha_k}$, for $k = 1, \ldots, M$, which makes the phase θ_k in Equation (5.77)

$$\theta_k = -\arg \alpha_k - \arg \delta_k, \tag{5.85}$$

and the power P_k of relay k

$$P_k = B^2 |\alpha_k|^2 \frac{|\alpha_k|^2 P + N}{|\delta_k|^2}, \tag{5.86}$$

where B must be chosen to satisfy the power constraint $\sum_{k=1}^M P_k \leq c(M)$, i.e., B^2 can be determined as

$$B^2 = \frac{c(M)}{\sum_{k=1}^M |\alpha_k|^2 \frac{|\alpha_k|^2 P + N}{|\delta_k|^2}}. \tag{5.87}$$

To simplify the expression for the distortion, we first point out that

$$\begin{aligned} \langle \alpha, \beta \rangle &= Ba(M) - B\alpha_0^2 + \alpha_0 \\ \langle \beta, \beta \rangle &= B^2 a(M) - B^2 \alpha_0^2 + 1. \end{aligned}$$

Using this in Equation (5.83), the distortion can be expressed as

$$D_1 \stackrel{def}{=} P\frac{B^2(a(M) - \alpha_0^2)N + \alpha_0^2 P + N}{B^2(a(M) - \alpha_0^2)^2 P + (a(M) - \alpha_0^2)B(2\alpha_0 P + BN) + \alpha_0^2 P + N}. \tag{5.88}$$

Plugging in B from above yields the claimed result. □

Remark 5.4 (optimal (multi-letter) decoding) *The two transmission steps of our coding scheme take place simultaneously. This acts like a convolutional code. Hence, the optimum decoder must consider all outputs* $Y[n]$

simultaneously; it cannot operate on a single-letter basis. This does not seem to lead to simple expressions for the achieved distortion, and asymptotically, the improvement of the optimal decoder over the single-letter decoder should not be expected to be large: For many cases, the interfering term $\alpha_0 X[n+1]$ in Equation (5.78) does not influence the asymptotic behavior.

Proof of Theorem 5.8. Recall

$$C_{BC} = \log_2\left(1+\frac{a(M)P}{N}\right). \tag{5.89}$$

We compare this to R_1. Let us write out as follows:

$$\begin{aligned}
2^{C_{BC}-R_1} &= \\
&\frac{B^2(a(M)-\alpha_0^2)N+\alpha_0^2P+N}{B^2(a(M)-\alpha_0^2)^2P+2B(a(M)-\alpha_0^2)\alpha_0P+B^2(a(M)-\alpha_0^2)N+\alpha_0^2P+N}\cdot \\
&\cdot\frac{a(M)P+N}{N} \\
&= \frac{PNB^2a^2(M)+N^2B^2a(M)+(\alpha_0^2P+N)Pa(M)-}{PNB^2a^2(M)+(N-2\alpha_0^2P)NB^2a(M)+2\alpha_0PNBa(M)+} \\
&\frac{-\alpha_0^2N^2B^2+N(\alpha_0^2P+N)}{+(\alpha_0^4PN-\alpha_0^2N^2)B^2-2\alpha_0^3PNB+N(\alpha_0^2P+N)}.
\end{aligned}$$

Plugging in from the power allocation, Equation (5.87), we can replace $B^2 = c(M)/b(M)$. Let us then multiply both numerator and denominator by $b(M)$ to obtain

$$\begin{aligned}
&2^{C_{BC}-R_1} \\
&= \frac{P^2Nc(M)a^2(M)+PN^2c(M)a(M)+}{P^2Nc(M)a^2(M)+(N-2\alpha_0^2P)PNc(M)a(M)+} \\
&\frac{+(\alpha_0^2P+N)P^2a(M)b(M)-}{+2\alpha_0P^2N\sqrt{c(M)}a(M)\sqrt{b(M)}+(\alpha_0^3P^2N-\alpha_0^2PN^2)c(M)-} \\
&\frac{-\alpha_0^2PN^2c(M)+PN(\alpha_0^2P+N)b(M)}{-2\alpha_0^4P^2N\sqrt{c(M)b(M)}+PN(\alpha_0^2P+N)b(M)} \qquad (5.90) \\
&= \frac{1+\frac{N}{P}\frac{1}{a(M)}+\frac{\alpha_0^2P+N}{N}\frac{b(M)}{a(M)c(M)}-\alpha_0^2\frac{N}{P}\frac{1}{a^2(M)}+}{1+\frac{N-2\alpha_0^2P}{P}\frac{1}{a(M)}+2\alpha_0\frac{\sqrt{b(M)}}{a(M)\sqrt{c(M)}}+\left(\alpha_0^3-\alpha_0^2\frac{N}{P}\right)\frac{1}{a^2(M)}-} \\
&\frac{+\frac{\alpha_0^2P+N}{P}\frac{b(M)}{a^2(M)c(M)}}{-2\alpha_0^4\frac{\sqrt{b(M)}}{a^2(M)\sqrt{c(M)}}+\frac{\alpha_0^2P+N}{P}\frac{b(M)}{a^2(M)c(M)}}. \qquad (5.91)
\end{aligned}$$

