Computers and Telecommunications Networks

MICHAEL PURSER

BLACKWELL SCIENTIFIC PUBLICATIONS

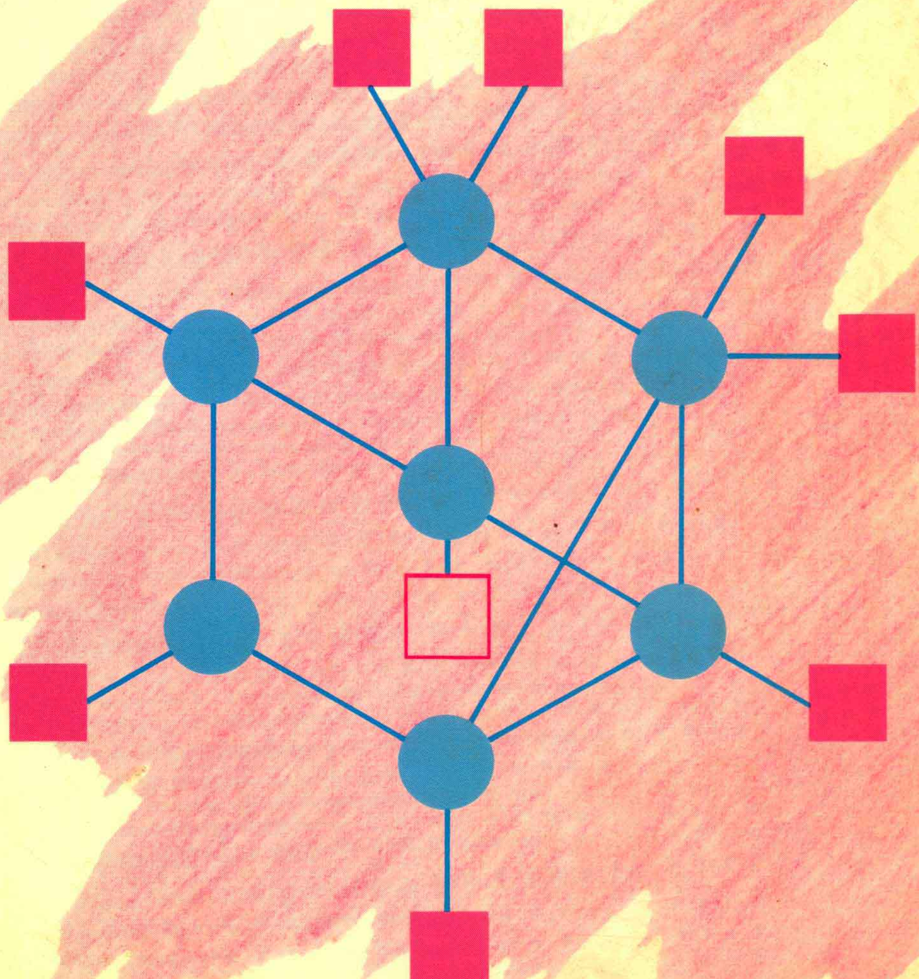

Computers and Telecommunications Networks

MICHAEL PURSER
Department of Computer Science
Trinity College
Dublin

BLACKWELL SCIENTIFIC PUBLICATIONS
OXFORD LONDON EDINBURGH
BOSTON PALO ALTO MELBOURNE

Editorial offices:
Osney Mead, Oxford, OX2 0EL
(*Orders:* Tel. (0865) 240201
8 John Street, London, WC1N 2ES
23 Ainslie Place, Edinburgh, EH3 6AJ
52 Beacon Street, Boston
Massachusetts 02108, USA
667 Lytton Avenue, Palo Alto
California 94301, USA
107 Barry Street, Carlton
Victoria 3053, Australia

First published 1987

Photoset by Enset (Photosetting), Midsomer Norton, Bath, Avon.
Printed and bound in Great Britain by Billing & Sons Limited, Worcester.

DISTRIBUTORS

USA and Canada
Blackwell Scientific Publications Inc
PO Box 50009, Palo Alto
California 94303
(*Orders:* Tel. (415) 965–4081)

Australia
Blackwell Scientific Publications (Australia) Pty Ltd
107 Barry Street,
Carlton, Victoria 3053
(*Orders:* Tel. (03) 347 0300)

British Library
Cataloguing in Publication Data

Purser, Michael
Computers and telecommunications networks.
1. Computer networks 2. Telecommunication
I. Title
621.38'0413 TK5105.5

ISBN 0-632-01765-1
ISBN 0-632-01767-8 Pbk

Library of Congress
Cataloging-in-Publication Data

Purser, Michael, 1937–
Computers and telecommunications networks.

