Preface

In large measure the traditional concern of communications engineers has been the conveyance of voice signals. The most prominent example is the telephone network, in which the techniques used for transmission multiplexing and switching have been designed for voice signals. However, one of the many effects of computers has been the growing volume of the sort of traffic that flows in networks composed of user terminals, processors, and peripherals. The characteristics of this data traffic and the associated performance requirements are quite different from those of voice traffic. These differences, coupled with burgeoning digital technology, have engendered a whole new set of approaches to multiplexing and switching this traffic. The new techniques are the province of what has been loosely called computer communications networks.

The subject of this book is the mathematical modeling and analysis of computer communications networks, that is to say, the multiplexing and switching techniques that have been developed for data traffic. The basis for many of the models that we shall consider is queueing theory, although a number of other disciplines are drawn on as well. The level at which this material is covered is that of a first-year graduate course. It is assumed that at the outset the student has had a good undergraduate course in probability and random processes of the sort that are more and more common among electrical engineering and computer science departments. (For the purpose of review, but not first introduction, the required background material is given in a pair of appendices.) The material in the text is developed from this starting point. The objective is to develop in the student the ability to model and analyze computer communication networks. We also seek to impart a critical appreciation of the literature in the field.

In a book at this level, it is inevitable that the choice of the particular material to be included in the text is heavily influenced by the author's own research and professional experience. However, as the book evolved, the work of others assumed a more and more prominent role. In fact, one of the most rewarding aspects of writing the book was learning and appreciating the fine work that others have done, and getting a sense of how the field has progressed. We hope that we have conveyed some of this to the reader. Our appreciation of a good deal of the material outside our immediate interests was gained to a significant degree through a number of first-rate survey papers in the field. We shall gratefully acknowledge these papers at the appropriate points in the text.

The level of the course for which this book is intended places certain limitations on the material that could be covered. In several instances important work could not be included simply because students could not be expected to have the appropriate mathematical training. In other cases the analysis was simply too involved. In spite of these difficulties, we derive great satisfaction from the fact that we were able to include almost all of the work which we consider to be seminal in the field. We tried very hard to present this work in a form that would be palatable to the serious though inexperienced student. A number of exercises illustrating the material are included.

Since the focus is on modeling and analysis, somewhat less attention is given to the details of implementation and operation. However, we attempted to include enough of this material so as to place the mathematical models in the proper context. In the course of the discussion we direct the reader to a number of authoritative sources on implementation and operation.

There is more material in the text than can be covered in an ordinary graduate course. In order to assist the instructor in selecting material, we point out several natural groupings. Chapters 1, 2, and 3 serve as an introduction to the rest of the text. After completing Chapter 3 one could go on to Chapter 4 or to Chapter 10. The main results in Chapters 4, 5, and 6 are obtained through the imbedded Markov technique, thereby providing a unifying theme for the three chapters. The results in Chapters 6, 7, 8, and 9 are pertinent to the modeling of a particular form of computer communication network—the local area network. In fact, all of the well-known accessing techniques are considered. In order to cover Chapters 6 through 9, some material on the imbedded Markov chain is also necessary.

Chapters 10, 11, and 12 also provide a convenient grouping. Chapters 10 and 11 rely heavily on Jackson network theory. The linkage between Chapters 11 and 12 is that they treat the so-called "higher-level" protocols, in particular flow control and routing. Although the final chapter on network layout stands somewhat apart from the others, there are points of contact between network layout and routing. This latter is the subject of Chapter 12.

In concluding these opening remarks we would like to express our deep gratitude to a number of people for their contributions to this work.

Jerry Foschini and John Silvester read an earlier version of the text and made a number of valuable suggestions for its improvement. Their diligent effort and encouragement were significant contributions. Thelma Hyland typed the manuscript. She was the soul of patience through the numerous revisions. A picture is worth a thousand words, and a thousand thanks are due to Miao Duen-Zhi for her assistance with the figures. The book evolved from notes of a course taught at McGill University over the period 1979–1983. A number of students suffered through the earlier (and later) versions of the notes. For their forbearance we express our deep gratitude. As mentioned earlier, the material in the book reflects our research. Our interest in computer communications began at Bell Labs, and was stimulated by Bob Lucky and Jack Salz. We would also like to acknowledge our first collaborator in the field, Dave Sherman. Finally, the author would like to pay tribute to his wife, Florence, and to his children, Mary, Ann, Jemmy, and Martin. They endured his absence and preoccupation during the long writing process.

