2000

Bandwidth management and quality of service

Adalbert Engel
Edith Cowan University

Bandwidth Management and Quality of Service

by

Adalbert Engel

BE, GradDipComp

A Thesis Submitted in Partial Fulfilment of the

Requirements for the Award of

Master of Science (Computer Science).

At the Faculty of Communications, Health and Science

Edith Cowan University, Mount Lawley.

Date of submission: 14 December 2000

ABSTRACT

With the advent of bandwidth-hungry video and audio applications, demand for bandwidth is expected to exceed supply. Users will require more bandwidth and, as always, there are likely to be more users.

As the Internet user base becomes more diverse, there is an increasing perception that Internet Service Providers (ISPs) should be able to differentiate between users, so that the specific needs of different types of users can be met. Differentiated services is seen as a possible solution to the bandwidth problem. Currently, however, the technology used on the Internet differentiates neither between users, nor between applications.

The thesis focuses on current and anticipated bandwidth shortages on the Internet, and on the lack of a differentiated service. The aim is to identify methods of managing bandwidth and to investigate how these bandwidth management methods can be used to provide a differentiated service. The scope of the study is limited to networks using both Ethernet technology and the Internet Protocol (IP). The study is significant because it addresses current problems confronted by network managers.

The key terms, Quality of Service (QoS) and bandwidth management, are defined. "QoS" is equated to a differentiating system. Bandwidth management is defined as any method of controlling and allocating bandwidth. "Installing more capacity" is taken to be a method of bandwidth management.

The review of literature concentrates on Ethernet/IP networks. It begins with a detailed examination of definitions and interpretations of the term "Quality of Service" and shows how the meaning changed over the last decade. The review then examines congestion control, including a survey of queuing methods. Priority queuing implemented in hardware is examined in detail, followed by a review of the ReSource reserVation Protocol (RSVP) and a new version of IP (IPv6). Finally, the new standards IEEE 802.1p and IEEE 802.1Q are outlined, and parts of ISO/IEC 15802-3 are analysed.

The Integrated Services Architecture (ISA), Differentiated Services (DiffServ) and MultiProtocol Label Switching (MPLS) are seen as providing a theoretical framework

for QoS development. The Open Systems Interconnection Reference Model (OSI model) is chosen as the preferred framework for investigating bandwidth management because it is more comprehensive than the alternative US Department of Defence Model (DoD model).

A case study of the Edith Cowan University (ECU) data network illustrates current practice in network management. It provides concrete examples of some of the problems, methods and solutions identified in the literary review.

Bandwidth management methods are identified and categorised based on the OSI layers in which they operate. Suggestions are given as to how some of these bandwidth management methods are, or can be used within current QoS architectures.

The experimental work consists of two series of tests on small, experimental LANs. The tests are aimed at evaluating the effectiveness of IEEE 802.1p prioritisation. The results suggest that in small Local Area Networks (LANs) prioritisation provides no benefit when Ethernet switches are lightly loaded.

DECLARATION

I certify that this thesis does not incorporate without acknowledgment any material previously submitted for a degree or diploma in any institution of higher education; and that to the best of my knowledge and belief it does not contain any material previously written by another person except where due reference is made in the text.

Signature_

Date 6 - 3 - 2001

ACKNOWLEDGMENTS

I would like to thank my Supervisor, Dr S. Paul Maj, for his positive leadership and his patient, heavy labour in reading and annotating the earlier versions of this thesis.

My thanks also go to fellow student, David Veal, for spotting, and supplying, what turned out to be some of my most useful references.

Thanks are also due to Geoff Lourens and Taib Jaffar from the School of Computer and Information Science (SCIS) support group. To Geoff for his timely help with the configuration of the BayStack switch; to Taib for his kindness and patience: Taib, I will definitely return your BayStack 450 Switch Manual, which I borrowed "for a few days" 18 months ago.

Last but not least to my son, Carl, who got me out of many tight corners during my battles with Microsoft Office 2000.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

1.0 INTRODUCTION

1.1 Background

As the Internet gains increasing importance as a tool for information retrieval and commerce, the demands on it continue to increase. There are no official figures for Internet growth and usage. However, according to figures compiled by Coffman and Odlyzko (1998, p.2):

> *"Traffic and capacity of the public Internet grew at rates of about 100% per year in the early 1990s. There was then a brief period of explosive growth in 1995 and 1996. During those two years, traffic grew by a factor of about 100, which is about 1000% per year (sic). In 1997, it appears that traffic growth has slowed down to about 100% per year."*

This enormous growth has been absorbed by the Internet by virtue of it being a decentralised and scalable system.

The Internet is based on packet-switching networks that use the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol stack. The Internet Protocol (IP) operates at Layer 3 of the Open Systems Interconnection Reference Model (OSI model). It relies on the services provided by the Local Area Network (LAN) technologies operating at Layers 1 and 2. The dominant LAN technology is "Ethernet" as defined in ISO/IEC 8802-3. This standard defines the mechanisms, Collision Sense Multiple Access (CSMA) and Collision Detection (CD), by which Ethernet frames access the Ethernet bus.

At Layer 3, IP generates packets which share the router links on the Internet, independent of the source of the packets. Similarly, at Layer 2, frames from different applications and/or users are competing, on an equal basis, for access to shared Ethernet buses. As a result, '*The Internet is the ultimate classless society. Internet routers simply don't distinguish between an online tax return, a million-dollar extranet supply line deal and an IRC flame war*' (Tebbutt, 1998).

The inability of IP and Ethernet to discriminate between packets and frames from different streams makes the provision of differentiated services difficult.

Differentiated services may be required because the multitude of users now on the Internet is made up of different types of users, with different needs: casual users are likely to be looking for cheap e-mail and web access, whereas users with real-time applications, such as video conferencing might well be prepared to pay for higher bandwidth. However, the technology currently in use on the Internet does not provide for differentiation between users, or between applications. It is likely that the needs of some casual users, as well as some high-bandwidth users are not being met.

The presence of high-bandwidth users has been noted by Coffman and Odlyzko (1998, p.14): *"The 100% annual growth rates for the Internet were the result of an increase in the number of users as well as increased traffic from existing users."* The increased traffic from existing users is due to the availability of new bandwidth-hungry applications such as web browsers and real-time audio and video. The popularity of audio and video applications is raising bandwidth requirements on a per-user basis to unprecedented levels: Minoli and Schmidt (1999, p.139), for example, expect that some business applications incorporating video will require between 1.5 to 6 Mbits/sec per user. Importantly, these applications are also time-sensitive.

