Arizona State University High Density Collection
S 2 02 09 10 03 16

ARIZONA STATE UNIVERSITY
A15048138244

Comcast, Ex. 1034
Comcast v. Entropic
IPR2024-00431

Computers in Communication

*A15048 138244

Computers in Communication

Gordon Brebner
Department of Computer Science
University of Edinburgh, Scotland

TK
5105.5
.B73X
1997
Science

THE McGRAW-HILL COMPANIES

London • New York • St Louis • San Francisco • Auckland
Bogotá • Caracas • Lisbon • Madrid • Mexico • Milan
Montreal • New Delhi • Panama • Paris • San Juan • São Paulo
Singapore • Sydney • Tokyo • Toronto

Published by
McGraw-Hill Publishing Company
Shoppenhangers Road, Maidenhead, Berkshire, SL6 2QL, England
Telephone 01628 502500
Fax 01628 770224

British Library Cataloguing in Publication Data

The CIP data of this title is available from the British Library, UK

Library of Congress Cataloging-in-Publication Data

The CIP data of this title is available from the Library of Congress, Washington DC, USA



Copyright © 1997 McGraw-Hill International (UK) Limited. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of McGraw-Hill International (UK) Limited.

Typeset by the author

Printed and bound in Great Britain at the University Press, Cambridge

Printed on permanent paper in compliance with ISO Standard 9706

To Rosemary

CONTENTS

Preface			**xiii**
1	**Introduction**		**1**
1.1	Background		1
1.2	Uses of computer communications		4
	1.2.1	Computer-oriented communication	4
	1.2.2	Telecommunications	8
	1.2.3	Radio and television broadcasting	10
	1.2.4	Summary of uses of computer communications	11
1.3	Physical links		12
	1.3.1	Physical media and their properties	14
	1.3.2	Physical communication services	20
1.4	How computers communicate		26
	1.4.1	Information, time and space	26
	1.4.2	Agreement and implementation	27
	1.4.3	Human influences	29
1.5	Chapter summary		31
1.6	Exercises and further reading		32
2	**Information**		**35**
2.1	Introduction		35
2.2	Agreement on information types		37
	2.2.1	Absolute and relative information types	38
	2.2.2	Examples of communicated information types	38
2.3	Implementation of required information types		45
	2.3.1	Transformation of information	45

	2.3.2 Representation of information	65
2.4	Chapter summary	74
2.5	Exercises and further reading	75
3	**Time**	**80**
3.1	Introduction	80
3.2	Agreement on time periods	83
	3.2.1 Absolute and relative time measurement	83
	3.2.2 Examples of communication time periods	85
3.3	Implementation of required time periods	92
	3.3.1 Segmentation and concatenation	92
	3.3.2 Flow control	97
	3.3.3 Acknowledgement and error handling	103
3.4	Chapter summary	107
3.5	Exercises and further reading	108
4	**Time Packages**	**110**
4.1	Introduction	110
4.2	Unsegmented time package	112
	4.2.1 Connectionless services	112
	4.2.2 Examples of unsegmented time packages	114
4.3	Simple handshake time package	116
	4.3.1 Examples of simple handshake time packages	118
4.4	Multi-stage handshake time package	120
	4.4.1 Examples of multi-stage handshake services	121
4.5	Connection-oriented time package	125
	4.5.1 Connection-oriented services	126
	4.5.2 Examples of connection-oriented time packages	129
4.6	Chapter summary	134
4.7	Exercises and further reading	135
5	**Space**	**138**
5.1	Introduction	138
5.2	Agreement on spaces	140
	5.2.1 Identifier schemes	140
	5.2.2 Absolute and relative spaces	147
	5.2.3 Examples of communication spaces	149
5.3	Implementation of required spaces	150
	5.3.1 Filtering and switching	150
	5.3.2 Splitting and multiplexing	159
5.4	Chapter summary	165

5.5	Exercises and further reading		166
6	**Message Broadcasting Networks**		**168**
6.1	Introduction		168
	6.1.1	Information basics	169
	6.1.2	Time basics	171
	6.1.3	Space basics	172
6.2	Multipeer channel implementation		173
	6.2.1	Common medium	174
	6.2.2	Chain or ring	174
	6.2.3	Star or tree	176
6.3	Communication multiplexing		178
	6.3.1	Isolated	179
	6.3.2	Permission-based	181
	6.3.3	Reservation-based	182
	6.3.4	Physical division	183
6.4	Examples of guided technology network implementations		184
	6.4.1	Ethernet	185
	6.4.2	Token ring	190
	6.4.3	FDDI and FDDI-II	194
	6.4.4	Token bus	200
	6.4.5	DQDB	204
	6.4.6	100 BASE VG-AnyLAN	208
	6.4.7	Discussion of examples	211
6.5	Unguided technology network implementations		212
6.6	Chapter summary		218
6.7	Exercises and further reading		218
7	**Message Switching Networks**		**222**
7.1	Introduction		222
	7.1.1	Information basics	224
	7.1.2	Time basics	225
	7.1.3	Space basics	227
7.2	Channel implementation		229
7.3	Communication multiplexing		235
	7.3.1	Isolated	237
	7.3.2	Permission-based	238
	7.3.3	Reservation-based	238
7.4	Examples of physical switching network implementations		239
	7.4.1	Switched ethernet	240
	7.4.2	Telephone system circuit switching	240
	7.4.3	X.25 packet switching	241
	7.4.4	Frame relay	244

	7.4.5	ATM cell relay	248
7.5	Examples of public switching network services		253
	7.5.1	PSTN	254
	7.5.2	CSPDN	254
	7.5.3	PSPDN	254
	7.5.4	N-ISDN	255
	7.5.5	SMDS	255
	7.5.6	B-ISDN	256
7.6	Example of a high-level switching network implementation		258
7.7	Chapter summary		263
7.8	Exercises and further reading		263
8	**Inter-networks**		**267**
8.1	Introduction		267
	8.1.1	Information basics	269
	8.1.2	Time basics	269
	8.1.3	Space basics	271
8.2	Inter-network switching		274
8.3	Examples of inter-network implementations		276
	8.3.1	IEEE 802 LAN bridging	276
	8.3.2	The Internet	280
	8.3.3	ISO inter-networking	302
8.4	Chapter summary		306
8.5	Exercises and further reading		307
9	**Case Study 1: Accessing the World Wide Web**		**311**
9.1	Introduction		311
9.2	The problem: information, time and space issues		314
9.3	Reliable end to end communication using TCP/IP		324
9.4	Home computer and the Internet		334
9.5	Within the Internet		340
9.6	WWW server and the Internet		347
9.7	Overall communications		349
9.8	Chapter summary		355
9.9	Exercises and further reading		355
10	**Case Study 2: Controlling a manufacturing device**		**359**
10.1	Introduction		360
10.2	The problem: information, time and space issues		362
10.3	FullMAP implementation		367
10.4	MiniMAP simplified implementation		384
10.5	Chapter summary		385

10.6 Exercises and further reading 386

11 Case Study 3: Making a video telephone call 388

11.1 Introduction 388
11.2 The problem: information, time and space issues 391
11.3 Management of an ATM connection 397
11.4 Using an ATM connection 402
11.5 Overall B-ISDN implementation 405
11.6 Chapter summary 408
11.7 Exercises and further reading 408

12 Standardization 411

12.1 Introduction 412
12.2 Standards bodies 414
12.2.1 International Organization for Standardization (ISO) 415
12.2.2 International Telecommunication Union 417
12.2.3 Internet Society 418
12.2.4 Professional and manufacturer bodies 420
12.2.5 Summary of standardization bodies 423
12.3 ISO standards 424
12.3.1 Structure 424
12.3.2 Services and protocols 430
12.3.3 Management 433
12.3.4 Other standards 434
12.4 Internet standards 435
12.4.1 Structure 435
12.4.2 Protocols 438
12.5 ITU-T standards 441
12.6 IEEE 802 standards 444
12.7 Proprietary standards 445
12.7.1 IBM System Network Architecture 446
12.7.2 Digital Network Architecture 447
12.7.3 AppleTalk 448
12.7.4 Xerox Network Systems and descendents 449
12.8 Chapter summary 449
12.9 Exercises and further reading 450

Index 454

CHAPTER
ONE

INTRODUCTION

The main topics in this introductory chapter are:

- benefits of communication
- history of human and computer communication mechanisms
- present-day uses for computer communications
- physical underpinning of computer communications
- the main principles of computer communications

1.1 BACKGROUND

Communication is beneficial for the human race. By communicating with one another, information can be shared — past experience, current affairs, predictions of the future — from here, there and everywhere. Also, resources and expertise can be shared, by communicating with the right people. Only hermits are noted for their ability to live satisfactorily in the absence of any communication with other people.

Similar observations apply to computers too. Individual computers are capable of gathering, processing, storing and distributing information, under the direction of humans. They are not only found in distinctive boxes with keyboards, mice and screens attached, or in large cabinets bristling with flashing lights and whirling tape drives, as seen in ageing science fiction films. Their basic information-handling capabilities can be harnessed for controlling other machinery, and so they are also hidden inside things like wristwatches, microwave ovens, central heating systems, factory production line equipment and nuclear power plant safety systems.

There are three main areas where benefits can be expected if one computer is able to communicate with others:

- it can get information that is stored by other computers;
- it can get other computers to do specialized work; and
- it can communicate with humans that use other computers.

The benefits need not only be in one direction — this computer can also export its own information, its specialized abilities and access to its human users. The outcome is a beneficial sharing of resources.

At present, it is usually the case that computers inside personal items like wristwatches, or inside domestic equipment like microwave ovens, do not communicate with others. However, this situation is on the point of changing, given that communication between the more recognizable types of computers has proved to be very useful, and that appropriate communication technologies are becoming available. It might also seem that a conventional home computer, or a single computer tucked into a musty office, is island-like, cut off from the world community of computers. However, this is a delusion, since such computers usually have an obliging communication mechanism: human beings transferring the latest fruits of the computing trade on floppy disk or compact disk.

This book is concerned with computer communications where there are no human middlemen, so that computers can converse directly with one another. Although human participation in the role of intermediary is being eliminated, it should not be forgotten that computers only communicate because humans have instructed them to and, moreover, that this is possible because humans have instructed the computers how to communicate.

The techniques covered arise from communications between conventional types of computer. However, they are equally applicable to the world of the future, where there will be things like intelligent houses with communication not just between domestic appliances but also with the fabric of the building itself.

The act of communication is not always easy for humans. For example, it is not feasible for every person in the world to communicate effectively with any other person whenever desired. Differences in culture, availability and physical location cause problems. Things get even tougher if one broadens communication to include other species of animal or plant, never mind any alien life forms that might visit or be visited.

Similar problems affect communication between computers. Because of this, the solutions used are rather similar to those developed by humans over the millenia, and so are usually familar to the neophyte from normal human experience. This is good, since it means that explanations of techniques used can be well motivated by appropriate examples from human communications.

Aside from reconciling differences between communicating parties, a further problem is how communication is physically achieved. For humans, there is a mixture of movement and media. People may move to make communication easier, or even possible at all. For example, one may climb a mountain to consult a guru, travel to work or attend a concert. Once the communicating parties are in suitable positions, communication might take place using sound through the air, visible sign language, telephones, television or sending items through the postal system. For computers, there is usually no movement involved in communication, robotics still being in its relative infancy. Because of this, computers communicate using media that physically reach the computers, and that are capable of transmitting computer conversations. Of course, this does not rule out the possibility of humans moving the computers to places where they can be reached by media.

The main themes of the book are concerned with these communication problems. The next two sections of this introductory chapter provide the enclosing framework for these themes. First, Section 1.2 surveys the main practical uses made of communications between computers. Such uses provide the rationale for indulging in computer communications in the first place, and indicate what are the main requirements of computer communications. After this, Section 1.3 surveys the physical media that can be used for the transmission of information between computers. This indicates the physical limits that apply to the process of communication and which, ultimately, affect what can be offered to the users of computer communication facilities.

Following Sections 1.2 and 1.3, Section 1.4 introduces the main themes of the book. The features of communications are classified under three main headings:

- **information:** the type of information that is communicated;
- **time:** when, and how quickly, a communication takes place; and
- **space:** which computers, and inter-connecting channels between computers, are involved in a communication.

There are two main problems associated with communication: achieving agreement between computers on the nature of communications; and implementing the required communications using available physical communication media. The processes of achieving agreement and implementation require a lot of human decision making prior to communications being possible. This entails human communications of a rather specialist kind. The section ends by introducing the ways in which these technical and political human communications have a major influence on what is possible and not possible in the computer world.

1.2 USES OF COMPUTER COMMUNICATIONS

Present-day uses for computer communications arise from a convergence between two different worlds. The first is a computer-centred world, where computers existed, and then it became convenient to inter-connect them. The other is a human-oriented world, where communication facilities existed, and then it became convenient to computerize these facilities. In the latter world, there was also convergence between telecommunications facilities largely used for inter-personal communication, such as the telephone, and broadcasting facilities largely used for entertainment, such as television.

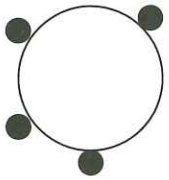
One term often used to describe the fruits of this general convergence is the **global village.** It is interesting to note that, when this phrase was coined by Marshall McLuhan in 1964, it was based on an extrapolation of existing human-oriented facilities, and did not envisage the future involvement of computer communications. However, it is still equally apt to embrace computer-based facilities. The technology shift is captured in another, more modern, term: the **information superhighway**. This refers to the collection of communication technology and information technology that will be used to underpin the global village of the future. Just as roads underpin the movement of people and goods, so the information superhighway will underpin the movement of information.

To understand the current uses made of computer communications, and to point ahead to future developments, it is useful to conduct a brief historical survey of how different communication systems emerged and then converged. This follows in the next three sub-sections, which cover developments on the computer front, on the telecommunications front, and on the broadcasting front, respectively. A central concern is the demands that are made by these differing types of systems if they are to be realized using computer communication facilities.

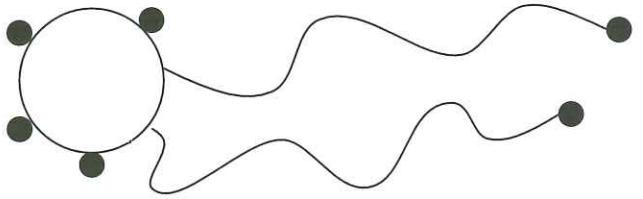
To quantify these demands, time can be measured as usual, in seconds, or in fractions or multiples of seconds. Information can be measured using **bits** (short for 'binary digits'). The quantity of information available in one bit is that required to distinguish between two possible values. For example, one bit of information is enough to distinguish between 'on' and 'off', 'black' and 'white' or 'yes' and 'no'. A more precise definition of both information, and the bit as a measure of information, is given in Chapter 2. The bit is like any other unit, so it is convenient to talk of the kilobit (kbit), which is 1000 bits, and the megabit (Mbit), which is 1 000 000 bits. Note that these units are powers of ten, rather than the powers of two (2^{10} and 2^{20}) sometimes used in computing circles.

1.2.1 Computer-oriented communication

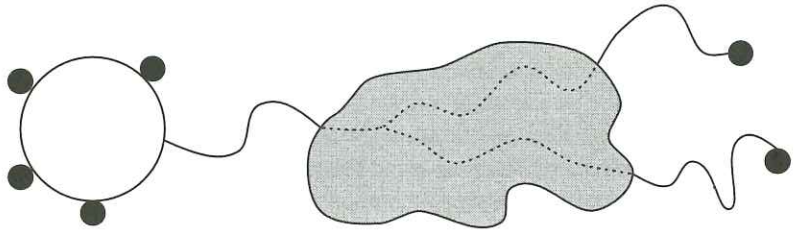
In the earliest days of computers, the roomful of boxes containing the component parts of the computer was the centre of attention. Information was supplied to a computer, and retrieved from a computer, by people who were physically present in the computer room. This is illustrated schematically in Figure 1.1(a). In the



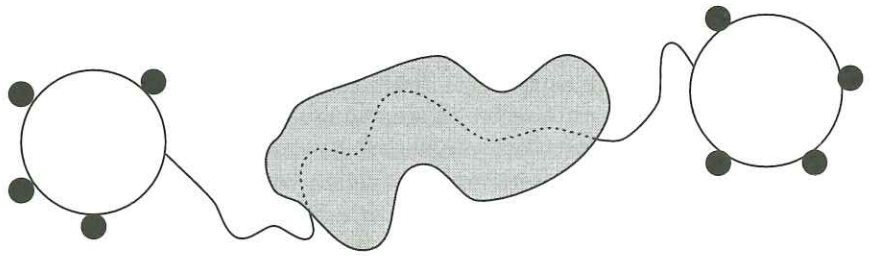
(a) computer with only directly-connected input/output devices



(b) computer with remotely-connected input/output terminals



(c) computer with terminals connected via communication service



(d) computers connected via communication service

Figure 1.1 Evolution of computer-oriented communications

unusual event that there was any communication between computers at all, this was done by carrying paper tapes or magnetic tapes produced by one computer along to be read by another computer. Thus, people were the servants of computers as much as computers were the servants of people.

The first developments in communications were designed to make life easier for people, by eliminating the need to be physically adjacent to the computer when interacting with it. One of these was the introduction of a remote terminal — a typewriter-like device with a keyboard for input and a teleprinter for output — connected by a cable to the computer. This is illustrated in Figure 1.1(b). In essence, terminals were just extra peripheral devices for the computer, the only difference being that they were a longer distance away. Because of this, and the fact that terminals were electro-mechanical devices, they operated at very slow speeds in computer terms: a few characters input per second and around 30 characters output per second.

This basic capability evolved into a model of interaction with computers that survives to the present day. There are two major improvements. One is in terminal technology, with video screens replacing printing, giving a consequent large increase in output speed, and also with terminals containing computers as their controllers (or, indeed, computers just emulating the behaviour of terminals). Beyond terminals, full WIMP (Window, Icon, Menu and Pointer) interfaces can also be used, if the computer is sophisticated enough to interact via such an interface.

The other improvement is the one of interest to computer communications. This is that the direct physical cable between terminal and computer can be replaced by any communication channel capable of passing information between a terminal and a computer. Useable communication channels include links through the conventional telephone system or through specialized computer networks. This is illustrated in Figure 1.1(c). The **Internet** has become the best-known specialized mechanism for providing links between computers located in all continents of the world. In 1997, it was estimated that around 20 000 000 computers could make use of the Internet. The overall effect presented to a human user is still of a direct physical link between terminal and computer, but this is an illusion.

The increased distancing of terminal from computer presents two agreement problems that must be solved to enable communications. The first is that a wide variety of terminals (or computers masquarading as terminals) may be used, and each terminal's behaviour must be reconciled with that expected by the computer. This issue, and its resolution, is discussed more fully on page 44 in Chapter 2.

The second problem concerns the human using the terminal. This person issues commands to the computer, or inputs information to the computer, through the keyboard. Given that the terminal is behaving as though it is a real peripheral device of the computer, its user must know exactly how to interact with that particular type of computer. This is acceptable if only one computer is ever used via the terminal, but the whole point is that communications advances now allow all sorts of computers, all over the world, to be used. No single person can be

expected to know how to interact with all of the different computers that are accessible.

A solution to this problem is for a particular user just to interact with his or her own familiar computer, and then that computer interacts with any other computers of interest on the user's behalf. This introduces some automation of the task faced by the human terminal user. In early versions of such facilities, the user's computer just pretended to be a human user of the distant computers. That is, it transmitted information that appeared to come from a terminal keyboard (albeit with a rather faster typist than normal) and then received back information by pretending to be a terminal printer or screen. This sort of interaction allowed files to be sent to, or fetched from, distant computers, electronic mail to be sent or received, or processing jobs to be given to distant computers and their results to be retrieved. These operations were directed by the human user.

In time, such procedures became refined, to eliminate the unnecessary humanization of the dialogue between computers. That is, new agreements to specify appropriate direct interactions between computers replaced the more verbose means used by humans. The resulting situation is illustrated in Figure 1.1(d). The new agreements covered things like transferring files, sending electronic mail and submitting jobs. These matters are discussed in more detail in Section 2.2.2. A further effect of such advances is that users gradually got more insulated from the vagaries of particular types of computer. For example, the same user interface could be used for something like electronic mail, regardless of what type of computers were involved in the mail transmission. Further extensions of the insulation process led to **distributed systems,** where the existence of a collection of computers is completely hidden from a user. Thus, it appears that a service is being provided by one single uniform computer system.

In distributed systems, it is fairly common for the computers involved to take on **server** and **client** roles. Servers have special capabilities, for example, storing particular information, performing particular processing of information, or having particular input or output devices attached. Clients can make use of servers by issuing requests, and by getting responses back. Overall, this is similar to the ways that humans directly interact with computers. However, in a distributed system, the computer client-server relationships are hidden from human users. The World Wide Web (WWW), considered in detail in Chapter 9, is an example familiar to many, in which the illusion of a world of information is presented to a user. In fact, the user's computer acts as a client, and computers throughout the world that store particular WWW pages of information act as servers for this client.

In summary, the evolution of computer-oriented communications has gone from a situation where a terminal was connected to a computer via a physical cable to a situation where a computer can be connected to numerous other computers via indirect channels supplied by complex communication systems. The types of information shared are more complex and the quantities are potentially huge. The speeds are related to computer rates rather than human rates. Early terminal links required the communication of only 300 bits per second. Two high speed com-

puters might be able to cope with the communication of 1 gigabit (1 000 000 000 bits) per second nowadays. When pushed to one extreme, a collection of communicating computers can be placed into one box, with high speed electronic interconnections, to form a **supercomputer**: a new, high speed computer consisting of many individual computing elements. This is the point at which computer communications meets the subject area of parallel computation.

1.2.2 Telecommunications

Human communication via a medium is long-established. In the mid-nineteenth century, the **telegraph** became an economical method for the electrical transmission of information represented in Morse code, or similar, over long distances. In style, Morse code is not far removed from the bit-focused approach to information followed by computers. The difference is that information is expressed in a three-valued form (dot, dash and pause) rather than a two-valued form. The major problem with the telegraph was the need for an expensive physical cable between the communicating parties. The development of telegraph systems often occurred in parallel with railways, with telegraph cables being located alongside the railway tracks. The medium was used fairly wastefully, because the speed of transmission was limited to the speed at which people could press Morse keys to send information and the speed at which people could listen to the transmission to decode the information. Thus, the restriction came from human frailty, rather than any fundamental physical limitation of the cabling.

Around the turn of the twentieth century, the **telephone** became available. Like the telegraph, a large investment in cabling was necessary. However, unlike the telegraph, the telephone was an analogue device, in that human speech was directly converted to electrical waveforms for transmission. Experience of using the telephone until recent times confirms the main problem of analogue transmission: dubious effects on the electrical signal are reflected as audible crackles, whines, etc. on reception. Nowadays, most telephone systems employ digital transmission systems, with speech being represented using a series of bits. This is because it is rather easier to reconstruct a series of bits from a damaged electrical signal than directly encoded speech. A complex collection of automated national and international telephone exchanges provides a near-worldwide telephone communication system, going a fair way towards a goal of allowing every person in the world to speak to any other. Transmission is not exclusively via cabling, with radio transmission becoming increasingly used by mobile telephones.

The telephone system can be used as a vehicle for computer-oriented communication, for example, as a way of connecting a terminal to a computer. This does not involve computers speaking to one another in a human way over the telephone. Instead, electrical waveforms of a similar style to those used to encode speech are transmitted, but they in fact encode series of bits of computer-style information.

The modern telephone system has evolved to be like a distributed computer system. The digital exchanges are just special-purpose computers, and the links between them are similar to links between computers. To an extent, a fully digital service is offered to end users of the telephone system, through the **Integrated Services Digital Network** (ISDN) service, and other offerings. However, in the main, the typical subscriber is still supplied with a traditional analogue-style service. This is implemented using the underlying digital facilities. Thus, if such a subscriber uses the service for computer communication, the information is needlessly passed through an analogue form before being transmitted digitally.

One extension to the conventional telephone is to make it into a videophone, where a picture of the caller is transmitted as well as the speech of the caller. It is feasible to supply a videophone service over a normal telephone line, but the picture quality is fairly crude. A further catch is that few people possess videophone equipment. However, the main problem is that much more information must be transmitted to encode pictures and so, to achieve better quality, a better communication facility is needed. For speech alone, it is enough to transmit about 10 kbits of information per second. However, for video, around 60 kbits of information per second is a lowest limit, with 1500 kbits per second needed for decent quality.

Other telecommunication services have been developed in addition to the telephone, geared to transmitting textual-style information rather than direct human communications. One of these is **fax** (facsimile), which can be used to transmit an image of a text page from one telephone user to another. This involves a digital representation of the page in terms of bits of information being transmitted between two fax machines, encoded as a telephone-style electrical signal. The fax machines are really just special-purpose computers, which conduct a conversation over the analogue telephone system. Many computers can now emulate the behaviour of fax machines, so that they can transmit page images to other computers or fax machines. The number of bits of information needed to represent a faxed page varies according to the pattern of black on white on the page. However, on average around 200 kbits are sufficient for one page. More details of fax representation are given on page 53 in Chapter 2.

A further family of information transmission services are those based on **telex**. These are fairly close relatives of computer-oriented communications that involve textual information being sent between computers and terminals or between computers. Telex itself is a long-established telegraph service that allows textual messages to be transmitted between subscribers. There are about 1.2 million subscribers, mostly businesses, around the world. The messages can be composed of upper case letters, numerical digits and some punctuation symbols. They are transmitted in a digital form, but the transmission rate is only 50 bits of information per second, which equates to only 10 characters per second. A more modern offering is **teletex,** which allows a much larger character set, including graphical symbols and word processor style facilities for composing pages. It has a much

faster transmission rate of 2400 bits per second, but this is still slow in modern computer communications terms.

Both telex and teletex are simple mechanisms for transmitting digital information between two specialized machines. **Videotex** is a variation that is akin to the client-server model for distributed computer systems. It is used over the normal telephone system. A user of videotex has a special video terminal attached to the telephone line, and this acts as the client. It is possible to telephone computer databases worldwide, and then conduct searches for information. The databases are the servers. One place where videotex is often seen is in travel agents, where video terminals are used to interrogate holiday and travel booking systems over the telephone. Videotex is an example of a **value-added service** operated by telecommunications providers. In addition to a raw communication facility, extra services are provided on top. This type of service illustrates a convergence of telecommunications with computer-oriented communications.

1.2.3 Radio and television broadcasting

Around the turn of the twentieth century, at the same time as the telephone was becoming available, a free worldwide communication medium was being harnessed: the electro-magnetic spectrum. This allowed the use of physical phenomena such as radio waves, rather than expensive cabling, for transmission. The first application was a wireless version of the telegraph. Although it seems strange with hindsight, the obvious use of this medium — broadcasting — was not realized until after the First World War. To an extent, this was due to an attitude that communication technology was only relevant for closed commercial and military applications, rather than the general public. The average person was still expected to travel in order to communicate directly. A further reason was that the cost of transmitting and receiving equipment was high, although there were no wiring charges. Since these early days, of course, there has been a massive growth in public radio and television broadcasting services.

Until the last years of the twentieth century, the sound and picture information of radio and television has been transmitted as an analogue signal. These signals are received using aerials, and then decoded. Digital transmission, with sound and pictures encoded using series of bits was only in its infancy in 1997. However, this is undoubtably the transmission style of the future, marking a convergence of radio and television with computer-oriented communications.

Compared with the transmission demands of the telephone and videophone, those of radio and television are much higher. In the case of radio, this is because the whole range of audible sounds must be transmittable, rather than just human speech. In addition, extras like stereo sound are also desirable. The information rate needed for high fidelity audio transmission is about 100 kbits per second — ten times the rate for telephony. In the case of television, pictures are larger, contain more detail and involve more motion over time. The information rate needed for television-quality video transmission is about 15 Mbits per second — again, ten

times the rate for decent-quality video telephony. Note that these rates are for digital transmission, where extensive computer processing of the information can compress it to fit these rates. Without such processing, the required rate for raw video transmission is very much higher.

Just as videotex allows the retrieval of digital information using the telephone system, so **teletext** is a similar facility associated with television. Note the subtle naming difference between the television world's *teletext* and the telecommunication world's *teletex*. Teletext has far less scope for interaction, compared with videotex. All of the available information is continuously broadcast in sequence, and then the television receiver has to filter out the information of no interest. This is in contrast to videotex, where a user explicitly initiates searches for information of interest. However, like videotex, teletext illustrates a convergence with computer-oriented communications.

The rise of cable television illustrates a convergence between television and telecommunications. The same physical technology can be used to deliver both telephone and television services to an end user. It can also be used to deliver computer-oriented communication services, thus completing a convergence of communication facilities with rather different histories. Although analogue transmission is still used for cable television, in time, with the advance of digital television, the raw physical cable will become the bearer of a digital information bit carrying service, which is then used to support different value-added communication services. This will represent each user's slip road to the information superhighway.

1.2.4 Summary of uses of computer communications

From the point of view of a user, the interface to the converged world of computer communications, telecommunications, radio and television should allow easy access to information, communication facilities and entertainment. These will be underpinned by the techniques associated with computer communications today. Indeed, already, devices such as telephones and television sets are really just special-purpose computers from a technical communications point of view.

The blending of computers and everyday devices will continue, to a point where, from a user point of view, explictly visible computers are not likely to be present. A user will only see 'intelligent' devices; these will be internally underpinned by computers. The devices might be examples of **wearable computing,** embedded in clothes or spectacles (or even in the body), or of **ubiquitous computing,** embedded in numerous objects all around the user. One particular new type of device will be a **knowledge robot** — a device that can speculatively gather information that is likely to be of interest to its user, an improvement on conventional facilities that only gather specific information when ordered to.

One key component of a converged information interface is the use of **multimedia,** where a blend of sound, still pictures, video sequences and computer-encoded text can be used to present information. Another key component is in-

Table 1.1 Information sizes (in bits)

Type of information	Number of bits
Teletext page	7000
One second of speech	10 000
Computer-stored page of text	12 000
Character-based video screenful	15 000
One second of hi-fi audio	100 000
Typical faxed page	200 000
One second of quality videophone	1 500 000
Colour photograph	4 000 000
One second of entertainment video	15 000 000

teraction, where the user is not just a passive recipient as for radio and television. As a first step, **video on demand** is seen as an extremely important facility of the future. This allows users to select films for immediate viewing — akin to going to a local video library — rather than having to fit in with the whims of broadcasters. Later, it is possible to imagine that interactive films will be produced, allowing viewers to participate. Such facilities are rather akin to the interactive multi-player computer games, already available through computer communication facilities.

The demands placed on computer communication facilities will be great. From an information point of view, a rich variety of types of information must be dealt with, with some of these types being very large in terms of the number of bits of information required. Table 1.1 summarizes some approximate minimum information sizes, many of them quoted earlier in this section. From a time point of view, ideally communications must happen quickly enough that there is no unacceptable delay to a user (for example, within half a second for speech over a telephone) and then continue at a fast enough rate to ensure high quality. From a space point of view, ideally any computer-based device in the world should be able to communicate with any other, and adequate channels must be available to support these communications. Most of this is technically feasible, but only at extortionate cost. Therefore, the practical decision is on what facilities can be provided at a reasonable cost.

1.3 PHYSICAL LINKS

Just as human communication devices like the telephone and television are underpinned by physical media, including electrical cables and broadcasts in the electro-magnetic spectrum, so are computer communications. This book is not concerned with the physical details of how such media are put to work, but only

with their observable behaviour: essentially, how good they are for transferring information. This encompasses issues like transfer rate, reliability and, of course, cost.

