A Slotted Ring Test Bed for the Study of ATM Network Congestion Management

by

Phil Holifield

Thesis submitted to the University of Central Lancashire in

partial fulfilment of the requirements for the degree of

Doctor of Philosophy

February, 1995

The work presented in this thesis was carried out in the Department of Electrical and Electronic Engineering, University of Central Lancashire, Preston, England, in collaboration with the School of Engineering, University of Durham, Durham, England.

6. CONCLUSIONS AND FURTHER WORK

The slotted ring has clear benefits for the transmission of real-time traffic using the ATM. The ORWELL protocol delivers fairness between transmitting stations but at an overhead related to the reset mechanism, and the Di-allocation largely determining this overhead. The Reset Interval, measured at each network station, can provide an indication of the level of traffic on the network, but it is affected by the traffic distributions on the ring. It has also been shown [4] that the statistical distribution of traffic carried by the network also influences the Reset Interval. The slotted ring LAN can be operated without a reset mechanism, classifying slots only as full or empty, in which case a suitable connection acceptance mechanism must be found. Measuring the numbers of full or empty slots passing a network station is being investigated as a potential indicator of network loading and of effective available bandwidth at the station.

7. REFERENCES

1. G. M. Woodruff, R. Kositpaiboon. "Multimedia Traffic Management Principles for Guaranteed ATM network performance," IEEE Journal on Selected Areas in Communications, Vol.8, No.3, pp.437-446, April 1990.

2. J. Chauhan, T. J. King, A. C. Micalef. "Specification of the Orwell protocol", British Telecom Technical Publications, May 1990.

3. R. M. Falconer, J. L. Adams. "ORWELL: A Protocol for an Integrated Services Local Network," British Telecom Technology Journal Vol.3, No.4, pp.27-35, October 1985.

4. P. Holifield. "A Slotted Ring Testbed for the study of ATM Local Area Network Congestion," PhD thesis under preparation, University of Central Lancashire, 1994.

Declaration

I declare that while registered with the University of Central Lancashire for the degree of Doctor of Philosophy I have not been a registered candidate or enrolled student for another award of the University of Central Lancashire or any other academic or professional institution during the research programme. No portion of the work referred to in this thesis has been submitted in support of any application for another degree or qualification of any other University or Institution of learning.

Signed ....P. Holifield..........

Table of Contents

Acknowledgements

Firstly I would like to thank Professor T. J. Terrell, my Director of Studies, for his support and guidance through all stages of this research project, to which he has generously devoted much time.

I would like to thank my second supervisor Professor P. Mars of the School of Engineering at Durham University for the technical guidance and direction he has given to the project.

Thanks are due to Mr. A. Stewart for his assistance in the manufacture of equipment, and his helpful advice on many technical matters.

I would also like to take this opportunity to thank Professor N.G. Burrow, Head of Department of Electrical and Electronic Engineering at the University of Central Lancashire, who has made resources available as required throughout the duration of this research.

Dedication

To my parents for their encouragement in all things academic

Abstract

A Slotted Ring Test Bed for the Study of

ATM Network Congestion Management

by

Phil Holifield

This thesis addresses issues raised by the proposed Broadband Integrated Services Digital Network which will provide a flexible combination of integrated services traffic through its cell-based Asynchronous Transport Mode (ATM). The introduction of a cell-based, connection-oriented, transport mode brings with it new technical challenges for network management. The routing of cells, their service at switching centres, and problems of cell congestion not encountered in the existing network, are some of the key issues.

The thesis describes the development of a hardware slotted ring testbed for the investigation of congestion management in an ATM network. The testbed is designed to incorporate a modified form of the ORWELL protocol to control media access. The media access protocol is analysed to give a model for maximum throughput and reset interval under various traffic distributions. The results from the models are compared with measurements carried out on the testbed, where cell arrival statistics are also varied. It is shown that the maximum throughput of the testbed is dependent on both traffic distribution and cell arrival statistics.

The testbed is used for investigations in a heterogeneous traffic environment where two classes of traffic with different cell arrival statistics and quality of service requirements are defined. The effect of prioritisation, media access protocol, traffic intensity, and traffic source statistics were investigated by determining an Admissible Load Region (ALR) for a network station. Conclusions drawn from this work suggest that there are many problems associated with the reliable definition of an ALR because of the number of variable parameters which could shift the ALR boundary. A suggested direction for further work is to explore bandwidth reservation and the concept of equivalent capacity of a connection, and how this can be linked to source control parameters.

Chapter 1

Introduction

1.1 Background

Demands on the telecommunications network are increasing at great pace. The information technology revolution of the nineteen eighties has resulted in an explosion of software applications which extends from CAD tools running on high performance workstations to the word processors and spreadsheets found in many homes, and used by people of all ages. The development of multimedia authoring packages, CD ROM drives, and the processing power of a new generation of personal computers means that the impact of information technology on the lives of all members of society is set to increase. The flexibility of the present telecommunications network to carry the large amounts of data involved in high bandwidth video applications, and to efficiently combine this in multi-media applications involving voice traffic, and computer data traffic, is very limited.

Technology in the field of telecommunications has not stood still, and the transmission capacity of optical fibres means that giga-bit per second data rates can be achieved on a single fibre. VLSI manufacturing using CMOS devices now allows fast switching of digital information in telecommunications switching centres. These developments have been brought together in the introduction of a global telecommunications network known as Integrated Services Digital Network (ISDN) which is significant because it has for the first time allowed the network user direct access to the digital network.

Even as ISDN is being implemented world-wide, a new strategy for dealing with the future demands on the telecommunications network is being investigated. The future network must be able to provide high bandwidth for video transmission, be flexible to provide for new services, and be easily managed. The new network is called Broadband-ISDN and differs fundamentally from ISDN because all data transmitted across the network is segmented into cells, this being known as Asynchronous Transmission Mode (ATM). At the beginning of this project B-ISDN was at its conceptual stages, but has progressed to national field trial in several countries. The need for efficient transmission protocols and switching techniques which can be applied to the ATM cells has been evident in the research carried out in to Broadband ISDN over the last six years.

1.2 Aims of the Project

One particular protocol developed by British Telecom Research Laboratories in 1984 was the subject of some interest both as a protocol for switching ATM cells, and as a protocol for a Local Area Network for Integrated Services traffic. The protocol was called ORWELL, and much of the work described in this thesis is a performance evaluation of ORWELL, and an analysis of the claims made for it as switching protocol for B-ISDN.

The aims of this project were:

- to build a hardware test bed based on a slotted ring and utilising the ORWELL protocol
- to investigate the behaviour of the ORWELL MAC protocol under varying traffic intensities, traffic mixes, traffic arrival patterns, and asymmetrical loading of the network.
- to characterise the network and generate analytical models of its performance
- to investigate methods of congestion prediction and avoidance
- to implement schemes for access control, and prioritisation with two classes of traffic, and analyse network performance under these schemes

1.3 Overview of the Thesis

A review of the current state of research in B-ISDN is carried out in chapter 2. This provides a background to the work carried out in this project and details the issues which are seen as fundamental to the operation of the B-ISDN network.

In chapter 3, the ORWELL protocol is described in detail, as are the hardware and firmware design issues for the testbed, including the preparatory modelling of an network station.

In chapter 4, traffic types and models of traffic behaviour for an ATM network are reviewed. The modelling of traffic generation on the testbed is described, and the software processes required for cell delay and cell loss analysis of received data are discussed.

In chapter 5, an analytical study of the performance of the testbed is carried out under a number of conditions of traffic loading, and for various traffic generation models. Mathematical models are derived which can be tested against experimental results from the testbed.

In chapter 6, experimental results are compared with the analytical models previously developed. The value of the ORWELL reset rate as a measure of free bandwidth is discussed.

In chapter 7, the introduction of two traffic classes to the test-bed is considered. Issues of prioritisation and access control, which enable an acceptable load area to be defined, are discussed. Indicators of congestion and an access control mechanism are discussed.

Conclusions are drawn, and further work is discussed, in chapter 8.

1.4 Summary

This thesis contributes to research in B-ISDN networks by providing analytical modelling of the throughput and reset interval of a slotted ring local area network testbed utilising an ORWELL based media access protocol under a number of traffic distributions. The analytical models are compared with results from simulations carried out on the testbed. These and other simulations allow conclusions to be drawn about the suitability of ORWELL as a media access protocol for a slotted ring ATM cell switching centre. The study of the ORWELL protocol leads on to a discussion of methods of congestion management in an ATM network. The benefits of different strategies for congestion management are considered and illustrated by results from simulations using the testbed. Conclusions are drawn from these investigations and the direction of further work is indicated.

Chapter 2

A Review of Broadband ISDN

2.1 Introduction

In this chapter a review of the development of B-ISDN is undertaken. The review focusses on the rationale for the introduction of a Broadband-ISDN network, and the technological issues facing the network designers and the committees recommending standards. The background to the work carried out in this project is explained in terms of the interest in media access protocols, hardware testbeds, and ATM networks.

2.2 Broadband - ISDN: the Future Network

Broadband Integrated Services Digital Network (B-ISDN) will be the next generation telecommunications infrastructure, carrying multimedia high bandwidth services over wide areas. The move towards such a network is driven by increasing customer demands for high bandwidth applications found in visual communications, together with the expansion of data communications of widely varying bandwidths to be integrated in a common telecommunications environment [1], [2], [3], [4], [5]. Technologically the emergence of B-ISDN is underpinned by the development of optical fibre networks capable of transporting data at the rates necessary for high bandwidth applications, and high speed VLSI circuits for packet switching and routing [6], [7], [8], [9], [10].

2.2.1 Origins of Broadband - ISDN

The concepts for B-ISDN have evolved from the narrowband ISDN (or simply ISDN) network, the standards for which were laid down in the CCITT I-series recommendations of 1984 [11]. Comite Consultatif International Telegraphique et Telephonique (CCITT) is the international standards body responsible for public telecommunications, now known also as International Telecommunications Union (ITU). Networks utilising these standards are being implemented world-wide in the early 1990s, they are based on the digitised telephone network which is characterised by a 64 kbit/s channel. The channel bit rate of 64 kbit/s is derived from 3.4kHz voice transmission requirements (8 bit sampling at a frequency of 8kHz).

ISDN has provided benefits for the user and the network provider alike, including: a common user-network interface for access to a variety of services, enhanced (out-of-band) signalling capabilities, service integration and new services. The ISDN concepts were based on copper distribution facilities and existing digital switch technology. To obtain higher bit rate transmission under ISDN, 64 kbit/s channels can be combined with groups of circuits being switched in common, but the maximum bit rate that can be offered is 2Mbit/s. This bit rate is insufficient for some of the high bandwidth applications detailed in Table 2.1 [1].

From the start of CCITT discussions on B-ISDN, in January 1985, it was considered necessary to look towards recommendations for a future broadband network based on optical fibre transmission and next generation switching technologies which would support connections of widely varying bit rates, from high resolution video services to low bit rate data, employing a common network interface.

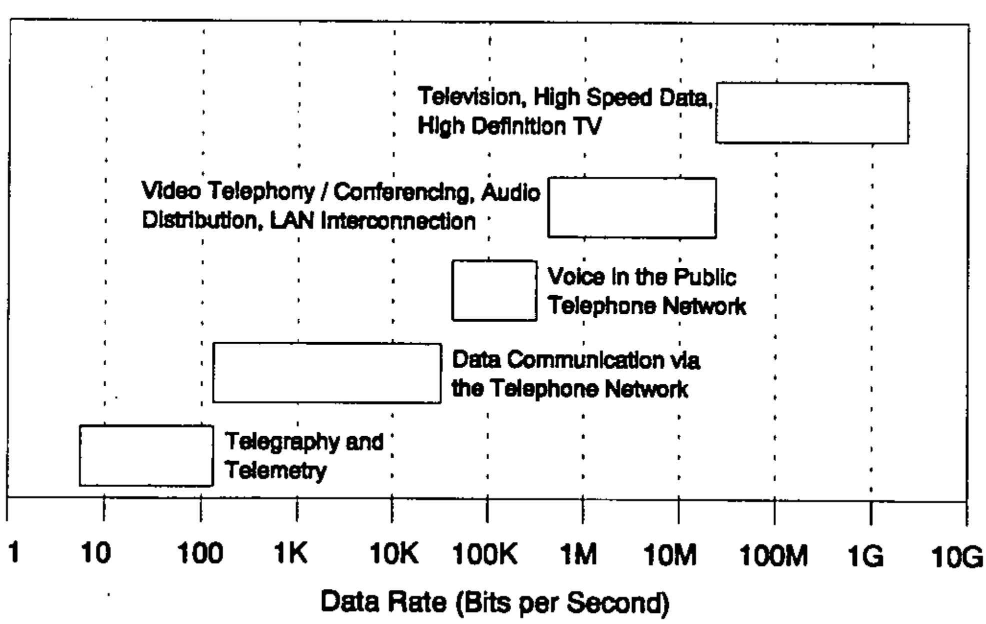

Figure 2.1. Data Rate Spectrum for some B-ISDN Services

2.2.2 Services and applications for B-ISDN

It is intended that the B-ISDN network will eventually form the infrastructure for public telecommunications, hence a major form of traffic for the network will be telephony. The characteristics of telephone traffic are well known through many years of research [12]. What is less well known is the nature of new traffic sources and applications which in time may come to dominate the required bandwidth of the network. The data rate spectrum of various services and networks is shown in figure 2.1 [2].

A multitude of potential broadband applications and services have been proposed for the broadband network. The data rates for these applications vary over 9 or 10 orders of magnitude, and not all data rates are constant. Whilst voice traffic may be a regular 64 kbit/s stream, data traffic consisting of commands, requests for information and file transfers is inherently bursty in nature, and many video applications will use some form of data compression such as variable bit rate codecs to conserve bandwidth. Some characteristics of broadband services are listed in Table 2.1 [1]. Burstiness is defined here as peak bit rate / mean bit rate.

Table 2.1. Characteristics of broadband services

Service	Bit Rate(Mbit/s)	Burstiness
Data Transmission (connection based)	1.5 to 130	1 to 50
Data transmission (connectionless)	1.5 to 130	1
Document transfer / retrieval	1.5 to 45	1 to 20
Videoconference / Video telephony	1.5 to 130	1 to 5
Broadband videotext / video retrieval	1.5 to 130	1 to 20
TV distribution	30 to 130	1
HDTV distribution	130	1

With such traffic requirements in mind the CCITT recommendations for Broadband Aspects of ISDN [13] adopt an interface model to the B-ISDN network based on a breakdown of all traffic entering the network into data cells, which are of a standard length and structure and may be used to transport data of whatever type of service through the network. Such a transport mechanism is known as Asynchronous Transfer Mode (ATM) to distinguish it from the existing Synchronous Transfer Mode (STM) which works on the basis of reserving time slots in a larger time frame structure for each connection made across the network. ATM can allocate bandwidth flexibly to services and applications of widely different bit rate and traffic characteristics. Most importantly it offers a common user interface to all traffic types so that the development of new applications and services is expedited by a standard network interface.

2.2.3 Asynchronous Transfer Mode

Early discussions on B-ISDN revolved around defining a user-network interface structure [14]. There was general agreement on re-using the narrowband ISDN concept of a common multi-services interface, but the nature of the interface was debated between advocates of STM and those of ATM. The STM or circuit switching option was cited as being efficient for existing systems, particularly telephony, but ATM was eventually chosen because:

- A single universal ATM packet fabric could handle many types of services, such as voice, data and video,
- Circuit based approaches become cumbersome as the number of channels grouped together for high bandwidth applications is increased and,
- Manufacturers and users wished to exploit the potential of ATM technology.

Having decided on ATM, a cell-based technology, the size of the cell had to be determined. A fixed size packet rather than variable sized one was adopted because of the ease of switching at broadband speeds, the ease of fair handling of data streams, and more predictable cell delay variation. In June 1989 CCITT study group XVIII defined a 48 byte payload with 5 byte header, see figure 2.2.

Cell Header 5 Bytes	Cell Payload 48 Bytes

Figure 2.2 CCITT Defined ATM Cell

The adoption of ATM as the transfer mode for B-ISDN has brought about its own problems in terms of transmission, switching and network management. Some of these issues, and the current state of research into them, are dealt with in section 2.2.

2.2.4 Optical Networks for B-ISDN Transmission

The development of powerful and economic optical transmission equipment, offering virtually unlimited bandwidth, is another driving force behind the B-ISDN network. Today's optical fibre transmission can provide low attenuation, hence long distances between repeaters, high transmission bandwidths up to gigabit/s levels, a low rate of transmission errors, and good resistance to interference [15], [16]. Such properties are essential if the network is to offer high quality real-time services such as video telephony, video conferencing, and TV distribution. Much work is being pursued regarding Passive Optical Networks (PON) for the distribution of high-bandwidth traffic to customers' premises. The Asynchronous Passive Optical Network (APON), has been studied as a basis for the provision of telephony and integrated services [17], [18].

The other main technological driving force behind B-ISDN is the development of high speed packet switches using the latest in CMOS and BiCMOS VLSI technologies [19], [20]. Processes down to 0.6 micron are now possible [9] allowing the switching of 600Mbit/s data streams required for B-ISDN. Power dissipation is a major factor limiting the level of integration of ATM switches, and for this reason, and other economic reasons, modular switches based on identical sub-switch elements are seen as a desirable route for research. Operational trials of various ATM switches have taken place [21], [22] and research into better methods and implementations is ongoing. Switching technology is discussed in section 2.3.

2.3 B-ISDN Architecture

2.3.1 ATM Cell Transmission

The ATM Cell header contains the information necessary to route an individual cell across the network. In contrast to STM where the connection is maintained by reserving a constant position in the transmission frame (Position Multiplexing), ATM cells each carry enough information in the cell header to perform cell routing, and hence maintain an end-to-end connection (Label Multiplexing). Figure 2.3. contrasts these two methods.

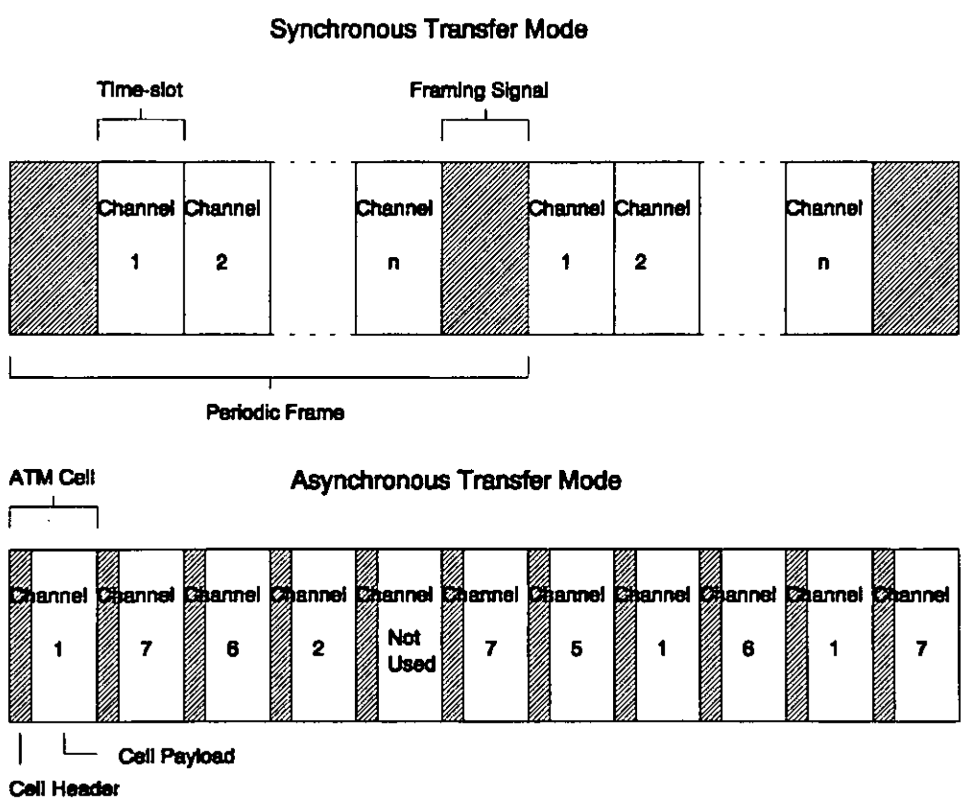

Figure 2.3. STM and ATM transmission principles

A label in the ATM header is structured in two fields to identify the Virtual Channel (VC) and Virtual Path (VP). ATM has thus combined the cell based nature of the packet switched networks such as X.25 [23] with the telephony based circuit switching techniques in the virtual channel concept. Connections are established during the signalling phase, and virtual channel and virtual path information is determined at the switching points. The existence of a virtual path as well as a virtual channel facilitates switching, as several virtual connections can be grouped together as a virtual path over sections of the transmission path. A full description of the ATM cell header can be found in the CCITT recommendations and other works [24], [1].

2.3.2 ATM Switching

The development of ATM cell switches continues with the development of VLSI manufacture and is proceeding at a rapid pace. Given the existence of optical networks capable of transmitting data at gigabit/s rates, the technological hurdle of being able to switch this data without introducing excessive delay or cell loss in the transmission path must be overcome. Switch architecture has been addressed by several authors [1], [25], [26], and four basic types can be defined. These are matrix type switching elements, central memory switching, bus type switching, and ring type switching.

Matrix switching elements are based on crosspoint switches, where cells arriving from a number of input streams are switched to a number of output streams. As cells will sometimes arrive on input streams at the same instant, or be competing for the same output streams, a certain amount of buffering for these packets is required [27], [28]. In matrix switching elements, buffering of ATM cells is provided at input, output and at the crosspoints themselves [29], [30], [31], [32]. The sizing of these buffers and schemes of arbitration between cells contending for access to the buffers are subjects of research [33], [34], as are the VLSI processes required to implement them.

The common memory switch uses a common memory element for all buffering, and hence has reduced memory needs compared with a matrix using physically separate buffers. This type of switch has been used in the PRELUDE experiment [35], and within the RACE (Research and development of Advanced Communication in Europe) project 1012, the Sigma switch [36] is used, which employs a common memory structure. The European RACE project has involved many educational and commercial research establishments in B-ISDN research and in the study of ATM switching for Broadband ISDN [37].

The ring type switching element is of particular relevance to this thesis. A slotted ring is used to transfer packets from input controllers to output controllers, as shown in figure 2.4. ATM cells are placed in the rotating slots by input controllers, and are transferred to their destination output controllers. When a slot has been released it can be re-used immediately increasing the effective bandwidth of the ring. Destination release of slots gives rise to problems of fairness of access to the ring's bandwidth which may be resolved by the ring protocol. Several rings can be operated in parallel to form a torus. The ORWELL protocol [38] employs such an approach for high speed packet switching, and is used as the basis for the medium access protocol of the slotted ring test bed designed for this thesis and described in chapter 3.

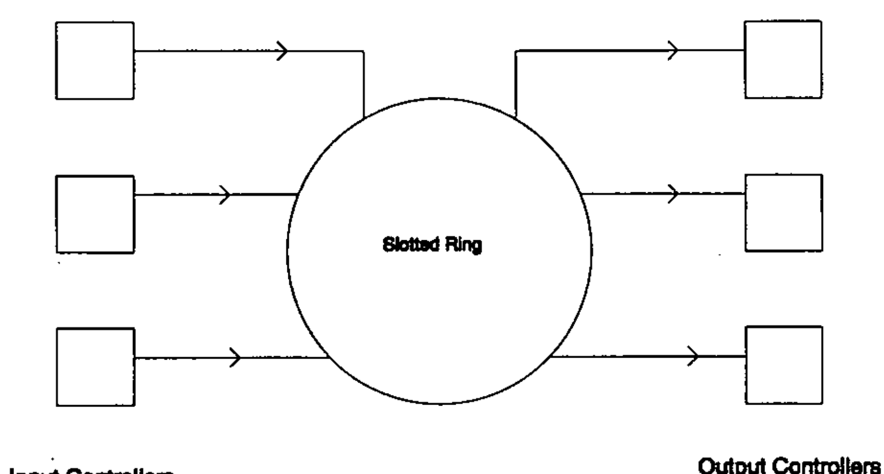

Figure 2.4. Ring type ATM switch

2.3.3 Quality of Service Concept Applied to B-ISDN

In order to provide a service which customers want to use, the network provider must be able to guarantee a certain Quality Of Service (QOS) to a potential user. The parameters by which the QOS is defined are the maximum rate of cell loss, and the maximum delay time of cells transmitted across the network. The variation of cell delay time or 'cell delay jitter' may also be an important parameter for the network user. Because of the variety of applications envisaged for the B-ISDN network, there will be some traffic types which have a requirement for low delay times and some traffic types which are delay tolerant. In general, real-time voice and video services are delay sensitive whereas computer communications are delay tolerant.

In several studies on the allocation of bandwidth in ATM networks [39], [40], [41], the possibility of more than one QOS class being offered to the network user has been explored, and various schemes of bandwidth allocation have been analysed and simulated. The concept of a number of QOS classes brings with it other implications for network management the first of which is prioritising the service of cells belonging to one QOS class over those belonging to another, which might occur where there is contention for access to a buffer at a switching point or multiplexer. Various studies have suggested that the overall capacity of the network can be increased by prioritising delay sensitive traffic over delay tolerant traffic [42], [43], [44]. In such situations the management of buffer length and prioritisation schemes can be used to maintain service to all QOS classes. For these reasons the ATM cell header has been allocated one bit for priority control. The bit will indicate if the cell can be discarded in favour of other cells in a congested network.

2.3.4 Traffic Analysis and Modelling

In order to perform analysis on the ATM network, certain assumptions have to be made on the behaviour of traffic sources, and the likely characteristics of future traffic sources. In general it is not possible to use 'real data' when simulating networks, although some such data has been gathered [45]. This is because the traffic carried on present networks cannot be assumed to be similar to that of future networks, and because the form of tomorrow's applications and services is not yet known.

Analytical models have been developed to represent both individual and multiplexed traffic sources, to model a cell arrival stream at an ATM switching node as a means of analysing network performance. In the case of high bandwidth traffic sources such as video codecs, models incorporating autocorrelation features over line and frame repetition periods have been developed [46], [47] to give a more realistic behaviour pattern for performance analysis. In many situations, Markov Models [48], [49], [50] have been used to generate cell arrival statistics for performance analysis. Generally, in simulations, traffic sources are characterised by a peak cell rate generation, a mean cell rate generation, and some measure of burstiness, i.e. the distribution of cells with time, see figure 2.5.

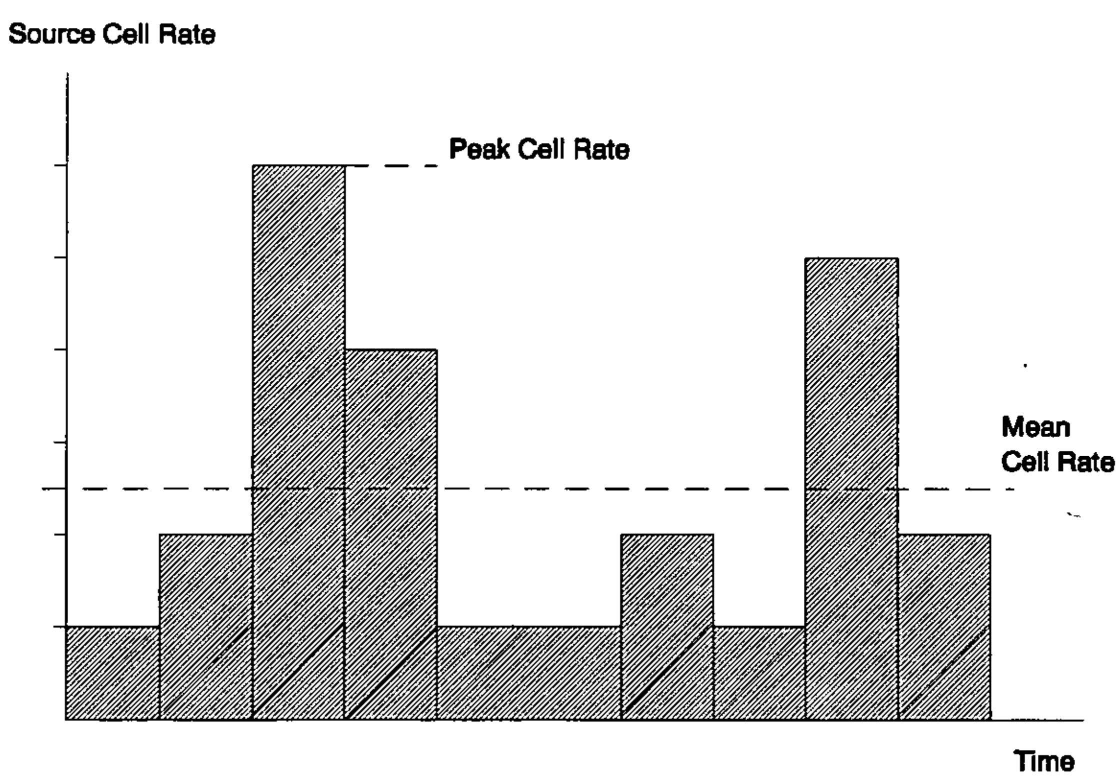

Figure 2.5. Peak rate, and mean rate of a bursty ATM traffic source

The burstiness property of ATM traffic means that not all sources on an ATM network will be transmitting at their peak bit rate at all times. This presents the possibility of network efficiency gain by statistical multiplexing. Instead of reserving the peak transmission bandwidth for each connection supported by the network, a somewhat reduced bandwidth will actually be required due to the averaging of peak rate bursts of cells from different sources at different times. Although statistical multiplexing is a desirable feature that ATM can offer over STM, the primary consideration of the network provider is to guarantee the particular QOS offered, so that the efficiency gain of statistical multiplexing must be of secondary importance.

2.3.5 Congestion and Admission Control

Every network has a limit to its capacity for transporting information, and the continued acceptance of more and more connections by an ATM network will inevitably lead to excessive cell delay, buffer overflow and cell loss. The effect is known as congestion, and the result is that the QOS provided to the user is degraded. Congestion management and avoidance is another major field of current research in B-ISDN. Avoiding congestion is more effective than trying to manage the network after congestion has occurred [51], [52], [53]. In avoiding congestion a mechanism for admission control of new connection requests or call acceptance is implemented, and is based on controlling the number of connections of each QOS class that may be accepted on to the network at any time. Call acceptance is determined during the signalling or call set-up phase of the connection. The acceptance of a new connection is dependent on the additional traffic not affecting the QOS offered to existing traffic. The ability of a network to accept a further connection will depend on the mix of traffic types of existing connections, as well as the actual numbers. This gives rise to the concept of an admissible load region for a network [43], where the combination of current connection QOS classes will determine admission to the network. This concept is described in chapter 7 for the slotted ring test bed , when two traffic types with different arrival profiles are simulated.

Much of the performance analysis of ATM networks has been concerned with admission control policies [54], [55], [56], congestion control, and prioritisation. Traffic control procedures are not currently standardised by CCITT, and the extent to which they should be standardised has not yet been decided. In some areas, defined standards of how the network will behave under overloading would be useful to both the network provider and the network user so that applications can be designed for optimal performance over a variety of network loading conditions. One area important to controlling the behaviour of input traffic is 'Source Policing', where the input of cells to the network is monitored and controlled.

2.3.6 Policing ATM Sources

The call set-up phase of a B-ISDN connection is a process of negotiation between the potential user and the network where, if the connection is accepted, the user will contract to keep to certain bounds of peak and mean bit rate for data transmitted, and the provider will guarantee a QOS in terms of cell delay and loss for the connection. In order for the system to work fairly both sides must keep to the agreed contract, but the provider must ensure the user does not exceed the agreed data input for the sake of other users, and the integrity of the entire network service. The concept of source policing thus arises, where the network ensures that each input connection keeps to the bounds of cell input which have been agreed.

Several mechanisms for source policing have been investigated in recent research [57], [58]. Methods such as the 'Leaky Bucket' [59], [60] are used to control the peak input rate by incrementing a counter every time a cell is received, and blocking the arrival of more than a certain number of cells in a given time period. Other methods based on window schemes [61] have the same effect. In general it is easier to control the peak rate of cell input to the network than it is to control the mean rate and the length of bursts. This is because of the longer time required to average measurements for mean rate, and the difficulty of determining the beginning and end of bursts. Standards are not yet developed in this area, and are dependent on definitions of traffic profiles which are also not yet clear.

2.4 B-ISDN Standards, Implementations, and Trials

2.4.1 Standards

The existing set of standards for B-ISDN are in terms of CCITT recommendations detailed in Table 2.2.

Table 2.2 CCITT recommendations on B-ISDN

Num	Title	Description
I.113	Vocabulary of terms for broadband aspects of ISDN	Definition of terms used for B-ISDN
I.121	Broadband aspects of ISDN	Principles of B-ISDN
I.150	B-ISDN ATM functional characteristics	ATM cell header functions for routing, priority, flow control.
I.211	B-ISDN service aspects	Classification and description of B-ISDN services, synchronisation, timing.
I.311	B-ISDN general network aspects	Networking - routing, switching, Signalling, Traffic control
I.321	B-ISDN protocol reference model and its application	Overview of functions in layers and sublayers, extension to I.311
I.327	B-ISDN functional architecture	Basic architecture model, network capabilities
I.361	B-ISDN ATM layer specification	ATM cell structure. Cell header and its coding

I.362	B-ISDN ATM adaption layer (AAL) functional description	Principles of the AAL, sublayering, segmentation and reassembly (SAR)
I.363	B-ISDN ATM adaption layer (AAL) specification	Specifications for 4 AAL protocol types
I.413	B-ISDN user-network interface	Reference configurations, Interfaces at 155.520 Mbit/s (622.080 Mbit/s)
I.432	B-ISDN user-network interface - physical layer specification	Electrical / optical parameters, interface structure (SDH frame), Header error control
I.610	OAM principles of the B-ISDN access	Operation And Maintenance principles

The comparatively rapid pace of standards development [14], has been driven by the telecommunications equipment manufacturers and network service providers who are keen to take advantage of the potential of high bit rate services becoming an everyday fact of life, and consequently being able to profit from this. ATM standards have been developed from narrowband ISDN standards, and incorporate interfaces to existing digital transport mechanisms. Particularly notable is recommendation I.432, which is concerned with the interface to existing optical fibre transport mechanisms known as Synchronous Digital Hierarchy [62], [63], [64]. It is this interface which will provide transport for large numbers of ATM cells concentrated together, and because the standard is used for the current generation of transport equipment, ATM facilities can be added piecemeal to the existing network. A description of the B-ISDN architecture and protocol can be found in the standards in table 2.2 and in several review articles [65], [66], [67].

2.4.2 Evolution of the Broadband Network

The need for broadband provision is driven by trends in business and society, and the potential for profit in supplying these needs. The ability to introduce broadband capabilities at an economic cost is determined by advances in the key technologies of high-speed integrated circuits, and fibre optics, together with the development of techniques needed to apply them. The move towards a fully integrated broadband network will be evolutionary [68], [69] and ATM facilities will be introduced in stages over a period of time. An evolutionary process is envisaged because of the continuing development of applications, technologies and ATM standards. At this point in time the nature of future applications cannot be known, though there is interest from the network operators especially in generating new ideas [70], these must be given the chance to establish themselves. Because of the limitations on capital and installation investments for new broadband equipment, broadband services will be introduced as demand increases from customers.

The initial stages of broadband introduction will be assisted by the conversion of telecommunications transport equipment to the Synchronous Digital Hierarchy (SDH) standard which is now ongoing. This standard offers advantages to current circuit switched narrowband ISDN transport, and is seen as the most effective transport protocol for B-ISDN and ATM.

The addition of ATM facilities to the existing network might be in areas such as high-speed LAN interworking, high-resolution image transfers, and interactive multi-media communications. Initially ATM will be offered as a permanent virtual circuit operating as a separate sub-network, though sharing the transport infrastructure of narrowband ISDN. This ATM overlay network will grow as demand grows, first to a switched ATM network, and then to a fully ATM-based network. The overlay approach will enable network operators to gain experience in management of the ATM network, as it expands under customer demand. Introduction of ATM private networks are also likely to drive up the usage of ATM facilities in the public network. In the long term an ATM core network is envisaged serving a wide range of cell relay and circuit applications. The early stages of the development of a B-ISDN network or sub-network are already to be seen through the activities of network operators.

2.4.3 Introduction of B-ISDN Networks

In North America AT&T has introduced ATM compatible equipment for high speed packet switching known as Switched Multimegabit Data Service (SMDS) and frame relay, an update of X.25 packet switching [71]. In Europe, there have been field trials of limited scale ATM networks in Germany and Switzerland [72] and Belgium, France and Spain [73] involving ALCATEL.

Other network providers such as British Telecom have developed their own ATM based Local Area Network (LAN) platforms to implement the technologies required for broadband networks and to gain experience in the operation and maintenance aspects as well as network management [74], [75]. Work done in the Universities also includes sophisticated ATM networks. Columbia University in New York has operated a programme called MAGNET for a number of years [76], [77].

Although much of the research into ATM techniques is carried out by analytical means and simulation, a valuable insight into network operation can be gained through the implementation of ATM networks.

2.5 The Slotted Ring Test Bed

2.5.1 Background

At the beginning of this project in 1989, the development of Broadband ISDN networks was at an early stage. Many of the issues that have been defined in the

CCITT recommendations of Table 2.1 were still undecided, and the choices which have since been made, were at that time still open to debate. There was considerable interest in slotted ring technology for switching ATM traffic and for use as a Local Area Network for ATM traffic. Integrated Local Area Network (ILAN) projects had existed for some years and two projects in particular influenced the direction of this research project, MAGNET [76] and ORWELL [78]. More recently work in Japan on the ATMR protocol [79] has taken place but here the MAGNET and ORWELL ILANs are considered as they had a bearing on the development of this project.

2.5.2 The MAGNET Project

MAGNET is a project initiated in the Center for Telecommunications Research, Columbia University, New York, in 1985. The stated objectives of the MAGNET project were to: "address fundamental issues and identify open questions arising in the design of the architecture of a Network Testbed for Integrated Local Area Network (ILANs) supporting video, voice, data, graphics and facsimile". These objectives have been pursued by a combination of theoretical investigations [76], [80], simulation studies [76], and exploratory designs and implementations [81], [82].

MAGNET is a slotted ring based ILAN linking multimedia integrated workstations known as EDDY, which provide access to real-time services such as televideo, video conferencing, and telephony, and access to data communication services by computer and facsimile.

It was proposed to characterise the MAGNET network performance under various proposed control schemes, and a set of evaluation guidelines were established to investigate the behaviour of the network under different traffic types and mixes. The effect of multiplexing techniques and scheduling schemes were investigated to gain a better understanding of media access protocols and the "fairness" problems associated with a slotted ring. From these studies a set of adaptive protocols has been suggested, controlled dynamically by a controller known as WIENER [77], adapting to network traffic load and changing user environment.

MAGNET was designed and implemented as a fibre optic slotted ring with two wavelength division multiplexed 100Mbit/s channels. The bandwidth capability of the system allows it to carry real-time video channels as well as voice and data. Different traffic classes have been defined to express the different requirements of traffic types and a Quality of Service has been defined for each class in order to evaluate network performance under changing conditions.

The concept of a hardware testbed which would allow a study of the Media Access Control (MAC) protocol, as well as higher layer functions such as prioritisation of traffic classes and congestion management, was attractive at the start of this study. Particular interest in a specific MAC protocol, ORWELL, developed by British Telecom Research Laboratories, led to the decision to base the slotted ring testbed on the ORWELL protocol.

2.5.3 The ORWELL Protocol

ORWELL was developed as a high speed LAN / MAN media access protocol in 1984 [83] at British Telecom Research Laboratories (BTRL). The main aim of the protocol was to provide an integrated services capability with complete flexibility of bandwidth allocation between traffic types. A slotted ring topology was chosen to obtain reduced delay time and increased throughput over token ring protocols by utilising destination release of slots. A slotted ring [84] is a ring with a pattern of circulating slots, each capable of holding a cell of data. Stations on the ring transmit cells by putting them in an empty slot which circulates to the destination station where the data is received, the slot being re-used or left empty, as shown in figure 2.6.

The established media access protocols for rings at that time were token-passing, and the Cambridge slotted ring. These were not developed for integrated traffic, and their performance has been shown to be inferior to ORWELL for dealing with delay sensitive services [85], [86].

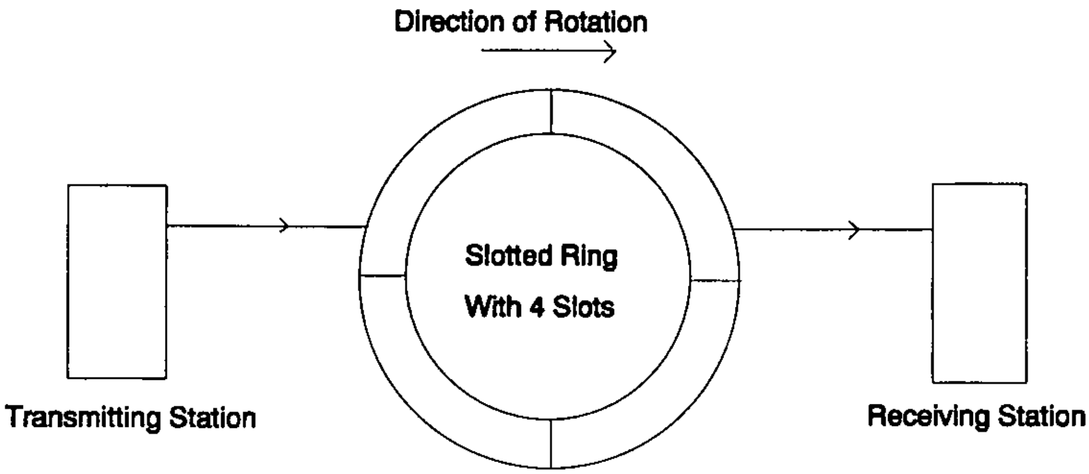

Figure 2.6. Slotted ring topology

As with all shared media protocols, ORWELL introduces a mechanism to ensure fair allocation of bandwidth between stations competing for access to the ring. The ORWELL protocol allows each station to transmit a pre-determined number of data cells before it must pause and wait for a reset signal. When all nodes are paused or have no data to transmit, the reset signal is generated to restore all stations to full transmit capability. Full details of the ORWELL protocol are discussed in chapter 3.

From the initial work on ORWELL [78], and after the standardisation of the ATM structure by CCITT [24], a full specification of ORWELL was produced making it ATM compatible [38]. Some analytical work and simulation studies on the performance of the protocol have been carried out [85]. An optical fibre based test bed has been developed at BTRL for the evaluation of many issues surrounding ATM networks. This test bed employs the ORWELL protocol although much of the reported work [74], [75] does not explicitly deal with the behaviour of the protocol. In the development of ORWELL, it is clear that the protocol could be used as the basis for either an ILAN [87], [88], or as a ring type switch for ATM cells.

2.6 Summary

A review of the concepts and research issues involved in B-ISDN is necessary to put the thesis in to context. B-ISDN has been developed as a solution to the integration of traffic sources of widely differing bandwidths and statistical characteristics using a common transmission medium and a common user interface. The testbed was developed to provide a better understanding of the ORWELL protocol and its application either as an ATM switching mechanism, or as a Local Area Network for integrated services traffic. The development of a hardware testbed as opposed to software simulations enables network parameters to be varied, and different simulations to be performed quickly. The ring protocol was also easily varied and customised. The disadvantage of the hardware testbed as a simulator was the lack of flexibility of network configuration, which was limited to a slotted ring. In spite of this limitation, the testbed has proved to be a useful tool in the analysis of the ORWELL protocol and the problems of congestion management in ATM networks. The design of the testbed and modifications made to the ORWELL protocol will be discussed in chapter 3.

Chapter 3

Testbed Architecture

3.1 Introduction

In this chapter, the design of a slotted ring testbed for the investigation of the ORWELL protocol, and of methods of ATM network congestion avoidance are described. Initially, a number of LAN media access protocols are compared and contrasted to show that the ORWELL protocol operating on a slotted ring has potential as a protocol for real-time traffic transmission. The choice of testbed architecture is then explained, and the hardware of each station's interface to the slotted ring is described in detail. The software processes which cause ATM cells to be generated and analysed are described, and the interface between these processes and the ring interface hardware is explained.

Some features of the development of the testbed are also discussed, such as the use of a PASCAL simulator program, the use of programmable logic devices to implement the ORWELL protocol, and the parallel programming language OCCAM used to program Transputers.

3.2 Media Access Protocols

In order to define an access protocol for the testbed, the protocols of existing Local Area Network (LAN) structures were studied, and their suitability for an Integrated LAN assessed.

In the allocation of bandwidth amongst various stations sharing a transmission channel such as a ring or bus, an initial decision must be made to allocate the bandwidth resource in a fixed manner, with each station separately allocated *1/N* of the total bandwidth, where *N* is the number of stations, or to allocate bandwidth as requested by stations. Since at any given time few of the *N* stations may be transmitting, much of the available bandwidth which might otherwise be allocated to these stations is wasted in a fixed bandwidth-allocation scheme. Bandwidth allocation on demand to those stations with data to transmit leads to more effective use of the channel capacity and the following LAN topologies and protocols all employ demand-based allocation of bandwidth. The function of Media Access (MA) protocols is to ensure the available bandwidth is allocated efficiently, and fairly.

The division of integrated services LAN bandwidth between various traffic types or traffic classes can also be performed in a fixed manner or a flexible manner. The Fiber Distributed Data Interface (FDDI) protocol [89], [90] which has been implemented as a high speed communications backbone network uses a fixed division of resources for isochronous (real-time) traffic and non-isochronous traffic. Studies have shown [91] that dynamic allocation of bandwidth resources between classes of traffic has advantages over fixed allocation, and given the unknown nature of future B-ISDN traffic, a dynamic allocation method seems sensible. For the testbed dynamic allocation of bandwidth between classes of traffic was chosen. Some of the existing Media Access protocols for LANs were investigated, and their suitability for integrated services, and in particular real-time traffic is discussed.

3.2.1 Random Access Protocols

The major application of random access protocols are for bus-based LANs, although these stem from the ALOHA protocol developed at the University of Hawaii for a ground-based radio system [84]. This work has been developed by many researchers and subsequently more refined protocols are in place for Ethernet [92] and IEEE 802.3 [84] protocols. The concept of ALOHA and other random access protocols is that each station on the network transmits a frame of data when it has data to send, and the receiving station acknowledges successful transmission with an acknowledgement message. With more than one station able to transmit at any time there is the possibility of transmissions colliding and both data frames being corrupted, if this occurs no acknowledgement will be sent. A station will wait for a period of time after sending its data frame to receive an acknowledgement that the data was received, and if the frame was lost it simply waits for a random length of time before re-transmitting. The fairness criteria is assured by the random access of each station to the available bandwidth. The ALOHA protocol works well if there is low usage of the channel, but throughput of good packets actually decreases as the level of attempted transmissions is increased over a certain point, due to the collision of data transmissions. The protocol has been analysed [84], [93] and for a Poisson distribution of arrivals at each station, the relationship between throughput and offered traffic is

$$S = Ge^{-2G} \tag{3.1}$$

Where S is the average number of successful transmissions per packet transmission time, and G is the average number of attempted packet transmissions per packet transmission time. The function reaches a maximum at $G=0.5$, where $S = \frac{1}{2e}$

Refinements of ALOHA for bus-based systems include the addition of a listening capability firstly to detect if another station is transmitting, and secondly to listen to its own transmission and detect if a collision has occurred by another station commencing its transmission simultaneously. Networks with these features are known as Carrier Sense Multiple Access with Collision Detection (CDMA/CD) systems and are widely used on LANs.

