

COMPUTER NETWORKS



Prentice/Hall International editions

ANDREW S. TANENBAUM

COMPUTER NETWORKS

ANDREW S. TANENBAUM

*Vrije Universiteit
Amsterdam, The Netherlands*



Prentice/Hall International, Inc.

Library of Congress Cataloging in Publication Data

Tanenbaum, Andrew S. (date)
Computer networks.

(Prentice-Hall software series)

Bibliography: p.

Includes index.

1. Computer networks.	I. Title	II. Series.
TK5105.5.T36	001.64'404	80-23283
ISBN 0-13-164699-0		

PRENTICE-HALL SOFTWARE SERIES
Brian W. Kernighan, advisor

©1981 by Prentice-Hall, Inc., Englewood Cliffs, N.J. 07632

All rights reserved. No part of this book
may be reproduced in any form or by any means
without permission in writing from the publisher.

10 9 8 7 6 5 4 3

This edition may be sold only in those countries
to which it is consigned by Prentice-Hall
International. It is not to be re-exported
and is not for sale in the U.S.A., Mexico
or Canada.

Printed in the United States of America

Editorial/production supervision: Nancy Milnamow
Manufacturing buyer: Joyce Levatino

PRENTICE-HALL INTERNATIONAL, INC., *London*
PRENTICE-HALL OF AUSTRALIA PTY. LTD., *Sydney*
EDITORA PRENTICE-HALL DO BRASIL, LTDA, *Rio de Janeiro*
PRENTICE-HALL OF CANADA, LTD., *Toronto*
PRENTICE-HALL OF INDIA PRIVATE LTD., *New Delhi*
PRENTICE-HALL OF JAPAN, INC., *Tokyo*
PRENTICE-HALL OF SOUTHEAST ASIA PTE. LTD., *Singapore*
WHITEHALL BOOKS LTD., *Wellington, New Zealand*
PRENTICE-HALL, INC., *Englewood Cliffs, New Jersey*

COMPUTER NETWORKS

To Suzanne, Barbara, Marvin, and Sweetie π

PREFACE

As computers have become smaller, cheaper, and more numerous, people have become more and more interested in connecting them together to form networks and distributed systems. At first these connections were made in an ad hoc way, typically with each computer regarding the other ones as terminals. However, in the past decade, a substantial body of knowledge has developed on the subject of computer networking, so that future networks and distributed systems can be designed systematically.

The key to designing a computer network was first enunciated by Julius Caesar: Divide and Conquer. The idea is to design the network as a sequence of layers, or abstract machines, each one based upon the previous one. By reducing the study of the whole to the study of its parts, the subject becomes more manageable. This book uses a model in which networks are divided into seven layers. The structure of the book follows the structure of the model to a considerable extent.

Chapter 1 provides an introduction to the subject of computer networks in general and layered protocols in particular. Chapter 2 treats a number of algorithms and heuristics for designing the topology of a network. Chapter 3 begins the study of the seven-layer model with the bottom layer (the physical layer); it covers the architecture of data communication systems, including the telephone system and satellites. Chapter 4 is primarily concerned with data link layer protocols, algorithms for reliably transmitting data over unreliable lines. Chapters 5, 6, and 7 all deal with the network layer; Chap. 5 treats point-to-point

networks, Chap. 6 treats satellite and packet radio networks, and Chap. 7 treats local networks. Chapter 8 studies the transport and session layers, in particular, end-to-end protocols and internetworking. Chapter 9 is about the presentation layer, including cryptography, text compression, virtual terminal protocols, and file transfer protocols. Chapter 10 provides an introduction to some application layer issues, primarily distributed data bases and distributed operating systems. Chapter 11 contains a reading list and bibliography.

The book is intended as a text for juniors, seniors, and graduate students in computer science, electrical engineering, and related disciplines. The only prerequisites are a general familiarity with computer systems and programming, although a little knowledge of elementary calculus and elementary probability theory is useful, but not essential. Since the amount of material in the book may be too much for a one semester course, depending on the level of the students, I have made a serious attempt to make each chapter relatively independent of the other ones. This way an instructor is free to skip whichever chapters he chooses. Chapter 2, in particular, which is long and highly technical, may be omitted without loss of continuity by students who are not concerned with the topology design problem. However, if Chap. 2 is omitted, I would recommend having the students look at Eq. (2-9) and Eq (2-10), since they are used in a few places in subsequent chapters to analyze network performance.

