Packet Switched Networks

THEORY AND PRACTICE

RICHARD BARNETT AND SALLY MAYNARD-SMITH

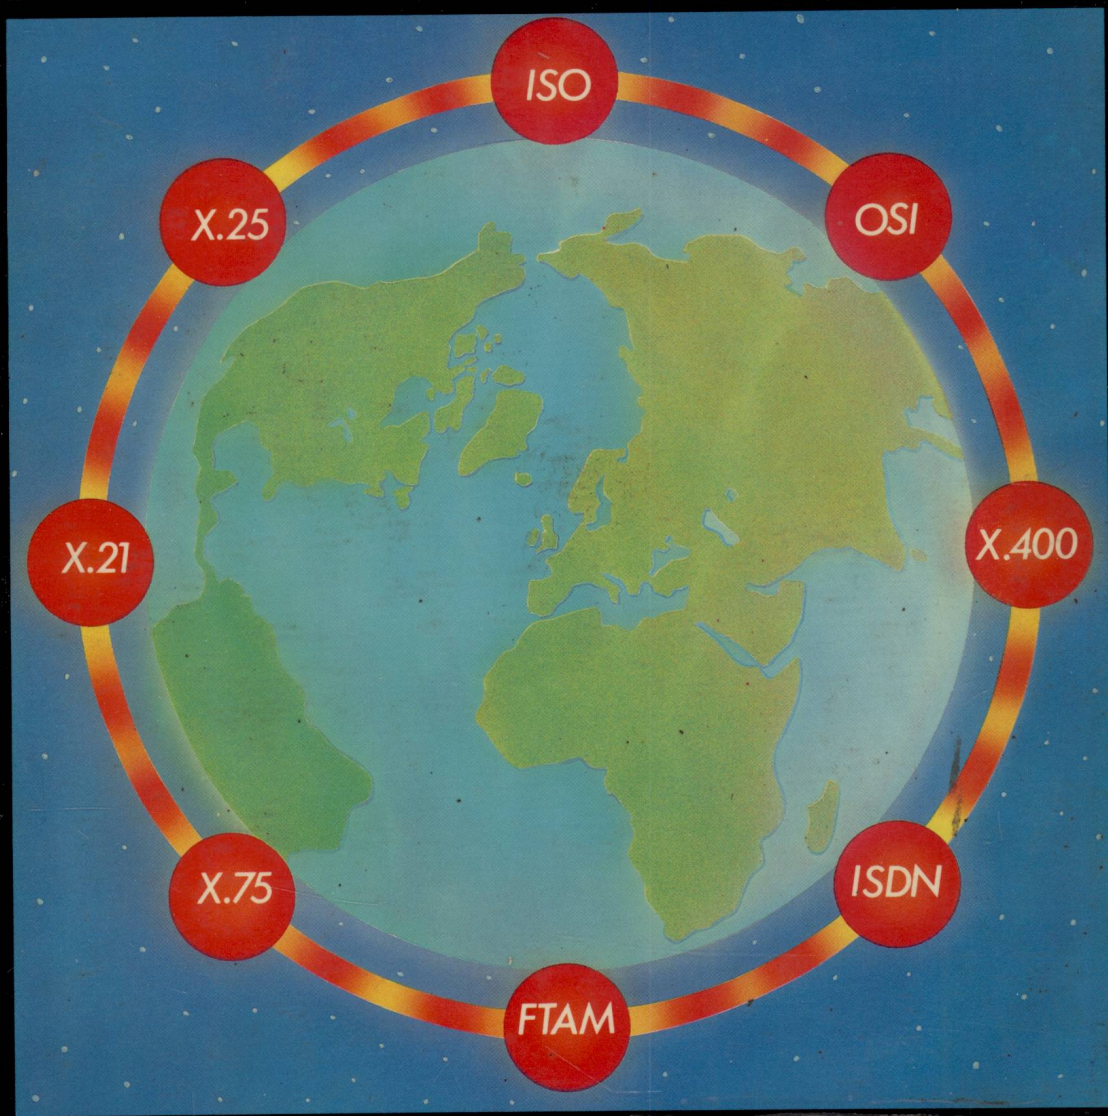

SIGMA PRESS

PACKET SWITCHED NETWORKS- Theory and Practice

Richard Barnett and Sally Maynard-Smith

SIGMA PRESS

Copyright © Richard Barnett and Sally Maynard-Smith, 1988.

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 prior written permission.

First published in 1988 by
Sigma Press 1 South Oak Lane, Wilmslow, SK9 6AR, England.