Although such bandwidth requirements, could feasibly be met for a limited number of users connected to a recently upgraded, switched Local Area Network (LAN), it is less likely that the Wide Area Network (WAN) links will be able to meet the aggregated bandwidth requirement resulting from even a moderate number of such users. To illustrate: a current-generation Ethernet switch, would have ports rated at 100 Mbits/sec and a typical backplane capacity of 2.5 Gbits/sec. Such a switch would be able to handle simultaneously 24 users each requiring a constant bit rate of 6 Mbits/sec. However, an optical fibre WAN link, typically rated at 1 Gbits/sec, would only be able to handle the aggregated load from $\left\lfloor \frac{1000}{24 \times 6} \right\rfloor = 6$ such switches.

Nevertheless, in spite of the constraints imposed by WAN links, Foo, Hui and Yip (1999) have shown that real-time audio and video can be transmitted across today's Internet. Foo et al. built their own communication system comprising a transmitter and a receiver communicating over Internet links. As will be detailed later, the system incorporates a combination of advanced techniques including data compression, buffers and audio silence deletion.

Other examples of special communication systems using the Internet are commercial telephony and video-conferencing systems currently on the market. Like the communication system built by Foo et al., these commercial systems are designed to make maximum use of the limited bandwidth available on the Internet. It is a case of the users' computers doing the work, rather than the network. Where these special communications systems are used, the bandwidth shortage is not being solved, just bypassed at the cost of high overheads and computing power provided by the user.

In view of the current and the anticipated bandwidth shortage, Internet observers like Van Houweling conclude that differentiation between users will be necessary: *"The best-effort Internet is inherently unable to guarantee with any reliability the performance necessary to run even the more moderately ambitious advanced networked applications envisioned today"* (VanHouweling, 1999, p.3).

Indeed, many research activities are currently aimed at providing Ethernet/IP systems with mechanisms for differentiating between users. The Internet Engineering Task Force (IETF) and the Institute of Electrical and Electronic Engineers (IEEE) have been able to approve several of these mechanisms and, to this effect, have recently issued a series of standards such as 802.1Q – Virtual Bridged Local Area Networks and 801.1p – Traffic Class Expediting and Dynamic Multicast Filtering, both to be discussed later.

Others Internet observers, such as Odlyzko (1998, p.18) contend that providing high bandwidth for all users might be sufficiently economical to make it the better solution. Whichever solution dominates, efficient bandwidth management methods will form part of either solution.

1.2 The Problems

The thesis focuses on two network management problems:

Problem 1: There is never enough bandwidth.

Bandwidth shortage is a common problem faced by Internet Service Providers (ISPs) and network managers. Clark (1996, p.681) expresses it as follows: *"...experience has shown that there never seems to be enough bandwidth. As more bandwidth is provided, users conceive new applications that consume it."*

With the number of Internet users still increasing and with bandwidth-hungry applications, such as audio and video, becoming more popular, bandwidth shortages are likely to remain.

Problem 2: Different users and/or applications require different levels of service.

The need to differentiate between users and/or applications poses a problem for managers of Ethernet/IP systems because Ethernet and IP are basically incompatible with the QoS concept. An IP network is connectionless. IP packets are routed dynamically as links become available. IP does not distinguish between packets from different flows. Thus at Layer 3 of the OSI model, no differentiation is made between users.

Similarly, at Layer 2, frames from different applications and/or users are competing for access to Multi Access Control (MAC) media, usually an Ethernet bus. The IEEE 802.3 standard does not distinguish between frames from different users.

1.3 Research Objectives

The objectives of this research are to -

1. Investigate bandwidth management methods.

The aim is to address Problem 1 above, by identifying methods and possible solutions for optimising bandwidth utilisation. Methods to be investigated include network segmentation, Virtual LANs (VLANs) and a range of queuing algorithms.

2. Evaluate IEEE 802.1p prioritisation as a mechanism for differentiated services.

The aim is to address Problem 2 above, by investigating how IEEE 802.1p prioritisation can be used as part of a differentiating system that provides end-to end Quality of Service (QoS).

This second part of the research project is based on the assumption that it may not be possible to meet future bandwidth demands simply by installing additional capacity. Here it is assumed that QoS, in some form or other, is unavoidable, and that its introduction is only a matter of time. This is indeed suggested by the bulk of current literature, although there are other views – refer 2.4 'QoS vs the "Bigger pipe" Approach'.

Scope of the Study: The investigation is carried out in the context of networks that use Ethernet technology together with the Internet Protocol (IP). Unless otherwise stated, discussion and analysis will refer to Ethernet/IP systems. Where alternative data transmission systems, such as Asynchronous Transfer Mode (ATM), alternative Layer 3 protocols such as IPX (Novell's Internetwork Packet Exchange), or alternative LAN technologies such as Token-ring, are under consideration, they will be specifically mentioned.

1.4 Significance of the Study

The Internet is now of importance in commerce, education and entertainment. The unprecedented and unpredicted demands appear to be difficult to meet. This study addresses two current networking issues:

1. Bandwidth management,

2. QoS for LANs

It should be of interest to Internet service Providers (ISPs) and network managers in both the public and private sectors.

1. Bandwidth management

Conserving bandwidth remains an important issue for network managers and ISPs because of *"an increase in the number of users as well as increased traffic from existing users"* Coffman and Odlyzko *(1998, p.14)*. Improving bandwidth management is seen as a key issue by Adams:

> *"Better bandwidth management, opposed to building extra capacity, has to be the answer. Bandwidth management not only increases a network's usable capacity, it also provides carriers with increased knowledge of network traffic patterns, resulting in more intelligent use of the network."* (Adams, 1998, p.34)

With improved bandwidth management, the utilisation of a network is reduced, rendering the network more readily available to the devices connected to it.

If bandwidth management could be improved to an extent where a significant amount of additional bandwidth is made available to users, then the need for differentiated services would be smaller and the "Bigger pipe" Approach (which rests on the premise that enough bandwidth can be made available to all users - refer 2.4 'QoS vs the "Bigger pipe" Approach') would be seen as a more realistic alternative.

2. QoS for LANs

The working committees of the Internet Engineering Task Force (IETF) have concentrated on QoS models which supplement the IP protocol and which are intended

to provide QoS across WANs. However, LANs represent an essential step in the path to the Internet. QoS for LANs is currently a recognised field of research:

> *"The quality models for local area networks and wireless networks are still somewhat unclear, at least when assessed from the viewpoint of Differential Services. These issues definitely require more research and development to attain a consistent QoS structure throughout all major packet networks."* (Kilkki, 1999, p.293)

1.5 Definitions of Key Terms

Bandwidth is "*the range of signal frequencies that a circuit passes*" (Halsall, 1993, p.50). Bandwidth imposes a limit on the transfer of information. In fact, it can be seen from the Nyquist formula –

$$C = 2B\log_2 M$$

where C = maximum data transfer rate (bits/sec)

B = bandwidth (cycles/sec)

M = levels per signalling element

that the maximum data transfer rate of a line or data transmission system is proportional to the bandwidth.

A digital signal being transmitted along a line is subject to –

- attenuation
- distortion due to the harmonics making up the signal being attenuated by different amounts
- delay distortion or intersymbol interference[1].