People who have no formal computer science background, but who have had some industrial experience with minicomputers, assembly languages, system programming languages, operating systems, or data communication should have little trouble understanding the book. Even programmers or managers unfamiliar with these areas can probably understand a substantial part of the book.*

Many people have helped me with this book. I would first like to thank Wiebren de Jonge for tens, if not hundreds, of hours of questions, discussions, arguments, and polemics about nearly every page of the manuscript. His impatience with my occasional indolence has shamed me into rewriting sections over and over until I finally got them right.

I would also like to thank Peter Apers, Dick Binder, Imrich Chlamtac, Dave Clark, George Conant, Ira Cotton, Rudy Cypser, Yogen Dalal, Dixon Doll, Phil Enslow, Howard Frank, Bill Franta, James P. Gray, Paul Green, Jan Hajek, Doug Jensen, Steve Johnson, Haim Kilov, Leonard Kleinrock, Simon Lam, Tony Lauck, Mike Liu, Alex McKenzie, Dave Morgan, Bob Morris, Holger Opderbeck, Robert Ryan, Paul Santos, Phil and Debbie Scherrer, Mischa Schwartz, John Shoch, Johan Stevenson, Carl Sunshine, Jim van Keulen, Stu Wecker, Barry Wessler, and Sylvia Wilbur for their help. Their collective assistance has greatly improved the book in many ways. My students, especially Dick Biekart, Herman Gerbscheid, and Jan de Ruitter have also been very helpful.

To all the members of the Computing Science Research Center at Bell

*A manual with the problem solutions is available from Prentice-Hall.

Laboratories, thanks for many stimulating discussions during my visit there in the summer of 1979. I would especially like to thank Sandy Fraser for arranging things, Brian Kernighan and Mike Lesk for initiating me into the mysteries of the CAT, and Lorinda Cherry for causing a certain large minicomputer to comment extensively about my diction.

At this point authors normally thank their typists. Not wanting to break with this worthy tradition, I would like to thank Andy Tanenbaum for his fast and accurate typing of innumerable versions of the manuscript. Somewhat more to the point, I would like to thank Ken Thompson, Dennis Ritchie, and the other members of the Computing Science Research Center at Bell Laboratories for developing the UNIX[†] operating system. The UNIX text processing tools made the typing and retyping of the 1,231,788 characters comprising the final manuscript a pleasure instead of a chore.

IBM and DEC have permitted me to use material from their SNA and DECNET protocol manuals, respectively, for which I am grateful.

Last but not least, I would like to thank Suzanne, Barbara, Marvin, and Sweetie π , Suzanne for her patience, especially since she knew what she was getting in for this time, Barbara for keeping my terminal free of peanut butter, Marvin for arriving at a propitious moment, and Sweetie π for being very quiet while I was writing.

ANDREW S. TANENBAUM

**COMPUTER
NETWORKS**

CONTENTS

PREFACE xiii

1	INTRODUCTION	1
1.1	THE USES OF COMPUTER NETWORKS	3
1.1.1.	Network Goals	3
1.1.2.	Applications of Networks	5
1.2	NETWORK STRUCTURE	7
1.3	NETWORK ARCHITECTURES	10
1.3.1.	Protocol Hierarchies	10
1.3.2.	Design Issues for the Layers	14
1.4	THE ISO REFERENCE MODEL	15
1.4.1.	The Physical Layer	16
1.4.2.	The Data Link Layer	17
1.4.3.	The Network Layer	17
1.4.4.	The Transport Layer	18
1.4.5.	The Session Layer	19
1.4.6.	The Presentation Layer	20
1.4.7.	The Application Layer	21
1.5	ARPANET, SNA, DECNET, AND PUBLIC NETWORKS	21
1.5.1.	Introduction to the ARPANET	22
1.5.2.	Introduction to SNA	23
1.5.3.	Introduction to DECNET	26
1.5.4.	Introduction to Public Networks and X.25	28
1.6	OUTLINE OF THE REST OF THE BOOK	29
1.7	SUMMARY	30
2	NETWORK TOPOLOGY	32
2.1	INTRODUCTION TO THE TOPOLOGY DESIGN PROBLEM	32
2.1.1.	Formulation of the Problem	32
2.1.2.	Hierarchical Networks	34
2.2	CONNECTIVITY ANALYSIS	36
2.2.1.	Introduction to Graph Theory	36
2.2.2.	Cuts and Network Flow	40
2.2.3.	The Max-Flow Algorithm	44

