LOCAL AREA NETWORKS

and their applications

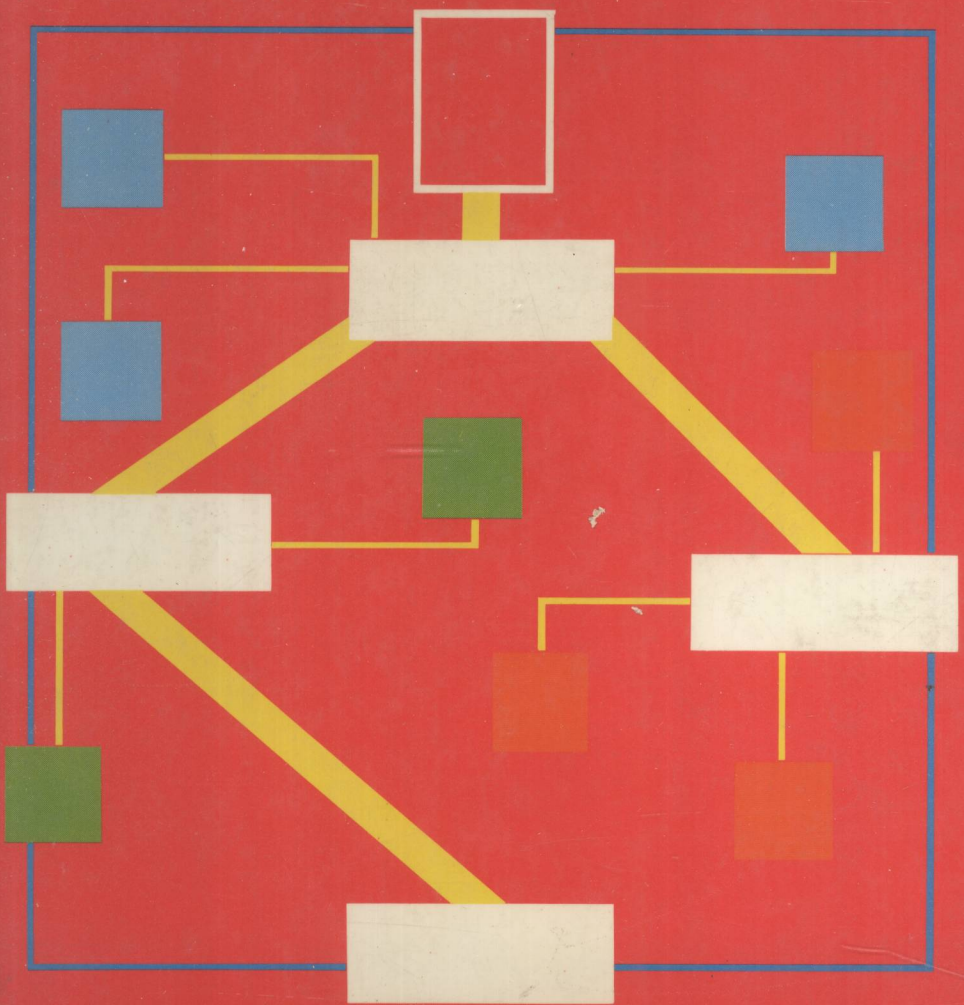

BRENDAN TANGNEY DONAL O'MAHONY

LOCAL AREA NETWORKS

and their applications

Brendan Tangney **Donal O'Mahony**
Trinity College, Dublin

PRENTICE HALL
New York London Toronto Sydney Tokyo

First published 1988 by
Prentice Hall International (UK) Ltd,
66 Wood Lane End, Hemel Hempstead,
Hertfordshire, HP2 4RG
A division of
Simon & Schuster International Group

Printed and bound in Great Britain by
A. Wheaton & Co. Ltd, Exeter

Library of Congress Cataloging-in-Publication Data

Tangney, Brendan, 1960–
Local area networks and their applications/
Brendan Tangney, Donal O'Mahony.
p. cm.
Bibliography: p.
Includes index.
ISBN 0-13-539578-X
1. Local area networks.
I. O'Mahony, Donal, 1961–
II. Title.
TK5105.7.T36 1987
004.4'8—dc19

British Library Cataloguing in Publication Data

Tangney, Brendan
Local Area Networks and Their Applicatons.
1. Local area networks (Computer networks)
I. Title II. O'Mahony, Donal
004.6'8 TK5105.7

ISBN 0-13-539578-X
ISBN 0-13-539560-7 (PBK)

1 2 3 4 5 92 91 90 89 88

PREFACE

In the early 1970s, a trend began in computing away from large centralized mainframes towards smaller departmental minicomputers. The first microprocessor system was introduced around 1971 and since then cheaper and more powerful machines have appeared on the market. This trend gathered momentum and has developed into a definite movement towards single-user workstations.

One of the disadvantages of this is that facilities that were available on larger, centralized systems were lacking in the new workstation approach. These included the ability to share information, send messages from one user to another, and, most importantly, to share access to peripheral devices, e.g. printers, disks and so on.

