

TP274
M381

8053189

NETWORKS AND DISTRIBUTED PROCESSING



**SOFTWARE,
TECHNIQUES,
AND ARCHITECTURE**

JAMES MARTIN



E8053189

PRENTICE-HALL, INC., Englewood Cliffs, N.J. 07632

Library of Congress Cataloging in Publication Data

MARTIN, JAMES, (date)

Computer networks and distributed processing:
software, techniques, architecture.

Bibliography: p.

Includes index.

1. Electronic data processing—Distributed
processing. 2. Computer networks. I. Title.

QA76.9.D5M37 001.64 80-20652

ISBN 0-13-165258-3

Computer Networks and Distributed Processing:
Software, Techniques, and Architecture
James Martin

©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.

Editorial/production supervision by *Karen J. Clemments*
Jacket design by *Alon Jaediker*
Manufacturing buyers: *Joyce Levatino and Gordon Osbourne*

Printed in the United States of America

10 9 8 7 6 5 4 3 2


PRENTICE-HALL INTERNATIONAL, INC., *London*
PRENTICE-HALL OF AUSTRALIA PTY. LIMITED, *Sydney*
PRENTICE-HALL OF CANADA, LTD., *Toronto*
PRENTICE-HALL OF INDIA PRIVATE LIMITED, *New Delhi*
PRENTICE-HALL OF JAPAN, INC., *Tokyo*
PRENTICE-HALL OF SOUTHEAST ASIA PTE. LTD., *Singapore*
WHITEHALL BOOKS LIMITED, *Wellington, New Zealand*

488500



COMPUTER

**COMPUTER NETWORKS
AND
DISTRIBUTED PROCESSING:
SOFTWARE,
TECHNIQUES,
AND ARCHITECTURE**

A  BOOK

TO CORINTHIA



PREFACE

At certain times a new technology emerges which is destined to change the nature of data processing: stored program control in the early 1950's, magnetic tape in the mid-1950's, large on-line storage in the 1960's, and the use of terminals. Today a technology of immense importance is spreading: *distributed processing in which intelligent machines in different locations cooperate by means of networks*.

This new and vital direction is throwing DP management into one of the worst dilemmas they have faced in the history of computing.

The choices that analysts and managers must make in this new area are complex and have long-lasting implications. Furthermore the financial implications of making the right choice are great. Some corporations appear to the author to be taking the wrong course. He has attempted to estimate the eventual cost of this. There can be little doubt that it will cost millions of dollars in some corporations, in abandoned approaches, redesign, and program rewriting. And this does not count the lost opportunities and inability to obtain information needed by management because of the incompatibility and nonconnectability of separate systems, minicomputers and intelligent terminals.

Large computers, minicomputers, and microcomputers in intelligent terminals are becoming interconnected into all manner of configurations. The corporation of the near future will be laced with networks which handle not only its data processing but also its word processing, mail and message sending. End users on the shop floor, in the sales offices, in the planning departments—in fact everywhere—are beginning to perceive what is happening and are demanding a piece of the action. They want a minicomputer, or an intelligent terminal, or access to distant data bases.

Unfortunately the technology of distributed networks is very complicated. It is easy to perceive the idea of networks, but very difficult to understand the technical subtleties. The subtleties can and must be hidden from the end user. The user perceives merely the dialogue or simple procedure that is provided for him. But that requires network architectures.

Network architectures specify the protocols or sets of rules which are needed to interconnect computers. These protocols are highly complex because there must be precise cooperation between distant intelligent machines.

Herein lies the dilemma. *The various architectures are incompatible.* The network architectures of IBM, Univac, DEC, NCR, Hewlett Packard, etc., are entirely different. Machines from different manufacturers cannot in general be hooked together (in spite of claims to the contrary) without bypassing the architecture and losing its advantages.

To make matters worse, major common carriers are developed or have developed *their* architecture for networks. By far the most spectacular of these is AT&T's A.C.S. (Advanced Communications Service), formally called BDN (Bell Data Network). However, different common carriers have different architectures; for example, Telenet, SBS (the communications satellite subsidiary of IBM, Comsat and Aetna), and PTT's in Europe, Japan and elsewhere (P.T.T. refers to a national telephone administration usually government controlled.)