2.2.4.	Disjoint Paths	47
2.2.5.	Monte Carlo Connectivity Analysis	54
2.3	DELAY ANALYSIS	56
2.3.1.	Introduction to Queueing Theory	57
2.3.2.	The M/M/1 Queue in Equilibrium	59
2.3.3.	Networks of M/M/1 Queues	62
2.4	BACKBONE DESIGN	67
2.4.1.	The Design Process	67
2.4.2.	Generating Starting Topologies	70
2.4.3.	Flow and Capacity Assignment	72
2.4.4.	Perturbation Heuristics	75
2.5	LOCAL ACCESS NETWORK DESIGN	80
2.5.1.	Assigning Sites to Concentrators	80
2.5.2.	The Concentrator Location Problem	83
2.5.3.	The Terminal Layout Problem	84
2.6	SUMMARY	87

3 THE PHYSICAL LAYER

91

3.1	THE THEORETICAL BASIS FOR DATA COMMUNICATION	91
3.1.1.	Fourier Analysis	91
3.1.2.	Bandwidth Limited Signals	92
3.1.3.	The Maximum Data Rate of a Channel	95
3.2	THE TELEPHONE SYSTEM	96
3.2.1.	Who's Who in the Telecommunication World	97
3.2.2.	Structure of the Telephone System	98
3.2.3.	The Local Loop	100
3.3	TRANSMISSION AND MULTIPLEXING	103
3.3.1.	Frequency Division and Time Division Multiplexing	103
3.3.2.	Digital Transmission	104
3.3.3.	The X.21 Digital Interface	108
3.3.4.	Communication Satellites	110
3.3.5.	Circuit Switching and Packet Switching	114
3.4	TERMINAL HANDLING	119
3.4.1.	Polling	119
3.4.2.	Multiplexing versus Concentration	121
3.4.3.	The Packet Assembler/Disassembler	122
3.5	ERRORS	125
3.5.1.	The Nature of Transmission Errors	125
3.5.2.	Error-Correcting Codes	126
3.5.3.	Error-Detecting Codes	128
3.6	SUMMARY	133

4	THE DATA LINK LAYER	136
4.1	ELEMENTARY DATA LINK PROTOCOLS	136
4.1.1.	Some Declarations Needed by the Protocols	137
4.1.2.	An Unrestricted Simplex Protocol	141
4.1.3.	A Simplex Stop-and-Wait Protocol	143
4.1.4.	A Simplex Protocol for a Noisy Channel	145
4.2	SLIDING WINDOW PROTOCOLS	148
4.2.1.	A One Bit Sliding Window Protocol	151
4.2.2.	A Protocol with Pipelining	153
4.2.3.	A Protocol That Accepts Frames Out of Order	157
4.3	EXAMPLES OF THE DATA LINK LAYER	165
4.3.1.	The Data Link Layer in the ARPANET	165
4.3.2.	The Data Link Layer in SNA and X.25	167
4.3.3.	The Data Link Layer in DECNET	172
4.4	ANALYSIS OF PROTOCOLS	174
4.4.1.	Protocol Efficiency	174
4.4.2.	Protocol Verification	177
4.5	SUMMARY	183
5	THE NETWORK LAYER I: POINT-TO-POINT NETWORKS	187
5.1	VIRTUAL CIRCUITS AND DATAGRAMS	187
5.1.1.	The Service Provided by the Network Layer	188
5.1.2.	Comparison of Virtual Circuit and Datagram Service	189
5.1.3.	The Internal Structure of the Subnet	192
5.1.4.	Comparison of VCs and Datagrams within the Subnet	195
5.1.5.	Independence of Subnet Service and Subnet Structure	196
5.2	ROUTING ALGORITHMS	197
5.2.1.	Flooding	198
5.2.2.	Static Routing	199
5.2.3.	Centralized Routing Algorithms	200
5.2.4.	Isolated Routing	202
5.2.5.	Distributed Routing Algorithms	205
5.2.6.	The Topology Update Problem	207
5.2.7.	Hierarchical Routing	211
5.2.8.	Broadcast Routing	213
5.3	CONGESTION	215
5.3.1.	Preallocation of Buffers	216
5.3.2.	Packet Discarding	217
5.3.3.	Isarithmic Congestion Control	219
5.3.4.	Flow Control	220
5.3.5.	Choke Packets	221
5.3.6.	Deadlocks	222