Next, we argue that under the stated assumptions, each summand both in the numerator as well as in the denominator of M in Equation (5.91) converges. By assumption, $1/a(M) \to \theta_a$, and hence $1/a^2(M) \to \theta_a^2$. This implies the following:

$$\lim_{M\to\infty}\frac{b(M)}{a^2(M)c(M)} = \left(\lim_{M\to\infty}\frac{1}{a(M)}\right)\left(\lim_{M\to\infty}\frac{b(M)}{a(M)c(M)}\right) = \theta_a\tau. \tag{5.92}$$

Similarly,

$$\lim_{M\to\infty} \frac{\sqrt{b(M)}}{a(M)\sqrt{c(M)}} = \left(\lim_{M\to\infty} \sqrt{\frac{1}{a(M)}}\right) \left(\lim_{M\to\infty} \sqrt{\frac{b(M)}{a(M)c(M)}}\right) \quad (5.93)$$

$$= \sqrt{\theta_a \tau}, \quad (5.94)$$

which also implies that $\sqrt{b(M)}/(a^2(M)\sqrt{c(M)}) \to \theta_a\sqrt{\theta_a\tau}$. Finally,

$$\lim_{M\to\infty} \frac{\alpha_0^2 P + N}{N} \frac{b(M)}{a(M)c(M)} = \tau\frac{\alpha_0^2 P + N}{N}, \quad (5.95)$$

leading to the following relationship:

$$\lim_{M\to\infty} 2^{C_{BC}-R_1} = \log_2 \left(\frac{1 + \frac{\alpha_0^2 P+N}{N}\tau +}{1 + \frac{N-2\alpha_0^2 P}{P}\theta_a + 2\alpha_0\sqrt{\theta_a\tau} + \left(\alpha_0^3 - \alpha_0^2\frac{N}{P}\right)\theta_a^2 -} \right.$$

$$\left. \frac{+\frac{N}{P}\theta_a - \alpha_0^2\frac{N}{P}\theta_a^2 + \frac{\alpha_0^2 P+N}{P}\theta_a\tau}{-2\alpha_0^4\theta_a\sqrt{\theta_a\tau} + \frac{\alpha_0^2 P+N}{P}\theta_a\tau} \right) \quad (5.96)$$

which concludes the proof. □

Proof of Theorem 5.10. Recall that

$$C_{MAC} = \log\left(1 + \frac{(P + c(M))d(M)}{N}\right). \quad (5.97)$$

We now compare this to R_1. Let us write out as follows:

$$2^{C_{MAC}-R_1} =$$

$$= \frac{B^2 N(a(M) - \alpha_0^2) + \alpha_0^2 P + N}{B^2(a(M) - \alpha_0^2)^2 P + (a(M) - \alpha_0^2)B(2\alpha_0 P + BN) + \alpha_0^2 P + N} \cdot$$

$$\cdot \frac{(P + c(M))d(M) + N}{N}$$

Plugging in again from the power allocation, Equation (5.87), we can replace $B^2 = c(M)/b(M)$. Let us then multiply both numerator and denominator by $b(M)/c(M)$ to obtain

$$2^{C_{MAC}-R_1} =$$

$$\frac{Na(M)c(M)d(M) - N\alpha_0^2 c(M)d(M) + PNa(M)d(M) - PN\alpha_0^2 d(M) +}{PNa^2(M) - 2PN\alpha_0^2 a(M) + 2PN\alpha_0^4 + 2\alpha_0 PNa(M)\sqrt{\frac{b(M)}{c(M)}} -}$$

$$\frac{+N^2 a(M) - N^2\alpha_0^2 + (\alpha_0^2 P + N)\left(b(M)d(M) + P\frac{b(M)d(M)}{c(M)} + N\frac{b(M)}{c(M)}\right)}{-2\alpha_0^3 N\sqrt{\frac{b(M)}{c(M)}} + N^2 a(M) - \alpha_0^2 N^2 + (\alpha_0^2 PN + N^2)\frac{b(M)}{c(M)}}$$

Dividing both the numerator and the denominator by $a^2(M)$ yields the desired form. By assumption,

$$\lim_{M\to\infty} \frac{c(M)d(M)}{a(M)} = \tau_1 < \infty, \quad \text{and} \quad \lim_{M\to\infty} \frac{b(M)d(M)}{a^2(M)} = \tau_2 < \infty. \quad (5.98)$$

Under these assumptions, we find that

$$\lim_{M\to\infty} \sqrt{\frac{b(M)}{c(M)}}\frac{1}{a(M)} \quad = \quad \sqrt{\tau_2\theta_c\theta_d}. \tag{5.99}$$

Finally,

$$2^{C_{MAC}-R_1} =$$

$$\frac{N\tau_1 - N\alpha_0^2\tau_1\theta_a + PN\tau_1\theta_c - PN\alpha_0^2\tau_1\theta_c\theta_a +}{PN - 2PN\alpha_0^2\theta_a + 2PN\alpha_0^4\theta_a^2 + 2\alpha_0 PN\sqrt{\tau_2\theta_c\theta_d} -}$$

$$\frac{+N^2\theta_a - N^2\alpha_0^2\theta_a^2 + (\alpha_0^2 P + N)(\tau_2 + P\tau_2\theta_c + N\tau_2\theta_c\theta_d)}{-2\alpha_0^3 N\sqrt{\tau_2\theta_c\theta_d}\theta_a + N^2\theta_a - \alpha_0^2 N^2\theta_a^2 + (\alpha_0^2 PN + N^2)\tau_2\theta_c\theta_d}.$$

which concludes the proof. □

Conclusion

Denn das Gemeine geht klanglos zum Orkus hinab.
– FRIEDRICH SCHILLER, *Nänie (1799)*

To code, or not to code: that is the question.

