

Cambridge University Press

978-1-107-11645-0 - Principles of Digital Communication: A Top-Down Approach

Bixio Rimoldi

Frontmatter

[More information](#)

Principles of Digital Communication

This comprehensive and accessible text teaches the fundamentals of digital communication via a top-down-reversed approach, specifically formulated for a one-semester course. It offers students a smooth and exciting journey through the three sub-layers of the physical layer and is the result of many years of teaching digital communication.

The unique approach begins with the decision problem faced by the receiver of the first sub-layer (the decoder), hence cutting straight to the heart of digital communication, and enabling students to learn quickly, intuitively, and with minimal background knowledge.

A swift and elegant extension to the second sub-layer leaves students with an understanding of how a receiver works. Signal design is addressed in a seamless sequence of steps, and finally the third sub-layer emerges as a refinement that decreases implementation costs and increases flexibility.

The focus is on system-level design, with connections to physical reality, hardware constraints, engineering practice, and applications made throughout. Numerous worked examples, homework problems, and **MATLAB** simulation exercises make this text a solid basis for students to specialize in the field of digital communication and it is suitable for both traditional and flipped classroom teaching.

Bixio Rimoldi is a Professor at the Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, where he developed an introductory course on digital communication. Previously he was Associate Professor at Washington University and took visiting positions at Stanford, MIT, and Berkeley. He is an IEEE fellow, a past president of the IEEE Information Theory Society, and a past director of the communication system program at EPFL.

Cambridge University Press

978-1-107-11645-0 - Principles of Digital Communication: A Top-Down Approach

Bixio Rimoldi

Frontmatter

[More information](#)

“This is an excellent introductory book on digital communications theory that is suitable for advanced undergraduate students and/or first-year graduate students, or alternatively for self-study. It achieves a nice degree of rigor in a clear, gentle and student-friendly manner. The exercises alone are worth the price of the book.”

Dave Forney

MIT

“*Principles of Digital Communication: A Top-Down Approach*, 2015, Cambridge University Press, is a special and most attractive text to be used in an introductory (first) course on “Digital Communications”. It is special in that it addresses the most basic features of digital communications in an attractive and simple way, thereby facilitating the teaching of these fundamental aspects within a single semester. This is done without compromising the required mathematical and statistical framework. This remarkable achievement is the outcome of many years of excellent teaching of undergraduate and graduate digital communication courses by the author.

The book is built as appears in the title in a top-down manner. It starts with only basic knowledge on decision theory and, through a natural progression, it develops the full receiver structure and the signal design principles. The final part addresses aspects of practical importance and implementation issues. The text also covers in a clear and simple way more advanced aspects of coding and the associated maximum likelihood (Viterbi) decoder. Hence it may be used also as an introductory text for a more advanced (graduate) digital communication course.

All in all, this extremely well-structured text is an excellent book for a first course on Digital Communications. It covers exactly what is needed and it does so in a simple and rigorous manner that the students and the tutor will appreciate. The achieved balance between theoretical and practical aspects makes this text well suited for students with inclinations to either an industrial or an academic career.”

Shlomo Shamai

Technion, Israel Institute of Technology

“The Rimoldi text is perfect for a beginning mezzanine-level course in digital communications. The logical three layer – discrete-time, continuous-time, pass-band – approach to the problem of communication system design greatly enhances understanding. Numerous examples, problems, and MATLAB exercises make the book both student and instructor friendly. My discussions with the author about the book’s development have convinced me that it’s been a labor of love. The completed manuscript clearly bears this out.”

Dan Costello

University of Notre Dame

Cambridge University Press

978-1-107-11645-0 - Principles of Digital Communication: A Top-Down Approach

Bixio Rimoldi

Frontmatter

[More information](#)

Principles of Digital Communication A Top-Down Approach

Bixio Rimoldi

School of Computer and Communication Sciences

Ecole Polytechnique Fédérale de Lausanne (EPFL)

Switzerland



CAMBRIDGE
UNIVERSITY PRESS

Cambridge University Press
978-1-107-11645-0 - Principles of Digital Communication: A Top-Down Approach
Bixio Rimoldi
Frontmatter
[More information](#)

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

Cambridge University Press is part of the University of Cambridge.