Bibliography: p.
Includes index.
1. Telecommunication systems.
2. Data transmission systems.
3. Computer networks. I. Title.
TK5102.5.P87 1987 004.6 86-31727

ISBN 0-632-01647-7
ISBN 0-632-01648-5 (pbk.)

Computers and Telecommunications Networks

PROFESSIONAL AND INDUSTRIAL COMPUTING SERIES

Consulting Editors

DAN SIMPSON
BSc, FIMA, FBCS
Alvey Directorate

RICHARD VERYARD
MA, MSc
James Martin Associates

The aim of this series is to address various practical issues and topics in information technology. Modern computing involves a growing number of specialist areas; it is increasingly difficult to keep one's knowledge current. The emphasis is on providing up-to-date coverage of a subject, rather than theoretical novelty. These books should be widely read by practising professionals in data processing and information technology, as well as by students on vocational courses, in order to maintain and improve professional standards of computing in industry and commerce. The authors themselves are drawn either directly from industry or are academics with first-hand industrial/commercial experience.

Other titles in the series

VERYARD: *Pragmatic Data Analysis*
TAYLOR: *PICK for Users*
JENNINGS: *Practical Data Communications: Modems, Networks and Protocols*
WALSH: *Computer Users' Data Book*
SPENCER: *Computing Standards: A practical guide for data processing management*

Preface

This book addresses the topic of the 'convergence' between data processing and telecommunications. It is written largely from the standpoint of the computer person who finds himself more and more concerned with telecommunication systems, as the means of linking terminals in general and computers in particular, and as subjects of control and management by computers.

Telecommunication engineers have often been considered to be slow to grasp the implications of computers and the power and flexibility of control by software. However, this charge can hardly be levelled against them today, although it may well have been true in the past. On the contrary, it is the computer expert who is likely to be almost totally ignorant of well-established disciplines such as traffic theory; and who has very little appreciation of the complexity, sheer high speed and large-scale parallel processing which characterize, for example, digital switches. The book attempts to remedy this situation, by presenting networks of many sorts in a form comprehensible to computer people.

In Chapter 1 basic network concepts and traffic theory are presented. In Chapters 2 and 3 circuit-switched and store-and-forward (SAF) data networks are reviewed, respectively. The aim of these two chapters is to review and clarify basic switching techniques. Chapter 4 looks at network design and routing techniques. Chapter 5 specifically considers the area of convergence between data and other network services. The role of PBXs and LANs in interfacing to networks is reviewed, and Integrated Services Digital Networks (ISDNs) are introduced. Chapter 6 discusses the Open Systems Interconnection Reference Model (OSI/RM). It is presented late in the book so that the need for such an architectural framework (however defective) can be appreciated. It was felt that an earlier introduction might be tedious and incomprehensible. Chapter 6 terminates by discussing telematic and Message Handling Systems in the OSI/RM context. Finally, Chapter 7 looks towards future Integrated Broadband Communication Networks (IBCNs) and other developments such as the provision of security in networks.

Where reference is made to international standards, those of CCITT have been used preferentially. This is not to minimize the importance of ISO, ECMA and other standardization bodies. However, it is felt that in the real world of telecommunications, wherever the inspiration may come from, it is the CCITT standard which nearly always prevails. (A glossary is provided explaining the meaning of 'CCITT' 'OSI/RM' and other abbreviations, as well as several standard technical terms.)

The book presupposes a familiarity with basic data communication topics such as synchronous transmission and line procedures, but a detailed knowledge of such subjects is not necessary. It also presumes a certain command of mathematics in one or two places, but the reader who finds this disconcerting may skip them without significant loss.

It should also be pointed out that the book is concerned with 'active' switching performed by 'intelligent' networks which consciously route traffic. 'Passive' switching as used in many Local Area Networks (LANs) is assumed to be familiar to readers and is not described in detail.