Undoubtedly, "not to code" is very appealing when it leads to an optimal cost-distortion trade-off: it involves the smallest possible delay and complexity. For two examples, it has long been known that "not to code" is the answer: The transmission of an iid Gaussian source across an additive white Gaussian noise channel, and the transmission of a binary uniform source across a binary symmetric channel.

In this thesis, we have shown that there are many more communication systems where "not to code" is the answer. For point-to-point source-channel communication systems, we provide a simple matching condition that identifies the source/channel pairs for which uncoded transmission is already optimal. This matching condition can be extended easily to take into account arbitrary codes: The optimal match then involves, on top of the source/channel pair, also the encoding and the decoding function. We call this the *measure-matching* condition.

"Not to code" is also the answer in certain communication scenarios beyond the simple point-to-point case. We illustrate this for the case of feedback, and of simple network topologies, including single-source broadcast and multiple description networks. For the case of single-source broadcast, when the matching conditions are satisfied, "not to code" is the answer in a much stronger sense than in the simple point-to-point scenario since the separation-based design performs strictly suboptimally.

In the capacity analysis of large Gaussian relay networks, "not to code" turns out to be the answer, but in a different sense: the relays perform uncoded forwarding. This is a capacity-achieving mode of operation in the limit as the number of relays tends to infinity. This result is extended to a sensor network situation, illustrating again that the separation-based design performs strictly suboptimally, while uncoded transmission achieves an optimal cost-distortion trade-off.

Perspectives

Uncoded transmission turned out to be an interesting tool in the analysis of a number of communication problems. This is in the first place due to its simplicity: Its performance and properties can usually be analyzed easily. There are yet many other communication problems to which uncoded transmission has not or not sufficiently been applied yet, including:

- **Feedback communication systems.**

  Initial results are presented in Section 3.5, but there is a more general theory. Particularly tempting in this context is the examination of systems with memory and feedback. There is no general simple expression for the capacity of a channel with memory and feedback; but it may be feasible to determine simple cost-distortion results using joint source-channel coding. More precisely, it may be possible to extend the measure-matching conditions to general systems with feedback.

- **Neural communication.**

  It seems that information theory is not matched to natural communication systems in the first place: information theory studies communication systems *irrespective of delay and complexity*, while these issues must be expected to be crucial in natural systems such as neural communication. Consequently, these communication systems should not be expected to perform optimally in the information-theoretic sense. However, as the optimality of uncoded transmission illustrates, a system can be information-theoretically optimal and yet of very low delay and complexity, if only the involved measures are favorably matched. This argument was outlined in Chapter 4, and much remains to be done. For example, it seems that evolution had all the time to implement a favorable match. It will be very interesting to study the question whether it *actually did* attempt to achieve the information-theoretically optimal match.

- **Communication networks.**

  For communication networks, joint source-channel coding is particularly interesting for two (related) reasons: First, the optimal cost-distortion trade-off is not known in general. Second, the source/channel separation theorem does not extend to networks, or joint source-channel coding *can* beat capacity in the sense that it may achieve a better cost-distortion trade-off than the one achievable by combining optimal distributed compression with capacity-achieving (network) channel codes.

This thesis shows that the measure-matching conditions can be useful for both of these questions. First, our results permit us to determine the optimal cost-distortion trade-off at least for certain simple topologies, including certain instances of single-source broadcast (Section 5.2) and a particular Gaussian sensor network topology (Section 5.4.5). Are there other (simple) network topologies for which the optimal cost-distortion trade-offs can be determined, potentially by extension of the measure-matching conditions? Second, we provide concrete illustrations of the fact that the separate source and channel code design may lead to strictly suboptimal performance (Section 5.4.6).

More explicitly, for general network situations, designing source and channel codes *jointly*, rather than separately from each other, gives strict gains in terms of the cost-distortion trade-off. However, to harvest these gains, the source *and* channel descriptions must be known at the time of code design, i.e., the modularity of the code design is lost. This modularity is very interesting for general purpose networks such as the internet; but it may be less crucial for specialized networks. Rather, in some cases, reducing the power (at fixed fidelity of the data reconstruction) may be much more important than having the source-channel modularity. A prime example is the following:

 – **Sensor networks.** Typically, at the time of code design, the description of the essential source and channel properties would be known. Consider e.g. a sensor network that measures environmental temperature data. Such data has a very particular structure. As illustrated in Section 5.4.6, sensor networks are a key application of uncoded transmission as studied in this thesis: it can save power and/or improve the accuracy of data representation with respect to applying optimal distributed data compression followed by capacity-approaching codes.