It furthers the University’s mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781107116450

© Cambridge University Press 2016

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2016

Printed in the United Kingdom by TJ International Ltd. Padstow Cornwall

A catalog record for this publication is available from the British Library

Library of Congress Cataloging in Publication data

Rimoldi, Bixio.

Principles of digital communication : a top-down approach / Bixio Rimoldi,
School of Computer and Communication Sciences, Ecole Polytechnique
Fédérale de Lausanne (EPFL), Switzerland.

pages cm

Includes bibliographical references and index.

ISBN 978-1-107-11645-0 (Hardback : alk. paper)

1. Digital communications. 2. Computer networks. I. Title.

TK5103.7.R56 2015

621.382–dc23 2015015425

ISBN 978-1-107-11645-0 Hardback

Additional resources for this publication at www.cambridge.org/rimoldi

Cambridge University Press has no responsibility for the persistence or accuracy of URLs for external or third-party internet websites referred to in this publication, and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

Cambridge University Press

978-1-107-11645-0 - Principles of Digital Communication: A Top-Down Approach

Bixio Rimoldi

Frontmatter

[More information](#)

**This book is dedicated to my parents,
for their boundless support and trust,
and to the late Professor James L. Massey,
whose knowledge, wisdom, and generosity
have deeply touched generations of students.**

Cambridge University Press
978-1-107-11645-0 - Principles of Digital Communication: A Top-Down Approach
Bixio Rimoldi
Frontmatter
[More information](#)

Contents

<i>Preface</i>	<i>page</i> xi
<i>Acknowledgments</i>	xviii
<i>List of symbols</i>	xx
<i>List of abbreviations</i>	xxii
1 Introduction and objectives	1
1.1 The big picture through the OSI layering model	1
1.2 The topic of this text and some historical perspective	5
1.3 Problem formulation and preview	9
1.4 Digital versus analog communication	13
1.5 Notation	15
1.6 A few anecdotes	16
1.7 Supplementary reading	18
1.8 Appendix: Sources and source coding	18
1.9 Exercises	20
2 Receiver design for discrete-time observations: First layer	23
2.1 Introduction	23
2.2 Hypothesis testing	26
2.2.1 Binary hypothesis testing	28
2.2.2 <i>m</i> -ary hypothesis testing	30
2.3 The <i>Q</i> function	31
2.4 Receiver design for the discrete-time AWGN channel	32
2.4.1 Binary decision for scalar observations	34
2.4.2 Binary decision for <i>n</i> -tuple observations	35
2.4.3 <i>m</i> -ary decision for <i>n</i> -tuple observations	39
2.5 Irrelevance and sufficient statistic	41
2.6 Error probability bounds	44
2.6.1 Union bound	44
2.6.2 Union Bhattacharyya bound	48
2.7 Summary	51

viii	Contents
2.8	Appendix: Facts about matrices 53
2.9	Appendix: Densities after one-to-one differentiable transformations 58
2.10	Appendix: Gaussian random vectors 61
2.11	Appendix: A fact about triangles 65
2.12	Appendix: Inner product spaces 65
2.12.1	Vector space 65
2.12.2	Inner product space 66
2.13	Exercises 74
3	Receiver design for the continuous-time AWGN channel: Second layer 95
3.1	Introduction 95
3.2	White Gaussian noise 97
3.3	Observables and sufficient statistics 99
3.4	Transmitter and receiver architecture 102
3.5	Generalization and alternative receiver structures 107
3.6	Continuous-time channels revisited 111
3.7	Summary 114
3.8	Appendix: A simple simulation 115
3.9	Appendix: Dirac-delta-based definition of white Gaussian noise 116
3.10	Appendix: Thermal noise 118
3.11	Appendix: Channel modeling, a case study 119
3.12	Exercises 123
4	Signal design trade-offs 132
4.1	Introduction 132
4.2	Isometric transformations applied to the codebook 132
4.3	Isometric transformations applied to the waveform set 135
4.4	Building intuition about scalability: n versus k 135
4.4.1	Keeping n fixed as k grows 135
4.4.2	Growing n linearly with k 137
4.4.3	Growing n exponentially with k 139
4.5	Duration, bandwidth, and dimensionality 142
4.6	Bit-by-bit versus block-orthogonal 145
4.7	Summary 146
4.8	Appendix: Isometries and error probability 148
4.9	Appendix: Bandwidth definitions 149
4.10	Exercises 150

