Chapter 1

Introduction

For many years voice telephony remained the dominant service supported by the telecommunications networks throughout the world. To support a single class of traffic only a single switching mechanism was required: circuit switching, which was well suited to the characteristics of voice traffic. Within the past twenty years the growth of the computer industry has led to the requirement that telecommunications networks offer computer communications services, originally aimed mainly at file transfer and terminal to host interconnection. Although the circuit switched telephony network has been used to support such applications it has proved inefficient and inadequate thus we have seen the development of packet switched networks to support computer communications traffic. Recently, with the rapid drop in the cost of computer technology has come the requirement to support many more communications services. Examples include voice, high quality audio, broadcast quality video, compressed video, image, facsimile, text and many forms of data transfer. The existing practice of providing a separate network, with its own switching mechanism, for every class of traffic cannot be extended to support the demand forecast for many new communications services. The most flexible solution is that of an integrated communications network with a single switching mechanism capable of handling all classes of traffic. Fast packet switching has been suggested as a possible switching mechanism to support integrated services.

If fast packet switching is to support the growth in demand for communications services for the foreseeable future, two major problems require attention. First the design and implementation of a fast packet switch must be investigated, capable of expansion from small sizes of switch up to structures of very high capacity, to handle the expected growth in traffic (especially if video services are contemplated). Secondly the switch must be capable of satisfying the delay requirements of the delay sensitive classes of multi-service traffic. These form the two major issues addressed in this dissertation.

1.1 Objectives

There are many possible approaches to the design of a fast packet switch. The work presented in this dissertation investigates a design which leads to a very simple hardware implementation. A simple implementation offers flexibility, a wide range of potential applications and operation at both conventional speeds and also at possibly very high speeds. The design features a number of parameters that have an effect upon the performance of the switch. Several techniques are also described to enhance the performance beyond that of the basic design. The first objective is thus to characterise the effect of the various parameters upon the performance of the switch. This allows the selection of the preferred design parameters.

The second objective is to characterise the performance of the switch for multi-service traffic. This objective has proven more difficult to satisfy mainly because of the problem of adequately defining the characteristics of multi-service traffic. The performance of the switch for telephony voice traffic, which is well characterised, has been investigated in detail. A simple model of multi-service traffic has also been used to investigate a statistical upper bound on the delay performance of the switch. The effect of packet length on the delay and throughput performance has been measured and some observations made on the packet loss probability due to buffer overload.

A third objective of this work has been the implementation of the two fundamental components of the switch design in current gate array technology. This provides an insight into the complexity of the design and its suitability for implementation in the various available logic families. It also lays the foundation for the development of an experimental model of the switch which will be required for future experimental work in the use of fast packet switching techniques for communications applications.

1.2 Outline

For a concise presentation of the design of the Cambridge Fast Packet Switch and of the major results discussed in detail in this dissertation the reader is referred to [118] or [117]. The work has also been presented at the conference "IEEE Infocom '88" [116] and at the 'European Telecommunications Workshop' [115].

The remainder of this chapter presents an introduction to the growth of telecommunications networks and discusses the requirement for the integrated support of multi-service traffic. Chapter two considers the switching mechanisms capable of supporting this requirement within a high capacity switch implementation and presents the argument for selecting fast packet switching. Some of the fundamental characteristics of a fast packet switch are also discussed. In chapter three a simple classification of fast packet switch designs is introduced followed by a review of many of the major designs of fast packet switch available in the literature. Many of these designs are based upon the use of a multi-stage interconnection network which also forms a central feature of the design of the Cambridge Fast Packet Switch. Chapter four therefore presents a review of interconnection networks concentrating upon those most relevant to the design of a fast packet switch. The design of the Cambridge Fast Packet Switch is presented and discussed in chapter five. Chapter six presents the results of a simulation study of the performance of the switch fabric in order to gain an insight into the effect of the various switch parameters on performance and to select the appropriate switch design parameters. Chapter seven introduces the requirements that multi-service traffic imposes upon the switch and presents the simulation results of various aspects of a simple model of multi-service traffic applied to the switch. A detailed simulation study of the switch performance for telephony voice traffic in the presence of saturated data traffic is also investigated as a specific example of multi-service traffic of practical interest. In chapter eight the details of the experimental hardware implementation of the two major components of the switch design in current gate technology are presented. Performance measurements of the hardware model are compared to the simulation results and the extension of the model to a full-scale switch implementation is discussed. Finally, chapter nine summarises the insight gained from this study, introduces some ideas for the support of multicast traffic, and discusses some of the problems involved in the networking of fast packet switches that remain for further study.