The networking architectures of the common carriers are fundamentally incompatible with those of some of the computer manufacturers. AT&T's impressive ACS is in head-on collision with IBM's mainstream direction in networking: SNA (Systems Network Architecture). One reason for this is that both perceive networking as *their* territory. AT&T thinks that data networks, like telephone networks, ought to be provided by the carrier. Data concentrators, switches, front-ends, intelligent terminals and terminal cluster controllers are considered the data equivalent of telephone concentrators, PBX's, local switching offices, trunk switching offices and telephone instruments. IBM, on the other hand, derives great revenue from its terminals, cluster controllers, 3705's, network equipment and the many instructions executed in the mainframe to support it. It does not want anyone else dictating how its distributed machines should be interconnected. A great territorial fight is in the making.

Meanwhile DP analysts and executives must make decisions about their corporate networks and how to implement distributed processing.

The end users are clamoring for distributed processing. They have the minicomputer salesman beating at the door. DP management knows that the scattered minicomputers must be hooked together—that 60% of the data stored in one location will be needed in other locations. Sometimes they ignore the problem and hope that new software will emerge that will tie the fragments together. They hear the ARPANET connected incompatible machines. They hear that new standards are emerging like HDLC and X.25, X.3, X.28 and X.29.

The new standards will help, but not completely and not quickly. There is no resemblance between the major architectures from manufacturers and X.25, X.3, etc. Even if there were, neither ARPANET nor the current standards would solve the most important of the incompatibility problems—the Layer 4 control mechanisms described in this book. The best that can be hoped for is a clumsy bridge at the transport subsystem level, but many of the worst problems are external to the transport subsystem.

What should a DP manager do?

He might close his eyes and do nothing. In this case he may deny his organization the very considerable advantages of the new machines. The end users may rebel, as they are doing in many places, and get their own machines.

He may allow a proliferation of incompatibility without counting the future cost.

He could select one manufacturer's distributed architecture and stick with it. There is much to be said for this if the product line in question is sufficiently diverse. The manufacturers with major distributed processing architectures have planned a "migration path" into the future in which many new products will conform to the architecture and can be installed with the minimum disruption to existing application programs. But he may fear being locked into one manufacturer.

He could insist that the computer network conform to the new standards, especially X.25, but this would lock out most of the products of major manufacturers, which are not compatible with X.25.

He may plan a judicious mixture of these approaches, deciding where compatibility is vital, where it can be dispensed with, and where it is possible to build a bridge between different incompatible systems. He may select a manufacturer's architecture such as SNA for a certain set of applications, and a common carrier architecture like X.25 or ACS for all other interconnections.

Whatever he does, he should plan and understand the future implications of his choices. He or his staff need to understand the technical tradeoffs in distributed processing networks. The cost of *not* understanding could be extremely high.

This book explains the technology of computer networks and distributed processing architectures. It will enable readers to understand the issues and perceive what is important in an architecture design. It indicates how the architectures are likely to evolve in the future. This information is vital in making decisions about how to implement distributed data processing.

Part 1 of this book discusses the different types of networks and distributed configurations.

Part 2 explains the concepts of architectures for distributed systems, and what are desirable features of the different layers of the architectures.

Part 3 describes the mechanisms that are used in networking, and particularly those in specifications for the new common carrier (PTT) networks such as the X.25 networks, their X.3/X.28/X.29 interface, and the Bell System ACS network.

Part 4 explains the requirements and mechanisms for dealing with errors, failures of different types, recovery, deadlocks, privacy and security.

Part 5 discusses the future directions of network architectures.

The subject is complex. Many readers will benefit from a fast scan through the book before detailed reading of the chapters which particularly concern them.