Contents	ix
5 Symbol-by-symbol on a pulse train: Second layer revisited	159
5.1 Introduction	159
5.2 The ideal lowpass case	160
5.3 Power spectral density	163
5.4 Nyquist criterion for orthonormal bases	167
5.5 Root-raised-cosine family	170
5.6 Eye diagrams	172
5.7 Symbol synchronization	174
5.7.1 Maximum likelihood approach	175
5.7.2 Delay locked loop approach	176
5.8 Summary	179
5.9 Appendix: $\mathcal{L}_2$, and Lebesgue integral: A primer	180
5.10 Appendix: Fourier transform: A review	184
5.11 Appendix: Fourier series: A review	187
5.12 Appendix: Proof of the sampling theorem	189
5.13 Appendix: A review of stochastic processes	190
5.14 Appendix: Root-raised-cosine impulse response	192
5.15 Appendix: The picket fence “miracle”	193
5.16 Exercises	196
6 Convolutional coding and Viterbi decoding: First layer revisited	205
6.1 Introduction	205
6.2 The encoder	205
6.3 The decoder	208
6.4 Bit-error probability	211
6.4.1 Counting detours	213
6.4.2 Upper bound to P_b	216
6.5 Summary	219
6.6 Appendix: Formal definition of the Viterbi algorithm	222
6.7 Exercises	223
7 Passband communication via up/down conversion: Third layer	232
7.1 Introduction	232
7.2 The baseband-equivalent of a passband signal	235
7.2.1 Analog amplitude modulations: DSB, AM, SSB, QAM	240
7.3 The third layer	243
7.4 Baseband-equivalent channel model	252
7.5 Parameter estimation	256

x	Contents
7.6	Non-coherent detection 260
7.7	Summary 264
7.8	Appendix: Relationship between real- and complex-valued operations 265
7.9	Appendix: Complex-valued random vectors 267
7.9.1	General statements 267
7.9.2	The Gaussian case 269
7.9.3	The circularly symmetric Gaussian case 270
7.10	Exercises 275
	<i>Bibliography</i> 284
	<i>Index</i> 286

Preface

This text is intended for a one-semester course on the foundations of digital communication. It assumes that the reader has basic knowledge of linear algebra, probability theory, and signal processing, and has the mathematical maturity that is expected from a third-year engineering student.

The text has evolved out of lecture notes that I have written for EPFL students. The first version of my notes greatly profited from three excellent sources, namely the book *Principles of Communication Engineering* by Wozencraft and Jacobs [1], the lecture notes written by Professor Massey for his ETHZ course *Applied Digital Information Theory*, and the lecture notes written by Professors Gallager and Lapidoth for their MIT course *Introduction to Digital Communication*. Through the years the notes have evolved and although the influence of these sources is still recognizable, the text has now its own “personality” in terms of *content*, *style*, and *organization*.

The *content* is what I can cover in a one-semester course at EPFL.¹ The focus is the transmission problem. By staying focused on the transmission problem (rather than also covering the source digitization and compression problems), I have just the right content and amount of material for the goals that I deem most important, specifically: (1) cover to a reasonable depth the most central topic of digital communication; (2) have enough material to do justice to the beautiful and exciting area of digital communication; and (3) provide evidence that linear algebra, probability theory, calculus, and Fourier analysis are in the curriculum of our students for good reasons. Regarding this last point, the area of digital communication is an ideal showcase for the power of mathematics in solving engineering problems.

The digitization and compression problems, omitted in this text, are also important, but covering the former requires a digression into signal processing to acquire the necessary technical background, and the results are less surprising than those related to the transmission problem (which can be tackled right away, see Chapter 2). The latter is covered in all information theory courses and rightfully so. A more detailed account of the content is given below, where I discuss the text organization.

¹ We have six periods of 45 minutes per week, part of which we have devoted to exercises, for a total of 14 weeks.