1.3 Growth of Communications Networks

Local Area Networks

A local area network (LAN) connects computers, workstations, terminals, printers and related peripheral equipment across a distance of up to approximately 1 km [72]. Although research into local area network design began over ten years ago, to some extent stimulated by distributed computing applications [109], major commercial exploitation has developed over the last five years or so. As the cost of computer based equipment has fallen, so the requirement for interconnection within the local area has grown with the major area of commercial application being office automation. The vast majority of traffic currently carried on commercial LANs consists of file transfer and interactive data traffic but much research has been aimed at supporting other classes of traffic on the LAN especially voice traffic [121, 106, 38, 144, 3]. The local area network is privately owned and maintained which has encouraged work aimed at the integration of the LAN and the private telephone exchange (PABX) with varying levels of success [42, 151, 152, 46].

The traffic capacity of current commercial LANs is in the region 1 – 10 Mbits/sec. With the geographical constraint of about 1 km, LANs must be interconnected via bridges to extend the capacity and area of coverage. Work on the transparent interconnection of LANs via bridges is well advanced [10, 133]. The extent to which the traffic capacity may be increased by bridging between multiple LAN segments is limited. Thus to interconnect large numbers of LANs and to support emerging wideband services, such as image and video, high speed local area networks are being developed for use as a backbone network [68, 67, 90, 129, 130]. These networks offer a bandwidth of the order of 100 Mbits/sec. This represents a considerable bandwidth for conventional computer communications applications but if the cost of workstations that support graphics and image applications continues to fall, and with the possible growth of video applications, a requirement to interconnect such high speed networks may develop. Unless traffic is highly localised, the interconnection of high speed LANs in a mesh topology, using simple bridges, will not greatly increase the capacity of the overall network. Thus the use of a fast packet switch as a high capacity multi-port bridge, to support the interconnection of high speed LANs, (and also as a high speed local area network itself,) forms a possible area of application for the use of fast packet switching technology. Furthermore, the high speed LANs that require interconnection need not be local to each other. Thus fast packet switching provides a mechanism whereby widely separated local area networks may be interconnected, at high capacity, to give the impression of a single virtual LAN spanning a very large area.

Metropolitan Area Networks

The metropolitan area network (MAN) is a public network which may be considered as an extension of the high speed local area network to encompass an urban area with a diameter of up to about 50 km [81]. It must also be capable of supporting multi-service traffic with reasonable efficiency and particularly the voice service. A number of network designs have been proposed [138, 62, 98], but the design most likely to be selected for public service is based on a dual bus arranged as a loop. It has the property that access to the network emulates the action of a single queue although sources are distributed across the network [103, 119]. The integration of multi-service traffic is currently proposed in a hybrid manner. Integration is achieved at the access and transmission level but separate circuit and packet switches are used for compatibility with the existing digital telephony network. Packet traffic, however, is fragmented and transported within short, fixed length packets which permits evolution to the support of multi-service traffic upon a single integrated switching mechanism should this prove desirable.

A single network segment cannot possibly support the evolving needs of an entire urban community thus many segments must be interconnected by means of a switch [29]. Thus as the demand for packet based communications traffic grows we find another application for a fast packet switch that offers the flexibility to support growth to a very high capacity and the ability to handle multi-service traffic.

Public Wide Area Networks

For many years the public telephone network has been evolving from analogue to digital transmission and switching techniques [127]. In the developed countries penetration of digital techniques into the trunk transmission and switching network is now very high and attention is being focussed upon the local telephone network. Digital access from the subscriber to a digital local exchange is forecast to encourage the development of new telecommunications services at an acceptable cost. The integrated services digital network (ISDN) [125] promises to provide integrated access through a common, standard interface to both circuit and packet switched networks. The circuit switched channels will offer a bandwidth of 64 kbits/sec and access to the packet network will be at up to 64 kbits/sec. The basic rate interface will offer two 64 kbits/sec circuit switched channels with a 16 kbits/sec packet switched signalling channel. In Europe, primary rate access will offer 30 circuit switched channels with a 64 kbits/sec packet switched signalling channel. Wideband access to (N×64 kbits/sec) channels is being considered for later implementation and standards have been developed for packet based user-network signalling.

