Telecommunications

Protocols and Design

John D. Spragins
Clemson University
with Joseph L. Hammond
Clemson University
and Krzysztof Pawlikowski
University of Canterbury

Addison-Wesley Publishing Company
Reading, Massachusetts ▪ Menlo Park, California ▪ New York
Don Mills, Ontario ▪ Wokingham, England ▪ Amsterdam ▪ Bonn
Sydney ▪ Singapore ▪ Tokyo ▪ Madrid ▪ San Juan ▪ Milan ▪ Paris

Sponsoring Editor: Tom Robbins
Production Supervisor: Helen M. Wythe
Production Coordinator: Amy Willcutt
Art Consultant: Joseph K. Vetere
Copy Editor: Stephanie Argeros-Magean
Illustrators: Scientific Illustrators
Text Designer: Hunter Graphics
Cover Designer: Marshall Henrichs
Manufacturing Manager: Hugh Crawford

Library of Congress Cataloging-in-Publication Data

Spragins, John D.
Telecommunications : protocols and design / John D. Spragins, Joseph L. Hammond, Krzysztof Pawlikowski.
p. cm.
Includes bibliographical references.
ISBN 0-201-09290-5
1. Computer networks. 2. Computer network protocols. I. Hammond, Joseph L. II. Pawlikowski, Krzysztof. III. Title.
TK5105.5.S67 1991
621.382—dc20 89-27327
CIP

Reprinted with corrections August, 1994

9 10 11 12 13-DOC-00 99 98 97

Preface

Although telecommunications is one of the most important and rapidly growing fields of current technology, there does not seem to be a fully adequate text for a first course on telecommunication networks at the upper level undergraduate or first year graduate level. This book has been written to fill the need for such a text for students in electrical engineering, computer engineering, computer science, and related disciplines. At Clemson University the material is taught at the senior or graduate level. In addition to serving as a university text, the book is suitable for a course taken by practicing engineers or computer scientists, and as a reference book.[1] The primary background required is a level of technical expertise that is typical of the students indicated—that is, a general knowledge of computer systems, plus mathematics through calculus, and a first course in probability or statistics.[2]

The book discusses a wide variety of problems encountered in designing telecommunication networks and presents common techniques to solve them. A good understanding of telecommunication protocols is essential for anyone working on the design of telecommunication networks, and their study is a major emphasis of this text. Nontechnical factors, such as legal, regulatory, and economic factors, plus standards activities, are also discussed since they have major impacts on design of telecommunication networks. The emphasis is on basic principles, and on motivations for designs, rather than on encyclopedic coverage of the state of the art. When practical, within constraints imposed by the background assumed for students, a quantitative approach is used. The field cannot be fully covered in a quantitative manner, however, since aspects that do not lend themselves to quantitative treatment are important. For example, only limited parts of the design of protocols can be treated quantitatively.

The basic approach to presenting this material is to describe problems that must be solved during the design of telecommunication networks and then to discuss typical approaches used to solve these problems. Chapters that concentrate on typical problem-solving techniques alternate with chap-

---

[1] Footnotes have been used liberally to present precise details useful for those using the book as a reference. Many of these can be ignored during a first reading, or in a first course based on the text.

[2] Students who do not have the first course in probability or statistics should be able to understand most of the material. Subsections that need to be omitted for classes with such students are listed here; their omission will not significantly impact coverage of other topics.

ters presenting standard telecommunications protocols employing these techniques. In some instances, a chronological sequence of approaches to a problem is given to illustrate the manner in which new designs evolve from earlier ones in order to solve problems with the earlier versions or extend their capabilities. More than 240 problems are included; a few are design problems for which a variety of reasonable solutions are possible. A detailed solutions manual is available from the publisher to instructors adopting the text for classroom use.

The field of telecommunications is evolving so rapidly that writing a text to cover the field has proved to be an extremely challenging task. Despite extensive efforts at standardization, there are major differences between different telecommunication architectures, aggravated by inconsistencies in documentation, terminology, and so forth. We have attempted to be more consistent in terminology than the standard literature, at times deviating from terminology in manuals to achieve consistency and make it easier to appreciate similarities and differences between architectures. This may make it slightly more difficult to learn the accepted terminology in the manuals for a particular architecture, but we hope it may prove to be an initial step in achieving more uniform terminology to be used in the field.[3]

The book covers communications aspects of the eight telecommunication architectures we feel are most important: the OSI Reference Model; U. S. Department of Defense architecture; IBM's Systems Network Architecture (SNA); Digital Equipment Corporation's Digital Network Architecture (DNA); the architecture for local area networks (LANs) being developed by the Institute of Electrical and Electronics Engineers; General Motors' Manufacturing Automation Protocol (MAP); Boeing Computer Services' Technical and Office Products System (TOP); and the Integrated Services Digital Network (ISDN) architecture being developed by telephone companies throughout the world.