[1] According to Halsall (1993, p.34) and Tanenbaum (1996, p.109) delay distortion is due to the harmonics travelling at different speeds and thus being delayed by different amounts. This idea appears to contradict the axiom that all electromagnetic radiation travels at the same speed in the same medium. The author suspects that delay distortion is due to the harmonics being subject to different phase shifts.

If the bit rate of the transmitted data is increased, while the bandwidth of the line remains fixed, then distortion of the digital signal will reach a point where the signal becomes unintelligible.

Bandwidth Management. For the purpose of this thesis "bandwidth management" is taken to include all hardware and software methods of controlling and allocating bandwidth.

Increasing bandwidth by adding more capacity will also be considered a method of managing bandwidth. This is in line with the approach taken by Determan who writes: *"QoS allows bandwidth management without adding more capacity."*(1999, p.12)

Determan is implying that on occasions, such as when QoS is not available, bandwidth management is carried out by adding more capacity. In other words, Determan considers "adding more capacity" a method of bandwidth management.

In the classification of bandwidth management methods proposed in Chapter 5, "adding or reducing capacity" will be classified as a Layer 1 bandwidth management method.

Quality of Service. In this thesis, Quality of Service (QoS) will refer to a differentiating system. For example, "to provide QoS" will mean to implement a differentiating system. This is the meaning adopted by authors such as Breyer and Riley – refer 2.3.1 "What is Quality of Service?".

2.0 REVIEW OF LITERATURE

2.1 Scope of the Literary Review

In keeping within the scope of the study (see 1.3 "Research Objectives"), this review of literature concentrates on developments in the Ethernet/IP area. References to Asynchronous Transfer Mode (ATM) are included to put Ethernet/IP into perspective, rather than to provide a detailed review of ATM.

As most of the research in bandwidth management and QoS is currently taking place on the Ethernet/IP front, the bulk of recent (1999 and 1998) literature on bandwidth management and Quality of Service (QoS), including papers and standards issued by the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronic Engineers (IEEE), is in fact written in an Ethernet/IP context. This review therefore also reflects a current literary trend.

2.2 The Internet Today

Growth of the Internet varies between the public and private sectors. Coffman and Odlyzko (1998, p.2) estimate that private line networks are growing 15% to 20% per year, while the public sector is growing by at least a 100% per year.

In the absence of a differentiated service, and with bandwidth limited, the Internet is currently struggling to meet the demands from real-time audio and video. Foo, Hui and Yip (1999, p.212) describe the current situation on the Internet as follows:

> *"the Internet presents a relatively harsh environment for real-time services especially for those that require large bandwidths such as video data....The potential high transmission delay and data packet loss of the Internet environment that is characteristic of a packet switched network without resource reservation mechanisms, have made real-time communications services difficult."*

One approach to meeting the demand is to compress the data so that fewer packets need to be send for a given amount of data. An illustration of this approach is a communications system built by Foo et al. for communicating over existing Internet links. The system comprises a transmitter, which acquires data and prepares it for transmission, and a receiver, which manages the received packets. To reduce the bandwidth requirements of the communications system, Foo et al. not only use data

compression but a combination of techniques including dynamic rate control, buffers, data packet lost replacement, audio silence deletion and dynamic video frames reconstruction.

Foo et al. have shown that it is possible to transmit real-time audio and video across the Internet. However, it appears that the transmissions so far are across local links only. Work on various gateways, including one that aims to integrate the Internet with the Public Switched Telephone Network (PSTN), is in progress.

The work by Foo et al. is an example of Internet users taking the initiative, by providing supplementary systems designed to make do with the Internet's limited bandwidth. Another example of these initiatives is the purchase of commercial telephony and video-conferencing systems such as Microsoft's "NetMeeting" (Foo et al., 1999, p.212).

In evaluating these systems, one must consider their current high complexity and cost. It should also be noted that it is the computers at each end of the line that are doing most of the work, not the network. The bandwidth shortage is not being resolved, only accommodated at the cost of additional computing power. A more basic, and more promising solution for resolving the bandwidth shortage may be to optimise bandwidth utilisation by providing different levels of service to different users.

2.3 Quality of Service

Quality of Service is not a new concept. QoS has always been an integral part of ATM. Prior to the recent Ethernet/IP developments mentioned above, QoS literature, if not written in a general, issues-based manner, was often written in the context of ATM. This is no longer the case. Most of the recent QoS literature is written in an Ethernet/IP context. Topical QoS issues and the authors dealing with them include –

QoS routing: (Apostolopoulos, Guerin, Kamat, & Tripathi, 1998),

(Faloutsos, Banerjea, & Pankaj, 1998);

Integrated Services Architecture: (Wolf, 1999), (Minoli & Schmidt, 1999);

Differential Services: (Stoica & Zhang, 1999), (Wroclawski, 1999).

In fact, recent literature reflects a revival in the popularity of Ethernet/IP relative to ATM. Thus "Quality of Service" is now discussed more often in the context of Ethernet/IP rather than in an ATM context.

2.3.1 What is Quality of Service?

The phrase "Quality of Service" is used literally by Halsall (1993, p.17): *"The transport layer offers a number of classes of service which cater for the varying quality of service (QOS) provided by different types of network."* It seems that as far as Halsall was concerned in 1993, "quality of service" in the context of networks, simply referred to the quality of the service provided by a network. Furthermore, to provide a way of quantifying "quality of service", Halsall defined a "QOS parameter":

> *"The quality of service (QOS) parameter comprises two lists of parameters: one the desired and the other the minimum acceptable parameters expected from the network for the connection being set up. These include (packet) transit delay, residual error probability, priority (if applicable), cost (charge for the call) and specified (rather than arbitrary) route."* (1993, p.384)

Thus Halsall sees "quality of service" as a method of quantitatively specifying the quality of the service provided by a network. Tanenbaum too uses the phrase literally: *"Each service can be characterized by a quality of service"* (1996, p.23). However, Tanenbaum, uses it in the context of ATM:

> *"Quality of service is an important issue for ATM networks, in part because they are used for real-time traffic, such as audio and video. When a virtual circuit is established, both the transport layer (typically a process in the host machine, the "customer") and the ATM network layer (e.g., a network operator, the "carrier") must agree on a contract defining the service"* (1996, p.460)

Taken together, the last two quotations mean that according to Tanenbaum, quality of service characterises the service provided by the ATM network layer to the transport layer.

However, according to Breyer and Riley (1999, p.501): *"QoS refers to the capability to reserve bandwidth through a network."* Significantly, Breyer and Riley are not talking about the literal meaning of "quality of service", but rather about a system that reserves bandwidth.