For any particular medium, its communication abilities can be categorized under the headings of information, time and space. The assumption for information is that any medium used is capable of transmitting bits, i.e., two-valued quantities, either individually or in a sequence. This service must be offered as a basis for computer communications, and nothing more complex is necessary. Therefore, the bit forms the atomic unit of information for communications. The details of how bits are transmitted, as signals, waveforms or whatever, are not relevant here. The service need not necessarily be completely reliable, in that some transmitted bits may be lost completely, changed in value or duplicated. However, it must be sufficiently useable that it can form a basis for reliable bit transmission.

From the point of view of time, there are two main measures of interest. The first is the **latency** of the medium, that is, the time taken between a bit being transmitted and being received. The fundamental physical limitation is the speed of light, which is approximately 3×10^8 metres per second. For example, this means that information cannot travel a distance of 1 km in less than 3.3 microseconds. Even if the raw medium has a latency as low as the limit imposed by the speed of light, the actual latency is likely to be increased by delays in the interfaces used to convert bits of information into signals for a medium.

The other important time measure is the **rate** of the medium, that is, the rate at which individual bits in a sequence can be transmitted. The subject of *discrete* information theory is covered in some detail in Section 2.3.1. Fundamental results from *continuous* information theory, a topic not covered in this book, can be used to derive limits on the maximum transmission rate possible over particular media, given things like the frequency range of waveforms used and inherent error rates.

Finally, from the point of view of space, physical media must act as channels for information being sent between computers. A fundamental property is the area or distance that can be spanned by a particular medium. Most media introduce a deterioration in the quality of tranmission as distance increases, to the point where practical bit transmission is no longer possible. Very often, this effect is related to the transmission rate — roughly, the faster the rate, the shorter the distance that can be spanned by most media.

Aside from area and distance, two distinctions can be drawn between the inherent properties of media. One distinction is between:

- **unicast**; and
- **broadcast**

media. The former allow one computer to send information to one other computer. This is the style of transmission familiar from the telephone system. The latter allow one computer to send information to a whole collection of other computers,

some of which may choose to receive the information and others may not. This style is familiar from broadcast radio and television.

This first distinction should not be confused with a different distinction. It is between:

- **guided**; and
- **unguided**

media. The former involve some type of cabling that is used to physically guide transmitted bits along their channel. The latter involve transmission 'through the air' with the effects spreading over a wider area, which includes the destination(s) of the channel. Examples of the less obvious cases are radio telephones for unicast, unguided transmission, and cable television for broadcast, guided transmission. The term **wireless** is also used to refer to unguided media, reflecting the fact that no wires are involved. This category includes radio, which traditionally was synonymous with the word 'wireless'.

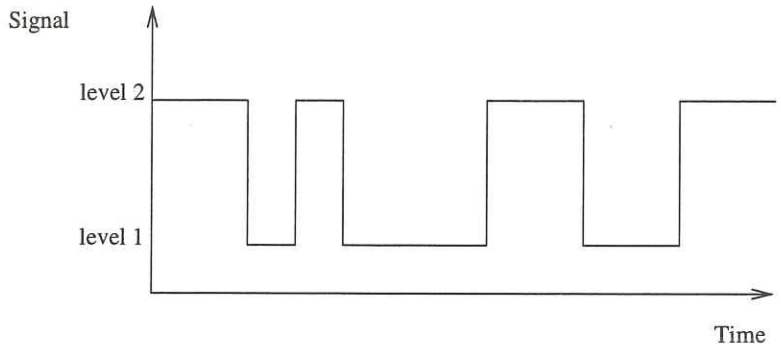
Section 1.3.1 below contains a brief survey of the most common types of raw physical transmission media, focusing on their bit transmission abilities. However, this is not the full story for real-life computer communications. In practice, there are legal and technical restrictions on the ways that media can be used. For example, there are no general rights that allow guided media to be laid across private property. In most countries, only national telecommunications organizations have the right to lay long distance cabling, or cabling along public roads. As another example, there are wide restrictions on transmission in the electro-magnetic spectrum. Different frequencies are allocated to radio and television companies, the military and other specialized uses. As a result, for communications outside small privately owned areas, services offered by approved operators must form the basis for computer communications, rather than direct use of raw media. Such services are discussed in Section 1.3.2 below.

1.3.1 Physical media and their properties

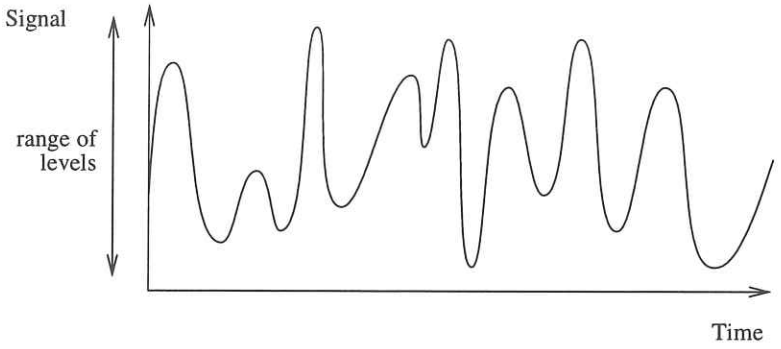
The basic characteristic of practical physical media is that, when in use, they involve the continuous transmission of a signal from one point to other points. This signal is used to encode sequences of one or more bits that form the information to be communicated. A distinction can be drawn between:

- **digital**; and
- **analogue**

signals. A digital signal has a small number of different possible values, for example, two different values for a binary digital signal. Over time, the value of the signal changes in discrete steps. An analogue signal has values drawn from a continuous range and, over time, the value of the signal varies over this range.



(a) digital signal with two different levels



(b) analogue signal over range of levels

Figure 1.2 Examples of transmission signals

Figure 1.2 shows examples of these two types of signal viewed over a period of time.

While 'digital' and 'analogue' are familiar terms in the computer world, two alternative terms are often used loosely in the computer communications world: **baseband** and **broadband.** The looseness of usage comes from the fact that these terms have precise meanings in the telecommunications world, referring to the way in which frequency bands are used. However, this is not an immediate concern here, although a little more discussion of this point is contained on page 65 in Chapter 2. One further caveat is that the term 'broadband' is now often just used as an adjective meaning 'high speed', in contrast to another technical term, 'narrowband', which is now often used just to mean 'low speed'.

The most-used media for computer communications are examined below in increasing order of *potential* (rather than typical) transmission rate. This has the side-effect of presenting unguided media first, then guided media. The main benefit of unguided media over guided media is one of mobility. The positioning of computers is not constrained by the need to be adjacent to inter-connecting cabling. If the equipment needed for transmitting and receiving over an unguided medium is small and light enough, then mobility translates into true portability.

There are few suprises below, since all of the media should be familiar to the reader from everyday human communication systems. The presentation omits two more special-case suggestions that have arisen from the Internet community. One is Waitzman's 'avian carrier' system, proposed on 1 April 1990, which involves bits being transferred on a roll of paper attached to one leg of an avian carrier. The other is Eriksson's 'accoustical transmission medium' system, proposed on 1 April 1996, which involves bits being transmitted using Morse code on sound waves emitted by the beeper of a personal computer.

Radio

Radio waves are an unguided, broadcast, analogue medium. The two different bit values are represented by differing transmitted radio waveforms. In the past, radio has been a Cinderella of the computer communications trade, to a large extent restricted to recreational activities by amateur radio enthusiasts. One notable exception was the pioneering use of radio to link computers distributed over various Hawaiian islands, which has impacted on several general aspects of computer communications. However, radio is now 'going to the ball', with the advent of portable computers and mobile communications. Radio offers a free medium, and flexibility of attachment for computers.

Radio transmission is used over fairly modest distances, for example, within a town or city, or over a smallish rural area. Over such distances, transmission rates are slow — as little as 50 kbits per second. However, over much shorter distances like 100 metres, rates of 10 Mbits per second are feasible. This allows a little flexibility, combined with respectable (i.e., computer-like) speeds. Each transmitted bit arrives at the speed of light over a direct path, which means the shortest latency possible. The problem with radio over any distance is interference from geographical features like mountains and lakes, not to mention amateur radio enthusiasts, and this can lead to rather high probabilities of bits arriving at a receiver damaged. Whole sequences of bits might be lost in unfavourable atmospheric conditions. This is far less of a problem over short distances indoors.

Infra-red

Infra-red transmission involves higher frequency waveforms than radio, and is useful for communication over short distances. In the home, it is a familiar medium for things like remote television controls. A main difference from radio is that infra-

red waves do not pass through solid objects, and so its use is confined to within rooms. Infra-red is not normally usable outdoors, since the sun's rays interfere. One advantage of the in-room use is that infra-red transmissions are not regulated, unlike radio, which involves the central assignment of radio frequencies to users.

Infra-red is an unguided medium that is naturally unicast, but can be used in either unicast or broadcast modes. For unicasts, an infra-red beam is just directly pointed at a receiver. For broadcasts, a diffuser is used to spread out the infra-red transmission so that it can be received throughout a room.

Digital transmission is used for slower infra-red transmissions, up to about 2 Mbits per second. This essentially involves switching the infra-red beam off and on to represent the two different bit values. For faster transmission rates, around 10 Mbits per second, analogue transmission has to be used, as for radio. Different waveforms are used to represent the different bit values. The main reliability issue involved in infra-red transmission is ensuring that the beam is not interrupted by a person or object standing in the way. This has the effect of causing the complete loss of a bit stream for the duration of the obstruction.

Microwave

Microwaves reside between radio and infra-red in the electro-magnetic spectrum. While infra-red is an attractive alternative to radio over very short distances, microwaves are usually used over long distances. However, some use is made of microwaves in the **Industrial Scientific and Medical** (ISM) waveband over short distances. One waveband is allocated to ISM worldwide, and transmissions can be made locally with no further regulation.

Microwaves travel in a straight line, and are narrowly focused, and so are an unguided unicast medium. Analogue transmission is used. Operating directly between two points, microwaves can be used for terrestial communications over distances of 50 kilometres or more. Until the advent of fibre-optic cables, discussed below, microwave transmission formed the backbone of the long-distance telephone system for several decades.

Microwaves are also used for communication between the earth and **geostationary satellites.** By bouncing signals off satellites, global distances can be covered. Direct transmission is not possible because of the curvature of the earth. When using such a satellite, most of the microwave's voyage is outside the atmosphere, and so interference is reduced and reliability is improved. Within the atmosphere, microwaves are prone to disruption, for example, by passing through rainfall. Geostationary satellites are positioned about 36 000 km above the equator, and revolve at the same rate as the earth, which means that they appear stationary from the ground. Positioning of such satellites is closely regulated.

In the 1990s, the first proposals for collections of **low-orbit satellites** were made. These orbit at an altitude of only 750 km, and so would not be geostationary. Instead, a sufficiently large number of satellites are in orbit that at least one is over any particular area at all times. The Iridium low-orbit satellite project, proposed

by Motorola, and targeted at mobile telephones, involves 66 satellites (and six spares). The first three satellites were launched at the beginning of 1997, and the system was scheduled to come into service in 1998. The Teledesic project, proposed by Microsoft and McCaw Cellular Communications, and targeted at high speed computer communications, involves 840 satellites (and 84 spares), and is scheduled to supply a service by 2001.

Transmission rates of around 50 Mbits per second are possible with microwaves over long distances. However, when a geostationary satellite is involved as an intermediary, the latency in receiving each transmitted bit is around 250 milliseconds because of the round-trip distance of 72 000 km to and from the satellite. This latency is distinctly high in computer terms. It is much reduced if a low-orbit satellite is used instead, since the round-trip distance is reduced by a factor of around 50.

Copper cable

In contrast to the above unguided media, copper cable is the traditional, and most widespread, guided transmission medium. This transmits electrical signals that represent bits. Either digital or analogue transmission can be used, with the terms 'baseband' and 'broadband' normally being used for this medium. Digital transmission is used over relatively short distances, say up to 1 km, at which point interference effects impact on reliability. Analogue transmission can be used over longer distances, up to 100 km, since anologue signalling is more robust to interference.

Copper cables are packaged in two main forms, both familiar in the typical home. The first is **twisted pair,** where two wires are twisted together in a helix like a DNA molecule — for example, this is used to connect telephones to the exchange. The other is **coaxial cable,** where a woven cylindrical wire surrounds a stiff core wire — for example, this is used to connect television sets to aerials.

The most common type of twisted pair cable is **unshielded twisted pair** (UTP). This cable comes in various quality grades, the two most common being Grade 3 and Grade 5, which determine the feasible transmission rates and distances. Usually, four twisted pair cables are bundled together, a practice originating from telephone wiring. This is exploited in some computer wiring configurations, for example, a 100 Mbits per second transmission rate over 100 metres can be obtained by using three parallel Grade 3 UTP cables. The same parameters can be achieved using a single Grade 5 UTP cable. The contrast to UTP is **shielded twisted pair** (STP), where the twisted pairs are enclosed within a protective shield to reduce interference. STP cable has better transmission characteristics than UTP, but much greater care has to be taken when installing it.

Twisted pair is cheaper than coaxial cable. However, while both types offer 100 Mbit per second transmission rates over short distances, coaxial cable allows higher transmission rates than twisted pair over longer distances. For example, only about 4 Mbits per second over a distance of 1 km is feasible with twisted pair,

whereas 100 Mbits per second is still feasible over this distance with coaxial cable. The electrical signals travel at about two-thirds the speed of light, so the latency suffered for each bit is not as low as possible, but is still acceptable. Reliability is significantly better than for the various unguided media above, as long as the signals are not subject to any external electrical interference.

Fibre optic cable

Fibre optic cable, which transmits light waves used to represent bits, is the guided technology that is becoming the norm for computer communications. The cable consists of a glass fibre, within a plastic coating to shield it from external light sources. Bits are transmitted digitally, with the absence or presence of a pulse of light used to represent bits of information. Fibre optic cable is now used extensively within the long-distance telephone system, replacing microwave transmission over long distances. It is also finding its way into many streets, albeit stopping short of individual homes, as cable television companies go about their business. Fibre optic cable is very cheap to manufacture, and has significant advantages over electrical cable. It can be used over distances of up to 100 km, and is almost totally reliable since the light signals are not affected by any external electrical or magnetic influences.

In 1997, rates of around 100 Mbits per second were commonplace, with 1000 Mbit per second transmission rates also being used in practice. In theory, rates of millions of Mbits per second are feasible. Each bit arrives at around two-thirds the speed of light, a similar latency as for electrical cable. The feasible transmission rates pose problems for conventional computers because transmission is slowed by the conversion between the electronic signals used by computers and the light signals used for transmission. One possibility under investigation is optical computing, where photonics are used as the internal basis of computers, rather than electronics.

Discussion

The basic properties of the media described above are summarized in Table 1.2, which gives typical present-day parameters, selected from wide possible ranges. The high cost assessment for satellite microwaves comes from the cost of launching the satellite. Once a satellite is in orbit, using it for bouncing microwaves is free. Note that assessing the unguided media as 'free' refers to the media themselves, and does not include the cost of transmitting and receiving equipment.

There is a place for all of these media types to support modern computer communications. For the future, radio and infra-red are likely to dominate mobile computer communications, operating over relatively short distances to allow connectivity between computers and fixed base stations connected to guided media. Fibre optic cabling is already becoming the dominant medium for longer distance communications, and is increasingly used for shorter distance guided communication as well. It has the advantages of speed and reliability over electrical cabling,

Table 1.2 Properties of raw communication media

Medium	Area	Mbits /second	Latency	Reliability	Cost
Radio	100 m 20 km	10 0.05	Minimal	Variable	Free
Infra-red	room	2	Minimal	Good	Free
Terrestial μwave	50 km	50	Minimal	Variable	Free
Satellite μwave	global	50	High	Good	High
Twisted pair	100 m 1 km	100 4	Low	Good	Cheap
Coaxial cable	1 km	100	Low	Good	Moderate
Fibre optic	10 km	1000	Low	Excellent	Cheap

as well as being cheaper, smaller and lighter. Although fibre optic cable has taken over some of the roles originally envisaged for satellites, there will still be a case for using satellites for use in terrain where cabling is hard, or where mobility or broadcasting over a wide area is a desirable feature.

1.3.2 Physical communication services

For communications outside a privately owned area, it is normally necessary to use physical links provided by a public telecommunications carrier (also known as a 'bearer', a 'common carrier' or a 'public carrier'. The organizations that act as carriers are known by various generic names, including **TELCO** (short for 'telecommunications company'), **PNO** (short for 'public network operator') and **PTT** (short for 'post, telegraph and telephone' administration). Until the mid-1980s, it was the case that a single carrier dominated in each country. In the USA, this was AT&T (operators of the Bell telephone system). In other countries, this was usually a branch of the government concerned with all types of communications (as the PTT name suggests).

In 1984, the telecommunications market in the USA was transformed, when AT&T was divided into many smaller component companies which, in turn, encouraged the growth of realistic competitors. Two examples, which have grown into serious long-distance carriers in the USA are MCI and SPRINT. MCI's name stood for 'Microwave Communications Inc' originally, reflecting its use of microwave transmission — however, it has now switched to fibre optic cable. The 'SPR' in SPRINT's name stands for 'Southern Pacific Railway', reflecting its use of cables run alongside existing railway tracks on railway land. In other countries, privatization and deregulation have also led to the growth of a competitive market. This includes telecommunication services being offered by companies involved in providing cable television.

Below, standard facilities that are made available by public carriers are outlined. These range from the normal telephone, through Integrated Services Digital Network (ISDN) services as an advance on the telephone, to dedicated digital transmission lines leased from the carriers. All of these are guided unicast media from a computer communications point of view. The telephone involves analogue transmission, and the others involve digital transmission. Note that this section only includes raw transmission capabilities provided by public carriers. Other, value added, services are described at other points in the book.

Telephone system

The worldwide telephone system, now sometimes known as the **Plain Old Telephone System** (POTS), was designed to transfer human speech. It allows a telephone user to make a call to the huge number of other telephones located around the world. This includes mobile telephones that are linked to the main wiring infrastructure of the POTS using radio transmission.

The POTS transmits analogue electrical waveforms representing speech, with wave frequencies in the range 400 to 3400 hertz being allowed to pass through the system. Although the user interface is analogue, long distance links now use digital transmission, as do many local parts of the system. This aspect of the telephone system is outlined below. There is more discussion of the POTS itself in Section 7.5.1.

Here, the interest is in how the analogue user interface can be used to support computer communications. In fact, this is little different from directly using media such as radio or broadband coaxial cable. Different signal waveforms are used to represent the two bit values. A device called a **modem** (short for 'modulator demodulator') is used to perform the conversion. In the case of the telephone, a waveform representing a sequence of bits must always have its frequency within the allowable 3000 hertz range, otherwise it will not pass through the POTS intact. This fact places fundamental limitations on the rate at which information can be transmitted.

For typical parameters of the telephone system, information theory indicates that the maximum rate is around 35 000 bits per second. Modems able to operate at a 33 600 bit per second rate are available. This is essentially a practical upper limit for modem technology over analogue telephone links and, in fact, such modems may need to operate at lower rates (e.g., 28 800 bits per second) over typical telephone links. Modems capable of a 56 000 bit per second rate began appearing at the end of 1996, but these rely on there being an all-digital telephone link; such links have a 64 000 bit per second theoretical upper limit. Other modems are advertized with higher rates than information theory allows, but these are achieved by compressing the bit stream before it is actually transmitted through the POTS.

Another feature of the POTS as a computer communication medium is that it has a fairly high error rate. This arises from the fact that perfection is not deemed necessary in order for humans to make sense of speech transmitted over the

telephone. People can ignore hisses, crackles and the loss of occasional fragments of speech. In bit terms, a telephone link might have an error rate as bad as 1 in 10 000 bits sent being damaged, although modern digital links are far better. This compares with a figure of 1 in 10^{12} for many direct media, such as coaxial or fibre optic cable. Thus, computers must take considerably more precautions for error detection and correction when using a telephone link. Appropriate techniques are described in Section 3.3.3.

Integrated Services Digital Network (ISDN)

As mentioned above, the internal workings of the telephone system are now predominately digital in nature. When this is extended to provide a digital interface to the end user, the resulting service is called the **Integrated Services Digital Network** (ISDN). The name reflects the fact that the service can be used to transmit not just speech, but also other types of information. Anything that can be encoded using digital bits can be transmitted using ISDN. As with the telephone system, calls can be made to any other subscriber that is connected to the ISDN system. A further benefit of all-digital transmission is that the reliability of bit transmission is dramatically increased.

Design work on the original ISDN service began in 1984, with the aim of establishing a worldwide system by the beginning of the 21st century. This happened before it was clear that there would be a huge demand for both high speed computer communications and real-time video communications in the future. Thus, the service was rather modest in terms of bit transmission rate. It is now known as **Narrowband ISDN** (N-ISDN), to reflect its modesty. The N-ISDN basic service, designed to replace a normal telephone link, offers two 64 Kbit per second digital channels for voice or data transmission, together with a 16 Kbit per second control channel. Each 64 Kbit per second channel is suitable for transmitting speech in digital form, so essentially N-ISDN offers the equivalent of two telephone lines instead of one.

N-ISDN services are now offered in many countries. To allow use of N-ISDN when a remote computer supports it, and the normal telephone service otherwise, devices called **hybrid modems** are available. These allow a computer to use a single piece of equipment to interface with the telephone system, and to gain speed benefits from N-ISDN when possible.

Unfortunately, because of the slow rate at which the details of N-ISDN were agreed, it had already become technically obsolete before a significant number of products had become available. Not only are the bits rates too slow for computer and video communications, but they are also far higher than necessary for speech communication, given rapid advances in the art of speech compression.

The result is that **Broadband ISDN** (B-ISDN) has been developed, with the intention of producing a new system that is robust against future technological developments. The basic standard B-ISDN rate is to be 155 Mbits per second, with higher rates possible. This gives a service that is comparable to most of the

raw media mentioned in the previous pages. The major underpinning technology for B-ISDN is described in detail in Section 7.5.6. Its worldwide implementation presents major technical challenges — in contrast to N-ISDN, which used fairly safe and well-proved technology. In the first instance, B-ISDN services will only be available to fixed-position users, but work is underway towards a high speed mobile B-ISDN service by 2005. This involves extension of the technology used for mobile telephones.

Leased lines

The telephone system and N-ISDN offer slow information transmission rates to the end user, with B-ISDN offering the prospect of a dramatically improved service in the future. If a user requires faster transmission rates at present, then the norm is to make use of a **leased line** provided by a public carrier. This gives a direct link between two fixed points, with transmission rates of up to 45 Mbits per second being available in North America and Japan, and up to 34 Mbits per second being available elsewhere in the world. The difference in rates is due to the slightly different digital transmission mechanisms used in the two different parts of the world. These will be described in outline here.

The first point to note is that the digital transmission channels used by the public carriers operate at bit rates far higher than those seen by individual end users. The basic idea is that a collection of end user calls share the same transmission channel. User calls are given transmission rights in rotation, each one transmitting for a very short period before allowing the next one in turn to transmit. After each one has had a turn, the cycle repeats again. The process is designed so that each user call can transmit at its required bit rate. For example, if 100 users share a 10 Mbits per second channel, then each user is able to transmit up to 100 kbits per second.

The benefit of this approach for the public carrier is that it needs only to provide a total bit rate high enough to support the maximum number of simultaneous user calls that can be active at one time. It does not use a separate physical channel for each different user, with a large proportion of these channels being idle at any given time.

In North America and Japan, the basic unit of digital transmission is known as the **T1 carrier**, which operates at a rate of 1.544 Mbits per second. For telephones, this is used to support 24 separate calls. The value of the analogue signal for each call is sampled 8000 times per second, and each value is represented by seven bits of digital information. This allows 128 signal values to be differentiated — in essence, the analogue signal is turned into a 128-level digital signal with level changes allowed 8000 times per second. There is a little more discussion of how this sampling is done on page 70 in Chapter 2. An extra bit of control information is added to each sample value, giving eight bits per call. The total of 192 ($= 24 \times 8$) bits of call information, plus one extra control bit to give 193 bits in total, is

Table 1.3 Standard carrier channel types

Carrier name	Mbits per second	Shared use
T1	1.544	–
T2	6.312	4× T1
T3	44.736	6× T2
T4	274.176	6× T3
E1	2.048	–
E2	8.848	4× E1
E3	34.368	4× E2
E4	139.264	4× E3
E5	565.148	4× E4

transmitted 8000 times per second, i.e., every 125 microseconds. This leads to the overall 1.544 Mbit per second rate.

In other parts of the world, the basic unit is known as the **E1 carrier**. This is broadly similar, except that there are 32 channels, rather than 24, and eight bits of sample information are used rather than seven. The overall transmission rate is 2.048 Mbits per second.

In turn, the basic T1 and E1 carrier channels can be combined to share faster channels. Table 1.3 shows other members of the T*n* and E*n* families. Note that the higher members of the families have rates that are a little larger than the exact total of the rates that can share these members. This is to allow for the transmission of extra bits of control information. Apart from T1 and E1 themselves, T3 and E4 are the most common types found in practice for leased lines.

These families, where higher transmission rate channels are shared by lower transmission rate channels, are known as the **Plesiochronous Data Hierarchy** (PDH). The term 'plesiochronous' means 'nearly synchronous'. This refers to the fact that the exact bit transmission rates on each of the lower rate channels might differ fractionally, perhaps one is a tiny bit faster than another or a tiny bit slower. The transmission mechanisms used for the PDH carriers make allowances for these clocking variations. Such differences in transmission rate arise from the fact that a different clock is used on each channel. The problem of reconciling different views on absolute time is one of the major topics examined in Chapter 3.

The problem of inconsistent clocks in PDH stems from the fact that it was an evolutionary technology, which had to take into account existing channels with separate clocks. For modern transmission over fibre optic cabling, it was possible to install new systems that have a single master clock. This gives the **Synchronous Data Hierarchy** (SDH) instead. The work on SDH originated from a standard developed by Bellcore in the USA called **SONET** (pronounced 'sonnet', and short

Table 1.4 SONET carrier channel types

Carrier name	Mbits per second
STS-1	51.84
STS-3	155.52
STS-9	466.56
STS-12	622.08
STS-18	933.12
STS-24	1244.16
STS-36	1866.24
STS-48	2488.32

for Synchronous Optical NETwork). This more poetic name is usually applied to the technology, rather than 'SDH', since the two only differ in minor ways.

The basic SONET STS-1 channel operates at 51.84 Mbits per second. This allows the transmission of 810 8-bit values every 125 microseconds. It is sufficient to carry one T3 channel, with some extra padding, or up to 32 separate T1 channels. There is then a hierarchy of channel rates based on STS-1, with a member named 'STS-n' being able to carry n different STS-1 channels. For example, an STS-3 channel can carry three STS-1 channels, then an STS-12 channel can carry four STS-3 channels. The defined SONET rates are shown in Table 1.4. The different rates chosen cover a range that includes values desirable to different constituencies in Europe, Japan and the USA, reflecting the fact that SONET and SDH were designed to be unifying standards for high speed digital transmission. For each defined STS-n bit channel, there is a corresponding OC-n optical channel that is used for the actual transmission. Thus, for example, an OC-3 optical channel is used to implement an STS-3 bit channel.

For SDH, the defined channels and their rates are the same, except that the hierarchy starts at the STS-3 rate, rather than the STS-1 rate. Thus, the basic SDH STM-1 channel operates at 155.52 Mbits per second, and then the other members of the family corresponding to those shown in Table 1.4 are STM-3, STM-4, STM-6, STM-8, STM-12 and STM-16. Note that the STM-1 basic rate is the same as the standard rate chosen for the B-ISDN service — this is no coincidence. However, as will be explained in detail in Section 7.5.6 and in Chapter 11, B-ISDN is not implemented directly using SDH channels. The 'STM' name stands for **Synchronous Transfer Mode,** that is, continuous transmission at fixed time periods. However, an alternative — **Asynchronous Transfer Mode** (ATM) — lies between the B-ISDN service and the physical STM transmission system. This adds a considerable amount of flexibility.

1.4 HOW COMPUTERS COMMUNICATE

From Sections 1.2 and 1.3, it might seem that there is not much to say about about what lies between the requirements of users and the physical transmission services available. However, this is certainly not the case, and indeed forms the subject matter of the rest of this book. If computers are not involved, then it is reasonably practicable to supply user services directly using physical media, as illustrated by the traditional telephone system and radio and television systems. However, when computers are introduced into the picture, there is much more scope for introducing subtlety into the communication system, and also into what can be done using the communication system. This turns the resulting provision into a distributed information processing system.

In this section, there is a short summary of the main concepts that underpin the material in the rest of the book. The material is not organized following an approach that is traditional in almost all books on computer communications: begin with the capabilities of the physical transmission medium, and then show how to add successive layers of enrichment, so that eventually it is possible to construct a service that is adequate to directly support the requirements of a distributed information processing system. Such an approach is founded on two practical considerations. One is the practical engineering of computer communication systems. The other is a conceptual model that is often used by people when they are agreeing on how computer communications should work — a topic covered in detail in Chapter 12.

These practical considerations are both important. However, this book is designed to focus on the more fundamental principles of the subject. That is, the nature of computer communications, the general problems that arise, and then the ways in which the problems can be solved. The book does not start from the solutions, and then try to find some principles. A beneficial side-effect of this approach for the beginner is that most of the principles have close analogies from human communications, and so are familiar from everyday conversation and interaction.

The three main strands of the book are introduced in Sections 1.4.1, 1.4.2 and 1.4.3 below. The first is concerned with the nature of communications. The second is concerned with how communication can be achieved. The third is concerned with human influences on the first two strands.

1.4.1 Information, time and space

A first strand of the book is the identification of three separate components of a communication between computers: **information, time** and **space**. These have already been referred to in this chapter, and correspond to the 'what', 'when' and 'where' of communication.

Information has the starring role, in that the sharing of information is the whole purpose of communication. The nature of information is the topic of Chapter 2.

Much of the content of Chapter 2 is not just of interest from a communications perspective. It is also relevant to a general study of how computers process and store information. Another way of putting this is that information can have an existence independent from any act of communication.

Time refers to how a communication proceeds over the course of a time period. This includes issues like when a communication time period begins and ends, whether communication is continuous over the time period, the time taken for information to travel between computers, and the rate at which information is communicated. Chapter 3 is concerned with the fundamental problems associated with time. Then, Chapter 4 introduces the most common ways in which the temporal behaviour of computer communications is organized. This involves combining various issues considered independently in Chapter 3. Note that time is a concept completely dependent on communications, unlike information and space. It captures the nature of particular acts of communication.

Space refers to the computers involved in a communication, and the channel used to connect them. This first requires some means of uniquely identifying computers and channels. There is a natural measure for time — seconds etc. — but, for space, a natural measure like physical position is not very appropriate since computers are likely to move around. The role of the channel in any communication is to supply a means by which information can flow between computers, in the way demanded by the pattern of information sharing in the communication. For example, in the simplest case, with only two computers sharing information with each other, a channel must allow information to be transmitted in each direction between the computers. Chapter 5 is concerned with the fundamental problems associated with space. Like information, spaces can exist independently of particular acts of communication.

It is convenient to imagine an external observer of a communication. After the communication has happened, the observer will know both the static components of the communication — the information communicated and the space it was communicated over — and the dynamic component — the timing of the act of communication. This particular frame of reference is important. First, the exact nature of a communication might not be known before it happens or during its happening, since many communications are evolutionary in nature. Second, the exact nature of a communication might not be known to any of the computers involved at any stage, since many communications are distributed in nature.

1.4.2 Agreement and implementation

For the three components of communications, there are two issues that are important to each:

- **agreement**; and
- **implementation**.

Agreement is concerned with reconciling differences of viewpoint between computers on the components of communications. This characterizes the distinctive feature of computer communications, that separates the area from the study of individual computer systems. There are two basic ways in which agreement can be achieved: absolute and relative. With absolute agreement, communicating parties operate within some pre-agreed context. With relative agreement, communicating parties reach agreement when an act of communication actually takes place. For example, when a communication begins, one party may just agree to cope with the style of communication being used by another party. However, note that all relative agreement arrangements have some sort of enclosing absolute agreement. This specifies exactly how the relative arrangements will work.