Jeremiah F. Hayes
Concordia University
Montreal, Quebec

Contents

2.7. Local Area Networks

2.7.1. General Considerations

The second major class of networks that we shall consider are *local area networks* (LANs), which are data networks having limited geographical area, usually within a kilometer.† Networks confined to a single office building, shopping center, or university campus are prime examples of LANs. The emergence of this sort of network is part of the general growth of computer and digital technology; however, the introduction of office automation and distributed processing systems has furnished additional impetus. In both of these applications LAN techniques play a significant part.

In the past, local area networks were defined in terms of geographical extent and data rate; however, in view of the rapid growth of the technology, a more useful definition may be in terms of usage and configuration. In providing a common channel for a number of users in a limited geographical area the emphasis is upon ease and flexibility in providing access. Due to the limited geographical area, bandwidth is not the critical commodity that it is in networks covering a wide area. Therefore, access to the network can be simplified at the cost of bandwidth. A second point is that data networks covering a large area require redundancy with respect to connectivity in order to ensure operation in the face of failures. For example, the ARPA net requires at least two paths between source–destination pairs. Because of the limited geographical extent of the typical LAN, it is in something of a protected environment and this sort of redundancy is unnecessary. This simplifies the topology since only a single path need be provided between source–destination pairs.

2.7.2. Topology

In current practice three basic topologies are prevalent in local area networks: the bus, the ring, and the star (see Figures 2.20a–2.20c). A fourth topology, the tree (see Figure 2.20d), is used in related systems and may be considered as a form of the bus configuration.

The bus topology is appropriate to transmission media such as coaxial cable or radio which allow what are, in effect, high-impedance taps. In principle, these taps do not affect the medium, and a large number of stations can be connected. Each of these stations can broadcast simultaneously to the others. The bus topology is particularly appropriate for the random accessing techniques that we shall be discussing presently.

†Surveys of LAN techniques are contained in Refs. 41–43. See also Ref. 44.

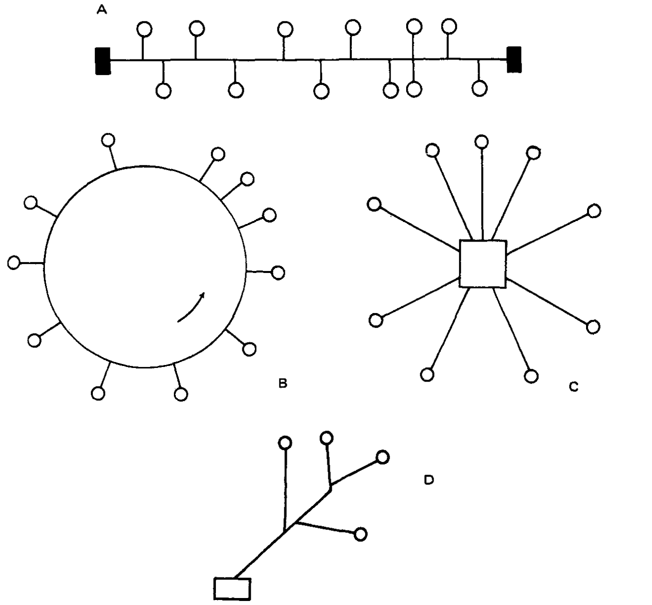

Figure 2.20. (A)–(D) Ring, bus, star, and tree configurations.

The ring topology is a sequence of point-to-point links with flow in one direction around the ring. As we shall see, there are several accessing techniques for the ring system. In all of these there is a delay due to processing at each of the stations. For reasons of reliability, there are provisions to bypass stations if they become inoperative.

In both the bus and the ring topologies the control of traffic is distributed. A topology in which control is concentrated is the star. In this case all of the traffic in the system is switched in a central "hub." The stations access the hub through high-speed lines.