To the extent possible, our discussions of telecommunication architectures have been based on original sources, rather than on secondary sources. Some secondary sources we have consulted contain disturbing numbers of inconsistencies. Most of the figures and tables from other sources have been adapted to be in a format consistent with the rest of the text, and sometimes to make corrections to the originals. The text also includes corrections to some common errors in the literature, such as statements of Shannon's channel capacity theorem that do not mention accuracy of data transmission (the key point of the theorem).

---

[3] Despite our attempt to use more uniform terminology than the literature, we faced problems in trying to create a consistent index, since there are still numerous points where different terminology has been used for essentially the same concept. In some cases we have resolved this by including cross references (e.g. references from fragmentation to segmentation and vice versa), while in other instances we have included index references for pages where basic concepts are discussed, regardless of whether the exact words in the index are used on the pages or not.

Table 1 Sections to Be Included in Protocols Course.

3.4– 3.10	
4.1– 4.10	
7.1– 7.7	7.9
9.1– 9.9	
10.1–10.4	
11.1–11.6	
12.3–12.5	

Table 2 Sections to Be Included in Design Problems Course.

2.1–2.7
3.1–3.5
5.1–5.4
6.1–6.6
8.1–8.12

Essentially all of the material in the book can be covered in a two-semester course. For a one-semester course, some selectivity in material covered is needed. Suggested course outlines and guidelines for selecting material from individual chapters to be covered are given below. An excellent course sequence would start with a course based on this text followed by a course explicitly oriented to telecommunication network performance modeling, using texts such as [SCHW87] or [BERT87]. Our current program at Clemson University follows this sequence.

The major single challenge in using this book as a text for a course lasting less than a full year is selecting an appropriate subset of the material to be covered. There are at least two possible one-semester courses that could be based on this material: one primarily devoted to telecommunication network protocols and one primarily devoted to standard approaches to solving the telecommunication network design problems. Each course would start with a general outline, such as that in Chapter 1; additional sections to be included are listed in Tables 1 and 2. Some additional selectivity may be necessary to limit the number of protocol architectures covered in the protocols course or the number of design approaches treated in the design problems course and guidelines for this are given below. The sections and subsections listed for the two courses are virtually disjoint, but skimming material in the complementary course would help students appreciate the significance of topics in the selected course. Instructors who wish to emphasize mathematical modeling techniques should also consider including the sections and subsections in Table 3. Advanced sections and subsections, indicated by □ in the margin,

Table 3 Material Emphasizing Mathematical Modeling Techniques.

2.4–2.7
3.25–3.2.6
4.11
5.3–5.4
6.3–6.5
8.4–8.6
11.1

Table 4 Material Requiring Mathematical Sophistication.

2.4.3–2.4.4	2.4.6	2.5.4	
3.2.6			
4.11			
5.3–5.4			
6.3.6–6.3.7	6.3.11	6.4.4	6.5
7.8			
8.6	8.12		

may be omitted for courses with students who have limited mathematical sophistication. These sections and subsections are listed in Table 4.

Some instructors teaching the "protocols course" may want to limit the number of different protocol architectures covered. If so, we have the following suggestions.

- The first five sections of Chapter 3 contain material essential for a good understanding of telecommunication networks, though parts of Section 3.5 on the structure of networks can be skimmed. Treatment of Sections 3.6–3.10, which introduce the various network architectures discussed in the text, can be limited to cover only those architectures selected for coverage in a particular course.
- A detailed presentation of only one of the three physical layer interfaces treated in Chapter 4, with brief sketches of the other two, is adequate for a first course. The interface chosen can be a matter of personal preference: EIA-232D/RS-232C/V.24 is by far the most common, but it is based on obsolescent technology; X.21 is the standard for the OSI Reference Model architecture but is less common; I.430/I.431 represent the state of the art but are just coming in. A treatment of all three is an excellent introduction to how telecommunications has evolved as designers learn from experience.
- The computer interfaces and communications channel interfaces treated in Sections 4.9 and 4.10 are essential, but are not often considered to be

part of networking architectures; neither is officially part of the standard architectures we treat. Instructors wishing to limit their treatment to these networking architectures may omit this material, but we have found that it clarifies essential steps in the communications process.