For information, it includes establishing agreement on the type of information that is being communicated. For time, it includes agreement on the time periods for communications and the rates at which communication takes place. For space, it includes agreement on identifiers for computers and channels, and the sets of computers involved in particular communications. These matters are discussed in Chapters 2, 3 and 5 for information, time and space respectively. The meta-problem of how humans agree on how computers should communicate is covered in the final chapter, Chapter 12.

Implementation is concerned with how the desired information, time and space properties of a communication can be realized using some other type of communication, usually something simpler and more physically practical. This is not just the issue of how the communication can be implemented at all, but is also concerned with the *quality* of the implementation. This encompasses issues like speed, latency, reliability, security and cost. Implementation problems are, of course, familiar from the study of computer systems in general. There, a reasonably attractive interface is provided to a human user. However, ultimately, this must be implemented using the very simple electronic logic circuits within the computer. In computer communications, there is a similar process, in that a communication service that is attractive to a user must be implemented in terms of a raw capability that can move bits of information between computers.

For information, two main implementation issues are how complex information types can be represented in terms of simpler types, and how information can be transformed to raise the quality of its communication. These points are addressed in Chapter 2.

For time, the fundamental implementation issue is concerned with how complex communications proceed over time. It may be possible to implement a single communication using a series of simpler communications over separate time periods. Conversely, it may be possible to combine several separate communications into one simple communication that is easier to carry out. Another implementation technique is to adjust the actual timing of a communication to improve its quality. These matters are discussed in Chapter 3, and then practical examples of common implementation packages are described in Chapter 4.

For space, the main implementation problem is concerned with producing channels with the required behaviour, given that physical media have a limited repertoire. This leads to the major area of **computer networking**, where collections of computers and channels act in combination to implement required communication channels. Another implementation technique is to combine several channels so that they share the same channel or, alternatively, to split one channel so that it uses several channels in parallel. The basic space implementation issues are discussed in Chapter 5.

The important space-related topic of computer networking is covered in considerable depth in this book. The extent to which the computers in a network must play an active part in the implementation of communications varies. Chapters 6, 7 and 8 introduce networks in three stages, each requiring a larger contribution from computers as network components, not just as network users. These are accompanied by examples covering the most common types of network occurring in practice. The most famous example included in Chapter 8 is the **Internet**.

Throughout Chapters 2 to 8, implementation is concerned with small steps from the more ideal to the more practical. The first four of these chapters deal with implementation steps for only one of the information, time or space components, as appropriate. The latter three chapters involves a synthesis of information, time and space implementation steps. The ultimate goal of implementation is to completely realize a required user communication in terms of physical communication facilities. To show how this can be done, using small implementation steps, Chapters 9, 10 and 11 contain three case studies, each illustrating a complete implementation process. The details of the case studies have been chosen to reflect applications of contemporary interest, implementation techniques that are of the most frequently encountered types, and physical facilities characteristic of the present day (or future).

1.4.3 Human influences

Chapters 2 to 5 present the fundamental principles of computer communications, then Chapters 6 to 8 introduce somewhat more practical issues related to computer networking, and finally Chapters 9 to 11 describe complete engineering solutions to specific computer communications problems. Underpinning the practical solutions to the fundamental problems is another process: human communication about how to achieve agreement and implementation within computer communications.

Chapter 12 is concerned with this process, which is an example of **standardization**. For many everyday objects, it is useful to have national and international agreements on standards, e.g., for screw sizes or for car tyres. Computer communications is an important special case in that, without standardization of some kind, computers cannot be made to converse unless they are completely identical in all respects.

The main products of standardization are **communication protocols**. These are collections of rules that govern the ways in which computers interact with one

another when communicating. The reason for using protocols is to allow communications with particular requirements to be implemented using communications that are more practically realizable. Thus, a protocol is a collection of prior agreements that is designed to assist in the implementation process. These agreements may involve information, time and/or space matters. The word 'protocol' is used because it matches the human usage, where protocols may be official formulae or bodies of diplomatic etiquette, that is, formalized rules of behaviour between people. This is exactly what is needed for communicating computers.

Before Chapter 12, which is the final chapter of the book, is reached, most chapters make reference to techniques and technologies that have been standardized. In order to avoid confusion, the names of the four most important standardization forces will be mentioned here:

- the Institute of Electrical and Electronics Engineers (IEEE);
- the International Organization for Standardization (ISO);
- the International Telecommunication Union, Telecommunications sector (ITU-T); and
- the Internet community.

These bodies crop up frequently in the text, but the names of some of the other bodies discussed in Chapter 12 also appear. In the case of the Internet community, there are also frequent references to the Internet Request For Comments (RFC) documents, which constitute an easily accessible, and very interesting, collection of literature on computer communications.

Finally, in this introduction, a word of warning is necessary over terminology. Human agreements on how computers should converse have to be precise, otherwise they would be unsuitable for computer use. However, the vocabulary used by humans to discuss computer communications is rather less than precise. Often, different words are used for the same concept. Sometimes, the same word is used for different concepts. This can be very confusing, not just for the novice but also for the expert.

It is common for textbooks to discuss the concepts of a particular area of computer communications using the vocabulary normally used within the area. This book takes an opposite line. The same term is used for a concept, wherever it crops up. Potentially, there is a problem, in that the reader may have vocabulary problems when going on to read specialist literature for particular areas. However, it is felt that the benefits of consistency and uniformity throughout the book far outweigh this concern.

One example will serve to illustrate this point. In this book, the word **message** is used to denote any unit of information that is shared between computers. For example, at the simplest level, a message might just be a single bit of information. In the literature, at least eight other words or terms are used in various different contexts: cell, datagram, data unit, frame, packet, segment, slot and transfer unit. It is interesting that the term 'data unit' originated in an international standardization

effort that sought to invent a neutral term, different from all those in use already. However, it is a bit of a mouthful, so 'message' has been preferred as a more natural term to use here.

1.5 CHAPTER SUMMARY

Communication is beneficial to the human race, and can also be beneficial for computers. It enables the sharing of information, and of information processing facilities, as well as giving access to human computer users.

Future usage of computer communications is centred around the notion of a worldwide information superhighway. This has its roots in three existing areas: computer-oriented communications, which arose as a means of connecting computers and their human users; telecommunications, which arose as a means of connecting human users; and radio and television, which arose as a broadcast education, information and entertainment service. A convergence of the three technologies has occurred, as a result of the omnipresence of computers inside everyday objects and devices. The resulting provision, incorporating interactive, multimedia access to information, makes far higher demands on computer communication systems than has been the case in the past.

Computer communications are supported by physical media linking the computers. These may be guided or unguided, unicast or broadcast. The main media types are radio, infra-red, microwave (including satellites), copper cable and fibre optic cable. These have differing properties, in terms of information transfer rate, distance spanned and reliability. For communications outside a small, privately owned area, the services of public carriers are usually required. These include the telephone system or more computer-oriented services such as ISDN and leased lines. For the future, Broadband ISDN is set to become the standard digital information transmission provision.

To bridge the gap between the required user information service and the physical media, the techniques of computer communications are required. The key components of any particular communication are the information shared, the time period of the communication, and the space of computers and channels involved. Carrying out the communication requires agreement on the nature of the communication between the computers and channels, and the implementation of the required communication in terms of communications that can be physically realized. Communication protocols are central to this — these are rules for how communications proceed, agreed by human designers. The construction and selection of protocols, and other areas of agreement on computer communications, involves standardization processes. These impose constraints on what is possible and not possible in practical computer communications.

1.6 EXERCISES AND FURTHER READING

1.1 Survey the ways in which you communicate with other people, noting the differing purposes of communications with different people.

1.2 Make a list of all of the computers that you use at home, at work, or elsewhere. Remember to include the computers hidden inside everyday objects. Consider any benefits that would arise if any of the computers was able to communicate with any of the others.

1.3 Suppose that you wish to communicate with someone who is standing beside you, but that neither of you speaks the same language. How might communication take place?

1.4 Discuss the ways in which technologies of all kinds have made the notion of a 'global village' possible.

1.5 Collect references to the 'information superhighway' from newspapers and magazines. What do these popular media usually mean when they use this phrase?

1.6 If you have access to the Internet, compare remote `telnet` access to a computer with direct WIMP access to the same computer. Why is `telnet` access still in existence?

1.7 Discuss possible benefits of client-server distributed computer systems, compared with stand-alone computer systems.

1.8 Fax has become a very popular telecommunications service. Why has this happened, given that there are alternatives such as teletex or computer-based electronic mail?

1.9 If a cable television company operates in your area, find out the range of services that are offered. Do these extend beyond television?

1.10 If you have access to the World Wide Web, search for extra information on any topics in this chapter that are of interest to you. How easy do you find such searching, and how useful are the results?

1.11 Outline the facilities that you would want your own personal knowledge robot to have. How easy do you think it would be to implement these facilities?

1.12 Around 65 000 films have been made since the introduction of cinema. Do you think it is feasible to operate a video on demand service that allows any user to see any film that they want at any time?

1.13 Radio waves, infra-red and microwaves are three examples of electromagnetic phenomena used for bit transmission. How useful do you think visible light would be as an unguided transmission medium?

1.14 Obtain information about the various low-orbit satellite proposals, in particular find an explanation of why the Iridium project involves only 66 satellites, whereas the Teledesic project involves 840.

1.15 Compare the information sizes given in Table 1.1 and the information transfer rates for media given in Table 1.2. Comment on the types of multimedia user

facilities that could be provided if the raw media were directly used to transfer information between users.

1.16 Find out who the main public telecommunications carriers are in your area. What services do they offer to the end user, and how much do these services cost? In particular, is an ISDN service available to the domestic telephone user?

1.17 If you have access to a modem, look at the list of facilities and options it offers to its user. How many of these are comprehensible to a person who knows little about the details of computer communications? What is the maximum information transfer rate possible?

1.18 From Table 1.3, work out the number of bits per second that are overheads for control purposes when (a) 24 T1 channels share a T3 channel; and (b) 64 E1 channels share an E4 channel.

1.19 The OC-3 carrier is shared by three OC-1 carriers. There is a variant called 'OC-3c' that is used as an unshared carrier. The total user information rate of OC-3 is 148.608 Mbits per second, whereas the user information rate of OC-3c is 149.760 Mbits per second. This compares with the raw bit rate of 155.520 Mbits per second. Calculate the number of control information bits that are used in each 125 microsecond cycle by (a) OC-3; and (b) OC-3c.

1.20 Look at everyday objects to see if they have been subject to national or international standardization. What compromises do you think were made in order to standardize any particular object?

Further reading

At this stage, there is no need for further reading on the core areas covered by the rest of the book. However, some readers may be interested in finding out more about the surrounding context: the uses of computer communications, and/or the physical support for computer communications.

For the reader who is interested in the uses of computer communications, there is a variety of possible sources. First, on computer-oriented communication, the wide range of books about the Internet and, specifically, about the World Wide Web, convey a reasonable impression of how computer communication facilities are presented to users currently. However, any reader who is exposed to a computing environment with communication facilities should encounter specific technical descriptions targeted at users. For general issues involving the telephone system, radio and television, and the convergence of these technologies with computer technologies, more specialized reading is available. *Multimedia Networking* by Agnew and Kellerman (Addison-Wesley 1996) deals with this convergence, and the emergence of interactive, digital, multimedia communications.

For the reader who is interested in the physical underpinning of computer communications, there is a great deal of technical literature. This book deliberately skims over the details of how bits can be transmitted; other books focus entirely on this issue. Basic technologies, such as electronics, photonics and radio, are

covered in the electrical engineering or physics literature. *Telecommunication System Engineering* by Freeman (Wiley 1996) covers the technological bases for telecommunications and computer communications.

For more details of the techniques used for transmitting binary information using digital and analogue signals, and in fact for more technical details of many topics introduced in this book, two general purpose textbooks of interest are *Data Communications, Computer Networks and Open Systems* by Halsall (Addison-Wesley 1996) and *Data and Computer Communications* by Stallings (Prentice-Hall 1997).

CHAPTER

FIVE

SPACE

The main topics in this chapter about space are:

- sharing of information by computers over a channel
- identifiers for computers and channels
- absolute and relative agreement on space between computers
- broadcast, multicast and unicast channels
- computers as communication filters and switches
- introduction to computer networks and inter-networks
- splitting and multiplexing of communications

5.1 INTRODUCTION

In this chapter, the model of communication is extended into its final form. In earlier chapters, a communication is seen as an act that shares information between two computers via a channel, within some time period. This has a significant limitation: communications take place in a little private world that contains only

two fixed computers and one channel between them. Such an assumption is a large restriction in terms of modelling communication between all computers of the world.

In order to deal with this, it is important to describe how information sharing is carried out in **space**. That is, which computers and channels are involved. The issues that arise are similar to those already seen when considering both information and time. First, it is necesssary to have agreement on where information sharing takes place. Second, it is necessary to have implementation that allows required communication spaces to be created using channels that allow different — usually simpler — communication spaces. The implementation work may not just be concerned with achieving communication, but also with achieving it at some particular level of quality.

Human experience

To make the discussion on space clearer, consider the following examples drawn from everyday human experience:

- a television programme
- a tourist sending postcards to friends
- a telephone call

These examples are in a similar vein to those used in Chapter 3, but are slightly varied to illustrate various space-related points effectively.

The space for all of these examples will be regarded as consisting of a collection of communicating people, together with some sort of communication medium. In the first example, the people are the programme presenters at a television station and the viewers of the programme, and the medium is a television broadcast. In the second example, the people are the tourist and his or her friends, and the medium is the postal system. In the third example, the people are the two callers, and the medium is the telephone system. In all three cases, the media involve some kind of technology. However, for other examples such as the presentation by a speaker or the tourist information desk from Chapter 3, the medium would just be direct human speech.

Computer communications principles

For the remainder of this book, the term 'a communication' will be used to mean an act of information sharing among a group of two or more computers via a channel, occurring within some time period. For any particular communication, the computers involved are described by a set containing a unique identifier for each one. The identifiers are unique within the enclosing universe of computers under consideration. This universe is visible in its entirety to an external observer, but not

necessarily to the individual computers. The channel also has a unique identifier among all possible channels. A communication is illustrated in Figure 5.1.

The information-passing capabilities of channels must be extended under the new definition of a communication. Rather than just being able to transfer information in one or both directions between two computers, a channel must be able to pass information in whatever ways necessary to achieve the required information sharing among the group of computers. Each communication space involves just *one* channel. With this definition, the channel can be regarded as the component that carries out the communication. Like the computers, the channel can have an existence independent from particular communications and so is, in this sense, a static component of the communication.

Just as the term 'message' is used consistently throughout this book when information is being discussed, so the terms 'computer' and 'channel' are used consistently when space is being discussed. Before proceeding further, it is worth noting some synonyms used in practice for 'computer' and 'channel' when describing practical computer communication systems. Other terms for computer include: **device**, **end system**, **host** and **station**. Later in this chapter (on page 153), a list of synonyms for specialized computers forming integral parts of communication systems is included. Other terms for channel include: **cable**, **circuit**, **interface**, **link**, **medium**, **path**, **pipe**, **tunnel** and **wire**. The word 'connection' may also be used to mean 'channel' in some practical contexts, but it is used consistently throughout this book to mean a particular type of time package, as described in Section 4.5.

5.2 AGREEMENT ON SPACES

For a communication to take place, the computers and channel must agree on rules that give a shared view of the space. There are two main problems. First, there must be an agreed scheme for identifying computers and channels. Second, given an identifier scheme, there must be agreement upon which computers and channel form the space.

5.2.1 Identifier schemes

With time, there is a natural physical measurement system available, and this could give a direct identifier scheme for time periods. For space, it is not so straightforward. A natural physical system would be one based on geographical position, using a suitably accurate positioning system. It is technically feasible to find very accurate positions for objects, for example, using the satellite-based Geographical Positioning System (GPS). However, such a scheme is not particularly appropriate, because physical position is not a very good way of deriving unique identifiers for computers or channels that are usable over a non-trivial time span. This is particularly so these days, with the advent of mobile computing, since it

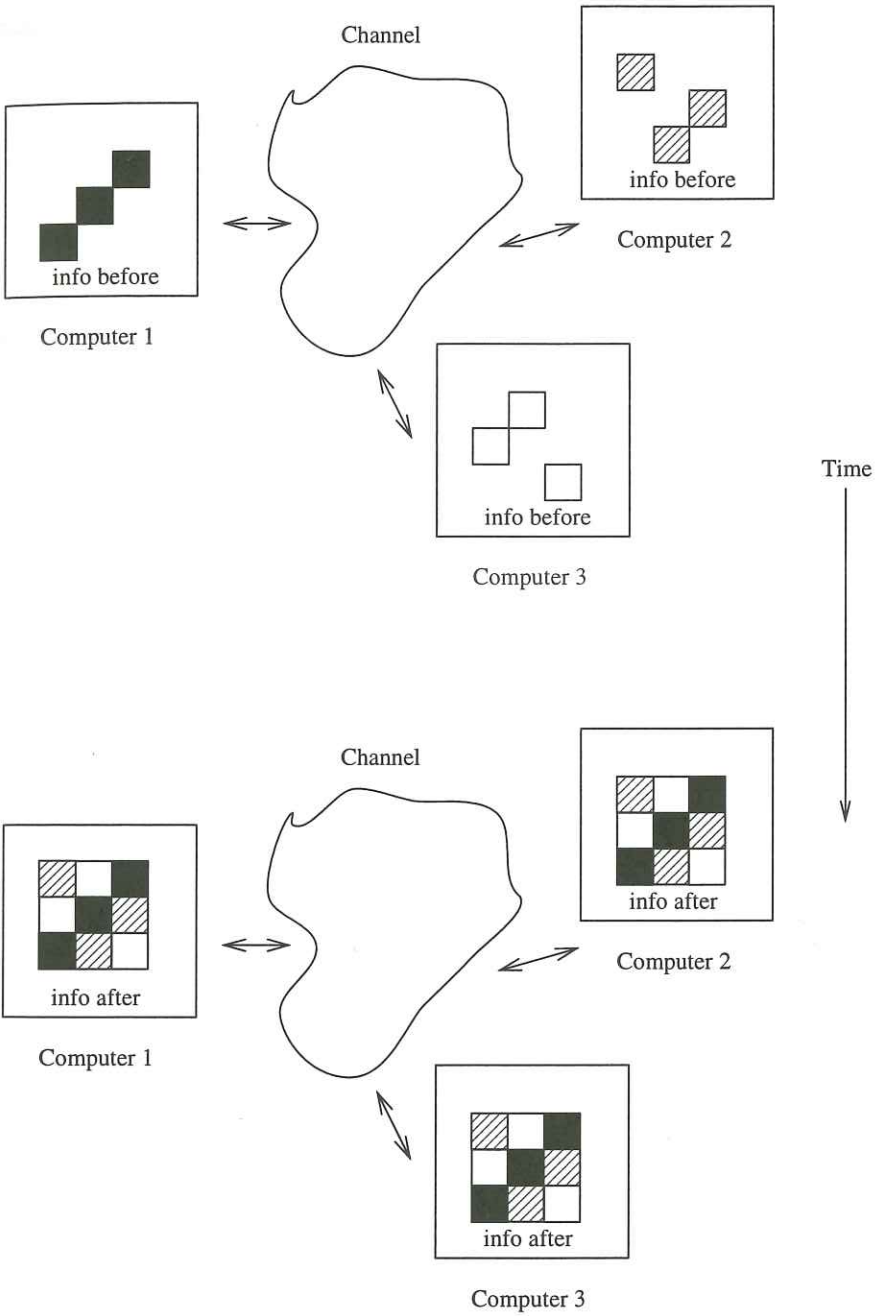


Figure 5.1 A communication

is no longer the case that computers stay in the same place semi-permanently. As an alternative, some type of identifier scheme chosen for human and/or computer convenience is used.

Human experience

Tourist sending postcards to friends

For the postal system, people are identified by their names and addresses. The effect is that individual people have a unique identifier that can be used to direct mail to them. Here, each friend has an identifier, as does the tourist, based on his or her current location. The medium is a part of the worldwide postal system, and a possible unique identifier could be derived from the address of the postbox used by the tourist together with the addresses of the friends. The exact way in which names and addresses appear varies from country to country, and indeed also according to individual taste. However, most addresses consist of several components which have geographical significance. These might include apartment/house name or number, street name, town/city name and country name.

Telephone call

A natural way of identifying telephone callers is just to use their two telephone numbers (assuming that each number is only used for one call at a time). The medium is a part of the worldwide telephone system, and a possible unique identifier could just be derived from the two telephone numbers. Telephone numbers are structured to have several components. In very many cases, this structuring has a geographical significance. Thus, the first component of an international telephone number usually identifies a country. Subsequent components can then be chosen in any way by the particular country, without having to worry about clashes with numbers in other countries. As an example of within-country structuring, in many cities, a second component identifies the city, a third component identifies an area of the city, and finally a fourth component is a unique identifier within that area of the city.

One problem with both the telephone number scheme and the postal address scheme is that, when people move, their telephone number and address usually have to change. That is, their location in the identifier structure changes.

Computer communications principles

Flat or hierarchical

The identifier scheme may be either **flat** or **hierarchical**. In a flat scheme, identifiers have no added structure, and so allocating a new unique identifier involves picking a string that has not been used before. In a hierarchical scheme, identifiers consist of a series of components that position them inside a hierarchy. The hierarchy is chosen for administrative and/or technical convenience. Postal addresses

and telephone numbers are both examples of this. For example, a unique 10-level postal address usable by aliens might be:

```
491 Princes Street
Edinburgh
Scotland
United Kingdom
Europe
Earth
Solar System
Milky Way
Local Group
The Universe
```

Allocating a new unique identifier only involves picking a string that has not been used before in the appropriate local part of the hierarchy. This is rather easier in the locally flat scheme than it is within a globally flat scheme, unless the total number of identifiers being used is small.

Names or addresses

When the identifiers have a mainly human significance, for example, organizational, functional or positional, they are usually referred to as **names**. When the identifiers have a mainly computer or communications significance, they are usually referred to as **addresses**. To ensure uniqueness, it is often the case that addresses are fairly long and convoluted. This is not a major problem for computers to handle. However, if humans are involved as users, then some means of allowing more human-friendly ways of identification than just the computer identifiers is required. Therefore, there may be both names and addresses for particular computers or channels. Translation of names to addresses (or vice versa) is usually achieved by a **directory** service. The general idea is familiar from telephone directories where, using a personal name, plus perhaps address or occupation, it is possible to look up the unique telephone number for a person.

Computer communications examples

IEEE 802 identifiers

IEEE 802 identifiers is used in IEEE (Institute of Electrical and Electronics Engineers) standard local area networks, which are discussed in detail in Chapter 6. Identifiers have 48 bits, with the first two bits being **00**, to indicate that the remaining bits are uniquely assigned under a scheme administered by the IEEE. This is shown in Figure 5.2. In principle, the identifier scheme is flat. However, because blocks of identifiers are assigned to different equipment manufacturers by the IEEE, there is a two-level hierarchy in practice. Each network interface has a fixed identifier for its lifetime, with a first component indicating the manufacturer and the second component being a serial number from that manufacturer.

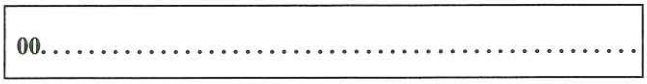


Figure 5.2 IEEE standard 48-bit identifier

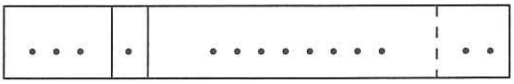


Figure 5.3 Basic structure of an ITU-T X.121 standard identifier

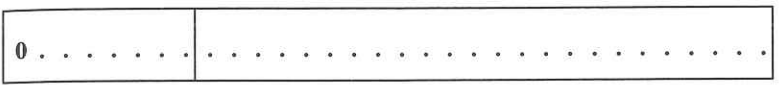
ITU-T X.121 identifiers

As explained later in Section 5.3.1, a computer network is a collection of computers and channels, together with agreements on how they inter-communicate. ITU-T X.121 identifiers are used in long distance computer networks operated by state or private telecommunications authorities. This particular type of network is discussed in detail in Section 7.4.3. An X.121 identifier consists of up to 14 decimal digits and, unsurprisingly, is reminiscent of a telephone number in appearance. The basic structure is shown in Figure 5.3. The first three digits are a country code, and the fourth digit is a network number within a country. For countries expected to have more than 10 internationally connected networks, more than one country code is allocated (e.g., the USA has codes 310 to 329). The remaining digits can be allocated in any way by a particular network. As an example, the next few digits might be used for a town or city code, as in telephone numbers. The final two (13th and 14th) digits, if present, are usually used as an identifier for a particular process within the computer identified by the first 12 digits of the identifier.

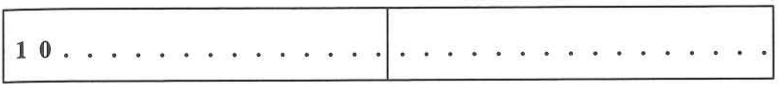
Internet Protocol identifiers

The Internet Protocol (IP) is mentioned as an example in Chapter 4, and is described in detail in Section 8.3.2. The Internet is an example of an inter-network. The characteristic of an inter-network is that it is a collection of networks of differing types. Inter-networks, in general, are the subject matter of Chapter 8. A particular feature of an inter-network is that the component networks might have different standard identifier schemes for computers. Thus, an inter-networking service must use an over-arching common identifier scheme that allows inter-working between the different networks. IP has such a scheme associated with it.

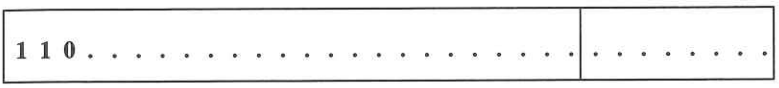
For IP version 4, which was used from the beginnings of the Internet and only began to be very gradually replaced by IP version 6 in 1996, identifiers are 32 bits (four bytes) long. Each computer in the Internet has a unique identifier, and this has two hierarchical components. The hierarchy reflects the fact that the Internet is composed of a collection of computer networks. Thus, in an Internet identifier for a computer, the first part is a network identifier and the second part is a computer



(a) Class A address format: 8 network and 24 computer address bits



(b) Class B address format : 16 network and 16 computer address bits



(c) Class C address format: 24 network and 8 computer address bits

Figure 5.4 IP identifier structure

identifier within the network. There are three basic types of identifier, Class A, Class B and Class C, as illustrated in Figure 5.4. Identifiers beginning with the three bits **111** are used for other purposes, mentioned in Section 8.3.2. The two-level identifier hierarchy allows network numbers (or blocks of network numbers) to be allocated to organizations, which can then allocate unique computer numbers within the network.

The designers of the identifier scheme allowed for up to $2^7 = 128$ large networks with up to $2^{24} = 16\,777\,216$ computers (Class A), up to $2^{14} = 16\,384$ medium networks with up to $2^{16} = 65\,536$ computers (Class B) and up to $2^{21} = 2\,097\,152$ small networks with up to $2^8 = 256$ computers (Class C). At that time, this seemed reasonable, since the world view was one of a relatively small number of large or medium size networks and a large number of small networks.

However, in the Internet, the problem since has been a fast growth in the number of networks containing 300 or more computers. Class B identifiers must be used for such networks, and there are relatively few available. For the next generation Internet Protocol version 6, identifiers are increased to a 16-byte length, and have a richer hierarchical structure. This scheme is designed to satisfy all envisaged demand for identifiers, as well as making allocation easier. This new form of IP identifier is described in Section 8.3.2.

ISO NSAP identifiers

The ISO standard connectionless inter-networking service, which is the ISO analogue of the Internet IP service, is described in Section 8.3.3. ISO NSAP (Network Service Access Point) identifiers for computers are used in this service, and

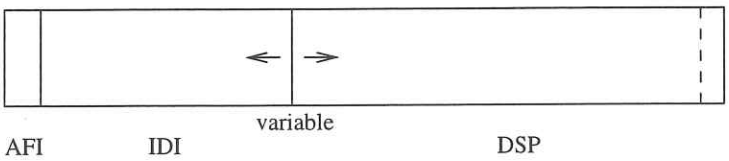


Figure 5.5 Basic structure of an ISO NSAP identifier

the standard protocol that implements it. The scheme used does not begin from scratch, as with IP version 4 identifiers. Instead, it unifies various existing identifier schemes used for both computer networks and telecommunications networks.

A full identifier can be up to 20 bytes or 40 decimal digits long. The basic structure is shown in Figure 5.5. All NSAP identifiers have a hierarchical structure with three components. There may be further structure within the second and/or third components. The first component consists of two decimal digits, which form an **Authority and Format Identifier** (AFI). This indicates which identifier scheme is used for the second component of the identifier, and the authority responsible for allocating identifiers within this scheme. It also indicates whether the third component is expressed in decimal or binary notation. There are seven identifier schemes included, coded by decimal numbers as follows:

Authority and scheme	Decimal	Binary
ITU-T public data network (X.121)	36	37
ISO data country codes (ISO 3166)	38	39
ITU-T telex (F.69)	40	41
ITU-T telephone (E.163)	42	43
ITU-T ISDN (E.164)	44	45
ISO international codes (ISO 6523)	46	47
Local non-standard scheme	48	49

Note that X.121, an earlier example, is one of the schemes included. The ISO data country codes scheme has a code for each country in the world, whereas the international codes are used for projects that span more than one country. The local scheme identifier is included to allow private identifier systems to be used.

The second component of an NSAP identifier is the **Initial Domain Identifier** (IDI). This is an identifier within the scheme specified by the AFI. For example, if the AFI is 36 or 37, then the IDI is a 14 decimal digit X.121 identifier. If the AFI is 38 or 39, the IDI is a three decimal digit country code, e.g., 840 for the USA. Other schemes also have appropriately lengthed IDI components. In the case of AFIs 48 and 49, the IDI component is null.

The third component of an NSAP identifier is the **Domain Specific Part** (DSP). This contains the actual identifier for a computer within an identifier scheme administered by an authority specified by the combination of the AFI and IDI fields. The maximum length of the DSP component depends on the length of

the IDI field, since the total length of an NSAP identifier cannot exceed 40 decimal digits (which corresponds to 160 bits, since each decimal digit is equivalent to four bits).

By convention, the final byte (or, equivalently, the final two decimal digits) of the DSP is an identifier for a process within a computer identified by the previous digits of the DSP. Apart from this, the DSP usually has a hierarchical structure, which follows the internal structure of inter-networks and networks. There is more discussion of such hierarchies on page 305 in Section 8.3.3, and also on page 379 in Chapter 10.

5.2.2 Absolute and relative spaces

Given that an identifier scheme is available, the communicators must agree on the space used. In a simple situation, such as the earlier model of communication, the agreement is implicit: there, the space consisted of the two computers and the connecting channel. In other, very static, circumstances, there might be permanent absolute agreements on the set of computers that communicate and on the channel used. However, apart from at the lowest levels of communication, such agreements are unusual. Mobile computing makes permanent agreements even less likely. Instead, the space is usually agreed upon when a communication actually begins, or is about to begin.

Human experience

Television programme

For a programme broadcast by a television station, the channel is the only component of the space that must be known by all of the communicators. The people watching the programme will usually know which television station is transmitting it, since television channels are normally only used by a single station. However, this is not necessary. Instead, viewers may learn the identity of the station from the transmission itself. The television station will not normally know who is actually watching the programme. This is because the information flow is a unidirectional broadcast. In summary, the communication space is relative to the actual communication, and participants have partial knowledge of it. In more specialized cases, such as closed-circuit television, the space might be absolute, and communicators will have exact knowledge of it.