Furthermore, Breyer and Riley (1999, p.72) see QoS as an alternative to providing more bandwidth: *"Raw bandwidth is a pretty good substitute for QoS"*. In other words, "QoS", being a system that reserves bandwidth, is seen as an alternative to another system, which aims to provide "raw bandwidth" for all users. Breyer and Riley call this latter system the "bigger pipe" approach.

Similarly, Determan (1999) sees "QoS" as a differentiating system: "*QoS is a way to prioritize and allocate resources in network devices....QoS allows bandwidth management without adding more capacity*". Like Breyer and Riley, Determan considers "QoS" an alternative to the "bigger pipe" approach.

Other authors that see "QoS" as a differentiating system include Maamria (1998) and Held (1997):

> Maamria (1998, p.43): "*QoS is also defined as the ability to provide differentiation between customer and application types.*" That is, "QoS" is here seen as a system able to differentiate between different users and between different applications.
>
> Held (1997, p.206): "*QoS is a term that comes from ATM...it means being able to negotiate a certain, guaranteed service quality prior to communicating.*" Thus, Held introduces another aspect – guarantees. In his view, "QoS" not only differentiates, but also provides a guarantee.

It can be concluded from the above quotations that there is no agreed definition of "QoS". The meaning of "QoS" varies from author to author. There is also evidence that the meaning has undergone some evolutionary change: originally authors saw "QoS" as a quantitative measure of the quality of service provided by a network, particularly an ATM network. The view was that QoS would guarantee a service with respect to one or more performance parameters.

In contrast, many recent authors see "QoS" as a differentiating system and they consider "QoS" to be an alternative to non-differentiating systems such as the "bigger pipe" approach. In this thesis, "QoS" will refer to a differentiating system. As such it will be compared and contrasted to alternative systems such as the "bigger pipe" approach.

2.3.2 What Quality of Service is not

Contrary to what the phrase implies, Quality of Service cannot provide better service for all. Van Jacobson (1998) puts it as follows: *"QoS does not create bandwidth and, since bandwidth allocated to some is bandwidth not available to others, QoS does not guarantee that you get better service."* (p.26)

As stated by Maamria above, QoS is about differentiating between users and application types. "Quality of Service" is a misnomer and would be more aptly referred to as "Differentiated Services". Indeed, one of the new IETF working groups tasked with "QoS" research has been named the "Differentiated Services" (DiffServ) working group.

2.3.3 The Scope of QoS

Although there seems to be a general agreement amongst current authors that QoS is a differentiating system, relatively little is said about the criteria on which this differentiation is to be based. Breyer and Riley state that *"Guaranteed availability of a minimum amount of bandwidth is a key ingredient of QoS"* (Breyer & Riley, 1999, p.44). Of course the "guaranteed availability of a minimum of bandwidth" will apply only to users prepared to pay for it.

Breyer and Riley (1999, p.61) see hardware-based prioritisation and RSVP as significant steps forward: *"The new IEEE 802.1p and IETF RSVP standards go a long way to providing QoS."* However,

> *"Ethernet does not have any built-in quality–of- or class-of-service guarantees. Recent efforts, such as the IEEE 802.1p and IETF RSVP standards, will improve Ethernet's capabilities in this respect, but true QoS will remain elusive."* (p.51)

Thus, according to Breyer and Riley, even when prioritisation is backed up by RSVP (ReSource reserVation Protocol), "true QoS" will not be achieved.

There is as yet no accepted definition of "true QoS". The importance of the QoS parameters in defining a particular QoS is perhaps best explained by Teitelbaum:

> *"A fundamental dimension of any application's QoS requirements is the set of transmission parameters about which assurances are needed. The transmission parameters most commonly mentioned as requiring assurances are bandwidth and latency."* (1999, p.7)

The QoS parameters associated with the Integrated Services Architecture (ISA) service categories are detailed under 3.1 "QoS Control Architectures".

2.3.4 QoS and the Role played by the User

Another aspect of QoS about which relatively little is said in current literature, is the constraint placed on the user. It is generally recognised that where a Service Level Agreement (SLA) exists, the network operator provides a guarantee, that the specified limits on the QoS parameters will not be exceeded. However, guarantees cannot be provided in isolation: the performance of a data channel, with a given bandwidth, will depend on the nature of the load. No matter what QoS parameters limits are specified, the actual delay and jitter will depend on the number of users and their applications.

It follows that the user too has a part to play, by keeping the load profile within certain limits. As indicated by Minoli (1999, p.409), applications need to be "controlled". The controls placed on applications run under ISA are detailed in "QoS Control Architectures" below.

2.3.5 Types of QoS

If QoS (i.e., a differentiating system) is to be implemented, which type should be used? As the type of QoS that can be chosen depends on the data transmission system in use, the argument here boils down to which data transmission system should be used. The dominant data transmission systems are -

- Asynchronous Transfer Mode (ATM) and
- the combination of Ethernet technology and Internet Protocol, which in this thesis is referred to as the Ethernet/IP system.

Huitema comes out strongly in favour of the Ethernet/IP system: "*ATM networks try to provide a full range of quality of services at the expense of extreme management complexity*" (1997, p.192).

In addition, Huitema sees ATM as a potentially expensive solution for the user. Noting that ATM uses a range of bit rates, he expects that users would be charged at different rates for different applications. In particular, a user would be charged every time an

application establishes a connection. In contrast to ATM QoS, Ethernet/IP QoS envisages a single charge rate per communication channel, which would be more economical for the users, because they would be able to buy bandwidth in bulk.

Minoli and Schmidt take a balanced view: "*ATM supports QoS very well, for example in an end-to-end pure ATM network. However, ATM must be used in the context of other embedded protocols, such as IP; hence QoS has to be ultimately understood in that environment*" (Minoli & Schmidt, 1999, p.237). Minoli and Schmidt are anticipating a third type of transmission system: a combination of ATM and Ethernet/IP, such as the hybrid system at Edith Cowan University (ECU) - 4.2.3 Extent of ATM.

Internet observers increasingly seem to accept that future networks are likely to be hybrid rather than purely ATM or Ethernet/IP. Minoli and Schmidt (1999, pp.377-391) describe several methods by which ATM supports IP. These include –

- LAN emulation (LANE)
- Classical IP-over-ATM (CIOA) and
- Multiprotocol over ATM (MPOA).

The above methods provide continuity between (otherwise incompatible) ATM and Ethernet/IP networks, but they do not integrate the two disparate types of QoS associated with ATM and Ethernet/IP. QoS integration is tackled by Wolf (1999), who specifically looks at ways ATM QoS and Ethernet/IP QoS can be combined to provide end-to-end QoS:

> "*We believe that interaction approaches for the QoS architectures of the Internet and ATM are necessary because both worlds will co-exist for a couple of years. Since they tend to increasingly serve the same applications due to the pertaining convergence process of data and telecommunications, they have to interwork with each other to fulfil application demands.*" (Wolf, 1999, p.57)

2.4 QoS vs the "Bigger pipe" Approach

Traditionally, the Internet has always provided a single class of service. As stated by Tebbutt (1998, p.384), *"the Internet is the ultimate classless society"*. The notion that different users and different applications require different treatment, and that the Internet cannot remain "classless" seems self-evident. Seen in this light, a differentiating system (i.e. QoS) appears to be the obvious solution to future demands on bandwidth. However, the complexity and overheads entailed in providing end-to-end QoS have led some Internet observers to question whether it is the best solution.