British Library Cataloguing in Publication Data

Barnett, R.
Packet-switched networks : theory and practice.
1. Computer networks 2. Data transmission systems 3. Packet switching (Data transmission)
I. Title II. Maynard-Smith, Sally
004.6'6 TK5105.5

ISBN (Sigma Press): 1-85058-095-2

Distributed by
John Wiley & Sons Ltd., Baffins Lane, Chichester, West Sussex, England.

Printed by Interprint Ltd, Malta

Cover design by Professional Graphics, Warrington, UK

Acknowledgments: the following are trademarks or registered trademarks:
DEC - Digital Equipment Corporation; DECNET - Digital Equipment Corporation; Ethernet - Xerox Corporation; IBM - International Business Machines Corp; INTEL - Intel Corporation; IPSS - British Telecommunications plc; KiloStream - British Telecommunications plc; MegaStream - British Telecommunications plc; NETBIOS - International Business Machines Corp; Packet SwitchStream - British Telecommunications plc; PC AT - International Business Machines Corp; PDP-11 - Digital Equipment Corporation; Prestel - British Telecommunications plc; VAX - Digital Equipment Corporation; WANG - Wang Laboratories Inc; Z80 - Zilog Corporation.

The authors apologise for any inadvertent failure to acknowledge other trademarks.

CONTENTS

试读结束：需要全本请在线购买： www.ertongbook.com

Part 1:

Network Basics

'Packet Switched Networks: Theory and Practice' has been split into four parts: Part 1 consists of a basic overview to networking; Part 2 covers protocols used in packet switched networks; Part 3 is more practical and deals with private packet switched networks; Part 4 takes a look at future developments.

This book is not intended to be an implementor's guide, but an introduction to packet switched networks. In all cases where the reader requires more information the relevant standard documents are suggested as further reading. Where views are given, they are those of the authors and not necessarily held by everyone else.

Network Basics

Part 1 starts with a guide to networking in general aimed mainly at beginners. It goes on to cover the basic concepts of packet switched networks. Finally, it covers the components of a packet switched network.

Packet Switched Network Protocols

Part 2 goes into much more detail. It covers the protocols used in packet switched networks, looks at International Standards, UK standards and X.25 (a packet switched network protocol) in detail, and higher level protocols and security.

Private Packet Switched Networks

Part 3 looks at private packet switched networks in practice, giving useful information for those installing them. It covers such things as what to look for when purchasing packet switching equipment.

Future Developments

Part 4 looks at some possible future developments, those likely to be around in the near future and others more hypothetical in nature. This part is fairly subjective - the future, of course, is unknown (to some extent anyway)!

1

A Beginner's Guide To Networking

1.1 Introduction

In this section we describe, very briefly, the major types of networks, to give some general purpose background to networking before concentrating on packet switched networks.

A large amount of networking jargon will be introduced in this section. It is not essential to be able to remember the meanings of all of the acronyms as there is a glossary at the end of the book. However, when dealing with networking, the jargon will surface and it is useful to have an idea of what the different terms mean.

1.2 Why Are Networks Required?

Imagine that networks do not exist - a situation that was almost true not too many years ago. Figure 1.1 shows the situation, a personal computer sitting on a desk. The personal computer is able to run a lot of software and do some very clever things, but the user is limited purely to the capabilities of the personal computer.

Suppose the user needs access to some information or software on another machine? If the machines are roughly compatible, then physically transferring floppy disks can provide a solution to the information transfer problem. What if the machines do not have compatible disks?

Most computers worthy of the name have a serial interface. This is sometimes referred to as an 'asynchronous interface' or an RS-232 port. Whatever the name, if 'terminal emulation' software is run on the personal computer, the serial port can be connected to a serial port of another computer. The personal computer can then act as though it is a terminal and operate the other computer. Many terminal emulation packages allow files to be transferred across the serial interface. A good example of a terminal emulation/file transfer package is the program called Kermit, which is available for most computers and ideal for transferring information between different types of computers.

What are the problems with this sort of set-up? Firstly, the computers must be close together, as the RS-232 electrical signalling cannot be driven over very long distances. Secondly, the transfer is limited by the maximum speed at which the serial interface

试读结束：需要全本请在线购买： www.ertongbook.com

can transfer data. This is often only 9600 bits per second. Since there are usually ten bits to a byte of data, this means that the link will be limited to 960 bytes per second. At this rate, a one megabyte file would take a minimum of around 20 minutes.