Tourist sending postcards to friends

The tourist is aware of all of the communicators in the space, since he or she is choosing to whom to send postcards. The tourist also selects the channel, by posting the cards in one or more postboxes and addressing each card to identify a mailbox for delivery. Thus, the tourist has a complete view of the space. The space is relative to the actual act of communicating. Each recipient has only a partial view of the space, obtained from the communication. This view shows

that the recipient and the tourist are communicators (unless the tourist sends an anonymous postcard, of course). The postcard is unlikely to reveal which other people were recipients of postcards. A recipient also does not have complete knowledge of the channel, only the fact that it included his or her mailbox. Given that the information flow is unidirectional, the postcard recipients do not require any further knowledge of the space.

Telephone call

The communication space for a normal telephone call is relative to the call starting. The caller knows who the called is, by selecting the number dialled. The called's telephone might automatically supply the number of the caller, or the caller is likely to identify himself or herself at the beginning of the call. Given that both parties know each other's telephone number, then both parties have exact knowledge of the telephone channel. Even if the human recipient does not know the caller's number, the telephone equipment does, and so a bidirectional information flow is possible.

Computer communications principles

As in the case of time, there is some subtlety in the simple definition of a communication. Again, this occurs because it is not made clear how the space is measured by the participating computers. There are two possible mechanisms. The first is an absolute scheme, where the communicators are described as a particular collection of computers and a channel. The second is a relative scheme, where the communicators are just described as 'the computers and channel involved in the communication.' In summary, in the first case, the communication depends on the space and, in the second case, the space depends on the communication.

In general, it is not easy to have pre-arranged absolute sets of communicating parties. Until a communication actually starts, it is not easy to decide exactly who the communicators are. Thus, most spaces are relative to the actual communication starting. In the absence of some absolute map of the communication space, each computer will have its own local version. This need not be a perfect record of the actual space. For example, if a computer has a completely passive role in the communication, i.e., it only experiences an information inflow, then it might not need to know exactly where the information is coming from. Further, when the implementation of complex spaces is considered later in this chapter, it will become clear that a computer might only need to know about a subset (perhaps of size one) of the communication space, in order to participate fully. Compared with time agreements, agreement on maps of spaces are normally coarser grain, that is, change far less frequently than agreements on time periods.

5.2.3 Examples of communication spaces

Two computers linked by a physical channel

When two computers are directly linked by a physical channel, there is an absolute communication space. It consists of the two computers and the channel. There is no need for any further knowledge to be derived from actual communications. This was the situation in the model of communication used until the current chapter.

Two computers using a connection-oriented service

In this example, two computers communicate using a connection-oriented communication service, akin to the telephone system, rather than using a permanent direct physical channel. Before the connection is established, that is, before the actual communication begins, there is no absolute map of the space. Knowledge of the space is acquired by the communicators in much the same way as happens with the two parties in a telephone call. It is normal for the computer starting up the connection to send its identifier to the other computer as part of the establishment process. The only significant difference from a telephone call is that the identifier scheme is likely to be richer. This is because several processes within a particular computer are likely to be participating in different communications at one time. Therefore, a unique identifier for the channel must be derived not just from identifiers for the two computers, but also extended identifiers for the particular processes within each computer.

Two computers linked using a connectionless service

This example is similar to the previous one, except that the temporal nature of the communication service is different: it is connectionless rather than connection oriented. In space terms, sending a message using a connectionless service is almost the same as communicating using a connection-oriented service. That is, when a computer sends a message, it is fixing the communication space. When a computer receives a message, it then becomes aware of the space and the other communicator. The communication space ceases to exist as soon as the message has been sent, just as the space of a communication using a connection-oriented service ceases to exist when the connection is closed. The information flow is unidirectional, so there is no particular need for the receiving computer to acquire any knowledge of the communication space. In fact, in this example, there might be multiple recipients of the message, if the connectionless service allows it. However, any particular recipient need not necessarily know whether or not there are others.

5.3 IMPLEMENTATION OF REQUIRED SPACES

Agreement on spaces is only concerned with how the communicators agree on the collection of computers, and the channel, for a communication. However, it is very often the case that there is no direct physical realization of the communication space. One implementation method, which is covered in Section 5.3.1, is to involve computers actively in order to restrict or extend the set of computers that are linked by a channel. A second method, which is covered in Section 5.3.2, is to reduce or increase the number of channels that are used to support channels required for communications. Note that both of these implementation methods refer to spaces, rather than individual communications. This is because space, like information, is a static component of communications that can have an independent existence.

5.3.1 Filtering and switching

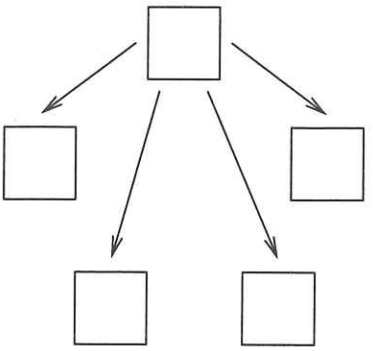
First, it is useful to introduce three technical terms used to describe different primitive types of information flow between computers in a space:

1. **broadcast:** information flows from one computer to all of the other computers
2. **multicast:** information flows from one computer to a subset of the computers
3. **unicast** or **point-to-point:** information flows from one computer to another

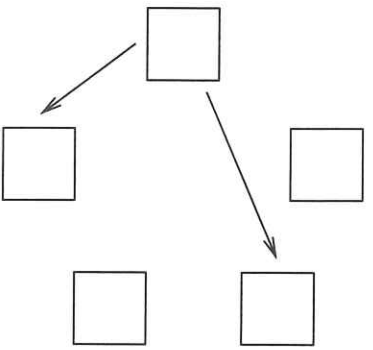
These are illustrated in Figure 5.6. Of course, broadcast and unicast are just extreme special cases of multicast. These three types of communication are the most common in practice, since they make a useful simplification that aids understanding, implementation and use of communications, and the channels that enable them.

Virtually all physical communication channels fall into one of two categories. One is channels that support broadcasting over a collection of three or more computers. The other is channels that support unicasting over a collection of two computers. This categorization was introduced in Section 1.3, where different physical media were classified as being either broadcast or unicast. To repeat a warning in that section, based on everyday human experience, it is tempting to think of radio or microwave as broadcast and copper or fibre optic cables as unicast. This is how they appear in conventional applications like radio and television, and telephone systems. However, one should also recall mobile telephones and cable television as counter-examples.

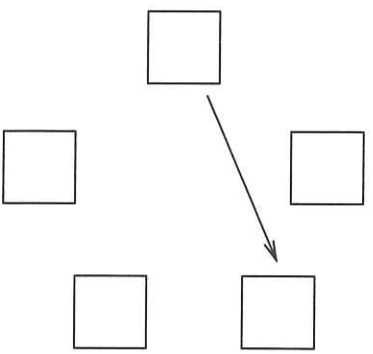
A further, related, primitive type of information flow is the **anycast**. This has the information flow of a unicast, but also has some characteristics of a multicast. It involves communication between one computer and *any one* of a subset of the computers. To an external observer, an anycast appears like a unicast. To the sending computer, an anycast appears like a multicast. The idea is that communication takes place with one of a collection of computers, but it does not matter which one, because all perform a similar function. For example, the communication might be a request-response handshake between a client and any one of a collection of equivalent servers.



(a) Example of broadcast communication



(b) Example of multicast communication



(c) Example of unicast communication

Figure 5.6 Broadcast, multicast and unicast information flows

Communications that involve more complex 'm to n' information flows must typically be implemented using communications with these more primitive types of 'one to n' information flow. This can be done by segmentation of such communications into sub-communications that have simpler information flows.

It is first necessary to look at the problem of implementing communications with one type of primitive information flow, by using one or more communications that involve different primitive information flow(s). Such implementation work is needed if a suitable channel is not available to allow the required communication to take place.

Human experience

Television programme

The idealized communication is of a television station multicasting to all the programme viewers. Unless the television station is operating in a fairly small area, this is unlikely to be the reality. Each viewer tunes into a broadcast transmission. However, this is likely to be from one of several distribution centres. The transmission may come from a UHF transmitter, from a satellite or from a cable centre. The television station transmits the programme to the various distribution centres, again using a variety of technologies. Thus, the implemented multicast communication may involve more than one sub-communication, and these sub-communications might be unicast or broadcast in nature. The whole process is coordinated by the television companies and the television distributors, who may or may not be the same organization.

Tourist sending postcards to friends

Here, the communication involves the tourist multicasting to the group of friends. There is no direct multicasting. Instead, a separate postcard is sent to each friend. Further, each postcard is not directly carried from postbox to mailbox. Instead, it is transferred by various postal staff through various postal sorting offices. Each step can be viewed as a unicast sub-communication. Thus, there are two stages: the multicast is implemented as a set of unicasts; and each unicast is implemented as a sequence of unicasts. The whole process is coordinated by the international postal system, which involves the cooperation of different national organizations.

Telephone call

The telephone communication is of a unicast nature. However, there is not usually a direct physical link between a pair of telephones. Instead, a link is made using a sequence of physical links through various telephone exchanges. Therefore, the required unicast communication is implemented using a sequence of unicast sub-communications. The whole process is coordinated by telephone companies or organizations, one or more depending on the parties involved in the telephone call.

Computer communications principles

Computers as filters or switches

Given that a computer is able to participate in communications using one or more channels, there are two activities it can carry out in order to assist in implementing communications requiring channels that are not directly available. The first is **filtering**, which means selectively ignoring some of the information that is communicated to it. This can be used to reduce the number of computers involved in a communication. The second is **switching**, which means selectively communicating further some of the information that is communicated to it. This can be used to increase the number of computers involved in a communication.

Computers and channels forming networks

For computers to assist in implementing communications, agreements on information, time and space are necessary. These guide the filtering and/or switching actions performed by the computers. The norm is to have agreement packages among a collection of computers, that result in the implementation of services providing channels that allow certain information flows between these computers. Any service that supports unicast, multicast and/or broadcast communications can be expected to allow any computers in the collection to be the source and the destination(s). That is, allowable information flows along channels are symmetrical from any viewpoint among the computers in the space.

A collection of computers and channels that uses general agreements on information, time and space matters among the computers in order to implement idealized channels using the actual channels is called a **computer network**.

The computers involved in a network can have a dual role. They are responsible for implementing the protocols associated with network agreements on information, time and space. However, they can also be users of the network service. The whole point of creating the idealized channels is to allow communication between computers for general purposes. In a single computer system, there may be a natural division between where network duties and general-purpose duties take place, for example, between hardware and software, or between operating system and user process.

Some computers may be solely concerned with implementing the network, and have no other duties. In such cases, they will be termed **filters** or **switches** in this book, depending on role, to reflect their more specialized networking roles. There are various synonyms used for such computers when describing practical communication networks. These include: **bridge**, **gateway**, **intermediate system** (used in contrast to 'end system' — a computer that just uses the network service), **node**, **protocol converter**, **relay**, **repeater**, **router** and **store-and-forward system**.

Note carefully that this implementation process can be repeated many times over. The channels made possible by one or more computer networks can be used, in turn, to construct another computer network that can provide different

channels. Because of this, extra terminology creeps in. For example, the term **sub-network** (or just **subnet**) is used to describe a network that is a component, with other other sub-networks, of a larger network. The term **inter-network** (or just **internet**) is used to describe a network that is constructed from a collection of smaller networks.

Computer communications examples

Unicast or multicast channel implemented using a broadcast channel

Suppose that there is a network that enables arbitrary broadcast channels among its computers. This can be used as the basis for implementing arbitrary unicast or multicast channels. To do this, the information to be shared over such channels is transmitted to every computer in the network using the broadcast capability. Computers then filter information received over the broadcast channel. Specifically, there must be an information agreement that all communicated information includes a way of specifying the identifier(s) of the computer(s) that are intended to receive the information. Then, all computers look at the identifier information. Only those with appropriate identifiers accept the rest of the information, and so participate in the communication. The other computers ignore the communication. This style of operation is usually termed 'promiscuous' — all computers can see all communicated information, but have to be trusted not to abuse the privilege.

Unicast channel implemented using different unicast channels

Suppose that there is a collection of computers, with unicast channels between certain pairs of computers. The computers and channels may or may not be organized into networks — it is not important. Given an appropriate arrangement of channels, this can be used as the basis for a network that can implement arbitrary unicast channels. It is necessary to establish information and space agreements so that unicast channels can be implemented between pairs of computers that do not already have direct unicast channels. This is feasible as long as it is possible to find a route formed from direct channels leading from one computer, via intermediate computers, to the other computer. Then, a series of sub-communications can take place to implement the required communication. The original computer, and the intermediate computers, are involved in switching, in order to forward information. To do this correctly, the communicated information must include information about the intended final destination of the information. Intermediate computers filter out the rest of the communicated information. To construct such a network, there must be at least one possible route between every pair of computers in the network. If there is more than one possible route, then there is scope for improving the quality of a communication, by choosing a route that is quicker, cheaper or more reliable, for example.

Broadcast or multicast channel implemented using unicast channels

Broadcast or multicast channels can also be implemented on a network of the type just described. One simple way is just for the information source to conduct a unicast communication with each recipient individually. In the case of broadcasts, the source must have some means of knowing the identity of all computers in the network, so that it can communicate with them. A better way is to enhance the switching behaviour of computers. In this case, the communicated information must include information about the identifiers of the computers that are to receive it. Then, the original computer, and any intermediate computers, conduct more restricted sub-communications over multicast (possibly just unicast) channels with one *or more* direct neighbours that can assist in implementing the full communication. That is, the remaining broadcast or multicast communication is partitioned between channels to neighbours at each stage. The effect is of a tree of channels, rooted at the information source, and with recipient computers as leaves and possibly nodes.

Networks and graphs

The language, both verbal and pictorial, of **graphs** is sometimes convenient for discussing spatial aspects of networks. Graphs have **vertices** (sometimes called 'nodes') and **edges** (sometimes called 'arcs') that connect pairs of vertices. A **directed graph** has **directed edges**, that connect ordered pairs of vertices, with the underlying idea of moving from the first vertex of the pair to the other.

In the case of computer networks, a network can be regarded as a directed graph, with computers represented by vertices and unicast channels represented by directed edges. The direction of edges corresponds to the information flow of channels. There need be no extra representational apparatus for multicast or broadcast channels, since these can be modelled by multiple unicast channels, as seen in the previous example. A graph representation of a small example network is shown in Figure 5.7. Note that all edges are bidirectional, which is normally the case in practice. That is, if it is possible to have a unicast channel in one direction, it is possible to have one in the other direction.

Given directed graph representations, some of the ideas from the examples above fit in well with concepts from the world of graphs. A **path** between vertex s and vertex t in a directed graph is a sequence of directed edges

$$(s, v_1), (v_1, v_2), \ldots, (v_{k-1}, v_k), (v_k, t)$$

for some k. A graph is said to be 'connected' if there is a path between all pairs of vertices. Thus, paths and connectedness are directly relevant to constructing networks supporting arbitrary unicast channels, as in the second example above. Connectedness means that there are routes between all pairs of computers. Figure 5.8 shows an example of a (contorted) path between the top-left and bottom-right vertices of the graph introduced in Figure 5.7.

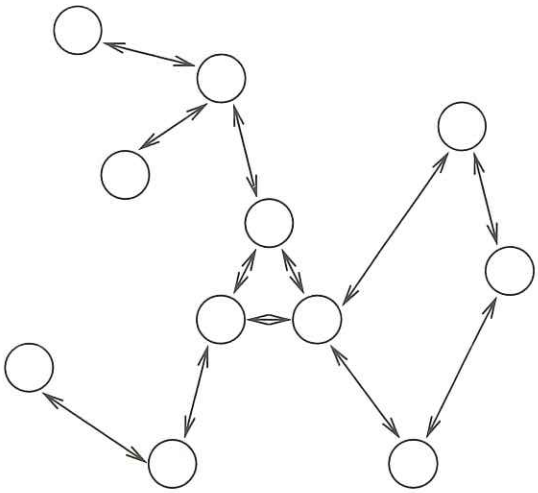


Figure 5.7 Example graph representing network

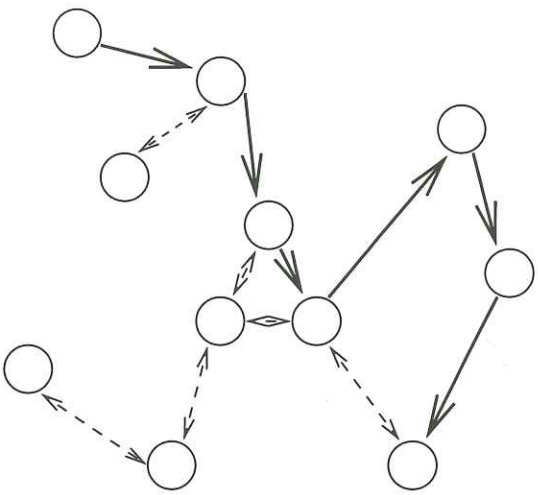


Figure 5.8 Directed path in example graph

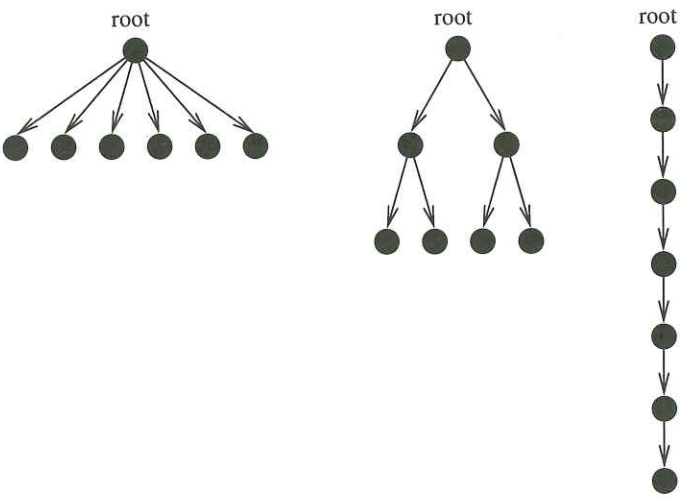


Figure 5.9 Three examples of trees with seven vertices

Trees can also be defined within graphs. A tree in a directed graph is a subgraph composed of a set of ℓ vertices and $\ell - 1$ directed edges between vertices in this set. Each edge is directed at a different vertex. The single vertex that does not have an edge directed at it is the root of the tree, and there must be a path from the root vertex to every other vertex in the tree. Figure 5.9 shows three examples of trees containing seven vertices.

A **spanning tree** of a graph is a tree that includes all the vertices of the graph. Trees and spanning trees are relevant when constructing networks that support multicasts and broadcasts, as in the third example above. Figure 5.10 shows an example of a spanning tree rooted at the leftmost vertex of the graph introduced in Figure 5.7.

There are two main advantages in using graphs as an abstraction of computer networks. One is as a way of hiding detail. For example, instead of having each vertex representing one computer, it is possible to have each vertex representing a collection of computers, perhaps the members of a component network. Edges then represent channels between computers in one collection and computers in another collection. Thus, graphs can be used as a tool for representing different levels of network implementation.

The other main purpose is to give access to the large collection of algorithms and data structures that the computer science community has made available for solving graph-related problems. There are standard methods not just for finding paths and trees in graphs, but also for finding cheapest paths and trees when cost measures are associated with edges.

Examples include Dijkstra's algorithm and the Bellman-Ford algorithm, both of which, given a vertex, find the shortest paths from that vertex to all other

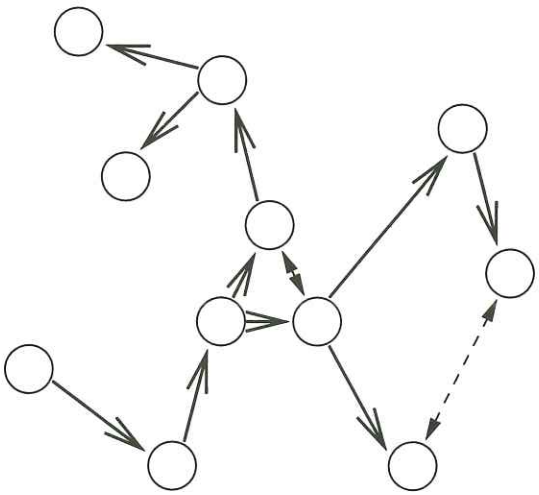


Figure 5.10 Directed spanning tree of example graph

vertices. The Bellman-Ford algorithm is more general than Dijkstra's algorithm, because it allows edges costs to be negative. However, its run time is proportional to the number of vertices multiplied by the number of edges, whereas Dijkstra's algorithm's run time is proportional to the number of vertices squared. These algorithms are useful when designing and operating networks, as will be seen in Chapters 7 and 8. It is not necessary to know the details of how the algorithms work in order to understand how networks work but, for interested readers, the details of Dijkstra's algorithm can be explored through Exercise 5.13.

Chapters 6, 7 and 8 are concerned with the detailed information, time and space agreements that are necessary in order to make unicast, multicast and broadcast networks work. The chapters are differentiated by the complexity of the switching task that has to be undertaken by computers in order to implement the required channels. Filtering also features, but not as a central concern. In essence, switching is a process of receiving a message as a result of one communication, and then selectively forwarding it using another communication.

In the networks of Chapter 6, computers either do no switching, or do decisionless switching in the sense that messages received on one channel are always transmitted on one or more other channels. It is assumed that communications over component channels have similar information, time and space characteristics. The effect of the non-existent or unsubtle switching is that the basic capability of the network is to supply broadcast channels, because no computer in the network introduces selectivity in its switching. Therefore, such networks are entitled **Message Broadcasting Networks**. This style of networking has its roots in physical channels that enable broadcasting.

In the networks of Chapter 7, computers do selective switching. Messages received on one channel are selectively transmitted on one or more other channels. It is assumed that communications over component channels have similar information, time and space characteristics. The effect of the switching is that, depending on the type of selectivity chosen, the network can supply unicast, multicast and/or broadcast channels. Focusing on the method of implementation, such networks are entitled **Message Switching Networks**. This style of networking has its roots in physical channels that enable unicasting.

Finally, in the networks of Chapter 8, computers again do selective switching. However, the switching goes beyond the simple relaying of messages from one channel to another. A switching computer must also deal with an assumption of significant differences in the information, time and space characteristics of communications on its different channels. This situation usually arises because the component channels are, in turn, implemented by networks with different information, time and space agreements. For this reason, and following the normal convention of the computer communications trade, the networks of the chapter are entitled **Inter-networks**.

Before embarking on an in-depth study of networking, however, there is another basic space implementation technique to consider. This has various uses, including the sharing of channels and the introduction of parallel transmission over channels.

5.3.2 Splitting and multiplexing

The second basic implementation technique for space involves the complementary procedures of **splitting** and **multiplexing**. These refer to the ways in which several channels can be used to implement one channel, and in which several channels can be implemented by sharing one channel.

Human experience

Television programme

Continuing from the last section's example (on page 152), suppose that the television programme is a transmission of a musical concert. The quality of sound in television broadcasts is usually inferior to the stereo quality available on radio. Therefore, pictures and sound are often broadcast simultaneously over the two different media. This is a simple example of splitting: one channel is implemented using two independent channels.

The particular television programme is only one of several or many that can be received at any time. However, these are all delivered using a single physical medium, whether it is electro-magnetic radiation or cable. This is a simple example of multiplexing: several independent channels are implemented using a single channel. In this case, **frequency division multiplexing** (FDM) is used: each television programme is broadcast in a different frequency band.

The frequency bands are all contained within the overall frequency range of the medium. Thus, the different communications do not interfere with each other.

Tourist sending postcards to friends

A tourist who wished to save money on postal costs when visiting a far-away destination could make use of multiplexing. He or she could place all of the postcards into an envelope and send it to one friend, then ask the friend to distribute the postcards. This would be a typical example of why multiplexing is used. That is, to share an expensive communication channel between several channels. Of course, multiplexing also happens at a higher level in this example. The tourist's postcards are only a few among very many items in the postal system. At various stages in the journey, postcards will travel with many other items in a single mail container.

Telephone call

Multiplexing is a common feature of the telephone system. At most stages, many calls share the same physical links. For example, one caller might be within an organization with a private switchboard. Then, all of the organization's calls are likely to share the same physical link to the external telephone system. As another example, many national or international telephone calls share physical links between cities or countries at any one time. The use of carrier channels for doing this is described in the discussion of leased lines in Section 1.3.2.

Computer communications principles

Just as segmentation and concatenation are opposites for one another for time period implementation, so splitting and multiplexing are opposites of one another for space implementation.

Splitting

An alternative term for splitting is **downward multiplexing**. The set of communications that involves one channel is partitioned into subsets of communications that involve different channels. The set of computers involved in the space of each communication is not affected. This process is illustrated in Figure 5.11. The benefit of splitting is that it allows the introduction of parallelism — several communications can proceed at the same time if they have independent channels. This allows the overall quality of the communication system to be improved, for example, by making individual communications faster.

Multiplexing

An alternative term, matching the alternative term for splitting, is the longer **upward multiplexing**. The sets of communications that involve several different channels are combined into a set of communications that involves one channel.

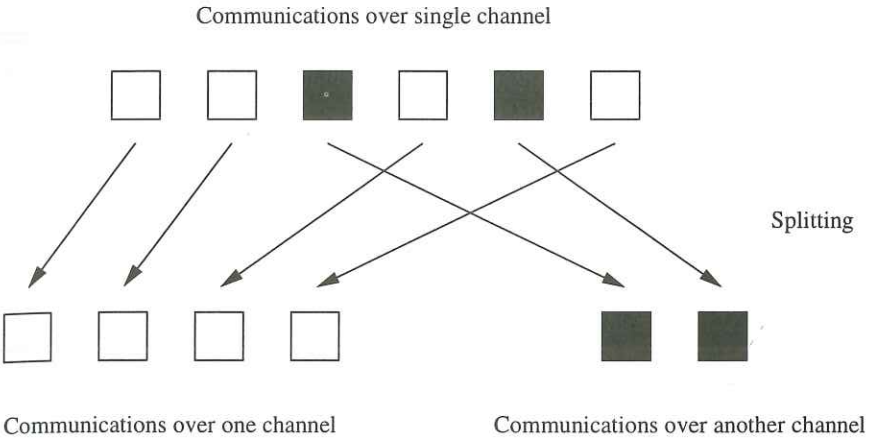


Figure 5.11 Splitting a channel over several channels

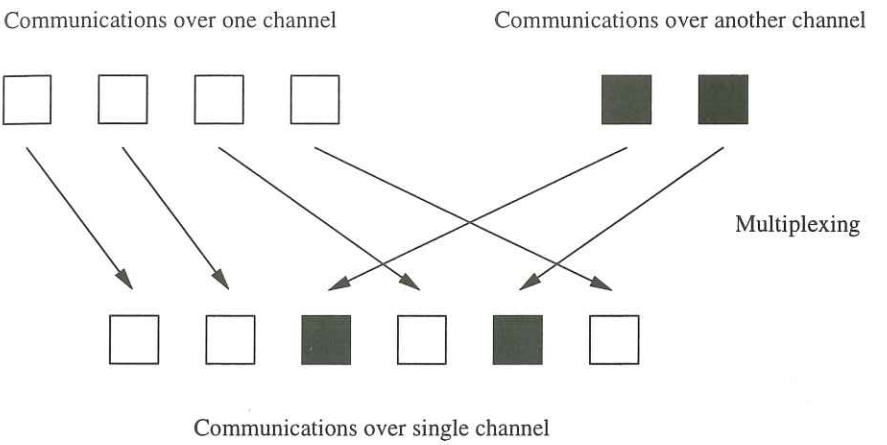


Figure 5.12 Multiplexing several channels on to a single channel

The sets of computers involved in each of the original sets of communications must be the same. This process is illustrated in Figure 5.12.

One frequent restriction on multiplexing is that all of the communications in the resulting set must have disjoint time periods, since the channel only allows one communication at a time. A simple counter-example is a bidirectional channel between two computers: this allows two independent unidirectional communications to take place at the same time, in opposite directions. However, this is not a major exception to the general rule, since the channel is really just a pair of independent unidirectional channels.

If attempted multiplexing does lead to a clash, then **contention** is said to occur. In general, to achieve the disjoint property, each communications that uses a channel being multiplexed must have its time period adjusted appropriately so that it does not overlap with others. This preliminary time implementation step is akin to flow control, but is carried out for the benefit of the channel rather than the computers. It relies on the multiplexed communications being asynchronous.

The advantage of multiplexing is that it allows the sharing of channels — several channels can make use of the same channel. This gives a cost saving, compared with providing multiple parallel channels. There may also be cost savings for the computers themselves, since they only have to deal with one channel at a time, rather than several.

Computer communications examples

Simplex and duplex channels

When a channel is used to connect two computers, there are several different possible communication behaviours that it might allow. In the model of a communication system introduced at the beginning of Chapter 2, a channel allows information to flow in either direction between the computers, at the same time if necessary. This is known as a **full duplex channel**. However, to simplify many of the communications described so far, they have had a unidirectional information flow. A channel that can only support communications in one direction is known as a **simplex channel**. In between these two types of channel, there is a third type: the **half duplex channel**. This type allows information to flow in either direction, but not in both directions at the same time, One use of a half duplex channel is for handshake-style communications.

In space terms, the different types of channel — simplex, half duplex and full duplex — are increasingly powerful. That is, a simplex channel can be implemented easily using either of the others, and a half duplex channel can be implemented easily using a full duplex channel. It may be possible to implement the features of a full duplex channel using a half duplex channel, essentially by multiplexing two unidirectional channels onto the single half duplex channel. However, to do this, it is necessary to segment any bidirectional communications into unidirectional sub-communications that have disjoint time periods. A half duplex or full duplex channel can be implemented easily using two simplex channels, one for each direction of information flow. This is an example of splitting.

Multipeer channels

A **multipeer** channel (alternatively termed a 'multipoint' channel) is a channel connecting n computers, that supports all of the n possible broadcast information flows. Typically, but not necessarily, at most one flow can be in progress at one time. However, the important point is that this channel can support the n different flows in any sequence over time. In Chapter 6, a dominating issue is how n separate

broadcast channels, one for each possible source, can be multiplexed onto a single multipeer channel.

In the opposite direction, the example of splitting a duplex channel over two simplex channels can be generalized to give a way of splitting a multipeer channel over a set of broadcast channels. If there are n different sources for the multipeer channel, then n broadcast channels are needed, each providing the information flow needed for a different source.

One example of this situation in practice is the **multicast group,** which is a collection of computers that send multicasts to each other. A multicast group of n computers requires n one to $(n - 1)$ multicast channels to implement the complete information flow of the group. Thus, over the n computers involved, the multicast group is just a multipeer channel implemented using broadcast channels within the group. In the context of an enclosing network, each of these broadcast channels is implemented using a multicast channel provided by the network.

Piggybacking

The discussion of connection-oriented time packages in Section 4.5 was restricted to a service offering only a unidirectional information flow from one computer to the other. That is, the service included just a simplex channel. However, to implement the service, the full duplex nature of an underlying channel was needed in order to send back control information. Real connection-oriented services provide a full duplex channel to their users, to allow a bidirectional information flow, i.e., information sharing, between computers.

From the point of view of the communication service, the duplex channel provided to a user is implemented as two simplex channels that are essentially independent. The only dependency between the channels is the fact that they are created and destroyed simultaneously, when the connection begins and ends, respectively. This means that, in principle, two duplex channels are appropriate for implementing the two simplex channels, each used for user information in one direction and control information in the other direction. However, in practice, there is only *one* underlying duplex channel, and so multiplexing is necessary. That is, a user information communication channel in one direction must be multiplexed with a control information communication channel for the other direction, and vice versa.