In terms of *style*, I have paid due attention to proofs. The value of a rigorous proof goes beyond the scientific need of proving that a statement is indeed true. From a proof we can gain much insight. Once we see the proof of a theorem, we should be able to tell why the conditions (if any) imposed in the statement are necessary and what can happen if they are violated. Proofs are also important because the statements we find in theorems and the like are often not in the exact form needed for a particular application. Therefore, we might have to adapt the statement and the proof as needed.

An instructor should not miss the opportunity to share useful tricks. One of my favorites is the trick I learned from Professor Donald Snyder (Washington University) on how to label the Fourier transform of a rectangle. (Most students remember that the Fourier transform of a rectangle is a sinc but tend to forget how to determine its height and width. See Appendix 5.10.)

The remainder of this preface is about the text *organization*. We follow a *top-down* approach, but a more precise name for the approach is *top-down-reversed with successive refinements*. It is *top-down* in the sense of Figure 1.7 of Chapter 1, which gives a system-level view of the focus of this book. (It is also top-down in the sense of the OSI model depicted in Figure 1.1.) It is *reversed* in the sense that the receiver is treated before the transmitter. The logic behind this reversed order is that we can make sensible choices about the transmitter only once we are able to appreciate their impact on the receiver performance (error probability, implementation costs, algorithmic complexity). Once we have proved that the receiver and the transmitter decompose into blocks of well-defined tasks (Chapters 2 and 3), we *refine* our design, changing the focus from “what to do” to “how to do it effectively” (Chapters 5 and 6). In Chapter 7, we *refine* the design of the second layer to take into account the specificity of passband communication. As a result, the second layer splits into the second and the third layer of Figure 1.7.

In Chapter 2 we acquaint ourselves with the receiver-design problem for channels that have a discrete output alphabet. In doing so, we hide all but the most essential aspect of a channel, specifically that the input and the output are related stochastically. Starting this way takes us very quickly to the heart of digital communication, namely the decision rule implemented by a decoder that minimizes the error probability. The decision problem is an excellent place to begin as the problem is new to students, it has a clean-cut formulation in terms of minimizing an objective function (the error probability), the derivations rely only on basic probability theory, the solution is elegant and intuitive (the maximum a posteriori probability decision rule), and the topic is at the heart of digital communication. After a general start, the receiver design is specialized for the *discrete-time* AWGN (additive white Gaussian noise) channel that plays a key role in subsequent chapters. In Chapter 2, we also learn how to determine (or upper bound) the probability of error and we develop the notion of a sufficient statistic, needed in the following chapter. The appendices provide a review of relevant background material on matrices, on how to obtain the probability density function of a variable defined in terms of another, on Gaussian random vectors, and on inner product spaces. The chapter contains a large collection of exercises.

In Chapter 3 we make an important transition concerning the channel used to communicate, specifically from the rather abstract discrete-time channel to the

more realistic *continuous-time* AWGN channel. The objective remains the same, i.e. develop the receiver structure that minimizes the error probability. The theory of inner product spaces, as well as the notion of sufficient statistic developed in the previous chapter, give us the tools needed to make the transition elegantly and swiftly. We discover that the decomposition of the transmitter and the receiver, as done in the top two layers of Figure 1.7, is general and natural for the continuous-time AWGN channel. This constitutes the end of the first pass over the top two layers of Figure 1.7.

Up until Chapter 4, we assume that the transmitter has been given to us. In Chapter 4, we prepare the ground for the *signal-design*. We introduce the design parameters that we care about, namely transmission rate, delay, bandwidth, average transmitted energy, and error probability, and we discuss how they relate to one another. We introduce the notion of isometry in order to change the signal constellation without affecting the error probability. It can be applied to the encoder to minimize the average energy without affecting the other system parameters such as transmission rate, delay, bandwidth, error probability; alternatively, it can be applied to the waveform former to vary the signal’s time/frequency features. The chapter ends with three case studies for developing intuition. In each case, we fix a signaling family, parameterized by the number of bits conveyed by a signal, and we determine the probability of error as the number of bits grows to infinity. For one family, the dimensionality of the signal space stays fixed, and the conclusion is that the error probability grows to 1 as the number of bits increases. For another family, we let the signal space dimensionality grow exponentially and, in so doing, we can make the error probability become exponentially small. Both of these cases are instructive but have drawbacks that make them unworkable solutions as the number of bits becomes large. The reasonable choice seems to be the “middle-ground” solution that consists in letting the dimensionality grow linearly with the number of bits. We demonstrate this approach by means of what is commonly called pulse amplitude modulation (PAM). We prefer, however, to call it symbol-by-symbol on a pulse train because PAM does not convey the idea that the pulse is used more than once and people tend to associate PAM to a certain family of symbol alphabets. We find symbol-by-symbol on a pulse train to be more descriptive and more general. It encompasses, for instance, phase-shift keying (PSK) and quadrature amplitude modulation (QAM).