- **Modularization of networks.**

  It would be tempting to modularize a network of channels: rather than treating it as one big network, to chop it up into a set of point-to-point channels. This leads to suboptimal performance, as is well known and illustrated also in this thesis. Clearly, designing the network codes jointly for the entire network is a rather involved optimization task; no simple solutions are known to date. One key question is whether there is a good (yet simple) criterion by which intermediate nodes recode their received signal. In the case of the Gaussian relay networks studied in Section 5.4, it turned out that uncoded transmission was good enough (asymptotically as the number of relays tends to infinity). In extension of this, is there a measure-matching condition for intermediate nodes in a network? Or at least in special classes of networks?

- **What is information?**

In the ergodic point-to-point case, information can be associated with *bits* in the sense that the capacity-cost and the rate-distortion function describe the (information-theoretically) optimal system entirely. In that sense, any question can be answered in terms of bits. The same does not hold for networks: There, information cannot simply be identified by bits. This was revisited in various places in this thesis, including Sections 1.6, 5.1, and 5.4.6. Certain new results were found in this thesis; but the question remains: If it is not bits, *what* is information in networks? It remains to be seen whether the concept of measure-matching can help to resolve this question.

If the goal of a thesis was to produce just one piece of advice, ours should probably be to always keep in mind that question: "To code, or not to code?"

Bibliography

Books

[1] T. Berger. *Rate Distortion Theory: A Mathematical Basis For Data Compression.* Prentice-Hall, Englewood Cliffs, NJ, 1971.

[2] D. Bertsekas and R. G. Gallager. *Data Networks.* Prentice Hall, 2nd edition, 1991.

[3] T. M. Cover and J. A. Thomas. *Elements of Information Theory.* Wiley, New York, 1991.

[4] I. Csiszár and J. Körner. *Information Theory: Coding Theory for Discrete Memoryless Systems.* Academic Press, New York, 1981.

[5] R. M. Dudley. *Uniform Central Limit Theorems.* Cambridge University Press, 1999.

[6] L. R. Ford and D. R. Fulkerson. *Flows in Networks.* Princeton University Press, Princeton, NJ, 1962.

[7] R. G. Gallager. *Information Theory and Reliable Communication.* Wiley, New York, 1968.

[8] V. K Goyal. *Single and Multiple Description Transform Coding with Bases and Frames.* SIAM, Philadelphia, 2002. To be published.

[9] F. Jelinek. *Probabilistic Information Theory.* McGraw Hill, New York, 1968.

[10] R. J. McEliece. *The Theory of Information and Coding.* Encyclopedia of mathematics and its applications. Addison-Wesley, Reading MA, 1977.

[11] H. V. Poor. *An Introduction to Signal Detection and Estimation.* Springer, New York, 1994.

[12] F. Rieke, D. Warland, R. R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the neural code.* MIT Press, Cambridge, MA, 1997.

Articles

[13] R. Ahlswede, N. Cai, S.-Y. R. Li, , and R. W. Yeung. Network information flow. *IEEE Transactions on Information Theory*, IT–46(4):1204–1216, July 2000.

[14] A. Albanese, J. Bloemer, J. Edmoonds, M. Luby, and M. Sudan. Priority encoding transmission. *IEEE Transactions on Information Theory*, IT–42(6):1737–1744, November 1996.

[15] V. Anantharam and S. Verdú. Bits through queues. *IEEE Transactions on Information Theory*, IT–42(1):4–18, January 1996.

[16] J. J. Atick. Could information theory provide an ecological theory of sensory processing? *Network 3*, pages 213–251, 1992.

[17] T. Berger. Multiterminal source coding. *Lectures presented at CISM Summer School on the Information Theory Approach to Communications*, July 1977.

[18] T. Berger. *E-Mail message*, July 2000.

[19] T. Berger. BIT — Biological Information Theory. *Invited Lecture, DSC Research Day,* EPFL, Lausanne, Switzerland, July 2001.

[20] T. Berger and R. W. Yeung. Multiterminal source encoding with one distortion criterion. *IEEE Transactions on Information Theory*, IT–35:228–236, March 1989.

[21] T. Berger, Z. Zhang, and H. Viswanathan. The CEO problem. *IEEE Transactions on Information Theory*, IT–42:887–902, May 1996.

[22] P. P. Bergmans. Random coding theorems for broadcast channels with degraded components. *IEEE Transactions on Information Theory*, IT–19(2):197–207, March 1973.

[23] A. B. Carleial. Multiple-access channels with different generalized feedback signals. *IEEE Transactions on Information Theory*, IT–28:841–850, November 1982.

[24] P.-N. Chen and F. Alajaji. Optimistic Shannon coding theorems for arbitrary single-user systems. *IEEE Transactions on Information Theory*, 45(7):2623–2629, November 1999.

[25] T. M. Cover. Broadcast channels. *IEEE Transactions on Information Theory*, IT–18:2–14, January 1972.

[26] T. M. Cover and A. A. El Gamal. Capacity theorems for the relay channel. *IEEE Transactions on Information Theory*, IT–25(5):572–584, September 1979.