To improve efficiency, it is not always necessary for the two types of communication to be disjoint in time. Most protocols for implementing connections use information-carrying messages that also have space for control information for the other direction of transfer. This is known as **piggybacking**, to convey the idea of control information hitching a lift along the channel. As an example, all TCP messages carry both the serial number of the information present and also the serial number and window size for the other direction. In fact, there is only one type of message, so all acknowledgements are piggybacked, on zero-size information messages if necessary. The LAPB protocol also allows piggybacking,

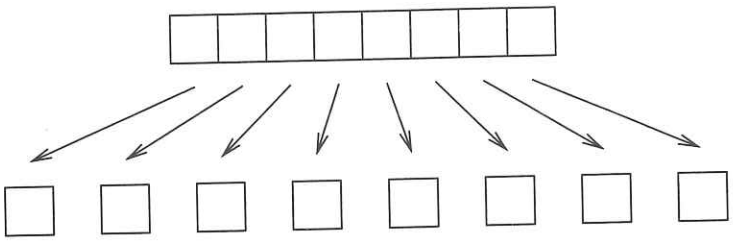


Figure 5.13 Splitting a character communication channel

although there are explicit control information datagrams as well. However, the TP4 protocol does not include piggybacking as a feature.

Sharing a single character

On page 88 in Chapter 3, the asynchronous method of transmitting eight-bit character representations could not be segmented in time further, since the 10 necessary bits are transmitted synchronously. An alternative way of transmitting each character is to use segementation and splitting. If eight separate binary channels are available, then each character communication can be segmented into eight bit communications, and then the eight bits can be transmitted in parallel, with a significant time saving. This is illustrated in Figure 5.13. In general, this approach can be used for larger strings of bits. For example, 32-bit or 64-bit buses are used within computers or between computers and peripheral devices. Parallel channels with synchronized data transfer are not normally used for communication between computers, both because of timing problems and because of expense.

Time division and statistical multiplexing

A normal requirement for multiplexing is to avoid contention. That is, to ensure that multiplexed communications have disjoint time periods. As already mentioned, achieving this is a flow control problem, in which the need for flow control arises from the channel rather than from the communicating computers. In some cases, flow control required by the computers might be adequate to allow multiplexing. For example, if there is a master computer and other slave computers, then the time-disjoint communications enforced by the master will ensure that there is no contention for the channel.

For the classical multiplexing problem of sharing one actual unidirectional channel between several required unidirectional channels, there are two main approaches to flow control — one synchronous, the other asynchronous. With **time division multiplexing** (TDM), the required channels take turns to use the actual channel, each being allowed a fixed period of time. Thus, the communications being multiplexed have synchronous time periods which are disjoint. This arrangement ensures that the multiplexing is fair, but is wasteful if some of the communications are irregular and so waste their guaranteed time slots. TDM has a

long history in the telecommunications world, and the methods of sharing a carrier channel described in Section 1.3.2 illustrate this.

An alternative, which is the most common in computer communications, is **statistical multiplexing**. This treats the communications to be multiplexed as being asynchronous. Ideally, communications on the required channels occur during disjoint time periods, in which case there is no problem. However, if there is contention, the multiplexing scheme is able to adjust the time periods to make them disjoint. This is done by delaying communications until a later time when the shared channel is free, using buffering associated with the channel. The effect is that a transmitting computer begins a communication at its choice of time, but that the receiving computer sees the communication beginning at a rather later time. The 'statistical' nature of the multiplexing refers to the fact that the buffering is adequate to cope with almost all cases of contention. In some extreme cases, such as communications taking place on all required channels at the same time, it may not be able to cope and information will be lost.

ISO standard transport service

As mentioned on page 133 in Chapter 4, the ISO standard connection-oriented transport service may be implemented by one of five standard transport protocols: TP0, TP1, TP2, TP3 or TP4. Three of these — TP2, TP3 and TP4 — allow for multiplexing. This capability is the main feature that distinguishes TP0 and TP2, and TP1 and TP3. Several channels provided by the transport service may be multiplexed onto a single channel provided by the underlying service. This allows a possibly expensive channel to be shared and so to be used more productively. The TP4 protocol also allows splitting. A channel between two computers using the transport service might be split between several different channels provided by an underlying service. This allows several communications to proceed in parallel, using separate channels. The parallelism is not synchronous, so the protocol must ensure that correct sequencing of information is preserved.

5.4 CHAPTER SUMMARY

A communication involves information sharing among a collection of computers using a single channel. The computers have to be given unique identifiers. Identifier schemes may be either flat or hierarchical, the latter making allocation of unique identifiers easier. All of the computers involved in a communication have to have some knowledge of the map of the space. This might result from an absolute agreement, where the space is fixed before communication. More usually, the map of the space is relative to the communication. That is, computers acquire knowledge of the space when a communication begins, or as it continues.

The communications that can be carried out using computer communication systems are usually forced to have one of three forms:

- unicast — one computer transmits information to another
- multicast — one computer transmits information to several others
- broadcast — one computer transmits information to all others

and most physical channels support either unicast or broadcast transmission directly. In order to implement any particular pattern of information sharing, an appropriate channel has to be implemented using the types of channel available.

A main way in which this can be done is by using computer networks, which are collections of computers and channels between the computers. Protocols, involving agreement on information, time and space issues, are needed in networks. Using these, networks harness the capabilities of the computers in order to provide desirable channels using the available channels. The two main tasks carried out by computers in networks are filtering and switching. The capabilities of networks can be categorized by the complexity of the switching tasks carried out, giving message broadcasting networks, message switching networks and inter-networks.

To implement channels efficently, it may be necessary use different underlying channels that link the same set of computers. This may involve splitting, where communications over one channel are partitioned over several different channels. Alternatively, it may involve multiplexing, the opposite of splitting, where communications over several different channels share the same single channel.

5.5 EXERCISES AND FURTHER READING

5.1 For everyday communications that you have with other people, or with groups of people, identify the communication spaces, state how much you individually know about the spaces, and explain how you acquire that knowledge.

5.2 Investigate the history of the Global Positioning System (GPS), and find examples of its usage in everyday life.

5.3 For a computer system that you use, find out any addresses or names that it has for communication purposes, and investigate whether these are from a flat or hierarchical identfier system.

5.4 The Internet **Domain Name Service** (DNS) is a distributed mechanism for mapping names to addresses. Look at RFC 1034 and RFC 1035 to discover how it works, but do not attempt to understand all of the details.

5.5 Read RFC 1884, which describes the IP version 6 addressing scheme.

5.6 Find out why the designers of the IP version 6 addressing scheme decided not to just use the existing ISO NSAP identifier scheme.

5.7 Discuss computer communication applications where it is necessary to know the exact composition of the communication space, and others where it is not.

5.8 Devise examples of how multicasts could be made available in a natural way for human communications.

5.9 Suppose that you live in a house shared with others. Give an example of how you might act as a filter, and as a switch, when collecting a mail delivery for the house.

5.10 For each computer network that you have access to, list the number of computers involved, the types of channel that the network implements, and the names of protocols used for agreement within the network.

5.11 The Internet is the most famous example of an inter-network. How many networks are within the Internet at the present time?

5.12 Graphs are a convenient abstract representation of computer networks. Give examples of other real-life systems that could be represented using graphs.

5.13 Look up the details of Dijkstra's shortest path algorithm, which should be found in any standard book on algorithms. For the example in Figure 5.8, use the algorithm to find the shortest path from the top-left vertex to the bottom-right vertex, assuming that every edge has the same cost (= 1, say).

5.14 Describe clearly the difference between splitting (described in this chapter) and segmentation (described in Chapter 3).

5.15 Describe clearly the difference between multiplexing (described in this chapter) and concatenation (described in Chapter 3).

5.16 The description of piggybacking involved multiplexing because it forms part of the implementation of a service that, essentially, provides its user with two unidirectional simplex channels. Therefore, the user has to perform splitting of its information-sharing communications to use these channels. Explain why, if the service user could directly make use of the full duplex channel that underlies the service, then piggybacking of acknowledgements would just be an example of concatentation.

5.17 Give examples of practical situations where time division multiplexing would be efficient.

5.18 Investigate the ISO standard Transport Protocol, to discover other differences between TP0–4, not mentioned in this chapter, or in Chapter 4.

Further reading

This chapter is intended to be self-contained — drawing on existing human experience of space issues in communication, in order to introduce the main ideas impacting on computer communication. The next three chapters, Chapters 6–8, expand the coverage of how computer networks are designed and used. Later on, Chapters 9 to 11 include material concerned with the role of space in particular areas of computer communications. More details about graphs, and graph algorithms, appear in textbooks on algorithms, for example, *Introduction to Algorithms* by Cormen, Leiserson and Rivest (MIT Press 1990).

CHAPTER

SIX

MESSAGE BROADCASTING NETWORKS

The main topics in this chapter about message broadcasting networks are:

- information, time and space basics of message broadcasting
- local, and metropolitan, area networks
- implementing a multipeer channel using simple channels
- multiplexing broadcasts onto a multipeer channel
- guided technology networks: ethernet, token ring/FDDI, token bus, DQDB and 100 BASE VG-AnyLAN
- unguided technology networks using radio and infra-red transmission

6.1 INTRODUCTION

As Chapter 5 explains, there are increasingly complex ways in which computers can contribute switching abilities to assist in supporting a computer network. This

chapter concentrates on the simplest cases: no switching is done; or non-selective switching is done. Because of the lack of discrimination, such networks can be regarded as **message broadcasting networks**, which have the basic characteristic that every communication that takes place has a space that contains all of the computers in the network. Therefore, if n computers are involved in the network, then n different one to $(n - 1)$ broadcast channels are made available by the network. In addition, broadcasting networks can supply unicast and multicast channels as well. These are implemented using filtering by computers, as described on page 154 of Chapter 5.

Message broadcasting networks may be implemented using broadcast-style channels, or using unicast-style channels. Reflecting the fact that the computers play a modest networking role in the switching sense, broadcasting networks are normally low-level implementations, in that the channels used are physical channels, rather than channels created by further lower-level network implementations. Therefore, this chapter focuses on such networks. However, in principle, the underlying ideas could be used to implement higher-level broadcasting networks using channels provided by other lower-level networks.

Traditionally, message broadcasting networks have been nearly synonymous with **Local Area Networks** (LANs), which are networks with a physical radius of no more than one or two kilometres. This is because the high communication speeds needed to make broadcasting a realistic possibility were only attainable over short distances. However, this assumption is changing with advances in technology. **Metropolitan Area Networks** (MANs), which are networks with a physical radius that can span a city, are often organized as broadcasting networks. Meanwhile, the techniques used for message switching networks, traditionally associated with longer distance networks and the subject of Chapter 7, are now being used in some LANs and MANs, as well as in experimental Home Area Networks (in the home) and Desk Area Networks (on the desk). Thus, it is now more useful to have a classification based on networking style, rather than based on physical distance spanned.

6.1.1 Information basics

The messages transmitted by broadcasting networks contain at least three basic components:

- source identifier: the identifier of the sending computer
- destination identifier(s): the identifier(s) of the intended recipient computer(s)
- information: the information being communicated by sending the message

Note that, for low-level broadcasting networks, messages are usually called 'packets'. In fact, this terminology is normal for all types of networks based on physical channels, including the types described in Chapter 7. Thus, the terms 'packet

broadcasting network' and, particularly, 'packet switching network' are often encountered in practice.

The source identifier is included so that a recipient can tell from where the message originated. This is needed because the network supports n different communication spaces, where n is the number of computers in the network. Including the source identifier in each message acts as a convenient way for computers to deduce the communication space from the actual communication. A relative scheme like this is necessary unless the recipient can differentiate between communication spaces by the network service providing independent access to n different channels.

The destination identifier(s) are included so that recipients can tell whether or not the message is intended for them. This allows unicast or multicast communications to be implemented using the broadcasting network. A recipient just filters out any message that does not carry its own identifier.

The identifiers used for computers have network-wide significance, and are of a fixed size agreed for the particular network. One example would be the 48-bit identifiers defined in the IEEE 802 series standards for LANs, and mentioned on page 143 in Chapter 5. This identifier scheme allows worldwide uniqueness of identifiers. For private LANs, an alternative also allowed by the standard is to use 16-bit identifiers. However, it is not possible to mix the two identifier lengths in the same network. The use of fixed-size identifiers, and also locating the identifiers in a fixed position within each message, simplifies the handling of messages by computers.

The information component of a message is just a sequence of bits (in fact, more likely a sequence of bytes) as far as the network is concerned. In principle, the information component might have an arbitrarily chosen length. Some networks impose a fixed length on this component, so all messages have a uniform length. This, again, simplifies handling of messages. If not, almost all other networks impose a maximum length on this component, and hence on the overall message length. This is because, as Section 6.1.3 explains, the collection of broadcast channels is multiplexed onto a single multipeer channel. A restriction on length ensures a degree of fairness in this sharing of the multipeer channel. Some networks also impose a minimum length restriction on messages. The reasons for this are discussed in detail later but, in summary, it may be because either the broadcast channel has characteristics that may cause it to 'lose' very short messages, or the communication of very short messages is inefficient.

The standard network message format, which may include other components in addition to the above, will have a direct representation in terms of bits. This representation is an absolute agreement, fixed for all computers participating in the network. The representation of bits depends on how the broadcast channels are implemented. If a single underlying broadcast-style channel is used, then the representation will be fixed. However, it may vary if a collection of different unicast-style channels is used.

In summary, from an information point of view, there is absolute agreement on the type of information communicated: the standard network message. There is also absolute agreement on how this is represented in terms of bits.

6.1.2 Time basics

Two absolute time measurements are usually possible for message broadcasting networks based on physical channels. The first is the latency, that is, the delay between the time that a computer begins a transmission and the latest time that any computer in the network begins to receive the transmission. This is affected by the physical technology used to implement the network, as well as the physical distances involved in the network. For example, a network of computers in the same room, connected by fibre optic cables, will have a very low latency. A network of computers distributed around the world, connected by satellite links, will have a high latency.

The second absolute time measurement is the rate at which transmission can take place, measured in bits per second. Again, this is affected by the physical technology used to implement the network. In the case of a genuine broadcast technology, the rate is a measurable global parameter. Where the network is implemented using unicast technology, the rate is determined by the slowest link contributing to the implementation. This is measurable, since the links are part of a tightly-knit collection, and their performance is dictated by the demands of the network.

In combination with the absolute agreements on information, the duration of communication time periods becomes measurable. For a fixed message size, the time between starting to transmit a message and finishing the reception of the message at all computers, is measurable. It is equal, in seconds, to:

$$\text{latency} + (\text{message size}/\text{rate})$$

where the latency is measured in seconds, the message size in bits, and the rate in bits per second. This puts absolute time bounds on the time period for the communication of the information in the message. Where there are minimum and/or maximum size restrictions on message size, lower and/or upper bounds respectively can be put on the communication time period. Note that, although the duration of communication time periods can be absolute, the message communications are asynchronous, since they may start at arbitrary times.

In high speed networking circles, a further measure of interest is the product of the rate and the latency, usually called the **bandwidth*delay product**. This gives a measure, in bits, of the amount of information that can be simultaneously present on the channel. That is, it measures the number of bits that can be transmitted during the time it takes for one bit to travel from the transmitter to the receiver.

In terms of time packages, the service offered by message broadcasting networks is normally an unreliable connectionless service. There is no segmentation, so information is presented to the service in message-sized units, and each unit

results in a single communication of a message. The bits of a message are communicated synchronously. A simple handshake time package is an elaboration on the unsegmented time package that is sometimes found. With this, a message is broadcast with a single intended recipient. This is followed by an acknowledgement message being broadcast by the recipient, intended to reach the original sender.

Given that distances are usually short in broadcasting networks, and the quality of the physical media is high, the service offered is usually inherently very reliable. No further quality is added by the network, for example, to provide a totally reliable service. Because of this, message formats include an error-detection code component. This allows recipients of messages to detect when errors occur, the normal action being to discard the damaged message. The network itself takes no responsibility for any actions required after this. The effect is the same as if the network had lost the message completely, which is another possibility. Duplication of messages should not occur in individual broadcasting networks.

In summary, from a time point of view, there is absolute agreement on the duration of communication time periods, or at least for bounds on the duration of time periods. However, the beginning of time periods is relative to communications starting. Message communications are not segmented over time, and the bits comprising a message are communicated synchronously. There is neither time-related error correction provided by the network nor provision for time-related flow control by the communicating computers. However, note that there is time-related flow control for multiplexing channels. This is discussed in detail in Section 6.3.

6.1.3 Space basics

As has already been discussed earlier in this chapter, there are often some absolute physical constraints on the area occupied by a broadcasting network. These are necessary to achieve the required communication speeds using the available technology. The two dominant types are Local Area Networks (LANs) and Metropolitan Area Networks (MANs) (the latter sometimes called Campus Area Networks (CANs) in academic environments). The absolute space constraints do not extend as far as fixing the physical position of the computers in the network, but may impose restrictions on the minimum or maximum distance allowed between computers. These restrictions are related to the physical transmission properties of channels.

The identifier scheme for broadcasting networks is usually flat. The most common such networks are under the control of a single organization, and so it would be possible to allocate unique identifiers relatively easily. However, most networking equipment is bought with an identifier already built in. IEEE standard broadcasting LANs are an example of a globally unique flat identifier scheme across all such networks in existence. Strictly speaking, this is unnecessary so long as the networks are not connected together. Such interconnection will be considered in Chapter 8.

The network map presented to users is a simple absolute one. All communications involve *all computers* in the network. More specifically, any individual computer can transmit information to all of the other computers. If computers require information about the particular computers in the network, then this must be configured absolutely, or must be acquired using the broadcasting service provided. When such information has been obtained, computers can introduce selectivity into communications by appropriately setting up the destination identifier component of messages.

Overall, a broadcasting network offers the ability to conduct n different types of communications, where n is the number of computers in the network. Each of these types involves the same collection of computers: all of the computers in the network. The difference is in the information flows allowed by the channels. Each of the types corresponds to a different computer broadcasting information to all of the other computers. No other types of communication are possible. If required, they must be implemented using the basic broadcast channels of the network.

The implementation of the broadcasting network service consists of two steps. First, a particular type of channel has to be implemented: a multipeer channel of the type described on page 162 in Chapter 5. This channel allows all of the n possible broadcast information flows. For networks based on a physical multipeer channel, the name **medium** is very often used for this channel, reflecting its physical nature.

Second, the n different broadcast channels (i.e., one for each of the n possible broadcast information flows) must be multiplexed onto the single multipeer channel. The term **Medium Access Control** (MAC) is often used to describe this process when a physical multipeer channel is used. The multiplexing is trivial if the multipeer channel allows several broadcasts to take place simultaneously, but this is uncommon.

The next two sections cover these two implementation components, and are followed by examples of real broadcasting networks, to illustrate how all of the ideas fit together in practice.

6.2 MULTIPEER CHANNEL IMPLEMENTATION

It is necessary to implement a channel that allows broadcasting by any one computer to all of the other computers. This means that there must be appropriate physical connectivity between the computers. Such connectivity may be of a guided type or an unguided type, or a mixture of both. The way in which the channel is implemented depends on various matters, including the layout of the computers and the distances between them, and the types of physical communication medium that can be used. Here, three main types of multipeer channel implementation are described. All have the characteristic of a regular structure. That is, they involve physical connectivity between computers that is clearly part of a coordinated aim: to achieve broadcasting. Note that, although couched in

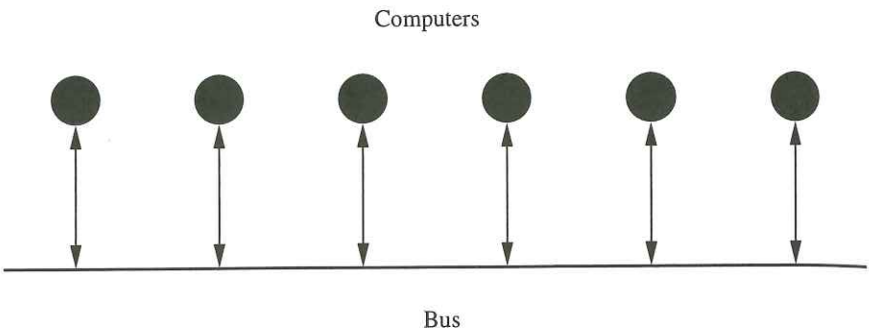


Figure 6.1 Multipeer channel using common medium

terms of physical connections, the ideas here translate to networks founded on higher-level connections with similar properties.

6.2.1 Common medium

A common medium is the simplest implementation approach. It is just a physical medium that is inherently multipeer in nature. That is, any computer is able to transmit on the medium and the transmission is received by all of the computers in the network. When a guided technology, such as an electrical cable, is used, the channel has a linear arrangement and visits each computer in turn. Such a channel is normally referred to as a **bus,** and is illustrated in Figure 6.1.

Guided technologies are only used over relatively short distances, e.g., a few kilometres at most, due to electrical limitations on the broadcasting of signals. For longer distances or, increasingly, for shorter distances also, unguided technologies, such as radio, can be used. In a recycling of a term that had become old-fashioned in a different context, broadcast networks with unguided channels are nowadays termed **wireless**.

In contrast with the other two types of channel implementation described below, the computers have a largely passive role. They are able to listen to the channel to detect broadcasts, and then to store copies of the broadcast information, without affecting the operation of the channel. The only active role is when initiating a new broadcast. No computer is required to play a part in maintaining or ending the broadcast by performing switching duties. This is a natural property of the medium.

6.2.2 Chain or ring

A chain is an approximation to a bus style of multipeer channel. The computers are arranged in a line (in the logical, rather than physical, sense) with a unicast channel between each pair of neighbouring computers in the line. The only difference for

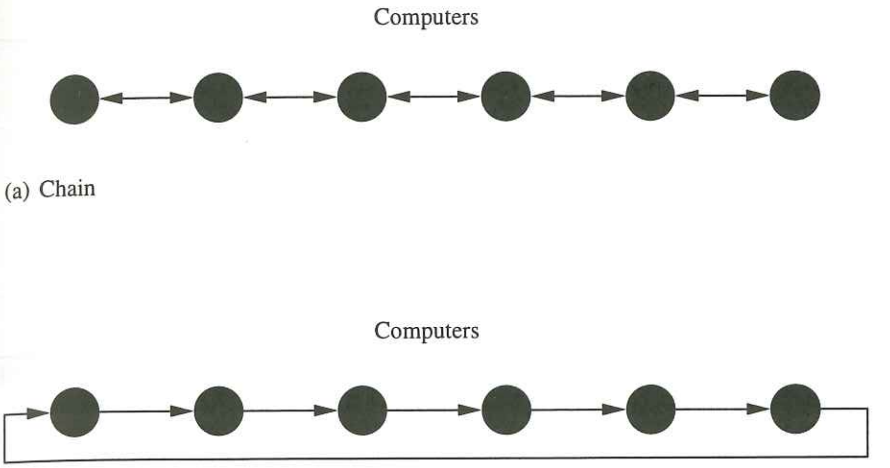


Figure 6.2 Multipeer channels using chain and ring arrangements

a ring is that the two computers at the end of the line are connected by a unicast channel, which turns the line into a circle. It is normal for the unicast channels in a ring to go in one direction only, since bidirectional operation is not required to achieve broadcasting. Figure 6.2 shows a bidirectional chain arrangement, and a unidirectional ring arrangement.

In a chain, when a computer broadcasts information, it must send it to both of its neighbours. The neighbours in turn pass it on to their neighbours. This continues until the information reaches the computers at the ends of the line. The effect is the same as for a bus, except that computers play an active role in propagating the broadcast information by carrying out non-selective switching. As the information passes through the intermediate computers, they can copy it as their own private copy of the broadcast.

In a ring, when a computer broadcasts information, it sends it to one of its neighbours. The neighbour in turn passes it on to its other neighbours. This continues until the information returns to the original broadcaster. Even in bidirectional rings, this simpler approach is preferred to one in which information is passed in both directions round the ring. It avoids some intermediate computer having to detect information arriving from both its neighbours, and then being selective in its switching by not passing the information on any further.

In principle, chain or rings can employ a variety of technologies for the unicast channels between computers. Unlike a genuine broadcast medium, there is no need for uniformity. However, the overall performance of the implementation will be determined by the worst unicast channel in the network. The fact that there is a collection of channels means that the extent of the network is only limited

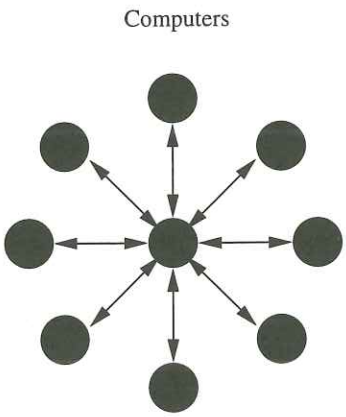


Figure 6.3 Multipeer channel using star arrangement

by a maximum distance between neighbouring pairs of computers, rather than a maximum distance between the two furthest points on the network. Thus, the technology is more suitable for Metropolitan Area Networks than a common medium technology.

The computers in such a network have an active switching role, in forwarding information to neighbours. However, this is normally done by low-level hardware that operates autonomously from the main computer (indeed, it may be separate from the computer), and so forwarding is not affected by loading on the computer or, more extremely, by the computer being switched off.

6.2.3 Star or tree

A star is an approximation to many types of wireless channel. A central computer, or a specialized switching device, acts as the centre of the star, and each computer in the network is connected to it by a duplex channel. A computer can broadcast a message by transmitting it to the centre, which in turn transmits it to all of the computers in the network along the separate channels. Thus, the centre performs non-selective switching to give it a genuine broadcasting capability, whereas the computers perform no switching and only have a single unicasting capability. Figure 6.3 shows a star arrangement.

Clearly, a star network depends completely on the central component. If this fails for any reason, then the network will not operate. The central switch has a reasonable level of complexity, since it must be able to read from a set of n input channels, where n is the number of computers in the network, and be able to write to n output channels, preferably simultaneously. This is in contrast to the modest switching required by an element in a chain or ring, which corresponds to the case $n = 1$ here.

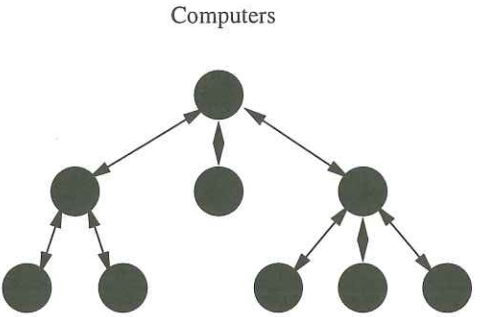
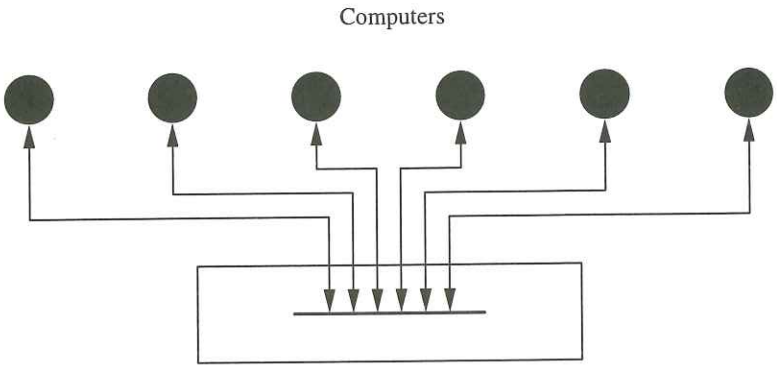


Figure 6.4 Multipeer channel using tree arrangement

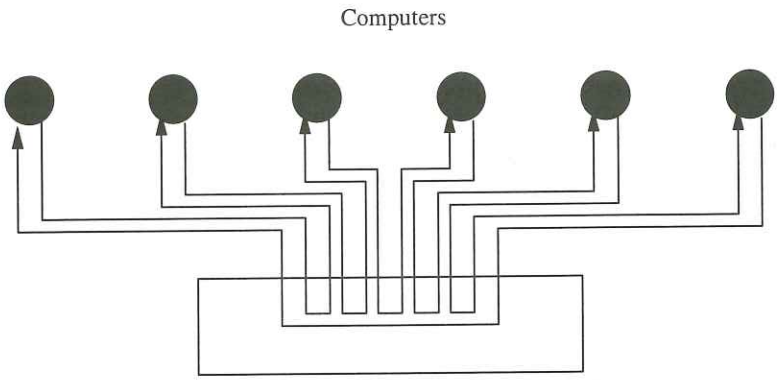
A generalization of a star network is to a tree network, which helps to simplify the complexity of the switching required. Groups of nearby computers are formed into mini-stars, the computers being the tree leaves, and the star centres being the nodes parenting the leaves. In turn, these star centres are formed into star networks, with the new star centres being the nodes at the next level up of the tree. Eventually, the top level of star centres is formed into a single star network, with a centre that is the root of the tree. Figure 6.4 shows a rearrangement of the star network of Figure 6.3 into a three-level tree. The advantage is that each component star network only has a small number of attachments: one duplex channel to the centre of its parent star network, and several duplex channels to the centres of its child star networks. Thus, to implement a broadcast, each star centre is only responsible for relaying the message to a small number of channels. In the event that one star centre fails, partial operation of the network is still possible.

A tree network is also advantageous for directly implementing unicast or multicast communications, when the communicating computers are located in a localized part of the tree. This is because it is not necessary to propagate messages throughout the whole network. However, this involves selective switching by the star centres, and so belongs more properly in Chapter 7.

From a physical point of view, many common medium or chain/ring style networks spanning short distances are organized in a star pattern. The common medium, with an interface for each computer, or the chain/ring, with a forwarding interface for each computer, is contained in a single box. Each computer is then connected to its own interface in the box. Thus, visually at least, the computers are connected to a central switching element. However, it is important to note that the box itself performs no switching function between the connections to the computers. It just houses the computer interfaces. Such a box is called a **hub**. Figure 6.5 illustrates hub arrangements for common medium and ring networks.



(a) Hub containing common medium



(b) Hub containing ring

Figure 6.5 Hub arrangements for common medium and ring networks

6.3 COMMUNICATION MULTIPLEXING

A message broadcasting network supplies a service that allows communications with different broadcast information flows to take place. However, the multipeer channel implementations described in the previous section do not allow arbitrary numbers of broadcast communications to happen at one time. In fact, the norm is for only one broadcast to happen at a time. Some of the multipeer channel implementations are based on independent unicast channels, and so could allow more than one broadcast at a time, because different messages could be simultaneously

transmitted on separate channels. In fact, this is usually avoided in practice, to reduce the complexity of the operational procedures.

Given that only one broadcast can happen at a time, all communications must occur during disjoint time periods, to avoid contention. Furthermore, to ensure fairness of access, most multiplexing schemes include absolute bounds on the length of time period for which one computer is allowed to make use of the multipeer channel.

Regulating access to the multipeer channel is essentially a process of regulating the beginning of communication time periods. Therefore, for the variety of mechanisms that are used in practice, the techniques used are mostly familiar from earlier study of flow control procedures in Section 3.3.2. Here, however, rather than enabling flow control between communicating computers, the mechanisms enable flow control between a transmitting computer and the multipeer channel. In fact, of the mechanisms considered below, only the first one is novel, in that it is not used for end-to-end flow control applications. The novelty comes from the fact that each computer carries out flow control in isolation, guided only by its observations of the channel.

6.3.1 Isolated

Isolated multiplexing of communications involves the individual computers using contention avoidance or detection techniques. These do not guarantee to prevent contention occurring, but seek to minimize it. The most obvious form of contention avoidance is for a computer to listen to the multipeer channel before attempting to try transmitting information. If a communication is already in progress, then it can defer its own communication until the channel is free. However, this does not prevent two or more computers both waiting until the channel is free, and then colliding by transmitting at the same time. One way of reducing the chances of this happening is to introduce a random element into the procedure. Rather than transmitting as soon as the channel is free, computers are obliged to wait for a random short period of time before trying to transmit. This results in a small degree of randomness in the starting time chosen for the repeat communication attempt, which helps in avoiding collisions as soon as the medium becomes free.