Random access protocols are not well-suited to an integrated services LAN because of the random loss of data packets due to collision, and the unbounded maximum transfer delay time that a packet might experience. Re-transmission of data packets is not feasible for high-speed real-time traffic. The framelength of the IEEE 802.3 CDMA/CD standard is at least 64 bits and a maximum of just over 1500 bits, so that loss of a frame could imply a significant loss of data. Although voice transmission over Ethernet has been implemented [131], a bus topology and random access protocol were discounted for the testbed because of these problems.

3.2.2 Token Ring

A ring-based network does not have the collision problem of a random access bus-based network because a ring is in fact a collection of point-to-point links. Data is sent from one point to the next and a circle is formed. As with a bus network, the frame of data being transmitted from one station to another must be addressed to the intended recipient so that other stations will not receive the data. In a ring LAN, stations will repeat data frames not addressed to themselves on to the ring so the destination station finally receives the data. In order to control allocation of the ring bandwidth for transmission, a polling mechanism known as token passing has been developed and is embodied in the IEEE 802.5 Token Ring standard [84]. The token is a specially encoded bit pattern which may be transmitted from one station to another around the ring. When a station has data to transmit, it seizes the token and transmits a data frame for a 10ms period of time known as the token holding time. After this, the station must relinquish the token and pass it to the next downstream station. If this station has data to send it will hold the token, otherwise it will retransmit to the next station and so on. The fairness criteria of the media access protocol is met by this token passing system, it also provides a ready indication of the intensity of traffic on the network since the token will cycle more quickly on a lightly loaded ring than on a ring where many stations have data to transmit. When a station has the token and is transmitting, data circulates around the ring to the destination station where it is copied, and arrives back at the source station which deletes the data from the ring. This is called source-deletion of data.

Performance analysis of token ring protocols can be found in [93], where a number of expressions for average transfer delay of packets are derived. The maximum delay time that a packet could experience is the token cycle time, since within each token cycle time, every ring station will transmit its queued data. A general expression of the cycle time *Tc*, for polling based networks will serve as an estimation of the maximum packet delay, so where M = number of stations on the network, N = average number of packets of mean length X bits stored at a station when it receives the token, R = the channel capacity in bits/s, and w = the token transfer time between stations

$$Tc = M\left(N\frac{X}{R}+w\right) \tag{3.2}$$

NX/R is the mean time required to empty the queue at a station

Relating this to the ring throughput which is defined as:

Throughput S = (average arrival rate / network capacity)

Assuming each station has the same mean arrival rate = λ (symmetric loading) then,

$$N = \lambda\ .Tc \tag{3.3}$$

$$S = M\lambda\ .(X/R) \tag{3.4}$$

hence,

$$Tc = Mw + STc \tag{3.5}$$

$$Tc = Mw/(1\text{-}S) \tag{3.6}$$

This shows that for a ring LAN employing token passing as a media access protocol, the maximum packet delay increases with the number of stations on the network and also with the network throughput. By controlling access to the network and thus limiting the throughput, delays can be kept within certain bounds [132], however a token passing media access protocol was ruled out for the testbed in favour of other ring media access protocols which looked more promising in limiting packet transfer delay times.

3.2.3 The Slotted Ring

The slotted ring shares the ring topology of a token ring but is distinguished by having one or more fixed length slots comprising of a constant number of bit positions which circulate continuously around the ring, see figure 3.1(a). The slotted ring is analogous to a set of rotating doors conveying people (packets) from inside a building to outside (source to destination) or vice-versa, see figure 3.1(b).

Each slot of the ring can accommodate one mini-packet of data which will itself consist of a header containing the destination station address, and a data payload. Since a circulating pattern of slots must be supported by the ring, the ring latency (the length of time taken for a bit to completely circulate the ring) will determine how many slots of a particular size can exist on the ring. For example a 10Mbit/s ring gives a bit-time of 0.1us. With a typical propagation delay of 5us per kilometre of cable, clearly 50 bits can be supported by each kilometre of cable. A ring with 10 stations each having a delay of 10 bits, connected by 2km of cable can support a 200 bit pattern which might consist of four 50-bit slots or three 64 bit slots with a gap of 8 bits. To increase the ring latency, shift registers may be used to pad out the number of bits supported by the ring allowing more slots to exist.

A number of slotted ring protocols have been developed one of which, the Cambridge Ring [94], is a forerunner of the ORWELL protocol, and has been analysed as a possible LAN for integrated traffic [86]. With the Cambridge Ring, access to slots from a station is controlled by three rules:

(i) a station can occupy only one slot on the ring

(ii) slots are cleared on their return to the source node (source release)

(iii) a station may not re-use the slot it has just cleared or the one immediately following.

Rules (i) and (iii) ease hardware constraints and rule (ii) ensures an even distribution of available bandwidth (fairness). The slot format for the Cambridge Ring is shown in figure 3.2.

There are a total of 38 bits of which 32 bits are the mini-packet contents, 16 bits of address information and 16 bits of data. Of the other 6 bits, 3 are used to implement the ring's media access protocol indicating status, such as when the slot is full or empty, and 3 are response bits from the receiving station to act as an acknowledgement to the transmitter.

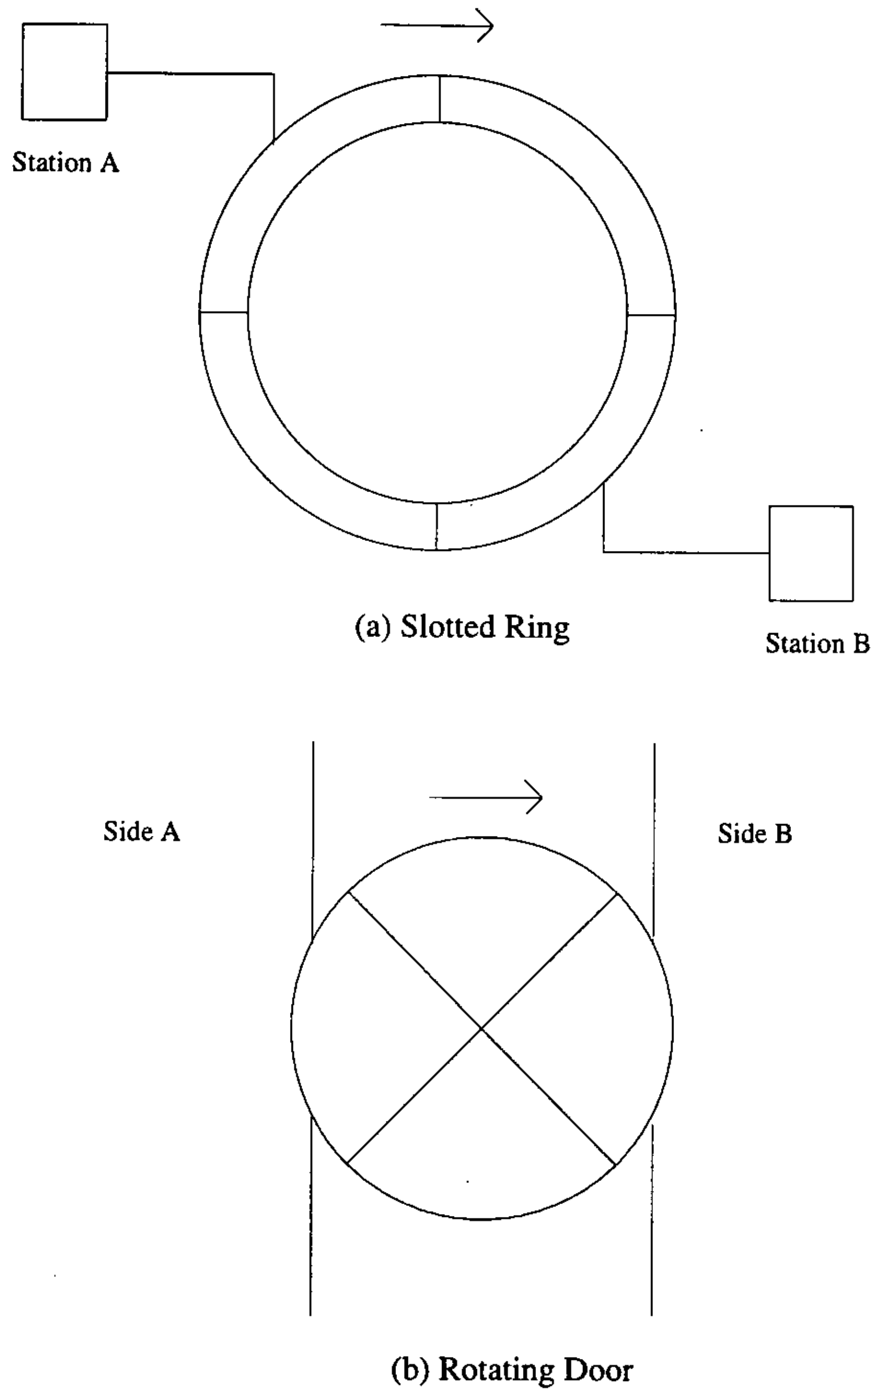

Figure 3.1 Slotted ring topology and rotating door analogy

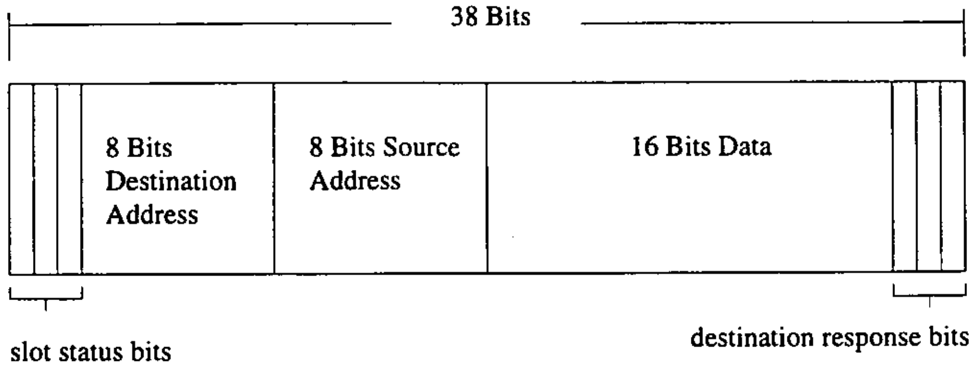

Figure 3.2 Slot Format for Cambridge Ring

Transfer delay analysis of packets carried over the Cambridge ring has been undertaken [86], but some of the problems highlighted for use as an ILAN are: a high overhead of header bits to data bits, the existence of quasi-stable states affecting the usable bandwidth of the ring, and source release of slots requiring them to completely circulate the ring before being re-used. There is no ready indication of traffic intensity on the network such as the token rotation time seen on a token ring. Studies comparing the Cambridge ring to ORWELL as an access protocol for integrated traffic have found ORWELL delivers a superior performance [85].

3.2.4 ORWELL protocol

The ORWELL media access protocol was developed at British Telecom Research Labs in 1984 by J.L. Adams and R.M. Falconer [78], [82] as a protocol for an integrated services LAN that could be adapted to any new traffic source, and as such it fitted in well with the development of ATM, particularly with the choice of a slotted ring capable of transmitting data in the form of mini-packets. The ORWELL slotted ring employs destination release of slots to improve access delays, but has to introduce a mechanism to ensure fair distribution of the network bandwidth among competing stations. The ORWELL fairness mechanism is based on a counter present at each station which is incremented whenever the station transmits a mini-packet into a passing empty slot. If the counter reaches an agreed maximum value for the node (D_i), the station enters what is known as the PAUSED state and further transmissions are temporarily inhibited. Downstream nodes are therefore allowed to access slots which might have been filled by the PAUSED station. The stations on the ring which have become PAUSED must be reset once all stations which have data to transmit have done so. This is achieved by a mechanism where a station which is PAUSED or IDLE (has no data to send) may not use an empty slot, but may mark it as a TRIAL slot. If the TRIAL slot arrives at a node which has data to transmit and is not PAUSED it will be overwritten and the data transmitted, however if the TRIAL slot circulates the ring and returns to its originating station the TRIAL slot is converted to a RESET slot. The RESET slot circulates the ring resetting all of the stations' counters (known as D-counters) to zero. Figure 3.3 illustrates the protocol.

The ORWELL protocol is suitable as an integrated services LAN protocol as :

1. A much tighter bound on the maximum packet transfer delay time can be obtained for an ORWELL mini-packet than for the other protocols considered, partly because the length of an ORWELL RESET cycle is much less than the token rotation time of a token ring, and all ORWELL stations gain access to the ring within a RESET cycle, rather than håving to wait for an entire token rotation cycle.
2. It provides a means to recognise traffic intensity on the network by means of the time between resets (reset interval) which will be low for light traffic loads and high for heavier loading.
3. It has the ability to handle unbalanced loads efficiently.
4. By destination release and reuse of slots, the ring can offer an effective bandwidth higher than the ring data transmission rate.

The reset interval, or its inverse the reset rate, is an indication of the available bandwidth remaining on the ring and as such provides a novel and fully distributed method for network load control.

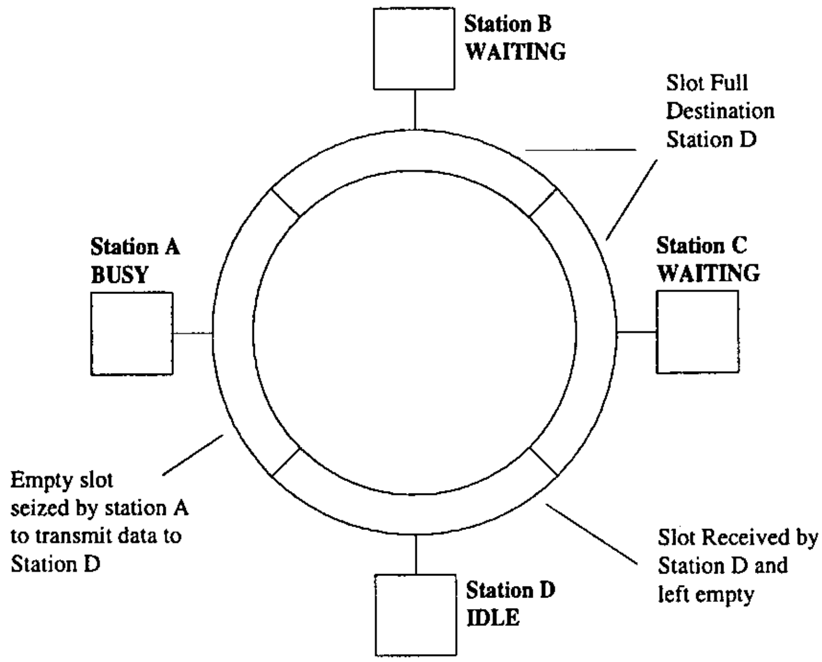

Station A is transmitting to Station D using all available slots and blocking access for stations B and C. Station D has no data to transmit and renders the received slots empty.

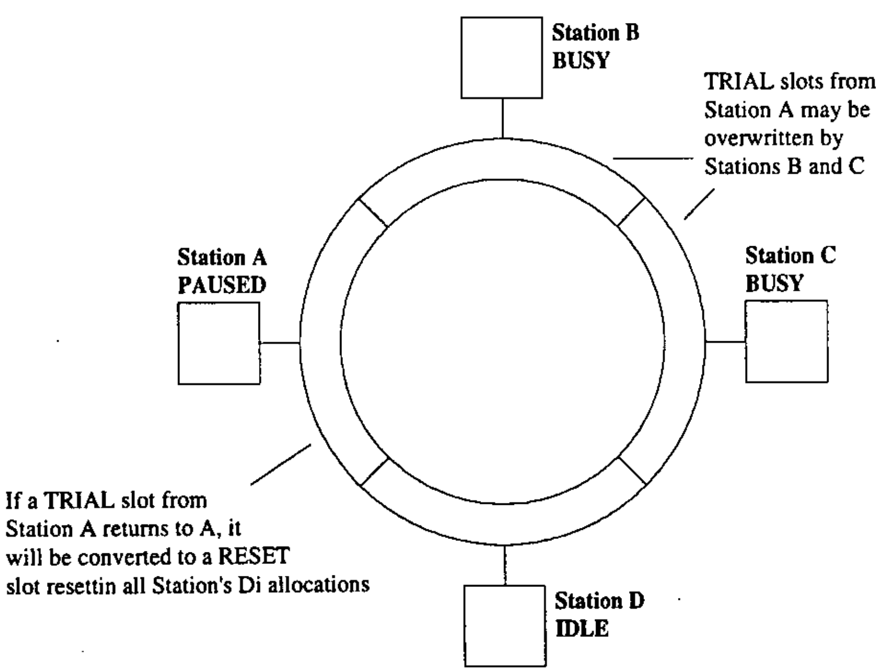

Station A has transmitted its full Di allocation of mini-packets, and has become PAUSED. The other Stations now have access to empty slots. Station A will convert EMPTY slots to TRIAL slots, and will convert its own TRIAL slot to a RESET slot if it succeeds in circulating the ring.

Figure 3.3 ORWELL Media Access Protocol Operation

The ORWELL protocol was chosen as the basis for the testbed media access protocol because of its suitability for integrated services traffic and the interest in ORWELL in the late 1980s. The information of ORWELL available in 1988 was a preliminary specification, and after ATM standardisation, ORWELL was specified as a fully ATM compatible protocol [38]. The testbed utilises the ORWELL concept but does not seek to implement the full specification due to limited resources, time and interest in other aspects of the investigation. An analysis of the testbed media access protocol performance is given in chapter 5.

3.3 Testbed Architecture and Protocol

3.3.1 Network Architecture

A diagram of the testbed architecture is given in figure 3.4. Four independent stations are connected to the four-slot ring. Each station consists of a ring-interface card, a packet-level-interface card, and a traffic generator/analyser.

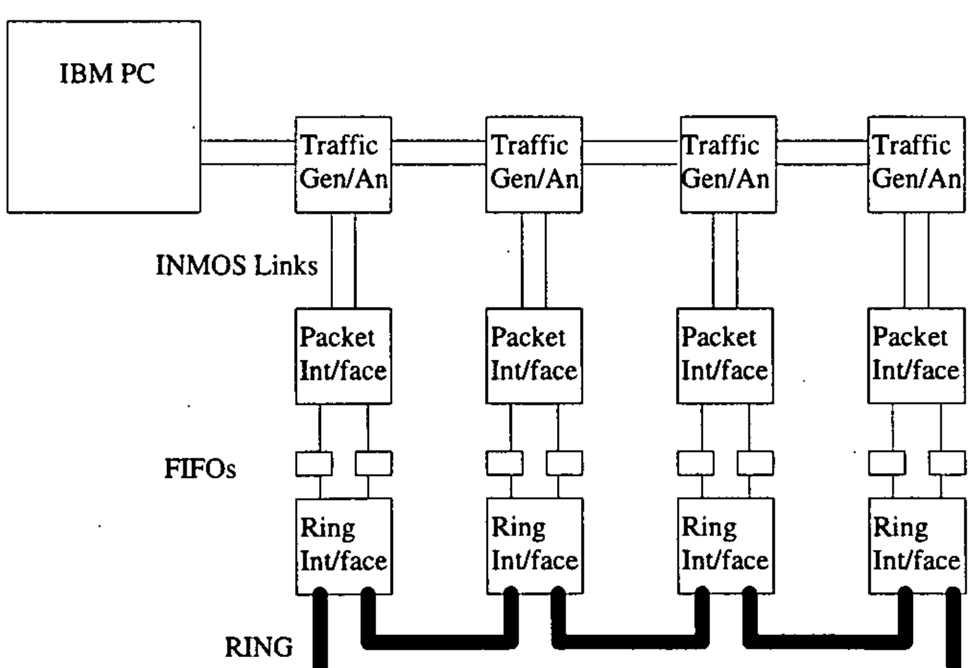

Figure 3.4 Testbed Architecture

The ring-interface card is a TTL-based printed circuit board with Programmable Logic Devices (PLDs) implementing the testbed protocol. The packet-level-interface is based on a T222 16-bit transputer used as a microcontroller to interface to the ring-interface card via input and output FIFO buffers, and the traffic generator/analyser which is a T800 32-bit transputer. The transputers for all of the four nodes are connected together via INMOS links [95] to a PC running a controlling program to collect and store results from the simulation exercises. The transputer programs are downloaded from the PC in a Transputer Development System (TDS) [96] environment at the start of a simulation. It is interesting to note that transputers are also used as hardware controllers on the MAGNET project [76]. They are particularly useful for real-time measurements because of onboard timers, and in the testbed they provide high processing power and a simple method of communicating results through serial links.

3.3.2 Physical Layer of the Testbed

The fabric of the testbed slotted ring is a shift-register contained on the ring interface card. Each station requires a 59-stage shift-register to perform its input and output operations on the ring. A further 126 stages are added at each station with the additional shift-register making a total of 185 stages, see figure 3.5.

The decision was taken to use shift-registers as the medium to support the slotted ring pattern rather than undertaking the design of interfaces to cable or fibre as this was considered not to be the main emphasis of the testbed development. The ring employs bi-phase (Manchester) coding of data so that future work could extend the ring to use cable or fibre, and station synchronisation to the incoming data would be possible. The second reason for encoding the data was to use a code violation [73] to signal the start-of-slot marker. As the slots circulate around the ring, and approach a station, the station must synchronise on the start of the slot to correctly read the header and receive the data packet contained in the slot. A unique identifier for the start of the slot can be created by using a binary sequence that cannot occur in normal data as shown in figure 3.6.

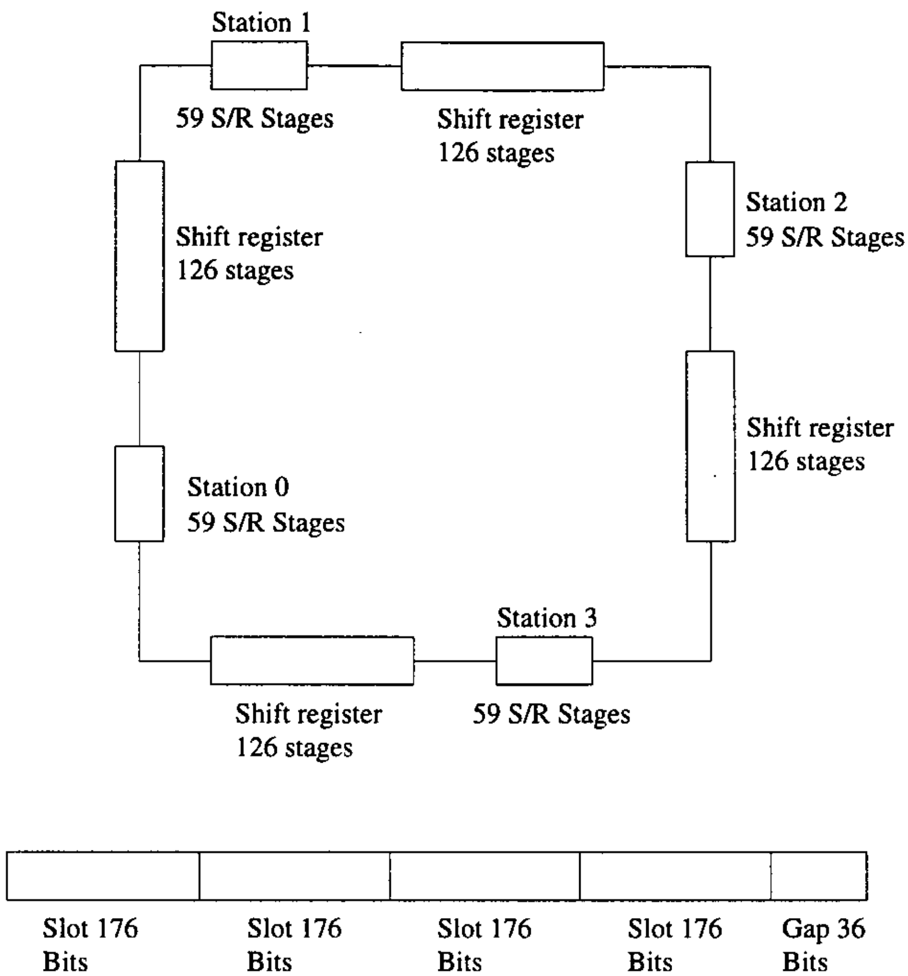

Figure 3.5 Organisation of Testbed Ring

The ring hardware is designed to operate up to 2MHz when encoded, or a data rate of 1 Mbit/s. In the simulations described in chapter 6 however, a data rate of 500 kbit/s was used to allow the traffic generator/analysers more processing time. This data rate gives a round ring propagation time of 740us.

3.3.3 Media Access Protocol of the Testbed

The structure of the testbed slot is shown in figure 3.7

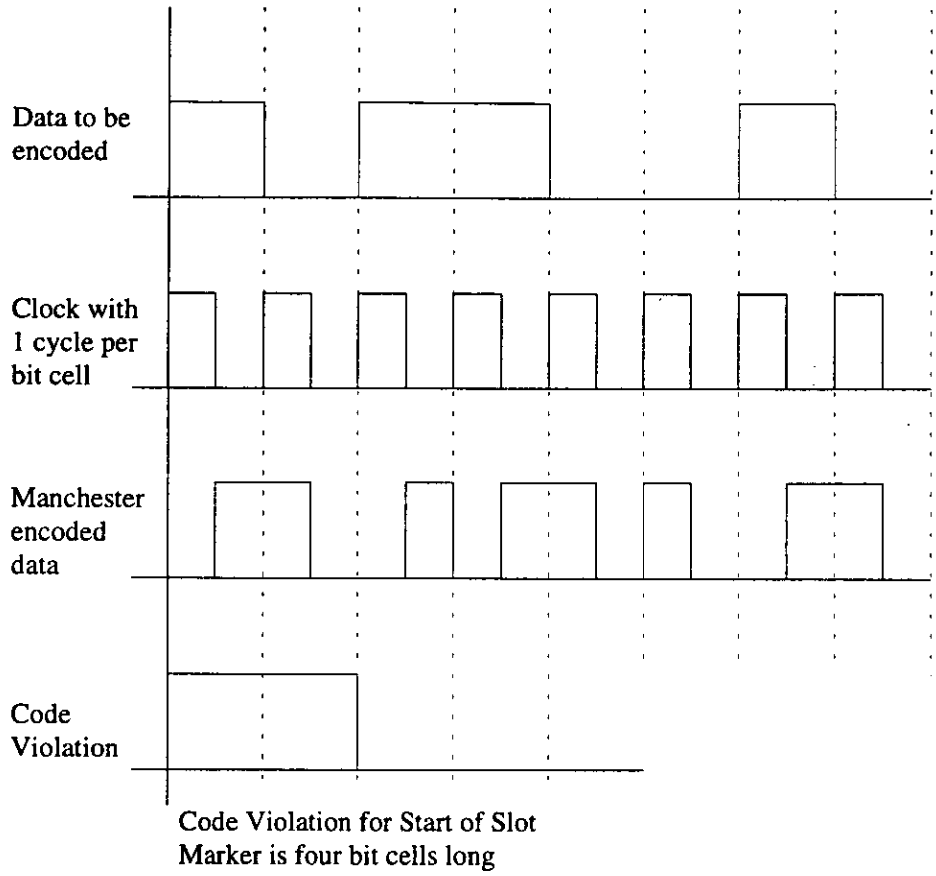

Figure 3.6 Manchester Code and the Start-Of-Slot Marker

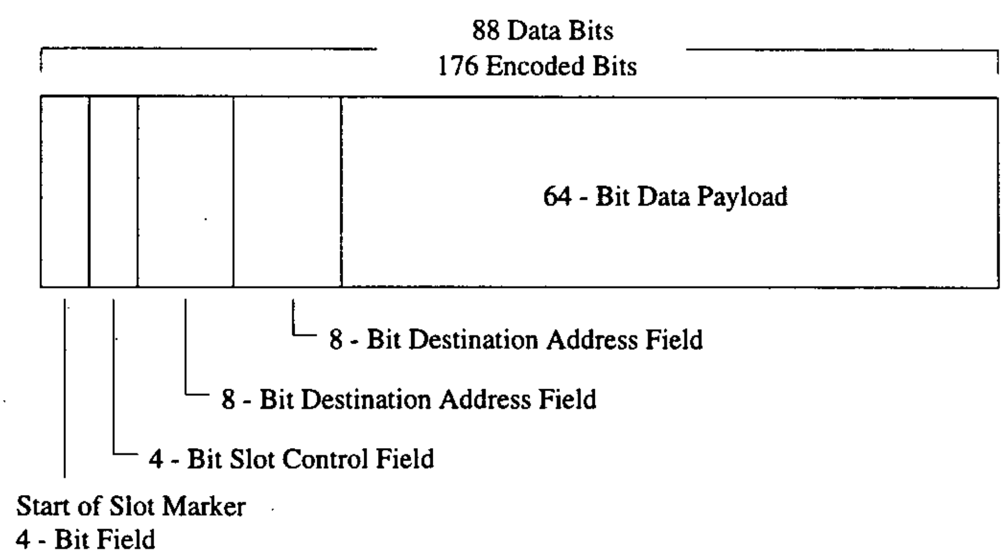

Figure 3.7 Testbed Slot Structure

The slot header consists of the 4-bit start of slot marker followed by 4 control bits. The slot will then contain the 80-bit packet of which the first 8 bits are to be the destination address. Figure 3.5 shows the slot structure of the whole ring which supports four slots, and a gap between the end of the fourth slot and the beginning of the first slot. In the figure, the bits are encoded data bits and are equivalent to one shift register cell each. The preliminary ORWELL slot structure was similar with a 20-byte packet and 4 control bits [83]. Later specification of ORWELL [38] gave it a 5-byte header and 48-byte payload in common with the ATM packet. It was not the purpose of the testbed to be fully ATM compatible so its specification was not changed. The 4 control bits labelled D, S, M, T, are used to implement the testbed's media access protocol. The bits D, and S, are used to define the type of slot, as shown in table 3.1.

Table 3.1 Slot Control Bits D and S

D	S	Description of Slot Type
0	0	Empty or Trial (depends on bit T)
0	1	Reset
1	0	Full

Bit T is used to indicate that the slot is a TRIAL slot, and bit M is used by the monitor to delete slots with corrupted destination addresses. If a FULL or RESET slot has its destination address corrupted, it will circulate the ring forever, unable to find its destination, and unable to be overwritten by any other station. To prevent this happening, the M bit is set to 1 each time a FULL or RESET slot passes the monitor station. If the slot reaches the monitor by circulating again it will be detected as its M bit will be '1' and it will be rendered EMPTY. A flowchart of the Media Access Protocol is shown in figure 3.8. It is similar to the ORWELL protocol in most respects but one important difference is that a station may re-use a slot which it has just received and stripped. Because the testbed protocol is implemented on PLDs, this

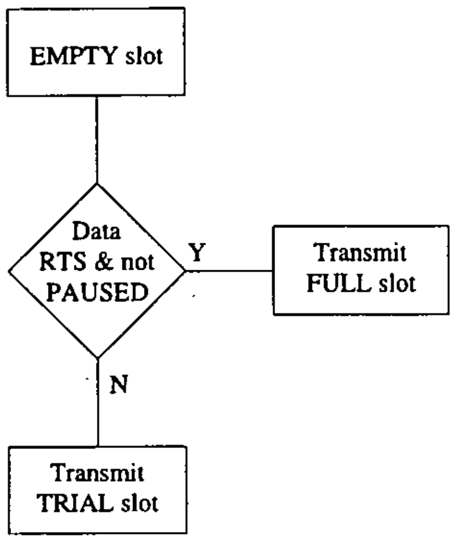

MONITOR action

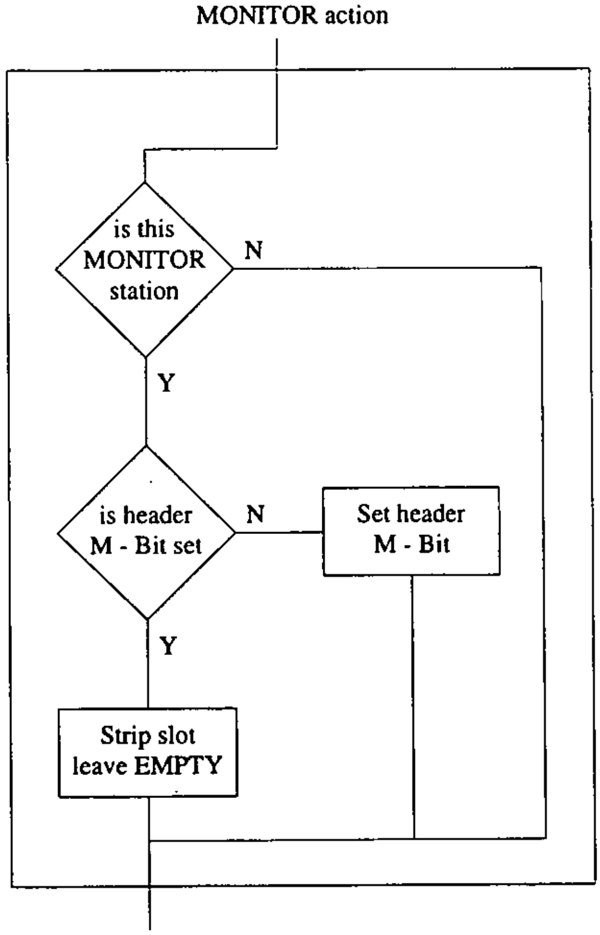

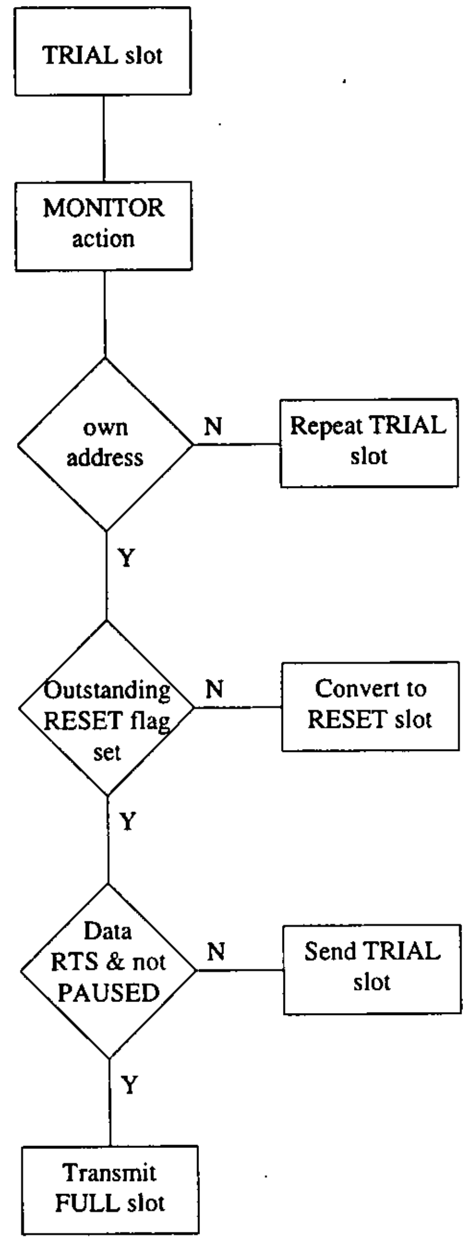

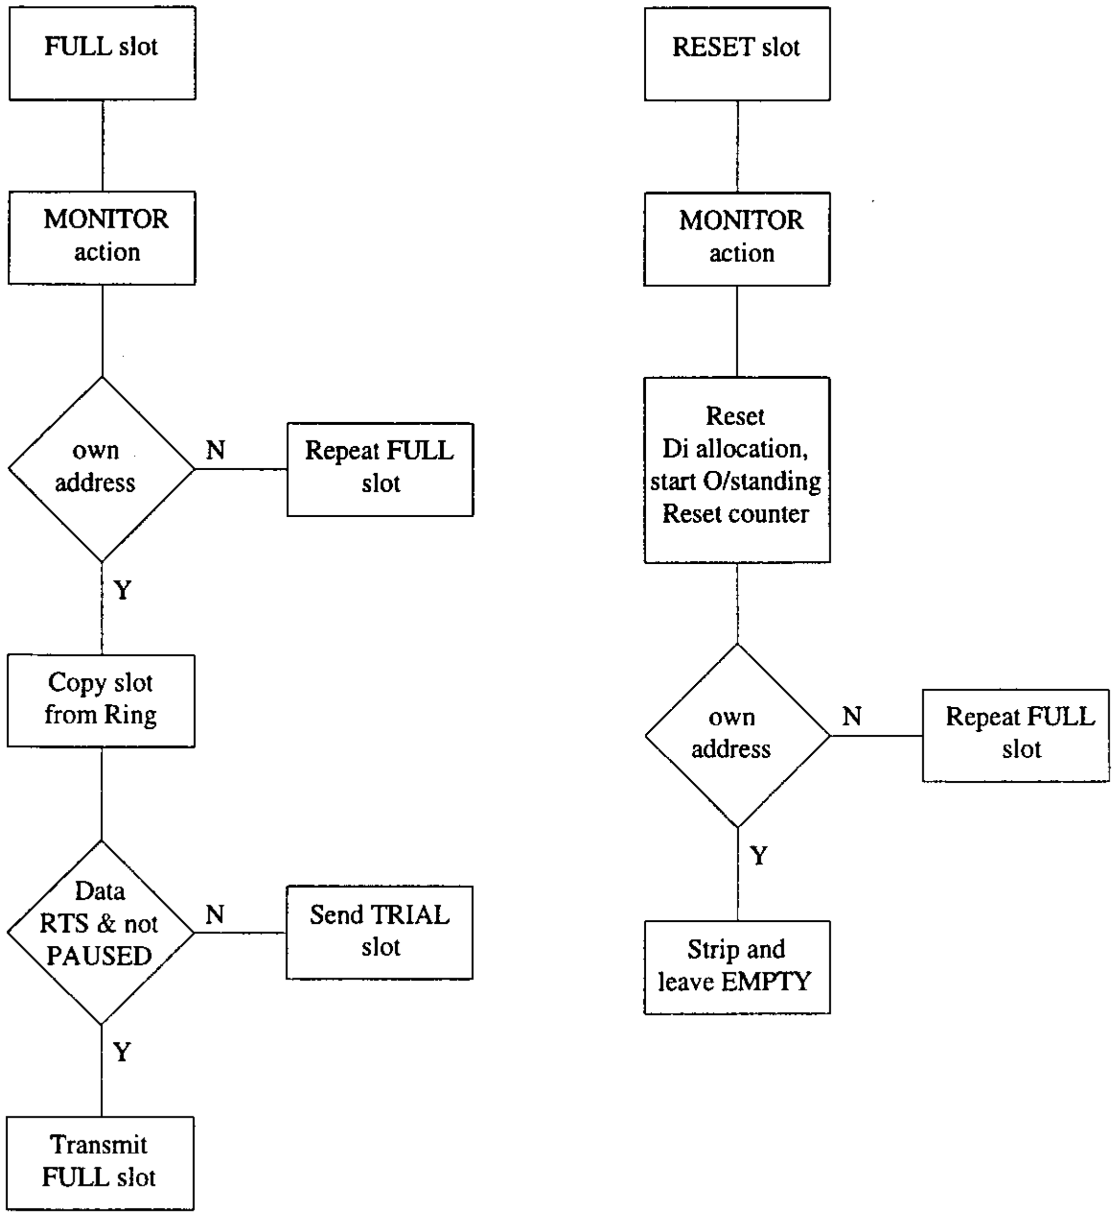

Figure 3.8 Flowchart of the Testbed Media Access Protocol

feature and others of the protocol can be changed with ease. The PLD software is written in a 'C' type language called TANGO [97]. TANGO is similar to the industry standard ABEL [98], both are logic compilers allowing a relatively high level description of logical operation. The source files for the PLDs are contained in the departmental document for this project [99].

A block diagram of the ring-interface card is shown in figure 3.9. The serial section contains the shift-register stages in line with the ring, it includes decoding and encoding of the Manchester encoded data, and interfaces directly to the FIFOs of the packet-level-interface card. The received slot header is decoded by the header-decoder block identifying the slot as FULL, EMPTY, TRIAL, or RESET, and comparing the cell address with the station address to see if the cell should be received. The new slot header is multiplexed into the serial stream for the transmitted slot.

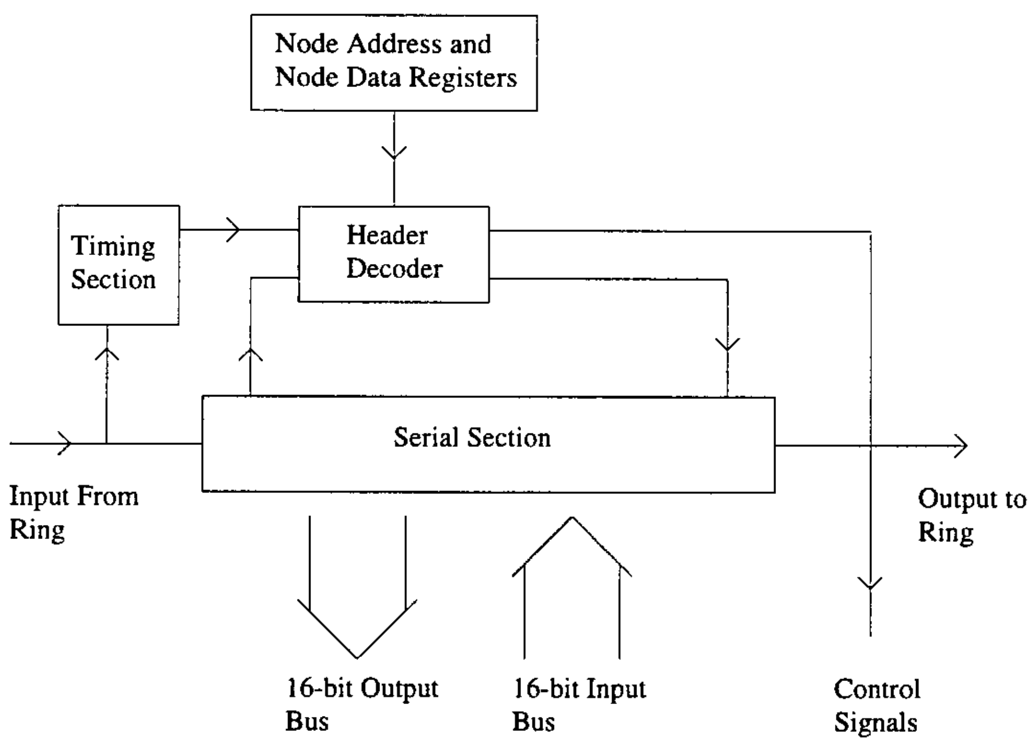

Figure 3.9 Block Diagram of the Ring Interface Unit

The timing section includes the start-of-slot detector and synchronises data transfers within the circuit. The header-decoder interprets the slot header and sets flags to control data transfers from ring to FIFOs and vice-versa, it also interacts with the various registers and counters contained in the register-section to allow the hardware to implement the media access protocol described. Some of the counters and registers are described below to illustrate how the protocol is implemented.

Node Address Register:	Contains the station's address for comparison with incoming slot destination address field.
Node Data Register:	Contains the station's allocated D-count (4-bits), the number of slots on the ring (3-bits), and a bit to identify the monitor station.
D-Counter:	Loaded from the Node Data Register and decremented each time a slot is filled until it reaches zero when the PAUSED flag is set.
Outstanding Reset Counter:	Used to prevent multiple resets occurring. When a reset is received, the Outstanding Reset Flag is set inhibiting further resets, and the Outstanding Reset Counter is loaded with the number of slots on the ring. It decrements to zero as each new slot arrives until, at zero, the Outstanding Reset Flag is reset.
Initialise Counter:	This counter is used by the monitor node to lay down a pattern of slots on the ring at initialisation.

Access to the hardware implementing the protocol allows the detailed analysis of slots arriving at a station by use of a logic analyser. This is a feature normally unavailable in software simulations, and consequently they do not offer this level of detail. A simulation of the node hardware was written in PASCAL prior to the hardware design of the ring-interface card. Many of the details of timing, and decoding of the slot header, were worked out at this stage. The simulation was extended as far as a single station on a three slot ring, but because of the length of time required to process at this level, the simulation was dropped in favour of a hardware testbed. As the prototype ring-interface card was built, it was tested by replacing software procedures from the PASCAL program by the prototype hardware interfaced through a PC prototype card using TURBO PASCAL 'port' procedures to read and write data to TTL latches [100]. A listing of the PASCAL simulator is given in the departmental document for this project [99]. A photograph of the ring-interface card is shown in plate 3.1.

3.4 Packet-level Interface for the Testbed

The Ring-Interface Card which implements the testbed's media access protocol (as described in the preceding section) interfaces the data on the ring to two 16-bit wide FIFO buffers. The packet-level interface card provides an interface between the FIFOs and the traffic generator/analyser software running on the T800 transputers as shown in the testbed architecture diagram figure 3.4. Connections are made between the T222 and T800 transputers using 10Mb/s serial INMOS links.

The functions required of the packet-level interface are:

- To initialise the ring-interface card and provide timing synchronisation to the traffic generator/analysers.
- To maintain packet level communications over the testbed ring,
- To ensure FIFO overflow does not occur,
- To monitor the ring's reset rate, and pass this information to the T800,
- To monitor the output buffer length (number of packets in the FIFO), and pass this to the T800.

3.4.1 Packet Transmit and Receive Processes

The packet-level interface card is based around an INMOS T222 transputer [95] which has a 16-bit data bus and four bi-directional hardware links. A block diagram of the card is shown in figure 3.10, and a photograph is shown in plate 3.2. The ring-interface cards and packet level interface cards combined in the testbed are shown in plate 3.3.