Finally, it is my pleasure to be able to record the support and encouragement I have received whilst writing this book both from my wife and family and from students in Trinity College, Dublin. The book was typed by Helen Smith, to whom I owe a great debt for her efficiency, patience and good nature.

Michael Purser
Trinity College, Dublin
1987

Contents

Chapter 1
Computers and Networks

1.1 WHAT IS A NETWORK ?

A computer network can be defined in various ways. One approach is to consider the network to be simply the 'physical network'. The physical network may be said to consist of all internal switching nodes,* their interconnecting links and the links leading to externally connected devices. The external devices themselves, computers and terminals collectively referred to as Data Terminal Equipment (DTE), are then considered as attached to, rather than forming part of, the physical network (Fig. 1.1).

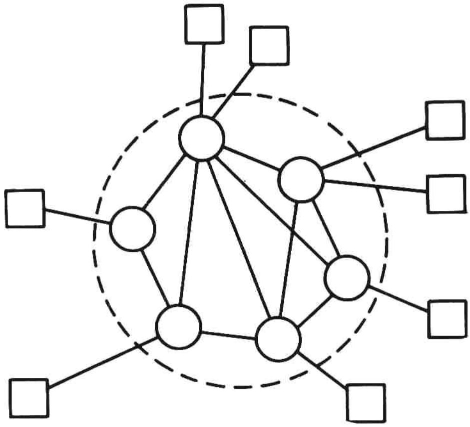

Fig. 1.1 A 'physical network' with switching nodes.

This minimal definition is convenient when we consider the typical control of a network, in which the physical network is operated by a Telecommunication Administration (TA, or more commonly 'PTT'), while the DTEs belong to users or subscribers.

*We call the points where links unite and traffic is switched 'Nodes' or 'Exchanges'. When the switching function is to be emphasized a Node is sometimes referred to as a 'Switch'.

Another, wider view would hold that the network consists not only of the physical network plus all attached DTEs, but of all the services available 'on' it—including those provided on attached DTEs. What is the network but the totality of all services that can be reached from any particular DTE?

Without considering further possible definitions, it is worth pointing out some of the problems with the above ones. For example, is a system for interconnecting many computers which contains no *switching* nodes a valid physical network? In such a system a sending DTE broadcasts a message, in which the selected destinations are identified over a common medium linking all DTEs; only the selected DTEs are supposed to receive it (Fig. 1.2). A minimal network of this type might connect only two DTEs. Again, is a 'star network' in which there is no switching, but all traffic is to or from one central DTE (Fig. 1.3) a valid physical network? In this book we shall not consider these networks which do not have switches, except superficially. However, they are generally regarded as valid networks, since many local area networks [1] and satellite-based networks use the broadcast-and-selected-destination technique; while the meshless (or netless) star network is one of the most familiar configurations for computer systems.

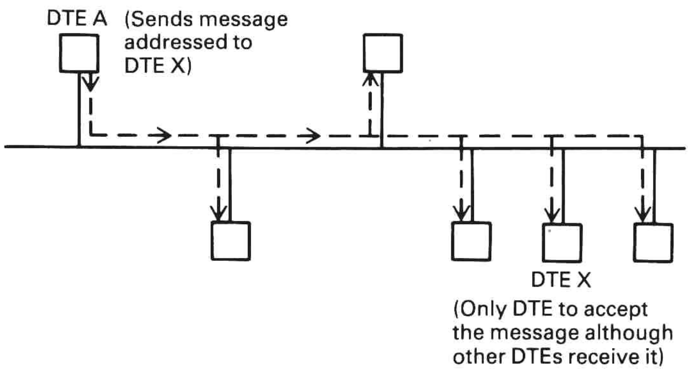

Fig. 1.2 A network without switching.

But, if we take the wide view that any system of interconnected computers and services is a network, then many networks apparently have no identifiable boundaries. They stretch worldwide; since from one

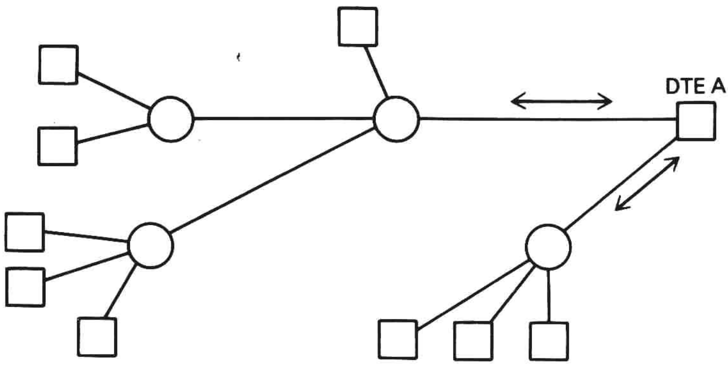

Fig. 1.3 A star network. All traffic in network is to/from DTE A.

terminal you can probably access literally millions of machines and services on the combined interconnected data, telephone, telex and other networks of the earth.