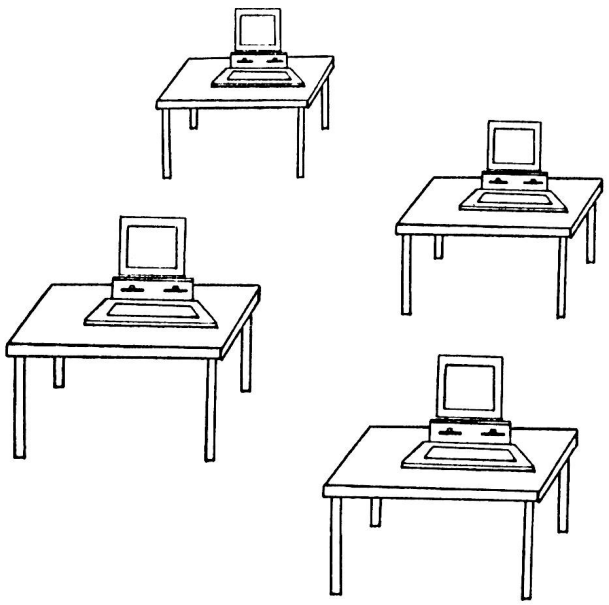

Figure 1.1 Life without a network

The first problem can be solved with devices called line drivers and modems. These allow much longer links between the computers. Most modems are designed to be connected to the Public Switched Telephone Network (PSTN) (i.e., the telephone system). Modems can not only automatically dial out but also answer incoming calls. The use of modems and the PSTN means that distance is not a problem. The two communicating computers can be on opposite sides of the world, as long as one can phone the other.

Modems tend to be limited in the rate that they transfer data due to the very low capacity of telephone lines. Modems that work at up to 2400 bits per second are available at reasonable prices but if higher speeds are required, the price goes up quite rapidly. At 2400 bits per second, a one megabyte file would take a minimum of around one hour and 20 minutes.

So we have seen that it is possible for two computers to communicate at moderately low speeds over any distance, provided that they are near a phone line. Suppose that a large office complex has several hundred, possibly different, personal computers that

often have to transfer data between themselves and possibly to computers located at a remote site. How can this be done?

Clearly, it is rather impractical to have directly wired connections between each pair of computers and rather wasteful to have to use the telephone system to allow the computers to communicate - not to mention the low transfer rates.

The solution is to use a network. All of the computers that need to communicate with each other are connected into the network. A link to the remote site is also connected to the network while another network at the remote site connects all of the computers there to the remote link.

Sounds good, doesn't it? If only life were that simple. There are a wide range of different network types and technologies, some good at some things, but bad at others. No network so far invented solves everyone's communication requirements.

Network technologies are often grouped into two different classes - LANs and WANs.

1.3 LANs and WANs

These are two acronyms are always being bandied about in the computing press but what are they and in which ways are they different?

LAN stands for Local Area Network. The 'Local' bit means that the network technologies in this group tend to be limited in the distance over which they can be used. A typical maximum distance between two computers connected to a LAN might be one kilometre. LANs tend to be used to interconnect computers within one building or else a set of buildings, physically very close to each other.

LANs can often support very high data transfer rates, frequently at several million bits per second. This allows users of computers connected to the LAN to share disks between the computers and allow very rapid transfer of information between them.

On the negative side, LANs are often very susceptible to hardware and software failures. Very few have decent network management facilities. (Network management facilities give information about how the network is performing, if there are any faults, etc. More about this later.)

WAN stands for Wide Area Network. The idea is that WANs take over where LANs leave off. WANs often cover extremely large distances, although they are being used more and more in LAN territory - over small distances. Packet switched networks fall into the WAN category as do some of the more esoteric networks such as packet radio networks.

The so-called 'national networks' are all packet switched networks. British Telecom's Packet SwitchStream service uses the X.25 packet switched network protocol.

Connection to Packet SwitchStream (more commonly known as PSS) is available nationwide. It interconnects with WANs in other countries, permitting PSS-connected computers to communicate with computers in most parts of the world. Another of the national networks in the UK is JANET (the Joint Academic NETwork) which allows free access for academic use to computers at any university site and soon the Polytechnics.

The main claim to fame of WANs is the interconnectivity they provide, as they can connect computers over large distances. They can also be quite resilient to hardware and software failures, and hence very reliable. Network management facilities can be very good and have to be where network availability is very important.