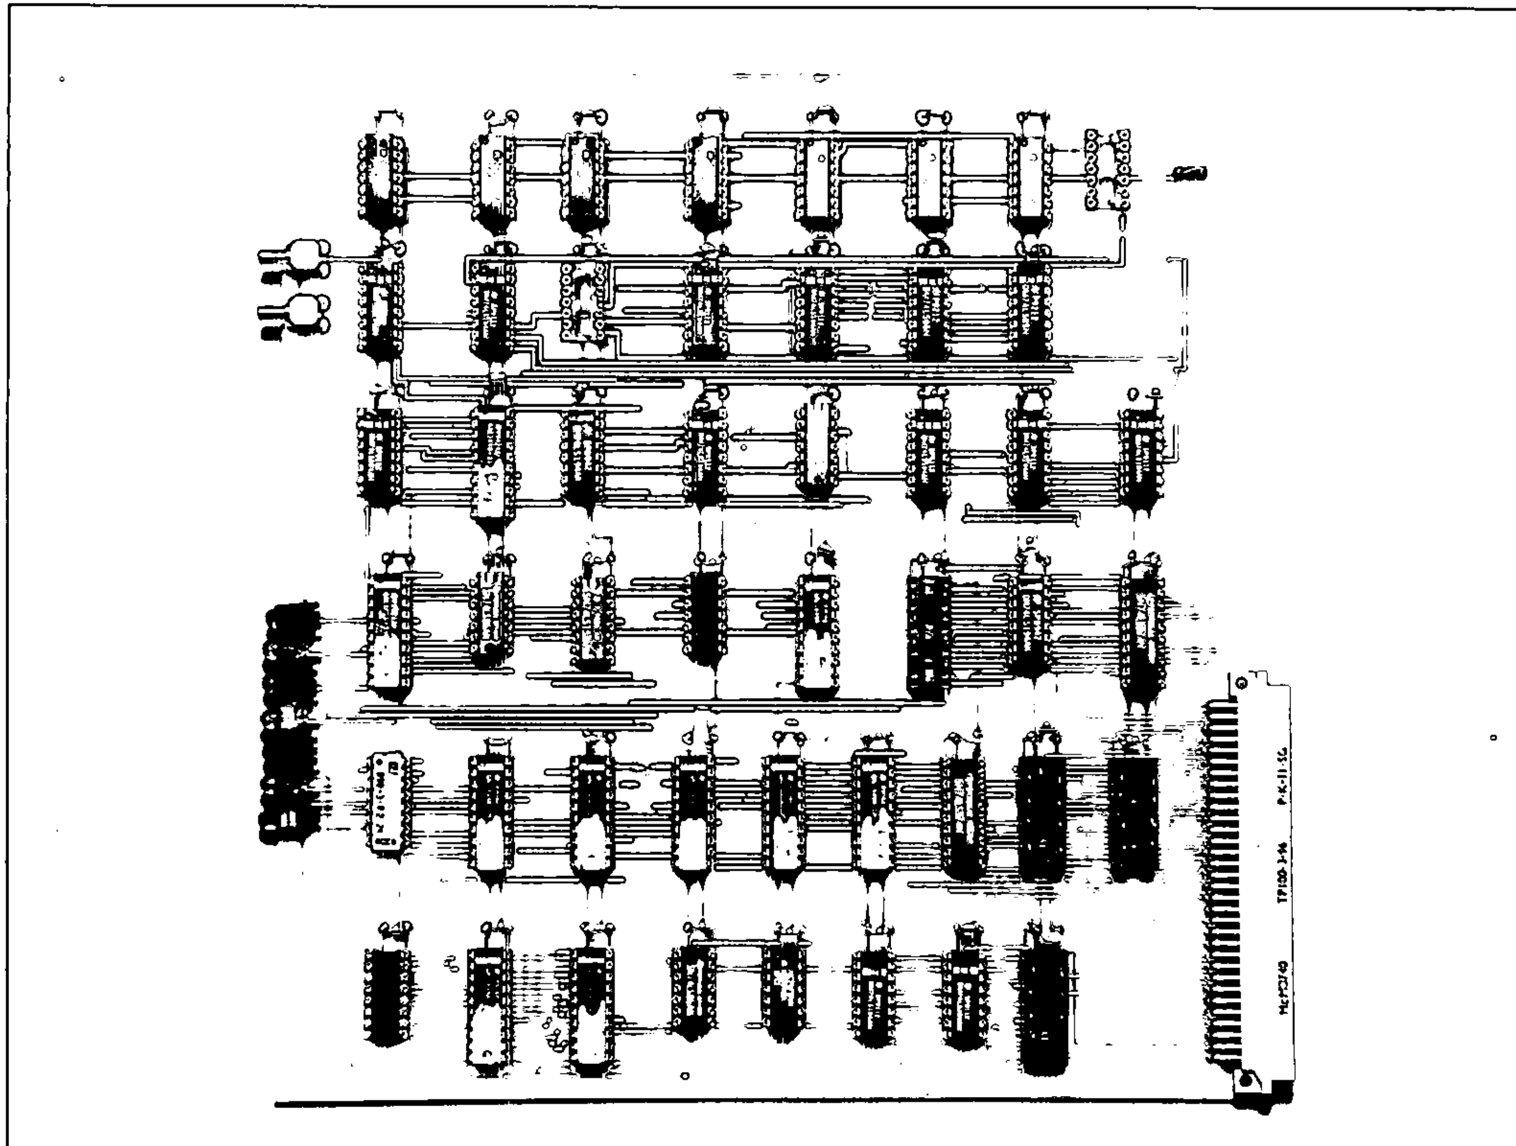

Plate 3.1 Ring Interface Card

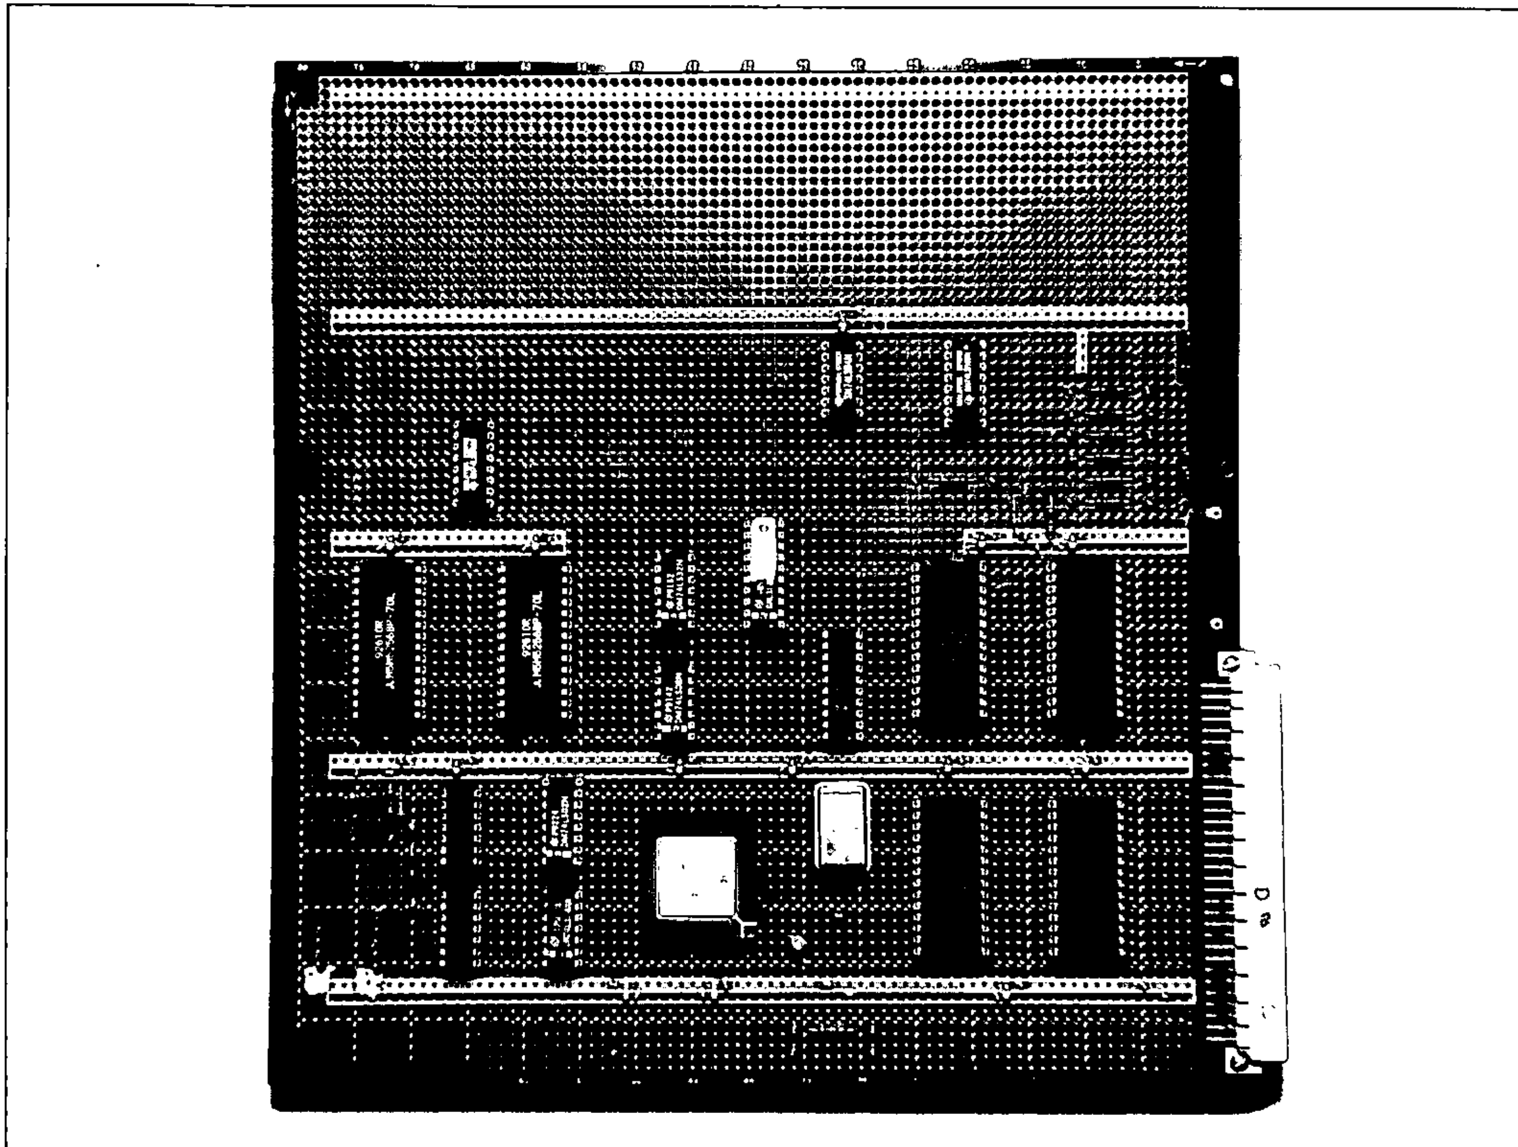

Plate 3.2 Packet-Level Interface Card

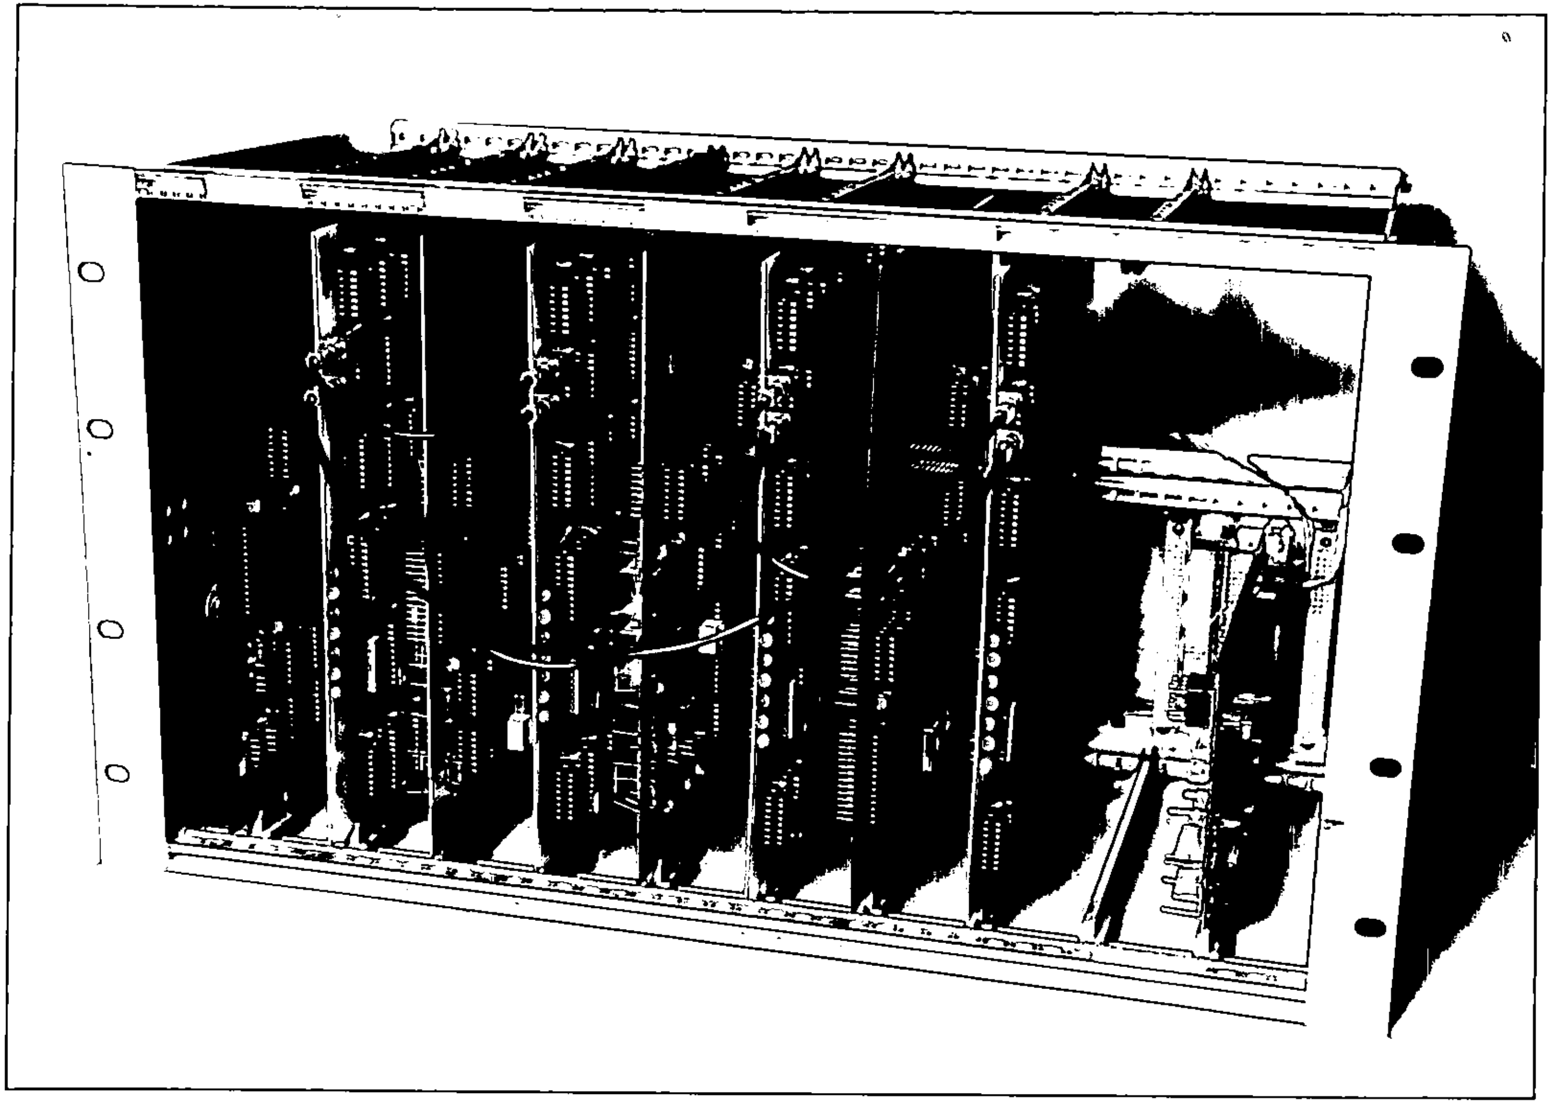

Plate 3.3 Slotted Ring Testbed

Only one of the four INMOS links is used to interface to the T800 traffic generator/analyser transputer which resides on a plug-in PC card [101]. Packets are transferred on these links as an array of 5 x 16-bit words, the first 8 bits of which give the address of the destination station. Further details of the packet structure are given in chapter 6.

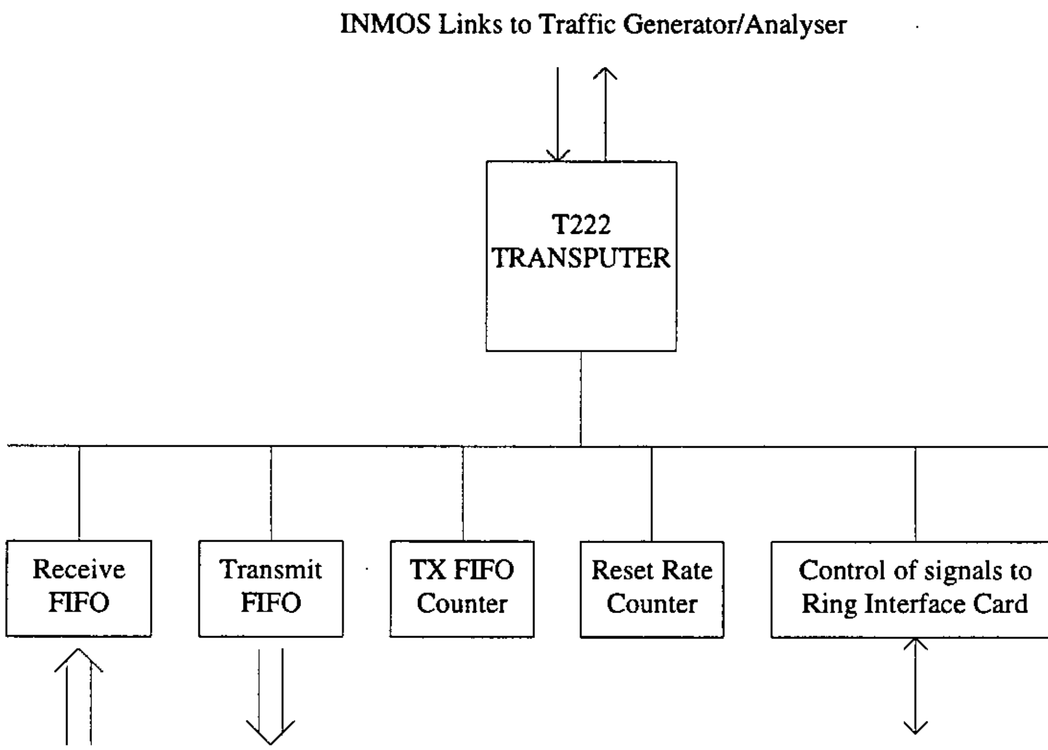

Figure 3.10 Block Diagram of Packet-Level Interface Card

At the start of a simulation period, the ring-interface card must be initialised with information in its node address register and node data register, and a pattern of slots laid down on the ring. The required information is passed down a pipeline of transputer links from the root transputer running the simulation control program. Transputer pipelines and programming in OCCAM are discussed in section 3.5. The controlling transputer generates a timing initialisation signal which is connected to the interrupt lines of all T222 transputers. This interrupt causes the T222 devices to communicate immediately with the T800 devices which are programmed to read their own internal timer values and store the result. In this way the T800 traffic generator/analysers are synchronised to within 1us of each other at the start of the simulation. The T222 also reads the reset rate counter and the output buffer counter to obtain initial values. The reset rate counter is incremented by the arrival of a valid reset slot, and the output buffer counter is incremented when a word is read from the buffer by the ring-interface card.

Once initialisation is complete, packet transmission can proceed. The T222 waits for the arrival of a packet on its INMOS link. Once a packet has arrived, the output buffer counter is read and compared with previous data to calculate the output buffer length. If the specified length is exceeded, the received packet is discarded as it is considered to be lost due to buffer overflow, otherwise the packet is written word by word into the output buffer. The effective length of the output buffer can be dimensioned using software and this can be used to control the maximum packet delay, as discussed in section 6.2.2. The addition of data to the buffer is noted for the calculation of output buffer length.

In parallel with the transmit process, a receive process polls the empty flag of the input buffer to detect if data has been written to it from the ring-interface card. If it detects data present it waits for one slot duration to ensure the entire slot has been written into the buffer and reads five 16-bit words from the buffer. The individual words are assembled into array format and output to the traffic analyser.

3.4.2 Reset Rate monitoring

The reset rate parameter is an indication of loading on the network. At low levels of applied traffic, most stations will be idle and emitting TRIAL slots which are converted to RESET slots when they circulate back to the originating node. As traffic levels increase, the probability that the TRIAL slots will be overwritten with data increases, and the time between resets increases, decreasing the reset rate. Each time a valid RESET slot is received at a node, a flag is set and this is used to increment the reset counter. To obtain a value for the reset rate, the counter is read at regular time intervals and the difference between readings used as the number of resets received in that interval. The interval between readings of the counter must be less than the (minimum reset interval) x (maximum counter value). If the minimum reset interval is taken to be the ring rotation time, the 8-bit counter is modulo 256, so the maximum interval for reading the counter is 740us x 256 = 190ms. The minimum interval for reading the counter depends on the degree of averaging required, but for a reset count of at least 1 to be guaranteed for a D_j allocation of 15, (the maximum possible allocation for this testbed), the interval should be 11ms. An interval of 30ms was selected for the testbed.

3.5 Programming in OCCAM

OCCAM is the natural language to choose when programming the transputer family of microcomputers since it was developed as a high-level language specifically for the Transputer. The language is designed to support parallel processes which can either run on the same processor by a user-transparent scheduling system and communicating via software channels, or on multiple processors communicating via hardware channels which are in fact 10 Mbit/s serial links (INMOS links). Transputers have on-board hardware timers accessible to the user, and an event input.

The OCCAM language defines three types of process: SEQ, PAR and ALT [133]. SEQ (sequential) processes are similar to the sequential programs written in non-parallel languages. They may contain assignments, conditional statements and loops. the PAR statement is used to make two SEQ processes run in parallel. For example:

```
PAR
  SEQ
    ...process 1
  SEQ
```

```
    ...process 2
```

is used to make process 1 and process 2 run in parallel. The execution of this stage of the program will not finish until both process1 and process 2 have terminated.
The ALT statement is used to accept inputs from a number of channels or timers on a first come first served basis.

Consider the following example:

```
ALT
  chan1 ? var1
    ...process var1
  chan2 ? var2
    ...process var2
```

In this case the channel which receives data first will be serviced first, and when its service procedure is completed, the competition will start again. Further information on OCCAM programming can be gained through 'An Introduction to OCCAM Programming' [102].

OCCAM is a high-level language, and is strongly typed, but it maps very well to the transputer's instruction set and is considered to be highly efficient. The ability to specify independent, parallel-running, transmit and receive processes was very useful for this application where packets are arriving from the ring-interface card and from the traffic generator in an unpredictable order and with an unpredictable distribution. In the packet level interface there are three parallel processes, as shown in figure 3.11.

The transmit and receive processes have been described in section 3.3, and the third process is the reset rate counter which is initiated by a 30ms interval timer. The reset rate monitor is given high priority because of the need for an accurate reset rate measurement. Communication between the three processes is carried out by OCCAM channels. Parallel processes operate asynchronously with each other, and may not have shared variables as the variable could be changed at unpredictable instants in time by different processes. To communicate, an outputting process sends a variable down a channel which can by inputted by the receiving process at a point in its program specified by the user. Only at channel communication statements do the two parallel processes synchronise, and once data has been transferred, they continue asynchronously. There is a danger of deadlock occurring when parallel processes are both trying to send data to each other or receive data from each other; they cannot synchronise, and cannot continue program execution until they do synchronise. This problem can be avoided by good programming.

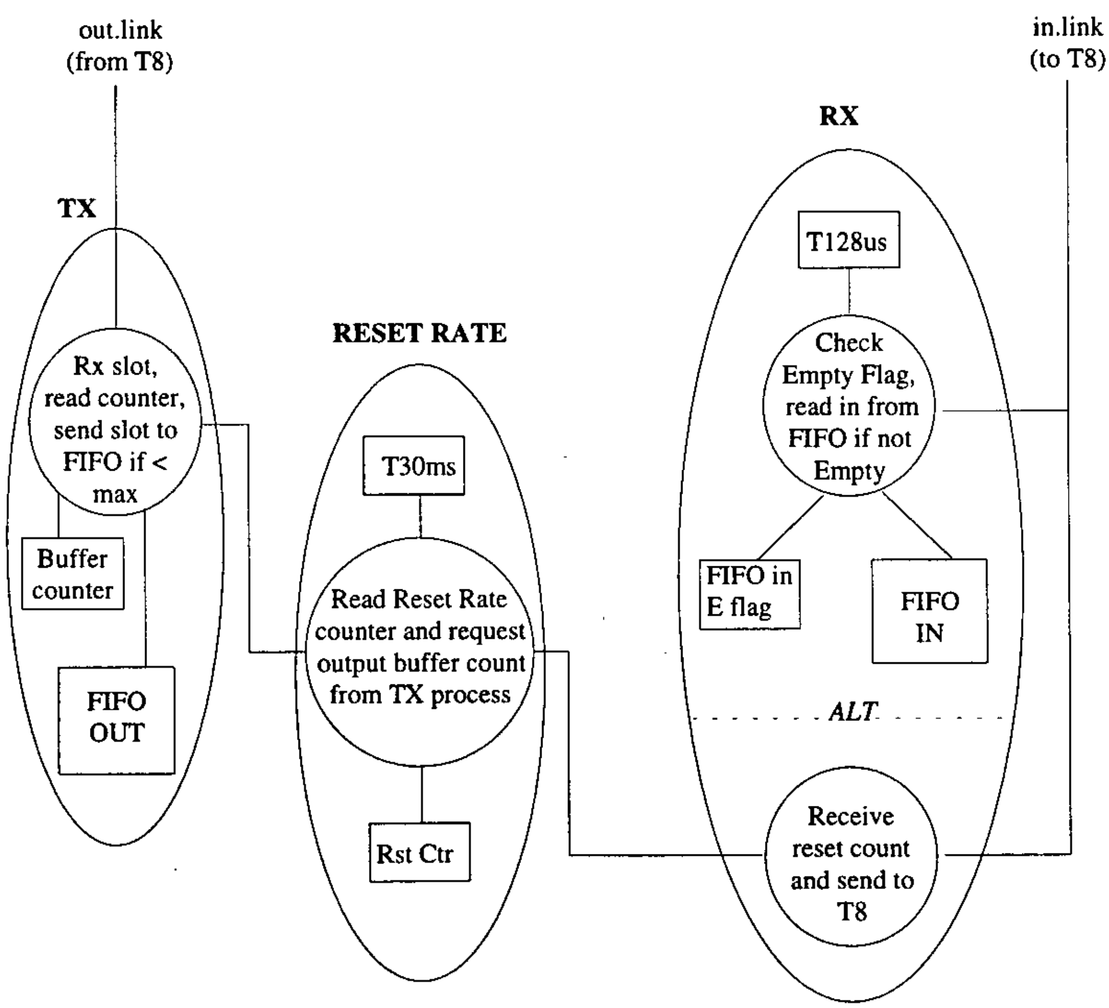

Figure 3.11 Packet-Level Interface Software Processes

In Figure 3.11, the communication channels to the T800 traffic generator/analyser are INMOS links. A software channel allows the reset monitor to pass its data to the traffic analyser by going through the receive process. In this case the receive process contains an OCCAM ALT statement to allow input from the reset rate monitor process and a timer controlled process looking for data arrivals. Whichever process communicates first will be served first. A full listing of the OCCAM software used in the simulator can be found in the departmental document [99].

3.6 Summary

The testbed architecture has been the major area discussed in this chapter. The choice of a slotted ring local area network, and ORWELL protocol, was considered to be appropriate for integrated services traffic. ORWELL had generated interest as an ATM network protocol, particularly as the reset rate offered a means of determining the level of traffic on the network independently at each station. The hardware interface between the ring and ATM cells was explained in detail, as were the software processes used to generate and analyse cell delay and loss. Measurement of the ORWELL reset-rate, and the output buffer occupancy was also described. The generation of traffic for the testbed is described in chapter 4, this process determines the rate and statistical distribution of cells generated for transmission on the network.

Chapter 4

ATM Traffic Characteristics and Traffic Modelling

4.1 Introduction

To carry out analysis or simulation on an ATM network in order to investigate some aspects of performance, it is necessary to make assumptions about the arrival pattern of ATM cells at the network inputs. It is desirable to be able to model either specific types of traffic likely to be present on such a network: voice, video or data traffic perhaps, or to model a multiplexed stream of ATM cells, which also requires knowledge of the characteristics of various traffic types. Because of the importance of accurate and efficient source modelling, a lot of work has been carried out in the field of characterising and modelling different traffic sources [45] [46] [47] [103].

The performance of an ATM network, and of a media access protocol such as ORWELL is dependent on the statistical characteristics of the traffic applied to it. Two obvious measures of traffic intensity are peak cell arrival rate, and mean cell arrival rate. These parameters alone do not fully characterise a traffic source, and some measure of the source burstiness is required. Several different methods of measuring burstiness are proposed in the literature, from cell rate variance [104], to maximum burst length [39]. The following section deals with the characteristics of various traffic sources, and models for these are outlined in section 4.2. The implementation of the traffic models on the network are discussed in section 4.3.

4.2 ATM Traffic Sources and Their Characteristics

In this section, the characteristics of packetized real-time data sources i.e. voice, and video sources are discussed together with the characteristics of non-real-time traffic such as file transfers, and LAN interconnection.

4.2.1 Voice Traffic

Voice traffic in an uncompressed form can be considered to be a constant 64 kbit/s data source. It is likely that on an ATM network, voice applications will make use of the fact that there are considerable periods of silence in voice transmission, particularly in a two-way conversation when there is normally a talker and a listener. The voice transmission can be characterised by periods of silence and by talkspurts, and bandwidth savings can be made by transmitting cells during talkspurts and not during silences. Two techniques are used to bridge periods of silence known as Hangover, and Fill-In. [105], [106]. Voice sources therefore, are likely to be either constant bit rate where there is no compression of data, or modulated constant bit rate, although more sophisticated methods of voice compression are available combining several techniques to reduce bandwidth requirements further [107]. The bandwidth of individual voice connections is relatively low compared to the bandwidth of a basic ATM path, and many connections are likely to be multiplexed in to one virtual path giving statistical averaging of the stream of cells. A 155.52Mbit/s (SDH1) transmission interface to ATM provides bandwidth for about 1200 pairs of 64kbit/s voice connections.

4.2.2 Video Traffic

Video traffic has a far higher bandwidth than voice traffic, in the tens of Mbit/s region before compression, and its statistical characteristics exhibit burstiness and autocorrelation features. Video codecs encode sequences of video images, and the ensuing data stream statistics depend on the type of codec used. In general there are three types of autocorrelation based on the repetitive nature of the video line, frame, and scene.

Line correlation occurs because one part of an image is highly correlated with data on the same part of the next line (spatial correlation). Frame correlation occurs because data at one part of an image is highly correlated with data on the same part of the next image (temporal correlation). Scene correlation occurs because sequences of images within a scene may be related to each other to a greater or lesser extent, more so in slowly moving scenes. If a video codec does not compress images it will provide a constant bit-rate output stream. The more highly compressing the codec is, the more variation in the output bit rate will occur, the statistical effect of which is a white noise like variation in the cell output rate from the codec. This white noise is a function of the type of coding scheme used, and is also subject to autocorrelation.

The amount of buffering which can be applied to the video signal determines how far the correlation and burstiness can be smoothed out. Non-frame buffered video codecs have line, frame, scene and white noise correlation. Frame buffered video codecs have scene and white noise correlation only, and multi-frame buffered video codecs remove some of the scene correlation. For distributed video the amount of buffering is relatively unimportant provided both user and sender have sufficient memory for the scheme being implemented. For interactive video such as video telephony and video conferencing, voice and video data streams are delay sensitive and buffering must be restricted. Figure 4.1 illustrates the autocorrelation of frame buffered and non-frame buffered video codec output.

The non-frame-buffered video exhibits peaks in its autocorrelation response at the frame repetition interval of 40ms. In both cases the scene correlation decreases in an exponential manner with time.

In characterising the traffic statistics of video codecs, the image coding methods of the codec must be taken into account. Some of these methods are interframe Differential Pulse Code Modulation (DPCM), intra frame DPCM, Conditional Replenishment (CR), Motion Compensation (MC), and various methods of transform coding, such as Discrete Cosine Transform (DCT). Run Length Coding (RLC), and Syntactic Coding are also used to compress the required bandwidth of the video signal.

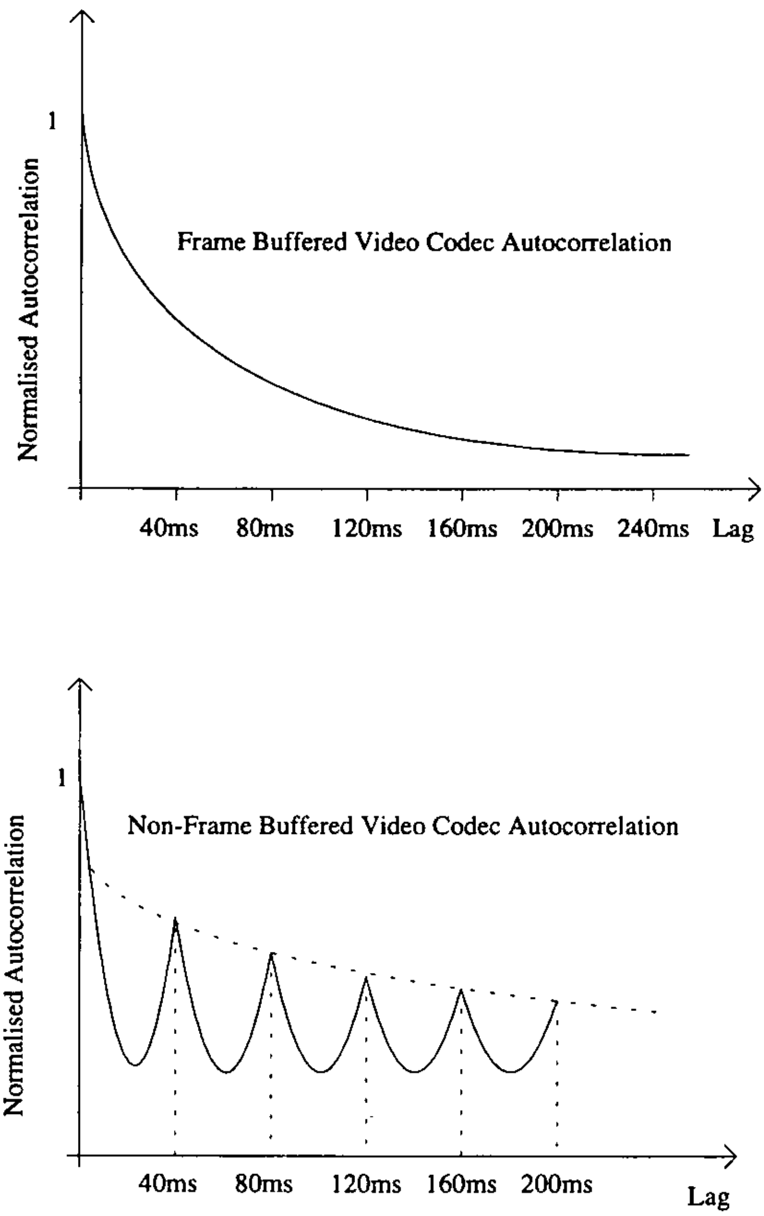

Figure 4.1. Normalised autocovariance for a Frame buffered Codec and a Non-Frame Buffered Codec.

Many studies of video coding methods have been pursued [108], [109], but ITU-T has recommended 2 video codec schemes for standardisation which are:

1. The conditional replenishment video codec, H.120 [110], [111], which detects and transmits areas of movement between frames.
2. The discrete cosine transform with motion compensation, H.261, MPEG [110], [112] which compresses the spatial frequency components of motion compensated successive frames.

These two video codecs exhibit line, frame and scene correlation. The objective of analytical models of video codecs is to introduce the autocorrelation features of the coding schemes at 40ms intervals for frame correlation and

40ms / (number of lines)

for line correlation. The autocorrelation has been shown to decrease approximately exponentially with temporal separation over a scene.

4.2.3 Data Traffic

The plethora of existing and future data sources, that is to say sources that are not time-delay critical, makes a general model difficult to generate. Low bit rate sources such as facsimile and electronic messaging are likely to be easily and efficiently handled by an ATM network, but the inter-connection of LAN networks and high-resolution image transfer require greater data rates. In the future it is predicted that the interconnection of computer based LANs, or multimedia LANs based on ATM techniques themselves, will be a major source of traffic [45], [113]. A study of Ethernet traffic was made by Shoch and Hupp [114] in 1979, and more recent data is available from Fowler and Leland (1989) [45]. The studies show that there is variation in traffic offered to a LAN of seasonal and hour to hour basis in the same way as telephone traffic has a busy hour. At this level, variations in traffic are of more interest to engineers dimensioning the transmission capacity of a network. Within any hour, the fluctuations in traffic arrival from minute to minute can be modelled by a Poisson type arrival process as has been the practice when analysing computer data networks. The crucial part of the study, and the most problematic for ATM network designers, is that at a milli-second level, traffic arrival is highly bursty being characterised by the arrival of packet trains. A packet train is a closely spaced sequence of traffic cells with the same source and destination. This type of bursty data could be expected, as for instance, files are downloaded from a file-server to a desktop computer.

The ratio of peak to mean traffic intensity when considered at high time resolutions may vary much more than that seen when averaging over an hour. A variation of network utilisation of 10% to 40% is not uncommon in day to day operation of a LAN. The peak traffic intensity to mean traffic intensity ratio in this case would be from 10 to 2.5. In the study by Fowler and Leland, it is reported that when considered over any 5 second interval the peak to mean ratio was 152:1, and taken over any 5ms interval the ratio was 715:1. This degree of burstiness cannot be modelled by a Poisson process.

The characteristics of data transmission depend very much on the actual sources and the applications running, however, high burstiness at the cell level is a key feature of this traffic type. Methods of providing a traffic model will be considered in section 4.3.

4.3 Modelling Traffic Sources

A traffic source model should be designed to have the same statistical characteristics as the real source for those characteristics expected to influence the performance analysis or simulations. By measuring the statistical characteristics of the model and comparing them to a real source the model's accuracy can be verified. In the development of source models it can be desirable to provide as general a model as possible which can cover a range of traffic streams by altering model parameters, rather than developing a new model for each source [115]. The complexity of the model may put restrictions on simulation or emulation depending on the processing power of the simulating machine or testbed. This section covers some of the methods used to simulate traffic sources.

4.3.1 The Bernoulli, Binomial and Poisson Processes

The Bernoulli process is a discrete time process where at each time interval there is a probability p that a cell arrival has occurred, and a probability q = (1-p) that no cell has arrived. The Bernoulli process is a memoryless process as the probability p is independent of any previous outcomes of the process. When n Bernoulli trials are conducted, the probability distribution formed is the binomial distribution. The probability density function of the binomial distribution is given by

$$f(x) = (n,c)p^x(1-p)^{n-x} \tag{4.1}$$

The mean of the distribution = np and the variance = $np(1\text{-}p)$ [116].

The inter-arrival times of a Bernoulli process are geometrically distributed. If the number of trials tends to very many and the time between trials tends to very small, the corresponding continuous time probability distribution function formed for the probability of k arrivals in the time interval t, where λ is the mean arrival rate is the Poisson distribution:

$$f_k(t) = \frac{(\lambda t)^k}{k!}e^{-\lambda t} \tag{4.2}$$

The interarrival times for the Poisson distribution are exponentially distributed, which can be seen by calculating the probability of zero arrivals, which decreases exponentially with t, that is

$$f_0(t) = e^{-\lambda t} \tag{4.3}$$

The Poisson distribution has been used to model call request arrivals in teletraffic analysis when there are a large number of independent uncorrelated sources [12], and in ATM systems is used in modelling cell arrivals from a large number of uncorrelated sources.

Both Bernoulli trials and the Poisson process are members of a group of stochastic processes known as renewal processes. The Bernoulli process can be extended to batch arrivals where either fixed size or variable sized batches are allowed to occur in each time interval. In this case the average arrival rate will be pN, where N is the mean batch size. This may be an appropriate way to model the arrival of packet trains associated with the transfer of computer data. Batch arrivals can also be dealt with by the Poisson process [117].

4.3.2 Markov Chains in Source Modelling

A Markov chain is a stochastic process linking a number of distinct states. Figure 4.2 shows an example of such a chain.

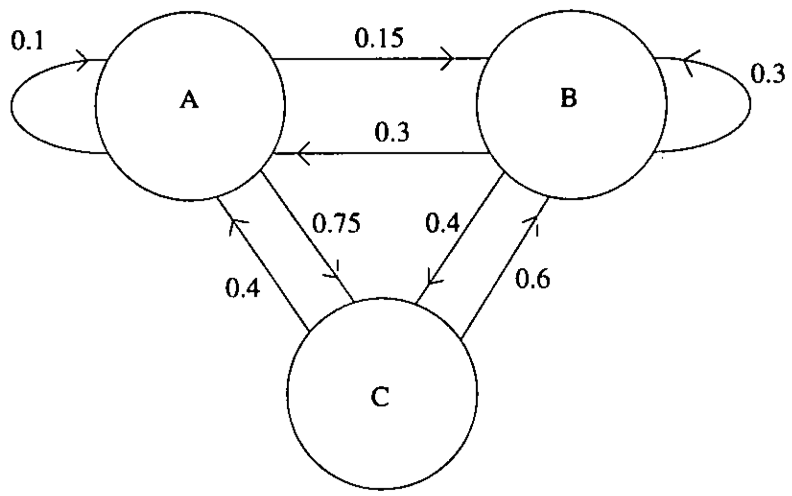

Figure 4.2 Markov Chain

The probability of transition from one state to another is defined on the arrows linking the states. If transitions can only occur at discrete intervals, the chain is known as a discrete time Markov chain. The probability that the next state will be any given state depends on the present state and the transition probability alone. The Markov process has the ability to remember only the present state. Analysis of the Markov chain [117] shows that there is a stationary equilibrium probability distribution for the probability of being in any one specified state in the chain, which is independent of the initial state after a length of time.

Markov chains have been used to assist in the modelling of traffic sources. The burst-silence model of a voice source has been proposed [41], [48] to model the talking and silent periods in a conversation using a two state Markov chain shown in figure 4.3.

One state represents talkspurts or bursts, and the other state represents silence. Within the talkspurt state, cells are emitted at a constant rate, and no cells are emitted during the silence state.

If The probability of transition from burst state to idle state = b, and

The probability of transition from idle state to burst state = a, and

the cell emission rate in the burst state = P,

then the average burst length B=1/b,

and the average cell rate A = aP/(a+b)

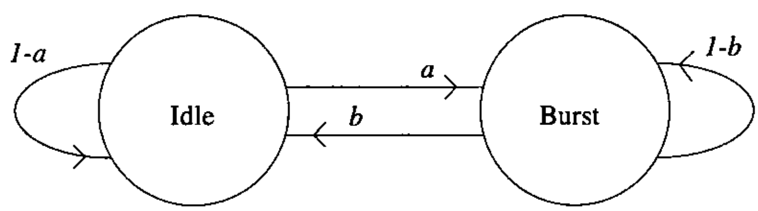

Figure 4.3 Two State Markov Model of a Voice Source

This model can be used to simulate sources of various levels of burstiness, and does not require high processing overheads making it suitable for the simulation of a number of multiplexed sources [41], [48], [118].

4.3.3 Auto Regressive Model for Video Codecs

Analysis of video codecs has shown that they can be characterised in the time domain by their bit-rate distribution, and bit-rate autocorrelation [119]. The bit rate distribution is approximated as a gaussian distribution, and autocorrelation reduces exponentially with time. A modelling function for such a video source is the Auto Regressive (AR) process which is defined as:

$$x(n) = x'(n) + E[x(n)] \qquad (4.4)$$

where,

$$x'(n) = \sum a(n)x'(n-m) + e(n) \qquad (4.5)$$

where $x(n)$ is the quantity of information generated in the nth frame; $E[x(n)]$ is expectation, $a(n)$ is the model parameter, M is the model order, $e(n)$ is a gaussian random process. Equation 3.4 shows that the information generated in the nth frame is the mean information quantity modulated by equation 4.5 which contains a Gaussian element for the Gaussian distribution, and an autocorrelated element based on a contribution from the previous M frames.

4.4 Traffic Generation in the Testbed

The first experiments to be performed on the testbed were to characterise the cell throughput and ORWELL reset interval, whilst varying network parameters such as *Di* allocation and symmetry of loading. At this stage there was not the intention of modelling a specific traffic source, but rather to investigate the network performance under a number of possible cell arrival patterns. ATM traffic streams can be characterised by their mean rate, peak rate, and a measure of burstiness. It was decided that in the network characterisation stage, traffic models which allowed these parameters of the offered traffic to be varied would be used.

4.4.1 Modelling the Cell Arrival Process

As has been described in section 4.1, the variation in statistical characteristics of ATM traffic sources is extremely wide from constant bit rate to highly bursty sources. The traffic arrival processes chosen to evaluate the testbed's performance are listed in Table 4.1.

At this stage, the modelling of autocorrelated sources was not considered because of the complexity of the traffic models. To obtain results comparable between different arrival patterns, the mean arrival rate for each cell arrival pattern was calculated.

Table 4.1 Traffic Models for Testbed Performance Evaluation

Traffic Model	Type of Source
Deterministic	Constant bit rate such as 64kbit/s voice streams
Bernoulli process	Traffic sources of low to medium burstiness such as multiplexed voice streams
Batch Bernoulli process	Bursty sources such as data traffic

For deterministic, Bernoulli, and batch Bernoulli processes, arrival intervals are discrete-time units (*Ta*). The mean cell arrival rate (λ *mean*) for these processes is,

Deterministic process: $\lambda \; mean = 1 / (Ta)$

Bernoulli process: $\lambda \; mean = P\{cell\ arrival\} / (Ta)$

Batch Bernoulli process: $\lambda \; mean = P\{batch\ arrival\}(mean\ batch\ size) / (Ta)$

4.4.2 Implementation of the Testbed Traffic Generation Process

The testbed traffic cell is shown in figure 4.4. The cell contains an 8 bit destination address, an 8 bit source address, a 32 bit time-stamp corresponding to the packet generation time, a 16 bit sequence number which is used to look for lost cells, and a 16 bit checksum to identify corrupted cells.

The OCCAM software processes for the T800 Traffic Generator/Analyser processor are shown in figure 4.5. The cell generation process is controlled by a timer, the time interval of which may be varied to change the traffic intensity of the discrete-time source models. Hence deterministic and Bernoulli type arrival patterns are readily implemented.

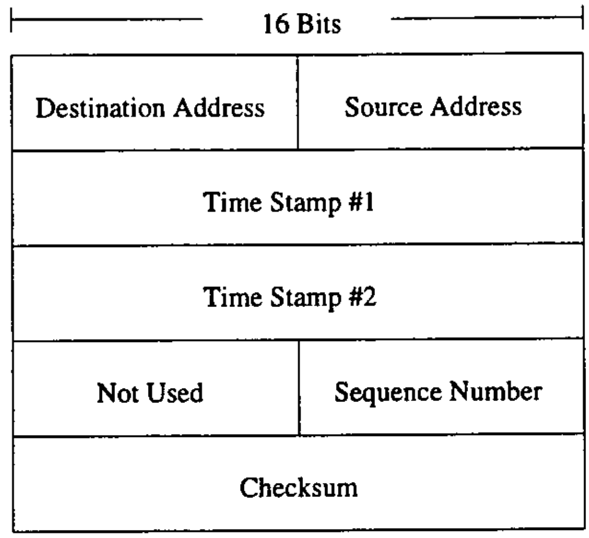

Figure 4.4 Testbed Cell Contents

4.4.3 Random Number Generation

One key requirement for the traffic generation process of the testbed is to be able to generate a random variable which can be used to determine cell arrival or non-arrival. Random number generators operate using a recursive relation in which the next number in the series is a function of the last one or two numbers. The series produced therefore requires at least one starting value which is known as the seed. The sequence produced is entirely deterministic, and the numbers produced are therefore pseudo-random although this is not necessarily a disadvantage as repeatability of a simulation experiment can be desirable. The series generated will have a cycle length where the same number pattern repeats itself. The desired properties of a generator function are as follows [116]:

1. It should be efficiently computable, requiring low processor time to generate.
2. The cycle length should be large so that the number sequence does not recycle.
3. The successive values should be independent and uniformly distributed. That is to say the next value should not be correlated to the present value, and there should be an even distribution of the generated numbers across all of the possible variable values.