The philosophy of the ISDN is to employ the existing copper cable between the subscriber and the local exchange for digital access. The possibility of gradually replacing the copper connection with optical fibre is currently under consideration with the opportunity of increasing the bandwidth between the subscriber and the local exchange to many hundreds of Mbits/sec. Such a network is referred to as the broadband ISDN (B-ISDN) [157, 91] and may be required to support services such as: video telephony; image, video and hi-fi audio retrieval services; and the distribution of high definition television [135, 9]. The economic feasibility of the evolution to a widespread broadband network is at present uncertain. However, active consideration is being given to developments in the switching and transmission technology that would be required.

Two approaches have been proposed for the realisation of the broadband ISDN: synchronous transfer mode (STM) and asynchronous transfer mode (ATM), also called new transfer mode [101]. STM is an extension of traditional circuit switching principles and provides channels of fixed bandwidth with a packet switched signalling mechanism. ATM, however, is based upon a fast packet switching mechanism which can provide channels with a bandwidth that is highly variable throughout the lifetime of a connection. In the parlance of broadband ISDN the term 'fast packet switch' implies the use of variable length packets, whereas fast packet switching with short, fixed length packets is called asynchronous time division (ATD). The Cambridge Fast Packet Switch is equally suited to handling both short, fixed length packets and variable length packets of any reasonable size without loss of efficiency. Thus the more general term 'fast packet switching' is used in this dissertation to cover both applications.

ATM offers the major advantage of flexibility over the STM approach. The traffic characteristics of future service requirements cannot be predicted in advance thus the more flexible the network, the easier it becomes for a network administration to offer new services. It is not certain, however, that all services may be supported across an ATM switching mechanism, e.g. distribution video. Also, ATM technology may not be fully developed within the timescale which may be required by early broadband implementations. Hybrid solutions have therefore also been proposed in which both STM and ATM are supported over the same access interface and transmission link [161]. One solution proposes STM channels at about 150 Mbits/sec for the distribution of entertainment video services with a full-duplex ATM channel also at 150 Mbits/sec for all other services [91]. Broadband ISDN therefore forms a further application of fast packet switching techniques with the dual requirement of very high capacity and the ability to support a large number of services.

Integration

Three distinct levels of integration may be recognised in the evolution of telecommunications networks: access, transmission and switching. The current ISDN provides integrated access to circuit and packet switching networks that remain totally separate. Thus no sharing of resources between the networks is possible, two networks must be separately maintained and the support of future services is limited by the characteristics of the individual networks. Integration at the transmission level continues to require separate circuit and packet switches but the bandwidth of the transmission links connecting the switches is shared dynamically between the two switching mechanisms. Switches of this nature are referred to as hybrid switches and much work has been done on the integrated transmission of circuit voice and packet data services [32, 48, 93, 155]. The circuit switched component offers low delay and low variance of delay which is a requirement of the voice service in current public networks [102] while the packet switched component handles bursty services.

Only when a single switching mechanism handles all classes of traffic is integration at the switching level achieved [146]. Such a network offers integration of access, transmission and switching and may be considered fully integrated. The greatest advantage of full integration is the flexibility to adapt quickly to the changing traffic requirements of new communications services. Other advantages include transmission efficiency, independence of the switching mechanism from the characteristics of the source traffic, and the need to support and maintain only a single integrated network. Fast packet switching offers one possible solution for a fully integrated network.

1.4 Multi-Service Traffic

A simple classification of multi-service traffic is presented in table 1.1 which is adapted from [84] and [83] which itself reflects current CCITT[1] thinking. The natural rate indicates the source bit rate of the traffic class and in some cases assumes a certain amount of compression to reduce redundancy in the signal. Some sources emit traffic continuously at a single bit rate but many exhibit bursty behaviour in which traffic is emitted in bursts interspersed with idle periods. The burstiness of a source is expressed as the ratio of the peak to average bit rates. If the communications channel is fast enough to avoid being a bottleneck then most forms of data traffic become bursty due to user behaviour and the need to share processors amongst applications. Most forms of non-data traffic are also bursty if coded by efficient signal processing technology as the information content of the signal varies with time. A further parameter of the source traffic is the holding time of a connection. The relationship between the range of bit rates and holding times for various classes of traffic and switching

[1] The International Telegraph and Telephone Consultative Committee.