Chapter 5 discusses how to choose the orthonormal basis that characterizes the waveform former (Figure 1.7). We discover the *Nyquist criterion* as a means to construct an orthonormal basis that consists of the T -spaced time translates of a single pulse, where T is the symbol interval. Hence we refine the n -tuple former that can be implemented with a single matched filter. In this chapter we also learn how to do symbol synchronization (to know when to sample the matched filter output) and introduce the eye diagram (to appreciate the importance of a correct symbol synchronization). Because of its connection to the Nyquist criterion, we also derive the expression for the *power spectral density* of the communication signal.


In Chapter 6, we design the encoder and refine the decoder. The goal is to expose the reader to a widely used way of encoding and decoding. Because there are several coding techniques – numerous enough to justify a graduate-level course – we approach the subject by means of a case study based on convolutional coding.

The minimum error probability decoder incorporates the Viterbi algorithm. The content of this chapter was selected as an introduction to coding and to introduce the reader to elegant and powerful tools, such as the previously mentioned Viterbi algorithm and the tools to assess the resulting bit-error probability, notably detour flow graphs and generating functions.

The material in Chapter 6 could be covered after Chapter 2, but there are some drawbacks in doing so. First, it unduly delays the transition from the discrete-time channel model of Chapter 2 to the more realistic continuous-time channel model of Chapter 3. Second, it makes more sense to organize the teaching into a *first pass* where we discover what to do (Chapters 2 and 3), and a *refinement* where we focus on how to do it effectively (Chapters 5, 6, and 7). Finally, at the end of Chapter 2, it is harder to motivate the students to invest time and energy into coding for the discrete-time AWGN channel, because there is no evidence yet that the channel plays a key role in practical systems. Such evidence is provided in Chapter 3. Chapters 5 and 6 could be done in the reverse order, but the chosen order is preferable for continuity reasons with respect to Chapter 4.

The final chapter, Chapter 7, is where the third layer emerges as a refinement of the second layer to facilitate *passband* communication.

The following diagram summarizes the main thread throughout the text.



Each chapter contains one or more appendices, with either background or complementary material.

I should mention that I have made an important concession to mathematical rigor. This text is written for people with the mathematical background of an engineer. To be mathematically rigorous, the integrals that come up in dealing with Fourier analysis should be interpreted in the Lebesgue sense.² In most undergraduate curricula, engineers are not taught Lebesgue integration theory. Hence some compromise has to be made, and here is one that I find very satisfactory. In Appendix 5.9, I introduce the difference between the Riemann and the Lebesgue integrals in an informal way. I also introduce the space of $\mathcal{L}_2$ functions and the notion of $\mathcal{L}_2$ equivalence. The ideas are natural and can be understood without technical details. This gives us the language needed to rigorously state the sampling theorem and Nyquist criterion, and the insight to understand why the technical conditions that appear in those statements are necessary. The appendix also reminds us that two signals that have the same Fourier transform are $\mathcal{L}_2$ equivalent but not necessarily point-wise equal. Because we introduce the Lebesgue integral in an informal way, we are not in the position to prove, say, that we can swap an integral and an infinite sum. In some way, having a good reason for skipping such details is a blessing, because dealing with all technicalities can quickly become a major distraction. These technicalities are important at some level and unimportant at another level. They are important for ensuring that the theory is consistent and a serious graduate-level student should be exposed to them. However, I am not aware of a single case where they make a difference in dealing with finite-support functions that are continuous and have finite-energy, especially with the kind of signals we encounter in engineering. Details pertaining to integration theory that are skipped in this text can be found in Gallager’s book [2], which contains an excellent summary of integration theory for communication engineers. Lapidoth [3] contains many details that are not found elsewhere. It is an invaluable text for scholars in the field of digital communication.