The first two properties are easy to establish, but the third requires a number of tests, and it is generally better to use an established generator with proven performance rather than to devise a new one. In the Traffic Generator/Analyser for the testbed, a type of Linear-Congruential Generator (LCG) known as a Multiplicative LCG has been used. The formula is

$$x_n = 7^5 x_{n-1} \bmod \left(2^{31} - 1\right) \tag{4.6}$$

The random variable is stored as a 64-bit integer, and a different seed is used at each station, spaced 100,000 sequence elements apart using a table of values [116].

4.4.4 Analysis of the Testbed Cells

As shown in figure 4.4, the testbed cells contain both source address and destination address. The addresses define the station and the traffic class if there is more than one type of traffic source model employed. At the receiving station, the traffic analyser will compare the time stamp of the received cell with the current time, and obtain a cell delay measurement. It will also check the sequence number to identify if any cells have been lost. The checksum is evaluated to ensure all received cells have no bit errors.

The analyser records mean and maximum cell delay times from each source, as well as lost cell numbers from each source and passes this information to the controlling transputer at the end of the simulation. The Packet-Level Interface transputer continuously passes information regarding reset rate and output buffer length to the traffic analyser through the INMOS link. This information is also recorded and passed to the transmit process so that it can be used in access control. At the end of a simulation period, the traffic generator process stops cell transmission and signals to the receive process that this has happened. The traffic analyser process continues to analyse the cells passed to it, but stops all averaging calculations. The controlling transputer sends requests to each T800 transputer in turn, for the stored data which it receives and files, using the host IBM PC filing system.

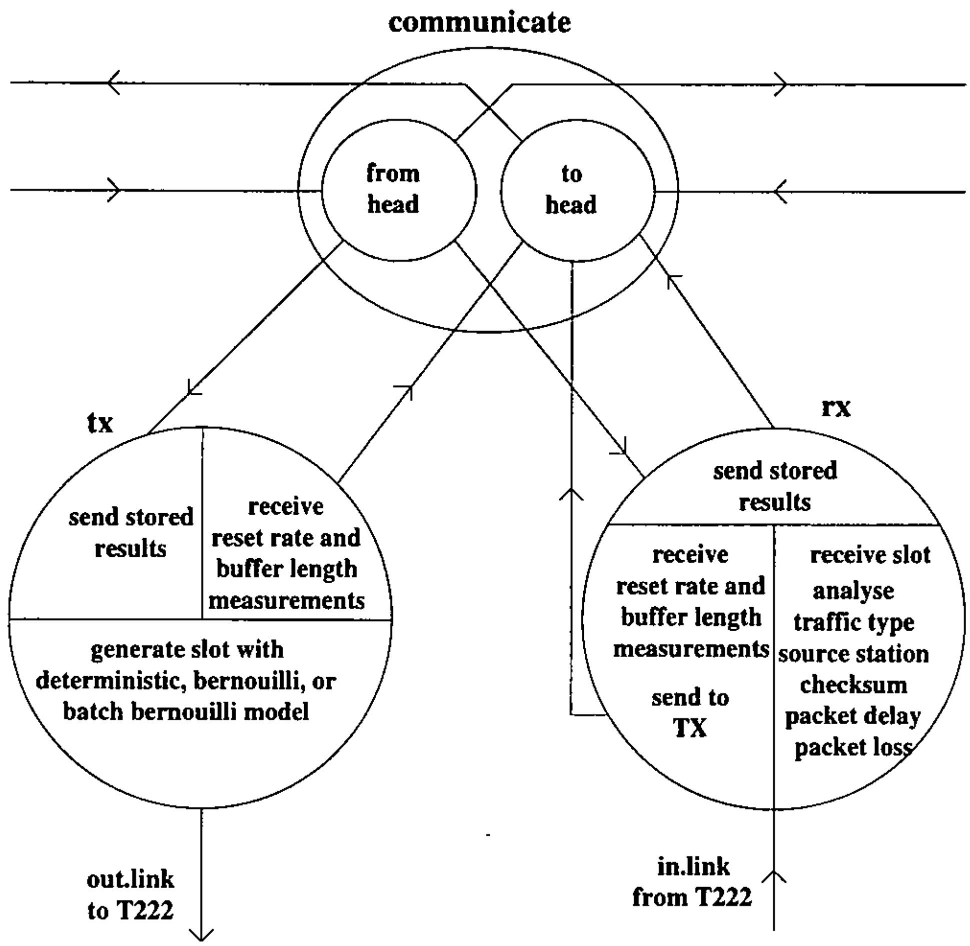

Figure 4.5 OCCAM processes for the Traffic Generator/Analyser

4.4.5 Control of the Simulation

Control of the testbed simulation is managed by the T800 transputer which is running under the control of TDS. A number of fixed length simulations can be carried out in each run of the controller program. In this batch of simulations, one or more parameters such as traffic intensity, or Di allocation, can be varied to give a set of related results. Because the testbed is a hardware implementation, with a throughput of up to about 10,000 cells per second, simulations involving the transmission of a million cells can be completed within two minutes, giving a high level of statistical consistency.

The results from each traffic generator/analyser are collected in arrays, and manipulated so that the cell delay and cell loss can be related to the station issuing those cells rather than the station receiving the cells. This facility is useful in analysing the effectiveness of access control and prioritisation of traffic classes to avoid cell delay and cell loss.

4.5 Summary of the Testbed Traffic Models

This chapter has shown how the testbed traffic generation process models the statistical characteristics of various ATM traffic sources. A model has been provided for constant bit rate services, multiplexed constant bit rate services, and bursty sources using respectively deterministic, Bernoulli and batch-Bernoulli processes. The level of burstiness in terms of peak cell rate to mean cell rate is variable, as is the mean cell arrival rate. These traffic sources are used in chapter 6 to characterise the testbed performance under differing traffic types, and in chapter 7 where two traffic classes are introduced to the testbed.

The analysis of the testbed under conditions of maximum loading is considered in chapter 5, and the results from this are compared to measurements taken on the testbed in chapter 6.

Chapter 5

Testbed Analysis

5.1 Introduction

In this chapter the Orwell protocol is analysed to derive a mathematical model for the maximum throughput, and maximum reset interval, of the slotted ring testbed under different traffic distributions. The predictions of these models are compared with another analysis of ORWELL given in [85], and will be compared with measured results from the testbed in chapter 6. The ORWELL reset rate mechanism is of particular interest since its use as an indicator of network loading has been proposed [78], [120], [121]. The analysis of the testbed at limiting values of ring throughput shows that the distribution of traffic on the network plays an important part in determining the maximum throughput and reset interval.

5.2 Stability Condition

The stability condition defines the maximum load that the ring can carry. The stability condition is dependent on traffic flow, so the cases of a single transmitting station, and of a symmetrically loaded ring configuration are considered here. In both cases cells generated at each station are randomly addressed to other stations on the network.

Some terms used in the analysis will first be defined.

R = Ring rate (bits/s), is the rate of transmission of data bits on the ring.

n = Number of stations on the ring.

S = Number of slots on the ring.

Ts = Slot duration time, the time between successive slots passing a point on the ring.

Tr = Ring Latency, is the time taken for a slot to circulate the ring.

RI = Reset Interval, is the time between resets.

5.2.1 Stability of the Ring with One Transmitting Station

Figure 5.1 illustrates the utilisation of slots by a single station transmitting to all other stations at its maximum bandwidth.

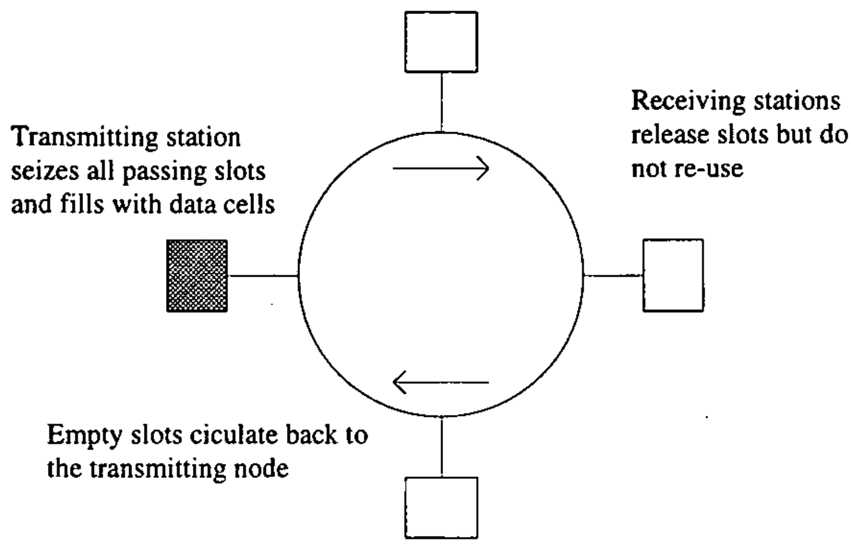

Figure 5.1. Single Transmitting Station

Each slot on the ring can be considered to be a server which services the queue of cells waiting at the transmitting station, one cell fitting exactly into one slot. Ignoring the effect of the reset mechanism, it can be seen that the interval between re-use of each slot i.e. the service time of each server (slot), is equal to *Tr*, the ring latency, as slots are only being used by the single transmitting station, and must completely circulate the ring before being re-used. In general the cell throughput, σ for the testbed is given by the number of servers divided by the service interval, τ :

$$\sigma = S / \tau \qquad (5.1)$$

The ring rate R can be expressed as a transmission rate of cells per unit time, R_{cell}, which, since each slot holds one cell, is calculated as:

$$R_{cell} = S / Tr \qquad (5.2)$$

Therefore for a single transmitting node, because $\tau = Tr$ it can be seen that,

$$\sigma = R_{cell} \qquad (5.3)$$

This result does not include the effect of the ORWELL reset mechanism which is considered in section 5.3.

5.2.2 Stability of Ring with Symmetric Loading

The configuration to be analysed is shown in figure 5.2.

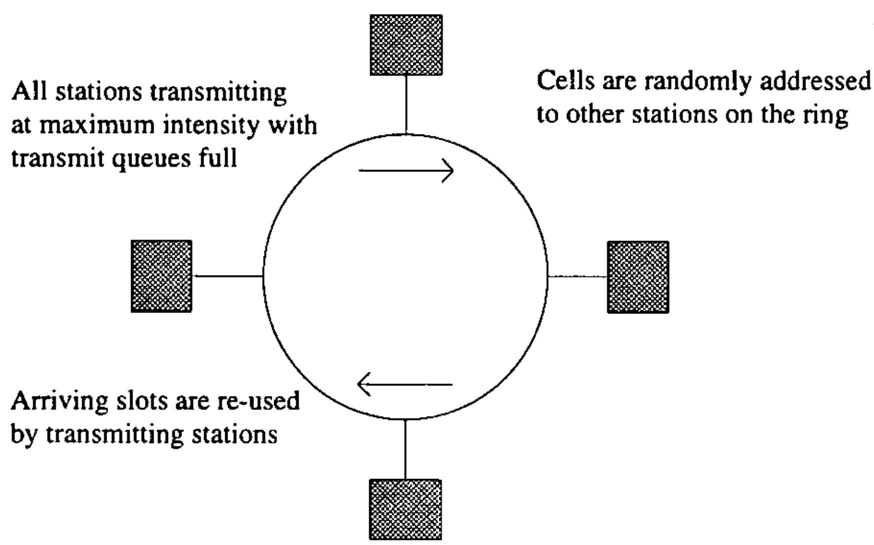

Figure 5.2 Stability condition under symmetric loading

In the symmetric case with random destination addressing of cells, the mean utilisation of each slot is the mean time taken to transport a cell from source to destination, since slots can be re-used by the destination station. The period of time that a slot is occupied in delivering a cell from source to destination is given by, τ , where on average:

$$\tau = Tr/2$$

hence using (5.1), and (5.2),

$$\sigma = 2.R_{cell} \tag{5.4}$$

If slots cannot be re-used by the receiving station and must instead be passed on empty to the next downstream station as in the full ORWELL protocol, the service time, τ is equal to the mean time taken to transport a cell from source to destination plus the time required for the slot to pass to the next station:

$$\tau = Tr/2 + Tr/n \tag{5.5}$$

Using equations 5.1 and 5.5,

$$\sigma = \frac{S}{Tr\left(\frac{1}{2}+\frac{1}{n}\right)} \tag{5.6}$$

$$\sigma = \frac{R_{cell}}{\left(\frac{1}{2}+\frac{1}{n}\right)} \tag{5.7}$$

The relative utilisation of the ring bandwidth is given by:

$$\frac{\sigma}{R_{cell}} = 2-\left(\frac{4}{n+2}\right) \tag{5.8}$$

Equation 5.8 is identical to that developed by Zafirovic-Vukotic and Niemeegers [85], using a multiple cyclic server model. Both results ignore the reset mechanism which will be considered in section 5.3.

In general the results above show that the condition for stability depends on the loading on the ring and will vary from R_{cell} to $2.R_{cell}$. This analysis assumes that stations are transmitting randomly and equally to all other stations. It should be noted that with particular loading patterns, the effective bandwidth of the ring could be greater than $2.R_{cell}$. In a traffic pattern where every node is transmitting to its nearest downstream neighbour, the effective bandwidth is $4.R_{cell}$. This unusual traffic flow makes good use of the transport capabilities of rotating slots, but there are other traffic flows which result in poor utilisation of the ring bandwidth, as shown in section 5.3.

The size of the ORWELL ring in terms of slots does not affect the results obtained in this section. Figure 5.3 shows two ORWELL rings, one of size 1000 bits and one of size 10,000 bits. The ring rate is 100Mbit/s in both cases. The throughput given by equation 4.1 is the same in both cases because the ratio *S/Tr* is the same. The ring rate is the factor which predominantly affects the maximum throughput of the slotted ring, rather than the number of slots or size of slots.

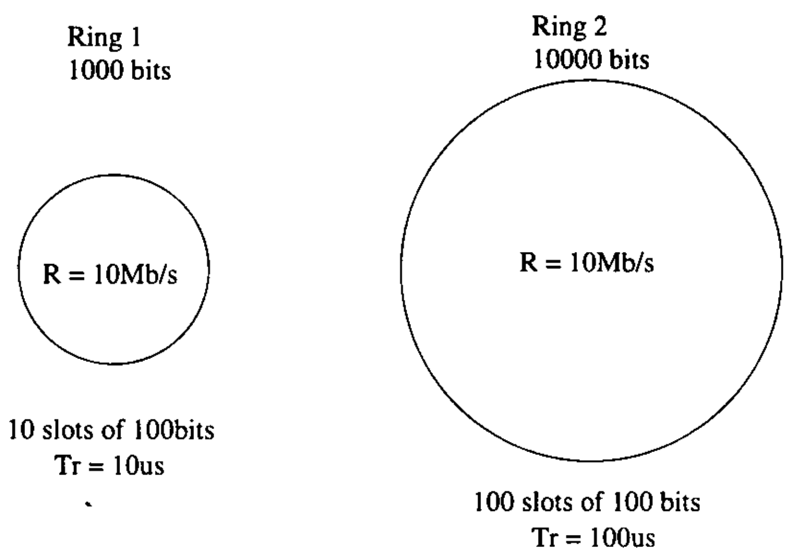

Figure 5.3 Comparing ORWELL Ring Sizes

In simulation studies [85] that are not possible with the fixed dimensioning of the testbed, it has been shown that the number of slots does slightly affect the performance of the ORWELL protocol, reducing throughput for low numbers of slots.

5.3 Analysis of the Reset Mechanism

The ORWELL protocol has been described in chapter 3. The reset mechanism ensures fair access to ring bandwidth, and can be used to sense the load on the ring, however it introduces an overhead in the transmission rate as TRIAL and RESET slots do not carry data. The following analysis determines the reset rate under extremes of loading.

5.3.1 Idle Ring

All stations when idle attempt to convert EMPTY slots to TRIAL slots and, if they receive their own TRIAL slot, will convert it to a RESET slot. When a station's Di-allocation is reset by receiving a RESET slot, its Outstanding Reset Timer (ORT) will be activated and it will not be allowed to convert its received TRIAL slots to RESET slots during this period. The minimum reset interval is thus the period of the ORT, which in the ORWELL protocol is specified as the Ring Latency *Tr*. In the ORWELL paper by Falconer and Adams [78] simulation results agree with the minimum reset interval being equal to the ring latency.

In the case of the testbed, when a station has been reset, it cannot be reset again within the ORT period which is *Tr*. The interval until this station is next reset depends upon which station on the ring originates the next reset, and how long the reset slot takes to arrive. On average if the next reset is equally likely to occur at any station, this time will be *Tr/2*, so the mean reset interval will be

$$RI_{(idle)} = Tr + Tr/2 \qquad (5.9)$$

5.3.2 Ring With One Station Transmitting at Maximum Intensity

When one station is transmitting at maximum intensity, equal to the ring rate *R*, all TRIAL slots reaching the transmitting station are filled with data cells until the station's Di-allocation is exhausted and it becomes PAUSED. Figure 5.4 shows that the TRIAL slot passing the station immediately after it becomes PAUSED must circulate to its originator, where it is converted to a RESET slot, and circulate back to the transmitting station to reset the Di-allocation. The following slot can then be used for cell transmission. The reset mechanism causes the transmitting node to pause for a period, T_{paused}

$$T_{paused} = Tr + Tr/S \tag{5.10}$$

The reset interval is the sum of the time spent transmitting T_{TX}, and the time spent paused, T_{paused}.

Time spent transmitting T_{TX} is the time required for *Di* slots to pass the transmitting station which is,

$$T_{TX} = Di.Tr/S \tag{5.11}$$

using equations 5.1, 5.10, and 5.11, the Reset Interval is,

$$RI = \frac{S + D_i + 1}{R_{cell}} \tag{5.12}$$

The throughput is the number of cells transmitted in the Reset Interval divided by the Reset Interval itself, hence using equations 5.10, 5.11, and 5.2, the maximum throughput is,

$$\sigma = \left(\frac{D_i R_{cell}}{S + D_i + 1} \right) \tag{5.13}$$

As *Di* tends to a large number, σ tends to R_{cell}, but for $Di = 1$, the reset overhead is very high, as only one slot in ($S + 2$) slots passing the transmitting station is carrying data.

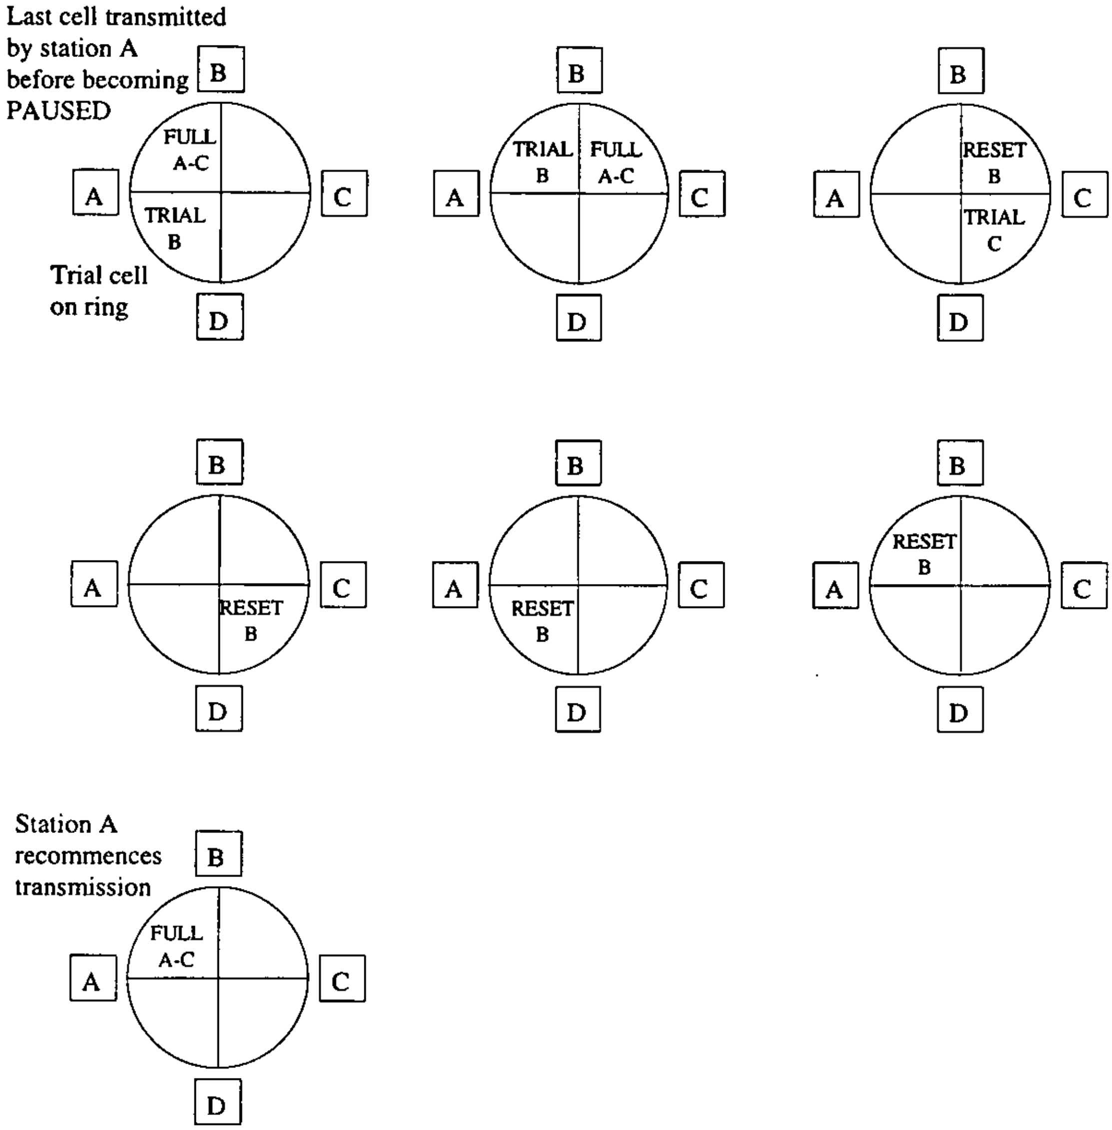

Figure 5.4 Reset mechanism with single transmitting station

5.3.3 Ring With All Stations Transmitting at Maximum Intensity

The reset mechanism is considered next when all nodes are transmitting at maximum rate, always having data ready to send. In this case, the assumption will be made that during the transmitting period, all of the slots circulating on the ring will be full all of the time, and that the nodes become PAUSED simultaneously. The reset process is then similar to that of an idle ring, requiring the full circulation of a TRIAL slot, and on average a further half circulation of the RESET slot before a station is reset. The mean service time for each cell carried is $Tr/2$, and the number of servers available is S. It is assumed that all nodes have the same Di-allocation.

The time required to transport $n.Di$ cells is:

$$T_{TX} = \left(\frac{n.Di.\left(\frac{Tr}{2}\right)}{S} \right) \tag{5.14}$$

The mean time each station is paused is

$$T_{paused} = Tr + Tr/2 \tag{5.15}$$

using equations 5.14, 5.15, and 5.2,

$$RI = \frac{nD_i + 3S}{2R_{cell}} \tag{5.16}$$

$$\sigma = 2R_{cell}\left(\frac{nD_i}{nD_i + 3S} \right) \tag{5.17}$$

Equation 5.17 illustrates that the ring throughput tends towards $2.R_{cell}$ when the product of n and Di increases, or as Di increases for a fixed number of stations.

5.3.4 Three Stations Transmitting to Fourth Station

Figure 5.5 shows the worst case traffic flow situation for the ORWELL ring. All stations are transmitting to one station which is idle. The level of traffic at each node is equal, except for the idle node, but the destination of the packets causes the ORWELL media access protocol particular problems.

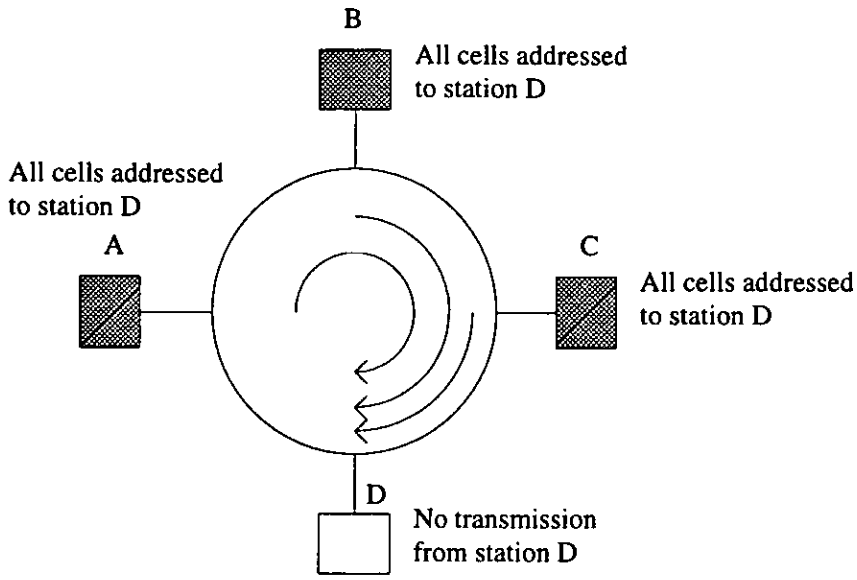

Figure 5.5 Asymmetric (worst case) loading of the testbed

The source of EMPTY slots in figure 5.5 is the single destination station (D). The slots will be used by station (A) until that station becomes PAUSED, the slots from (D) then pass through (A) to station (B) which will transmit its full Di-allocation, and so on. The sequence of cell transmission is shown in figure 5.6. The distribution of destinations makes the ORWELL ring behave like a polling system or token ring. The time required to transmit all of the Di-allocations is:

$$T_{TX} = (n-1).Di.Tr/S \qquad (5.18)$$

The time each station is paused before the next starts to transmit, plus the time for the first station to be reset and start transmitting again is:

$$T_{paused} = \left(\frac{Tr}{n}\right)((n-1)+2) \qquad (5.19)$$

using 5.18, 5.19, and 5.2

$$RI = \frac{D_i(n-1)+S\left(1+\frac{1}{n}\right)}{R_{cell}} \qquad (5.20)$$

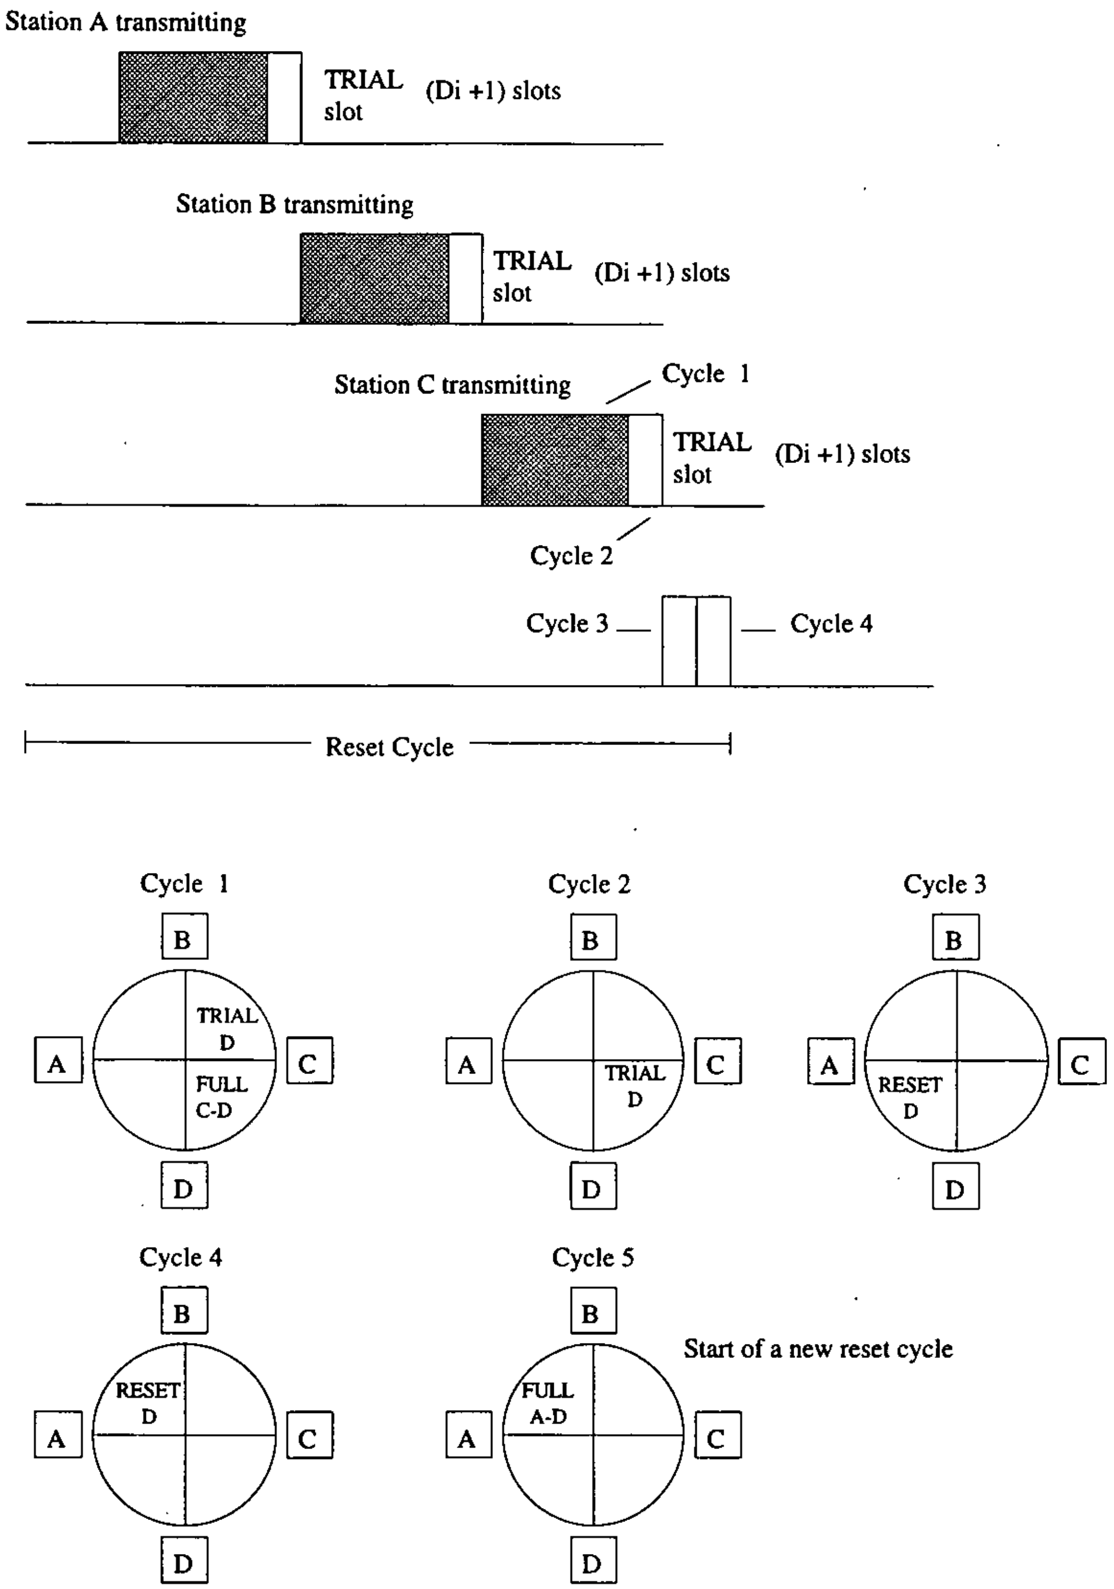

Figure 5.6 Reset cycle for the asymmetrically loaded testbed

and,

$$\sigma = \frac{R_{cell} D_i (n-1)}{D_i (n-1) + S\left(1 + \frac{1}{n}\right)} \tag{5.21}$$

The throughput tends to R_{cell} with increasing $n.Di$, which is similar to the throughput for a single station on the network (equation 5.13).

5.4 Summary

Table 5.1 summarises the results for section 5.3.

Table 5.1 Summary of Analytical Results for Section 5.3

Traffic Flow Description	**Reset Interval**	**Cell Throughput (σ)**
Single transmitting station to any destinations	$RI = \frac{S + D_i + 1}{R_{cell}}$	$\sigma = \frac{D_i . R_{cell}}{S + D_i + 1}$
All nodes transmitting at equal intensity	$RI = \frac{nD_i + 3S}{2R_{cell}}$	$\sigma = 2R_{cell}\left(\frac{nD_i}{nD_i + 3S}\right)$
All but one stations transmitting to other one	$RI = \frac{D_i(n-1) + S\left(1 + \frac{1}{n}\right)}{R_{cell}}$	$\sigma = \frac{R_{cell} D_i (n-1)}{D_i (n-1) + S\left(1 + \frac{1}{n}\right)}$

An analysis of the basic slotted ring mechanism demonstrates that for a single transmitting station the ring throughput tends to R_{cell}, the ring bandwidth expressed in cells per second, and for a symmetrically loaded ring with randomly addressed cells, the throughput tends to $2.R_{cell}$. By examining the behaviour of the ORWELL protocol and its reset mechanism in detail, the limiting values of throughput for three different traffic flows have been expressed in terms of R_{cell} - the rate of cell circulation, S - the number of slots on the ring, n - the number of stations on the ring and the Di value of each station.

An analysis of the reset mechanism has also provided a model for the reset rate at maximum throughput for each of the models. The reset rate is an indicator of the traffic intensity on the ring and is shown, by the results in table 5.1, to be dependent on the traffic flow existing on the ring.

The results obtained in this chapter will be tested against measured results from the testbed as a means of verifying the testbed performance in certain limited conditions. The testbed will then be used to provide more detailed information on the behaviour of the protocol under conditions where analysis is difficult or impossible.

Chapter 6

Results of the Testbed Simulations

6.1 Introduction

In this chapter simulations are performed on the slotted ring testbed and are compared to the analytical results obtained in chapter 5. The traffic models described in chapter 4 are used together with the testbed to study the effect of traffic arrival statistics on the performance of the slotted ring. The ORWELL reset rate mechanism is of particular interest since its use as an indicator of network loading has been proposed [78], [120], [121]. Measurements of the reset rate are therefore made for various traffic arrival distributions, traffic intensities, and traffic addressing distributions. Other performance metrics such as cell loss and cell delay times are measured during the simulations, and are related to the output buffer occupancy.

6.2 Simulation Results

To evaluate the testbed performance under different traffic distributions, and to verify the analytical results obtained in section 5.3, the testbed was subjected to the various traffic flows considered in that section. Homogenous traffic was used in this characterisation, and it is not connection-based, so traffic intensity is described in cells / second rather than numbers of connections. A Bernoulli traffic arrival process was used.

6.2.1 Ring with One Station Transmitting at Maximum Intensity

Simulation results were obtained for each analysis undertaken in section 5.3. The first case simulated was that of a single transmitting station. Figure 6.1 shows the measured maximum throughput and the calculated maximum throughput using equation 5.13 at each *Di* value, and figure 6.2 shows the measured and calculated maximum reset interval using equation 5.12. The results are very close and several data points overlay each other.

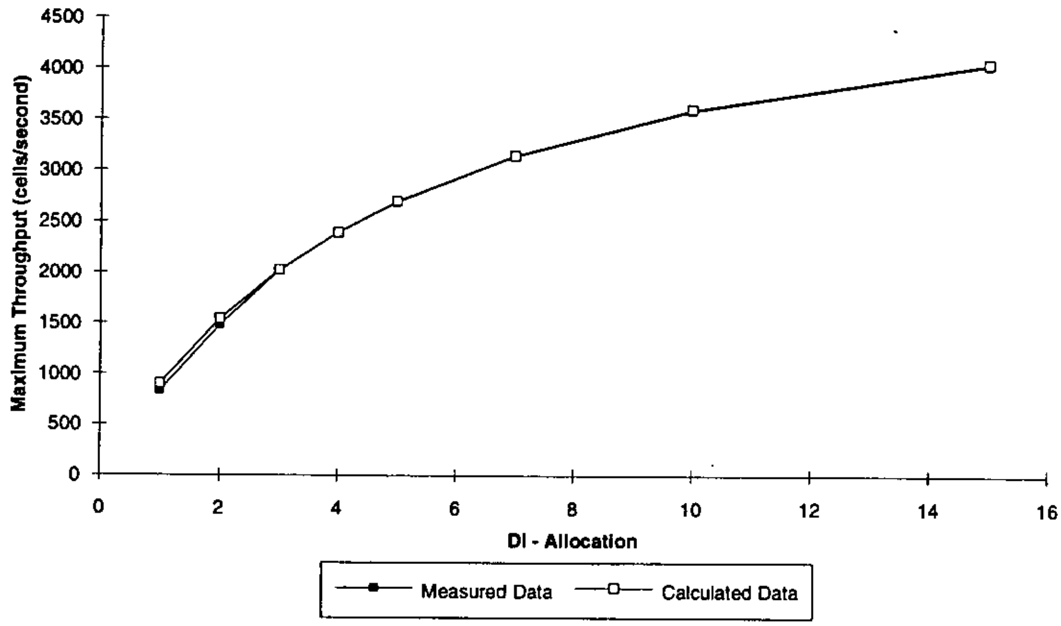

Figure 6.1 Maximum Throughput : *Di* Allocation for a single node

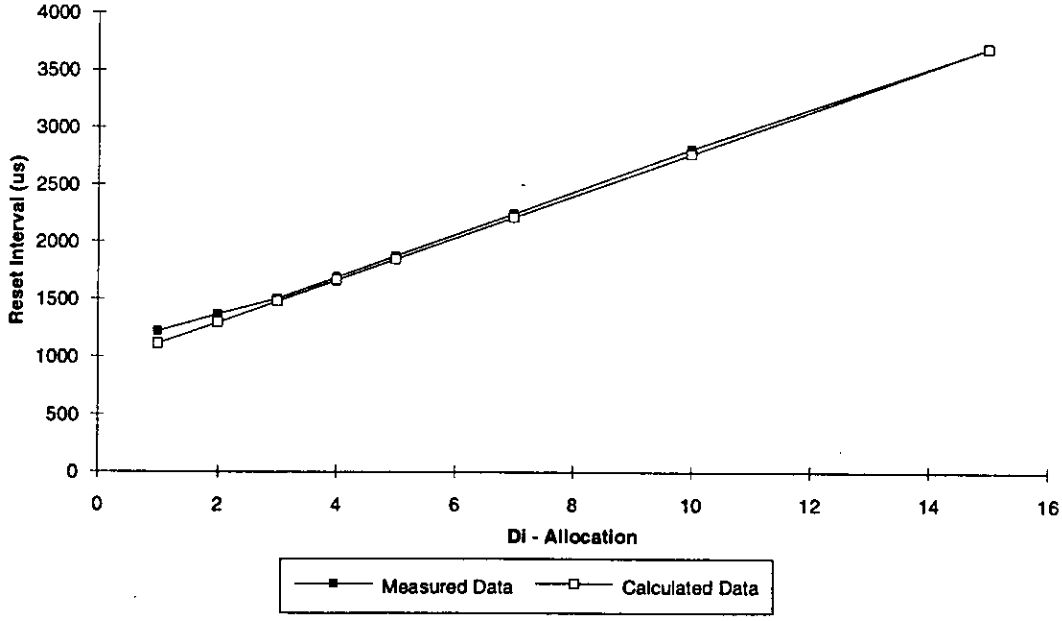

Figure 6.2 Maximum Reset Interval : Di allocation for a single station

Calculated and measured results show close agreement for this traffic configuration. The results diverge at the lowest Di allocations, where the station is paused for a greater percentage of time than it is transmitting data. From taking measurements on ten occasions, the variation around the measured maximum throughput or Reset Interval is within 1%.

6.2.2 Ring With All Stations Transmitting at Maximum Intensity

The analysis of the symmetrically loaded testbed is contained in equations 5.16 and 5.17. The plot of maximum throughput against Di allocation is shown in figure 6.3, and maximum reset interval against Di allocation in figure 6.4.

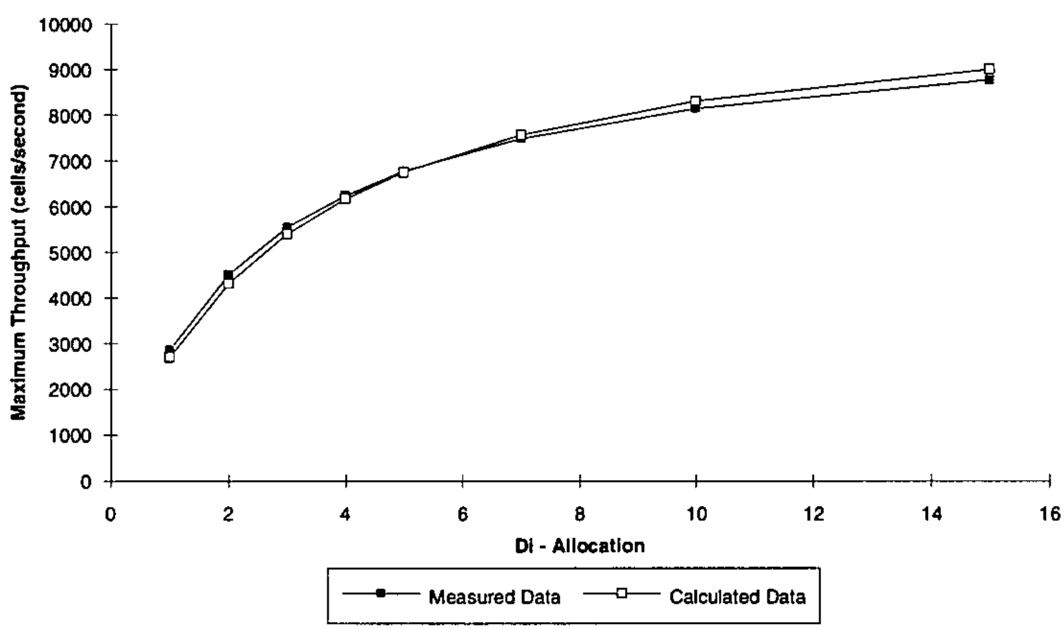

Figure 6.3 Maximum Throughput against Di Allocation for the Symmetrically loaded testbed

In figures 6.3 and 6.4, the calculated results and measured results show a discrepancy at low Di-allocations as they did with the single transmitting station in figures 6.1 and 6.2. At high Di-allocations there is a larger discrepancy in the case of the symmetrically loaded ring because of the assumption that all stations will become paused at the same instant. Equation 5.14 uses this assumption to determine the time that a station is transmitting in any reset cycle. As the Di-allocation increases, the probability that all stations will cease transmission simultaneously decreases, leading to a divergence of the calculated and measured values of about 5% at Di=15.

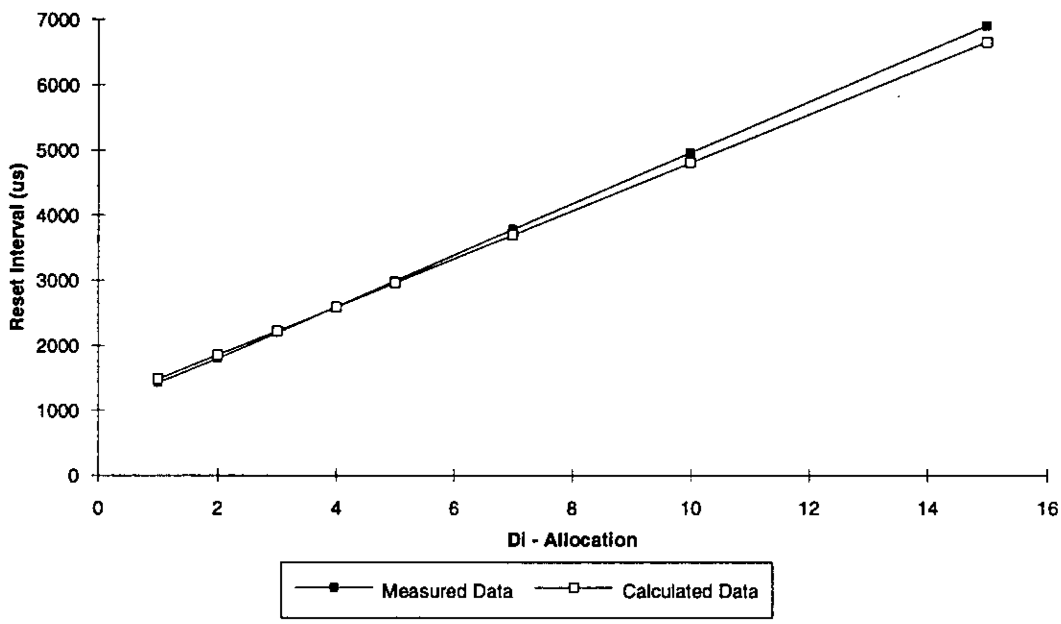

Figure 6.4 Maximum Reset Interval against Di-allocation for the symmetrically loaded testbed.

6.2.3 Three Stations Transmitting to Fourth Station

The measurements conducted in this case, are for the testbed traffic flow as described in figure 5.5. The analysis of this situation is given in equations 5.11 and 5.12. Figure 6.5 shows calculated and measured results for the maximum throughput against Di-allocation, and figure 6.6 shows maximum reset interval against throughput.