While computing devices were getting smaller, cheaper and more powerful, advances were also being made in communications technology. In particular a new generation of networks emerged which operated in a limited geographical area, but much more reliably and at higher data transmission rates than previous networks.

The introduction of these high speed **local area networks** (LANs) opened up the possibility of attaining the advantages of workstations without losing those of centralized mainframe systems. Information could flow between individual workstations at speeds which, to a large extent, hid the fact that they were not working on the same system.

This book examines these two aspects of local networks. Part One discusses the components that go to make up a local network. Part Two looks at the higher level software that is necessary to exploit the capabilities of a 'distributed' computing system, that is one composed of a number of individual machines connected by such a LAN.

The book is suitable for use as a higher level undergraduate and post-graduate textbook. The treatment of LANs is fairly detailed and the second section serves as an introduction to distributed computing. It will also provide valuable, up-to-date knowledge for computer professionals involved in planning and implementing LANs within their organization. No prior knowledge of networks is assumed, but a basic understanding of the field of data communications and computer operating systems will be helpful. The remainder of this preface explains the layout of the book in a little more detail.

Part One of the book begins by defining what exactly we mean by the term 'network', and how a 'local area' network is distinguished from other types. Chapters 2, 3 and 4 will cover the choices available in terms of communications media, the ways in which the machines are connected (topology), and how orderly transmission is carried out (the access method). As an illustration of the techniques covered, Chapter 5 will discuss how some of the better known LANs (e.g. Ethernet, IBM's Token Ring and the Cambridge Ring) are implemented using the previously described techniques. In Chapter 6 the concept of layered protocols is introduced and we examine some of the standards that have been developed in this area.

Part Two of the book concentrates on the details of LAN-based computer systems. It begins in Chapter 7 by analyzing what the essential differences are between running a centralized system and one made up of many interlinked machines. The concept of sharing resources between nodes on a network is introduced. Chapter 8 introduces the notion of the 'client/server' model, a widely used technique for organizing access to resources. Chapter 9 focuses on one very important type of server, namely 'file servers'. This is followed in Chapter 10 by an introduction to the concept of a 'distributed system', it explores some of the issues involved in building one, and looks at some example systems including CMDS, the Newcastle Connection and Locus. We complete the section by discussing how LANs can be used in the office and factory environments.

ACKNOWLEDGEMENTS

David Brownbridge, David Hutchison and Les Smith for comments on early versions of the manuscript. Maggie McDougall and Abigail Cooke for editing and production. Stephen and Neville for inspiration, our parents for opportunity. But above all, Tricia and Mary for patience, support, encouragement, etc.

CONTENTS

PART ONE

LOCAL AREA NETWORKS

CHAPTER ONE

WHAT IS A LAN?

A concise definition of what constitutes a local area network is difficult to produce. An easier approach is to first define what a computer network is, and then to distinguish LANs as a subset of these. When the term computer network was first used, it described any interconnections between computers. Since that time, three subclasses have emerged that are distinguished primarily by their geographical scope.

The first of these is the *wide area network* (WAN). This network spans a large area – possibly several continents. The second major type is the *local area network* (LAN) and, as the name suggests, it is confined to relatively small areas such as a building or a group of buildings, for example a university campus. A third type, which is just emerging at the time of writing, is the *metropolitan area network* (MAN). The scope of this class of network lies between LANs and WANs, i.e. spanning a small city or a town.

In naming these types of network, the main distinguishing factor would appear to be the size of the area covered. This factor has major effects on the technology used to implement the network, its administration and the type of applications that can be implemented on it. The ramifications of these factors will emerge in the course of this book.

1.1 WHAT IS A NETWORK?

Tanenbaum [1] defines a network as 'an interconnected collection of autonomous computers'. Two computers are said to be interconnected if they are capable of exchanging information. Central to this definition is the fact that the computers are *autonomous*. This means that no computer on the network can start, stop, or control another. This excludes from our definition any system where terminals or other peripheral devices are connected to a central host: for example, an IBM 370 configured with multiple 3270 terminals connected to it.

Another type of system that we exclude from our discussions is where two or more processors are connected together via either a common bus or an area of shared memory. While this may conform to the definition above, the problems involved differ significantly from other types of networks.

1.2 WORKING DEFINITION OF A LAN

Having defined what is meant by a network, the next task is to distinguish between LANs and WANs. LANs are different in the following important respects.

- The distance between the nodes is limited. There is an upper limit of approx 10 km, and a lower limit of 1 m.
- While WANs usually operate at speeds of less than 1 mega-bit per second, LANs normally operate at between 1 and 10 mbps. Using optical fiber technology, it is possible to achieve speeds of the order of hundreds of megabits per second.
- Because of the short distances involved, the error rates in LANs are much lower than in WANs. This extra reliability has an impact on both the protocols used in their operation and the range of applications that they can support.
- The distance limitations involved in LANs normally mean that the entire network is under the ownership and control of a single organization. This is in sharp contrast to WANs, where the network typically spans national boundaries. In this case, the network is normally operated by the country post and telecommunications authorities rather than by its users.