After scanning the book, a fast path could be taken as follows:

CHAPTER 2	The Trend to Distributed Processing
CHAPTER 3	Types of Distributed Systems
CHAPTER 4	Private Networks

CHAPTER 5	Public Networks
CHAPTER 6	AT&T's Advanced Communications Service
CHAPTER 7	Distributed Intelligence Time-Sharing
CHAPTER 9	Basic Network Functions
CHAPTER 11	Layers of Control
CHAPTER 12	Interfaces
CHAPTER 13	Physical Link Control
CHAPTER 14	The Transport Subsystem
CHAPTER 15	The Session Services Subsystem
CHAPTER 19	Packet Switching
CHAPTER 22	CCITT Recommendation X.25
CHAPTER 30	Errors and Recovery
CHAPTER 34	Deadlocks
CHAPTER 36	The Future of Network Architectures

To understand fully the mechanisms and tradeoffs, the whole book should be studied.

James Martin



CONTENTS

Preface *xi*

PART

THE PROMISE

- 1 The Promise of Computer Networks 2
- 2 The Trend to Distributed Processing 22
- 3 Types of Distributed Systems 36
- 4 Private Networks 65
- 5 Public Networks 87
- 6 AT&T's Advanced Communications Service 102
- 7 Distributed Intelligence Time-sharing 122
- 8 Types of Computer Networks 130

PART

**COMPUTER NETWORKS**

- 9** **Basic Network Functions** 138
- 10** **Logical, Virtual, and Transparent Facilities** 151
- 11** **Layers of Control** 158
- 12** **Interfaces** 181
- 13** **Physical Link Control** 202
- 14** **The Transport Subsystem** 210
- 15** **The Session Services Subsystem** 232
- 16** **Network Management** 250
- 17** **Standards, Agreements, and CCITT Recommendations** 260

PART

**NETWORK MECHANISMS**

- 18** **HDLC, SDLC, etc.** 276
- 19** **Packet Switching** 299
- 20** **Flow Control and Pacing** 312
- 21** **Routing Techniques** 326
- 22** **CCITT Recommendation X.25** 341
- 23** **The Packet Assembly/Disassembly Interface** 364
- 24** **Datagrams** 375
- 25** **Circuit Switching and Hybrid Networks** 387

- 26 Local Wideband Networks—Ethernet 408
- 27 Radio and Satellite Networks 418
- 28 Message Compression 440
- 29 Utilization Patterns of Computer Networks 452

PART IV ERRORS, FAILURES AND SECURITY

- 30 Errors and Recovery 464
- 31 Recovery at the Physical Link Layer 469
- 32 Transport Subsystem Recovery 480
- 33 Session Recovery 490
- 34 Deadlocks 505
- 35 Security and Privacy 517

PART V FUTURE

- 36 The Future of Network Architectures 544
- Index 559

PART I

THE PROMISE

1 THE PROMISE OF COMPUTER NETWORKS

In America's pioneering days, each log cabin in the wilderness was an isolated outpost in a disconnected world. The pioneering days of the computer are still with us. But a new technology is evolving that will change the computer world as much as America changed. It is the technology of distributed processing and computer networks. Instead of being isolated outposts, computers will use programs and data stored by computers in other locations.

The growth of the communications systems that will make this technology economical is taking place at the same time as the spread of vast numbers of minicomputers and microcomputers. The combination of these developments is fundamentally changing the patterns of data processing. Many jobs which used to be done on large, heavily shared computers can now be done on stand alone minicomputers or microcomputers. Data in central storage units can be shared by large numbers of dispersed users, many having a small computer which can process the data or provide an end-user dialogue which simplifies the access to information. Data of interest to one terminal user may reside in multiple distant computers, accessible via networks. When a computer program says GET, it may be accessing a file unit attached to a computer a thousand miles away.

PROTOCOLS

To enable the increasing variety of computers to communicate with one another, there must be rigorously defined protocols—i.e., sets of rules about how control messages and data messages are exchanged between machines, and how they control the communication process. Along with the protocol definition, the formats of the control messages and the headers and trailers of the data messages are likewise rigorously defined.

These formats and protocols become quite complex. It is desirable that there should be a widely accepted standard so that all manner of machines can communi-

cate. Unfortunately, several different sets of formats and protocols are emerging from different organizations, as we discuss later in the book.