The main drawback of WAN technologies has been the data transfer rates available. A typical connection to a WAN may operate at 9600 bits per second, the same rate as the serial interface on the average personal computer! Certainly, this is no competition for LAN technologies.

The situation is changing as WAN technology improves. Another important improvement is the high speed communication links now available. These operate over WAN type distances and can support rates up over two million bits per second if required.

Having given some background to networking we will give more detail on some of the different network types in more detail.

1.4 Ethernet

Ethernet is probably the most successful of all of the LAN technologies. In terms of topology, it is a 'bus network'. Figure 1.2 shows the general idea. All the devices are connected to the main cable, the Ethernet cable, via transceivers. Separate Ethernet cables can be joined together using repeaters or bridges. The cable itself is a special coaxial cable, as are the connectors and the transceivers.

Data can be transferred via an Ethernet network at ten million bits per second, a quite impressive speed. Since the two computers in communication can be up to a kilometre or so apart, Ethernet can be extremely effective for communications between computers in the same building or group of buildings.

How do the computers on the network decide which is able to use the cable? Because Ethernet is an example of a 'baseband' network, only one computer can use the network at any one time. So how do they do it?

In fact, the process is a good example of ordered chaos. When a computer wants to transmit some data on the Ethernet, it first 'listens' to what is happening on the cable to see if anyone else is using it. If so, it must wait. If not, it will start transmitting the data while still listening to the cable. This is done to check that the data being transmitted is

not being corrupted, the most probable reason for which is another computer on the network transmitting data at the same time. This is known as a 'collision'. If a collision occurs, then both computers must give up and try again later. They wait a random time interval before trying to send the data again. This process is known as CSMA/CD (Carrier Sense Multiple Access/Collision Detection).

What are the practical aspects of Ethernet? There are some physical details that need care. The Ethernet cable must be installed correctly, with all the paraphernalia that goes with it (taps, transceivers, terminators, etc.). Once installed, it should be fairly reliable, as there is comparatively little to go wrong. If something does go wrong, fault finding can be a little tricky, as it may not be obvious which of the various computers connected to the cable is causing the problem.

Monitoring of network performance is a fairly simple affair, as all devices connected to the cable see all of the data being transferred. The other side of this coin is that Ethernet is inherently prone to 'eavesdropping'. (Anybody can look at the data and extract things like passwords and sensitive information unless the data is encrypted at a higher level.)

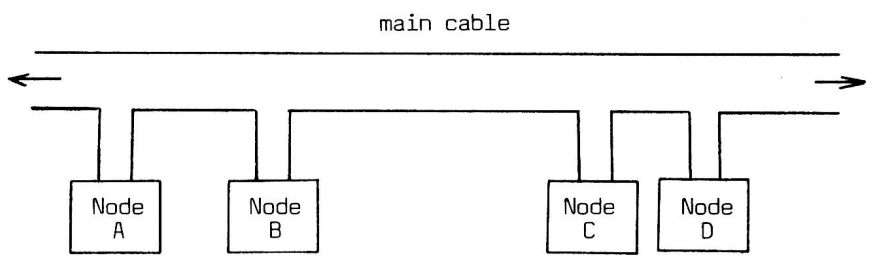

Figure 1.2 Bus network

Ethernet is good, unless a lot of computers are trying to use it at once (in other words, as long as the loading is not too high). When loading does increase, the total data throughput tends to decrease, because of the number of collisions that occur.

1.5 Broadband Networks

Networks like Ethernet are known as 'baseband' networks. All of the capacity of the cable is taken up in one fell swoop by a single node transmitting data. 'Broadband' networks are different in that they use the technique of 'frequency division multiplexing' to provide a number of simultaneous channels on the same physical cable. A typical cable may have a useful bandwidth in excess of 300MHz which, in theory at least, can support a large number of slower channels.

The concept is very much like the radio FM waveband. This is divided into a number of frequency channels that may or may not contain a signal (Capital Radio, for example). The main difference is that all of the signals are contained within the cable rather than broadcast into the atmosphere.

The main commercial example of a broad band network is the WANG network. Figure 1.3 shows the general idea. As far as a node on the network is concerned, there are two separate cables: one for received data; the other for transmitted data. It is, in fact, one cable looped at one end of the network.

In order to transmit or receive data on the broadband network, the network nodes must use a radio frequency (RF) modem to convert the data into a suitable form for the cable. One important issue is the allocation of frequency channels to nodes. If the node is to be able to utilise several different channels, then its modem must be 'frequency agile'. This means that the frequency at which the modem operates can be programmed. A central controller can then allocate one of several channels to a node, which programs its modem to operate at that frequency.