  This fact has no bearing on the performance of the LAN, although it does have a major impact on the way in which it is administered. The localized control of the facility greatly increases flexibility.

It can be seen from the above, that LANs differ from other types of network in that the area that they cover is limited. This means they can operate at high speeds and with very low error rates. These two properties are the main distinguishing features of LANs.

The first local networks were introduced into the academic world in the mid-1970s and, as the technology developed, they were adopted by the major computer manufacturers. By the early 1980s, most major hardware companies were offering LAN-based products. At the time of writing, the lower level aspects of LANs are starting to stabilize. Of the many approaches advocated by the researchers, a select few are becoming international standards, and are being adopted by computer users.

1.3 CLASSIFICATION OF LANS

There are many different network implementations that conform to the above criteria, each adopting a different approach. To distinguish between different

types of LAN, one usually examines three things: the media used to connect the nodes, the pattern of connection or *topology*, and the algorithms used to control access to the medium. The next three chapters will examine each of these in some detail before going on to see how companies and researchers have implemented them in Chapter 5.

1.4 REFERENCES

1 Tanenbaum, A.S., *Computer Networks*, Prentice Hall Inc., Englewood Cliffs, New Jersey, 1981.

CHAPTER TWO

TRANSMISSION MEDIA AND TECHNIQUES

As with any network, LANs must have a connecting medium of some sort to carry the information from node to node. Because wiring a building is a major expense to the potential LAN user, he must ensure that the system he uses fulfills certain criteria before embarking on the work. The medium chosen must be robust, that is, it must be immune to natural and environmental hazards, for example electrical noise, lightning and other forms of *electromagnetic interference* (EMI). For maximum availability of the network, the system must be serviceable. This means that mishaps such as cable breaks and node failures must be easily repairable or bypassed. The final cost of installation of the network will be affected by the cost of the medium, together with the cost of connecting each node to it.

Transmission speed is an important parameter of a network which will depend largely on the range of applications envisaged. These range from *terminal-to-host* traffic, for which speeds of below 1 kilobit per second are adequate, to applications such as graphic image and file transfer which need speeds of the order of megabits per second. The medium chosen will have to accommodate current as well as projected requirements over the system's expected lifetime of 15 years or more.

In order to appreciate the problems involved in medium selection, we must first look at the ways in which information is transmitted.

2.1 DATA TRANSMISSION TECHNIQUES

There are many ways of transmitting digital information through a medium. Making the choice between one technique and another is normally a question of striking a balance between performance, in terms of the speed and accuracy

of transmission, and cost. Certain parameters of the system are crucial in determining the former.

The normal way to transmit information through a medium is to vary an electrical signal at the transmitting end by some means, and detect these variations at the receiver. There are two major obstacles to successful reception: *attenuation* and *noise*. Noise can emanate from a variety of sources in the environment and serves to distort the signal. Attenuation is a measure of how much the strength of the signal is reduced in passing through the medium. It is proportional to the distance traveled and will be present to differing degrees depending on the frequency of the signal being transmitted. This dependency on frequency serves to further distort a signal as it passes through the medium.

For a particular medium, there will be a range of frequencies that can be transmitted through it without incurring significant attenuation. If the transmitter strays outside this range, reception will be difficult. In determining how much 'information' can be sent through the cable the most important aspect to consider is the width of this frequency range. This is known as the *bandwidth* of the medium.

Given that the transmission medium has a particular bandwidth, there are a variety of ways of transmitting information through it. In LANs, these are usually divided into *baseband* and *broadband* transmission. The details of these two methods will be described in the following sections.

2.1.1 Baseband transmission

When transmitting digital information, the simplest scheme is to have two voltage levels to represent 1s and 0s. Figure 2.1 shows how a number of bits have been encoded in this manner. More sophisticated schemes such as Manchester encoding (see Appendix A) can be used.

In baseband transmission, this voltage-encoded signal is applied directly to the medium. The signal is attenuated in its passage through the medium, causing the quality of the received signal to decrease with distance traveled.

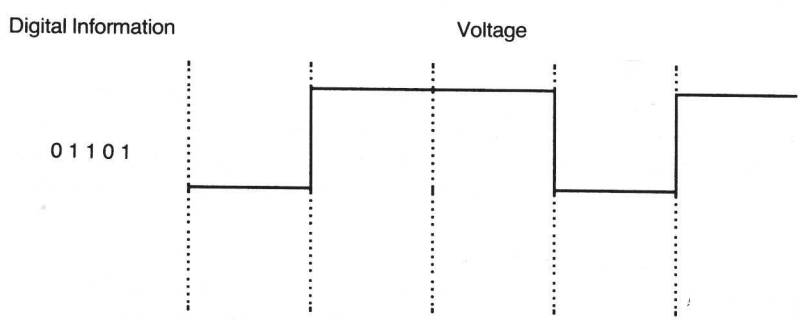

Figure 2.1 Digital information represented as a voltage