There are in fact two main schools of thought:

- QoS i.e., a differentiating system
- An undifferentiating system where the aim is to provide high bandwidth for all users, called the "bigger pipe" approach by Breyer and Riley (1999, p.72).

2.4.1 QoS

QoS is discussed in most recent texts dealing with the Internet or network engineering. These include (Breyer & Riley, 1999), (Minoli & Schmidt, 1999), (Minoli & Alles, 1996) and (Tanenbaum, 1996). Some recent papers dealing with QoS in the context of Ethernet/IP have been listed under "Quality of Service" above.

QoS is the sole topic at a series of annual workshops known as the IFIP International Workshop on Quality of Service. At the 5th workshop (Angin et al., 1997), held in May 1997, the theme was "Building QoS into Distributed Systems". The workshop included sessions on QoS Routing, traffic management, QoS management and QoS-based transport protocols. The topics suggest that researchers had reached the stage where they were able to focus on a number of key issues. The delivery of 20 experimental research papers, in addition to 20 speculative position statements, suggests that in 1977

QoS research had already reached a degree of maturity and was moving into development.[2]

The considerable research efforts that have recently gone into QoS reflect the dominant view that differentiation between users, that is QoS, will be necessary. For example, the recently launched Internet2 Project, which promotes collaborative research by 130 universities, 40 corporations and 30 other organisations, fully endorses the principles of QoS:

> *"From the beginning, quality of service (QoS) has been essential to the Internet2 vision. The best-effort Internet is inherently unable to guarantee with any reliability the performance necessary to run even the more moderately ambitious advanced networked applications envisioned today."* (VanHouweling 1999)

2.4.2 The "Bigger pipe" Approach

The "bigger pipe" approach aims to provide high bandwidth for all users. It is a simple solution that avoids the complexities of QoS. However, its success will depend on the amount of bandwidth that can be made available, which in turn will depend on the effectiveness of bandwidth management methods (refer 1.4 "Significance of the Study"). Nevertheless authors such as (Breyer & Riley, 1999), (Odlyzko, 1998) and (Bajaj, Breslau, & Shenker, 1998a) are optimistic about the "bigger pipe" approach and sceptical of the need for a differentiated service. They leave open the possibility that the "bigger pipe" approach may be the better solution.

Breyer and Riley (1999, p.72) consider that QoS is only needed under certain conditions:

> *"Raw bandwidth is a pretty good substitute for QoS - QoS is necessary when a network is overloaded and sporadic delays are a normal part of the operation of the network."*

Thus, according to Breyer and Riley, if enough bandwidth can be provided, then QoS may not be necessary. Whether enough bandwidth can be provided will depend on the price of bandwidth:

[2] At the time of writing, the author has not been able to find the proceedings for the 1998 and 1999 workshops.

> *"If prices do decrease sufficiently rapidly compared to traffic growth, then it might be economically optimal to continue with the present system of flat rate pricing, and to provide high-quality service to all packets."*(Odlyzko, 1998, p.18)

The simplicity of "flat rate pricing" and of providing the same service to all packets makes the "bigger pipe" approach attractive: *"Given the huge potential cost to the entire IT system of any modifications to the Internet, though, simplicity will surely be at a premium"* (Odlyzko, 1998, p.18). Odlyzko considers the current QoS schemes too complicated. Bajaj, Breslau and Shenker express similar sentiments, also indicating where the complexities will be:

> *"The question remains as to what benefits offering additional levels of service would provide. Offering multiple levels of service carries with it the cost of additional complexity to deal with signaling the priority level, merging reservations with different priority levels, and scheduling overhead."* (Bajaj, 1998a, p.67)

Last, an unusual argument in favour of the "bigger pipe" approach is put forward by Breslau and Shenker:

> *"A reservation-capable network will not deliver satisfactory service unless its blocking rate (the rate at which it denies reservation requests) is low, and at such provisioning levels best-effort networks will provide completely adequate service".* (Breslau, 1998, p.4)

This can be interpreted as meaning that the repeated denial of requests makes the reservation-capable network inefficient. Alternatively, Breslau and Shenker may be referring to human nature: users will not tolerate a system that frequently withholds service.

2.5 Congestion Control

At low utilisation a network can generally respond to increasing demands. That is, within a certain working range, the network's throughput will increase as more packets are injected into it. However, there comes a time when the network cannot handle any more packets. If more packets are injected, some packets will get lost and will have to be re-transmitted. At this point performance deteriorates markedly and may collapse altogether (Tanenbaum, 1996, p.374).

When congestion occurs the normal network traffic controls may not function properly, as when TCP acknowledgement packets are lost. In that situation the network's bandwidth is not fully utilised, and is in effect wasted. Thus congestion control can be considered to be a method of bandwidth management.

Methods of congestion control have been subject to extensive treatment by many authors. These include (Bajaj, Breslau, & Shenker, 1998b) on load-shedding, (Reardon, 1998) on traffic shaping and (Guerin, Kamat, Peris, & Rajan, 1998) on buffer allocation.

The two basic approaches to congestion control are Admission Control, which denies access to the network once congestion occurs, and traffic shaping, which makes the traffic flow less bursty. Traffic shaping in turn can be sub-divided as follows:

1. Queuing
2. Rate Control Mechanisms

In the past traffic shapers have been widely used in ATM networks (Tanenbaum, 1996, p.379). Traffic shapers for IP networks, on the other hand, are relatively new, but they are now available from many vendors (Reardon, 1998).

2.6 Queuing

In queuing, the traffic is separated into different queues. The packets in the queues are dealt with in accordance with a pre-set priority. Methods of queuing include the following:

1. Priority Queuing
2. Fair Queuing
3. Weighted Fair Queuing

2.6.1 Priority Queuing

In priority queuing the high-priority queues are emptied before any lower-priority traffic is transmitted. *"This approach works fine for bursty traffic, but if policies aren't properly set then low-priority traffic can be starved of bandwidth"* (Reardon, 1998).

Notably, neither the high-priority queues nor the low-priority queues are guaranteed any bandwidth. This shortcoming is addressed in Class-Based Queuing (CBQ), in which a high-priority queue is guaranteed a certain bandwidth. If the high-priority queue does not need all of its bandwidth, then other traffic can use the idle bandwidth.

Provided the high-priority queue is not guaranteed all of the available bandwidth, there will always be some bandwidth available for the low-priority traffic. Thus CBQ also avoids starvation.