Some formats and protocols relate just to the transmission process for transporting data from one machine to another via a complex network. Communications, however, consists of more than merely a data transport mechanism. It may be necessary for one machine to load a program into another, to specify how a remote file will be used, to convert from the procedures or characters of one machine to those of another, to specify security procedures, to specify how compressed data should be expanded or edited before printing, to determine how machine usage will be charged, and so on. In some cases an elaborate contract is drawn up between machines before they begin a session of interaction. These processes, which are separate from the transport process, are specified to a greater or lesser extent in the architectures for computer networks and distributed processing.

RESOURCE SHARING The purpose of many computer networks is to permit a far-flung community of users to share computer resources. Many such users now have their own minicomputers and calculators, so the shared resources have to be interesting enough to warrant access via a network. The facilities accessible by networks are in fact becoming more interesting at a rapid rate.

The remote computer may contain *software* which a user needs to employ. It may be proprietary software kept at one location. It may require a larger machine than any at the user's location. The distant computer may provide access to *data* which is stored and maintained at its location. Sometimes the remote machine controls a large or special printing facility. Sometimes the remote machine compiles programs which are used on smaller peripheral machines.

One of the best known resource-sharing networks is ARPANET, a United States network interconnecting more than fifty university and research computer centers (including a few in Europe). Many universities have interesting, and sometimes unique, computer facilities. These will become more valuable if they can be employed by a larger community of users. The idea behind ARPANET is that a professor or student at one university should be able to employ the facilities at any other university on the network. Usually he does this by means of a terminal, receiving responses from a distant computer almost as fast as if he were at the location of that machine.

A good example of a resource worth sharing is a program at M.I.T. for assisting with mathematics—MACSYMA. It is one of the largest programs written. It can solve simultaneous equations, factor polynomials, differentiate, and evaluate the world's worst integrals. A person needing this help can access it via ARPANET.

Public or semipublic networks can also provide access to highly specialized facilities. The *New York Times*, for example, has automated its archival stores of news stories. Hundreds of millions of abstracts of news items dating back to the early 1960's can be searched with a query language which permits unanticipated queries by linking multiple descriptive words. There are stock market systems which permit users to

search for stocks which meet specified criteria of company or financial characteristics—price, earnings, performance, and so on. Encyclopedia companies are working at computerizing encyclopedic bodies of information. A network enthusiast's vision is that vast numbers of information systems like these will become accessible, inexpensively, to network users.

Networks can provide access not only to information systems but also to large numbers of problem-solving facilities. Some corporations have a large number of intricate computerized tools for various types of engineering design or evaluation. These are made available to engineers throughout the corporation via a corporate network. There are cash flow models, market forecasting models, budgetary control models, and so on. One hospital has successfully operated a prediagnosis system for several years. A computer program interviews a patient to make a preliminary diagnosis of his ailment to decide whether he[†] should see a nurse, doctor, or specialist. Computerized medical diagnosis will not replace that of a doctor but it could provide worried persons with a quick indication of whether they need to see a doctor or to take some simpler remedial action. Similarly there are systems for job counseling, finding employment, booking theater seats, doing tax calculations, and so on.

In general there is such a rich diversity of computer applications and information banks that it is impossible to predict the way networks would be used if easy network access and a high level of sharing made them economically available.

Box 1.1 lists a few of the resources worth sharing via networks.

BOX 1.1 A few of the resources that are accessible via computer networks

- A computer for mathematical operations which solves simultaneous equations, factors polynomials, differentiates, and solves the most complex integrals (e.g., M.I.T.'s MACSYMA).
- Advanced word processing facilities which can handle multiple type fonts, professional print editing, and are on-line to printing equipment.
- Legal search systems which can search and display cases, patents, etc.
- Business directories; corporate annual reports; business analyses.
- Design tools for complex design processes, simulation, etc.

[†]The author has given much thought to the problem of avoiding words with a sexist connotation. It is possible to avoid words like "man," "manpower," "mankind," but to avoid the use of "he," "his" and "him" makes sentences clumsy. In this book, whenever these words appear, please assume that the meaning is "he or she," "his or her" and "him or her."