[27] T. J. Cruise. Achievement of rate-distortion bound over addidtive white noise channel utilizing a noiseless feedback link. *Proc IEEE (Letters)*, 55:1102–1103, April 1967.

[28] I. Csiszár. On an extremum problem of information theory. *Studia Scientiarum Mathematicarum Hungarica*, pages 57–70, 1974.

[29] R. R. de Ruyter van Steveninck and S. B. Laughlin. The rate of information transfer at graded-potential synapses. *Nature*, 379:642–645, February 1996.

[30] G. Dueck. Partial feedback for two-way and broadcast channels. *Information and Control*, 46:258–266, 1980.

[31] A. A. El Gamal and M. Aref. The capacity of the semideterministic relay channel. *IEEE Transactions on Information Theory*, IT–28:536, May 1982.

[32] P. Elias. Networks of Gaussian channels with applications to feedback systems. *IEEE Transactions on Information Theory*, 13(3):493–501, July 1967.

[33] W. H. R. Equitz and T. M. Cover. Successive refinement of information. *IEEE Transactions on Information Theory*, 37(2):460–473, March 1991.

[34] N. T. Gaarder and J. K. Wolf. The capacity region of a multiple-access discrete memoryless channel can increase with feedback. *IEEE Transactions on Information Theory*, IT–21:100–102, January 1975.

[35] R. G. Gallager. Capacity and coding for degraded broadcast channels. *Problemy Peredaci Informaccii*, 10(3):3–14, 1974.

[36] R. G. Gallager. Finding parity in a simple broadcast network. *Problemy Peredaci Informaccii*, 34(2):176–180, 1988.

[37] M. Gastpar, P. L. Dragotti, and M. Vetterli. The distributed Karhunen-Loève transform. In *Proc 2002 International Workshop on Multimedia Signal Processing*, St. Thomas, US Virgin Islands, December 2002.

[38] M. Gastpar, G. Kramer, and P. Gupta. Antenna-clustering strategies for the multiple-relay channel. *to be submitted to IEEE Transactions on Information Theory*, 2002.

[39] M. Gastpar, G. Kramer, and P. Gupta. The multiple-relay channel: Coding and antenna-clustering capacity. In *Proc IEEE Int Symp Info Theory*, page 136, Lausanne, Switzerland, July 2002.

[40] M. Gastpar and B. Rimoldi. Neural communication: An information-theoretic point of view (poster). In *2001 DSC Research Day*, EPFL, Lausanne, Switzerland, July 2001.

[41] M. Gastpar and B. Rimoldi. Information theory and biological systems (poster). In *2002 DSC Research Day*, EPFL, Lausanne, Switzerland, July 2002.

[42] M. Gastpar, B. Rimoldi, and M. Vetterli. To code or not to code. In *Proc IEEE Int Symp Info Theory*, page 236, Sorrento, Italy, June 2000.

[43] M. Gastpar, B. Rimoldi, and M. Vetterli. To code, or not to code: Lossy source-channel communication revisited. *submitted to IEEE Transactions on Information Theory*, May 2001. Revised July 2002.

[44] M. Gastpar, B. Rimoldi, and M. Vetterli. To code, or not to code: On the optimality of symbol-by-symbol communication. *EPFL DSC Tech. Rep. 026 (submitted to IEEE Transactions on Information Theory)*, May 2001.

[45] M. Gastpar and M. Vetterli. On the capacity of large Gaussian relay networks. *submitted to IEEE Transactions on Information Theory*, September 2002.

[46] M. Gastpar and M. Vetterli. On the capacity of wireless networks: The relay case. In *Proc IEEE Infocom 2002*, New York, NY, June 2002.

[47] T. J. Goblick. Theoretical limitations on the transmission of data from analog sources. *IEEE Transactions on Information Theory*, IT–11(4):558–567, October 1965.

[48] R. M. Gray. Conditional rate-distortion theory. Technical Report 6502-2, Stanford University, Stanford, CA, October 1972.

[49] R. M. Gray. On the asymptotic eigenvalue distribution of Toeplitz matrices. *IEEE Transactions on Information Theory*, IT–18:725–730, November 1972.

[50] M. Grossglauser and D. N. Tse. Mobility can increase the capacity of wireless networks. In *IEEE INFOCOM 2001, Anchorage, Alaska*, April 2001.

[51] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, IT–46(2):388–404, March 2000.

[52] P. Gupta and P. R. Kumar. Towards an information theory of large networks: An achievable rate region. In *Proc IEEE Int Symp Info Theory*, page 159, Washington DC, June 2001. (Also submitted to *IEEE Transaction on Information Theory*, 2001.).

[53] S. V. Hanly and D. N. C. Tse. Multiaccess fading channels—Part II: Delay-limited capacities. *IEEE Transactions on Information Theory*, IT–44(7):2816–1831, November 1998.

[54] T. Ho, M. Médard, and R. Koetter. A coding view of network recovery and management. In *Proc IEEE Int Symp Info Theory*, page 137, Lausanne, Switzerland, July 2002.

[55] M. Horstein. Sequential transmission using noiseless feedback. *IEEE Transactions on Information Theory*, IT–9:136–143, July 1963.