To avoid these problems we shall not attempt a precise definition of the term 'network'. However, the user should envisage a network being generally composed of the nodes and transmission links under the control of a single administration, using a more or less homogeneous technology; plus the DTEs immediately connected to it. When we wish to confine remarks to the switching and transmission portion we shall call it the physical or 'bearer' network; since it carries or bears the traffic, but is not, in general, a source or sink of traffic. When we want to emphasize the services (other than pure transmission and switching) available on or via the network, we shall explicitly use the term 'network services'.

1.1.1 Computers in networks

When the term 'computer network' is used, the common interpretation is that a network, in which the DTEs are computers, is meant. But in almost every network today the switching nodes are, if not computers, then largely controlled by computers. For the purposes of this book we regard these computer-based switches as comprising as important a subject of discussion as the attached DTEs.

Moreover, most services supplied 'on' the network, i.e. supplied by the administration operating the physical networks, are also based on computers. These computer-based services may be built into the nodes in software, or perhaps provided on an attached computer which is in every way like a DTE, except that it is owned and operated by the administration of the physical network and not by a subscriber (Fig. 1.4).

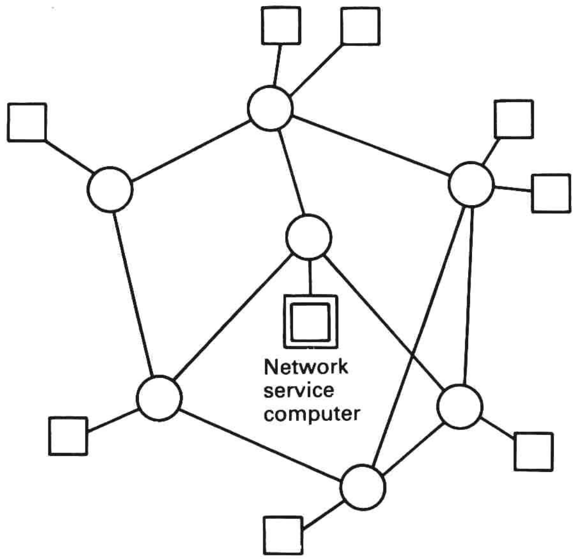

Fig. 1.4 A service 'on' the network provided on a special computer.

Yet another instance of computers which are not DTEs in networks is provided by the 'gateway'. A gateway is an interface between two distinct networks, in which some sort of conversion or mapping between the formats and functions of data valid for each of the two networks takes place (Fig. 1.5). The gateway enables traffic to travel usefully between the two distinct networks.

This book aims to discuss the use of computers in networks, with particular emphasis on networks used to carry data traffic, wherever those computers may be—inside or outside the physical network.

1.1.2 Different switching techniques

It has been stated that we are concerned with *switched* networks. A switched network is one in which traffic arriving on some link at a switching node is routed out on another (or possibly the same) link,

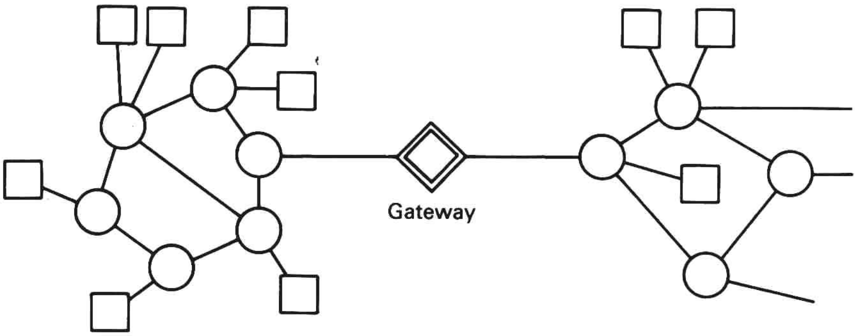

Fig. 1.5 A gateway between two networks.