Service Class	*Service*	*Natural Rate bits/sec*	*Burstiness*
Conversation	Telephony	4–64k	2–3
	Video Telephony	2–34M	1–5
	Interactive Data	1k–1M	>10
	Telemetry	<10k	>10
Mail	Voice Mail	4–64k	2–3
	Video Mail	2–34M	2–3
	Text	1k–1M	1–10
	Facsimile	10k–1M	1–10
	Mixed Mode	100k–10M	1–10
File Transfer	Bulk Data	1M	1–10
	Program Download	1M	1–10
	CAD/CAM	1–40M	1–10
Retrieval	Hi-Fi Audio	1–2M	(2)
	Video	2–34M	2–3
	Mixed Mode Document	100k–10M	1–10
	Data	1M	1–10

Table 1.1: Multi-service traffic characteristics.

mechanisms is illustrated in fig. 1.1 which is taken from [91].

Each class of source traffic also exerts various requirements on the performance of the communications network: set-up delay, bit error rate and information delay. The set-up delay is the time required to establish a connection across the network. Estimates of the bit error rate required to support the various services vary widely but for a data service a residual error rate of better than 10^{-12} may be required. For the various real-time services such as voice and video the delay requirements may not permit the use of an error detection and correction protocol but due to the redundancy of the signal a higher bit error rate may be tolerated. The information delay is the delay requirement from source to destination across the network of which several measures may be significant: mean, jitter and percentile. Some services may tolerate high and variable delay across the network. Other services, however, such as telephony, place stringent requirements upon the upper bound of delay and the delay jitter, (variance of delay across the network.) These must be maintained throughout the duration of the connection else the establishment of the connection should be refused. This requires that the network be aware of the bandwidth required by a connection request and be capable of ensuring that this bandwidth is available before granting the connection. This problem will be considered in greater detail in chapter 7. Further discussion of multi-service traffic and of the services under consideration for support by the broadband ISDN may be found in [135, 9, 8, 92, 156, 142].

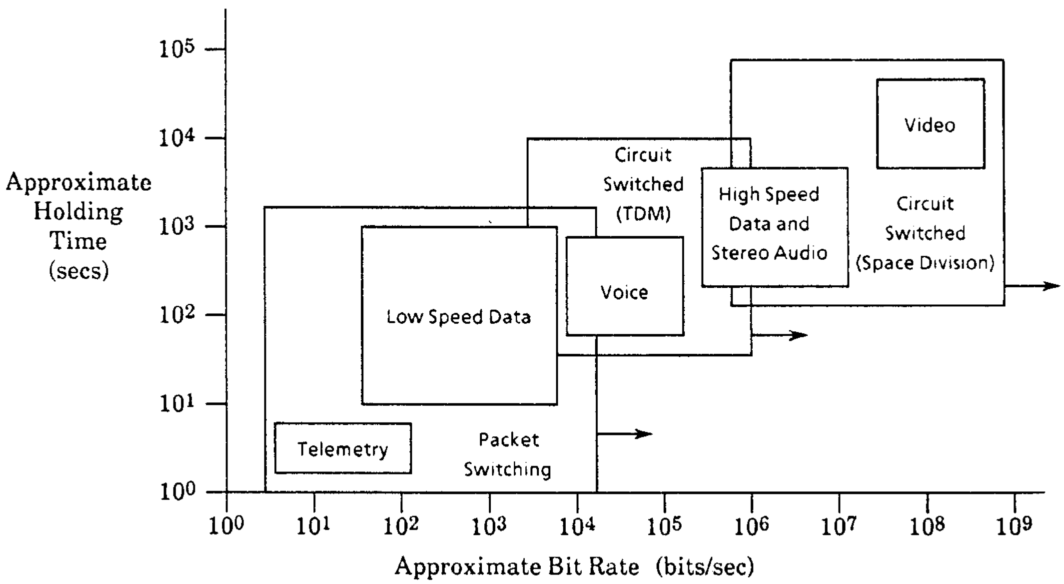

Figure 1.1: Relationship between bit rate and holding time for various classes of traffic and switching mechanisms.