The above mechanism allows collisions between information-carrying messages to occur, and this is part of its design. A variation is to use an extra protocol when computers begin transmitting. Rather than just transmitting the message that they wish to send, the computers instead transmit special control messages. These messages have different lengths for the different computers, and so have different communication time periods. The idea is that the computer that transmits for longest wins the competition for the channel, since the others drop out earlier. If the aim is just to eliminate contention, then each computer might have a different length allocated, and so the computer with the higest number always wins. However, if fairness is needed, then computers can use randomly chosen lengths.

This still leaves open the possibility of contention if two computers have chosen the same number.

Although avoidance techniques can reduce the chances of transmission contention occurring, it is not totally prevented. To deal with this, a contention detection method is used. This relies on a transmitting computer being able to detect that its transmission is being interfered with. Detection is done by receiving broadcast information at the same time as transmitting it. If a transmitting computer receives different information to its own, or receives garbled information, then it is alerted to the fact that another computer is transmitting at the same time. Once this abuse of the channel has been detected, the protocol requires that all transmitting computers cease transmitting. This may not be immediate, since there is a requirement for the computers to prolong the broadcast damage so that all computers in the network are made aware that a collision has occurred, and so can throw away any message prefixes already received successfully.

After a collision has been detected, the computers involved have acquired the information that they are in competition for the multipeer channel. Therefore, it is essential to employ a collision avoidance technique, otherwise the problem will be repeated at each future attempt. The standard technique involves waiting for a random period before trying again. Thus, the starting time of a repeated communication attempt is chosen with a random element. Note that this use of randomness is *necessary* after a collision, rather than just desirable as it is before first trying to transmit. This approach is known as the ALOHA method, in honour of its use by the pioneering University of Hawaii radio-based network of 1971, which was the ancestor of all message broadcasting networks. The introduction of randomness into the starting time chosen for communications means that no global control mechanism is needed.

The choice of possible random period lengths is important. If the choice of periods is too small, then further contention is likely. If the choice is too large, then communications might be delayed for unnecessarily long periods of time. There is an algorithm, known as **randomized exponential backoff**, for choosing appropriate random times before trying again. On the i-th occasion that a computer suffers a collision for a particular communication, it chooses a delay in the range

$$0, \ldots, 2^{i-1}$$

'time slots'. Thus, the first time, it chooses either zero or one, the second time, it chooses zero, one, two or three, etc. This continues up to some limit on i ($i = 10$ being a standard value), at which point the range is not increased further, to avoid delays becoming outrageously large.

A problem arises with this scheme if computers are able to transmit messages very quickly in succession. Once one computer has got access to the channel, it may be able to keep it while transmitting many messages one after another. After each of the messages, any competitor will collide with the computer's next message transmission. However, the computer that has just transmitted will choose a backoff period of either zero or one, but its competitors will be choosing a

period from an increasingly wide range. Thus, the chances are that the computer in possession will keep its access to the channel. This phenomenon is called **capture**. One solution to this is called the **Binary Logarithmic Arbitration Method** (BLAM), and involves all computers (not just colliding computers) maintaining the same backoff period range.

The length of a time slot is the maximum time that it takes for a transmitting computer to be sure that it has not suffered from a collision. Given this, and assuming that only one competing computer picks the lowest random number of slots, this computer is guaranteed not to collide with its rivals. The appropriate slot length is equal to the time for a message to propagate to the furthest point on the network, plus the time for any evidence of a collision to propagate back (plus a small safety margin). For a physical multipeer channel, this can be computed as an absolute time period, given knowledge of the distance spanned by the channel, and data transmission rates along it.

Contention avoidance or detection techniques are only really practicable for networks on which there is a fairly low loading and, further, that the load imposed by the computers is fairly randomly distributed over time. If this is not the case, then a lot of time and network capacity will be wasted by collisions, and one of the techniques described below will be more suitable. However, where loading circumstances are favourable, this type of technique has various points in its favour. First, it is completely decentralized, with computers making decisions locally, guided by the raw information that is being received from the channel. Second, it involves a simple protocol, which is always an advantage. Third, it allows computers to transmit information immediately (as long as no contention occurs), rather than be delayed by any arrangements needed under a flow control agreement.

6.3.2 Permission-based

Permission-based multiplexing depends on computers requiring some form of permission before being allowed to broadcast. Only one computer is allowed to have permission at a time, in order to guarantee that multiplexed communications have disjoint time periods. One way to achieve this is to have a master computer that polls the other computers in turn. However, this has the disadvantages of any centralized approach: the operation of the network is only as good as the operation of the master computer.

A less domineering alternative is to use ticket-based flow control, as described on page 101 of Chapter 3, where tickets are issued by the receiver. However, here, the central differences are:

- there is only one ticket in circulation;
- the ticket is issued on behalf of the multipeer channel rather than by the receiver of communicated information; and
- the ticket is immediately passed on to another computer by any computer that receives the ticket and has no immediate use for it.

When a computer receives the ticket, it is allowed to broadcast a message if it needs to and, after it has done so, it must pass the ticket to another computer in the network.

When such a network is first activated, then one computer must be responsible for generating the ticket. This computer might be fixed, under some absolute agreement for the network, or might be chosen dynamically by some form of appointment, competition or election protocol among the computers on the network. Thereafter, there must be a rule to ensure that the ticket is continuously circulated among all of the computers in the network, to give them an opportunity to transmit. For a network which seldom changes, and which contains computers that seldom fail, there may be a fixed ordering of the computers. However, in general, it is necessary to have a protocol that allows computers to be added or removed from the circulation list dynamically. Communication of the ticket, which is just a special form of message, must be by unicast communication between the current ticket holder and the next ticket recipient using the basic broadcast capability of the multipeer channel.

Advantages of a permission-based technique are that contention for the medium is eliminated and that the rights of computers to transmit can be regulated by the chosen ticket passing order. For example, if the ticket is passed around each computer in turn, then transmission rights are fairly shared out; this is a reasonable arrangement for a network which is highly loaded. As an alternative example, if it was arranged that one computer received the ticket every second time that it is passed on, then the effect would be that this computer was given priority for its transmissions.

Disadvantages include the fact that, before transmitting, a computer must wait for the ticket to arrive, although none of the other computers may actually need to transmit at that moment. A further disadvantage is the complexity of the arrangements for the ticket itself. As well as dealing with adding and removing computers, the ticket protocol must deal with computers that fail while holding the ticket. It must also deal with computers that, through error, create multiple tickets. However, these are not insurmountable problems, and such protocols exist and are used in practice, as will be seen in Section 6.4.

6.3.3 Reservation-based

Reservation-based multiplexing is an extension of permission-based multiplexing that is designed to reduce the time that a computer has to wait before being allowed to broadcast a message. A problem with ticket-style schemes, as noted above, is that a computer must wait its turn for the ticket, even if other computers before it have no need of the ticket. A reservation-based scheme is based on the style of flow control where a transmitter requests a ticket from the receiver when it needs to, or is about to need to, transmit.

In this case, any such protocol is complicated by the fact that receivers are not responsible for issuing tickets. Rather, tickets are issued on behalf of the multipeer

channel. One possibility is to designate a particular computer as a master, responsible for issuing tickets, and for the other computers to make reservations with this master. However, this makes the network completely reliant on the master computer. Ticket-passing, just described above, is an example where the responsibility for the control mechanism is distributed among all of the computers in the network. Ideally, a reservation-based scheme should similarly distribute responsibility for noting reservations and issuing tickets. The basic complication is that, when a reservation is made, all computers that might conflict with the reservation must be notified of it. Thus, some kind of broadcasting of each reservation is required. This is rather different from the unicasting of tickets.

A fatal flaw in this appears to be the fact that broadcasting of reservation messages should be allowed whenever necessary — which introduces exactly the same problem that reservation is meant to be solving! One solution is to generalize what happens with ticket-passing, by allocating short time slots on the multipeer channel that allow computers to make reservations, between the main items of broadcasting business. This solution anticipates the technique of time division multiplexing, discussed in the next section. Another solution is to employ piggy-backing as a multiplexing technique that combines reservation messages with existing broadcast messages. Then, computers receive notification of future reservations along with communications that are happening at the present.

On receiving a reservation message, it is essential that a computer takes it into account when formulating its own transmission plans. In order for a reservation scheme to be completely fair, each computer should maintain a count of the number of pending reservations. When it wishes to transmit, it should broadcast a reservation message and also note the value of this count. Then it should allow that number of other computers to transmit first. After this, assuming that all the other computers have received the same reservation messages, the computer can proceed with its transmission without fear of collision. Moreover, transmissions proceed in the same order that reservations were made. This mechanism assumes that all computers are allowed to have reservations pending. In cruder schemes, only one computer might be allowed to have a reservation pending, in which case the use of counters is redundant.

6.3.4 Physical division

The above schemes for multiplexing communications, so that they do not overlap in time, are very much in the asynchronous spirit. That is, the starting times of communications are adjusted as appropriate, under the control of the multiplexing policy. An alternative is to use a method of sharing the medium that is directly related to its physical properties, rather than to the communications required of it.

A very synchronous approach is **time division multiplexing**, already mentioned in Section 5.3.2. With this, the communication time available is divided into fixed-size time slots, and then these slots are allocated to the computers on the network. There will be periodic cycles, during which each computer has at

least one time slot available for its exclusive use. Time division multiplexing removes the need for distributed mechanisms to decide which computer is allowed to transmit next. It also allows guarantees to be made on the absolute time at which communications will takes place. For applications that require real-time information transmission, such guarantees are important.

The main disadvantage of time division multiplexing is its inflexibility. If the communication needs of the computers in the network are fixed and continuous, the scheme works well. However, in general, the scheme is wasteful if the communication needs of the computers vary over time. Time slots will be unused by some computers, while other computers will have to wait. Of course, if computers are frequently added to or removed from the network, there must be some scheme for adjusting the allocation of time slots. However, this would normally be under human control, rather than happening automatically during normal operation.

An improvement is to blend time division multiplexing with the other types of multiplexing. Some time slots are permanently reserved for things like computers with continuous needs, ticket-passing or reservation passing. The remaining slots are allocated using a dynamic multiplexing method.

There are other methods of multiplexing a physical broadcast channel, but these lie beneath the bit frontier, being concerned with how bits are transported over the physical medium. The main method is **frequency division multiplexing**, which can be used where bits are transmitted by the modulation of a waveform. Separate logical channels can be provided over the physical channel by allowing each to transmit within a different frequency range. This allows multiple communications to take place during the same time period. However, this benefit is not gained for free, since the communication rates possible depend on the available bandwidth. Thus, the raw data rate of the physical channel is divided up between the logical channels with, if anything, a reduction in the aggregate data rate.

Sometimes, the allocation of frequency bands to communications varies very dynamically. For example, in radio communication, frequency hopping — changing the frequency range used — may occur as often as for every bit transmitted, in an attempt to avoid interference from radio sources broadcasting on particular frequencies. Embroidering this further, **code division multiplexing** is a multiplexing method that relies on different computers having different pseudo-random frequency hopping sequences, and this allows sharing of the medium without fixed frequency ranges having to be allocated to the computers in the network. There is more discussion of this technique in Section 6.5.

6.4 EXAMPLES OF GUIDED TECHNOLOGY NETWORK IMPLEMENTATIONS

The six examples in this section cover the main message broadcasting networking technologies that are in use for LANs and MANs based on guided media. Two of the six — token ring and FDDI — have roughly similar operational principles,

but there are some distinctive differences. These two, and each of the other four, illustrate a different combination of multipeer channel implementation and communication multiplexing technique. Each combination has been designed to suit the practical environment in which the network is expected to operate: the physical medium used and the quality of broadcasting service offered.

6.4.1 Ethernet

Ethernet is a message broadcasting mechanism based on a guided common medium, with contention avoidance and detection used for multiplexing. Here, the term **ethernet** is used as a generic term for such networks. However, the name 'Ethernet' (with a capital 'E') has historic significance, being the name chosen for the pioneering network of this type, first developed by Xerox in the early 1970s. Therefore, strictly speaking, the name only refers to one particular product, although it has come to be used as a generic name for the technology.

A more accurate name is the less catchy **CSMA/CD,** standing for Carrier Sense Multiple Access with Collision Detection. The 'Carrier Sense' part refers to computers checking whether the channel is busy before trying to transmit. The 'Collision Detection' part refers to computers noticing when their transmissions collide with those of other computers, and then withdrawing and retrying gracefully. A further name used for ethernet, encapsulating the type of multiplexing and the type of channel, is the **contention bus**. The major influence on ethernet technology was the pioneering work by Xerox, joined later by DEC and Intel. This led to standardization by the Institute of Electrical and Electronics Engineers (IEEE), specifically the IEEE 802.3 standard for CSMA/CD networks. The following description refers to IEEE 802.3, but the general principles apply to any ethernet style network.

Information

The messages transmitted on an IEEE 802.3 ethernet have the format illustrated in Figure 6.6. The seven preamble bytes are present to allow a recipient to synchronize bit timing with the transmitter, and then the start byte signifies the beginning of the message contents. The destination and source identifiers are standard IEEE 802 identifiers, as described on page 143 in Chapter 5. After them, the two-byte length field gives the length in bytes of the information carried by the message. This cannot just be deduced from the overall length of the message minus the length of the fixed control fields, because the message may be padded to ensure that it is larger than the minimum length allowed for messages. The minimum length restriction is to allow collision detection, and will be discussed a little later. Excluding the preamble and start bytes, the minimum length of a message is 64 bytes, and there is a maximum length of 1518 bytes. After the information bytes and any padding, there is a 32-bit CRC using the CRC-32 generator to allow error detection by a recipient.

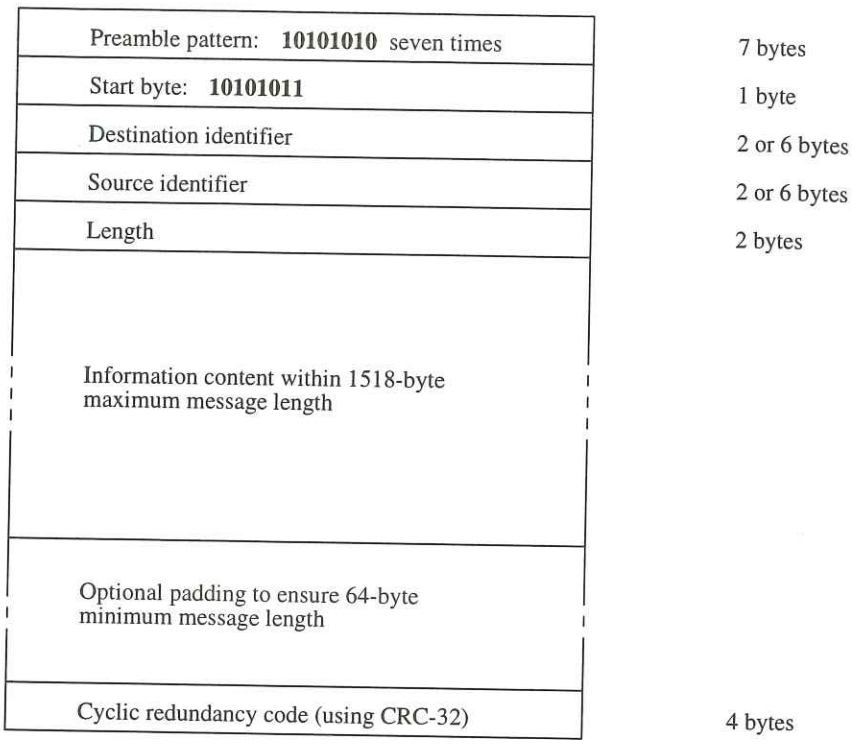


Figure 6.6 Format of IEEE 802.3 ethernet message

Time

The latency of communications and the transmission rate of communications vary between different varieties of 802.3 ethernet. They are determined by the maximum allowed length of the common medium and the bit rate of the medium. These parameters are incorporated in a naming convention for different varieties, for example, the original **10 BASE 5** standard operates at 10 Mbits per second ('10') and has a maximum channel length of 500 metres ('5'). The 'BASE' refers to the fact that baseband transmission is used. The normal rate for ethernets was 10 Mbits per second until the mid-1990s, at which point enhanced standards for a 100 Mbits per second rate began to come into use. The maximum channel length depends on the physical technology used. For example, 10 BASE 5 is for thick (0.5 inch diameter) coaxial cable, whereas **10 BASE 2** is thin (0.25 inch diameter) coaxial cable. The latter is often known as 'cheapernet' since it is cheaper to install than the former. Signals propagate along coaxial cable at approximately 200 metres per microsecond, so the latency of a 10 BASE 2 ethernet would be approximately 1 microsecond. Thus, the maximum communication time period on such a network would be: $10^{-6} + (8 * 1518)/10^7 = 0.0012154$ seconds.

The basic service of ethernets is an unreliable connectionless service. Any communication time period involves the continuous transmission of a message. The bits of each message are communicated synchronously, with Manchester encoding — described on page 87 in Chapter 3 — being used to maintain synchronization.

Space

As the maximum channel lengths quoted above may suggest, ethernets are used as Local Area Networks. Some extension in the distance covered is possible by using devices called **repeaters**. These connect together channels, replaying signals from one to another. For example, with 10 BASE 5, up to three repeaters can be inserted, which allows a distance of up to 2 km to be covered. The identifiers within a LAN are flat, with each computer being programmed with its own identifier. Thus, unicasts or multicasts can be implemented by filtering, each computer only accepting broadcast messages with a destination identifier matching its own. With the IEEE identifer scheme, a multicast has a destination identifier that is the number of a pre-assigned multicast group. There is also a special identifier meaning 'all computers on the network'. Messages broadcast with this identifier as their destination are treated as genuine broadcasts, and are accepted by all of the computers receiving the broadcast message.

Multipeer channel implementation

The ethernet multipeer channel is implemented by a guided common medium that is a bus. As already mentioned, coaxial cable is one possible technology, with the common cable visiting each computer on the network. There are alternative technologies, but these make use of a hub arrangement, rather than having a single medium snaking around the computers. The hub is a central box, with an interface for each computer in the network. The general arrangement of this was shown in Figure 6.5. The computers are connected to their interfaces in the box using a pair of connections: one for receiving and one for transmitting. When any computer transmits to the hub, the hub's function is to echo the transmission on all of the outgoing receiving connections to the computers. Thus, the effect is that the hub behaves like a common medium.

Possible technologies for the receiving and transmitting connections are twisted pair electrical cable or fibre optic cable. The standard for the former is **10 BASE T**, and for the latter is **10 BASE F**. Note the letters 'T' and 'F' replacing the number representing the maximum channel length, which is no longer an issue. The maximum length of the connection between a computer and the hub is 100 metres or 500 metres for the two standards respectively.

There are IEEE 802.3u standards for **fast ethernet**, that is, 100 Mbits per second ethernet. IEEE 802.3u is essentially just a version of 802.3 that has a 10 times faster transmission rate. The possible technologies are all hub-based — there is no version for a genuine physical bus. Therefore, **100 BASE T** is the generic

name for the higher speed follow-up to 10 BASE T. The physical technology used comes in three flavours.

The **100 BASE T4** standard involves using four Grade 3 UTP twisted pair cables (hence the 'T4') to achieve the 100 Mbits per second transmission rate. One cable is used just for receiving, one just for transmitting, and the other two for both receiving and transmitting. Each cable has a transmission rate of 33.33 Mbits per second. This gives the required combined rate, since the 100 Mbits per second channel is split between the slower channels. For higher quality Grade 5 UTP cabling capable of faster transmission rates, the **100 BASE TX** standard can be used. It involves only two cables, as for the 10 Mbits per second standard. One is used for receiving, the other for transmitting, both at the 100 Mbits per second rate. Finally, there is the **100 BASE FX** standard, for fibre optic cabling. The major difference from 100 BASE TX is that the distance between a computer and the hub can be up to 2 kilometres, rather than just 100 metres.

100 BASE T is the natural evolution of the hub-style 10 BASE T ethernet. Section 6.4.6 describes a rather different evolution from the 10 BASE T network, to provide a 100 Mbits per second network. This fully exploits the fact that, although logically a bus, the physical connectivity of the network is actually a star or tree. The first 100 BASE T products appeared in 1994, so it is a relative newcomer. However, its successor was not far behind, since the first **gigabit ethernet** products appeared in 1997. These are covered by IEEE 802.3z standards, and are 1000 BASE T ethernets. Fibre optic cable is the normal type of cabling, but four-cable or eight-cable Grade 5 UTP is also an option. As in the move from 10 BASE T to 100 BASE T, so the move from 100 BASE T to 1000 BASE T is largely concerned with the physical transmission of bits rather than new protocols.

Communication multiplexing

The ethernet multiplexing scheme involves contention avoidance and detection, as described in Section 6.3.1. The original Ethernet was the pioneer of this technique in a cable-based network. The use of a minimum message size is crucial to making the scheme work efficiently. If two computers at opposite ends of the ethernet channel both started transmitting very short messages at the same time, then each would have finished transmitting before hearing the other's message, due to the latency of the network. Some computers in between would detect a collision, at the point where the messages collided, but others would be unaware of the collision until after the message from the nearer computer had been successfully received. This is illustrated in Figure 6.7.

The minimum message size is chosen so that the time to transmit it is more than twice the latency of the network. That is, the message will still be being transmitted after the time it takes for the beginning of the message to traverse the whole bus (during which time other computers may start transmitting), plus the time it takes for any competing transmission to reach the first transmitting computer. The choice of a 64 byte minimum size for messages gives a minimum

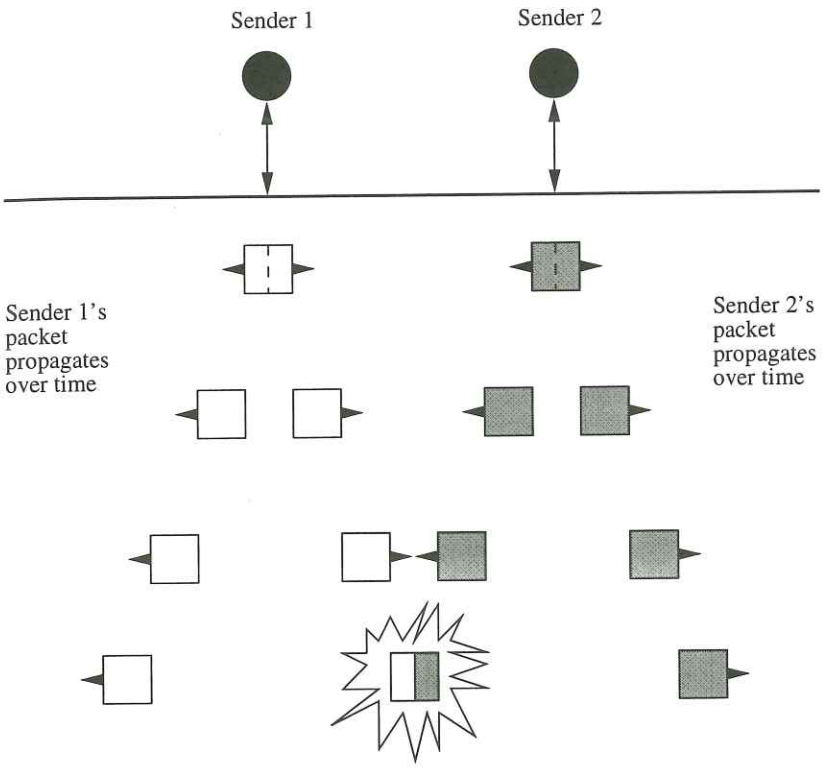


Figure 6.7 Ethernet collision invisible to transmitters

communication time of 51.2 microseconds on a 10 Mbits per second ethernet — enough for collisions to be detected on a coaxial cable up to 5 km long (if this was allowed).

Just as the transmission time must be long enough for collisions to be observed by all of the transmitting computers, it is also necessary to ensure that collisions are noticed by all of the computers in the network. It is possible that, if two transmitting computers notice a collision quickly, they will both stop transmitting before evidence of the collision has propagated through the network. This is illustrated in Figure 6.8. To ensure that the news of the collision travels, both stations transmit a short 32-bit 'jam signal' before going quiet. This is recognized by all computers attached to the bus as a signal to abandon any communications in progress. Then, using the binary exponential backoff algorithm, each competing computer waits for a random period of time before trying again.

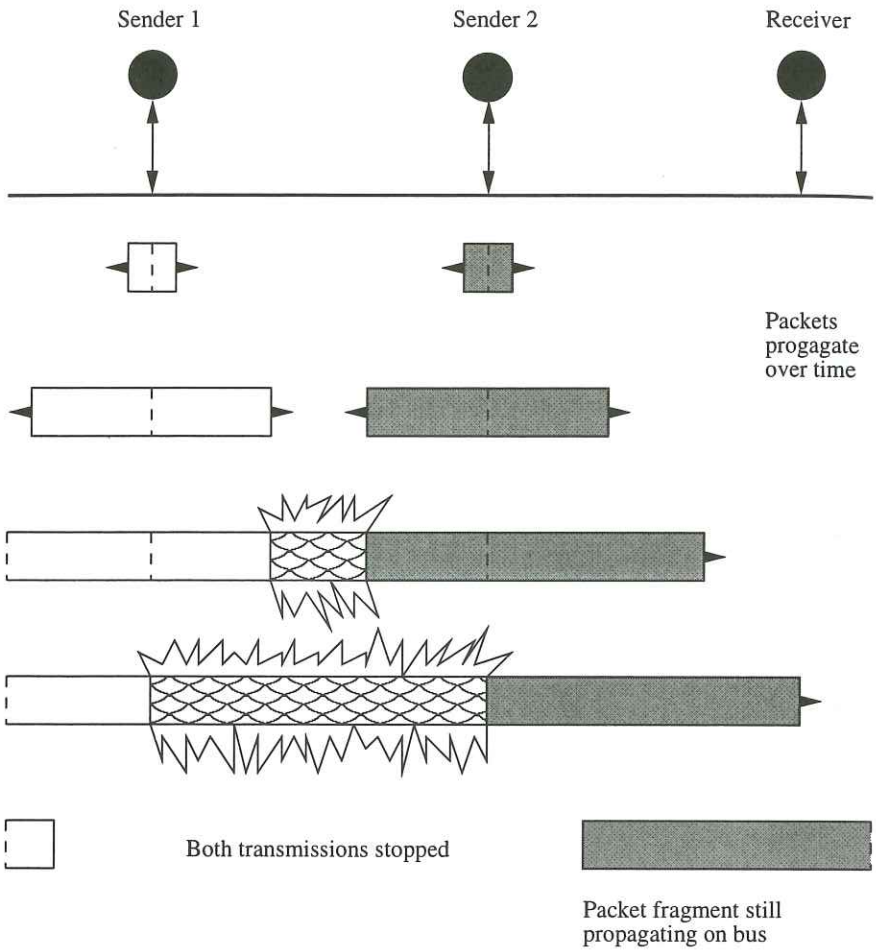


Figure 6.8 Ethernet collision invisible to a receiver

6.4.2 Token ring

Token ring is a message broadcasting mechanism based on a ring of simplex channels, with a ticket used for permission-based multiplexing. The term 'token' is used for the ticket used for multiplexing, hence the name given to this type of network. The first network of this type was developed by IBM in the late 1960s. IEEE developed the 802.5 standard for token rings, based strongly on the system developed by IBM. The following description refers to IEEE 802.5, but the general principles apply to any token ring style network. Note that, in 802.5 networks, a reservation-based mechanism is also used to elaborate the basic style of permission-based multiplexing.

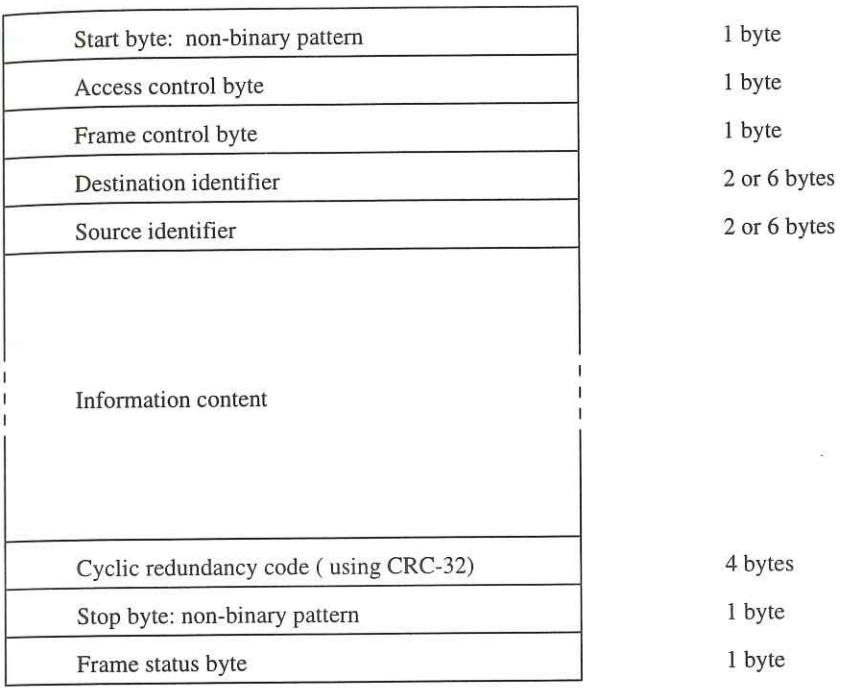


Figure 6.9 Format of IEEE 802.5 token ring message

Information

The messages transmitted on an IEEE 802.5 token ring have the form illustrated in Figure 6.9. The start byte signifies the beginning of the message contents. It contains four bits that do not have valid **0** or **1** values — this is possible because bits are Manchester encoded, and these special 'bits' have signal values that are not valid bit codings. Next, the access control byte contains various fields concerned with the multiplexing mechanism, as discussed below. The frame control byte differentiates between messages that are carrying information and messages that are for ring control purposes. The destination and source identifiers are standard IEEE 802 identifiers, as described earlier.

After the information bytes, there is a 32-bit CRC-32 cyclic redundancy code to allow error detection by a recipient. There is no minimum size restriction on messages. In principle, there need be no maximum size restriction but, in practice, to ensure fairness in sharing the channel, a maximum of 5000 information bytes is typical. After the CRC, there is a stop byte which, like the start byte, contains four special 'invalid bits', as well as two other bits that are discussed below. Finally, there is a frame status byte, which contains acknowledgement information, also discussed below.

Time

The latency of a token ring network depends on a number of things:

- the number of computers in the network;
- the speed at which bits are forwarded by each computer in the ring;
- the distances between computers; and
- the technology used for each unicast channel between computers.

Since it must be possible for a message representing the token to circulate round the rings, with all its bits present at one time, it is important to ensure that the latency of a ring is large enough to accommodate this.

The standard bit rates for an 802.5 token ring are 4 or 16 Mbits per second. For a 4 Mbits per second ring, given that the token message size is 24 bits (only a start byte, an access control byte and a stop byte), a token will be transmitted in 6 microseconds. If the latency is less than this, the token could not circulate independently when the ring is idle. For example, for a typical bit propagation speed of 200 metres per microsecond, the physical latency of a 1 kilometre ring would be only 5 microseconds — too short. To deal with this, it is necessary to introduce artificial delays in the ring interfaces to increase the latency sufficiently. This is only a problem for the special case of the token, since no other messages are allowed to circulate round the ring independently, without intervention by a computer.

The basic service offered is an acknowledged connectionless service. Messages are communicated in a time period with continuous synchronous communication of bits which, as remarked above, are represented using Manchester encoding to allow synchronization. It is possible for an asynchronous sequence of messages broadcast by a computer to be parts of one communication that has been segmented. The stop byte of each message contains an Intermediate (I) bit, which is set in messages that are intermediate in a sequence; the bit is cleared in the final message of a sequence.