QoS by Buffer Management

An alternative method of implementing a priority queue is to give it a larger buffer.

Guerin, Kamat, Peris and Rajan (1998) use a mathematical model to simulate a FIFO (first in, first out) queue. Guerin, et al. (1998) provide expressions that associate bit rate guarantees with buffer allocation. They conclude that bit rates can be guaranteed by simply using buffer management: *"We have established how rate guarantees can be provided by simply using buffer management. Exact expressions were provided that associate rate guarantees with buffer allocation in a simple FIFO queue"* (Guerin et al., 1998, p.39). Since rate guarantees are essential components of Service Level Agreements (SLAs), this type of buffer management could be used as a method of providing a differentiated service.

Normally use of buffers is limited to temporary storage. For example, in the communications system designed by Foo, Hui and Yip (1999): "*The purpose of buffer management is to cushion the out-of-order, late delivery and jitters experienced by the data packets.*" Guerin et al. go beyond the normal use of buffers by showing that buffers can be used as the sole mechanism for providing a differentiated service.

Guerin et al. also put forward a hybrid scheme, which combines buffer management and Weighted Fair Queuing (WFQ), and which, they consider, has "*some potential benefits*". In this hybrid scheme the FIFO queue is replaced by a number of queues served by a WFQ scheduler. The combination of buffer management and WFQ would therefore be a refinement to abovementioned buffer management scheme.

2.6.2 Fair Queuing

The Fair Queuing (Round Robin) algorithm is implemented by establishing multiple queues at each output line, one queue for each source. The router scans the queues in turn, takes one packet from a queue, forwards the packet and then moves on to the next queue. In this way Fair Queuing ensures that the output lines on a router are equally shared among the users. Sending more packets will not increase a user's throughput (Tanenbaum, 1996, p.388). However, the bandwidth available to a user at any one time will depend on the number of users at the time. Since the number of users can change, the bandwidth per user cannot be guaranteed.

2.6.3 Weighted Fair Queuing

Weighted Fair Queuing (WFQ) is an early priority algorithm intended to give more bandwidth to servers than to clients. For example, the weight may be allocated on the basis of the number of flows coming out of a machine, so that a server running several processes will get more bandwidth (Tanenbaum, 1996, p.389).

By allocating a weight to each user, WFQ guarantees each user a fixed portion of the channel bandwidth. That is, *"traffic is not only assigned to specific priority queues; it is also apportioned a share of bandwidth"* (Reardon, 1998). Guaranteed bandwidth per user means that WFQ can provide limited QoS (no control is exerted over delay and jitter, which will depend on the traffic pattern). However, if the traffic is made smooth by adding a leaky bucket algorithm, either in hardware or in the source's operating system, then it is possible to provide Guaranteed Service as defined for the Integrated Services Architecture (ISA).

WFQ is used by Guerin et al. (1998) to refine a QoS scheme based on buffer management – refer "Priority Queuing" above.

2.6.4 Random Early Detection

Random Early Detection (RED) is a refinement to the queuing process and could, in theory, be used in conjunction with any of the above queuing methods. Queuing schemes without early detection are obliged to simply drop all incoming packets once the queue is full. A RED scheme, on the other hand, drops packets randomly when the load has

exceeded a pre-set limit, even before the queue is full. In this way the onset of congestion is slowed down, if not avoided altogether. Should the buffer fill up regardless, then all incoming packets are discarded.

RED is a proactive countermeasure to congestion that anticipates the onset of congestion and counteracts it. It is considered a realistic application for inclusion in Differentiated Services networks (Kilkki, 1999, p.193). Even QoS sceptic Odlyzko considers WFQ and RED a useful way to provide unobstrusive QoS "*without destroying the exceedingly valuable stateless nature of the Internet*" (Odlyzko, 1998, p.12).

Packet Dropping Mechanisms

Dropping packets randomly (also known as "Uniform Dropping") means that all packets are treated equally with respect to dropping. A possible refinement for RED could be the use of priority dropping, where lower priority packets are dropped before higher priority packets. Priority dropping could conceivably be part of a QoS architecture where the packet loss ratio is of importance, for example, when the packet loss ratio is stipulated in the Service Level Agreement (SLA) between the service provider and user.

Bajaj, Breslau and Shenker (1998b) conducted mathematical simulation experiments in which they allocated a priority to the dropping of packets. They found that although priority dropping performed better than uniform dropping in all of their experiments, the benefits of priority dropping over uniform dropping were less than expected. This may have implications for other priority schemes such as priority queuing.

2.7 Rate Control Mechanisms

Whereas the majority of bandwidth management methods try to provide additional bandwidth to match an application's requirements, a rate control mechanism aims to adjust an application's bandwidth requirements to match the available bandwidth.

Such a mechanism is described by Bolot and Turletti (1994) who employed their rate adaptive coding algorithms to control packet video transmissions. Bajaj et al. (1998, p.132) found the method unsuitable for multicast transmissions, because the paths to some of the multiple receivers will have differing amounts of bandwidth available.

In any case, it would appear that the bandwidth requirements for real-time applications cannot be adjusted for long, unless very large buffers are used to overcome jitter.

As regards delay, it is not clear whether Bolot and Turletti's video transmissions were streamed (that is, tolerant of delay) or two-way (that is, not tolerant of delay, e.g. video conferencing). Delay, of course, would remain, regardless of the use of buffers.

2.8 Priority Queuing implemented in hardware

Whereas priority queuing was traditionally implemented by programming routers, the current generation of Ethernet switches is able to implement the priority queues in hardware. On an Ethernet switch complying with the IEEE 802.1p standard, the ports can be prioritised, so that Ethernet frames entering the switch through a prioritised port will be given priority status. In addition, all ports have high and low-priority output queues. Thus frames exiting the switch through any port can be separated according to user priority and directed to high or low-priority queues.

Frames passing through a switch are subject to minimal latencies, typically a few microseconds. This compares to approximately 200 microseconds for a typical software-based router (Breyer & Riley, 1999, p.194).

The speed advantages of hardware-based priority queuing have –

- allowed priority queuing to be used on a much larger scale than was previously possible
- resulted in priority queuing to be carried out in Layer 2 rather than Layer 3 of the OSI model

In addition, according to Breyer and Riley (1999, p.73) the availability of 801.1p-capable equipment will *"narrow the gap"* so that ATM QoS will have much less of an advantage over Ethernet/IP QoS.

An example of 802.1p-capable equipment is the Bay Networks BayStack 450 Ethernet switch, which is the one used for the experimental work in Chapter 6. As explained in the

manual (BayNetworks, 1998), the switch is programmable and can be configured for IEEE 802.1p prioritisation and IEEE 802.1Q tagging.