- A choice of DLC protocols in Chapter 7 should be based on the following considerations. Start-stop DLCs are the oldest and most rudimentary, but may still be the most common. Bisync was the first really successful synchronous protocol and is still widely used. ARPANET DLC involves interesting and useful modifications of Bisync but it is not widely used. DDCMP and the bit-oriented protocols represent the current generation of DLC protocols, with HDLC and variants most prominent in standards. A survey of all protocols treated gives a useful picture of how protocols evolve with experience.
- A choice of network layer protocols in Chapter 9 should be consistent with coverage of protocol architectures treated throughout a course. Portions of the discussion of X.25, such as the treatment of important X.25 based vendor networks, have been relegated to the appendix since their coverage is not essential.
- If it is not practical to cover all of the material on internetworking in Chapter 10, we suggest that at least one of the two types of bridges in Section 10.2 be discussed, along with at least an introduction to both types of Internet Protocols in Section 10.3. Brief treatments of the other topics would be worthwhile.
- The most important transport layer protocols, from Chapter 11, are TCP and the TP4 version of the OSI transport layer. The corresponding DNA and SNA layers should be covered if DNA and SNA are examined throughout a course.
- Changes to the ISDN architecture in Chapter 12 are currently being made more rapidly than changes in any other architecture treated. Hence, we recommend that treatment of this material be supplemented with up-to-date material from the current literature. On the other hand, this architecture to a large extent represents the direction in which future development of the telecommunications field is moving.

Corresponding comments about material to cover in a "design problems" course are as follows.

- If students already have a strong communications background, the communications material in Chapter 2 (Sections 2.4–2.6) may be briefly skimmed or even completely eliminated. Alternatively, instructors (especially those who are not engineers) may prefer to skim over this material as not being of major importance to the course they are teaching.
- We suggest that at least an overview of the major classes of medium access control algorithms in Section 5.2 be treated. Only a brief summary

of the performance modeling results in Section 5.3 needs to be included, however, especially if students have limited mathematical backgrounds.

- We suggest that the basic fundamentals of error control coding from Chapter 6 be treated, but a detailed understanding of the various types of codes discussed in Subsections 6.3.6–6.3.12 is not necessary.
- It may be impossible to cover all of the routing algorithms in Chapter 8 in a short period of time. Either the Bellman-Ford-Moore algorithm in Subsection 8.4.1 or the Dijkstra algorithm in Subsection 8.4.2 should be covered in reasonable detail, but the other could be treated more briefly. The Floyd-Warshall algorithm in Subsection 8.4.3 is more complex and less common, so it could be omitted. The sections on optimal routing and combined optimal routing and flow control are the most mathematically sophisticated sections of the entire text and are not necessary for an understanding of later material.
- At least the essentials of the sliding window and buffer allocation approaches to flow control in Sections 8.9 and 8.10 should be covered, but not all variants need to be examined. The variants to be covered may be selected by instructors according to their personal preferences.

A number of people have helped in the preparation of this book. Sections of the text have been written by Joe Hammond and Krys Pawlikowski, with the major portion written by John Spragins. All three have edited various sections of the text. Most of the book has been written at or around Clemson University, where the material has been taught in a course taken by both senior and graduate students. A major part of John's material was written during a semester's leave of absence from Clemson plus summers devoted to the book, with a minor part written during a summer (winter there) at the University of Canterbury in New Zealand and final details completed during a sabbatical at Hewlett Packard Laboratories, Bristol, England. Much of Joe's material was written during summers at his mountain cabin, and a large part of Krys' contributions were completed during a sabbatical (from the University of Canterbury) spent at Clemson University.

We would like to express our appreciation to the following referees for their careful reviews: Mostaffa H. Ammar, Georgia Institute of Technology; Robert E. Bradley, Digital Equipment Corp.; Mario Gerla, UCLA; Tony Hsiao, Purdue University; Thomas Robertazzi, SUNY at Stony Brook; and B. W. Stuck, Business Strategies. The manuscript has benefitted greatly from their reviews, plus reviews by others who remain anonymous.

Typing and other secretarial work by B. Rathz and R. Watkins are also appreciated. Others who have helped in preparation of the manuscript and in other ways include A. Aletty, S. S. Bijjahalli, M. Den, N. El-Khoury, H. Jafari, S. Kamat-Timble, W. P. Lyui, P. Moghe, M. Raman, J. Song, D. Stevens, H. R. Wang, J. Ward and W. P. Yeh. A number of persons at Addison-Wesley gave us valuable assistance, including Tom Robbins and Barbara Rifkind—editorial assistance, Stephanie Argeros-Magean—copyediting, Joseph Vetere—art coordination, and Helen Wythe—production supervision. Finally, I would like to thank my wife, Cathy, and my children, Katrina, Jennifer and Matthew for their encouragement and to apologize for the way manuscript preparation impacted the time I had available for them. Cathy's support has been especially valuable, including serving as primary breadwinner for our family during the half year when I took leave of absence from my job to work on the book. Joe and Krys also wish to thank their wives, Edith and Barbara, respectively, for their support.

J. D. S.
Clemson, S. C.

Contents

Chapter 7 Data Link Layer Protocols *303*

Chapter 8 Routing and Flow Control *355*