A fourth topology, the tree, has been used to distribute data over a wider area. Typically, in networks using the tree topology, remote stations access a central processor. A ubiquitous example of such a network uses multipoint private lines in common carrier networks. The tree topology would also be relevant for CATV networks used to transmit data, since these sorts of networks are closely related to LANs and the same modeling techniques are relevant.

The four topologies can be the basic building blocks for more complex topologies. For example, a system of interconnected rings has been proposed. In the current standard involving the bus topology a series of connected buses form a kind of rootless tree.

2.7.3. Transmission Media

In the foregoing discussion of topology we alluded to the transmission media used to implement the network. In current practice three kinds of transmission media are used in LANs: twisted pairs, coaxial cable, and optical fiber. The twisted pairs of copper wires are the same as those used in the telephone plant. Because of the readily available technology, there is a strong tendency to operate at the same rates as short-haul digital carrier systems, i.e., the T1 rate 1.544 Mbps. Higher speeds are possible as well. The nature of twisted pairs is such that it is basically a point-to-point medium. As mentioned earlier, coaxial cable can operate as a multiple access medium by use of high impedance taps. Rates of up to 10 Mbps are attained in commercial systems. Optical fiber is the medium with the highest transmission rate currently used in LANs. Data can be transmitted at rates of up to 50 Mbps using light emitting diodes. With Laser transmitters, rates in the range of hundreds of Mbps can be achieved without repeaters over distances compatable with LAN operation. At the present writing optical fiber is basically a point-to-point medium. However, the development of a low-loss optical tap could change this situation.

In the future other transmission media may play a role in LAN and related systems. Radio, for example, would allow more flexible operation than the media we have considered. A radio medium would allow a random accessing technique to be employed. A second widely available medium which could be adapted for two-way data transmission is the CATV network. We shall be considering this medium in some detail in the sequel.

2.7.4. Access Protocols

The fundamental purpose of the local area network is to allow data sources that are dispersed throughout an area to share a common transmission medium. Because of this dispersal, transmission capacity must be expended to coordinate the flow of traffic from each of the sources. The way that this is done is the function of the access protocol. A significant portion of this book is devoted to the modeling and analysis of these access protocols. The objective of the analysis is to evaluate performance in terms of delay, storage, and throughput.

2.7.4.1. Polling

The predecessors to the current local area networks were tree networks,† which provide communications between a number of remote stations and a central processor. Such an arrangement would be appropriate to a bank's credit checking system, for example. The coordination of traffic from the stations is effected by roll call polling. Each station is assigned an address. The central processor polls stations for messages by broadcasting the addresses in sequence. The reception by a station of its address is, in effect, a license to transmit. If a station has a message to transmit, it interrupts the polling cycle.

Polling systems are modeled and analyzed in Chapter 7 of the text. The salient result of this analysis is the effect of overhead on performance. Overhead in this case is the time required to poll all of the stations even if none of the stations has a message to transmit. In many applications a significant portion of the overhead is due to the establishment of communications between the remote stations and the central processor. In the case of voiceband modems, for example, this would require phase and timing recovery and equalizer training for the higher speeds.

An alternate technique, called hub polling, attempts to reduce this overhead. In this case the license to transmit is transferred between stations directly without going through a central processor. This technique is similar to those used in ring systems, which we shall be discussing in the next section.

2.7.4.2. Ring Protocols

From the perspective of current techniques, the first local area network was the ring system built by Farmer and Newhall.[45] This system is shown in Figure 2.21.‡ The flow of data around the ring is organized into the frame

†In Chapter 13 we shall be considering the layout of tree networks.
‡A survey of LANs with the ring topology is given in Ref. 46.

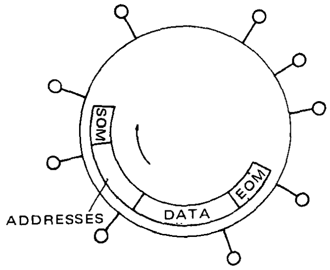

Figure 2.21. Farmer–Newhall loop.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | sales@docketalarm.com | 1-866-77-FASTCASE

fastcase® Smarter legal research.