2.8.1 The limits of 802.1p prioritisation

IEEE 802.1p prioritisation is a differentiating system and subject to the following limits:

- **IEEE 802.1p prioritisation does not work across WANs.** 802.1p prioritisation operates in Layer 2 of the OSI model and is limited to the local area network, although this may be a Bridged LAN (that is, a LAN consisting of two or more LAN segments joined by bridges or Ethernet switches). The priority information can be transmitted across bridges but not across routers (which operate at Layer 3) nor can it be transmitted into neighbouring Wide Area Networks (WANs). The priority information is lost once it enters a router. On its own, 802.1p prioritisation can only provide QoS for traffic that originated in the local area network.

  To address this problem the Ethernet/IP system relies on another, independent protocol, the Resource Reservation Protocol (RSVP) which operates in Layer 4. RSVP reserves bandwidth between a sender and a receiver by examining each link along the route. The receiver interrogates each network device along the route. If each device has spare bandwidth available the path is established, if not the sender is informed.

  According to Breyer and Riley (1999, p.61) *"the new IEEE 802.1p and IETF RSVP standards go a long way to providing QoS"*. This is because 802.1p prioritisation and RSVP complement each other: 802.1p prioritisation works in LANs, while RSVP can make reservations along WAN links. Thus 802.1p prioritisation, in conjunction with RSVP, should be able to provide end-to-end QoS comparable to that provided by ATM (Breyer & Riley, 1999, p.73).

- **802.1p prioritisation does not provide guaranteed performance.** 801.1p prioritisation uses a priority queuing algorithm. Unlike, say WFQ, priority queuing cannot guarantee bandwidth, much less delay or jitter. It seems that the ISA Controlled Load Service (refer 3.1.1 "Integrated Services Architecture") would be the best that 802.1p prioritisation can provide.

2.8.2 Effect of Load

Bajaj, Breslau and Shenker (1998a), while carrying out simulation experiments to determine whether prioritisation should be used for real-time traffic on the Internet, found that the usefulness of service priority depends, among other factors, on the –

- **Burstiness of the traffic**. *"With smoother background traffic and FIFO service, performance does not degrade until higher levels of utilization are reached. Therefore, the relative benefits of priority service are smaller and only occur at higher levels of utilization, making a weaker case for multiple levels of service."* (p.76)
- **Ratio of best-effort to real-time traffic in the network.** According to Bajaj et al, best-effort applications tend to absorb the delays and distortions caused by the high bandwidth demands of real-time applications.

This work by Bajaj et al. is an indicator that experimental results from LAN tests will depend to a large degree on the type of load that can be provided in the laboratory.

2.9 The New Internet Protocol

Although the current version (v4) of the Internet Protocol has been spectacularly successful, concerns about its adequacy were raised when the Internet growth rate approached the 100% per year growth rate experienced in the early 1990s. The main concern was the 32-bit address space, which would not be able to accommodate all the devices, such as television sets and mobile telephones, that might in future be connected to the Internet. Apart from addressing, other areas singled out for improvement were routing, security and the transmission of priority information.

The IETF started work on a new version of IP in 1990 and called for proposals. "*After much discussion, revision, and jockeying for position, a modified, combined version of the Deering and Francis proposals, by now called SIPP (Simple Internet Protocol Plus) was selected and given the designation IPv6*" (Tanenbaum, 1996, p.438).

The major improvements, as listed by Tanenbaum, are –

1. Source and destination addresses increased from 32 to 128 bits

2. Simplified header that allows routers to process packets faster

3. Better support for options, which speeds up packet processing

4. Security features include authentication and privacy

5. More attention paid to type of service. In particular, *"Packet headers support differentiated classes of service and even identify individual flows so that, for instance, delay-sensitive IP telephony can be handled differently from bulk newsgroup traffic"* (Tebbutt, 1998, p.92).

While IPv6 was under development, IPv4 did not stand still. In fact it was enhanced in two ways:

> *"The first is Network Address Translation (NAT), a technique that allows hundreds of users on a private IP network to share a single Internet address.... Together with Classless Inter-domain Routing (CIDR), a method for simplifying Internet routing, NAT has relieved IP addressing pressures to such an extent that experts now question whether IPv6 is necessary."*(Tebbutt, 1998, p.92)

Doubts about the future of IPv6 are reinforced by the fact that the new DiffServ QoS control architecture uses the IPv4 TOS (Type of Service) field rather than the "priority" and "flow label" fields in the IPv6 header (refer 3.1.2 "Differentiated Services Model").

Whichever IP version predominates, end-to-end Ethernet/IP QoS will depend heavily on the services provided by the Network Layer, and hence on the features built into the Internet Protocol.

2.10 The New Standards

Expectations of ATM's capabilities were increasing rapidly in 1996, but its subsequent development did not match the high expectations. *"For a few years ATM was on the rise, but with the advent of Gigabit Ethernet and numerous other innovations, ATM seems to be very much on the defensive"*(Breyer & Riley, 1999, p.41). According to Odlyzko the reason for ATM's slow acceptance may be the nature of the load:

"ATM was conceived with the idea, inspired by voice and multimedia, that traffic would consist of long-lived flows with reasonably defined bandwidth, latency, and jitter requirements. However, that is not what we have on our networks today. Most of the traffic consists of Web page downloads that are small, and what matters is how quickly the entire page is delivered." (Odlyzko, 1998, p.11)

This slower-than-expected development of ATM has allowed the Internet research community to take the initiative with new developments that enhance the Ethernet/IP system. The Internet Engineering Task Force (IETF) as recently as 1998, has –

1. extended the capabilities of LANs and interconnecting bridges. The new standard ISO/IEC 15802-3 - Media Access Control (MAC) Bridges (ISO/IEC Final DIS 15802-3, 1998), which supersedes ISO/IEC 10038:1993, supports the transmission of time-critical information across bridges from one LAN segment to another.

2. specified a method of creating Virtual Local Area Networks (VLANs) above the physical fabric of bridged LANs. The new IEEE standard 802.1Q – Virtual Bridged Local Area Networks, makes use of the extended LAN Bridging concepts and mechanisms that were introduced by ISO/IEC 15802-3 above.

3. IEEE 802.1Q defines a VLAN frame format for carrying information between VLANs. Significantly, this VLAN frame includes fields for VLAN identification and user priority. Thus user priority information can be carried end-to-end across 802.1Q VLANs, even though the underlying physical LANs (which may be using older frame formats such as 802.2) may not be able to support the transmission of user priority information.

4. standardised the prioritisation of traffic across a LAN. The new IEEE standard 802.1p - Traffic Class Expediting and Dynamic Multicast Filtering, now incorporated into ISO/IEC 15802-3, uses the priority field in the 802.1Q VLAN frame to signal user priority.

2.10.1 ISO/IEC 15802-3

1. The ISO/IEC 15802-3 Environment

In the environment defined by ISO/IEC 15802-3, IEEE 802 LANs are interconnected by MAC bridges to form a Bridged Local Area Network (the terms are defined under "15802-3 Terminology" below).