The last part of this preface is addressed to instructors. Instructors might consider taking a bottom-up approach with respect to Figure 1.7. Specifically, one could start with the passband AWGN channel model and, as the first step in the development, reduce it to the baseband model by means of the up/down converter. In this case the natural second step is to reduce the baseband channel to the discrete-time channel and only then address the communication problem across the discrete-time channel. I find such an approach to be pedagogically less appealing as it puts the communication problem last rather than first. As formulated by Claude Shannon, the father of modern digital communication, “The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point”. This is indeed the problem that we address in Chapter 2. Furthermore, randomness is the most important aspect of a channel. Without randomness, there is no communication problem. The channels considered in Chapter 2 are good examples to start with, because they model randomness without additional distractions. However, the choice of

² The same can be said for the integrals involving the noise, but our approach avoids such integrals. See Section 3.2.

such abstract channels needs to be motivated. I motivate in two ways: (i) by asking the students to trust that the theory we develop for that abstract channel will turn out to be exactly what we need for more realistic channel models and (ii) by reminding them of the (too often overlooked) problem-solving technique that consists in addressing a difficult problem by first considering a simplified “toy version” of the same.

A couple of years ago, I flipped the classroom in the following sense. Rather than developing the theory in class via standard ex-cathedra lectures and letting the students work on problems at home, I have the students go over the theory at their own pace at home, and I devote the class time to exercises, to detecting difficulties, to filling the gaps, and to motivating the students. Almost the entire content of the book (appendices apart) is covered in the reading assignments.

In my case, flipping the classroom is the result of a process that began with the conviction that the time spent in class was not well spent for many students. There is a fair amount of math in the course *Principles of Digital Communication* and because it is mandatory at EPFL, there is quite a bit of disparity in terms of the rate at which a student can follow the math. Hence, no single pace of teaching is satisfactory, but the real issue has deeper roots. Learning is not about making one step after the other on a straight line at a constant pace, which is essentially what we do in a typical ex-cathedra lecture.

There are a number of things that can improve our effectiveness when we study and cannot be done in an ex-cathedra lecture. Ideally, we should be able to choose suitable periods of time and to decide when a break is needed. More importantly, we should be able to control the information flow, in the sense of being able to “pause” it, e.g. in order to think whether or not what we are learning makes sense to us, to make connections with what we know already, to work out examples, etc. We should also be able to “rewind”, and to “fast forward”. None of this can be done in an ex-cathedra lecture; however, all of this can be done when we study from a book.³ Pausing to think, to make connections, and to work out examples is a particularly useful process that is not sufficiently ingrained in many undergraduate students, perhaps precisely because the ex-cathedra format does not permit it. The book has to be suitable (self-contained and sufficiently clear), the students should be sufficiently motivated to read, and they should be able to ask questions as needed.

Motivation is typically not a problem for the students when the reading is an essential ingredient for passing the class. In my course, the students quickly realize that they will not be able to solve the exercises if they do not read the theory, and there is little chance for them to pass the class without theory and exercises.

But what makes the flipped classroom today more interesting than in the past, is the availability of Web-based tools for posting and answering questions. For my class, I have been using Nota Bene.⁴ Designed at MIT, this is a website on

³ ... or when we watch a video. But a book can be more useful as a reference, because it is easier to find what you are looking for in a book than on a video, and a book can be annotated (personalized) more easily.

⁴ <http://nb.mit.edu>.

which I post the reading assignments (essentially all sections). When students have a question, they go to the site, highlight the relevant part, and type a question in a pop-up window. The questions are summarized on a list that can be sorted according to various criteria. Students can “vote” on a question to increase its importance. Most questions are answered by students, and as an incentive to interact on Nota Bene, I give a small bonus for posting pertinent questions and/or for providing reasonable answers.⁵ The teaching assistants (TAs) and myself monitor the site and we intervene as needed. Before I go to class, I take a look at the questions, ordered by importance; then in class I “fill the gaps” as I see fit.