[56] V. Kafedziski. The applications of random operators in Shannon theory - channel capacity, rate distortion and optimal coding. In *Workshop on Concepts in Information Theory*, Breisach, Germany, June 22-28 2002.

[57] T. Kailath. An applications of Shannon's rate-distortion theory to analog communication over feedback channels. In *Proc Princeton Symposium on System Science*, Princeton, NH, March 1967.

[58] A. H. Kaspi and T. Berger. Rate-distortion for correlated source with partially separated encoders. *IEEE Transactions on Information Theory*, IT–28(6):828–840, November 1982.

[59] R. Knopp and P. Humblet. Information capacity and power control in single-cell multiuser communications. In *Proc IEEE International Conference on Communications*, pages 331–335, Seattle, WA, June 1995.

[60] R. Koetter and M. Médard. An algebraic approach to network coding. In *Proc IEEE Int Symp Info Theory*, page 104, Washington DC, June 2001.

[61] R. Koetter and M. Médard. An algebraic approach to network coding. In *Proc Infocom*, New York, NY, June 2002.

[62] J. Körner and K. Marton. How to encode the modulo-two sum of binary sources. *IEEE Transactions on Information Theory*, IT–25(2):219–221, March 1979.

[63] V. Koshelev. Multilevel source coding and data-transmission theorem. In *Proc VIII All-Union Conf. on Theory of Coding and Data Transm.*, volume 1, pages 85–92, Vilnius, U.S.S.R., June 1978.

[64] V. Koshelev. Hierarchical coding of discrete sources. *Probl. Peredach. Inform.*, 16(3):31–49, 1980.

[65] V. Koshelev. An evaluation of the average distortion for discrete scheme of sequential approximation. *Probl. Peredach. Inform.*, 17(3):20–33, 1981.

[66] G. Kramer. *Directed Information for Channels with Feedback*. PhD thesis, Eidg. Techn. Hochschule (ETH), Zürich, Switzerland, 1998.

[67] G. Kramer. Capacity results for the discrete memoryless network. *IEEE Transactions on Information Theory*, 2002. Accepted for publication.

[68] G. Kramer and A. J. van Wijngaarden. On the white Gaussian multiple-access relay channel. In *Proc IEEE Int Symp Info Theory*, page 40, Sorrento, Italy, June 2000.

[69] F. Labeau. *Strategies for exploiting residual redundancy in joint source-channel coding.* PhD thesis, Université catholique de Louvain (UCL), Louvain-la-Neuve, Belgium, 2000.

[70] J. N. Laneman, D. N. C. Tse, and G. W. Wornell. Cooperative diversity in wireless networks: Efficient protocols and outage behavior. *submitted to IEEE Transactions on Information Theory*, January 2002.

[71] A. Lapidoth and P. Narayan. Reliable communication under channel uncertainty. *IEEE Transactions on Information Theory*, IT–44(6):2148 – 2177, October 1998.

[72] L. A. Lastras and T. Berger. All sources are nearly successively refinable. *IEEE Transactions on Information Theory*, IT–47(3):918–926, March 2001.

[73] J. L. Massey. Joint source and channel coding. In J. K. Skwirzinski, editor, *Communication Systems and Random Process Theory*, pages 279–293. Sijthoff & Noordhoff, Alphen aan den Rijn, Netherlands, 1978.

[74] J. L. Massey. Causality, feedback and directed information. In *Proc 1990 Int Symp Info Theory and Its Applications (ISITA'90)*, pages 303–305, Hawaii, U.S.A., November 1990.

[75] N. Merhav and S. Shamai. On joint source-channel coding for the Wyner-Ziv source and the Gel'fand-Pinsker channel. *submitted to IEEE Transactions on Information Theory*, May 2002.

[76] U. Mittal and N. Phambo. Duality theorem for joint source-channel coding. *IEEE Transactions on Information Theory*, IT–46(4):1263–1275, July 2000.

[77] J. K. Omura. Optimum linear transmission of analog data for channels with feedback. *IEEE Transactions on Information Theory*, IT–14(1):38–43, January 1968.

[78] Y. Oohama. Gaussian multiterminal source coding. *IEEE Transactions on Information Theory*, IT–43(6):1912–1923, November 1997.

[79] Y. Oohama. The rate-distortion function for the quadratic Gaussian CEO problem. *IEEE Transactions on Information Theory*, IT–44(3):1057–1070, May 1998.

[80] L. H. Ozarow. A problem with two channels and three receivers. *Bell Sys. Tech. Journal*, 59(10):1909–1921, December 1980.

[81] S. S. Pradhan. On rate-distortion function of Gaussian sources with memory in the presence of side information at the decoder. *Project Report, ECE 480. University of Illinois at Urbana-Champaign*, 1998.

[82] S. S. Pradhan, J. Chou, and K. Ramchandran. Duality between source and channel coding with side information. *UCB/ERL Technical, Memorandum No. M01/34, University of California, Berkeley*, December 2001. Submitted to *IEEE Transactions on Information Theory*.

[83] S. S. Pradhan, R. Puri, and K. Ramchandran. Mds source-channel erasure codes. *submitted to IEEE Transactions on Information Theory*, 2001.