Space

IEEE 802.5 token rings are intended as Local Area Networks. The next example (FDDI), in Section 6.4.3, illustrates the techniques extended to Metropolitan Area Network distances. The identifer scheme within a token ring LAN is flat, with each computer being programmed with its own identifier. The message broadcasting service can be used to implement unicasts and multicasts in the same way as for ethernet: computers filter out any messages with destination identifiers that do not match their own.

Multipeer channel implementation

The multipeer channel is implemented using the simplex channels between computers. A transmitting computer sends its message round the ring, in a series of hops between the computers. This allows an acknowledgement service to be provided very easily. The frame status byte at the end of each broadcast message

contains an acknowledge (A) bit and a copied (C) bit, which are both zero when the message sets out. In fact, there are two copies of these bits in the frame status byte, to increase reliability given that this byte is not covered by the CRC.

Any computer that recognizes a match of the destination identifier with its own identifier sets the A bits before passing the message on; further, if it makes a copy of the message, it also sets the C bits. Thus, when the message arrives back at its transmitter, this computer can tell (a) whether any computer recognized its identifier, and (b) whether any computer accepted the message. As a further acknowledgement mechanism, the stop byte in each message contains an error (E) bit, which is zero when the message sets out. Any computer that detects an error as the message circulates can set the E bit, and this allows the transmitter to tell whether any errors occurred during the broadcasting of the message.

As for ethernet, the actual ring connections may be located in a single box, which acts as a hub. Computers are then connected to interfaces in the hub by transmitting and receiving connections. Such an arrangement has the added benefit of protecting the ring against computer failures, since the hub can arrange automatic bypassing of any interfaces whose attached computers fail or are switched off.

Because a broadcasting computer is responsible for removing its messages from the ring after circulation, there is the possibility of orphan messages being doomed to roam the ring if their parent computer fails in some way during their progress round the ring. To deal with this problem, and various other administrative problems, the IEEE 802.5 standard includes the fact that one computer is designated as a ring master, responsible for such things as detecting orphan messages. For this purpose, the access control byte of each message contains a monitor bit, initially zero but set by the master computer when the message passes through. If a message arrives at the master computer with this bit already set, then it must be an orphan and is removed from the ring by the master. Because of this (unfortunate) reliance on a master computer, the ring standard includes further protocols for checking that the master computer is still working and for appointing a successor if not. These are implemented using special control messages distinguished by particular values in the frame control byte of the message.

Communication multiplexing

The basic token ring multiplexing scheme is permission-based, and a computer must receive a message representing the token on the ring before being allowed to broadcast a message. To send its first message, the token message is converted into an information-carrying message by changing one bit in the access control byte as it passes through. After sending its message, and receiving it back again when it has passed round the ring, the computer either sends another message or sends a message representing the token to its neighbour in the ring.

As described, this allows fair sharing of the multipeer channel, as long as computers are not allowed to hold the token for an excessive period of time. The

default limit on the token holding time is 10 milliseconds. The master computer is responsible for ensuring that the token keeps in circulation. If a computer fails while holding the token, then the token disappears. Given the maximum token holding time, and knowing the number of computers in the ring, the master computer is able to determine a maximum period of time that can pass before it should receive the token for its turn. If this time is exceeded, then it generates a replacement token.

As well as the basic permission-based mechanism, there is also a fairly elaborate reservation-based mechanism to allow computers to have different priorities of access to the ring. The access control byte of a token message contains a three-bit field giving the priority (0 to 7) of the token. When a computer wants to broadcast a priority p message, it must wait for a token with a priority value that is less than or equal to p. Any other tokens must be allowed to pass by, destined for computers with higher priority needs. The access control byte of all messages contains a three-bit reservation field, which allows computers to make claims for future use of the token.

The reservation field is set to zero when a message begins travelling round the ring. If a computer is waiting to transmit a message of priority p, and a message passes through with a reservation field value less than p, then it can set the reservation field to p before passing on the message. The effect is that, when the message arrives back at its original sender, this computer knows what priority is required for the token it reintroduces to the ring. Note that, after a high priority token has been used, this rule may have the effect of reducing the priority again to a lower value.

6.4.3 FDDI and FDDI-II

The Fibre Distributed Data Interface (FDDI) is a message broadcasting system based on the token ring. It makes use of optical fibre channels operating at 100 Mbits per second. Because of the higher data rate, there are several differences from the token ring mechanisms described in the previous section. Also, FDDI allows synchronous communications to take place at regular intervals, which involves a mechanism additional to the token-based mechanism for multiplexing asynchronous communications. This mechanism does not guarantee that synchronous communications will take place at precisely regular times, and a variant known as FDDI-II has been developed to allow this. The description of FDDI below focuses only on the ways in which FDDI and FDDI-II differ from the IEEE 802.5 standard for token rings.

Information

The format of FDDI messages is shown in Figure 6.10. Although the format can be largely understood in terms of bytes, the introduction of **nibbles** is necessary for a full understanding. A nibble is a four-bit unit, that is, half of a byte. The

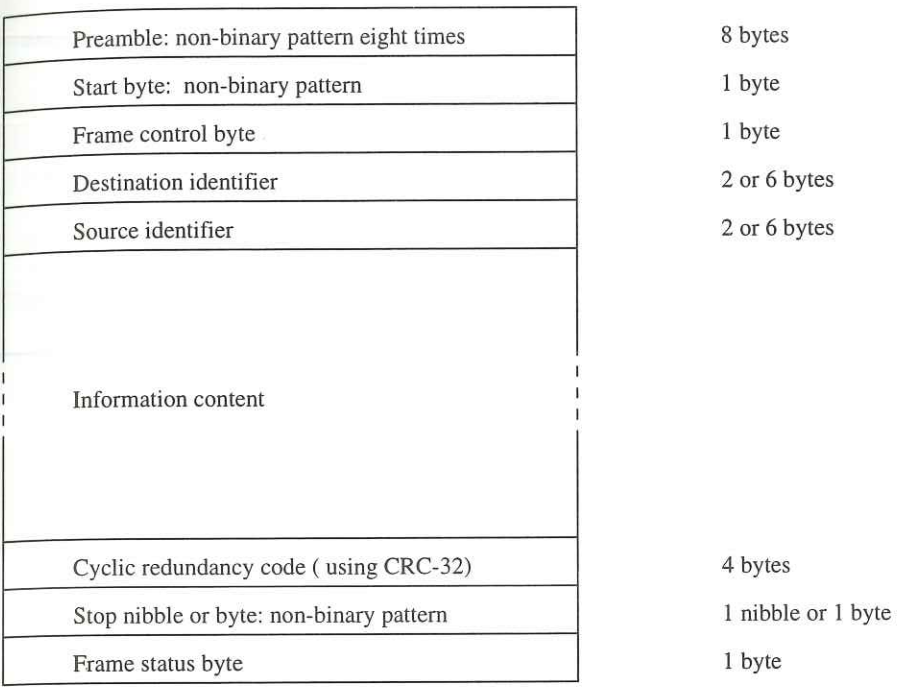


Figure 6.10 Format of FDDI message

use of nibbles is a significant difference for FDDI: for transmission, each four-bit nibble of a message is encoded as length five on-off pattern of light signals. This is different from Manchester encoding, but the 'five bit' representation of each nibble ensures that there are always transitions between on and off present, to assist synchronization.

The FDDI message has an eight-byte preamble to allow synchronization between transmitter and receiver at the higher transmission speed. There is no need for a token ring-style access control byte, since FDDI does not use priority and reservation fields, as will be seen below. The end delimiter is a stop nibble on normal messages, and is a stop byte on token messages. The maximum allowed message size is 4500 bytes.

Time

The latency of an FDDI network depends on the same type of issues as the latency of a token ring network. A significant factor is that the distances allowed between computers in the ring are significantly larger, ten- or twenty-fold. Thus, the latency is higher. For example, in a maximum size 100 km ring with 500 computers attached, the latency would be around 1 millisecond. The greatly improved time factor is that the data rate is 100 Mbits per second, which is 25 times higher than

the basic token ring rate. The basic service is an acknowledged connectionless service, like that of a token ring.

Space

FDDI is intended for use in Metropolitan Area Networks, since the fibre optic technology allows transmission over longer distances. Unlike IEEE 802 LANs, where either 2-byte or 6-byte computer identifiers may be used, but not both at once, FDDI allows a mixture. The frame control byte of each message contains a bit that indicates which identifier scheme is used in that message. This can be viewed as being a simple two-level hierarchical identifier scheme.

Multipeer channel implementation

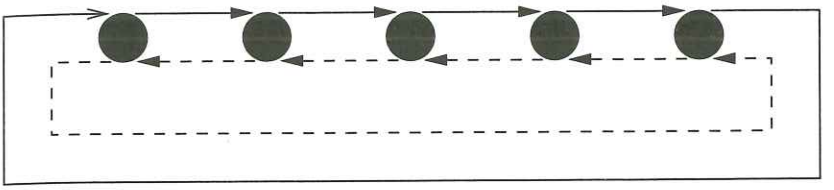
FDDI uses a ring formed by simplex channels between computers. The standard provides for dual rings, travelling in opposite directions. In normal usage, only one ring — the primary ring — is used, and in the same way as for the token ring. The other ring — the secondary ring — is not used. However, if one of the channels forming the primary ring fails, or one of the computers fails, then the two computers on either side of the failure point can loop back their incoming primary ring channels to their outgoing secondary ring channels, thereby forming an emergency ring with a sequence of secondary ring channels replacing the missing one or two primary ring channels. This is illustrated in Figure 6.11. It is allowable to have computers that have only single ring attachments. However, to ensure reliability, such computers must be connected to a dual ring via hubs that are able to isolate the computers in the event of their failure.

As with the token ring, the way of implementing broadcasts allows an acknowledgement service to be provided easily. Messages carry A, C and E bits, which can be set by computers around the ring. The only difference from the token ring standard is that the E bit is located alongside the A and C bits in the frame status byte at the end of the message.

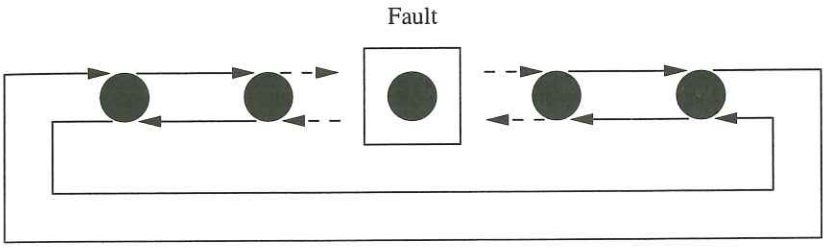
A significant difference from the token ring standard is that there is no master computer responsible for the token-passing mechanism. This, like the presence of a dual ring in FDDI, gives greater reliability against failures. In general, the monitoring responsibility is distributed among the computers in the ring. This is intimately connected with the way that token-passing works, but a side-effect is that messages which are not removed from the ring by their transmitter can either be discreetly removed by one of the other computers, or are removed through re-initialization of the ring following token loss.

Communication multiplexing

Because the latency of an FDDI ring is significantly longer than that of a token ring, there is a fundamental change to the token-passing mechanism. Rather than re-introducing the token when broadcasting of a message is complete, i.e., the message has travelled round the ring, the transmitting computer sends a token



(a) Normal operation



(b) Operation to avoid fault

Figure 6.11 Single FDDI ring formed from dual ring channels

message to its neighbour as soon as it has finished sending its message(s) to its neighbour. Thus, potential transmitters are not delayed by a time corresponding to the latency of the network. In particular, more than one message may be circulating the ring at one time, each message corresponding to a different broadcast message. This relaxation of the multiplexing conditions is possible in chain or ring-style networks because the multipeer channel is based on a collection of independent unicast channels, and so there is scope for multiple simultaneous broadcasting without contention occurring.

The rules that a transmitting computer must follow when holding the token are more complex, since provision is made for synchronous communications to be guaranteed reasonably accurate time periods. Each computer has a time allocation (possibly zero) that it is allowed to use for the transmission of synchronous communication each time it receives the token. Then, the rules for allowing other, asynchronous, communications are set up to ensure that synchronous communications can occur on a regular basis.

An important parameter is the target token rotation time (TTRT), which measures the maximum time that a computer should have to wait between occasions of receiving the token. All of the computers have the same value for the TTRT,

which is chosen from the range 4 to 165 milliseconds. The setting of the synchronous time allocations, and the setting of the TTRT value, is carried out by a network management computer. Computers in the network negotiate with it, and are informed by it, using a special protocol.

It can be shown that no computer ever has to wait more than two times the TTRT to receive the token. If this happens, the computer initiates steps to introduce a new token, assuming that the old token has been lost. In fact, it can be shown that the actual time for the token to circulate tends towards TTRT in normal operation. On receiving the token, a computer is allowed to use up to its time allocation for synchronous transmission. Then, if the token arrived before TTRT time units had elapsed since the computer's last use of the token, the computer can use the time difference for asynchronous transmission.

The rules for asynchronous transmission are further elaborated to allow different priorities. This is also based on the use of absolute time measurement, rather than the reservation-based mechanism used on the token ring. There are eight different priorities, each with a threshold value that measures a length of time. When it holds the token, a computer tries to transmit its messages in decreasing order of priority. The computer is only allowed to transmit a message if the remaining time allowed to it for asynchronous transmission is greater than the threshold value for the message's priority level. Thus, a computer may have to pass on the token early, rather than transmit lower priority messages. The effect of this is that higher priority messages can usually be sent in an unconstrained manner, whereas lower priority messages are sent in a fair manner when network capacity is available.

Although the basic mechanism gives a reasonable guarantee for the delay suffered by synchronous traffic (no more than TTRT), this is still not good enough for **isochronous** communications. These are required to convey traffic which is generated at regular time intervals and must be delivered at the same constant rate. The FDDI-II standard supports the above FDDI service as its **basic mode**. However, if all computers attached to the network have been upgraded appropriately, then FDDI-II also supports **hybrid mode**. In this mode, time division multiplexing is used to divide the multipeer channel among a collection of slower channels. Each of these can be used for either isochronous transmission or for normal FDDI-style transmissions.

The time division multiplexing is carried out by a master computer, which repetitively generates bit sequences called **cycles**, one every 125 microseconds. Note that this rate corresponds to the normal sampling frequency for human speech, which is no coincidence, since real-time speech communication is a significant application requiring isochronous communication. Given the 100 Mbits per second transfer rate, this frequency means that each cycle is a string of 12 500 bits. The format of a cycle is shown in Figure 6.12.

Within each cycle, after 29 nibbles of preamble and header field, there are 12 groups that carry information. The first byte in each group is used for a normal FDDI-style channel, and so a total of 12 bytes are carried for this channel in each cycle. The rest of the group is used for 16 different isochronous channels, with

Preamble: non-binary pattern five times	5 nibbles
Cycle header	24 nibbles
Group 0	129 bytes
Group 1	129 bytes
	8 * 129 bytes
Group 10	129 bytes
Group 11	129 bytes

(a) Overall cycle format

Dedicated packet byte	1 byte
Wideband channel 0	8 bytes
Wideband channel 1	8 bytes
	12 * 8 bytes
Wideband channel 14	8 bytes
Wideband channel 15	8 bytes

(b) Group format within cycle

Figure 6.12 Format of FDDI-II cycle

eight bytes carried for each. Thus, a total of 96 bytes is carried for each channel in each cycle. This arrangement allows the FDDI-style channel to carry 768 kbits per second, and each isochronous channel to carry 6144 kbits per second. The capacity of each isochronous channel can be further time division multiplexed to carry many slower channels, for example, one channel can carry 96 ISDN channels running at 64 kbits per second each.

The division between non-isochronous and isochronous channel use is flexible, because the bytes allocated to any isochronous channels that are not required can be reassigned to the FDDI-style channel, thus increasing its capacity in steps of 6144 kbits per second, to a maximum of 99.072 Mbits per second. Part of the header field of each cycle indicates which, if any, such reassignments apply to the cycle.

Preamble: **10101010** one or more times	1 or more bytes
Start byte: non-binary pattern	1 byte
Frame control byte	1 byte
Destination identifier	2 or 6 bytes
Source identifier	2 or 6 bytes
Information content	
Cyclic redundancy code (using CRC-32)	4 bytes
Stop byte: non-binary pattern	1 byte

Figure 6.13 Format of IEEE 802.4 token bus message

6.4.4 Token bus

Token bus is a message broadcasting system based on a common guided medium, with a ticket used for permission-based multiplexing. Thus, it merges the channel technology of ethernet with the multiplexing mechanism of token ring. As with the token ring, the word 'token' is used to refer to the ticket. Token bus networks are designed to be useful in situations where bus-style cabling between computers is already available, and guarantees are needed on the maximum delays suffered by any computer wishing to broadcast. This type of network was first developed as part of the Manufacturing Automation Protocol (MAP) project, which is discussed in detail in Chapter 10. The work evolved into the IEEE 802.4 standard for token buses. The following description refers to IEEE 802.4, which is the only significant type of token bus network in existence.

Information

The messages transmitted on an IEEE 802.4 network have the form illustrated in Figure 6.13. This is the same format as used by FDDI, except that there is no frame status byte at the end. Such a byte cannot be used on a token bus because messages are genuinely broadcast to passively listening computers, and there is no scope for computers to adjust frame status bits as messages pass by. The FDDI and IEEE

802.4 message formats can be viewed as the natural evolution and improvement of the token ring format, which was based on earlier pre-existing LAN products.

Time

The basic timing characteristics of a token bus are the same as for ethernet, in that all messages broadcast along a common medium. There are a range of standardized speeds for IEEE 802.4 token buses: 1, 5 or 10 Mbits per second for coaxial cable technology. There are also standards of 5, 10 and 20 Mbits per second for optical fibre technology. The coaxial cable standards specify broadband transmission, using a modulated signal with an appropriate bandwidth for the different speeds. This type of transmission is used so that the same cable can be shared by different transmissions, not necessarily all related to computer communications, at different frequencies. Relatively sophisticated electronics are required to ensure that the signals do not stray outside the proper frequency range. A cheaper alternative is to use carrierband transmission, which uses a modulated signal but does not share the medium with any other transmissions.

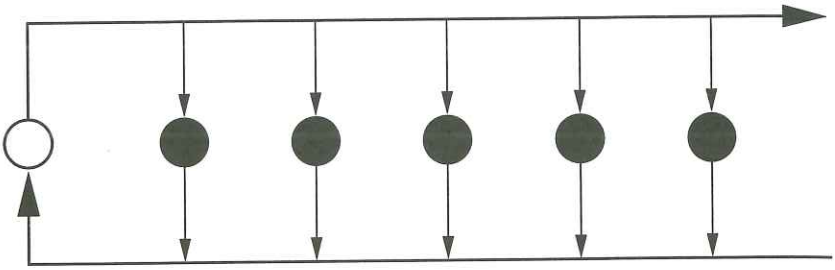
The basic service offered on a token bus network is an unreliable connectionless service, as for ethernet, with the transmission of each datagram being continuous over a time period.

Space

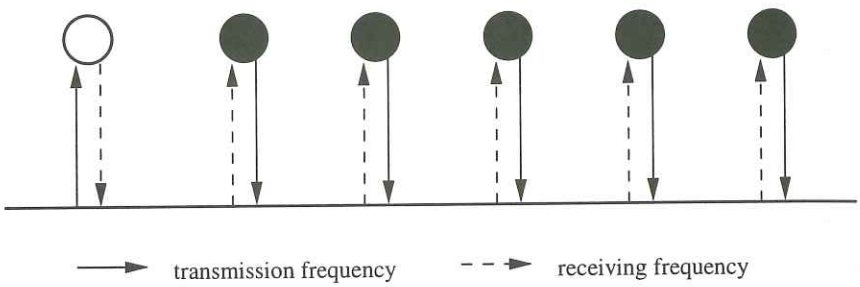
Like ethernet, token buses are used as local area networks. However, the fact that broadband signalling is used on coaxial cable media, and over better quality coaxial cable than that used for ethernet, means that longer distances can be spanned, for example, tens of kilometres. Also, it is possible to have more computers connected to a cable segment. The identifier scheme in token bus is the same as for ethernet, as is the way in which unicasts and multicasts are implemented in terms of the provided broadcast service.

Multipeer channel implementation

Broadband transmission is inherently unidirectional, since it is not possible to pass signals of the same frequency in both directions along the same cable. Because of this, a coaxial cable bus for an IEEE 802.4 network can be implemented in two different ways. These are illustrated in Figure 6.14. The first way has a pair of physical buses, one for each direction. All transmissions are made in one direction towards a **headend** at one end of the bus. The headend then repeats the transmissions along the other bus, from which the computers receive the signals. An alternative is to have a single bus, but to use two different signal frequencies. Transmissions are made at one frequency to the headend, and it then repeats the transmissions at a different frequency in the opposite direction, to be received by the computers in the network. This scheme is using frequency division multiplexing to implement the dual bus scheme.



(a) Headend repeater



(b) Headend frequency convertor

Figure 6.14 Two implementations of IEEE 802.4 bus

For optical fibre transmission, a single bus is not physically realizable. Instead, a hub configuration is used to form a logical bus. Computers are connected by fibres to the hub, and the only requirement is that the hub behaves like a bus: all computers hear all transmissions. This arrangement is the same as for hubbed ethernets, in particular the 10 BASE F standard. Note that broadband transmission is not used for optical fibre token bus networks.

Communication multiplexing

The token bus multiplexing scheme, as for token ring and FDDI, involves a token-based permission method. The central difference is that there is no physical ring to pass a token around. Instead, a logical ring is formed within the network, using unicast channels implemented on the multipeer channel. This ring need bear no resemblance to the ordering in which computers are physically located on the bus. All that is required is that the token circulates fairly around all of the computers

in the network. Whenever the token is passed from one computer to another, all computers receive the transmission. However, just as for normal unicast message transmissions, computers do not accept the token message unless it has their identifier as its destination identifier.

There is no master computer on a token bus network, in order to enhance reliability. Instead, there are complex protocols for managing the logical ring and the token in a distributed manner. This is a major disadvantage compared with ethernet. The essence of these protocols is the use of special types of message by computers to find new successors in the logical ring. When a token holder transmits the token to its current successor in the ring, it then listens for a next transmission. If the token transmission succeeds and the recipient is working normally, then the previous token holder will hear a message transmission by its successor. This will either be an information-carrying message or the token message being passed once again. If nothing is heard, or a damaged message is received, then the original token holder tries to pass the token once more. If this still does not work, the original token holder assumes that its successor has failed, and broadcasts a special message asking its successor's successor to respond, so that the logical ring can be relinked. If no response is received from the successor's successor, a more general procedure is followed to find a new successor.

This procedure is also activated at random intervals by token holders, in order to allow new computers to join the logical ring. It makes use of the fact that computer identifiers can be interpreted as numerical addresses. A special message is broadcast, inviting any computer with an address in the range between the message's source address and destination address to join the ring. Normally, the destination address is the address of the computer's successor in the logical ring. However, after a failure of its successor, and its successor's successor, the token holder puts its own address in the destination field, thus inviting all computers in the network to volunteer to be its successor. If this works, the result is that a two-computer ring is formed, and then it grows by soliciting further successors.

The problem with the invitation process is that more than one computer may respond at the same time. This is dealt with by a little ethernet-style process. If a collision is detected, the competitors wait for a small number of time slots before trying again. This is not a random number, but is determined by the least significant two bits of each computer's address. The first competitor to achieve a collision-free response to the token holder is admitted to the logical ring. A competitive process is also used to decide on the token holder when the network is initialized. In essence, each computer broadcasts for a period of time proportional to the logarithm of its network address. The result of this is that the highest-numbered computer becomes the token holder, and it can then solicit a successor to start forming the logical ring.

Like token ring and FDDI, token bus allows different priorities for message transmissions. The basic scheme used is like that of FDDI. There are four different priority levels, together with target token rotation times for each level. A computer transmits messages in descending order of priority, only allowing each priority

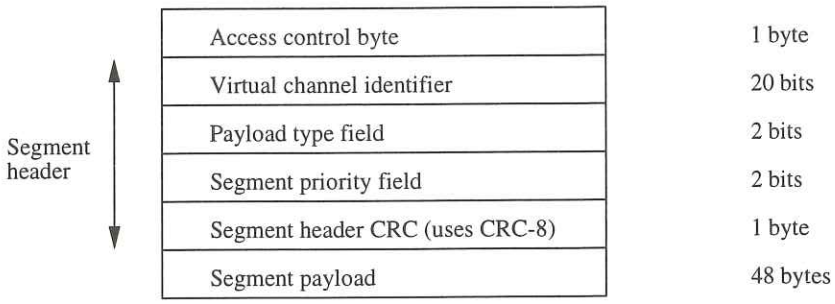


Figure 6.15 Format of DQDB message

level if the time used for transmission so far does not exceed the threshold for that priority level. The overall effect is that absolute guarantees can be made of the delay suffered by top priority messages. Such guarantees cannot be made for ethernet-style bus networks.

6.4.5 DQDB

The Distributed Queue Dual Bus (DQDB) is a message broadcasting system based on a pair of independent unidirectional buses. It is designed to be usable over a variety of types and speeds of physical channel. DQDB qualifies for inclusion in this chapter, because it does not involve any switching by computers. However, strictly speaking, DQDB does not agree with the chapter title, since messages are normally only broadcast in one direction along a bus, not in both directions. However, broadcasts are still possible, just by the sender carrying out two 'half broadcast' communications. DQDB was developed by IEEE as its 802.6 standard for Metropolitan Area Networks, and the typical type of channel would be a high speed offering from a common carrier.

Information

The format of a DQDB message is shown in Figure 6.15. All messages have the same length: 53 bytes. The first byte is used to control the access mechanism for the DQDB message broadcast capability. The remainder of the message is called a **segment**, and has a four-byte header, followed by 48 bytes of information. The choice of the word 'segment' reflects the fact that the information content consists of parts of a larger information unit chopped up to fit into the fixed-sized messages. The DQDB message is very closely related to the ATM cell, which is introduced in Chapter 2: 53 bytes long, with five bytes of control information followed by 48 bytes of real information. The close relationship is no coincidence.

Time

The operation of a DQDB network is controlled by a 125 microsecond clock. Each clock period, a sequence of messages is sent from one end of each bus to the other end. The exact number depends on the bit rate of the physical medium used. For example, over a 34.368 Mbits per second channel, nine messages can be transmitted in each period (including small amounts of extra control information between messages). The clock cycle length is the same as that for FDDI-II, and was similarly chosen to provide support for isochronous services because it is the same frequency as that used for telephone voice services.

The basic service offered by DQDB is an unreliable connectionless service, with the transmission of each fixed-length message being continuous over a time period. There are two classes of service: the pre-arbitrated (PA) service that gives guaranteed access to the multipeer channel for isochronous traffic, and the queued arbitrated (QA) service that gives access to the channel for other types of traffic.

Space

DQDB is primarily intended as a LAN interconnection facility. Thus, a DQDB network and a collection of connected LANs can be used to implement a larger network. The distance covered by a DQDB network may be in excess of 50 kilometres, and so can connect together a collection of LANs across a city. Thus, DQDB networks can be Metropolitan Area Networks (MANs). Because DQDB connects together computers representing LANs, rather than individual end-user computers, the typical number of connected devices is somewhat lower than for the networks already seen.

The mechanism used for identifying the source and destination computers for a communicated message cannot be directly decribed in terms of the basic message format. The actual mechanism varies, depending on the sort of service that is offered on top of the basic connectionless DQDB service for communication of 53-byte messages. There are three types: an isochronous service, a connection-oriented service and a connectionless service (allowing communication of arbitrary length messages).

The four-byte segment header contains a 20-bit **virtual channel identifier** (VCI) field. For the isochronous service and the connection-oriented service, this field identifies a particular connection between two computers, and so computers receiving such a broadcast message can regard this field as being a combined source and destination identifier. For the connectionless service, this field contains an all-ones value, and so does not give any spatial information. To determine the destination identifier of a message, it is necessary to look more deeply inside.

Each unit of data conveyed by the connectionless service may be up to 9188 bytes long, and is transmitted with a header that includes an eight-byte destination identifier and an eight-byte source identifier. Either 16-bit or 48-bit identifiers, as used in other IEEE 802 standards, or also 60-bit identifiers, as used in ISDN, can be carried. The data unit, together with its header (and also a trailer), must

be divided up to fit into a sequence of segments. Each segment used carries an identifier value that is unique to the data unit. Thus, computers receiving broadcast messages must look for segments that carry the first part of a data unit, and so contain the header, to check the destination identifier. If the destination identifier is relevant to a computer, then all subsequent segments with the same identifier are also relevant. In fact, the transmission of each data unit is akin to a little connection being established: the first segment gives the computer identifiers involved, then later segments just carry the identifier for the connection.

The exact details of data unit formats are not included here, since they are an issue for the next stage of implementation above the basic DQDB message transfer service. It is unfortunate that the connectionless DQDB provision has to be aware of the next-level implementation detail in order to determine the destinations of its messages.

Channel implementation

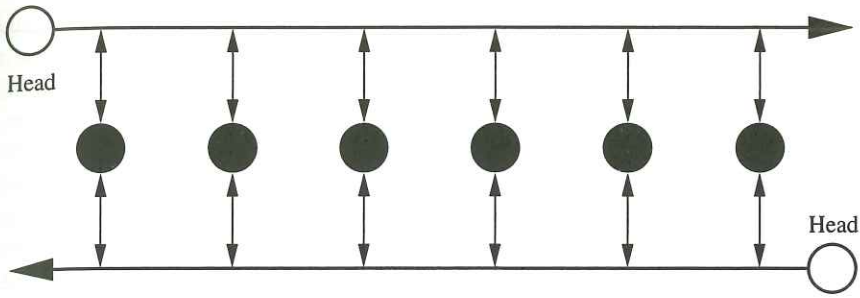
A DQDB network has two independent buses, one for transmission in each direction, so there is double the carrying capacity of a normal bus. At the beginning of each bus, there is a computer called the **head**, which is responsible for generating the messages that are sent along the bus. One difference from the other bus-style networks considered earlier is that computers can change messages as they pass by, rather than just passively read them. In this sense, 'dual chain' would convey a more accurate description than 'dual bus'.

Computers transmit information by placing it inside empty messages that pass by. A message is empty when it leaves the head, and until it reaches a computer that fills it. Thereafter, the message is full, and its contents can be read by any computer downstream on the bus. Thus, a half-broadcast effect is achieved. A source computer must ensure that information is transmitted on the appropriate bus(es), in order to reach its intended destination(s).

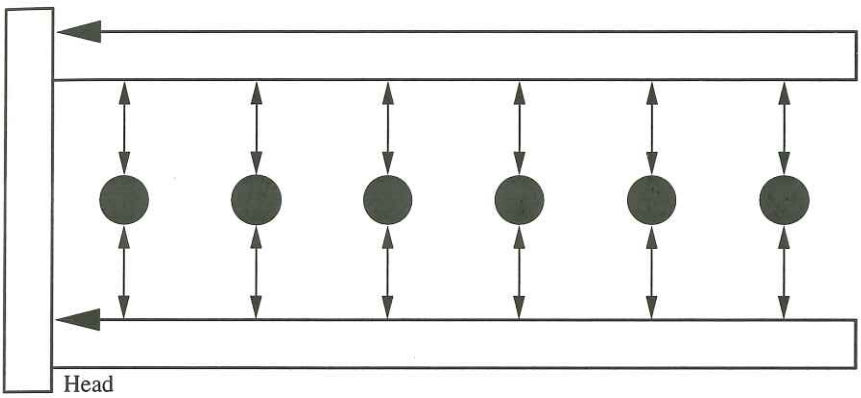
In an **open bus topology** network, the two heads are at opposite end of the network. In a **looped bus topology** network, the same computer is the head for both buses. Thus, the layout of the network is similar to the dual ring arrangement used for FDDI. The difference is that the head computer does not relay messages from one side to the other, and so there is no ring effect. The two topologies are shown in Figure 6.16. The advantage of a looped bus topology is that, should a failure occur in the network, the computers on either side of the fault can be configured as heads, with the normal head relinquishing head duties, and so an open bus topology network can be formed. This is also shown in Figure 6.16.