Most of the class time is spent doing exercises. I encourage the students to help each other by working in groups. The TAs and myself are there to help. This way, I see who can do what and where the difficulties lie. Assessing the progress this way is more reliable than by grading exercises done at home. (We do not grade the exercises, but we do hand out solutions.) During an exercise session, I often go to the board to clarify, to help, or to complement, as necessary.

In terms of my own satisfaction, I find it more interesting to interact with the students in this way, rather than to give ex-cathedra lectures that change little from year to year. The vast majority of the students also prefer the flipped classroom: They say so and I can tell that it is the case from their involvement. The exercises are meant to be completed during the class time,⁶ so that at home the students can focus on the reading. By the end of the semester⁷ we have covered almost all sections of the book. (Appendices are left to the student’s discretion.) Before a new reading assignment, I motivate the students to read by telling them why the topic is important and how it fits into the big picture. If there is something unusual, e.g. a particularly technical passage, I tell them what to expect and/or I give a few hints. Another advantage of the flipped classroom is never falling behind the schedule. At the beginning of the semester, I know which sections will be assigned which week, and prepare the exercises accordingly. After the midterm, I assign a MATLAB project to be completed in groups of two and to be presented during the last day of class. The students like this very much.⁸

⁵ A pertinent intervention is worth half a percent of the total number of points that can be acquired over the semester and, for each student, I count at most one intervention per week. This limits the maximum amount of bonus points to 7% of the total.

⁶ Six periods of 45 minutes at EPFL.

⁷ Fourteen weeks at EPFL.

⁸ The idea of a project was introduced with great success by my colleague, Rüdiger Urbanke, who taught the course during my sabbatical.

Acknowledgments

This book is the result of a slow process, which began around the year 2000, of seemingly endless revisions of my notes written for *Principles of Digital Communication* – a sixth-semester course that I have taught frequently at EPFL. I would like to acknowledge that the notes written by Professors Robert Gallager and Amos Lapidoth, for their MIT course *Introduction to Digital Communication*, as well as the notes by Professor James Massey, for his ETHZ course *Applied Digital Information Theory*, were of great help to me in writing the first set of notes that evolved into this text. Equally helpful were the notes written by EPFL Professor Emre Telatar, on matrices and on complex random variables; they became the core of some appendices on background and on complementary material.

A big thanks goes to the PhD students who helped me develop new exercises and write solutions. This includes Mani Bastani Parizi, Sibi Bhaskaran, László Czap, Prasenjit Dey, Vasudevan Dinkar, Jérémie Ezri, Vojislav Gajic, Michael Gastpar, Saeid Haghighatshoar, Hamed Hassani, Mahdi Jafari Siavoshani, Javad Ebrahimi, Satish Korada, Shrinivas Kudekar, Stéphane Musy, Christine Neuberg, Ayfer Özgür, Etienne Perron, Rajesh Manakkal, and Philippe Roud. Some exercises were created from scratch and some were inspired from other textbooks. Most of them evolved over the years and, at this point, it would be impossible to give proper credit to all those involved. The first round of teaching *Principles of Digital Communication* required creating a number of exercises from scratch. I was very fortunate to have Michael Gastpar (PhD student at the time and now an EPFL colleague) as my first teaching assistant. He did a fabulous job in creating many exercises and solutions.

I would like to thank my EPFL students for their valuable feedback. Pre-final drafts of this text were used at Stanford University and at UCLA, by Professors Ayfer Özgür and Suhas Diggavi, respectively. Professor Rüdiger Urbanke used them at EPFL during two of my sabbatical leaves. I am grateful to them for their feedback and for sharing with me their students' comments.

I am grateful to the following collaborators who have read part of the manuscript and whose feedback has been very valuable: Emmanuel Abbe, Albert Abilov, Nicolae Chiurtu, Michael Gastpar, Matthias Grossglauser, Paolo Ienne, Alberto Jimenez-Pacheco, Olivier Lévêque, Nicolas Macris, Stefano Rosati, Anja Skrivervik, and Adrian Tarniceriu.