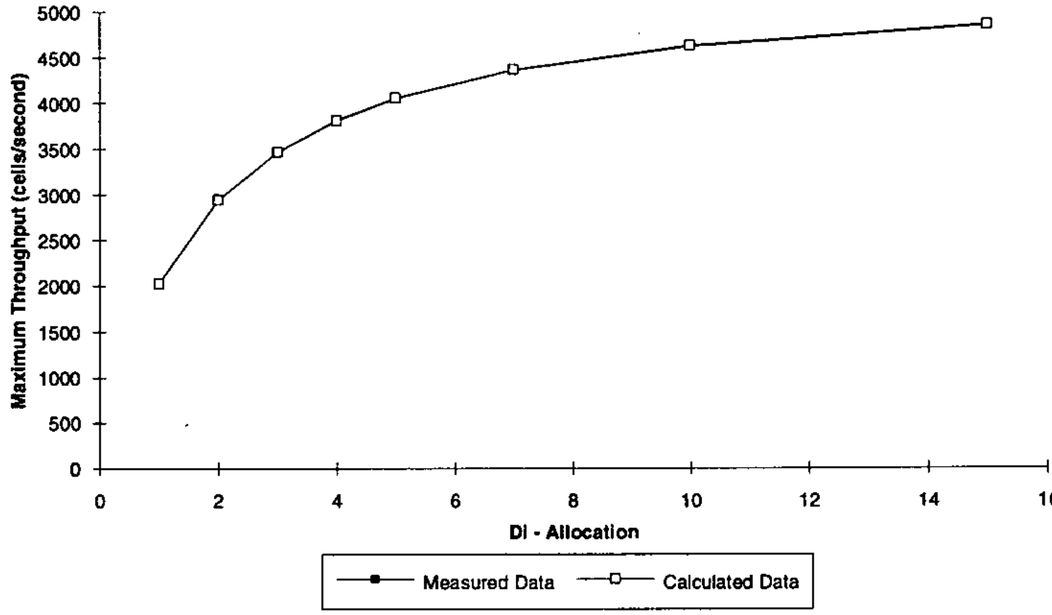

Figure 6.5 Maximum Throughput against Di allocation for asymmetrically loaded testbed

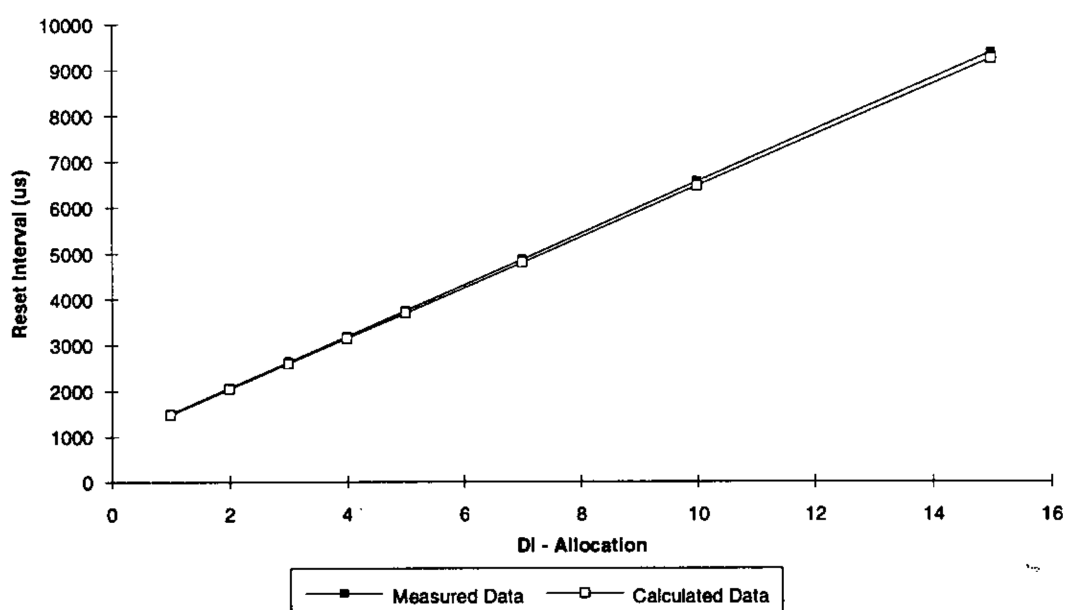

Figure 6.6 Maximum Reset Interval against Di allocation for asymmetrically loaded testbed

The measured and calculated results in the case of the asymmetrically loaded testbed show close agreement as the transmit mechanism is very well defined. The reason this traffic flow configuration is shown is that it represents the worst case flow for an ORWELL based protocol to control media access. Each transmitting station downstream of the first station is successively starved of service by the others leading to lower throughput and a greater reset interval than in the symmetrically balanced case. The behaviour of the ORWELL protocol under these conditions is similar to that of the token ring protocol, where each station has to wait for its upstream neighbour to finish transmitting before it can commence, hence only one station transmits at any one time.

6.2.4 ORWELL Reset Interval for 3 Traffic Processes

Sections 6.2.1 to 6.2.3 have compared analytical and measured results for the maximum throughput, and maximum reset interval of the testbed under 3 different traffic distributions. In order to take these measurements, the output buffers of the transmitting stations were maintained full resulting in significant cell loss due to buffer overflow. In normal operation, the slotted ring would be required to operate with a bounded cell loss rate so that a certain quality of service could be maintained to the traffic streams, not allowing the output buffers to fill up and cells to be lost. Figure 6.7 shows the variation of reset interval with ring throughput for each of the traffic distributions considered in the preceding sections for a fixed Di-allocation of 10.

It can be seen that the reset interval does increase with cell throughput for each of the 3 distributions, but the value of the reset interval for a given throughput depends on the actual traffic distribution, when the throughput is greater than the maximum throughput for a single station. Unless throughput is limited to the throughput for a single station, the reset interval or its inverse, the reset rate, does not appear to be a good indicator of cell throughput when traffic distribution is allowed to vary. It has been suggested, [78], [121], that the ORWELL reset rate can be used as an indicator of cell throughput for admission control purposes, however this work is based on symmetric loading of transmitting stations and random addressing of cells, and does not consider asymmetrical traffic distributions. The analysis and simulations in this section show that under changing traffic distribution, the reset rate could give a misleading indication of cell throughput.

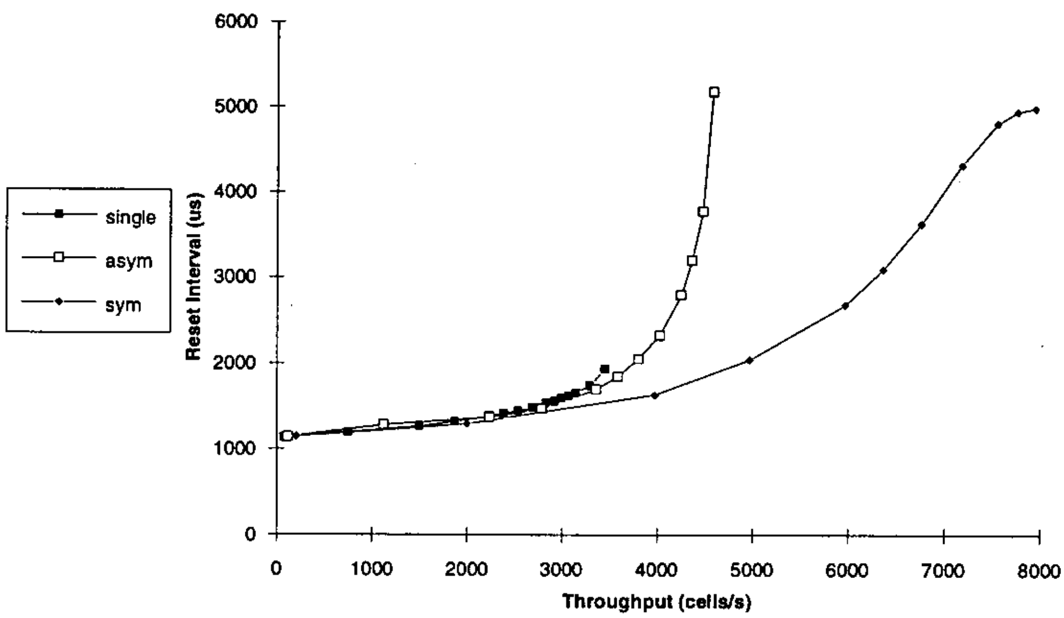

Figure 6.7 Reset Interval : Throughput for 3 Traffic Distributions with Di=10

6.3 Performance of the ORWELL Testbed With Different Traffic Models

Whilst it is useful and important to consider the maximum throughput of any given configuration of traffic on the testbed, the network would not be operated at these levels of loading because of the rate of cell loss and delay which would be incurred. To enable an operating region of the testbed to be defined, the regions of loading where output buffer overflow does not occur must be identified and the cell delay time must be characterised. In this section, the hardware testbed is used to show that the cell arrival pattern and traffic flow distribution have an effect on the reset rate, the cell delay time, and the cell loss rate.

6.3.1 Cell Throughput and Reset Interval

Having considered how different traffic distributions affect the potential of the testbed to carry traffic, the effect on the testbed performance of the traffic models discussed in chapter 4 was investigated. In these experiments, all stations were symmetrically loaded and transmitting randomly to the other stations on the testbed. The throughput and reset interval of the testbed were measured for a variety of Di-allocations under three of the traffic models considered in chapter 4. The first model is a deterministic arrival process where a single cell arrives at a regular, defined interval, which determines the traffic intensity. The second model is a Bernoulli arrival process with a 0.5 probability that a cell will arrive in any given interval, and the third model is the arrival of a fixed size batch of 4 cells arriving with a probability of 0.125 in each Bernoulli trial interval (batch Bernoulli arrival). The mean cell arrival rate was varied by altering the time interval between deterministic arrivals or Bernoulli trials. The graphs of throughput against offered traffic are shown for varying Di-allocations in figures 6.8, 6.9 and 6.10 and reset interval against throughput are shown for varying Di-allocation in figures 6.11, 6.12, and 6.13.

The graphs of throughput against offered traffic show that in each traffic arrival process the reset overhead limits the maximum throughput particularly at low Di-allocations. It can be seen that the same maximum throughput for a given Di-allocation is achieved for each traffic model when the output buffers always have data cells to transmit. This happens only when the cell arrival rate is greater than the stability condition and causes a high cell loss rate due to output buffer overflow. In the region where the output buffers do not overflow and there is no cell loss, throughput is greatest for a deterministic cell arrival process, since with this process the probability of the output buffer being empty at a time when an empty slot is available is the smallest.

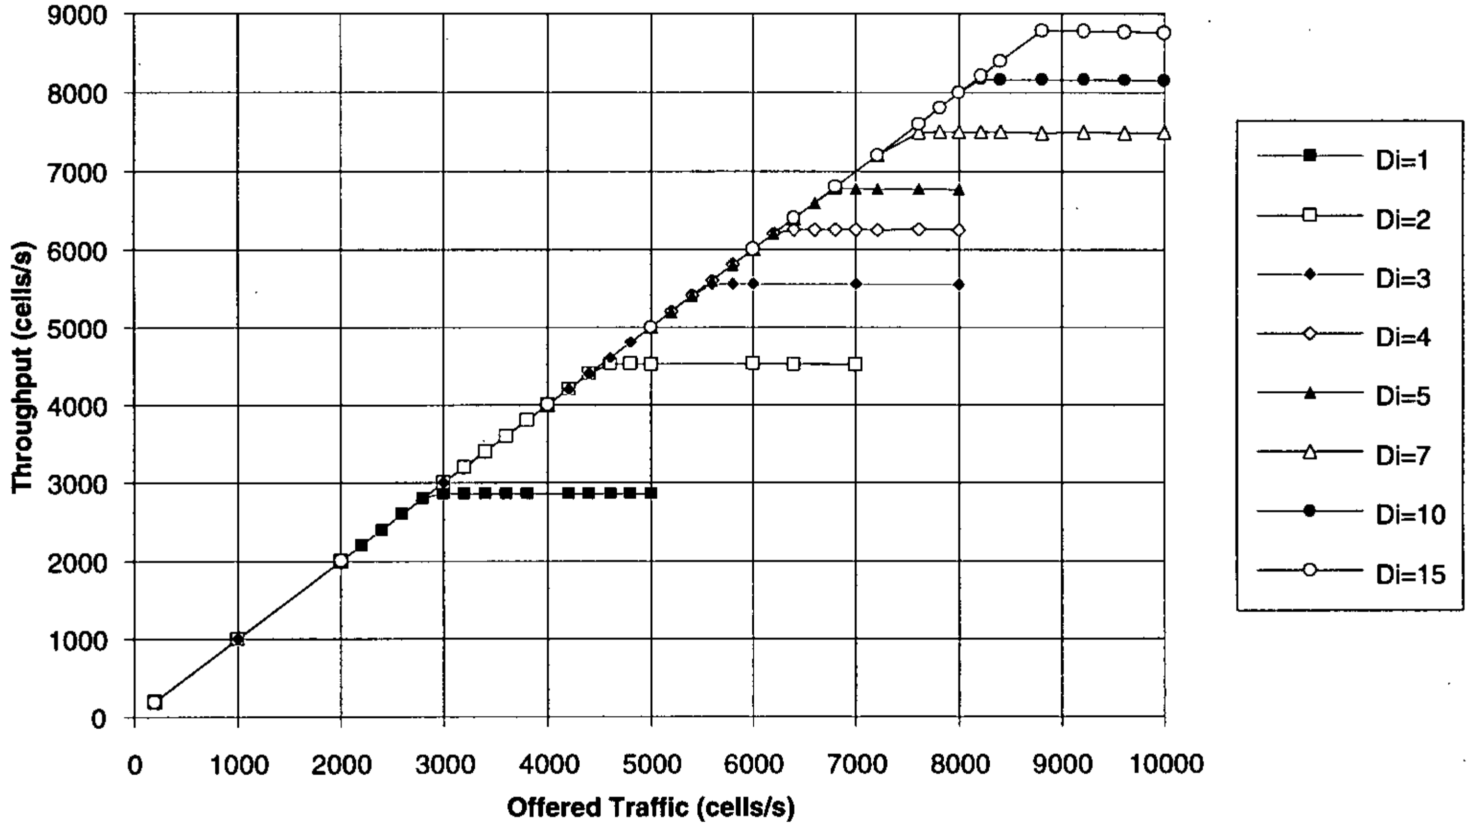

Figure 6.8 Throughput : Offered Traffic for Deterministic Cell Arrival

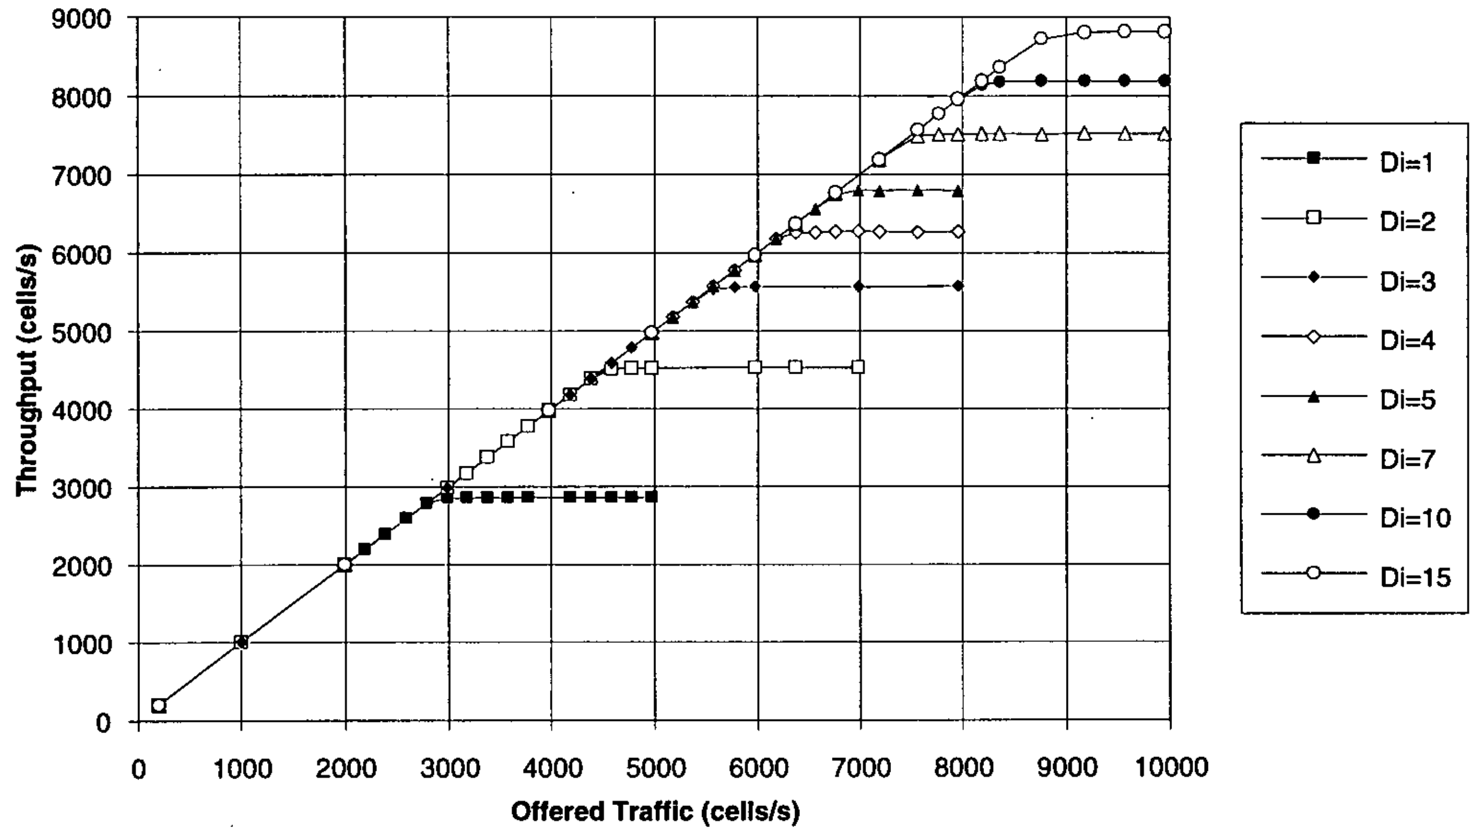

Figure 6.9 Throughput : Offered Traffic for Bernouilli Cell Arrival

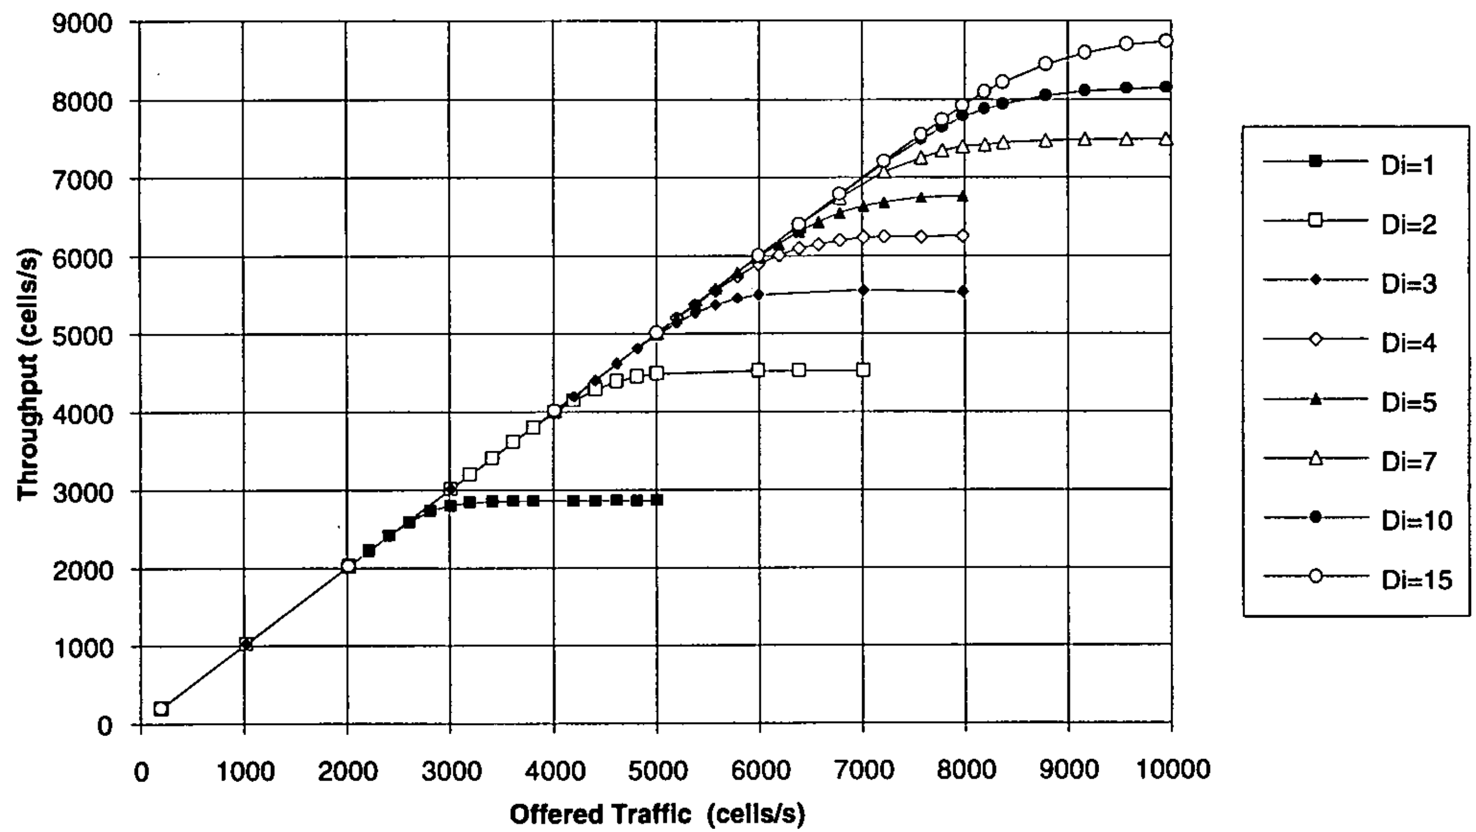

Figure 6.10 Throughput : Offered Traffic for Batch Bernouilli Cell Arrival

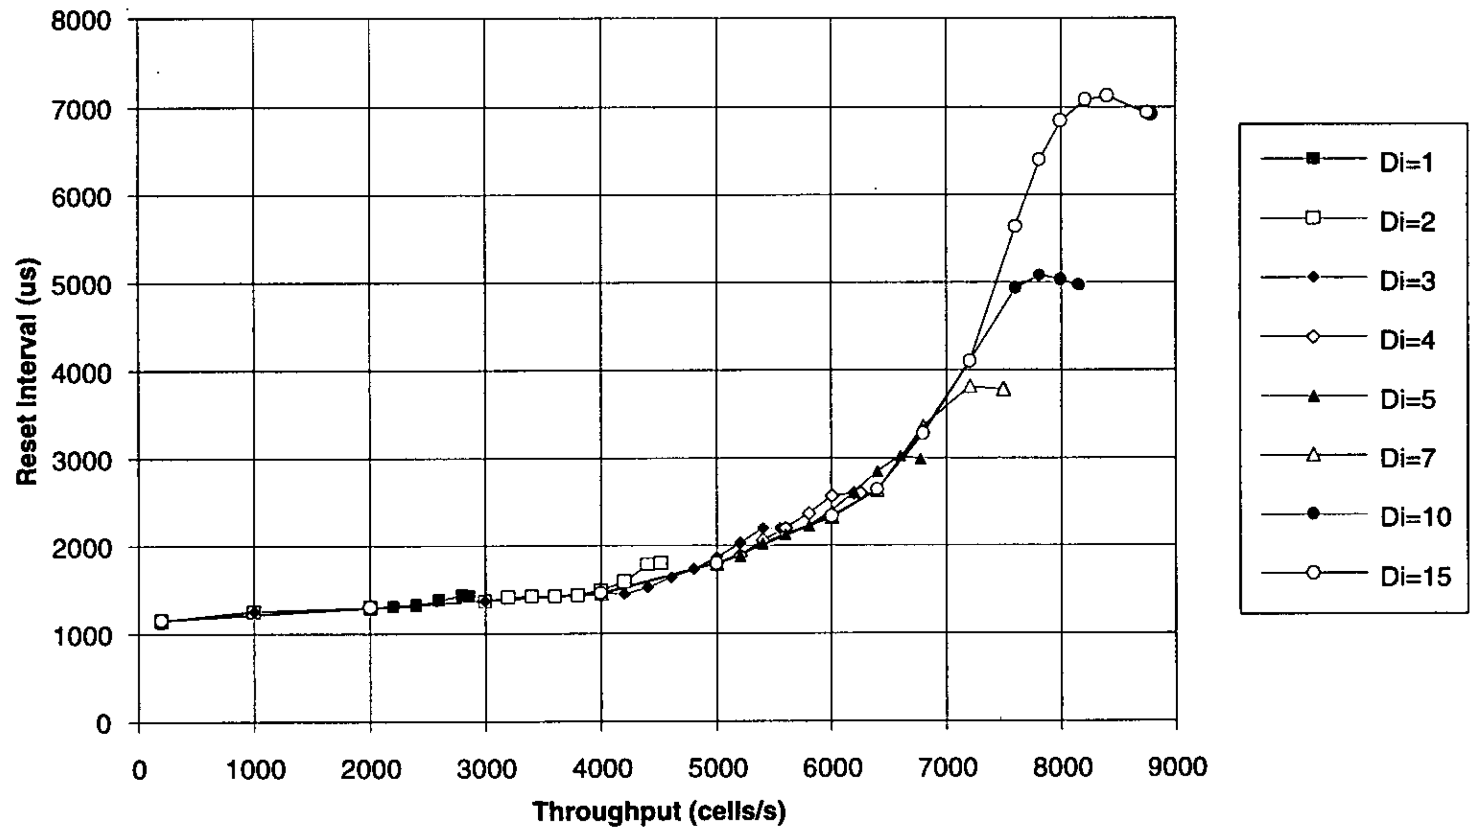

Figure 6.11 Reset Interval : Throughput for Deterministic Cell Arrival

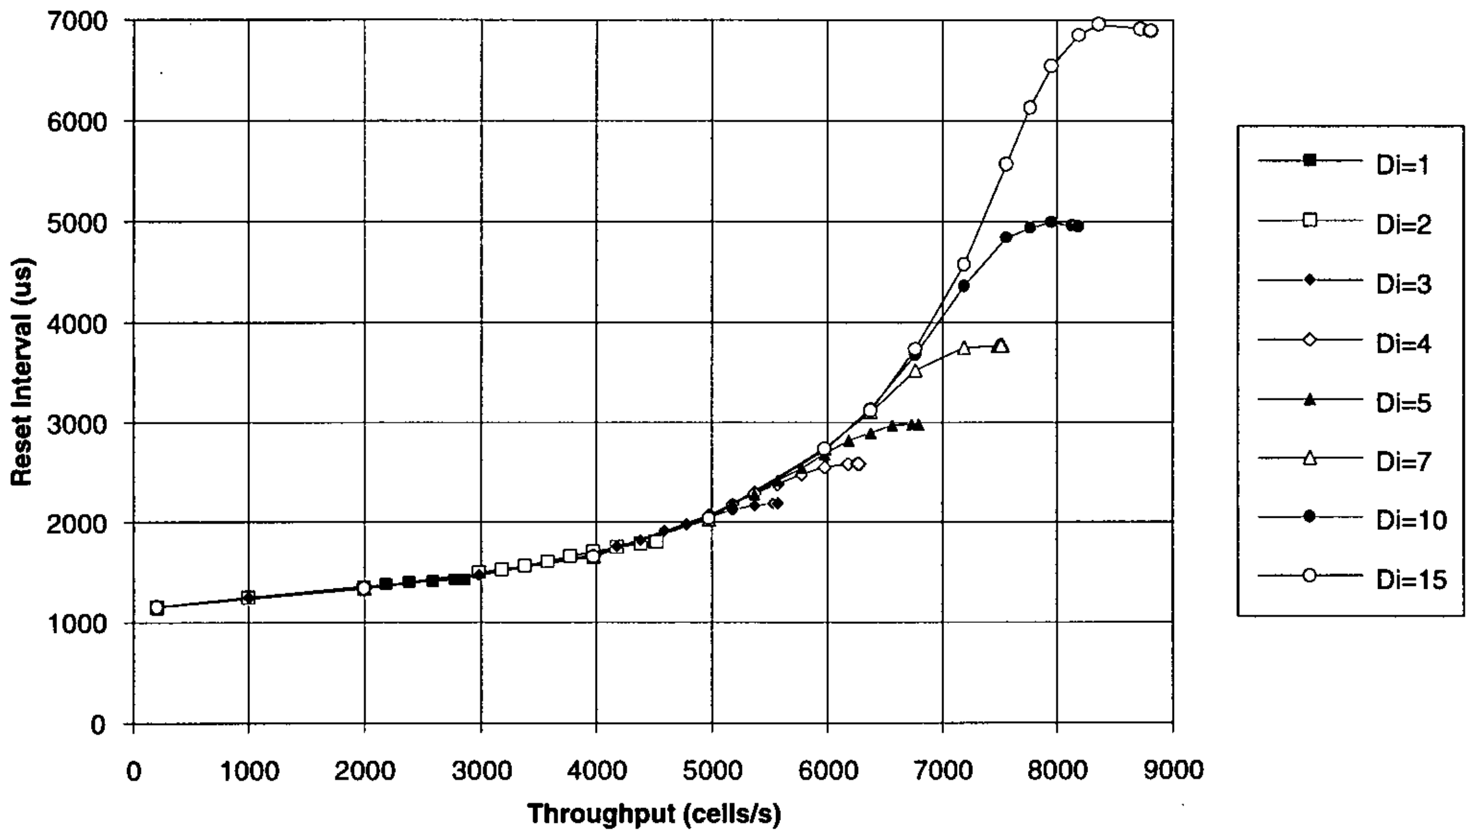

Figure 6.12 Reset Interval : Throughput for Bernouilli Cell Arrival

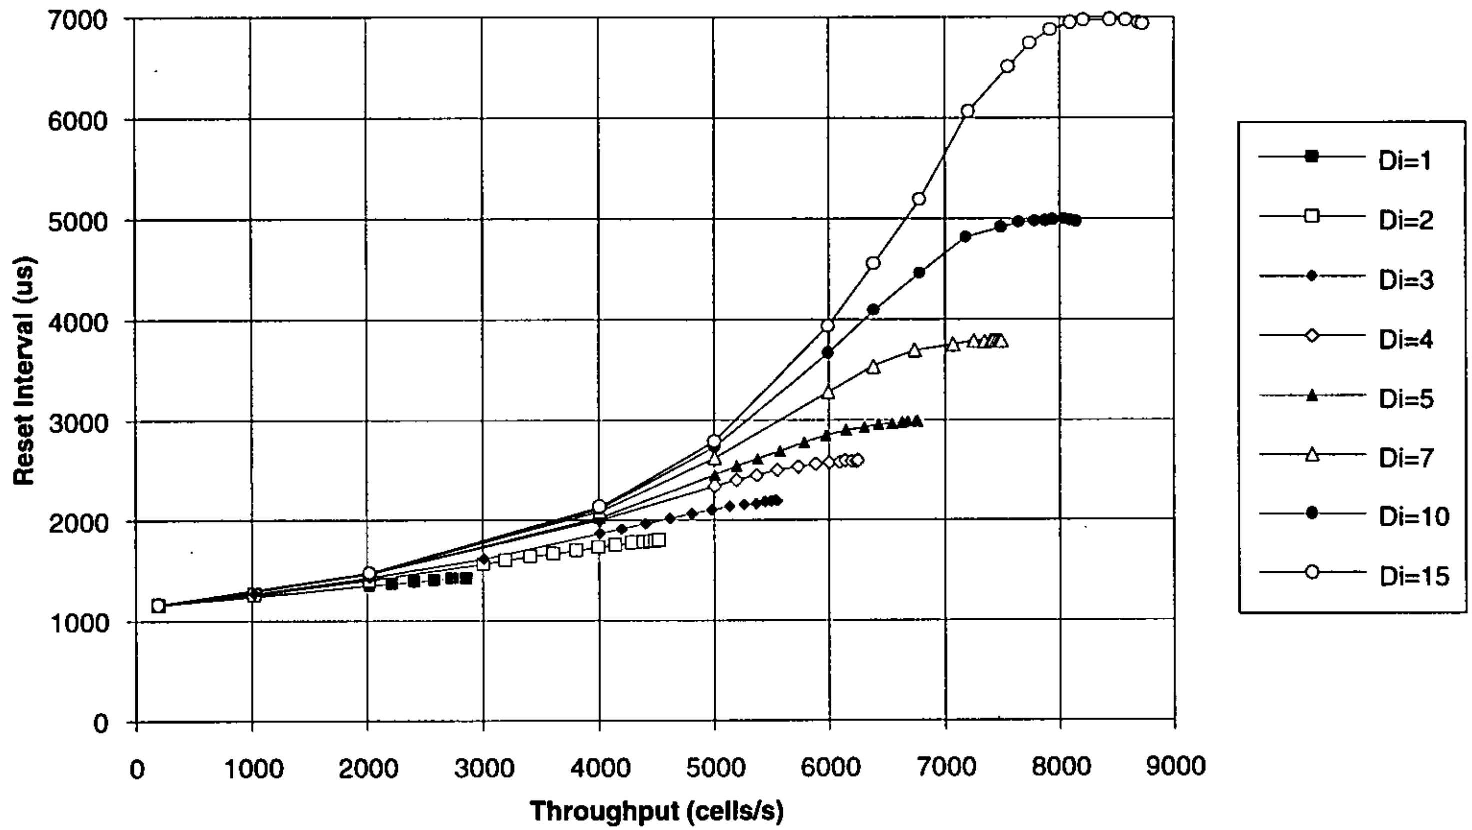

Figure 6.13 Reset Interval : Throughput for Batch Bernouilli Cell Arrival

The deterministic process of figure 6.11 shows that the same reset interval path is followed for all Di-allocations, but the maximum throughput and reset interval are limited by the overhead of the reset mechanism. As the number of cells being transmitted in each reset interval rises towards the maximum throughput, the reset interval increases slightly above the value obtained when there are always cells to transmit. This occurs because of a synchronisation effect between the time that successive stations generate cells, and the time that a slot takes to travel from one station to another. In these experiments, the stations do not generate slots at the same instant, they are offset in time, to minimise the effects described.

The Bernoulli and batch Bernoulli arrival processes shown in figures 6.12 and 6.13, also show a gradual change in the reset interval with increasing traffic load. The batch Bernoulli arrival process gives generally higher reset intervals for a given load because of blocks of cells being transmitted from a station, making the downstream stations wait for service. These graphs show that the overhead of the reset mechanism is very high at low Di-allocations, and so for the rest of this section of work all stations have been given the maximum Di-allocation of 15.

In figure 6.14, the plot of reset interval against throughput for each of the traffic models is shown, when $Di = 15$. In the legend, *det* represents deterministic traffic, *ber* represents Bernoulli traffic, and *bl4* represents batch Bernoulli traffic with a batch size of 4. At a throughput of 2000 cells/s, the reset interval varies by ±8% of the midpoint between the traffic arrival processes, at 4000 cells/s the variation is ±20%, and at 6000 cells/s the reset interval for batch Bernoulli arrival is almost twice the reset interval for deterministic arrival. At traffic intensities up to 6000 cells/s there is no loss of cells with any of the cell arrival processes. This graph demonstrates that the reset interval is affected by the traffic arrival statistics at each network station, even when the traffic distribution is kept constant.

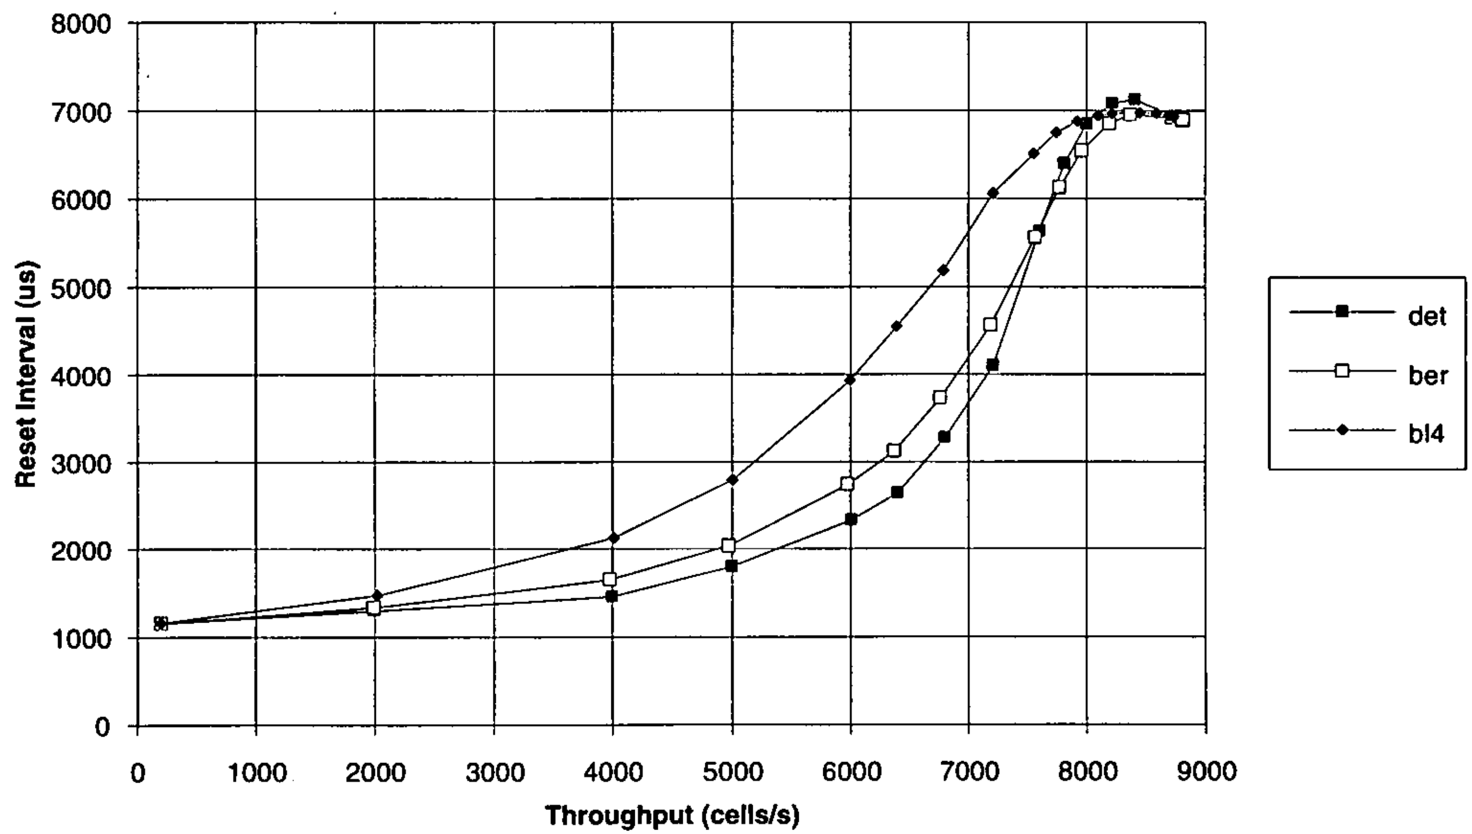

Figure 6.14 Reset Interval : Throughput for 3 Cell Arrival Processes (Di=15)

6.3.2 Cell Loss and Delay

For a symmetrical traffic flow distribution with random cell addressing, and a Di-allocation of 15 to each station, cell delay and output buffer length were plotted against cell throughput. Figure 6.15 shows mean cell delay, and maximum cell delay is shown in figure 6.16. The mean output buffer length is shown in figure 6.17, and maximum output buffer length in figure 6.18. Cell loss occurs because of buffer overflow, and figure 6.19 shows cell loss for the three traffic arrival processes.

Cell delay in fact consists of the following components: transmit overhead, output buffer queuing, propagation delay, input buffer queuing, and receive overhead. In this testbed, the overheads for transmit and receive are assumed constant, the input queuing is assumed to be a fixed value, and the mean propagation delay for randomly addressed cells is $Tr/2$. The mean cell delay at very low levels of transmitted cells is 1100us (+/- 10%), which can be considered to encompass fixed overheads and the mean propagation delay.

Because cell loss occurs due to buffer overflow, the dimensioning of the buffer will affect at what throughput cell loss occurs. Larger buffers will reduce cell loss for a given throughput, but increase maximum possible cell delay.

A comparison of the figure 6.15 with figure 6.16, and 6.17 with 6.18, shows that the cell delay on the testbed is largely dependent on the queuing delay. The maximum queuing delay can be estimated as the maximum output buffer occupancy in cells divided by the maximum throughput in cells/s per station. That is:

$$Delay_{max} = 50/2200$$

This calculation gives 22.7ms maximum delay, compared to the measured maximum delay of about 25ms. The estimated maximum delay is less than the measured value because it is based on the mean service time of the queue, equal to the maximum

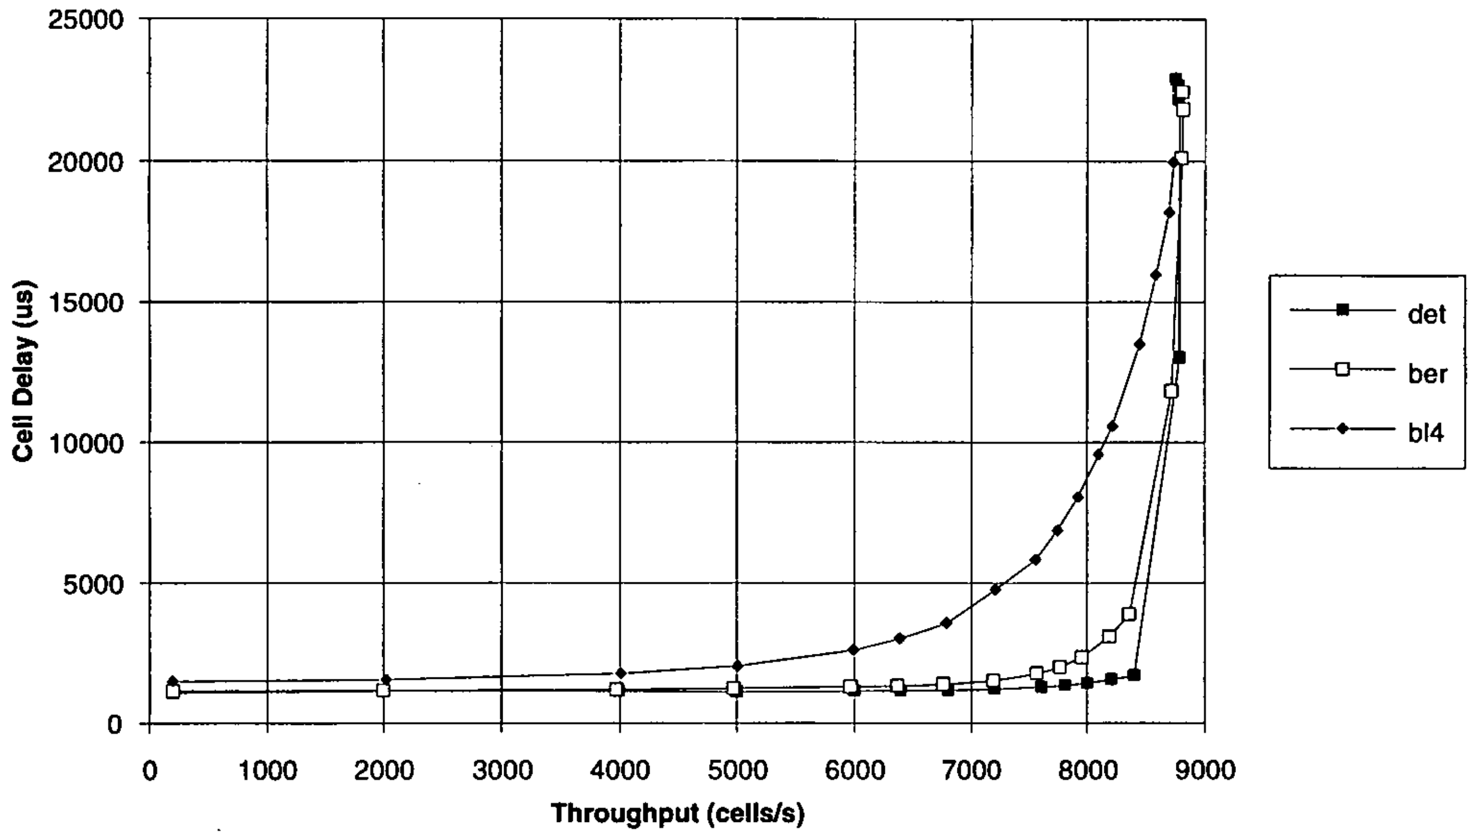

Figure 6.15 Mean Cell Delay : Throughput for 3 Cell Arrival Processes (Di=15)

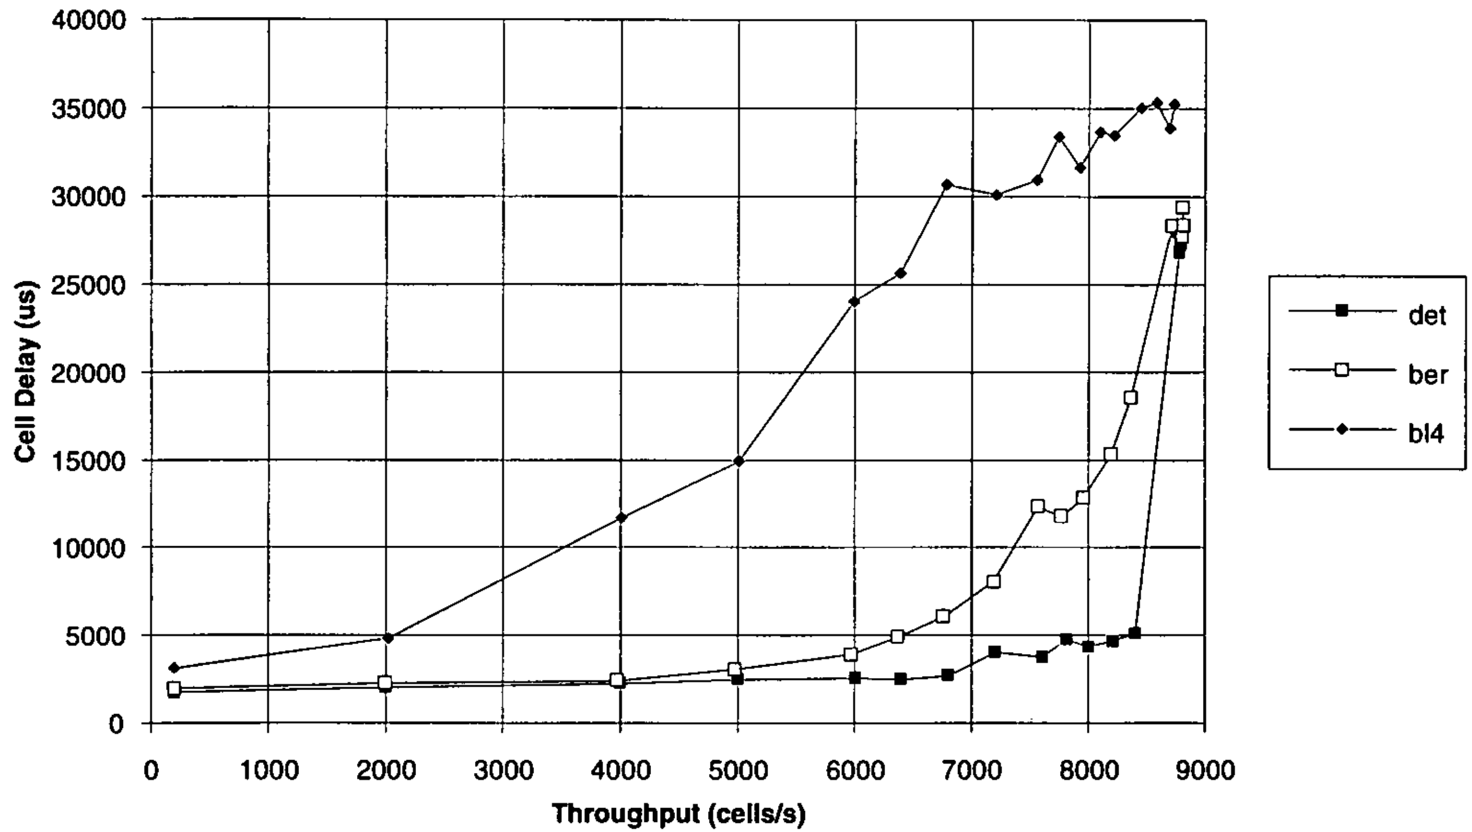

Figure 6.16 Maximum Cell Delay : Throughput for 3 Cell Arrival Processes

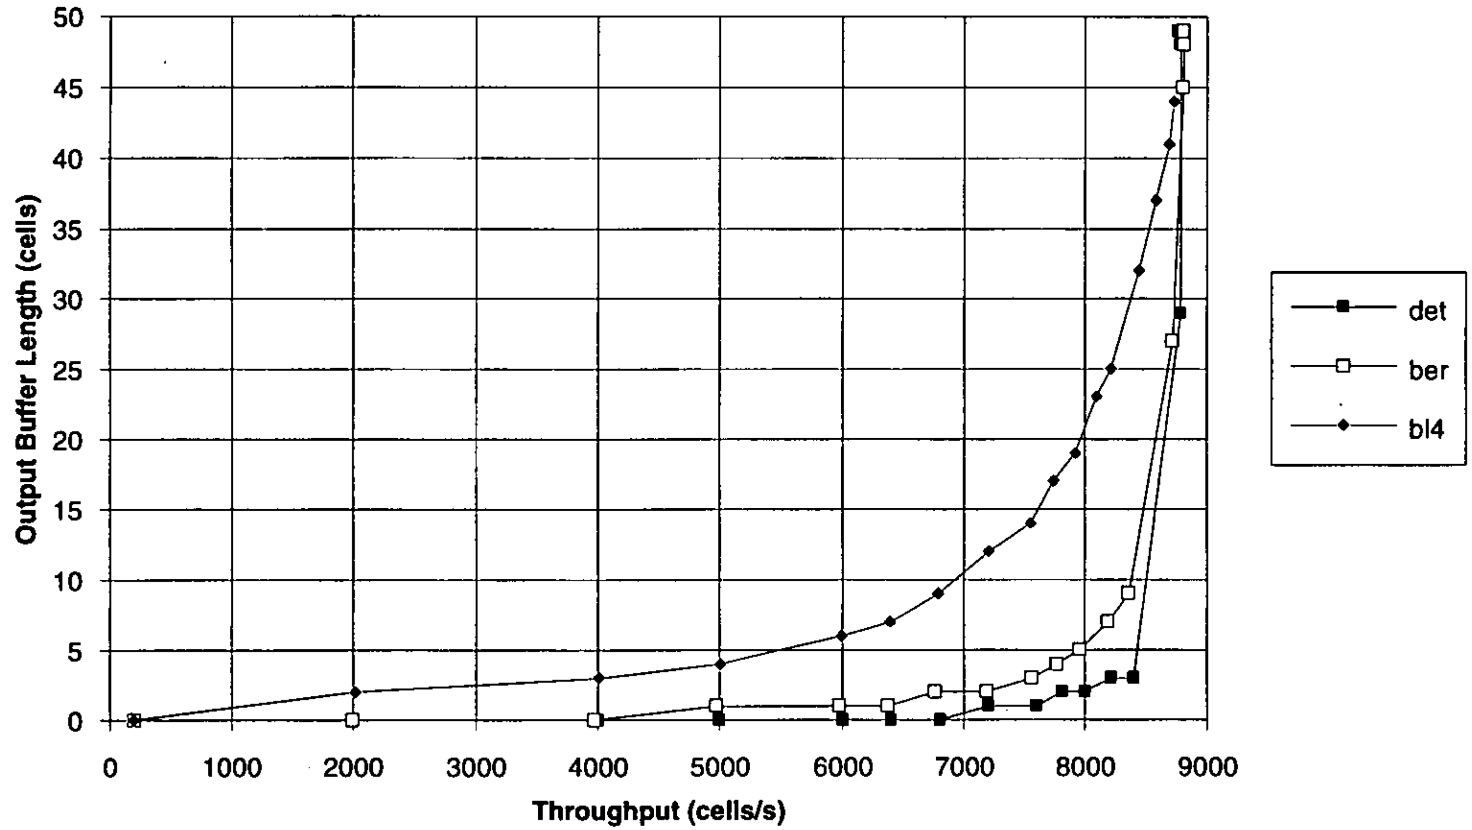

Figure 6.17 Mean Output Buffer Occupancy : Throughput for 3 Cell Arrival Processes