Multiplexing

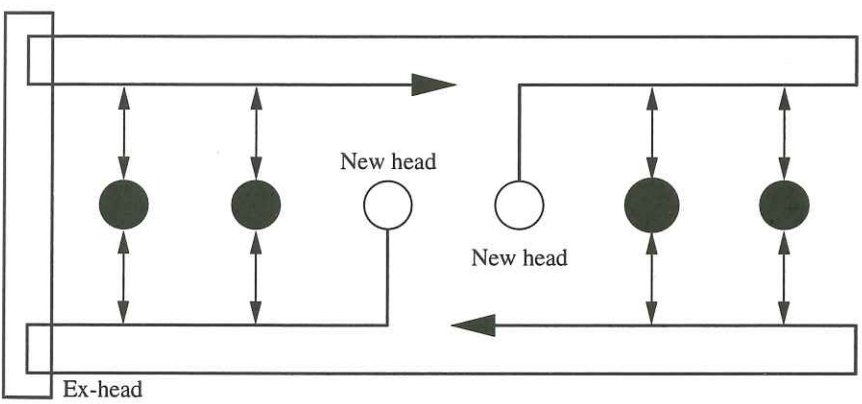
To support isochronous communications using the pre-arbitrated (PA) service, some of the messages transmitted by the head are pre-reserved for use by computers with PA connections established. Because messages are transmitted at a regular frequency, this gives the effect of time-division multiplexing. For other types



(a) Open bus topology



(b) Looped bus topology



(c) Looped bus topology reconfigured around failure point

Figure 6.16 Two implementations of DQDB bus

of communication, using the QA service, a reservation-based scheme is used to multiplex each bus. As the name Distributed Queue Dual Bus may suggest, this involves a distributed queue scheme, based on the ideas described earlier in the general discussion of reservation-based methods. In essence, each message transmitted by the head that is not pre-reserved can be regarded as having a ticket piggy-backed on it. When a message is used by a computer to transmit information, the ticket is removed before the message is passed on down the bus. This appears to be a scheme that is extremely biased towards computers that are near the head. However, the distributed queue reservation scheme stops this happening.

It is simplest to describe the reservation scheme for one bus only; a symmetrical and independent scheme applies to the other bus. Four bits in the access control byte at the beginning of each message are relevant. The 'busy' bit indicates whether or not the message is currently carrying information. When a message passes a computer with the busy bit zero, it is interpreted as being a ticket permitting transmission of a segment. Transmission is achieved by inserting the segment into the message and setting the busy bit. However, to achieve the distributed queuing effect, computers only make use of a ticket when it is their turn, as described earlier.

Three request bits are used to make reservations for the bus in the other direction. There are three bits because three levels of priority can be used, with a distributed queue for each priority. All queued reservations for a higher priority get precedence over queued reservations for a lower priority. Each computer can only make one request at any priority level at a time. To make a reservation for the bus at a particular priority level, a computer waits until a message passes on the other bus with the appropriate request bit zero, and then sets this request bit. After this, it must then allow enough empty messages to pass through, to satisfy all pending higher priority requests and any same-priority requests made earlier than its own.

In fact, the rule is a little more complicated than this. With the basic distributed queue scheme, it can be shown that, on highly loaded lengthy networks, computers nearer the head end of a bus get preferential treatment. To deal with this, a technique called **bandwidth balancing** is used. Instead of always exerting its right to use a passing empty message, a computer lets a certain proportion pass by. A parameter β is fixed for each bus. After using β messages, a computer is obliged to let the next empty message pass. The value of β must be in the range 1 to 64, with the default being 8. The IEEE 802.6 standard recommends that bandwidth balancing is enabled on any bus that spans a distance greater than the distance needed to accommodate a 53-byte message, for example, 2 km when a 44.376 Mbits per second transmission rate is used.

6.4.6 100 BASE VG-AnyLAN

In earlier examples, such as ethernet and token ring, one possible physical arrangement for the network is to use a hub, with individual computers connected

by paired channels to the hub. This has become a very widespread arrangement, since it allows centralized management and maintenance. The technology is sufficiently reliable that having a single central component is not a major concern — in the event of an occasional failure, one hub box can be speedily substituted for another. Given this, the IEEE 802.12 committee, which developed one of the higher speed follow-ons to the IEEE 802.3 10 BASE T ethernet standard, took the natural step of modifying the logical topology to align with the physical topology, i.e., having the form of a star or tree. This resulted in the **100 BASE VG-AnyLAN** standard. The '100' refers to a 100 Mbits per second transfer rate, and the 'VG-AnyLAN' refers to the use of Voice Grade cable (UTP twisted pair, used in existing 10 BASE T networks) and the ability to inter-work with existing 802.3 ethernet LANs and 802.5 token ring LANs. Thus, the standard offers a high speed upgrade path that can make use of existing LAN cabling, and also inter-work with existing LANs that are left intact. When older style networks, LANs in particular, are incorporated into newer networks, they are usually referred to as **legacy networks**. For brevity, the rest of the description will abbreviate the inelegant name 100 BASE VG-AnyLAN to just 'AnyLAN'. This is a shade more humanized than the alternative numerical name '802.12'.

Information

The messages broadcast in an AnyLAN network may have the format of either IEEE 802.3 ethernet messages or IEEE 802.5 token ring messages. However, one or other format must be agreed upon, and used throughout the network. The choice will be affected by the type of other LANs, if any, that are connected to the network.

Time

The raw data transfer rate of an AnyLAN network is 100 Mbits per second, as its full name indicates. In most existing hub networks, the maximum length of cable between a computer and a hub is 100 metres, so the typical latency between two computers via the hub will be less than 1 microsecond.

The basic service of AnyLAN is an unreliable connectionless service. Any communication time period involves the continuous transmission of a message. The bits of each message are transmitted synchronously, with each group of five bits being transmitted as a group of six bits. Each transmitted sextet has the same number of zeros and ones in order to maintain synchronization.

Space

In overall effect, the space characteristics of AnyLAN are similar to those of ethernet and token ring. A flat identifier scheme is used, and unicasts or multicasts are delivered to any computers that match the broadcast message destination identifier with their own identifier. Genuine broadcast messages have an identifier that is matched by all computers in the network.

The important difference is that genuine broadcasting of all messages does not happen. Although all messages are available to all computers in principle, the network prevents the delivery of irrelevant unicasts or multicasts to computers, rather than the computers themselves filtering out unwanted messages. This gives an added degree of security, since the computers are no longer required to be trustworthy. Because of this behaviour, there is not a true multipeer channel involved in implementing the network service. As for DQDB, redundant broadcasting is suppressed.

Channel implementation

In its simplest form, an AnyLAN network is connected as a star, with a switching device called a **repeater** in the centre. Note that this use of the word 'repeater' is different from its use in an ethernet context. Individual computers are connected to the repeater using voice grade cable. This cable in fact contains four UTP twisted pairs, and is of the same type described earlier in the brief summary of the 100 BASE 4T ethernet standard. An overall 100 Mbits per second transfer rate is achieved using splitting over the four pairs, each one having a 25 Mbits per second transfer rate.

In a more complex form, an AnyLAN network can be organized as a tree, with repeaters as nodes and computers as leaves. The connections between repeaters may use a variety of cabling technology, since these are not constrained by the need for backward compatability with earlier hub-style LAN cabling of computers. This can include either higher quality twisted pair or optical fibre. These types of cable allow distances of greater than 100 metres to be spanned.

When a computer wishes to broadcast a message, it sends it to its parent repeater. Then, the exact operation of the repeater depends on whether or not it is at the root of the tree (always the case when the network is a simple star). A root repeater selectively copies input messages received from one child to all of its children. All messages are copied to children that are repeaters. Only relevant messages, i.e., messages with appropriate destination identifiers are copied to children that are computers.

A non-root repeater copies input messages received from its children to its parent, and copies input messages received from its parent to all repeater children and to relevant computer children. Thus, although it may be physically organized as a tree, the network always functions logically as a star, with the root repeater at the centre. All messages are received by all repeaters, with selectivity of reception only occurring for individual computers.

Multiplexing

It is possible that more than one computer in an AnyLAN network might want to transmit a message at the same time. Permission-based multiplexing is used to ensure that only one computer transmits at a time. This does not need any distributed scheme for issuing tickets: repeaters use round-robin polling to arbitrate

between their children. As with the actual transfer of messages, polling is carried out on a logical star basis, centred on the root repeater. The root repeater polls each child in turn, to see if it needs to transmit. If the child being polled is a repeater itself, then it polls each of its children in turn. Thus, the overall effect is of the root repeater polling the leaves of the tree in turn.

In fact, a repeater does not explicitly poll each child in turn by asking whether the child wants to transmit and then awaiting a response. During, or at the end of, each message transmission, each child that wishes to transmit turns on a control signal to the repeater, indicating a request. Then, polling just consists of looking for the next request signal in the round-robin order. In a sense, each child is making a reservation; however, each time round, only one child has its reservation honoured, and the others must try again later.

AnyLAN allows two priorities of message transmission — normal priority for most messages and high priority for delay-sensitive messages. Therefore, there are actually two round robins. On each poll, a repeater checks the high priority round robin for requests and, if there are any, it grants the next request in order. If there are no high priority requests, the repeater checks the normal priority round robin for requests. Clearly, if there are many high priority messages, it is possible that normal priority messages might be severely delayed. To prevent this, there is a timeout of between 200 and 300 milliseconds set on the delay suffered by a normal priority message. If this expires, the message is treated as a high priority message by the repeater.

6.4.7 Discussion of examples

The examples of guided technology message broadcast networks illustrate a range of different multipeer channel implementations and communication multiplexing methods. Overall, the evolution of local area networks shows some evidence of a wheel turning full circle. Before the advent of LANs in which all computers are peers, such as ethernet and token ring, a 'local network' typically consisted of one powerful master computer managing a collection of smaller computers or terminals, in a logical star arrangement. With the introduction of hub-style cabling for LANs, the physical connections moved back towards the star style. Now, in examples like the 100 BASE VG-AnyLAN network, and switched ethernet, described in Section 7.4.1, the logical star arrangement has reappeared. However, the essential difference is that the centre of the star is no longer a master computer that is communicating with slaves. It is now just a dedicated switch with the sole purpose of facilitating communication. This means that the LAN still consists of a community of peer computers, and the dedicated centre can be made far more reliable than a general-purpose computer with extra networking roles could ever be.

Meanwhile, the traditional LAN technologies find applications in Metropolitan Area Networks, where the distances involved make distributed control mechanisms still desirable. Thus, FDDI is a natural extension of the token ring

mechanism. Ethernet is not particularly suited to extension over long distances, because its mechanism hinges on global sensing of the network state; thus, delays in acquiring sole use of the medium increase with distance. This is why DQDB has emerged as an alternative mechanism for bus-style MANs.

The other major trend in the development of message broadcasting networks has been the need to cater for traffic that is not asynchronous computer communication. Token ring and token bus allow some guarantees to be made on communication delays. However, these are not good enough for isochronous traffic, such as real-time speech and video. Therefore, a very traditional multiplexing scheme — time division multiplexing — has become popular again. As the examples of FDDI-II and DQDB show, modern time division multiplexing is blended with other multiplexing schemes, so that a mix of isochronous and asynchronous traffic can be supported efficiently.

6.5 UNGUIDED TECHNOLOGY NETWORK IMPLEMENTATIONS

The examples of message broadcasting networks given above all involved guided media — electrical cable or fibre optic cable. Wireless networks make use of unguided media, and have the dual advantages of avoiding cabling costs and of allowing the easy interconnection of portable computers. The media used are either radio waves or infra-red signals. These are inherently broadcast media, and so are directly suited to message broadcasting.

A variety of manufacturers have developed their own wireless LANs, each one rather different from the others. However, international standards began to emerge in the mid 1990s. The IEEE 802.11 standard — an addition to the well-established set of IEEE 802 LAN standards — is concerned with wireless LANs. In addition, and in consultation with IEEE, the European Telecommunications Standards Institute (ETSI) formulated a standard called HiperLAN (High Performance Radio LAN) for wireless LANs operating at around 10 times the rate of 802.11 standard LANs. In due course, as with wired LANs, these standards will prevail over proprietary solutions, allowing inter-operability of different manufacturers' equipment.

Broadcasting messages using radio is not a new phenomenon. Indeed, the multiplexing scheme used by ethernet had its root in the scheme used in a pioneering radio-based network of the University of Hawaii. There is also a long-standing standard protocol (AX.25) used by amateur radio enthusiasts for communication between computers. The difference with modern wireless LANs, compared with earlier radio networks, is that they provide much higher speeds which are comparable with guided media LANs, and also have much higher reliability. The higher speeds are facilitated by the use of higher frequency radio or by the use of infra-red. The higher reliability is achieved by the use of more subtle sig-

nalling techniques to lessen radio interference and by the use of error detection and correction techniques.

Information

The messages used for wireless broadcasting do not need any particular additional control fields. Starting and ending delimiters, together with source and destination identifiers, are sufficient along with the information itself and an error detection code. If the bit error rate suffered by messages is relatively high — for example, rates of up to 1 in 1000 bits being corrupted are not unknown where precautions are not taken to guard against radio interference — then it is necessary to keep messages short. Otherwise, it is probable that a high percentage of messages will suffer damaged bits. If such a consideration does not apply, then there may be an incentive to make messages as long as possible. For example, if CSMA/CD is used to multiplex the medium, then larger messages are better to minimize contention periods.

Time

The IEEE 802.11 standard offers raw data rates of 1 or 2 Mbits per second using radio or baseband infra-red, and also rates of 4 or 10 Mbits per second using broadband infra-red. The HiperLAN standard offers up to 20 Mbits per second using radio, with a genuinely high-speed rate of 155 Mbits per second planned eventually. The basic service of wireless LANs is an unreliable connectionless service. A communication time period involves the continuous transmission of the bits of a message. The timing of the individual bits depends on the transmission mechanism used.

For radio, **direct sequence spread spectrum** transmission involves the data bit stream being exclusive OR-ed with a higher rate pseudo-random bit stream, so that the resulting transmitted bit stream has the higher rate. This is illustrated in Figure 6.17. An alternative is to use **fast frequency hopping** transmission. This involves changing the signal frequency pseudo-randomly several times during each bit time. Both of these types of transmission mode are allowed by the 802.11 standard. They do not feature as part of the HiperLAN standard. Both types involve a decomposition of each bit time period. The reason for doing this is to lessen interference from other radio sources. Roughly speaking, the idea is to widen the frequency band used for transmission, but the technical details are not of mainstream importance here.

Space

The range of a 20 Mbits per second radio LAN is only 50 metres, which is distinctly low, compared with guided media LANs. However, a range of 800 metres is possible for a 1 Mbit per second radio LAN. Infra-red LANs are restricted to a single room, for the simple reason that infra-red does not pass through walls. The shorter distances possible with high-speed radio LANs and infra-red

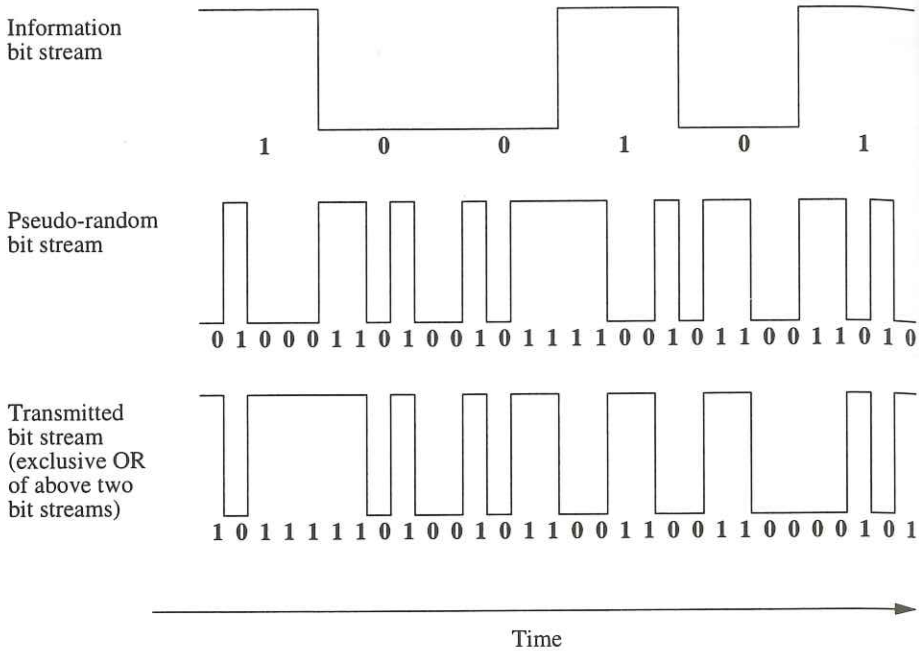


Figure 6.17 Example of direct sequence spread spectrum transmission

LANs are adequate to allow mobile communications. A portable computer can communicate with other computers in the same room or corridor that are connected to guided media LANs. When a more extensive radio-based network is required, then intermediate relay stations must be used to pass on radio transmissions from one computer to another. However, this is not really message broadcast networking any more (although, of course, messages are being physically broadcast); rather, it is message switching, which is the subject of Chapter 7.

Mobility is an issue that relates space and time, since computers may actually be in motion when communicating. The IEEE 802.11 and HiperLAN standards are both intended for slowly moving computers only, for example, at speeds no more than 10 metres per second. This is acceptable for LANs, where the assumption is that the computers are in a small area, probably indoors, and so the scope for movement is limited. Other arrangements are needed for WANs, where computers might be in fast moving vehicles.

No special identifier schemes are necessary for wireless LANs. A flat identifier scheme, as used in wired LANs, is perfectly adequate. In amateur message radio networks, one particular identifier scheme is used: computers are identified by the radio call signs of the radio hams who are carrying out the communication for the computers.

Channel implementation

Radio provides a natural broadcast technology. Infra-red is more natural as a point-to-point medium. However, the output of an infra-red source can be optically diffused, so that it is spread over a wide angle. By this diffusion, and then reflection from walls of a room, it is possible to achieve broadcasting. An alternative, for both radio and infra-red, is to use directional transmission in connection with a fixed roof-mounted reflector. A passive reflector just reflects the waves so that they reach the receivers of all computers. An active reflector repeats signals received by transmitting them back to the receivers. This means that the transmitters can use lower power signals than are necessary with passive reflectors.

Multiplexing

Wireless LANs suffer from the same multiplexing problem as wired LANs, and the main solutions either involve some form of contention avoidance or detection, or physical division. Permission-based or reservation-based schemes are less common, but an optional feature of the IEEE 802.11 standard involves a permission-based system operated by a master computer.

CSMA/CD schemes for wireless LANs are similar to the scheme used for ethernet. The only complication is that it is not possible to transmit and receive at the same time, and so a transmitter cannot hear collisions. To deal with this, at the beginning of message transmission, computers rapidly switch between transmitting and receiving, in a pseudo-random manner, for the first few bit time periods. If another transmission is heard during one of the receiving periods, a collision has been detected.

The multiplexing scheme used in HiperLAN is called **EY-NPMA** (Elimination Yield — Non-pre-emptive Priority Multiple Access), which is essentially a variant of CSMA/CD. The elaboration comes in the collision detection method. Rather than any computer immediately starting to transmit its information when the channels becomes free, computers transmit control messages of differing lengths. The general idea is that the computer that has chosen the longest length wins, since other competitors back off. The lengths are chosen pseudo-randomly, so the mechanism is fair. In fact, matters are a little more complex than this, since messages can have five different priority levels.

An elaboration of CSMA/CD is **CSMA/CA**: CSMA with Collision Avoidance. This involves a little extra politeness before transmitting. After sensing that the medium is free, computers wait for a short random period of time before transmitting. This is designed to lessen the problem of clashes when several computers are waiting for the medium to become free before transmitting. CSMA/CA is used in the IEEE 802.11 standard, in which it forms part of an overall multiplexing protocol: **Distributed Foundation Wireless Medium Access Control** (DFWMAC), which also deals with other problems suffered by wireless communication.

A main problem solved by DFWMAC is the fact that wireless communication might not involve a genuine multipeer channel. There may be computers that are

not able to hear 'broadcasts' made by other computers that are too distant. Thus, strictly speaking, the wireless network implements a set of 'multi-multicast' channels. This leads to the **hidden station problem**, which is illustrated in Figure 6.18. The problem occurs when one computer cannot hear that another computer is already broadcasting to part of the network, and so it begins to transmit. This transmission will collide with the one already in progress, in those parts of the network that can hear both transmissions.

DFWMAC involves a four-stage handshake for message transmission, the first two stages of which deal with the hidden and exposed station problems. The four stages are:

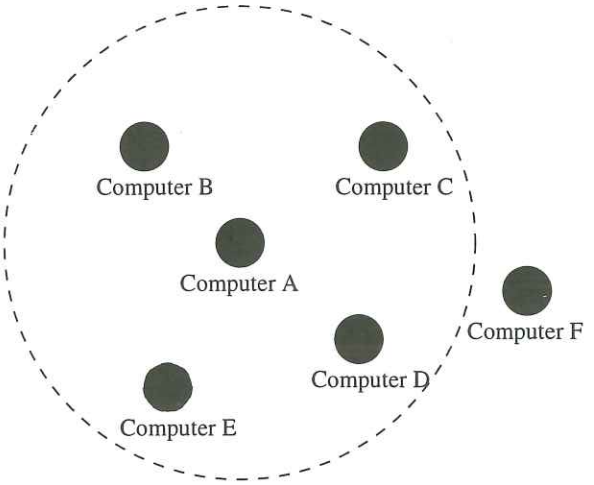
- message sender sends a Ready To Send (RTS) control message, containing the length of the actual information message, to that message's recipient(s);
- message recipient(s) send back a Clear To Send (CTS) control message, also containing the information message length;
- message sender sends the information message; and
- message recipient(s) send back an acknowledgement control message.

An alternative at the second stage is to send back a control message indicating the intended recipient is not able to receive; obviously, if it is completely switched off, no message will be returned. The protocol deals with the hidden station problem because, although a hidden computer will not hear the RTS message, it will hear the CTS message. The fact that the RTS and CTS messages contain the length of the actual message to be transmitted means that all listening computers are given warning of the length of the transmission to follow.

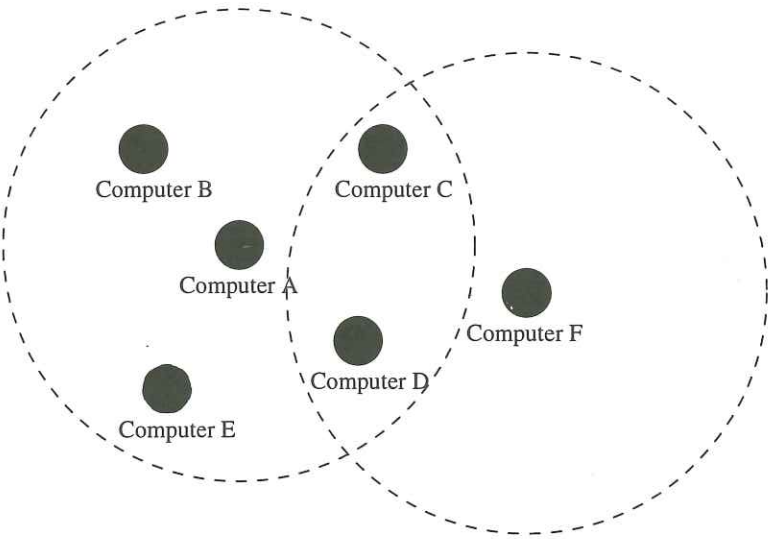
Further, but optional, complication is introduced in DFCMAC, in order to deal with communications that require guarantees on their starting time. A master computer can be present, and it is given priority access in the CSMA/CA scheme, that is, it can acquire channel use rights before any other computers make bids. Having acquired the channel, the master computer then issues polls to computers that have been put on its list of those requiring regular time periods for communications. When polled, each such computer can transmit without encountering contention. After polling is finished, the channel is made available for normal communications.

As alternatives to isolated multiplexing methods, time division multiplexing and frequency division multiplexing are two possible physical division schemes which can be used when all transmissions occur via a central base station. That is, computers transmit messages to the base station, which then broadcasts them to all of the network. Assignment of time slots or frequency bands may be done statically by the base station, or dynamically in response to requests.

A further possibility is **code division multiplexing**, for example, where frequency hopping is used. A pseudo-random sequence is used to control the hopping and, if a different sequence is assigned to each computer, and all the assignments are made known to all of the computers, then multiplexing of point-to-point channels is possible. A transmitter selects the pseudo-random sequence of its receiver, and uses it for the transmission. With care in the choice of pseudo-random se-



(a) Computer A is heard by Computers B, C, D and E, but not F



(b) Computer F collides with A at Computers C and D

Figure 6.18 Hidden station problem in wireless network

quences, it is unlikely that competing transmissions will use the same frequency at the same time. The disadvantage is the need for all computers to have complete knowledge of the others' sequences. Also, of course, this technique does not work for broadcasting to multiple receivers.

6.6 CHAPTER SUMMARY

Networks in which computers perform no switching, or only non-selective switching, generally offer a connectionless message broadcasting service. In practice, such networks occur as Local Area Networks or Metropolitan Area Networks. These networks are formed by directly using physical channels between computers. The physical media are normally guided at present, but use of unguided media is increasing to give wireless communication. Communication rates are in the megabits per second range: 10 or under for earlier networks, 100 or over for more modern networks.

One main implementation problem is constructing a multipeer channel using the underlying channels of the network. This might involve using an existing common medium directly, or constructing regular topologies from existing unicast channels. Topologies include chains, rings, stars and trees. The other main implementation problem is multiplexing broadcast channels onto the multipeer channel. This might involve a standard flow control technique, such as use of permissions or reservations. Alternatively, computers may act independently and aim for contention avoidance and/or detection. Finally, physical division multiplexing of time or frequencies may be used.

Various standard message broadcasting networks have been developed, to take account of differing channel availabilities, and different qualities of service required. The standardization efforts of IEEE have been a significant influence on what is used in practice. Examples of different networks types are: ethernet; token ring and its successor, FDDI; token bus; DQDB; 100 BASE VG-AnyLAN; and 802.11 wireless. There are many other variants, both developed experimentally and used in practice.

6.7 EXERCISES AND FURTHER READING

6.1 Explain how a room full of air can act as a message broadcasting network for a group of people within the room.

6.2 Discover whether you have access to any (a) LANs, and (b) MANs, and, for each, what computers are connected together.

6.3 In what ways might you find a Home Area Network useful?

6.4 What common features does the outline message format given on page 169 have with a letter sent through the postal system?

6.5 Suppose a message is transmitted over a phsical channel. Give an intutitive argument why the minimum length of the communication time period is equal to the bandwidth*delay product of the channel plus the message size in bits, divided by the channel rate.

6.6 If you are familiar with the internal architecture of computers, discuss the differences and similarities between an internal computer bus and a bus used as a common medium for a message broadcasting network.

6.7 Give examples of how bus, chain, ring, star and tree arrangements might underpin human communication systems.

6.8 Hub arrangements for common medium or chain/ring networks are popular with maintenance engineers and network managers. Why do you think this is so?

6.9 Compare the contention avoidance/detection strategy for isolated multiplexing with the ways in which people within a group would try to avoid speaking at the same time.

6.10 Find out about the pioneering University of Hawaii radio network that led to the ALOHA method for isolated multiplexing.

6.11 Suppose that three computers are in contention for transmission rights, and the randomized exponential backoff algorithm is being used to underpin isolated multiplexing. Either by evaluating probabilities, or by computer simulation, investigate how many time slots are likely to pass before all three transmissions are successfully carried out.

6.12 Give practical examples of network use patterns for which an isolated multiplexing approach is likely to be (a) effective, and (b) ineffective.

6.13 If you have read the book *Lord of the Flies* by William Golding, describe the permission-based mechanism used for multiplexing communications among the shipwrecked boys.

6.14 Repeat Exercise 6.12, but for a permission-based multiplexing approach.

6.15 A shop can serve only one customer at a time, and has an area for customers waiting to be served. If customers can stand/sit where they like in this area, but the aim is that customers are served in order of arrival, describe a reservation-based approach to multiplex the communications between customers and the server.

6.16 The term 'channel hopper' is sometimes used to describe a television viewer who continually switches from one channel to another. Explain why this process is more akin to time division multiplexing than to frequency division multiplexing.

6.17 Find out the main differences between the original Ethernet design and the final IEEE 802.3 ethernet standard.

6.18 Work out the minimum length that a communication time period can have on a 10 BASE 2 ethernet.

6.19 In an **empty packet ring** (sometimes called a 'slotted ring'), small empty messages circulate, rather than a token circulating. When wishing to transmit,

computers wait for an empty message, then fill it. Discuss the relative merits of the token ring mechanism and an empty packet ring mechanism.

6.20 Find out how a **register insertion ring** operates, and explain how it is contention-free, but does not involve either a permission-based or reservation-based approach.

6.21 Write down a list of differences between the 802.5 token ring mechanisms and the FDDI mechanism, and explain why each difference exists.

6.22 Give informal arguments to show why the typical token circulation time in an FDDI network tend towards the TTRT, and the maximum circulation time is never more than double the TTRT.

6.23 Suggest reasons for the format of an FDDI-II cycle, in particular, for the existence of the 12 groups within each cycle.

6.24 List the properties of a token bus network and, for each property, say whether it is inspired by ethernet or by token ring (or by neither). Explain why this inspiration is likely to have occurred.

6.25 Find out more about how the logical ring is maintained by the IEEE 802.4 protocol.

6.26 Investigate the history of the IEEE 802.6 standardization process that led to the selection of DQDB as the MAN standard.

6.27 Give an intuitive explanation of why bandwidth balancing is needed as an extra mechanism on length DQDB networks.

6.28 The 100 BASE VG-AnyLAN standard is designed to harmonize with existing legacy networks that, in particular, use IEEE 802 identifiers for computers. Discuss the benefits that would arise if a more hierarchical identifier scheme could be used instead.

6.29 Why might the timeout period for AnyLAN to upgrade normal priority messages to high priority have been chosen to be in the range 200–300 millseconds?

6.30 Conduct a survey of the different radio-based LAN technologies that are currently available.

6.31 Discuss the difference between contention avoidance and detection in networks based on guided, and on unguided, technologies.

6.32 The **exposed station problem** occurs when a computer hears a transmission in a wireless network, and needlessly holds off transmitting since its target recipient is outside the range of the transmission in progress. Explain how the four-stage handshake mechanism used in DFWMAC gives scope for avoiding the exposed station problem.

Further reading

The topics covered in this chapter are predominately associated with the topics covered in the LAN and MAN literature. One possible specialist textbook for

further reading is *Local and Metropolitan Area Networks* by Stallings (Macmillan 1997), although there are several others. The classic paper on Ethernet is "Ethernet: Distributed Packet Switching for Local Computer Networks" by Metcalf and Boggs, published in Communications of the ACM, **19**, July 1976. For information about a major influence, the University of Hawaii radio network, see "Development of the ALOHANET" by Abramson, published in IEEE Transactions on Information Theory, **IT-31**, March 1985. A full understanding of the issues behind the workings of more recent, and considerably more complex network types, such as FDDI and DQDB, requires a plunge into the communications research literature.