5.4	EXAMPLES OF THE NETWORK LAYER	225
5.4.1.	The Network Layer in the ARPANET	226
5.4.2.	The Network Layer in SNA	231
5.4.3.	The Network Layer in DECNET	235
5.4.4.	The Network Layer in X.25	237
5.5	SUMMARY	245

6 THE NETWORK LAYER II: SATELLITE AND PACKET RADIO NETWORKS 249

6.1	SATELLITE PACKET BROADCASTING	250
6.1.1.	Conventional Channel Allocation Methods	251
6.1.2.	Pure ALOHA and Slotted ALOHA	253
6.1.3.	Finite Population ALOHA	257
6.1.4.	Delay and Throughput of Slotted ALOHA	259
6.1.5.	Stability of Slotted ALOHA	265
6.1.6.	Controlled ALOHA	269
6.1.7.	Reservation ALOHA	271
6.2	PACKET RADIO	273
6.2.1.	The University of Hawaii ALOHA System	273
6.2.2.	Design Issues for Packet Radio Networks	277
6.3	SUMMARY	282

7 THE NETWORK LAYER III: LOCAL NETWORKS 286

7.1	CARRIER SENSE NETWORKS	288
7.1.1.	Persistent and Nonpersistent CSMA	289
7.1.2.	Ethernet	292
7.1.3.	Collision-Free Protocols	296
7.1.4.	Limited-Contention Protocols	300
7.2	RING NETWORKS	307
7.2.1.	Token Rings	307
7.2.2.	Contention Rings	311
7.2.3.	Slotted Rings	312
7.2.4.	Register Insertion Rings	313
7.3	SHARED MEMORY SYSTEMS	315
7.3.1.	Processor-Memory Interconnection	315
7.3.2.	Examples of Shared Memory Systems	317
7.4	SUMMARY	320

8	THE TRANSPORT AND SESSION LAYERS	324
8.1	TRANSPORT PROTOCOL DESIGN ISSUES	325
8.1.1.	Transport Service	326
8.1.2.	Addressing and Connection Establishment	335
8.1.3.	Flow Control and Buffering	338
8.1.4.	Multiplexing	343
8.1.5.	Synchronization in the Presence of Delayed Packets	345
8.1.6.	Crash Recovery	351
8.2	INTERCONNECTION OF PACKET-SWITCHING NETWORKS	353
8.2.1.	Gateways	354
8.2.2.	The Level of Interconnection	358
8.2.3.	The X.75 Model versus the Datagram Model	359
8.2.4.	Internetwork Packet Fragmentation	364
8.3	THE SESSION LAYER	368
8.4	EXAMPLES OF THE TRANSPORT AND SESSION LAYERS	369
8.4.1.	The Transport Layer in the ARPANET	369
8.4.2.	The Transport and Session Layers in SNA	377
8.4.3.	The Transport Layer in DECNET	380
8.5	SUMMARY	381
9	THE PRESENTATION LAYER	386
9.1	NETWORK SECURITY AND PRIVACY	386
9.1.1.	Traditional Cryptography	388
9.1.2.	The Data Encryption Standard	396
9.1.3.	The Key Distribution Problem	406
9.1.4.	Public Key Cryptography	410
9.1.5.	Authentication and Digital Signatures	413
9.2	TEXT COMPRESSION	417
9.2.1.	Encoding a Finite Set of Equally Likely Symbols	418
9.2.2.	Huffman Coding	419
9.2.3.	Context Dependent Encoding	420
9.3	VIRTUAL TERMINAL PROTOCOLS	421
9.3.1.	Classes of Terminals	421
9.3.2.	The Data Structure Model	423
9.3.3.	Design Principles	425
9.3.4.	An Example Virtual Terminal Protocol	426
9.4	FILE TRANSFER PROTOCOLS	429
9.5	EXAMPLES OF THE PRESENTATION LAYER	431
9.5.1.	The Presentation Layer in the ARPANET	431
9.5.2.	The Presentation Layer in SNA	433
9.5.3.	The Presentation Layer in DECNET	434
9.6	SUMMARY	435