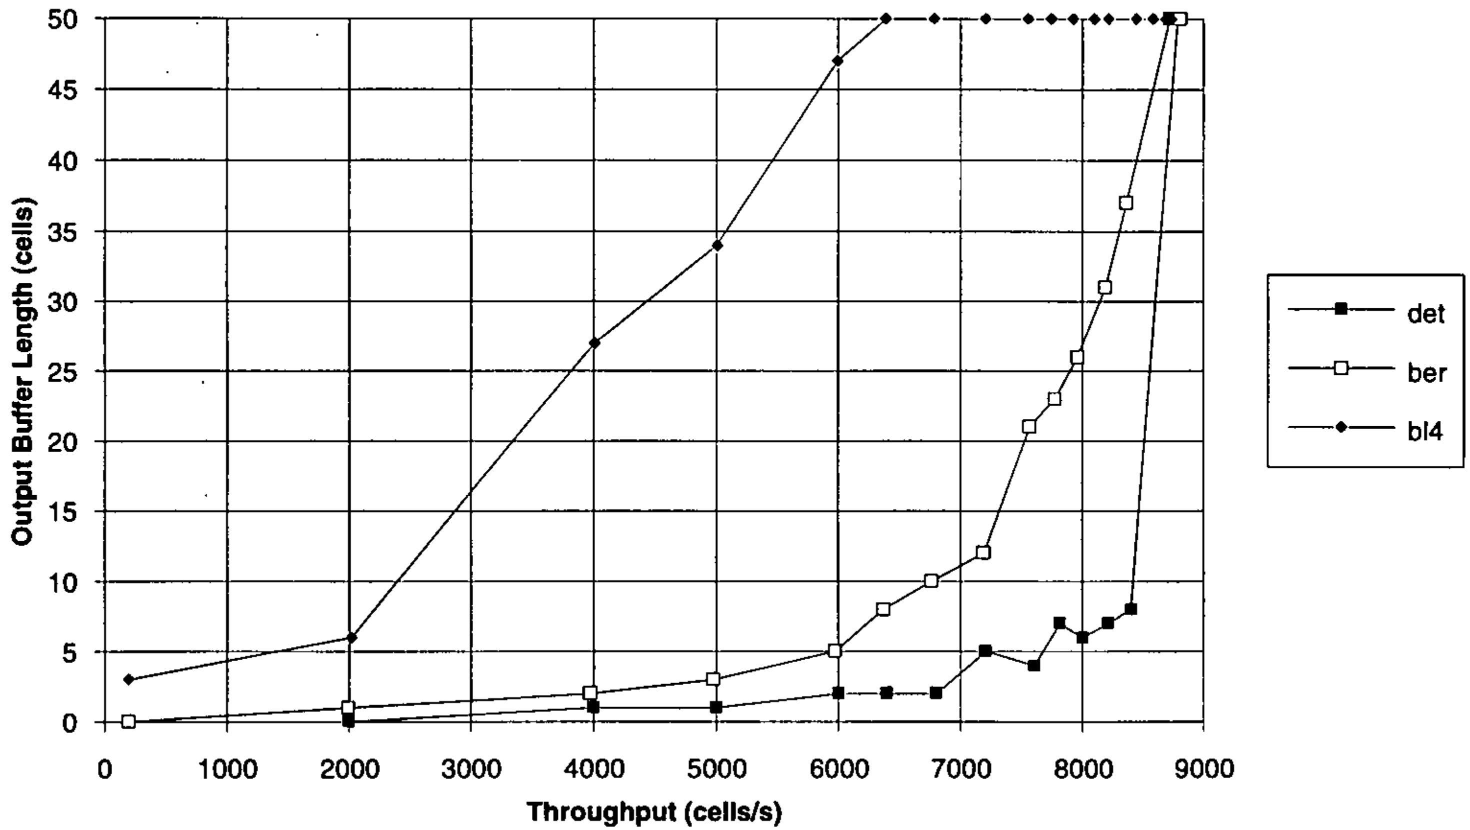

Figure 6.18 Maximum Output Buffer Occupancy : Throughput for 3 Cell Arrival Processes

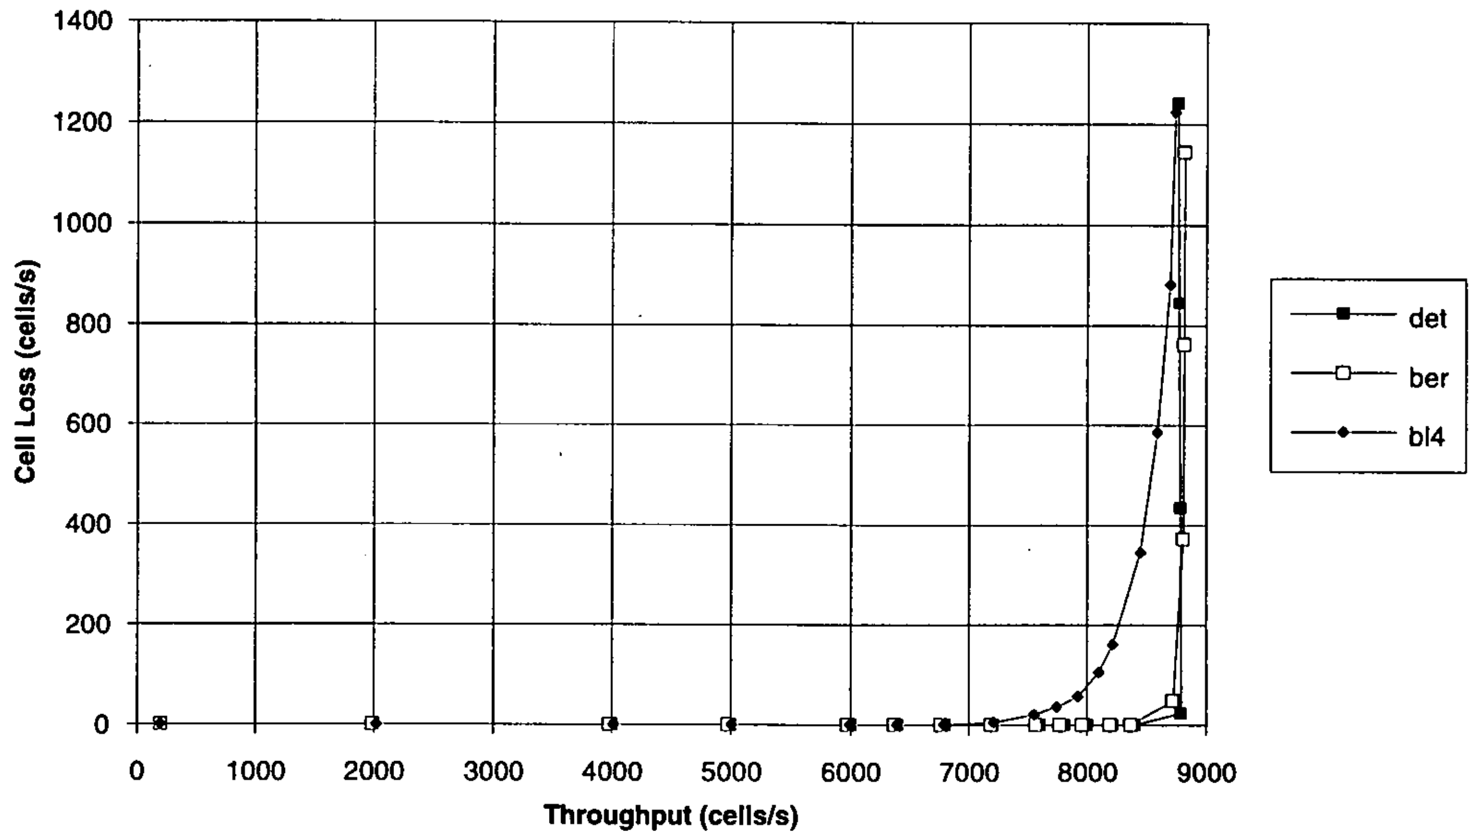

Figure 6.19 Cell Loss Rate : Throughput for 3 Cell Arrival Processes

throughput of cells. In the measured maximum delay, a genuine worst-case performance is measured.

The detailed analysis of cell delay requires the application of some results from queuing theory [117]. A queuing situation is defined by the traffic arrival process and the queue service process. In order to obtain maximum and mean waiting times in the queue, it is first necessary to define these processes. For the ORWELL testbed, and for ORWELL rings in general, the service process is to a large extent defined by the arrival process at other stations. The reason for this is that the occupation of the servers (slots) in transmitting cells from one station to another prevents them from being used by other stations. If cells arrive with a deterministic pattern, the probability that there will be no servers available and that the cell will have to queue is low. Figure 6.20 illustrates a deterministic arrival process, and the number of servers required at any one time does not exceed three. If however cells arrive in blocks at stations for transmission as shown in figure 6.21, the availability of servers will be reduced when blocks arrive at different stations almost simultaneously. When block arrivals do not coincide, the ring will remain idle, so some of the bandwidth is not used.

In figure 6.17 the graph of mean output buffer occupancy against throughput for the three chosen traffic arrival processes is shown. The rate of change from low output buffer occupancy to maximum buffer occupancy is greater for deterministic traffic than for Bernoulli distributed traffic, and batch Bernoulli arrival traffic.

The analysis of the queues is non-trivial as the system which is essentially a four queue four server system, has a statistical distribution of service which is dependent on the statistical distribution of the cell arrival, and does not fit conveniently into well known queuing theory, though this is a possible area for further work.

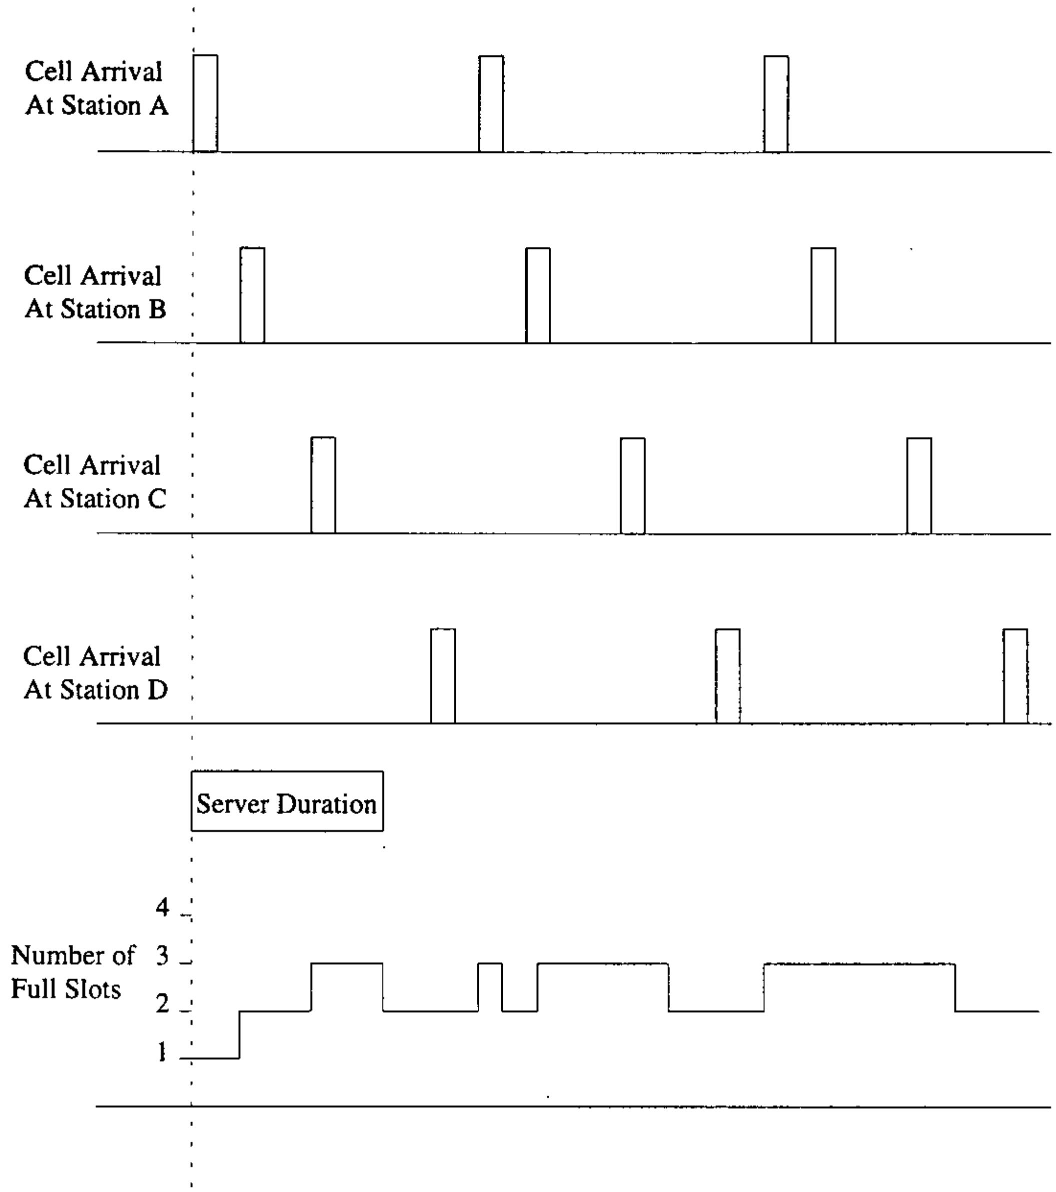

Figure 6.20 Occupation of servers for deterministic arrival process

6.3.3 Scalability of Results

The slotted ring testbed has a ring rate R, equal to 500Kbit/s, much less than practical ATM rates which may be 600 Mbit/s. Measurements taken on the testbed are scalable to larger and faster networks with the same topology and protocol by scaling the ring rate and the slot size. An example of this is that if the maximum throughput for the testbed is 8000cells/second, for a ring rate of 0.5Mbit/s, it would be 800,000 cells/second for a ring rate of 50Mbit/s. The cell size also affects the data throughput of the ring and can be taken in to account in scaling calculations.

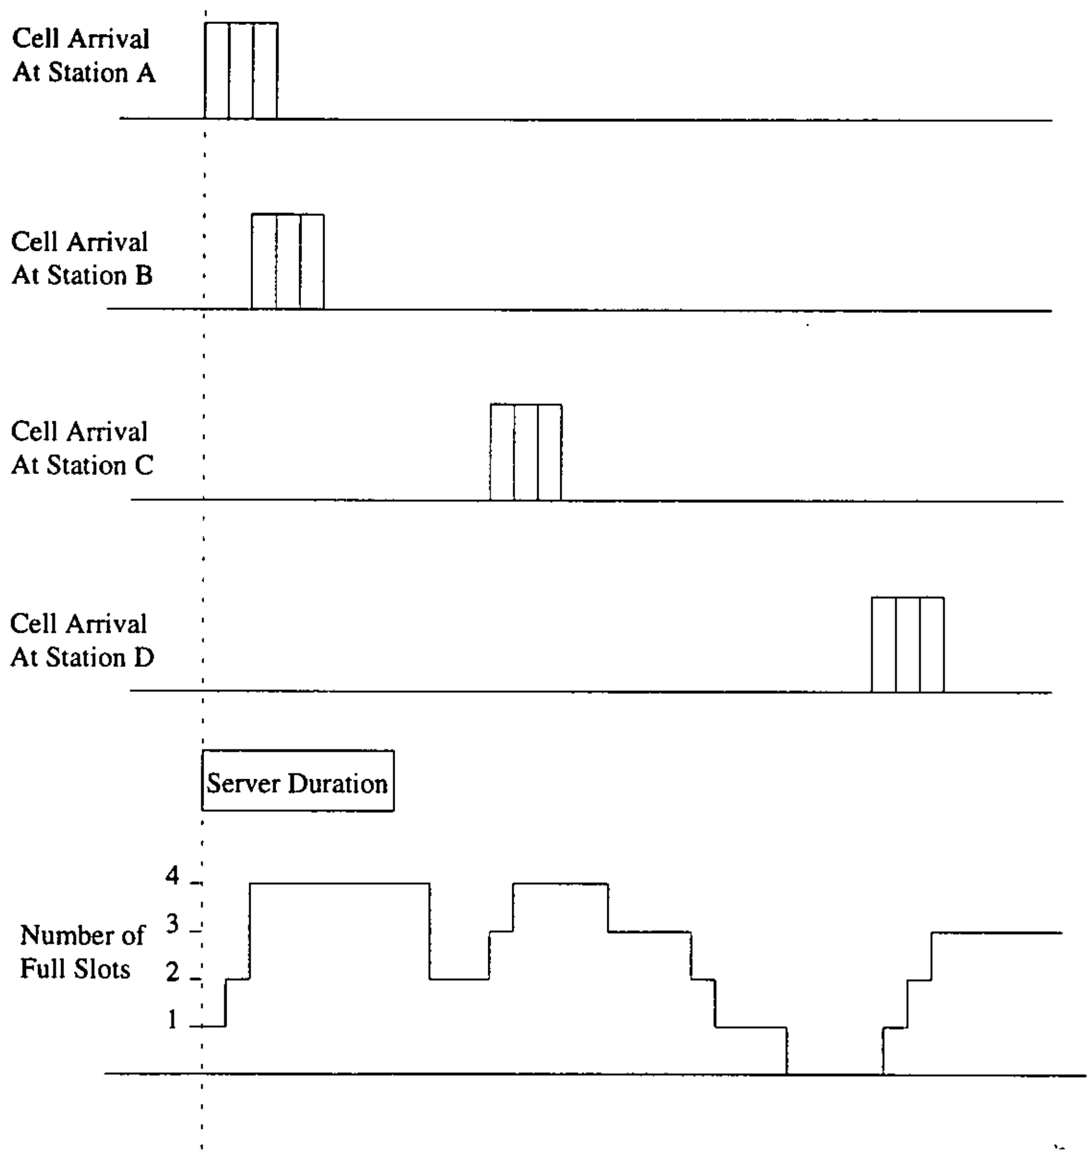

Figure 6.21 Occupation of servers for deterministic arrival process

The validity of scaling of results relating to delay calculations are more difficult to assess. Cell delay consists of fixed overheads such as processing time, and propagation delay, and variable overheads such as queuing delay. The cell delay variation (jitter) can be related to queuing delay which is dependant on buffer size and the cell rate R_{cell}. The main purpose of the testbed is to identify trends in performance which can be applied to larger and faster networks.

6.4 Slotted Ring Testbed Without ORWELL Protocol

The ORWELL protocol introduces an overhead in the transmission of data cells because of the reset mechanism. If this mechanism were to be eliminated, efficiency could be increased up to the theoretical stability condition discussed in section 5.2. The PLDs containing the ORWELL protocol as shown in figure 3.8 were replaced by PLDs implementing a much simplified protocol, allowing only two types of slot: full, and empty. Each station was allowed to fill empty slots, and re-use full slots received by the station. The performance of the new protocol, named the Full/Empty protocol, is compared to that of the ORWELL protocol, in table 6.1, and figure 6.22.

Table 6.1 compares the maximum throughput with output buffer saturated for the ORWELL protocol, the new Full/Empty protocol, and gives stability condition calculated in section 5.2. Both the 4 station configuration and the single station configuration are considered.

Table 6.1 Performance of Ring Protocols under Maximum Traffic Load

Configuration	**ORWELL Di=15**	**Full/Empty**	**Stability Criteria**
	Max.Throughput (Cells/s)	Max.Throughput (Cells/s)	Max.Throughput (Cells/s)
4 Stations	8811	10795	10810
1 Station	4049	5381	5405

Table 6.1 shows that throughput can be maximised by discarding the ORWELL reset mechanism and utilising only full and empty cells. If the intensity of traffic on the network can be controlled by connection admission control, the extra efficiency of the Full/Empty protocol increases throughput and decreases cell delay at higher traffic intensity. Figure 6.22 plots the output buffer occupancy of a station against throughput for a symmetrically loaded ring using a) the ORWELL protocol and b) full and empty slots only. The maximum output buffer occupancy is a good measure of maximum cell delay, as explained in section 6.2, and it can be seen from figure 6.22 that the Full/Empty protocol shows an improvement over ORWELL.

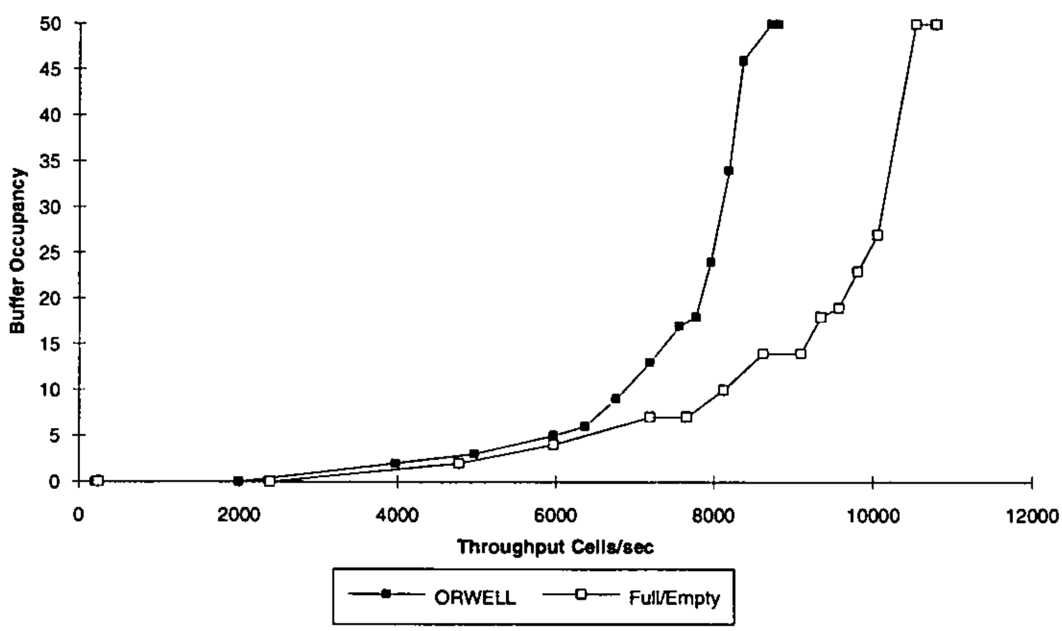

Figure 6.22 Output Buffer Occupancy : Throughput for ORWELL and Full/Empty Protocols

6.5 Summary of the Results Obtained Through Analysis and Simulation

The ORWELL reset mechanism has been introduced to ensure fairness in a slotted ring where a station has the ability to transmit one cell into every passing slot. Without the reset mechanism, stations downstream of a transmitting station could be starved of empty slots in which to transmit. The ORWELL protocol ensures that every station can transmit its Di-allocation of cells in a reset cycle, and the reset interval gives some indication of the loading on the ring. In this chapter and in chapter 5, analysis and simulation have shown that the proportion of time that slots are unable to carry data because of the reset mechanism is an overhead dependent on the Di-allocation, so the efficiency of the ring in carrying data cells improves as Di increases, and therefore the usable bandwidth of the ring increases. One of the main claims for the ORWELL protocol is that the reset interval or its inverse, the reset rate, can be used to measure the traffic level on the ring independently at each node. While it is true that the reset interval provides an indication of cell throughput, it is also dependent on the traffic distribution of the ring, and on the traffic arrival statistics at each station as shown in section 6.1.

In a connectionless data transmission environment, the fairness aspect of a LAN protocol is considered important to ensure access to the transmission media by all competing stations. In a connection oriented transmission environment however, the access to media is controlled by a connection admission system, and the most important criteria for accepting another connection is that quality of service can be guaranteed for new and existing traffic. It is therefore possible to operate a slotted ring with only full and empty slots, using a connection acceptance process to limit the traffic on the network, and dispense with the overhead of the ORWELL reset mechanism.

The maximum cell delay is likely to be a key indicator of the quality of service that a network can offer to real-time traffic. The maximum cell delay time is bounded by the time it is allowed to queue before being transported on the slotted ring network as shown in section 6.2. A transmitting station will not know the delay time of cells it is transmitting to a destination station, but will be able to sense if buffer overflow occurs at its own output buffer. Appropriate dimensioning of the transmit output buffer of a station provides a mechanism to bound maximum cell delay, and to detect if cells are being lost due to buffer overflow. The occupancy of the output buffer may also be used as an indicator of potential cell loss. The results obtained from the testbed allow a study of heterogeneous traffic to be made in chapter 7, including an investigation of congestion management strategies.

Chapter 7

ATM Congestion Management

7.1 Introduction

In a Synchronous Transmission Mode (STM) system such as the global TDM telephone network, connections are established from one user to another during the signalling or dialling process. If there is insufficient capacity in the network to support the desired connection, through a lack of time slots over a transmission route, or lack of capacity at a switching exchange, the call is blocked and the user receives the engaged tone to inform him of this. If a connection is established, that connection is guaranteed until one or other of the users clears the call. In this way, there is a natural mechanism, limited by the provision of lines and time-slots, to determine if a new call can be accepted or not. Once a connection is accepted, the lines and time-slots are reserved for that call alone, hence the new call does not affect the traffic carried by any existing calls.

The ATM network is a connection-based network, but does not have a clear and well defined mechanism such as availability of time-slots, for the acceptance of new connections, as does STM. The cell based ATM network is prone to congestion at its switching points, where traffic from various sources is concentrated, arriving on the input transmission paths, and being switched to various output transmission paths according to the cell header address label. Inevitably there will be contention amongst incoming packets for the same output path and buffering must be provided to ensure that cells can be accommodated. The introduction of cell buffering at the switching points introduces both delay in the transmission time for a cell as it is queuing for service, and the finite probability of buffer overflow and consequently cell loss. The primary aim of ATM network congestion management is to ensure that both cell delay and cell loss probability are kept within acceptable bounds.

If the number of connections accepted by an ATM network is continuously increased, a point will be reached when cell congestion occurs with associated buffer overflow, cell delay and cell loss. Admission control is therefore a key element to the congestion avoidance strategy. The analysis of B-ISDN networks to enable an admission control framework to be defined has been the subject of many papers [41], [54], [122], [123], the implementation of which relies on indicators of both network usage, and the effect on the network of accepting a new connection. Some proposed schemes are discussed in more detail in section 7.2.3.

In ensuring congestion avoidance, the network is obliged to police the rate of cells being transmitted into the network from sources where the connection has been accepted. This is an obligation on the network operator to prevent the abuse of an accepted connection, and to safeguard other users from the effect of excessive cell origination by a particular source. Section 7.2.4 discusses methods of source policing.

In addition to admission control, congestion can be avoided by use of a well chosen service discipline at the switching points. This involves classification of traffic into different classes, and offering a different quality of service to each class. The purpose of the classification is to enable one traffic class to be prioritised over another. The issues raised by heterogeneous traffic, and prioritisation will be discussed in sections 7.2.1 and 7.2.2.

7.2 Strategies for ATM Network Congestion Management

7.2.1 Heterogeneous Traffic Classification

In the B-ISDN, the variety of data sources and range of required bandwidths is very wide as has been discussed in chapter 2. Traffic sources may be classified into two main groups, namely those having stringent delay requirements (in general, interactive voice and video services), and those which are delay tolerant such as electronic messages, facsimile, and data file transfers. Services may also be classified as being cell loss sensitive or cell loss tolerant. No standards or CCITT recommendations have yet been developed for the classification of traffic sources, which in turn means that the quality of service that is guaranteed for each accepted connection has itself not been defined. The various studies carried out in a heterogeneous traffic environment make their own assumptions about the number of grades of service, and the quality of service attached to each grade of service.

The philosophy behind ATM as a transport mechanism for B-ISDN suggests that traffic should be handled in as simple and integrated fashion as possible which indicates that a low number of traffic classes should be provided to cover the range of cell loss and delay requirements of all services [41]. Although a single grade of service is not ruled out, the stringent delay requirements applicable to real-time voice and video services clearly do not apply to data traffic or stored voice and video services, and a scheme with at least two grades of service seems likely [51]. This type of traffic classification scheme will be used in the simulations carried out in section 7.3, it allows for the buffering of highly bursty traffic which is not delay sensitive, yielding an improved statistical multiplexing performance.

7.2.2 Prioritisation of Traffic Classes

When more than one traffic class is introduced to an ATM network, the question of prioritisation of cells of one class over those of another arises. For example, 2 cells may arrive simultaneously at a switch from different input paths both having the same destination path. The order in which they are serviced, and the priority in which they are discarded if the output line buffer is full, is decided by the scheduling or priority scheme operating at the switch. Some alternative priority strategies will be discussed in this section.

The priority queuing strategy is composed of the service discipline, used to select which cell is serviced first, and the buffer access control discipline which deals with new arrivals to the queue. Accordingly, there are two major categories of priority queuing depending on where the priority rules are applied [33], these are the service discipline (also known as time/delay discipline or input discipline) and the buffer access discipline (also known as space discipline or output discipline) as shown in figure 7.1.

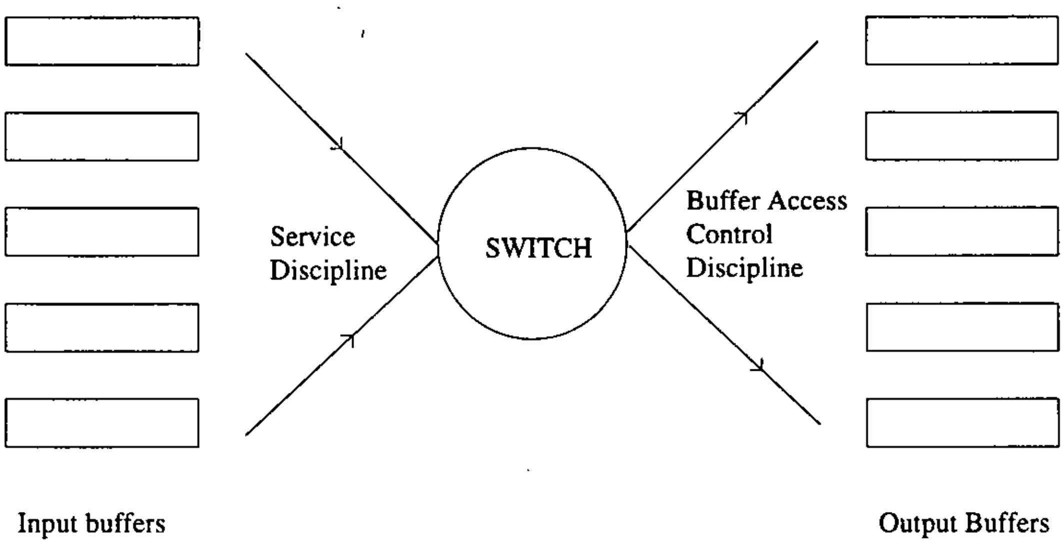

Figure 7.1 Priority Queuing Disciplines

There is a further classification of priority queuing strategies in to pre-emptive and non pre-emptive strategies. In the pre-emptive case, for a service discipline the cell with the highest priority is serviced immediately, even if another cell of a lower priority class is waiting. For the buffer access discipline in the pre-emptive case, cells of a higher priority may push out cells of a lower priority from the queue. The non pre-emptive cases do not allow this to happen.

Prioritisation is used to favour cells of one class over those of another class. The success of a priority queuing strategy depends on its ability to ensure that the benefits gained by the high priority class do not outweigh the penalties to which the low priority class is subjected.

The implementation of a buffer access control discipline priority queuing scheme depends on the physical organisation of buffers in relation to the input streams of cells at different priorities, and to the access to servers. Figure 7.2 shows four schemes of buffer organisation. The first is complete buffer sharing with pushout (CS+PO), where a single buffer is employed, both traffic streams joining at the head of line (HOL), but class 1 (high priority) cells may push out class 2 (low priority) cells if the buffer is full. The second method is partial buffer sharing (PBS), a limit is imposed on the buffering available to both classes, but class 2 cells may only join the queue when it is below a certain length. The third scheme is complete buffer partitioning, but with complete bandwidth sharing (CBP+CBWS), where prioritisation is achieved by dimensioning the buffer space for each priority class of cells separately. Finally a complete partitioning (CP) of priority classes is possible.

The analysis of priority queuing systems is more complex than for a single class of cells. Analysis also makes assumptions about the cell arrival process and cell service process. Simulation has also been widely used to give results on queuing which show that there are significant advantages in a priority queuing scheme where the QOS requirements of the classes in terms of loss and delay are significantly different [34].

7.2.3 Traffic Intensity and Equivalent Capacity

Because of the statistical multiplexing of connections in the ATM network, and the variable inter-arrival times between packets from a given source, the traffic intensity at any point in the network may vary considerably over a period of time. In order to make a prediction of the impact on existing traffic of accepting a new connection, the estimation of traffic intensity is an key factor. One option for the network designer is to allow sufficient capacity for each traffic source to be constantly transmitting at its peak cell rate, i.e. using non-statistical multiplexing. However, the bursty nature of some traffic sources means there are considerable efficiency gains to be made by statistical multiplexing if a reliable estimation of the equivalent capacity of the link can be made.

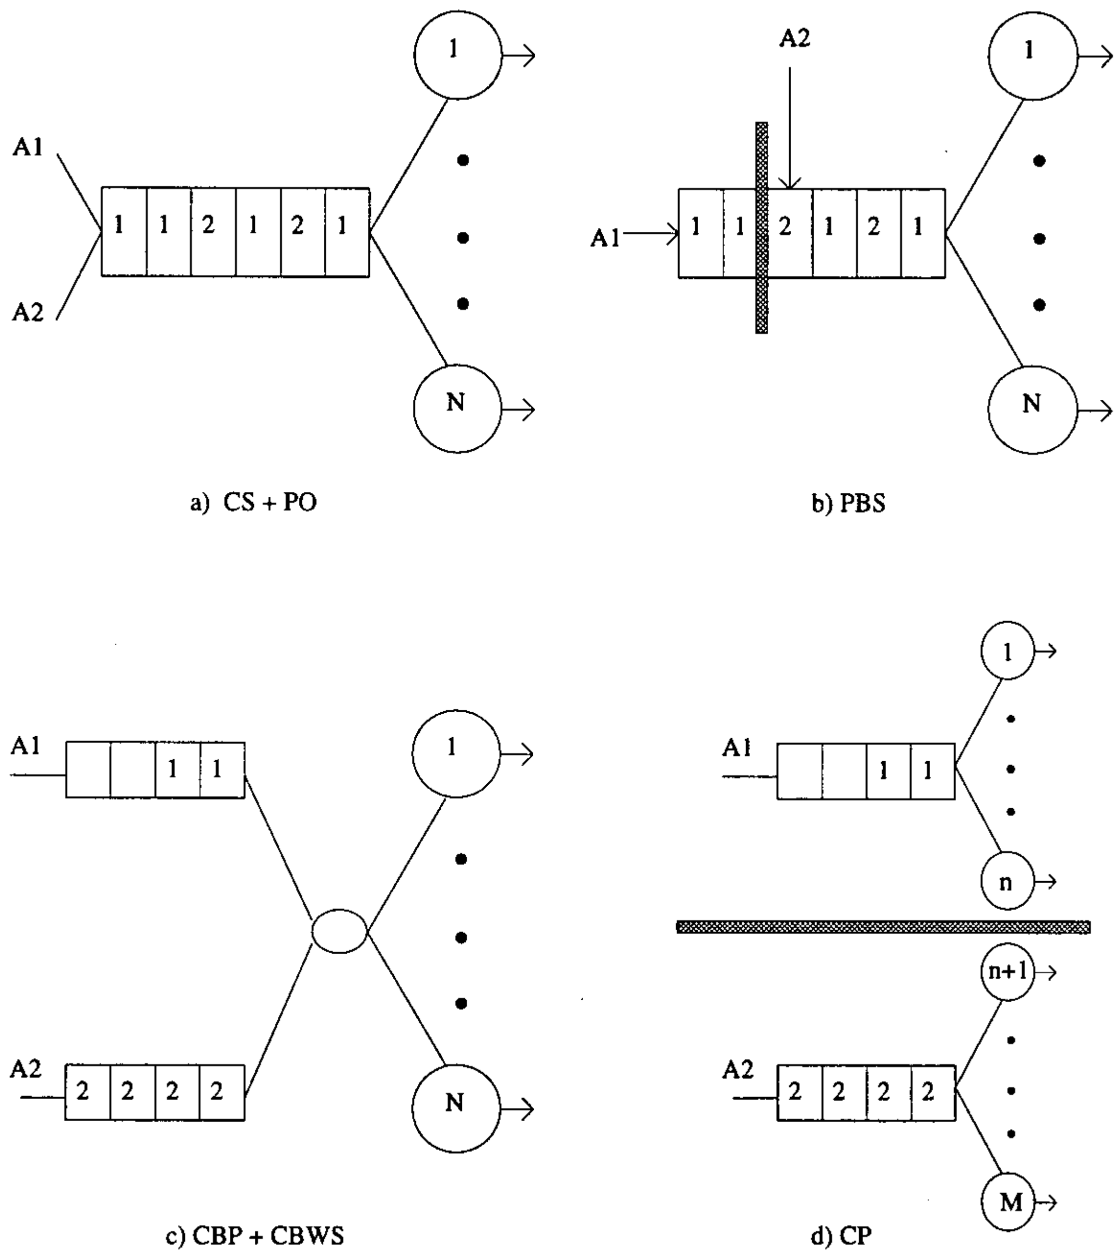

Figure 7.2 Buffer Allocation in Priority Queuing Schemes

Traffic intensity can be estimated from a knowledge of the number of connections accepted on a network link, and the characteristics of the traffic sources. The performance of a switching node can be characterised for various combinations of traffic type, and a set of tables or curves of acceptable connection numbers and types can be stored at the switch. This approach to bandwidth allocation known as a Schedulable Load Area, is discussed in section 7.2.5, and uses analysis or simulation to generate a series of curves to be used as guidelines when new call requests are considered by the network. The drawback of this method is the additional storage required to save the precomputed curves at network nodes, and whether the curves are susceptible to changes in the statistics of the traffic sources.

An alternative to using predetermined values to estimate the bandwidth requirement of a single connection, or multiplexed connections, is to determine the 'equivalent capacity' of the connection or group of connections [39]. The equivalent capacity of a set of connections multiplexed on a link may be defined as the amount of bandwidth required to achieve the desired QOS, in terms of buffer overflow probability, given the offered aggregated bit rate generated by the connections. It is a function of individual connection characteristics and available network resources such as buffers.

Equivalent capacity must employ a simple calculation without high processing overheads using measurable traffic parameters such as mean cell rate, peak cell rate, variance of cell rate, maximum burst length or other measures of cell distribution. Traffic flow measurements are undertaken by counting the cell arrivals in a certain time-window, to identify these parameters, and the time-window used to make the estimates is chosen to avoid autocorrelation of sources [103]. Decision making in the call admission process can be done in real-time, by calculating the upper bound cell loss probability for existing traffic and a worst case estimate of the new requested connection [54], [56]. Admission control by equivalent capacity bandwidth reservation is discussed in section 7.2.6.

7.2.4 Source Policing Mechanisms

During the call set-up phase of an ATM connection, the user negotiates a set of traffic parameters described in terms of peak cell rate, mean cell rate and burstiness with the network. It is however possible for the user to violate the agreed 'contract' by transmitting cells up to the maximum capacity of the User Network Interface (UNI) which would be unacceptable to the network operator both from the position of charging policy and the possible interference in terms of QOS with other users. To prevent violations of the traffic contract by accident or maliciously or for economic gain, a network function called Usage Parameter Control (UPC) is required, and is defined in CCITT recommendation I.311. UPC is more often described as source policing, and several methods have been developed to fulfil the function [57], [58].

The action taken by policing mechanisms on the detection of excessive cell input to the network may take several different forms. Cells may be discarded at the policing point, or may be marked for deletion if contention for buffer space occurs during transmission. This would mean that no penalty was incurred by the user unless a cell volume based charging scheme applied. Alternatively the call control function could be notified of detected violations, and the call could be terminated, or a punitive charge applied, although the reactive time of this measure might not be short enough to avoid network congestion. Some suggested mechanisms for source policing will be reviewed.

The Leaky Bucket Mechanism [59], [61], [124], consists of a counter which is incremented each time a cell is generated by the source, and is decremented in fixed time intervals to a minimum of zero as shown in figure 7.3. If the cell arrival rate exceeds the decrement rate, the counter value will rise, and once a predefined limit is reached cells will be marked or discarded until the counter value falls again. The Leaky Bucket scheme is easily implemented and can be exactly modelled by the *G/D/1-S* queuing system which describes a queue with a general traffic arrival process, single deterministic server, and a limited buffer length of S. The service time of the model is equal to the decrement interval of the Leaky Bucket, and the number of cells in the queue plus the server is equal to the counter limit N. Solutions for the stationary system are know which allows analysis of this mechanism [61].

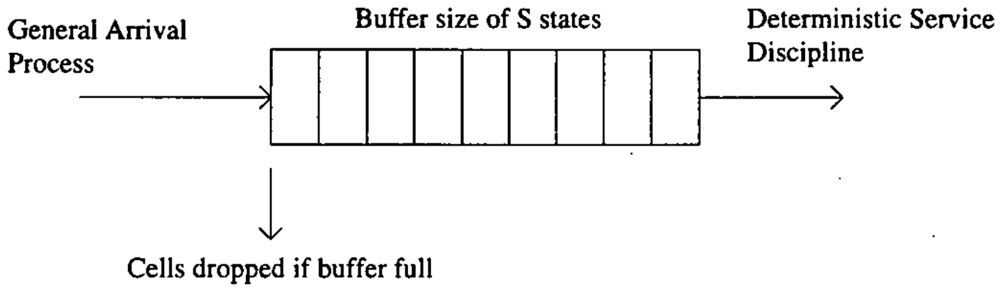

Figure 7.3 The Leaky Bucket Source Policing Mechanism as a G/D/1-S Queue

Various window schemes have been proposed for source policing such as Moving Window, Stepping Window and Jumping Window [125]. The window mechanisms work by accepting a number of cells up to a maximum of N cells during a predetermined window interval. Each scheme employs a different method of resetting the window interval. In the Jumping Window scheme each new interval starts immediately as the previous interval finishes, whereas the stepping window starts with the first new cell arrival after the end of the previous window interval. The performance of all of the window schemes are highly dependent on window interval length, and on the traffic arrival statistics.

7.2.5 Bandwidth Allocation using an Admissible Load Region

As described in section 7.2.3, some call admission strategies are based on the concept of predetermined information in the form of tables or curves which are used by a switching point to accept or reject new calls. In the case of homogenous traffic types where there is only one class of traffic, a scalar quantity giving the maximum number of calls to be accepted would suffice. In the case of two traffic classes, with differing cell distributions, and cell arrival rates, as well as differing QOS guarantees, a call acceptance region could be defined as described in figure 7.4, by off-line simulation or analysis. The existence of a reliable Admissible Load Region (ALR) depends on an effective bandwidth being attributable to all calls of a particular traffic classification, then all calls in that class can be treated identically.