Due to the rather impressive facilities required of the broadband network interfaces, they tend to be very expensive. Also, they require that the special cable be installed correctly. So far, broadband networks have found only limited application, mainly due to the high cost involved. In terms of function, ISDN (see section 1.12) promises to offer very high bandwidth channels with guaranteed throughput and so broadband networks may never be used extensively.

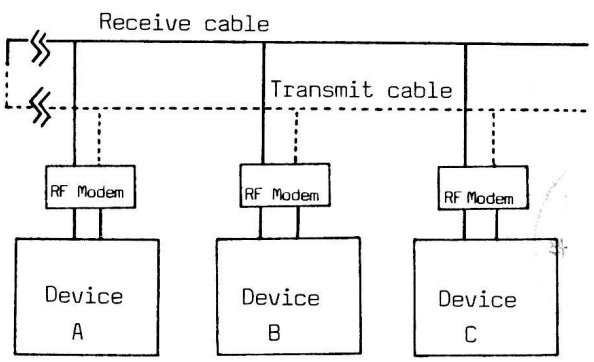

Figure 1.3 Broadband network

1.6 Token Passing

Token passing refers to the method in which nodes on a network gain access to the shared network cable. Usually this is in the form of a ring. Nodes are attached at

various points around the ring. Frames of data circulate around the ring, and are inserted and removed by the network nodes.

In a token passing ring, there is usually only one 'token' circulating. Without the token, no node can insert frames of data onto the ring.

To illustrate the operation of a token passing ring we will use the example of the token passing ring network IBM has chosen for its PC LAN network. In this case, the 24 bit token circulates around the ring when no data is being transmitted. When a node wishes to transmit data to another node, it 'captures' the token and marks it as busy. That node can then transmit a data frame onto the ring.

The frame of data carries with it an address which specifies the node that is the destination of the frame. When the destination node copies the data from the frame, it sets bits in the frame to indicate to the sender that it has correctly received the data that was in it. When the sender receives the modified frame again, it removes it from the ring and releases the token back onto the ring so that other nodes may transmit data.

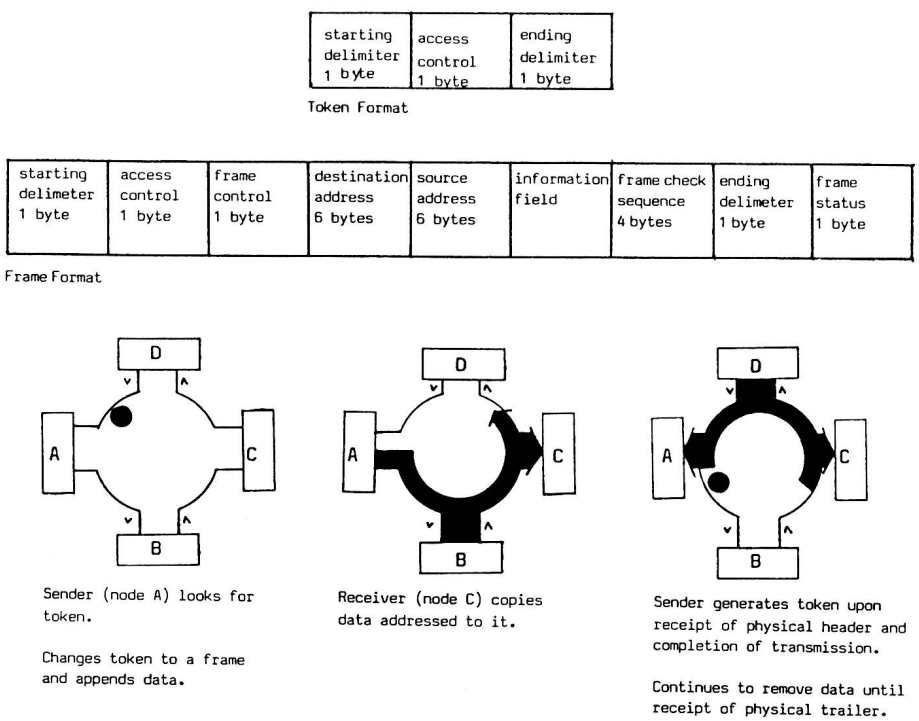

Figure 1.4 Token passing network

Figure 1.4 shows the structure of the token and data frames and gives an example of the operation of the token passing network. The IBM token passing ring operates at

试读结束：需要全本请在线购买： www.ertongbook.com