under the active control of the node. If we consider a data-switching node, it makes the decision as to how to route the traffic on the basis of information normally carried in the traffic data. That information may be included in each data unit which has to be switched, like the address on the envelope of a letter in the post. Alternatively, the information may be supplied at the start of a data exchange across the network, usually in the form of special data units or commands to establish a 'call', in a manner analogous to the telephone network. Once established, a call is a conceptual pipeline through the node, through which all traffic associated with the call is routed. The call must, of course, be cleared sometime.

These two distinct ways of switching go by different names in different circumstances, but may be broadly classified as 'message-switching' (like a letter) or 'call-switching' (like a telephone call).

The ways in which data are handled (as opposed to routed) on passing through a switch also separate into two broad categories. In one technique, called 'circuit-switching', there is effectively no buffering in the node, incoming data are routed out directly. Since there is no buffering, if continuous data streams are to be handled successfully, the rate of input and the rate of output must be the same. Circuit-switching is usually a subclass of call-switching, since it is hardly feasible to dispense with buffering if each data unit has to be examined and routed independently, as it would be if message-switching were used. There is in effect a direct correspondence between the circuit (the physical pipeline) and the call (the conceptual pipeline). Alternatively, the 'store-and-forward' technique may be used for handling data traffic. In this case, a data unit is stored in the node before being forwarded. Queues of such stored data units can build up in the node, depending on the availability of an outgoing link for forwarding. The store-and-forward technique can be used with call-switching, notably in the case of virtual call packet-switched data networks. It is almost always used with message-switching.

Store-and-forward is the technique which most obviously suggests itself to those familiar with computers. It does however present an immediate problem. How much buffering capacity should be provided in a node and what is to be done if it threatens to become exhausted? This problem does not exist with circuit-switching, but instead a question of efficiency arises. The capacity of the circuit, the physical pipeline, is wasted if the user does not maintain a continuous flow of data.

A compromise switching technique is that of packet-switching. Packets are small blocks of data, usually of a fixed maximum size (e.g. 128 octets). If a real message is larger than this size it is divided into the requisite number of packets. The packets are handled by the store-and-foward technique. If message-switching is used (each packet is self-contained) we refer to a packet as a 'datagram'. If call-switching is used, we have a conceptual pipeline named a 'virtual call' (VC). By placing restrictions on the maximum packet-size the buffering problems associated with store-and-forward are somewhat eased. By choosing store-and-forward rather than circuit-switching the waste of capacity when handling 'bursty' non-continuous data flows is obviated.

VC packet-switched data networks (PSDNs) are favoured by PTTs, for reasons we shall discuss later, for public PSDNs. Datagram PSDNs are more acceptable to academics and researchers for their greater flexibility, and for whom the concept of a 'call' is not relevant to the bearer network, but only to the users involved in the communication.

Circuit-switching is discussed in Chapter 2, store-and-forward switching in Chapter 3. In Chapter 4 some topics applicable to both techniques are presented, and in Chapter 5 new types of networks are discussed, in particular Integrated Services Digital Networks.

1.1.3 Interfacing to networks

A switching network exists in its own right. It is not 'transparent' to users. If we consider the star network of Fig. 1.3 we can see that such a network can be largely transparent since control is vested in the central computer. All other DTEs should be compatible with it and its physical extensions, such as communication links. However, Fig. 1.2 shows that when there is no single central control, if arbitrary DTEs are to be able to exchange data meaningfully, then some standards (which may be usefully called 'network standards') should exist. Such standards would at least define how to connect physically to the network, and how destination addresses are to be represented in messages so that all potential recipients can recognize them.

Thus even a 'passive' network, such as that of Fig. 1.2, will require standards, defining the interface to the network, which are not necessarily those normally used by the individual DTEs in the absence of a network.

In the case of networks with active switching the need for network interfacing standards is even greater. This is because the network is no longer transparent to users, except at a considerably higher level of communication. To establish contact with a DTE *across* the network an originating DTE must talk *to* the network itself. Specifically the originating DTE must send information to the nodes themselves, so that they may actively route the data.