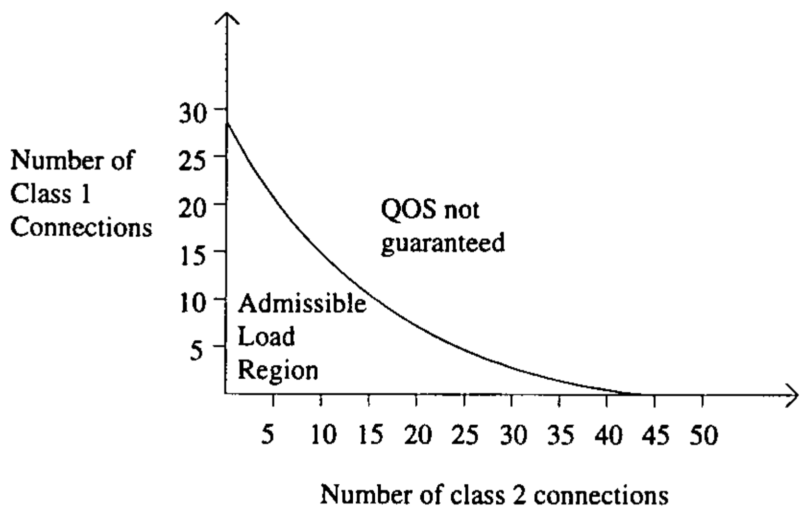

Figure 7.4 Admissible Load Region

For more than two traffic classes, a multidimensional surface is required to implement such a method of admission control. As the number of classes increase, this method become more cumbersome requiring more storage space, greater computational time and greater simulation time to determine the boundaries. A framework of adaptive admission control is proposed in [118] which describes three methods to store the call acceptance area and uses adaptive control to allow corrective action to be applied to the accept / reject boundary if traffic characteristics vary over time, and if initial off-line simulation is inaccurate.

The use of an admissible load region has been investigated in the MAGNET programme [42], [43], for a multi-class traffic system. The QOS constraints defining different traffic classes are specified in terms of cell loss probability, and maximum consecutive cell loss. The latter term is clearly significant in that single isolated cell losses are less consequential than bursts of consecutive cell losses.

7.2.6 Admission Control by Bandwidth Reservation

Whilst the definition of an ALR for all combinations of traffic classes provides one method of ensuring network access is controlled, the major drawbacks of this method are that discrete traffic groups must be defined, and the amount of stored data required at each switching point increases rapidly as the number of classes increases. Another method of allocating bandwidth in an admission control mechanism, is to use the calculated equivalent bandwidth of a call. In this case, each switching point on the route of the proposed connection would reserve a bandwidth equal to the effective bandwidth of the call for the period of the connection, and the switch will not allow the sum of reserved bandwidths to exceed its total capacity on the particular output connection requested by the call. Figure 7.5 illustrates this concept with a multiplexer fed by a number of ATM streams each having a certain equivalent bandwidth. The sum of the input line equivalent bandwidths must not be allowed to exceed the output link equivalent capacity. The equivalent capacity of the multiplexer output link will depend on the output buffer size and QOS guarantees given to the service.

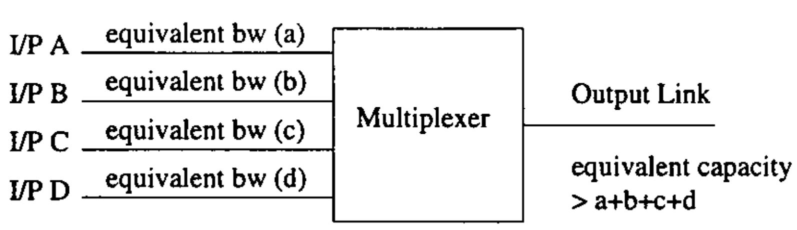

Figure 7.5 Reservation of Equivalent Bandwidth

One attraction of this kind of bandwidth allocation is that an equivalent bandwidth calculation based on peak and mean bit rates, and maximum burst length can be declared during signalling, and policed by the network. This means both the network provider and the user, are aware of the shaping of the input traffic stream should it violate the constraints of the policing mechanism. The user can determine the amount of buffering between the application and the input to the network, and choose a balance between speed of data transfer and cost of connection.

7.2.7 Analysis of Cell Loss Probability

The analysis of ATM queuing systems for cell loss probability is the subject of many recent papers [126], [127], [128], [129], [130], as well as being included in the majority of papers concerned with admission control and prioritisation. At any ATM switching point where cell buffering occurs, there is a certain probability of the incoming cell rate from various multiplexed traffic sources causing buffer overflow. The parameters which must be defined in examining the queuing model are: traffic source characteristics, number of input sources, buffer size, and service characteristics.

The cell arrival process depends on the traffic sources, and their characteristics which may be deduced from traffic flow measurements, or from assumptions about the statistical characteristics of the sources, and their number. The analysis of cell loss probability can be carried out using two main methods. The first is by exact methods where the solution of standard queuing problems such as the M/D/1-S queue yields exact probabilities of buffer overflow due to variation in cell inter-arrival times. A second method is known as the 'bufferless fluid flow' model [56], which ignores the short-term dynamics resulting from the discrete nature of the ATM cells and their asynchronous arrivals, but takes account of the burst-level statistics of traffic sources such as peak cell rate, mean cell rate, and maximum burst period. A combination of the two methods is sometimes used [122].

7.3 Simulation of Heterogeneous Traffic on the Testbed

In order to develop a scheme of access control for the slotted ring testbed, two classes of traffic were defined with different cell arrival processes and different QOS requirements, to simulate real-time and non-real-time traffic. The admissible load region of a station was calculated under different traffic intensities for each class of traffic, and under different conditions of network loading. The performance of prioritisation of real-time over non-real-time traffic was measured, as was the performance of the ORWELL protocol compared to the Full/Empty protocol.

7.3.1 Definition of Traffic Classes and Quality of Service

Two traffic classes are proposed in this study of the testbed, one representing delay-sensitive constant bit rate sources to be called Class 1, and one representing bursty, delay tolerant sources to be called Class 2. The method chosen to model these sources is simple and requires little computational overhead. In both cases a Bernoulli arrival process was used, with the probability of success i.e. a cell arrival, or batch of cell arrivals, being proportional to the number of connections admitted.

For Class 1 traffic the Bernoulli trial interval is 400us, and each connection accounts for 2% probability of a cell arrival, hence 50 connections ensures a cell arrival in the interval. This situation models a multiplexer connected to 50 input lines, checking for a cell arrival at each one in rotation at 400us intervals, as shown in figure 7.6. For Class 2 traffic, the batch size is fixed at 10 cells, and though the Bernoulli trial interval is again 400us, the probability of each connection yielding a batch arrival is 0.2%. The QOS demands of each traffic type must now be defined. Class 1 traffic has a maximum transmission time limitation imposed on it of 15 ms, so that if this transmission delay limit is exceeded the slot is considered to be lost. Class 2 traffic does not have a delay limitation, but does have a limited cell loss requirement. Both Class 1 and Class 2 traffic are guaranteed a cell loss rate less than 10^{-3}. By defining

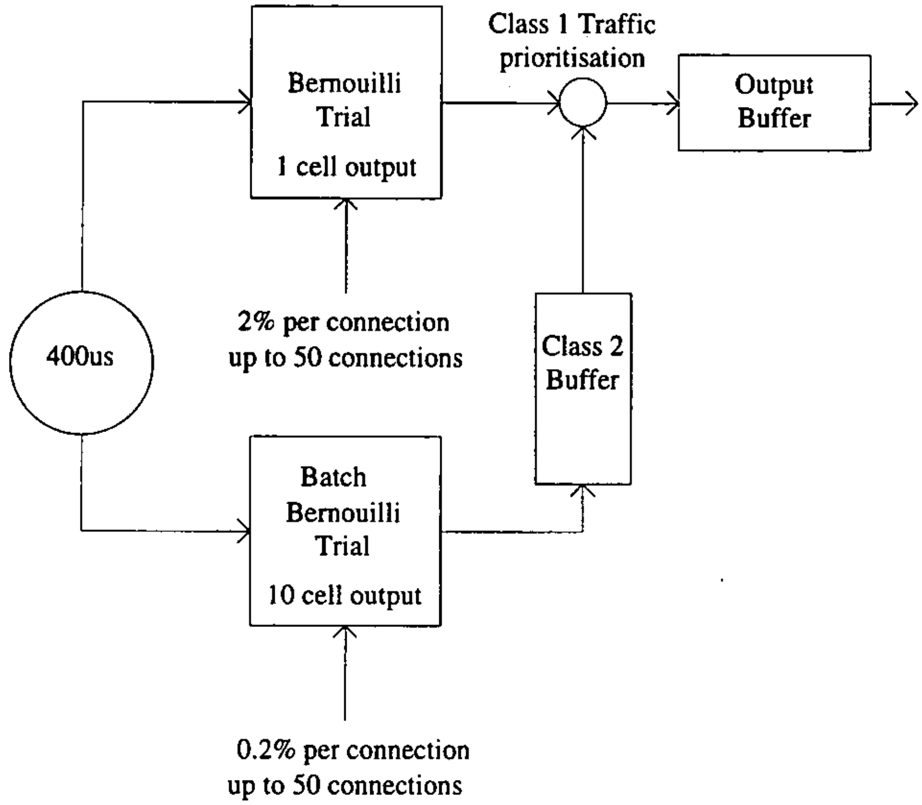

Figure 7.6 Traffic Models for Heterogeneous Traffic Simulations

such a cell loss limit, with Class 1 traffic having delayed cells added to the cell loss, an Admissible Load Region (ALR) in terms of connections accepted within QOS constraints, can be defined as a performance indicator for a testbed station under various buffering and priority schemes. Performance under varying network loading can also be considered.

7.3.2 Prioritisation and Buffering Schemes

In order to discuss prioritisation, the buffering scheme used for each traffic class at the testbed station must be described. Figure 7.7 shows two alternative schemes. Scheme (a) applies no prioritisation. Cells arriving individually or in batches from either traffic source in each 400us period are written in to the output buffer. If the output buffer is full the cells are lost. This is a complete buffer sharing scheme without prioritisation. Scheme (b) prioritises Class 1 cells over class 2 cells and introduces extra buffering for the bursty class 2 traffic. Cells from the class 2 buffer are only written into the output buffer if there are no class 1 cell arrivals, thus only one cell can be passed to the output buffer in a any 400us period. This is a partial buffer sharing scheme with prioritisation of class 1 cells over class 2 cells, but it is not a pre-emptive scheme.

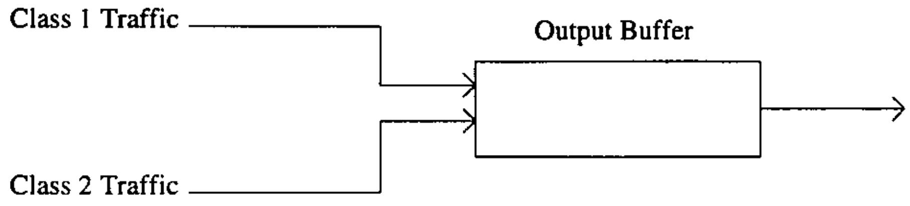

Scheme a) no prioritisation

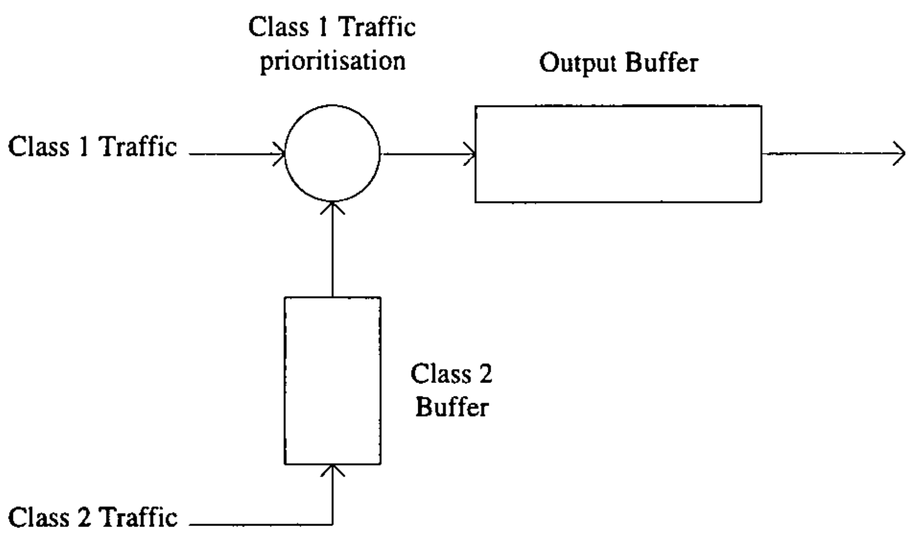

Scheme b) prioritisation of class 1 traffic

Figure 7.7 Testbed Station Output Buffering

The ALR graphs were obtained for both of the buffering schemes shown in figure 7.7. The network was operated with only the observed node allowed to transmit on to the testbed. Figure 7.8 shows the ALR for scheme (a) with no priority, and for scheme

(b) which employs priority. It can be seen that the introduction of priority classes improves the ALR of the station at high levels of bursty traffic, because of the additional buffering introduced to traffic which is not delay-sensitive.

7.3.3 ALR on a busy network

The level of traffic on the slotted ring testbed will affect the ALR for any station on the network. To observe this, the ALR for a single station was plotted with different levels of connections at the other network stations. Figure 7.9 shows the ALR region for 4 levels of traffic on the network. The ALR is given for node 0 in each case but the traffic transmitted at other stations is:

BG0 stations 1 ,2, and 3, transmit no traffic

BG30 stations 1 ,2, and 3, have 30 class 1 connections and 30 class 2 connections

BG35 stations 1, 2, and 3, have 35 class 1 connections and 35 class 2 connections

BG40 stations 1, 2, and 3, have 40 class 1 connections and 40 class 2 connections

The ALR shows the combination of connections which result in a cell loss rate of less than 10^{-3} at any station.

7.3.4 ALR Variation With Class 2 Batch Size

Figure 7.10 shows the ALR with a Class 2 batch size of 5 and of 10. It can be seen that the combination of traffic of dissimilar statistical characteristics presents problems for statistical multiplexing. When the batch size is equal to 5, the equivalent capacity of a buffered class 2 connection is similar to that of a class 1 connection which is shown by the linear way in which connections of either class are combined over most of the ALR. When the batch size is increased to 10, there is difficulty in combining the two traffic classes at high levels of class 2 traffic. Both the priority buffering scheme and the very bursty nature of the class 2 traffic contribute to this effect.

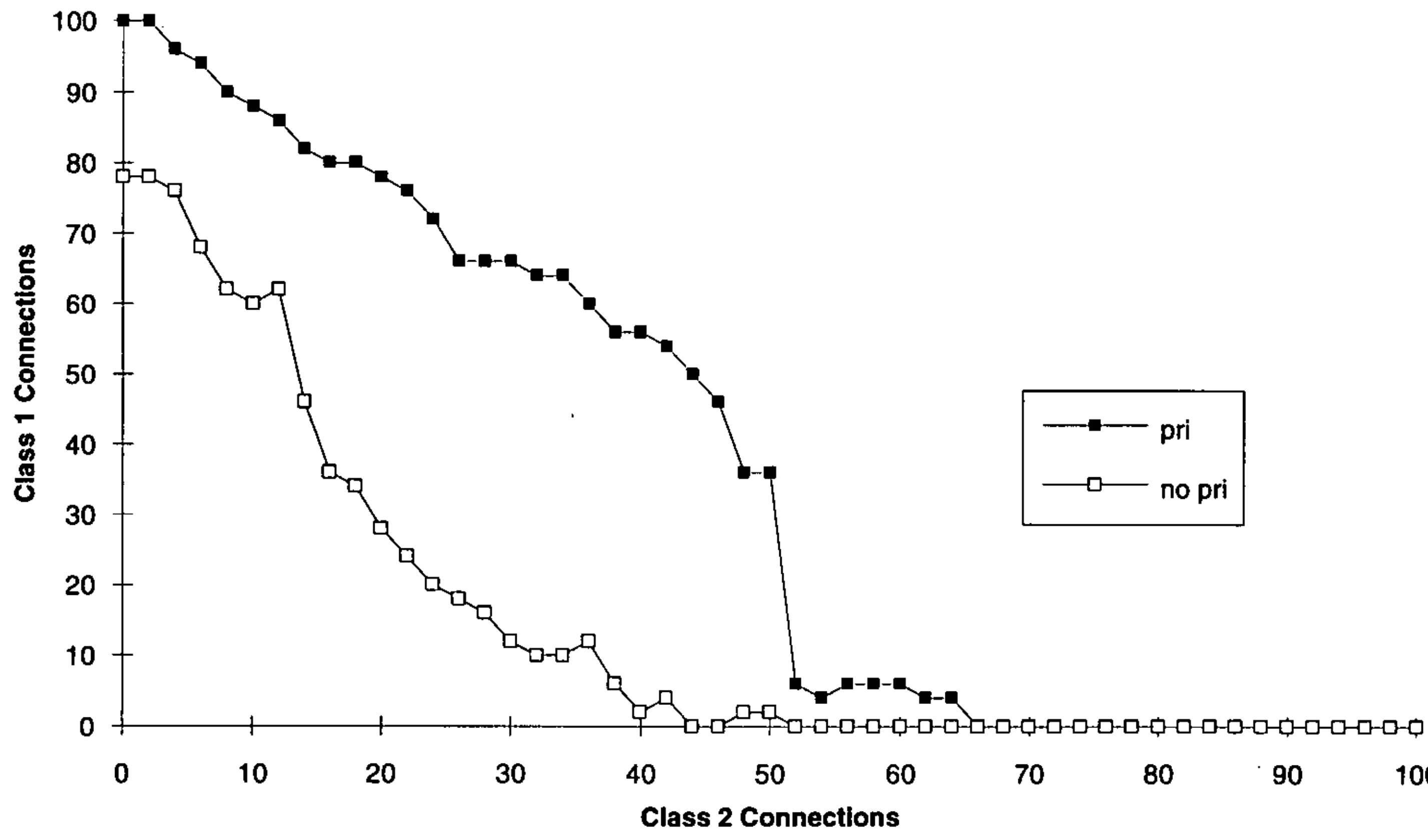

Figure 7.8 ALR for Prioritized and Non-prioritized Buffer Scheme

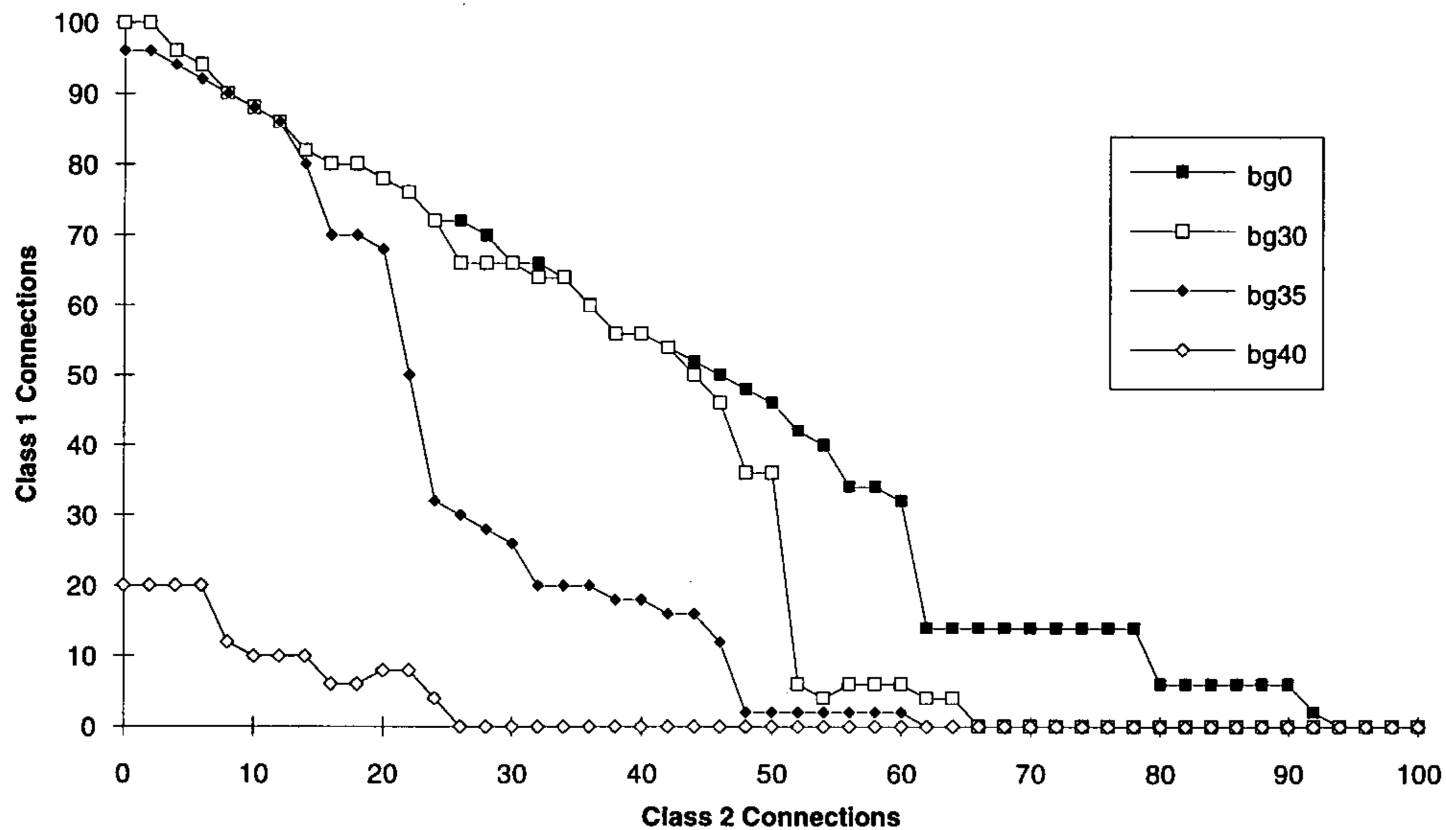

Figure 7.9 ALR for Network Station with 4 Levels of Background Traffic

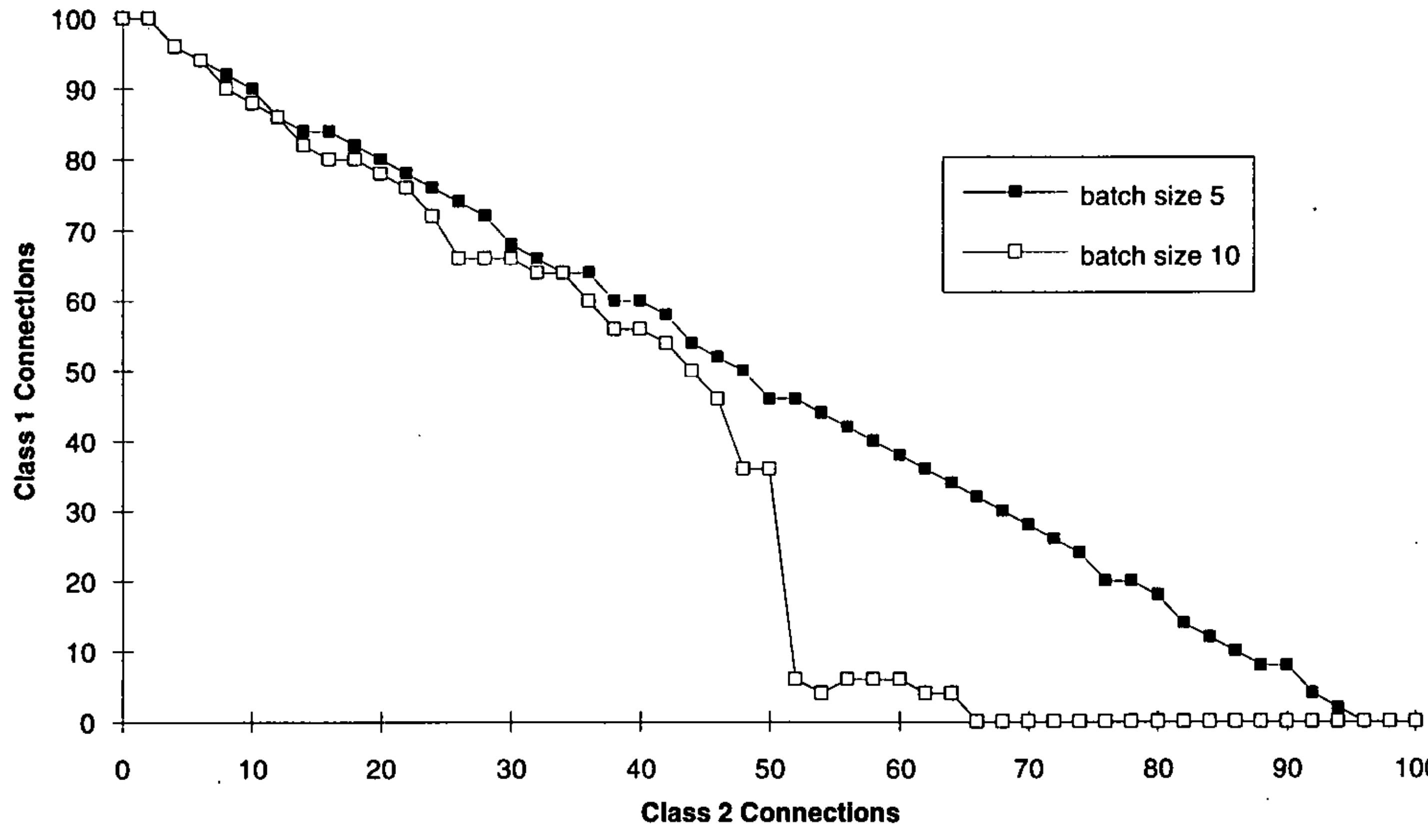

Figure 7.10 ALR for Class 2 Batch Size of 5 and 10

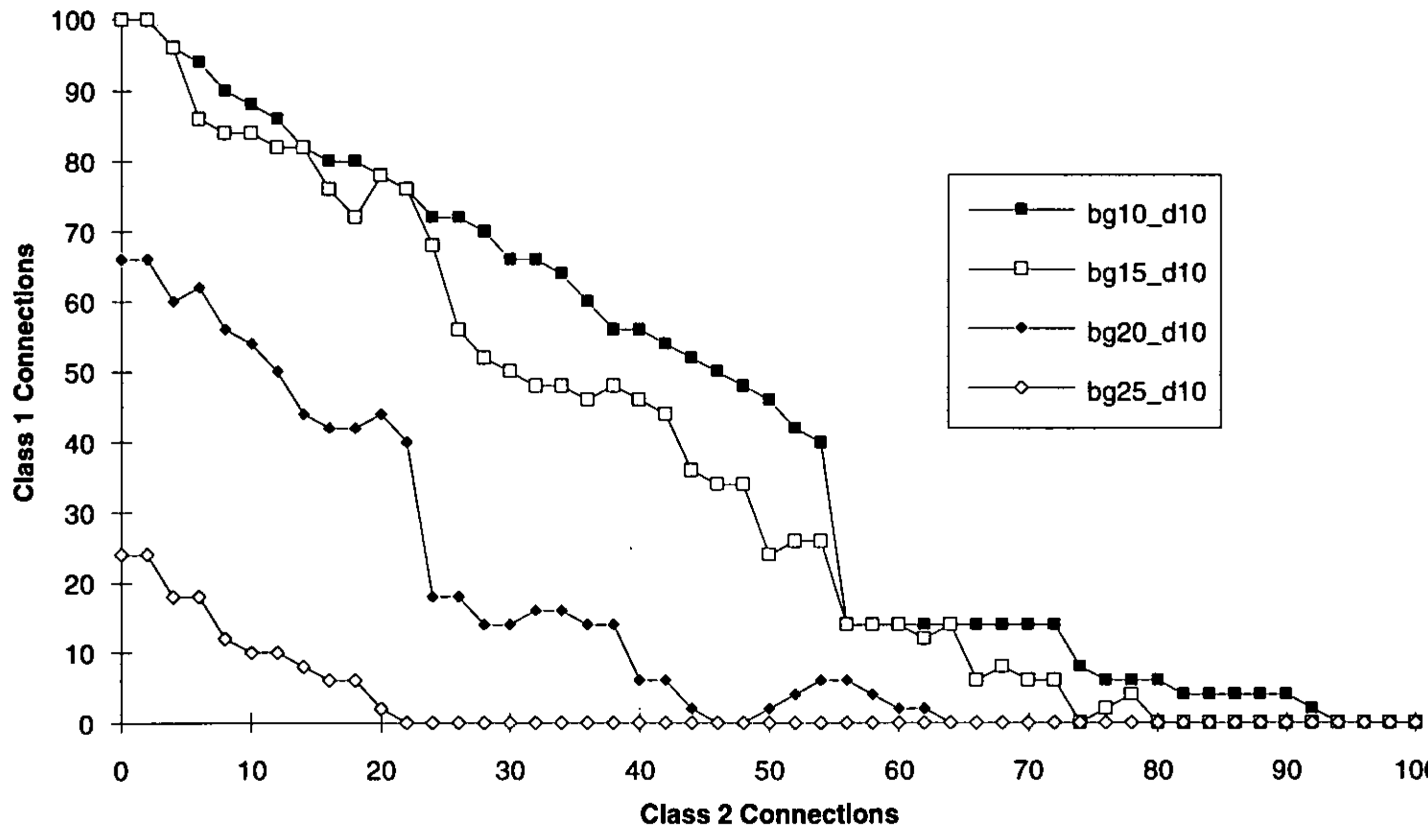

Figure 7.11 ALR with ORWELL Protocol (Di=10) and Full/Empty Protocol

7.3.5 ALR with ORWELL protocol and with no protocol

The ALR was plotted for the testbed using the ORWELL protocol with Di = 10, in figure 7.11. Various levels of traffic were applied to the network in a similar manner to figure 7.9 which relates to the testbed using the Full/Empty protocol. The ALR is given for node 0 in each case but the traffic transmitted at other stations is:

BG10 stations 1 ,2, and 3, have 10 class 1 connections and 10 class 2 connections

BG15 stations 1 ,2, and 3, have 15 class 1 connections and 15 class 2 connections

BG20 stations 1, 2, and 3, have 20 class 1 connections and 20 class 2 connections

BG25 stations 1, 2, and 3, have 25 class 1 connections and 25 class 2 connections

The ALR shows the combination of connections which result in a cell loss rate of less than 10^{-3} at any station. By comparison with figure 7.9, it can be seen that higher levels of throughput can be achieved without the ORWELL protocol.

7.4 Summary of Congestion Management Studies

In this chapter a study of bandwidth allocation strategies for an ATM network has been undertaken. The purpose of the strategies is to ensure that a QOS can be offered to all connections accepted by the network, and to ensure the maximum utilisation of available bandwidth. Classification of traffic into groups with different QOS requirements has been considered, and a variety of prioritisation and buffering schemes discussed. The results obtained from the testbed, though qualitative, show that prioritisation of traffic classes can lead to an improved ALR where the traffic classes have different QOS requirements, and different statistical cell arrival profiles.

Admission control is the principal means of controlling congestion on a connection based telecommunications network. During the call set up phase of a connection, sufficient bandwidth must be reserved to ensure the QOS of new and existing calls. The two approaches to reserving bandwidth described in the literature are by defining an ALR based on the admissible region of mixes of traffic class, or by estimating the equivalent bandwidth of each call, and the equivalent capacity of a link. The testbed has been used to demonstrate how an ALR can be determined for a heterogeneous traffic environment with two traffic classes.

The results of simulations performed on the testbed show that a prioritising output buffer improves the station ALR, where the QOS requirements for each traffic class are different.

The ALR boundary would be used to determine access to the network for a new connection. The simulation results show that the ALR boundary is determined by a number of different factors including: the overall level of traffic on the network, the mix of currently accepted connections at a station, the nature of the media access protocol, and the definition of each traffic class and its associated QOS. From the work described in chapters 5 and 6, it is also reasonable to believe that the overall mix of traffic connections, and the addressing of cells being transmitted by the slotted ring would affect the ALR boundary. Because of the uncertainty of the ALR boundary, and hence of the stored information relating to call acceptance or call blocking, other methods of controlling access to the network should also be considered.

Equivalent bandwidth, and equivalent capacity as a means of bandwidth reservation are very promising concepts for admission control in ATM networks. The allocation of equivalent bandwidth values rely on traffic characteristics which can be declared by the user and policed by the network operator. The determination of a system of parameters to be used to assess equivalent bandwidth requirements of a new connection, and the capacity of a network station to accept such a connection, form a current area of ATM network research, and an area of further work for this project.

Chapter 8

Conclusions and Recommendations for Further Work

8.1 Summary and Conclusions

The development of the Broadband-ISDN network is driven forward by two factors, technological capability and customer demand. Technology in the area of optical fibre transmission can provide virtually unlimited bandwidth, and VLSI switching can achieve data rates of giga-bits per second. Telecommunications network operators, which had previously enjoyed protected monopoly positions in many countries are now being forced into competition not only with each other, but with new providers of telecommunications services such as the cable television companies. The demand from customers for integrated telecommunications services including high bandwidth real-time services will ensure continued development of B-ISDN and its characteristic asynchronous transmission mode.

Since the start of this project in 1988, there have been many developments in the field of Broadband-ISDN. CCITT has introduced recommendations which standardise the ATM cell and other aspects of the transmission protocol. The aims of this project were to develop a hardware testbed with slotted ring configuration, which could be used to investigate the behaviour of the ORWELL protocol under a variety of traffic flows, and traffic mixes. The testbed could then be used to study problems associated with congestion management in an ATM network.

The testbed was developed as a platform for the performance evaluation of ORWELL as a protocol, and for the study of aspects of ATM network congestion. A modified and simplified form of the ORWELL protocol was used, based on the preliminary ORWELL specification, however the features specific to ORWELL such as the reset mechanism were preserved. Hardware and software were developed in a modular fashion so that the algorithms controlling the media access protocol, traffic generation and traffic analysis, were flexible and easily altered. Hardware was provided on the ring interface to monitor the ORWELL reset rate, and the output buffer occupancy of each station on the slotted ring.

The use of transputers in the design of the network stations enabled the processing power of a 32-bit floating point processor (T800) to be coupled to each station. The transputers were operated under the control of the Transputer Development System which enabled easy access to the code for development. As the project progressed, and the generation of data cells by different traffic generation models was attempted, the processing power of the T800 was a limiting factor in the rate at which cells could be generated.

Traffic modelling of voice, video, and data sources was studied, but the processing power of the testbed allowed a limited complexity of the cell generation algorithm. However, the testbed was able to show clearly how characteristics such as the ORWELL reset rate are affected by the cell generation algorithm in terms of the burstiness of the cell-based traffic, as well as its intensity.

The ORWELL protocol was analysed to give information as to the maximum ring throughput, and the reset-rate under different traffic conditions. The analytical results were compared to results gained through the testbed, and good agreement was found. Both the analytical results and the simulations show that the reset rate and throughput of the slotted ring depend on the symmetry of traffic flow on the network. The maximum throughput in cells per second can vary by a factor of 2:1 from an environment where all stations are addressing all other stations randomly, to one where stations are depriving other stations of bandwidth, when an asymmetrical traffic flow is created. Further simulations were used to evaluate the performance of the testbed under differing cell generation algorithms and showed that the reset rate was dependent on traffic intensity, and also on the statistical distribution of arriving cells. In papers published on the ORWELL protocol it has been suggested that the reset rate could be used to indicate impending network congestion, however such work is based on symmetrical traffic flows on the ORWELL ring, and deterministic cell generation as found in telephony. The results from the testbed show that the reset rate can vary widely if traffic flow is asymmetric, or if the statistical distribution of arriving cells changes.

The overhead of the ORWELL reset mechanism was shown to be dependent on the Di-allocation given to the network stations, so that when low Di-allocations were used the ring throughput was limited as a high proportion of slots were designated as TRIAL and RESET slots, and therefore not able to carry data cells. A protocol using only full and empty cells was developed to replace the ORWELL protocol. Simulations using the new Full/Empty protocol showed improved performance in terms of throughput and delay over the ORWELL based protocol.

The studies of ORWELL, as a media access protocol, show that access of cells to the ring from the transmitting stations can be apportioned by using Di-allocation only when the station output buffers are kept full, and in this state, most resets occur because stations are forced to their paused state. Under these conditions, because the station output buffers are fully occupied, cell loss due to buffer overflow is likely to occur. In a connection oriented network where access control is used to ensure cell loss does not occur, output buffers are not allowed to overflow, and most resets are initiated from idle stations rather than paused stations. In these circumstances the ORWELL protocol offers no advantage over the simpler Full/Empty protocol for a slotted ring.

The wider aspects of ATM network management were considered in chapter 7. Congestion management consists of a number of strategies to ensure that the quality of service in terms of cell delay or cell loss is maintained for all users granted access to the network. Connection admission control is the foremost method of ensuring that a new connection does not prejudice the quality of service of existing connections, and relies on the reservation of sufficient bandwidth on each point-to-point connection across which the call is routed, for the duration of the call. Bandwidth reservation is complicated in a heterogeneous traffic environment because of the differing requirements of each call, and the differing cell generation characteristics. To simplify the problem, two approaches have been used, traffic classification into grades of service and the defining of an ALR, and the concept of equivalent bandwidth and equivalent capacity.

The division of traffic into grades of service each of which has different cell loss and cell delay requirements means that an admissible load region can be defined for a network path, or number of paths. The admissible load region is a stored chart of information of the numbers of concurrent calls of each grade of service that can be accepted whilst maintaining the quality of service to all calls. Two grades of service were defined to show how an admissible load region could be calculated for the testbed. The effect of prioritisation of one traffic type over another was shown, as was the effect on the admissible load region of one station of other traffic on the network.

The second method of bandwidth reservation is to attribute an equivalent capacity to each network link, and an equivalent effective bandwidth to each accepted connection. The equivalent bandwidth would take into consideration the intensity, peak rate, and burstiness of the traffic source. The sum of the equivalent bandwidth of each accepted call would not be allowed to exceed the equivalent capacity of the output link.

The main original contributions to knowledge described within this thesis are:

- The development of a Transputer-based slotted-ring testbed with a modular architecture utilising programmable logic to implement a modified ORWELL protocol. The parallel processing and real-time programming capabilities of OCCAM were used to obtain results from the testbed.

- The development of equations to describe the maximum reset interval and maximum throughput of the ORWELL protocol for three traffic flow configurations. An analysis of the behaviour of ORWELL was carried out under asymmetrical traffic distributions, and this was validated with results from the testbed. Such an analysis of the ORWELL protocol with asymmetrical traffic has not previously been published. The details of the analysis were presented as a paper at an international conference (see appendix A).

- An analysis of the performance of the modified ORWELL protocol under various traffic flows, and cell arrival patterns, to identify the effect of these factors on the reset rate as measured on the testbed, and to determine the suitability of the reset rate as an indicator of ATM network congestion.

8.2 Recommendations for Further Work

For admission control to be implemented in an ATM network, the potential user must declare appropriate parameters to the network during the call set up phase of a connection. These parameters would probably include mean cell arrival rate, peak cell arrival rate, and some kind of measure of burst length. Every accepted connection to the network must be policed by the network operator to ensure that the quality of service offered to all users is not jeopardised by excessive, and unexpected, cell input from one source. If the user declared parameters have a direct relevance to the policing mechanism, the way in which the cell arrival pattern will be shaped by the controlling function will be known by both user and operator, which is highly desirable.

Admission control using the concept of equivalent bandwidth and equivalent link capacity, relies on the estimation of cell delay distribution and cell loss probability from a knowledge of the cell arrival statistics at a multiplexing point. The worst case arrival pattern must be used in this estimation if a guaranteed quality of service is to be provided to all users.

In recommending further work, a more detailed study into policing mechanisms could be undertaken to select a system where the declared parameters are both easy for the user to define, and useful in estimating the delay distribution at a ATM multiplexing point. Analytical methods of estimating equivalent bandwidth of a connection based on the user defined parameters could be investigated.

The testbed used for this project is somewhat limited in processing power, especially for complex cell arrival algorithms. To improve its performance for larger scale testing of an ATM local area network, it would be advisable to separate transmit and receive processes at each station, allocating a powerful 32-bit processor for each function. The ring interface circuitry may be improved by use of FPGA devices for greater bandwidth, and easily reconfigurable hardware.

The testing of an admission control strategy using equivalent bandwidth allocation would be particularly interesting on a slotted ring, which is in essence a series of point-to-point ATM links where the availability of bandwidth at any station is determined by the cell arrivals at other stations.

It is suggested that although an upgrading of the processing power of the testbed would enable more qualitative studies of the congestion management strategies undertaken in chapter 7, the proposed study of effective bandwidth calculation and its relation to source policing mechanisms would be better carried out using software simulation to give flexibility of configuration, or a new hardware testbed dedicated to this function, to give faster and more detailed simulation.

References

[1] "Integrated Broadband Networks, Handel, R., Huber, M.N.," Addison Wesley, 1991.

[2] "Broadband Services and Applications, Goode, C.W.B.," Electrical Communication (ALCATEL), Vol.64, Number 2/3, 1990, pp.124 - 131.

[3] "Broadband Applications and Services Prospectus," DeMaio, M., Kafka, H.J., AT&T Technical Journal, November/ December 1993, pp.23-28.

[4] "Broadband Multimedia Applications Using ATM Networks," Armbruster, H., Wimmer, K., IEEE Journal on Selected Areas in Communications, Vol.10, No.9, December 1992, pp.1382-1396.

[5] "Strategic Impact of Broadband Telecommunications in Insurance, Publishing, and Health Care," Wright, D.J., IEEE Journal on Selected Areas in Communications, Vol.10, No.9, December 1992 pp.1369-1381.

[6] "Broadband ISDN Switching Capabilities from a Services Perspective," Spears, D.r., IEEE Journal on Selected Areas in Communications, Vol.5, No.8, October 1987 pp.1222-1230.

[7] "Physical Design Issues for Very Large ATM Switching Systems," Banwell, T.C. et Al., IEEE Journal on Selected Areas in Communications, Vol.9, No.8, October 1991 pp.1227-1238.

[8] "An ATM Switch Architecture for Provision of Integrated Broadband Services," Pattavina, A., IEEE Journal on Selected Areas in Communications, Vol.9, No.9, December 1991 pp.1537-1558.

[9] "Technology for Broadband Switching," Barri, P., Dierckx, R., Goubert, J., Eilenberger, G., Electrical Communication (ALCATEL), Vol.64, Number 2/3, 1990, pp.166-176.

[10] "High-Capacity Lightwave Local Area Networks," Henry, P.S., IEEE Communications Magazine, Vol 27, No.10, October 1989, pp.20-26.

[11] CCITT: COM XVIII- 228-E, Geneva, March, 1984.

[12] "Principles of telecommunication-traffic engineering," Bear, D., Peter Peregrinus, 1978 (2nd Ed.)

[13] "CCITT: Recommendation I.121. Broadband Aspects of ISDN," CCITT, Geneva, 1991.

[14] "Principles and Standards for Broadband ISDN," Prasanna, P.K., Levy, R., Swenson, J., AT&T Technical Journal, November/ December 1993, pp.9-14.

[15] "Trends in U.S. Broadband Fiber optic Transmission Systems," Cohen, L.G., IEEE Journal on Selected Areas in Communications, Vol.SAC-4, No.4, July 1986 pp.488-497.

[16] "Fiber-Optic Transmission Systems - Status and Trends in Japan," Kimura, T., IEEE Journal on Selected Areas in Communications, Vol.SAC-4, No.4, July 1986 pp.498-505.

[17] "ATM Access Through a Passive Optical Network," Ballance, J.W., Rogers, P.H., Halls, M.F., Electronic Letters, Vol.26 No.9, 1990.

[18] "The Provision of Telephony over Passive Optical Networks," Hoppitt, C.E., Clarke, D.E.A., British Telecom Technology Journal, Vol.7 No.2 pp.100-114, 1989.

[19] "Design and Technology aspects of VLSIs for ATM switches," Banniza, T.R., Eilenberger, G.J., Pauwels, B., Therasse, Y., IEEE Journal on Selected Areas in Communications, Vol.9, No.8, November 1991 pp.1255-1265.

[20] "CMOS VLSI Applications in Broadband Circuit Switching," Hayward, G., Gottlieb, A.M., Jain, S.K., Mahoney, D.D., IEEE Journal on Selected Areas in Communications, Vol.5, No.8, October 1987 pp.1231-1241.

[21] "Sunshine: A High-Performance Self-Routing Broadband Packet Switch Architecture," Giacopelli, J.N., Hickey, J.J., Marcus, W.S., Sincoskie, W.D., Littewood, M., IEEE Journal on Selected Areas in Communications, Vol.9, No.8, October 1991 pp.1289-1298.

[22] "Practical Implementation and Packaging Technologies for a Large-Scale ATM Switching System," Itoh, A., Takahashi, W., Nagano, H., Kurisaka, M., Iwasaki, S., IEEE Journal on Selected Areas in Communications, Vol.9, No.8, October 1991 pp.1280-1288.

[23] "CCITT: Recommendation X.25. 'Interface between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE) for Terminals operating in the Packet Mode and Connected to Public Data Networks by Dedicated Circuit'.," Blue Book, Fascile VIII.2, Geneva, 1989.,

[24] "CCITT: Recommendation I.150. 'B-ISDN ATM Functional Characteristics'," CCITT, Geneva, 1991.,

[25] "Broadband System Switching Architecture," Hiraide, K., Nakajima, S., Yasui, T., IEEE Journal on Selected Areas in Communications, Vol.5, No.8, October 1987 pp.1346-1355.

[26] "Architecture, Performance, and Implementation of the Tandem Banyan Fast Packet Switch," Tobagi, F.A., Kwok, T., Chiussi, F.M., IEEE Journal on Selected Areas in Communications, Vol.9, No.8, October 1991 pp.1173-1193.

[27] "ATM Switch Architectures with Input-Output-Buffering," Badran, H.F., Mouftah, H.T., Computer Networks and ISDN Systems Vol.26, 1994, pp.1187-1213.

[28] "Switching in the Broad-Band ISDN," Schaffer, B., IEEE Journal on Selected Areas in Communications, Vol.SAC-4, No.4, July 1986 pp.536-541.

[29] "Input versus Output Queueing on a Space-division Packet Switch," Karol, M.J., Hluchyi, M.G., Morgan, S.P., IEEE Transactions on Communications, Vol.35 No.12, December 1987, pp.1347-1356.

[30] "ATM Switches - Basic Architectures and their Performance," Rathgeb, E.P., THeimer, T.H., Huber M.N., International Journal of Digital and Analog Cabled Systems, Vol.2, no,4, October 1989, pp.227-236.

[31] "Self Routing Banyan Networks in an ATM-Environment," Huber M.N., Rathgeb, E.P., THeimer, T.H., Proceedings of the International Conference on Computer Communication, Tel Aviv, 1988, pp.167-174

[32] "High-Speed Digital Switching Technology Using Space-Division-Switch LSI's," Yamada, H., Kataoka, H., Sampei, T., Yano, T., IEEE Journal on Selected Areas in Communications, Vol.SAC-4, No.4, July 1986 pp.529-535.

[33] "Priority Queueing Strategies and Buffer Allocation Protocols for Traffic Control at an ATM Integrated Broadband Switching System, Lin, A.Y.-M., Silvester, J.A., IEEE Journal on Selected Areas in Communications, Vol.9, No.9, December 1991 pp.1524-1536.

[34] "Priority Assignment Control of ATM Line Buffers with Multiple QOS Classes," Takagi, Y., Hino, S., Takahashi, T., IEEE Journal on Selected Areas in Communications, Vol.9, No.7, September 1991 pp.1078-1091.

[35] "PRELUDE: An Asynchronous Time-division Switched Network," Coudreuse, J.P., Servel, M., Proceedings of the International Conference on Communications, Seattle, 1987, paper 22.2

[36] "An Experimental ATM Switching Architecture for the Evolving B-ISDN Scenario," Garetti, E., Melen, R., Arnold, A., Gallassi, G., Scozzari, G., Fox, A.L., Fundneider, O., Goldner, E.H., Proceedings of the XIII International Switching Symposium, Stockholm. 1990, Vol.IV, pp.15-22.