Acknowledgments

xix

I am particularly indebted to the following people for having read the whole manuscript and for giving me a long list of suggestions, while noting the typos and mistakes: Emre Telatar, Urs Niesen, Saeid Haghighatshoar, and Sepand Kashani-Akhavan.

Warm thanks go to Françoise Behn who learned L^AT_EX to type the first version of the notes, to Holly Cogliati-Bauereis for her infinite patience in correcting my English, to Emre Telatar for helping with L^AT_EX-related problems, and to Karol Kruzelecki and Damir Laurenzi for helping with computer issues.

Finally, I would like to acknowledge many interesting discussions with various colleagues, in particular those with Emmanuel Abbe, Michael Gastpar, Amos Lapidoth, Upamanyu Madhow, Emre Telatar, and Rüdiger Urbanke. I would also like to thank Rüdiger Urbanke for continuously encouraging me to publish my notes. Without his insistence and his jokes about my perpetual revisions, I might still be working on them.

List of symbols

$\mathcal{A}, \mathcal{B}, \dots$	Sets.
$\mathbb{N}$	Set of natural numbers: $\{1, 2, 3, \dots\}$.
$\mathbb{Z}$	Set of integers: $\{\dots, -2, -1, 0, 1, 2, \dots\}$.
$\mathbb{R}$	Set of real numbers.
$\mathbb{C}$	Set of complex numbers.
$\mathcal{H} := \{0, \dots, m - 1\}$	Message set.
$\mathcal{C} := \{c_0, \dots, c_{m-1}\}$	Codebook (set of codewords).
$\mathcal{W} := \{w_0(t), \dots, w_{m-1}(t)\}$	Set of waveform signals.
$\mathcal{V}$	Vector space or inner product space.
$u : \mathcal{A} \rightarrow \mathcal{B}$	Function u with domain $\mathcal{A}$ and range $\mathcal{B}$.
$H \in \mathcal{H}$	Random message (hypothesis) taking value in $\mathcal{H}$.
$N(t)$	Noise.
$N_E(t)$	Baseband-equivalent noise.
$R(t)$	Received (random) signal.
$Y = (Y_1, \dots, Y_n)$	Random n -tuple observed by the decoder.
j	$\sqrt{-1}$.
$\{ \}$	Set of objects.
$A^\top$	Transpose of the matrix A . It may be applied to an n -tuple a .
$A^\dagger$	Hermitian transpose of the matrix A . It may be applied to an n -tuple a .
$\mathbb{E}[X]$	Expected value of X .
$\langle a, b \rangle$	Inner product between a and b (in that order).
$\ a\ $	Norm of the vector a .
$ a $	Absolute value of a .
$a := b$	a is defined as b .
$\mathbb{1}\{S\}$	Indicator function. Its value is 1 if the statement S is true and 0 otherwise.
$\mathbb{1}_{\mathcal{A}}(x)$	Same as $\mathbb{1}\{x \in \mathcal{A}\}$.

List of symbols		xxi
$\mathcal{E}$	Average energy.	
$K_N(t + \tau, t), K_N(\tau)$	Autocovariance of $N(t)$.	
$\Re$	Used to denote the end of theorems, definitions, examples, proofs, etc.	
$\Re\{\cdot\}$	Real part of the enclosed quantity.	
$\Im\{\cdot\}$	Imaginary part of the enclosed quantity.	
$\angle$	Phase of the complex-valued number that follows.	

List of abbreviations

AM	amplitude modulation.
bps	bits per second.
BSS	binary symmetric source.
DSB-SC	double-sideband modulation with suppressed carrier.
iid	independent and identically distributed.
l. i. m.	limit in $\mathcal{L}_2$ norm.
LNA	low-noise amplifier.
MAP	maximum a posteriori.
Mbps	megabits per second.
ML	maximum likelihood.
MMSE	minimum mean square error.
PAM	pulse amplitude modulation.
pdf	probability density function.
pmf	probability mass function.
PPM	pulse position modulation.
PSK	phase-shift keying.
QAM	quadrature amplitude modulation.
SSB	single-sideband modulation.
WSS	wide-sense stationary.