Thus the existence of a switched network implies the existence of interfacing standards for it. All computers or other DTEs requiring connection to such a network must be adapted, by suitable hardware and software means, to be compatible with those standards.

Among such interfacing standards are to be included the well-known X.25 [2] and ISDN [3] recommendations. They are international standards, recommended by CCITT for use with specific types of networks throughout the world. The need for an interfacing standard for a given network has been made clear, but why have an *international* standard for all networks of that type? The answer is based on economics. Since DTEs require special hardware and software to interface to a given network, there are capital costs incurred by network users, as well as current costs for transmitting or receiving data across the network. It is obviously advantageous to users and manufacturers of DTEs alike if those capital costs can be kept as low as possible, by ensuring that the interfacing hardware and software are as widely applicable as possible.

1.1.4 Internal and internetwork standards

Just as it is clear that interfacing standards are required for a given network, it is also clear that there must be internal standards. It would be ludicrous to have to define a different procedure for the exchange of data between each pair of interconnected nodes in a network. Some internal standards must apply to the use of links and to the construction and behaviour of the nodes or switches themselves. It is not clear, however, that these internal standards should be applicable to other networks besides the one under consideration.

Nevertheless, if distinct networks are to be interconnected usefully (Fig. 1.5), it is desirable that their internal standards should not be too dissimilar; otherwise the gateway becomes unnecessarily complicated. After all, seen from the outside, a collection of interconnected networks is just a single large network. The commands given by a DTE to this network (in accordance with the interfacing standards) have somehow to be amenable for interpretation by *all* the nodes which might receive them. This will certainly be facilitated if such commands are forwarded internally in the network in some relatively standardized form. Examples of functions needing such standardization are addressing and routing, which are relevant to the combined networks.

In practice, although international standards do not exist for all aspects of internodal communication, standards do exist for multiplexing, signalling and addressing. More specifically, there are international standards, such as X.75 [4], for the interconnection of networks belonging to different administrations.

1.1.5 End-to-end communication

The aim of the network is to make the interconnection of DTEs possible. This is pointless if the DTEs which have been connected are incapable of understanding the messages received from each other. There is a need for the standardization of end-to-end messages (that is, messages between DTEs) across the network. At this 'higher' level we are in a sense back to the situation discussed earlier in which a transparent network is available, provided that the lower level interfacing standards have been obeyed.

Firstly, these higher-level, end-to-end standards must address the problem of concealing or masking any residual problems not taken care of by the network. Such problems might include loss of messages by the network. The end-to-end standards should define means for the detection of, and recovery from, such network errors. Secondly, agreement must be reached on the method for establishing communication between 'entities' or 'processes' within the corresponding DTEs; not merely between the DTEs themselves. A remote user at a terminal may wish to access a particular service in a computer across the network. How is that service selected? Does he have to use a different mechanism for all such services in different computers, or is there an (internationally) standardized procedure?

Between corresponding processes there must then exist agreement on the format and the meaning of end-to-end messages. This stipulation is anything but trivial, if one considers the very wide range of messages one could envisage sending between computers. For example, one might send a job file to be run on one remote computer; or one might receive a colour display for a graphics terminal from another such computer.

This requirement for end-to-end standardization has been addressed by the International Standards Organisation (ISO), among others. It is discussed in the general framework of the Open Systems Interconnection Reference Model (OSI/RM) [5] in Chapter 6.

1.1.6 Services and applications

The mention of the need for standardization on the format and meaning of end-to-end messages leads naturally to the topic of services and applications. After all, end-to-end communication is not an end in itself; it exists to enable services to be used or applications to be performed. It could be argued that networks are only incidental to services or applications, but this is a simplification. For example, an application which involves many simultaneous participants, such as a real-time teleconference, presupposes some sort of networking infrastructure. More importantly, the boundaries between the bearer network and the services supplied by means of it are not always clear. To see this one need only consider the Telex network, where the DTEs are usually integrally supplied by the TA.

A more up-to-date example of services closely allied to the supporting network is provided by that of Electronic Mail in which documents, in coded electronic form, are exchanged or shared between users via a network. The Message Handling Systems (MHS) of the X.400 [6] recommendation are a concrete case. The end-to-end service enabling documents coded in facsimile, teletex or other formats, to be exchanged is associated with explicit functions built into the supporting network, such as the 'holding' of traffic for users who are not on-line.