10	THE APPLICATION LAYER	440
10.1	DISTRIBUTED DATA BASE SYSTEMS	440
10.1.1.	The Relational Data Base Model	441
10.1.2.	The Relation Distribution Problem	446
10.1.3.	Query Processing	448
10.1.4.	Concurrency Control	452
10.1.5.	Crash Recovery	457
10.2	DISTRIBUTED COMPUTATION	460
10.2.1.	The Hierarchical Model	461
10.2.2.	The CPU Cache Model	462
10.2.3.	The User-Server Model	464
10.2.4.	The Pool Processor Model	466
10.2.5.	The Data Flow Model	468
10.3	NETWORK AND DISTRIBUTED OPERATING SYSTEMS	476
10.3.1.	Network Operating Systems	476
10.3.2.	Distributed Operating Systems	479
10.3	SUMMARY	482
11	READING LIST AND BIBLIOGRAPHY	486
11.1	SUGGESTIONS FOR FURTHER READING	486
11.1.1	Introduction	487
11.1.2	Network Topology	487
11.1.3	The Physical Layer	488
11.1.4	The Data Link Layer	488
11.1.5	The Network Layer I: Point-to-Point Networks	489
11.1.6	The Network Layer II: Satellite and Packet Radio Networks	489
11.1.7	The Network Layer III: Local Networks	490
11.1.8	The Transport and Session Layers	490
11.1.9	The Presentation Layer	491
11.1.10	The Application Layer	492
11.2	ALPHABETICAL BIBLIOGRAPHY	493
	INDEX	508

1

INTRODUCTION

Each of the past three centuries has been dominated by a single technology. The eighteenth century was the time of the great mechanical systems accompanying the industrial revolution. The nineteenth century was the age of the steam engine. During the twentieth century, the key technology has been information gathering, processing, and distribution. Among other developments we have seen the installation of worldwide telephone networks, the invention of radio and television, the birth and unprecedented growth of the computer industry, and the launching of communication satellites.

As we move toward the final years of this century, these areas are rapidly converging, and the differences between collecting, transporting, storing, and processing information are quickly disappearing. Organizations with hundreds of offices spread over a wide geographical area routinely expect to be able to inspect the current status of even their most remote outpost at the push of a button. As our capability to gather, process, and distribute information grows, the demand for even more sophisticated information processing grows even faster.

Although the computer industry is young compared to other industries (e.g., automobiles and air transportation), computers have made spectacular progress in a short time. During the first two decades of their existence, computer systems were highly localized, usually within a single large room. Not infrequently, this room had glass walls, through which visitors could gawk at the great electronic wonder inside. A medium size company or university

might have had one or two computers, with large ones having had at most a few dozen. The idea that within 20 years equally powerful computers smaller than postage stamps would be mass produced by the millions was pure science fiction.

The merging of computers and communications has had a profound influence on the way computer systems are organized. The concept of the “computer center” as a room with a large computer to which users bring their work for processing is rapidly becoming obsolete. This model has not one, but two flaws in it: the concept of a single large computer doing all the work, and the idea of users bringing work to the computer, instead of bringing the computer to the users.

The old model a single computer serving all of the organization’s computational needs is rapidly being replaced by one in which a large number of separate but interconnected computers do the job. These systems are called **computer networks**. The design and analysis of these networks is the subject of this book.

Throughout the book we will use the term “computer network” to mean an *interconnected* collection of *autonomous* computers. Two computers are said to be interconnected if they are capable of exchanging information. The connection need not be via a copper wire; lasers, microwaves, and earth satellites can also be used. By requiring that the computers be autonomous, we wish to exclude from our definition systems in which there is a clear master/slave relation. If one computer can forcibly start, stop, or control another one, the computers are not autonomous. An ILLIAC IV type of system, with one control unit and many slaves, is not a network. Nor is a large computer with remote card readers, printers, and terminals a network.

There is considerable confusion in the literature between a computer network and a **distributed system**. Enslow’s (1978) definition requires a distributed system to have a *system-wide* operating system, with services requested by name, and not by location. In other words, the user of a distributed system should not be aware that there are multiple processors; it should look like a virtual uniprocessor. Allocation of jobs to processors, processor scheduling, allocation of files to disks, movement of files between where they are stored and where they are needed, and all other system function must be automatic.

On the other hand, Liebowitz and Carson (1978) have said: “A distributed system is one in which the computing functions are dispersed among several physical computing elements.” Obviously this definition includes many systems that Enslow’s excludes.

In our view a distributed system is a special case of a network, one with a high degree of cohesiveness and transparency. In essence a network may or may not be a distributed system, depending on how it is used. However, in view of the lack of any generally accepted nomenclature, we will use the term “computer network” in a generic sense, to cover both computer networks and distributed systems.