[37] "Switching Block Studies, Performance Evaluation and Traffic Engineering for ATM.," Domingo, J., Michiel, H., et al., European

Transactions on Telecommunications and Related Technologies, Vol.5, No.2, 1994, pp.59-69.

[38] "Specification of the ORWELL protocol," Chauhan, J., King, T.J., Micalef, A.C., British Telecom Technical Publications, May 1990.

[39] "Equivalent Capacity and its Application to Bandwidth Allocation in High-Speed Networks," Guerin, R., Ahmandi, H., Naghshineh, M., IEEE Journal on Selected Areas in Communications, Vol.9, No.7, September 1991 pp.968-981.

[40] "Real-Time Scheduling with Quality of Service Constraints," Hyman, J.M., Lazar, A.A., Pacifici, G., IEEE Journal on Selected Areas in Communications, Vol.9, No.7, September 1991 pp.1052-1061.

[41] "Multimedia Traffic Management Principles for Guaranteed ATM Network Performance," Woodruff G.M., Kositpaiboon, R., IEEE Journal on Selected Areas in Communications, Vol.8, No.3, April 1990 pp.437-446.

[42] "Real-Time Scheduling with Quality of Service Constraints, Hyman, J.M., Lazar, A.A., Pacifici, G., IEEE Journal on Selected Areas in Communications, Vol.9, No.7, September 1991 pp.1052-1061.

[43] "A Separation Principle Between Scheduling and Admission Control for Broadband Switching," Hyman, J.M., Lazar, A.A., Pacifici, G., IEEE Journal on Selected Areas in Communications, Vol.11, No.4, May 1993 pp.605-616.

[44] "Local Allocation of End-to-End Quality-of-Service in High-Speed Networks," Nagaragan, R., Kurose, J., Towsley, D., IFIP Transactions C-Communication Systems 1993 Vol.15,

[45] "Local Area Network Traffic Characteristics, with Implications for Broadband Network Congestion Management," Fowler, H.J., Leland,

W.E., IEEE Journal on Selected Areas in Communications, Vol.9, No.7, September 1991 pp.1139-1149.

[46] "Characterisation of Video Codecs as Autoregressive Moving Average Processes and Related Queueing System Performance," Grunenfelder, R., Cosmas, J.P., Manthorpe, S., Odinma-Okafor, A., IEEE Journal on Selected Areas in Communications, Vol.9, No.3, April 1991 pp.284-293.

[47] "Models for Packet Switching of Variable-Bit-Rate Video Sources," Sen, P., Maglaris, B., Rikli, N.-E., Anastassiou, D., IEEE Journal on Selected Areas in Communications, Vol.7, No.5, June 1989 pp.865-869.

[48] "Modeling Packet Arrivals from Asynchronous Input Lines," Ramaswami, V., Latouche, G., 12th International Telegraffic Congress, Torino, 1989.

[49] "A Markov Modulated Characterization of Packetised Voice and Data Traffic and Related Statistical Multiplexer Performance," Heffes, H., Lucantoni, D., IEEE Journal on Selected Areas in Communications, Vol.SAC-4, No.6, September 1986, pp.856-868.

[50] "Approximate Analysis of an ATM Switching System with Bursty Arrivals and Finite Capacity," Jou, Y.F., Nilsson, A.A., Lai, F., IFIP Transactions C-Communication Systems 1993 Vol.15,

[51] "A Congestion Control Framework for ATM Networks," Gersht, A., Kyoo, J.L., IEEE Journal on Selected Areas in Communications, Vol.9, No.7, September 1991 pp.1119-1129.

[52] "A Multidimentional Framework for Congestion Control in B-ISDN," Ramamurthy, G., Dighe, R.S., IEEE Journal on Selected Areas in Communications, Vol.9, No.9, December 1991 pp.1440-1451.

[53] "Congestion-free Communication in High-speed Packet Networks," Golestani, S.J., IEEE Transactions on Communications Vol.39 December 1991.

[54] "Dynamic Call Admission Control in ATM Networks," Saito, H., Shiomoto, K., IEEE Journal on Selected Areas in Communications, Vol.9, No.7, September 1991 pp.982-989.

[55] "Integration of ATM Call Admission Control and Link Capacity by Distributed Neural Networks," Hiramatsu, A., IEEE Journal on Selected Areas in Communications, Vol.9, No.7, September 1991 pp.1131-1138.

[56] "A Call Admission Control Scheme for ATM Networks Using a Simple Quality Estimate," Murase, T., Suzuki, H., Sato, S., Takeuchi, T., IEEE Journal on Selected Areas in Communications, Vol.9, No.9, December 1991 pp.1461-1470.

[57] "The Policing Function in an ATM Network," Dennisen, F., Desmet, E., Petit, G.H., Proc.1990 Int. Zurich Sem. Digitital Communications, Zurich, Switzerland, March 1990, pp.131-144.

[58] "The Policing Function in ATM Networks," Rathgeb, E.R., Theimer, T.H., XIII International Switching Symposium, Stockholm, Sweden, May 1990, paper A8.4.

[59] "Effectiveness of the "Leaky Bucket" Policing Mechanism in ATM Networks," Butto, M., Cavallero, E., Tonietti, A., IEEE Journal on Selected Areas in Communications, Vol.9, No.3, April 1991 pp.335-342.

[60] "Distributed Source Control: A Network Access Control for Integrated Broadband Packet Networks," Ramamurthy, G., Dighe, R.S., IEEE Journal on Selected Areas in Communications, Vol.9, No.7, September 1991 pp.990-1001.

[61] "Modelling and Performance Comparison of Policing Mechanisms for ATM Networks," Rathgeb, E.P., IEEE Journal on Selected Areas in Communications, Vol.9, No.3, April 1991 pp.325-334.

[62] CCITT Recommendation G.707, "Synchronous Digital Hierarchy Bit Rates", CCITT, Blue Book, Fascicle III.4, Geneva, 1989.

[63] CCITT Recommendation G.708, "Network Node Interface for the Synchronous Digital Hierarchy", CCITT, Blue Book, Fascicle III.4, Geneva, 1989.

[64] CCITT Recommendation G.709, "Synchronous Multiplexing Structure", CCITT, Blue Book, Fascicle III.4, Geneva, 1989.

[65] "B-ISDN Architecture and Protocol," Kawarasaki, M., Jabbari, B., IEEE Journal on Selected Areas in Communications, Vol.9, No.9, December 1991 pp.1405-1415.

[66] "Network Aspects of Broadband ISDN," Soneru, M.D., Anderson, J., Kafka, H.J., AT&T Technical Journal, November/ December 1993, pp.29-38.

[67] "New Standards in Relation to the Broadband ISDN," Engel, A., Mobasser, B., Electrical Communication (ALCATEL), Vol.64, Number 2/3, 1990, pp.177-185.

[68] "Evolution Towards Broadband," Debuysscher, P., Glon, J-P., Langenbach-Belz, M., Electrical Communication (ALCATEL), Vol.64, Number 2/3, 1990, pp.260-268.

[69] "Alcatel's View on Broadband Deployment," Fontaine, B., Van Malderen, R., Dunogue, J., Electrical Communication (ALCATEL), Vol.64, Number 2/3, 1990, pp.276-281.

[70] "Traffic Analysis for a Custom Video Architecture," Mouchtaris, P.N., IEEE Journal on Selected Areas in Communications, Vol.10, No.9, December 1992 pp.1421-1426.

[71] "Evolution Toward Broadband," Kafka, H.J., AT&T Technical Journal, November/ December 1993, pp.15-22.

[72] "BERKOM Field Trial: Operational Results and Services," Domann, G., Electrical Communication (ALCATEL), Vol.64, Number 2/3, 1990, pp.269-275.

[73] "Broadband National Experiments in Belgium, France, and Spain," De Prycker, M., Paul, J.L., Campos, A., Electrical Communication (ALCATEL), Vol.64, Number 2/3, 1990, pp.218-228.

[74] "The Asynchronous Transfer Mode technology Platform - Initial Laboratory Trials," King, T.J., Rowley, R., Whalley, G.M., Clementson, J.M., Chauhan, J., British Telecom Technology Journal Vol.9 No.2 April 1991, pp.71-83.

[75] "Asynchronous Transfer Mode (ATM) - Practical Management Issues," King, T.J., Barker, T.J.A., Micallef, A.C., Mackay, K.W., British Telecom Technology Journal Vol.9 No.2 April 1991, pp.57-70.

[76] "MAGNET: Columbia's Integrated Network Testbed," Lazar, A.A., Patir, A., Takahashi, T., El Zarki, M., IEEE Journal on Selected Areas in Communications, Vol.SAC-3, No.6, September 1985, pp.859-871.

[77] "Real-Time Traffic Measurements on MAGNET II," Lazar, A.A., Pacifici, G., White, J.S., IEEE Journal on Selected Areas in Communications, Vol.8, No.3, April 1990 pp.467-483.

[78] "ORWELL: A Protocol for an Integrated Services Local Network," Falconer, R.M., Adams, J.L., British Telecom Technology Journal Vol.3 No.4 October 1985, pp.27-35.

[79] "ATMR: Asynchronous Transfer Mode Ring Protocol," Imai, K., Ito, T., Kasahara, H., Morita, N., Computer Networks and ISDN Systems Vol.26, 1994, pp.785-798.

[80] "A Reference Model for Integrated Local Area Networks," Lazar, A.A., Mays, M.A., Hori, K., Proceedings of the International Conference on Communications (ICC'86), Toronto, CANADA, June 1986, pp.531-536.

[81] "An Optical Fiber-Based Integrated LAN for MAGNET's Testbed Environment," Patir, A., Takahashi, T., El Zarki, M., Lazar, A.A., IEEE Journal on Selected Areas in Communications, Vol.SAC-3, No.6, September 1985, pp.872-881.

[82] "Packetised Video on MAGNET," Lazar, A.A., White, J.S., Proceedings of the SPIE 86, Boston, MA, September 1986.

[83] "ORWELL: A Protocol for carrying Integrated Services on a Digital Communications Ring," Adams, J.L., Falconer R.M., Electronic Letters 20 (23) (1984), pp.970-971.

[84] "Computer Networks", Tannembaum, A.S., Prentice Hall, 1988.

[85] "Performance Analysis of Slotted Ring Protocols in HSLAN's," Zafirovic - Vukotic, M., Niemegeers, I.G., Valk, D.S. , IEEE Journal on Selected Areas in Communications, Vol.6, No.6, July 1988.

[86] "A simulation Study of the Cambridge Ring with Voice Traffic," Falconer, R.M., Adams, J. L., Whalley, G.M. , Br Telecom Technol J Vol.3 No.2 April 1985

[87] "Integrated Services using a Baseband Local Network," Needham, R., Ades, S., IEE Conference Publication No.278, June 1987.

[88] "The Application of ATM techniques to the Local Network," Gallagher, I., Ballance, J., Adams, J., British Telecom Technology Journal Vol.7 No.2 April 1989.

[89] "The FDDI Optical Data Link," Burr, W.E., IEEE Communication Magazine, Vol.25, May 1986, pp.18-23

[90] "High Performance Networks - a Focus on the Fiber Distributed Data Interface (FDDI) Standard," Joshi, S.P., IEEE Micro Magazine, Vol6. June 1986, pp8-14.

[91] "A Performance Analysis of Integrated Communications Systems," Johnson, S.A., British Telecom Technology Journal Vol.3 No.4 October 1985, pp.36-45.

[92] "Ethernet: Distributed Packet Switching for Local Computer Networks," Metcalf, R.M., Boggs, D.R., Commun.. ACM, Vol.19 pp.395-404, July 1976.

[93] "Performance Analysis of Local Computer Networks," Hammond, J.L., O'Reilly, P.J.P., Addison Wesley, March 1988.

[94] "The Cambridge Digital Communications Ring," Wilkes, M.W., Wheeler, D.J., Proceedings of the Local Area Communications Network Symposium, Boston, MA. USA (1979).

[95] "The Transputer Data Book," INMOS, INMOS 1989

[96] "Transputer Development System," INMOS, Prentice Hall International 1988.

[97] "TANGO PLD Reference Manual," ACCEL Technical Assistance, May 1990.

[98] "ABEL Reference Manual," DATA I/O Corporation

[99] "Departmental Technical Document for Slotted Ring Test Bed," Holifield, P., Department of Electrical and Electronic Engineering, University of Central Lancashire, Preston, PR1 2HE, 1994.

[100] "TURBO PASCAL reference manual," Borland.

[101] "Transtech TMB08 User Manual," Transtech, England, 1990.

[102] "A Tutorial Introduction to OCCAM Programming," Pountain, D., May, D., BSP Professional Books, London, 1987

[103] "A Review of Voice, Data, and Video Traffic Models for ATM," Cosmas, J.P., Petit, G.H., Lehnert, R., et Al., European Transactions on Telecommunications and Related Technologies, Vol.5, No.2, 1994.

[104] "A Traffic Measurement Method and its Application for Cell Loss Probability Estimation in ATM Networks," Yamada, H., Sumita, S., IEEE Journal on Selected Areas in Communications, Vol.9, No.3, April 1991 pp.315-323.

[105] "Variable Rate Coding of Speech," Dubnowski, J., Crochiere, R., Bell Sys. Tech J., Vol.58, 1979, pp577-600.

[106] "Variable Rate Speech Coding for Asynchronous Transfer Mode," Nakada, H., Sato, K., IEEE COMM-38 No.3 March 1990, pp277-284

[107] "Variable Rate Speech Coding and Delay Analysis for Universal Transport Network," Sato, K., Nakada, H., Sato, Y., Proc. IEEE INFOCOM '88, New Orleans, USA, March 1988, pp. 1-10.

[108] "Universal ATM Video Coding Architecture," D'Agostino, S., David, J., Heiss, R., Verbiest, W., Electrical Communication (ALCATEL), Vol.64, Number 2/3, 1990, pp.186-193.

[109] "Variable Bit-Rate coding of Video Signals for ATM Networks," Kishino, F., Manabe, K., Hayashi, Y., Yasuda, H., IEEE Journal on Selected Areas in Communications, Vol.7, No.5, June 1989 pp.801-806.

[110] "New Video Coding Standard for the 1990's," Carr, M.D., Electronics and Communication Engineering Journal, June 1990, pp.119-124.

[111] "Statistics of Video Signals for Viewphone-type Pictures," Chin, H., Googe, J., Griffiths, R., Parish, D., IEEE Journal on Selected Areas in Communications, Vol.7, No.5, June 1989 pp.826-832.

[112] "Scene Adaptive Coder," Cein, W-H., Pratt, W.K., IEEE Trans. on Comm. Vol. COM-32, No.3, March 1984, p225-232.

[113] "Broadband ISDN: An Overview," Bosco, H.L., Stefanik, M.J., AT&T Technical Journal, November/ December 1993, pp.4-8.

[114] "Measured Performance of an Ethernet Local Network," Shoch, J.F., Hupp, J.A., Commun. ACM, Vol.23, No.12, December 1980, pp711-721.

[115] "Service Characteristics and Traffic Models in a Broadband ISDN," Bermejo, L., Parmentier, P., Petit, G.H., Electrical Communication (ALCATEL), Vol.64, Number 2/3, 1990, pp.132-138.

[116] "The Art of Computer Systems Performance Analysis," Jain, R., John Wiley & Sons 1991.

[117] "Queueing Systems Volume 1: Theory," Kleinrock, L., John Wiley & Sons 1975.

[118] "Adaptive ATM Call Access Control Using Learning Algorithms," Chen, J.R., Mars, P., School of Engineering and Applied Science, University of Durham, 1991.

[119] "Basic Characteristics of Variable Bit Rate Video Coding in ATM Environment," Nomura, M., Fujii, T., Ohta, N., IEEE Journal on Selected Areas in Communications, Vol.7, No.5, June 1989 pp.752-760.

[120] "Orwell," Adams, J.L., Computer Networks and ISDN Systems Vol.26, 1994, pp.771-784.

[121] "A Modelling Study of the ORWELL Ring Protocol," Mitrani, I., Adams, J.L., Falconer, R.M., Teletraffic Analysis and Computer Performance Evaluation, Elsevier Science Publishers B.V. (North-Holland), 1986, pp.429-438.

[122] "Statistical Multiplexing, Bandwidth Allocation Strategies and Connection Admission Control in ATM Networks," Mitrou, N.M., Kontovasilis, K.P., Kroner, H., Iversen, V.B., European Transactions on Telecommunications and Related Technologies, Vol.5, No.2, 1994, pp.33-47.

[123] "Bandwidth Allocation Algorithms for Packet Video in ATM Networks," Chowdhury, S., Sohraby, K., Computer Networks and ISDN Systems Vol.26, 1994, pp.1215-1223.

[124] "Performance Simulation of Leaky Buckets for ATM Network Access Control," LeMercier, M.., IFIP Transactions C-Communication Systems 1993 Vol.15,

[125] "Evaluation of UPC Functions for Peak Cell Rate Enforcement," Hemmer, H., European Transactions on Telecommunications and Related Technologies, Vol.5, No.2, 1994.

[126] "Simulation Analysis of Loss Characteristics in ATM Switching Systems by R1022," Habermann, R., European Transactions on Telecommunications and Related Technologies, Vol.5, No.2, 1994, pp.49-57.

[127] "An Analysis of Statistical Multiplexing in an ATM Transport Network," Saito, H., Kawarasaki, M., Yamada, H., IEEE Journal on Selected Areas in Communications, Vol.9, No.3, April 1991 pp.359-367.

[128] "Loss Performance Analysis of an ATM Multiplever Loaded with High-Speed ON-OFF Sources," Baiocchi, A., Blefari, N., Listanti, M. et al., IEEE Journal on Selected Areas in Communications, Vol.9, No.3, April 1991 pp.388-393.

[129] "Performance Evaluation of ATM Multiplexing Schemes for Bursty Traffic Sources in Lossy ATM Networks," De Sanctis, M., Costabile, R., IFIP Transactions C-Communication Systems 1993 Vol.15,

[130] "Queueing Systems for Modelling ATM Networks," Pujolle, G., Perros, H.G., IFIP Transactions C-Communication Systems 1992 Vol.5, pp.301-323.

[131] "Packet Voice Communication on an Ethernet local computer communications systems, Gonsalves," T.A., Proc. ACM SIGCOM '83, 1983, pp.178-185.

[132] "An Integrated Services Token-Controlled Ring Network," Wong, P-C., Yum, T-S., IEEE Journal on Selected Areas in Communications, Vol.7, No.5, June 1989 pp.670-678.

[133] "OCCAM2 User Manual", INMOS, 1988.

Appendix A

Published Papers

1. *A Slotted Ring Local Area Network for the Investigation of ATM Network Congestion Management*
   The Communications Networks Symposium, Manchester Metropolitan University, Manchester, England, 11th to 12th July 1994.

2. *Modelling and Performance Evaluation of a modified ORWELL protocol under symmetrical and asymmetrical Traffic Distributions*
   The European Symposium on Advanced Networks and Services, RAI Congress and Exhibition Centre, Amsterdam, The Netherlands, 20th to 23rd March 1995.

A Slotted Ring Test Bed for the Study of ATM Network Congestion Management

Phil Holifield and Professor Trevor J. Terrell
Department of Electrical and Electronic Engineering
University of Central Lancashire
Preston, PR12HE.
and
Professor Phil Mars
Department of Engineering and Computer Science
University of Durham
Durham, DH1 3LE.

Abstract

This paper describes a slotted ring test bed which uses the ORWELL protocol [1] to enable various forms of access control mechanism to be developed and evaluated so that ATM network congestion is minimised. The slotted ring test bed architecture is discussed, and test results are presented for the ORWELL reset rate and for mean packet delay times. The basis for continuing work on adaptive access control mechanisms is outlined.

1. Introduction

Broadband Integrated Services Digital Network (B-ISDN) is an emerging communication network presently at a research and development stage. BISDN is intended to provide multimedia services, including real time and non-real time traffic.

B-ISDN makes use of an Asynchronous Transfer Mode (ATM) as opposed to the Synchronous Transfer Mode (STM) employed by present ISDN systems. In STM systems a time slot is reserved for each active call in a hierarchical frame structure. This leads to bandwidth being reserved even when there is no data to be transmitted ie. when there is silence in a voice conversation, or when a data source is idle. Under ATM [2], all traffic is segmented into equal sized units called packets or cells each of which contains a header field with address information, and a data field. Cells can be dynamically allocated to users on demand which increases flexibility over STM and allows for statistical multiplexing when many sources are involved. Aspects of ATM to be studied are the probability of cell loss and the variation in cell delay times which will affect the quality of service that can be provided. These are heavily dependent on the traffic intensity of the network, and a good indication of network loading and remaining available bandwidth are required to control access to the network.

2. Slotted Ring Test Bed

The slotted ring test bed is a multi-transputer based system (see Figure 1), which runs under the control of the Transputer Development System software. Each node has a 32-bit T800 transputer to model packet generation, and to analyse and store results. Packets generated in the T800 are received by the 16-bit T222 Transputer via an Inmos link, and placed in the hardware output

FIFO. Slots received from the network are placed in the input FIFO by the Medium Access Controller (MAC), and are read by the T222 and passed to the T800 for analysis.

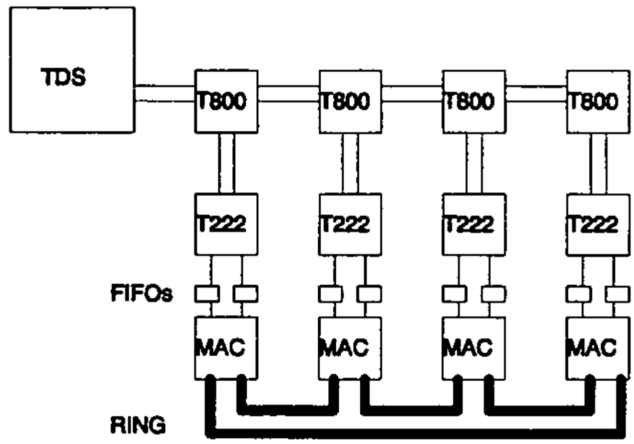

Figure 1. Architecture of Slotted Ring Test Bed

The packets used in the test bed have destination and source addresses in the 16-bit header field. The data field is occupied by a time stamp to monitor delay, a sequence to detect lost slots, and a checksum to detect corruption of data. Time stamping and analysis of delay are performed in the T800 devices which are synchronised by an external interrupt at the start of each simulation. The sequence and lost slot analysis are also performed by the T800 devices. The T222 devices act as microcontrollers to the MAC units which are implemented in hardware.

The four node network has been established with a data transmit rate of 500kb/s. There are four slots each of which contains an entire packet as shown in Figure 2.

Dest. Addr.	Source Addr.
Time Stamp	
Time Stamp	
Sequence Number	
Checksum	

Figure 2. Format of Packets

A slot completes a circuit of the ring in 744us. The mean time for a slot to travel from one node to the next is therefore 186us.

3. The ORWELL Protocol

The ORWELL protocol ensures fair access to the ring for all nodes by means of Trial and Reset slots. Each node is allocated a D-count which is the maximum number of packets it may transmit before pausing. A node which is idle or paused will issue a trial slot when it is unable to fill a passing empty slot. Any node with data ready to send and not paused can overwrite this trial, but if it circulates undisturbed to the originator, the node converts its own trial slot to a reset slot.

The reset slot circulates around the ring, resetting D-counts at all nodes, and is removed by the originator node.

The characterisation tests apply traffic to the network for a ten second period. During this time the number of slot-sized packets transmitted and received at each node is recorded, as are the ring's reset rate, the delay time of each packet, and the numbers lost due to buffer overflow.

4. Measurement of ORWELL Reset Rate, and Packet Delay

If there is little or no data applied to the ring, trial slots will be constantly in circulation and will be converted to resets very frequently. The minimum time between resets approximates to the time taken for a reset slot to circulate the ring. The outstanding reset timer will then prevent resetting for the next four slots so that the minimum time between resets is 5 slot-to-slot times ie. 5 x 186us = 930us. In reality, because of the interaction of the outstanding reset timer and the conversion of trial to reset slots at different points on the ring, the maximum reset rate has been measured at around 900 resets/second or approximately 6 slot-to-slot times (1116us) between resets.

Figure 3 shows the reset rate measured for increasing throughput with a poisson arrival process and random addressing. D-counts of 1,2,3,4,5,7,10 and 15 are shown on the same graph. From the maximum reset rate there is an almost linear reduction in reset rate for increasing throughput. The time between resets comprises of a fixed element: the reset overhead determined by the outstanding reset timer, and a variable element: the time waiting for a trial slot to circulate without being filled. As throughput increases the probability of a trial slot being filled also increases, therefore the probability that the trial slot will circulate completely and be converted to a reset slot decreases with a linear relationship to throughput. The D-count is not expected to have any impact on this because in the case of Figure 3. almost all resets are a result of idleness in the network and not of nodes being paused. The trial slots complete circulation of the ring due to natural breaks in the arriving traffic.

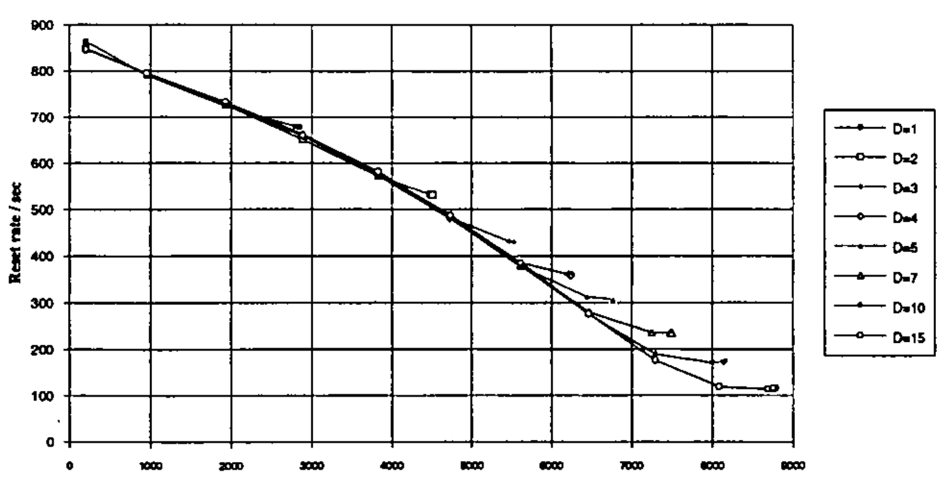

Figure 3. Reset Rate Variation with Throughput

To analyse the expected packet delay together with the probability of slots being lost due to buffer overflow it is possible to model each node of the network as an M/M/1 queue with an arrival rate equal to the throughput at that node, a service rate equal to the service rate at maximum throughput for that D-count, and a known maximum buffer length. To obtain delays comparable to those measured, 1100us must be added, which is the time required for transmission and reception of a packet when the throughput is at its minimum value and there is effectively no queueing. It is assumed that there is no variation of queueing in the receive FIFO i.e. minimum queueing at all times. Figure 4 shows how the mean packet delay varies with throughput.

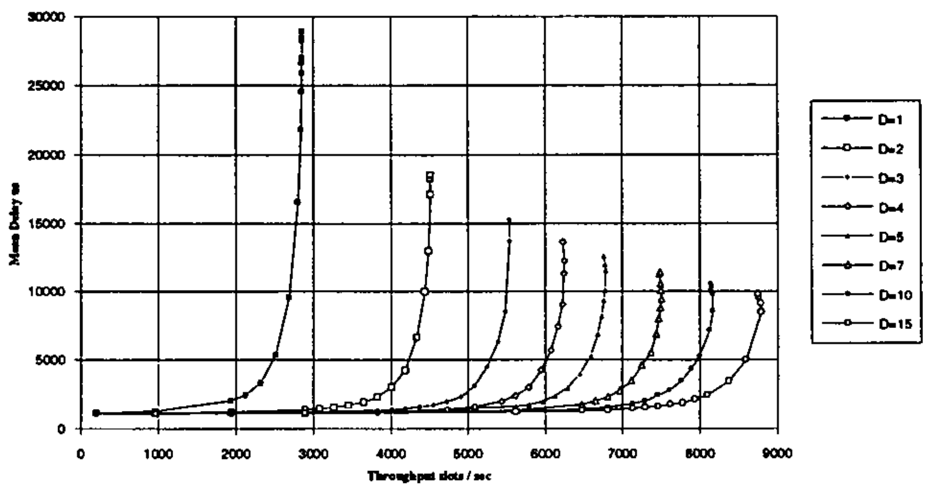

Figure 4. Mean Packet Delay Variation with Throughput

5. Requirements for Access Control

In order to minimise congestion in the ATM network [3] it is necessary to have an effective access control mechanism. The purpose of access control is to ensure that the required grade of service in terms of maximum delay and packet loss is guaranteed. To implement the access control mechanism, performance indicators [4] have to be determined and used to adaptively control access to the slotted ring. A global indicator of ring loading is the reset rate, whereas a local indicator is the node transmit buffer queue length. The test bed architecture shown in Figure 1 enables parameters measured at each node to be collected centrally and used as part of an adaptive access control algorithm. Development work is now on-going to produce a realistic traffic demand model to test various forms of access control mechanism.

6. References

[1] "Orwell: a Protocol for an Integrated Services Local Network", R.M. Falconer and J.L. Adams, *Br Telecom Technol J*, Vol. 3, No. 4, October 1985.

[2] "The Application of ATM Techniques to the Local Network", I. Gallagher, J. Balance, and J. Adams, *Br Telecom Technol J*, Vol 7, No. 2, April 1989.

[3] "Local Area Network Traffic Characteristics, with Implications for Broadband Network Congestion Management", H.J. Fowler and W.E. Leyland, IEEE Journal on Selected Areas in Communications, Vol. 9, No. 7, September 1991.

[4] *Performance Analysis of Local Computer Networks*, J.L. Hammond and P.J. O'Reilly, Addison Wesley, 1988.

Modelling and performance evaluation of a modified Orwell protocol under symmetrical and asymmetrical traffic distributions

Phil Holifield and Trevor J. Terrell

Department of Electrical and Electronic Engineering
University of Central Lancashire, Preston, PR1 2HE, UK.

Phil Mars

Department of Engineering and Computer Science
University of Durham, Durham, DH1 3LE, UK.

ABSTRACT

With the emergence of Broadband ISDN and the adoption of ATM as its transport mechanism there is a growing interest in slotted ring networks both for high speed packet switching and for multimedia LAN applications. This paper describes a slotted ring testbed which uses a modified Orwell protocol for media access, and which has enabled a study of the modified Orwell protocol under symmetrical and asymmetrical traffic flow distributions. The architecture of the slotted ring testbed is outlined, and analytical models for the protocol are presented for different traffic flow distributions. The analytical results are compared to experimental results obtained from the testbed. The suitability of the reset rate as an indicator of available bandwidth in an access control mechanism is considered, and it is shown that for constant traffic flow and traffic statistics the reset rate works well as an indicator of potential network congestion. Under changing traffic distributions and traffic arrival statistics, additional indicators of congestion are required. On-going work into access control for the test-bed in a integrated services environment is outlined.

1. INTRODUCTION

The Asynchronous Transfer Mode (ATM) has been adopted as the transport mode for Broadband ISDN. The advantages of ATM include a common cell-based user-network interface for all traffic, which lends itself to integrated services traffic and variable bit rate data sources. One of the major difficulties in the management of an ATM network is to maintain the Quality of Service (QOS) for all accepted connections in terms of cell loss and cell delay. The call acceptance or rejection procedure must be able to protect the QOS provided to existing traffic, as well as guaranteeing the QOS of new connections. To successfully implement a call acceptance mechanism, an indicator of the network loading is required, as is a method to determine whether the loading is close to the maximum permitted to maintain the defined QOS [1]. It is preferable that the indicator can be measured independently at each network station. This paper analyses a loading indicator for the ORWELL slotted-ring protocol, which is known as the reset interval, and assesses its value as an indicator of potential network congestion under varying traffic distribution.

2. THE SLOTTED RING TESTBED

Local Area Networks (LANs) for delay sensitive traffic, such as voice and video services, must be able to provide a quality of service which limits cell delay, and cell loss. The commonly used data LANs such as CSMA/CD and Token Ring have been used to transmit delay sensitive data, but both protocols have inherent problems. The CSMA/CD protocol being random access, is susceptible to cell collision, and hence data loss. It also suffers from a potentially unbounded transmission delay time as stations compete for access to the transmission medium. The Token Ring protocol is collision free but has a potential transmission delay of one token rotation period as each station must transmit in sequence, and only one station transmits at a time. The Slotted Ring consists of a number of circulating slots in a ring configuration each of which holds a packet of data. An Empty slot is filled with a data cell as it passes a station with data to transmit, and the Full slot is emptied at its destination station, and the cell is marked as empty. The slotted ring has the benefit that all stations can transmit continually, so reducing the delay for medium access, and that slots can be released at the destination, effectively increasing the ring bandwidth.

One disadvantage of the slotted ring configuration is that it is not inherently fair to all stations, and bandwidth starvation can occur if one station is downstream of a second, heavily transmitting, station. A protocol designed to introduce fairness to the slotted ring is the Orwell protocol [2]. This protocol employs a counter at each station on the network known as the D-counter. Each time the station fills a slot the D-counter is incremented until a value known as the Di-value is reached. The station is then Paused and cannot transmit more cells until it is Reset. Stations which are Paused or Idle, change Empty slots to Trial slots. A Trial slot may be filled by a station which is not paused, but if the Trial slot circulates the ring and returns to its originating station, it is converted to a Reset slot which circulates the ring resetting all D-counters to zero. In this way all stations may transmit up to Di cells in each reset cycle. The interval between reset cycles (Reset Interval) is dependent on network utilization, and has been proposed as a indicator of network loading [3].

The slotted ring testbed developed by the authors, uses a modified ORWELL protocol and has a data rate of 500 kbit/sec, with four slots and four stations. The packet size is 80 bits, consisting of a 16-bit header with destination and source addresses, and a 64-bit data block which includes a time-stamp, sequence number and checksum. Stations may re-use received slots. Data is generated and analysed by a T800 transputer at each station allowing individual cell delay times to be calculated, as well as numbers of cells transmitted, received, and lost. The Reset Interval is monitored independently by each station.

3. ANALYSIS OF THE ORWELL RESET INTERVAL FOR AN IDLE RING

Some terms used in the analysis will first be defined.
R = Ring rate (bits/sec), is the rate of transmission of data bits on the ring (500kb/s).
N = Number of stations on the ring (4).
S = Number of slots on the ring (4).

Ts = Slot duration time, is the time for successive slots to pass a point on the ring (185us).
Tr = Ring Latency, is the time for a slot to circulate the ring (740us).
RI = Reset Interval, is the time between resets.

All stations when idle attempt to convert EMPTY slots to TRIAL slots, and if they receive their own TRIAL slot, will convert it to a RESET slot. When a station Di allocation is reset by receiving a RESET slot, its Outstanding Reset Timer (ORT) will be activated and it will not be allowed to convert its received TRIAL slots to RESET slots during this period. The minimum reset interval is thus the period of the ORT, which in the ORWELL protocol is specified as the Ring Latency, Tr.

When a station's D-counter has been reset, it cannot be reset again within the ORT period which is Tr. The interval until this station is next reset depends upon which station on the ring originates the next reset, and how long the reset slot takes to arrive. On average if the next reset is equally likely to occur at any station, this time will be $Tr/2$, so the mean reset interval will be

$$RI_{(idle)} = Tr + Tr/2 \qquad (1)$$

4. ANALYSIS OF THE ORWELL RESET MECHANISM FOR THREE DIFFERENT TRAFFIC FLOWS

The ORWELL Reset Interval has been analysed for three different traffic flow configurations with stations transmitting at maximum intensity, that is to say the transmit buffer of an active station is always full and there are always cells ready to be transmitted if an empty slot is available. In order to maintain a full output buffer at each station there is a high loss rate of cells due to buffer overflow, but the analysis demonstrates how in the limit of operation, the traffic flow affects the reset interval of the ring, and the maximum throughput that can be achieved.

The following three traffic flows are considered:
1. A single transmitting station (single).
2. All four stations transmitting randomly to each other (symmetrical).
3. Three stations all transmitting to the fourth which is idle (asymmetrical).

These traffic flows have been chosen to demonstrate the variation in reset interval and throughput when the ring throughput is saturated.

The ORWELL Reset Interval is determined by the sum of the time required for each station to transmit its D_i allocation, and the time required for the reset mechanism to zero all D-counters. The throughput of cells transmitted by the ring from source station to destination station in each reset interval is equal to the sum of the D_i allocations of all active stations, since the fairness of the ORWELL protocol ensures all stations are allowed to transmit their full D_i allocation. A full description of the reset mechanism for each of the selcted traffic flows is given in [4].

The ring rate R can be expressed as a transmission rate of cells per unit time, R_{cell}, which, since each slot holds one cell, is calculated as

$$R_{cell} = S / Tr \tag{2}$$

4.1. Single transmitting station

The time that the single station is paused is given by

$$T_{paused} = Tr + Tr/S \tag{3}$$

Time spent transmitting, T_{TX}, is the time required for Di slots to pass the transmitting station, which is

$$T_{TX} = D_i(Tr/S) \tag{4}$$

using equations 2, 3 and 4

$$RI = \frac{S + D_i + 1}{R_{cell}} \tag{5}$$

hence as Di cell are transmitted, the maximum throughput, σ is

$$\sigma = \frac{D_i . R_{cell}}{S + D_i + 1} \tag{6}$$

As D_i tends to a large number, σ tends to R_{cell}, but for $D_i = 1$, the reset overhead is very high, as only one slot in $(S + 2)$ slots passing the transmitting station is carrying data.

4.2. Four Stations Transmitting Randomly

For this case it is assumed that every available slot is filled during the transmit period, and all stations become idle at once. The mean time required to transport a cell is $Tr/2$.

$$T_{TX} = \left(\frac{n . Di . \left(\frac{Tr}{2} \right)}{S} \right) \tag{7}$$

The mean time each station is paused is

$$T_{paused} = Tr + Tr/2 \tag{8}$$

using equations 2, 8, and 9

$$RI = \frac{nD_i + 3S}{2R_{cell}} \tag{9}$$

and

$$\sigma = 2R_{cell}\left(\frac{nD_i}{nD_i + 3S}\right) \tag{10}$$

Equation 10 illustrates that the ring throughput tends towards 2 times R_{cell} when the product of n and D_i increases, or as D_i increases for a fixed number of stations.

4.3. Three stations all transmitting to the fourth

The distribution of destinations makes the ORWELL ring behave like a polling system or token ring. The time required to transmit all of the D_i allocations is

$$T_{TX} = (n-1)D_i(Tr/S) \tag{11}$$

the time for the first station to be reset and start transmitting again is

$$T_{paused} = \left(\frac{Tr}{n}\right)((n-1)+2) \tag{12}$$

using equations 2, 12, and 13

$$RI = \frac{D_i(n-1) + S\left(1+\frac{1}{n}\right)}{R_{cell}} \tag{13}$$

and

$$\sigma = \frac{D_i(n-1)R_{cell}}{D_i(n-1) + S\left(1+\frac{1}{n}\right)} \tag{14}$$

The throughput tends to R_{cell} with increasing $n.Di$, which is similar to the throughput for a single station on the network.

5. COMPARISON OF CALCULATED AND MEASURED RESULTS

The measured results and analytical results show close agreement, with the largest discrepancy being about 5% in the case of four stations transmitting randomly. The difference between measured and calculated results in this case can be attributed to the assumption that all stations become paused at the same instant, which is less likely for higher Di allocations leading to a longer Reset Interval. The variation of Reset Interval against D_i allocation is shown in figure 1, and that of throughput against D_i allocation in figure 2. Measured results are indicated in the legend by (m), and calculated results from section 4, by (c).

The calculated and measured results in figures 1 and 2, are made at D_i allocations between 1 and 15, with the active stations transmitting at the maximum rate possible, throughput being limited by the bandwidth of the ring and the Orwell protocol. Figure 3 shows the measured Reset Interval for a fixed Di-allocation of 10 for each active station, at varying levels of throughput. The cell loss rate is less than 1 in 1000 at the maximum throughput shown in figure 3.

The analytical and experimental results in figures 1 and 2, show that while the ORWELL reset interval may be used as an indicator of throughput and of network loading for a constant distribution of traffic, for a changing distribution of traffic, the reset rate may not give an accurate indication of the level of traffic on the network. In a situation where the overall traffic flow is symmetrical and many relatively low bit rate connections are multiplexed, such as a switch for 64kb/sec voice data, the reset rate may be a good indicator of network loading as a stable traffic flow is likely to exist. Because stations measure the reset interval independently of any knowledge of how many other stations are transmitting, or what overall traffic flow exists on the network, the measured reset interval alone may give misleading information as to the loading of the slotted ring, and the bandwidth available for new traffic to be accepted on to the network. In a multimedia LAN where high bit-rate one-way transmissions may occur, the reset rate could give an incorrect indication of the loading on the network and on the probability that traffic congestion could occur. The value of the reset interval as an indicator of network loading for congestion management purposes is therefore questionable.

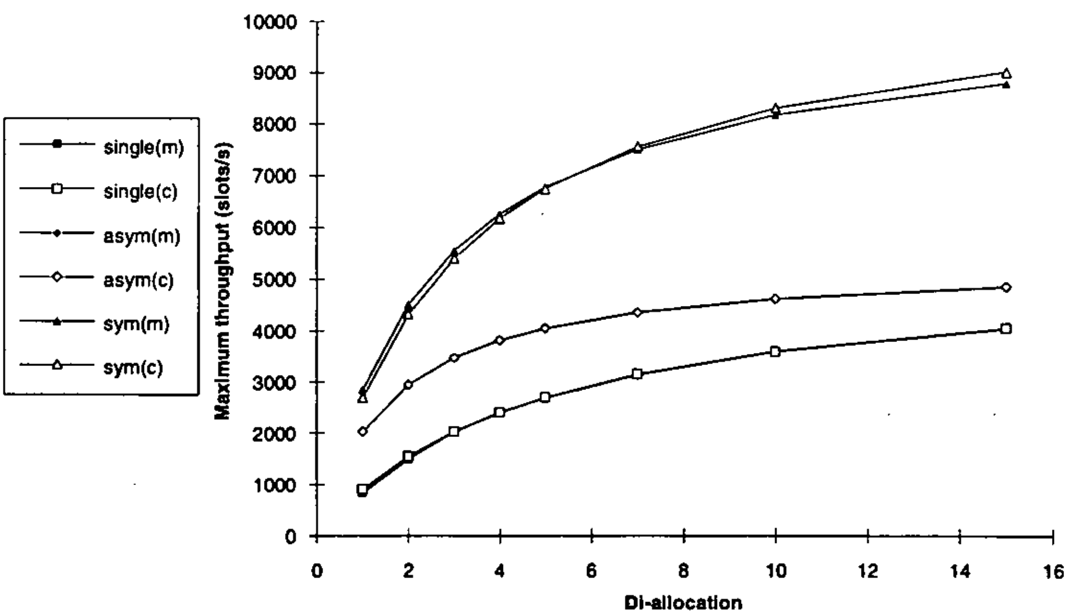

Figure 1. Throughput against Di-allocation for three traffic flows

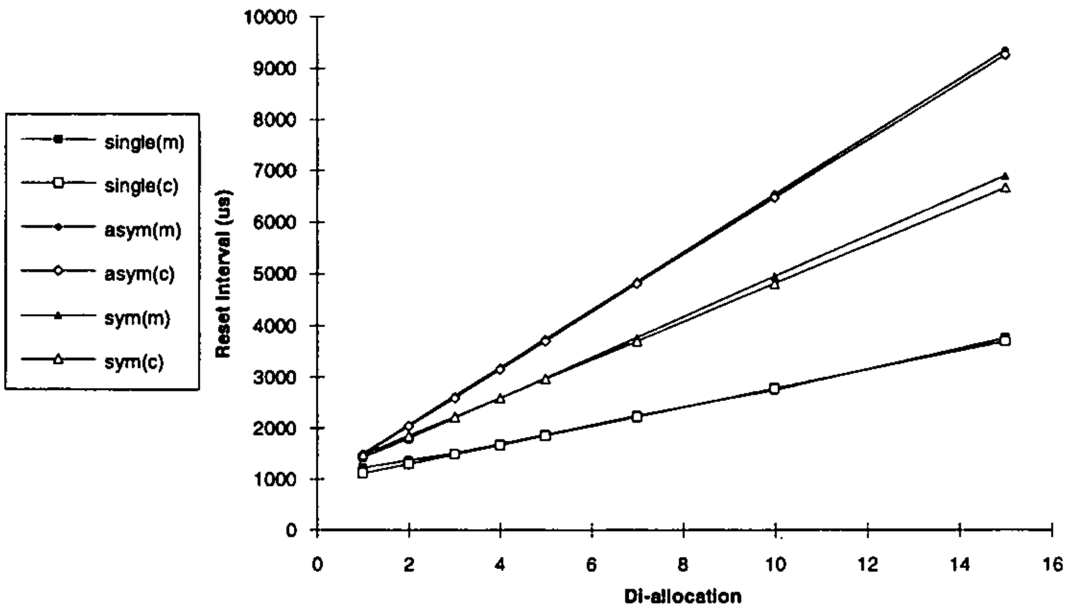

Figure 2. Reset Interval against Di allocation for three traffic flows

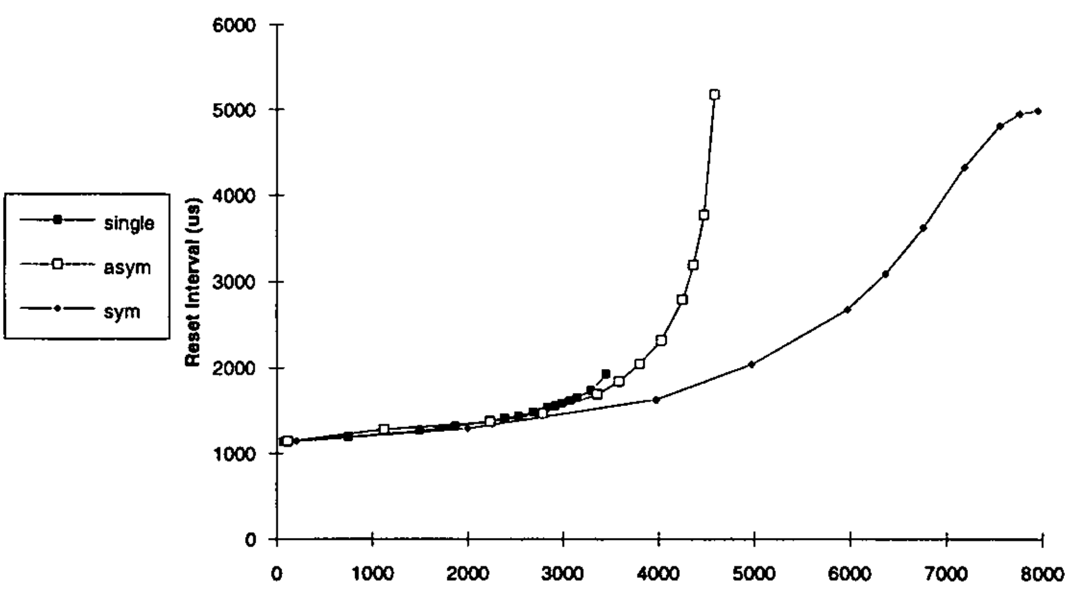

Figure 3. Reset Interval against throughput for three traffic flows with Di = 10