[84] K. Ramchandran and M. Vetterli. Multiresolution joint source and channel coding. In H. V. Poor and G. W. Wornell, editors, *Wireless Communications: A Signal Processing Perspective*. Prentice Hall, Upper Saddle River, NJ, 1998.

[85] Z. Reznic, R. Zamir, and M. Feder. Joint source-channel coding of a Gaussian mixture source over a Gaussian broadcast channel. *IEEE Transactions on Information Theory*, IT–48(3):776–781, March 2002.

[86] B. Rimoldi. Successive refinement of information: Characterization of the achievable rates. *IEEE Transactions on Information Theory*, IT–40:253–259, January 1994.

[87] B. Rimoldi. Beyond the separation principle: A broader approach to source-channel coding. In *4th International ITG Conference on Source and Channel Coding*, Berlin, Germany, January 2002.

[88] K. Rose. A mapping approach to rate-distortion computation and analysis. *IEEE Transactions on Information Theory*, IT–40:1939–1952, November 1994.

[89] I. Rubin. Information rates and data-compression schemes for poisson processes. *IEEE Transactions on Information Theory*, IT–20(2):200–210, January 1974.

[90] J. P. M. Schalkwijk. A coding scheme for additive noise channels with feedback—Part II: Band-limited signals. *IEEE Transactions on Information Theory*, IT–12:183–189, April 1966.

[91] J. P. M. Schalkwijk and L. I. Bluestein. Transmission of analog waveforms through channels with feedback. *IEEE Transactions on Information Theory*, IT–13:617–619, October 1967.

[92] J. P. M. Schalkwijk and T. Kailath. A coding scheme for additive noise channels with feedback—Part I: No bandwidth constraint. *IEEE Transactions on Information Theory*, IT–12:172–182, April 1966.

[93] B. Schein and R. G. Gallager. The Gaussian parallel relay network. In *Proc IEEE Int Symp Info Theory*, page 22, Sorrento, Italy, June 2000.

[94] A. Sendonaris, E. Erkip, and B. Aazhang. User cooperation diversity—Part I: System description. *IEEE Transactions on Communications*, 2002. Accepted for publication.

[95] A. Sendonaris, E. Erkip, and B. Aazhang. User cooperation diversity—Part II: Implementation aspects and performance analysis. *IEEE Transactions on Communications*, 2002. Accepted for publication.

[96] S. Shamai and S. Verdú. Capacity of channels with uncoded side information. *European Transactions on Telecommunications*, 6(5):587–600, Sept.-Oct. 1995.

[97] S. Shamai, S. Verdú, and R. Zamir. Systematic lossy source/channel coding. *IEEE Transactions on Information Theory*, IT–44(2):564–579, March 1998.

[98] C. E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. Journal*, 27:379–423, 623–656, 1948.

[99] C. E. Shannon. The zero-error capacity of a noisy channel. *IRE Transactions on Information Theory*, IT–2:8–19, 1956.

[100] D. Slepian and J. K. Wolf. Noiseless coding of correlated information sources. *IEEE Transactions on Information Theory*, IT–19:471–480, 1973.

[101] I. E. Telatar. Capacity of multi-antenna Gaussian channels. *Bell Labs Technical Memorandum*, June 1995. Also published in *European Transactions on Telecommunications*, 10(6):585-596, Nov.-Dec. 1999.

[102] D. N. C. Tse and S. V. Hanly. Multiaccess fading channels—Part I: Polymatroid structure, optimal resource allocation and throughput capacities. *IEEE Transactions on Information Theory*, IT–44:2796–2815, November 1998.

[103] E. C. van der Meulen. *Transmission of information in a T-terminal discrete memoryless channel*. PhD thesis, University of California, Berkeley, CA, 1968.

[104] E. C. van der Meulen. Three-terminal communication channels. *Adv. Appl. Prob.*, 3:120–54, 1971.

[105] S. Vembu, S. Verdú, and Y. Steinberg. The source-channel separation theorem revisited. *IEEE Transactions on Information Theory*, IT–41(1):44–54, January 1995.

[106] R. Venkataramani, G. Kramer, and V.K. Goyal. Multiple description coding with many channels. *submitted to the IEEE Trans. Inform. Theory*, August 2002.

[107] S. Verdú. The exponential distribution in information theory. *Problems of Information Transmission*, 32(1):86–95, 1996.

[108] S. Verdú and T. S. Han. A general formula for channel capacity. *IEEE Transactions on Information Theory*, IT–40:1147–1157, July 1994.

[109] P. Viswanath, D. N. C. Tse, and V. Anantharam. Asymptotically optimal water-filling in vector multiple-access channels. *IEEE Transactions on Information Theory*, IT–47:241–267, January 2001.

[110] H. Viswanathan and T. Berger. The quadratic Gaussian CEO problem. *IEEE Transactions on Information Theory*, IT–43(5):1549–1559, September 1997.

[111] F. M. J. Willems. The discrete memoryless multiple access channel with partially cooperating encoders. *IEEE Transactions on Information Theory*, IT–29:441 – 445, May 1983.

[112] A. D. Wyner. The rate-distortion function for source coding with side information at the decoder-II: General sources. *Information and Control*, 38:60–80, 1978.

[113] A. D. Wyner and J. Ziv. The rate-distortion function for source coding with side information at the receiver. *IEEE Transactions on Information Theory*, IT–22:1–11, January 1976.

[114] L.-L. Xie and P.R. Kumar. A network information theory for wireless communications: scaling laws and optimal operation. *submitted to IEEE Transactions on Information Theory*, 2002.

[115] R. Zamir. The rate loss in the Wyner-Ziv problem. *IEEE Transactions on Information Theory*, IT–42(4), 1996.

[116] Z. Zhang. Partial converse for a relay channel. *IEEE Transactions on Information Theory*, IT–34(5):1106–1110, September 1988.

[117] J. Ziv and M. Zakai. On functionals satisfying a data-processing theorem. *IEEE Transactions on Information Theory*, IT–19(3):275–283, May 1973.

Curriculum Vitae

Michael Gastpar
Institute of Communication Systems
Ecole Polytechnique Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland
Email: `Michael.Gastpar@epfl.ch`
Web: `http://lcavwww.epfl.ch/~gastpar`

Education

1999 - 2002 Doctorat ès Sciences, Institute of Communication Systems, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.

1997 - 1999 M. S. in Electrical Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA.

1995 - 1996 Erasmus exchange student at the University of Edinburgh, UK.

1992 - 1997 Dipl. El.-Ing., Eidgenössische Technische Hochschule (ETH), Zürich, Switzerland.

Professional Experience

2001 Summer Internship, Mathematics of Communications Research Division, Bell Laboratories, Lucent Technologies, Murray Hill, NJ, USA.

Awards and Honors

ETHZ — EPFL exchange fellowship (three-year fellowship)

Personal

Born in Bern, Switzerland. Swiss Citizen.

Publications

Journals and Manuscripts

1. M. Gastpar and M. Vetterli. On the capacity of large Gaussian relay networks. *submitted to IEEE Transactions on Information Theory*, September 2002.

2. M. Gastpar. The Wyner-Ziv problem with multiple sources. *submitted to IEEE Transactions on Information Theory*, June 2002.

3. M. Gastpar, B. Rimoldi, and M. Vetterli. To code, or not to code: Lossy source-channel communication revisited. *submitted to IEEE Transactions on Information Theory*, May 2001. Revised July 2002.

4. M. Gastpar, B. Rimoldi, and M. Vetterli. To code, or not to code: On the optimality of symbol-by-symbol communication. DSC Technical Report 026, EPFL, Lausanne, Switzerland, May 2001.

5. M. Gastpar and D. G. M. Cruickshank. Two-stage Wiener filter based cancellation receivers for DS-CDMA. *Electron. Lett.*, 32(9):805–806, 1996.

Conferences

1. M. Gastpar, P. L. Dragotti, and M. Vetterli. The distributed Karhunen-Loève transform. In *Proc 2002 International Workshop on Multimedia Signal Processing*, St. Thomas, US Virgin Islands, December 2002.

2. M. Gastpar, G. Kramer, and P. Gupta. The multiple-relay channel: Coding and antenna-clustering capacity. In *Proc IEEE Int Symp Info Theory*, page 136, Lausanne, Switzerland, July 2002.

3. M. Gastpar and M. Vetterli. On the asymptotic capacity of Gaussian relay networks. In *Proc IEEE Int Symp Info Theory*, page 195, Lausanne, Switzerland, July 2002.

4. M. Gastpar and M. Vetterli. On the capacity of wireless networks: The relay case. In *Proc IEEE Infocom 2002*, New York, NY, June 2002.

5. M. Gastpar. On the Wyner-Ziv problem with multiple sources. In *Proc Joint Conference on Communications and Coding*, Saas Fee, Switzerland, March 2002.

6. M. Gastpar, B. Rimoldi, and M. Vetterli. On source/channel codes of finite block length. In *Proc IEEE Int Symp Info Theory*, page 261, Washington DC, June 2001.

7. M. Gastpar, B. Rimoldi, and M. Vetterli. On optimal low-complexity source/channel coding. In *Winter School in Coding and Information Theory*, Ulm, Germany, December 2000.

8. M. Gastpar and Y. Bresler. On the necessary density for spectrum-blind nonuniform sampling subject to quantization. In *Proc IEEE Int Conf Acoustics, Speech, and Signal Processing*, volume 1, pages 348–351, Istanbul, Turkey, June 2000.

9. M. Gastpar, B. Rimoldi, and M. Vetterli. To code or not to code. In *Proc IEEE Int Symp Info Theory*, page 236, Sorrento, Italy, June 2000.

10. M. Gastpar and Y. Bresler. On the necessary density for spectrum-blind nonuniform sampling. In *Proc IEEE Int Symp Info Theory*, page 445, Sorrento, Italy, June 2000.

11. Y. Bresler, M. Gastpar, and R. Venkataramani. Image compression on-the-fly by universal sampling in Fourier imaging systems. In *Proceedings of the IEEE Information Theory Workshop on Detection, Estimation, Classification, and Imaging*, page 48, Santa Fe, NM, February 1999.

12. M. Gastpar and G. Kramer. A spectral criterion for feedback coding. In *Proc IEEE Int Symp Info Theory*, page 126, Boston, MA, 1998.