The Bridged LAN provides a service to the end stations. The service is called the Media Access Control (MAC) Service because it is provided by the MAC sublayer to the Logical Link Control (LLC) sublayer. The MAC Service provided by bridged LANs *"should not be significantly inferior to that provided by a single LAN"* (ISO/IEC Final DIS 15802-3, 1998, p.23). That is, if the MAC bridges comply with 15802-3, they will seamlessly interconnect the LANs so that end stations may be connected to any of the bridged LANs and still receive the same service - the Bridged LAN will behave like a single LAN.

The MAC Service includes the maintenance of Quality of Service. *"The Quality of Service parameters to be considered are those that relate to*

a) Service availability
b) Frame loss
c) Frame misordering
d) Frame duplication
e) The transit delay experienced by frames
f) Frame lifetime
g) The undetected frame error rate
h) Maximum service data unit size supported
i) User priority
j) Throughput" (ISO/IEC Final DIS 15802-3, 1998, p.23).

2. User Priority

"*The MAC Service includes user priority as a Quality of Service parameter....The user priority associated with a frame can be signalled by means of the priority signalling mechanisms inherent in some IEEE 802 LAN MAC types*" (p.26). However, not all IEEE 802 LAN MAC types can signal user priority. For this reason, the 802.1Q VLAN frame was defined. It can be used to carry user priority information regardless of the ability of individual LAN MAC types to signal priority.

The user priority parameter has a range 0 through 7 (3-bit field). *"The Bridge maps the user priority onto one or more traffic classes; Bridges that support more than one traffic class are able to support expedited classes of traffic"* (p.26). Since "Ethernet switches" may be considered multi-port bridges, this is clearly the mechanism that enables current-generation Ethernet switches complying with IEEE 802.1p, to differentiate between frames from different flows.

3. VLANs

Since the designation "bridges", used exclusively in the above standards, also applies to Ethernet switches (a switch is a multi-port bridge), the new standards are highly relevant to switched networks. In fact, the main purpose of the new standards is to define recognised methods for managing switched networks. Network managers who have migrated to a switched environment, can now implement prioritisation by creating VLANs and by using the newly-defined VLAN frame to carry priority information.

To transmit priority information it is necessary to use the VLAN frame. VLANs can be viewed simply as a means of implementing prioritisation, and the VLAN frame as a vehicle for carrying priority information.

As a VLAN can span more than one Ethernet switch, the VLAN frames can carry the user priority information across extensive geographical areas.

4. QoS

ISO/IEC 15802-3 supports QoS. However, being a Multiple Access Control (MAC) sublayer standard, it provides only one link in the chain of QoS support. This is the link between the MAC and Logical Link Control (LLC) sublayers.

RSVP provides QoS support in the Network Layer and should be able to access any service provided at the Data Link Layer, including the LLC and MAC sublayers. Whether RSVP is able to access the MAC Service provided by the MAC sublayer to the LLC sublayer could be a subject for further investigation.

5. 15802-3 Terminology

LAN: In the context of 15802-3, "LAN" stands for "IEEE 802 Local Area Network". "*IEEE 802 LANs (also referred to in the text as LANs) are Local Area Network technologies that provide a MAC Service equivalent to the MAC Service defined in ISO/IEC 15802-1. IEEE LANs include ISO/IEC 8802-3 (CSMA/CD), 8802-4 (Token Bus), ISO/IEC 8802-5 (Token Ring)*" (ISO/IEC Final DIS 15802-3, 1998, p.18).

An "IEEE 802 LAN" may be considered to be a "LAN segment". In this thesis the abbreviation "LAN" will retain its more general meaning, whereas an "IEEE 802 LAN" will be referred to as a "LAN segment".

Bridged Local Area Network: "A concatenation of individual IEEE 802 Local Area Networks interconnected by MAC Bridges" (ISO/IEC Final DIS 15802-3, 1998, p.18).

3.0 THEORETICAL FRAMEWORK

Internet research advances by means of Internet Engineering Task Force (IETF) working groups tasked with solving specific problems. The emphasis is on providing practical solutions within tight time frames rather than building theoretical frameworks. Nevertheless, the IETF has created a number of frameworks for identifying and controlling the development of QoS architectures.

3.1 QoS Control Architectures

Currently, three mainstream methods for implementing QoS in an Ethernet/IP system can be identified.

3.1.1 Integrated Services Architecture

The Integrated Services Architecture (ISA), also known as IntServ, was developed by the IETF to support real-time applications on the Internet. IntServ uses the Resource Reservation Protocol (RSVP) to provide QoS for Ethernet/IP networks. The following service categories are defined:

- **Guaranteed Service**. Allows the user to specify a maximum delay. This service is analogous to ATM's Constant Bit Rate (CBR). Suitable for real-time applications. Application controls include leaky bucket and Weighted Fair Queuing (WFQ).

- **Controlled Load Service**. Suitable for congestion sensitive applications. Application control uses leaky bucket.

- **Best-effort Service**. No differentiation (what is currently available from the Internet) (Minoli & Schmidt, 1999, p.409).

With IntServ every router along the path needs to examine every packet to determine the service category and to allocate the required resources. This means additional software is required on every router.

In addition, IntServ requires each router to process and maintain data on different flows. This makes IntServ a "stateful" architecture, which is less desirable than a "stateless" architecture. The need to store this data makes IntServ less scalable. For example, if there are identical flows emerging from one source (as in multicasting) then IntServ treats every flow individually, making the process inefficient.

Because of problems of this nature, IntServ's development has slowed and attention is focussing on alternative architectures.

3.1.2 Differentiated Services Model

Differentiated Services (DiffServ) is a more recent QoS control architecture than IntServ. It aims to overcome the problems with the IntServ.

The DiffServ mechanism differs from the IntServ mechanism in that DiffServ does not attempt to keep track of all flows at every node. Instead it uses a field in the packet header to specify how a particular router should treat a particular packet. That is, packets are routed on a per-hop basis rather than a per-flow basis.

A packet needs to carry enough information to specify the subroutines that need to be activated when the packet arrives at a router. This information is referred to as the Per-Hop Behaviour (PHB).

To specify PHB, DiffServ requires a 6-bit field known as Differentiated Services Code Point (DSCP) (Carpenter & Kandlur, 1999). DiffServ can use the 8-bit Traffic Class field in the IPv6 header. However, it can also use the under-utilised, 8-bit Type of Service (TOS) field in the IPv4 header. This has implications for the future acceptance of IPv6 (refer "The New Internet Protocol" below).

The difference between IntServ and DiffServ is expressed in broad terms by Wroclawski:

> *"Internet's RSVP and Integrated Services work, was originally driven by the perceived needs of real-time and multimedia applications. More recently, though, Internet researchers have taken a broader focus, looking for mechanisms that gracefully combine support for new applications, higher levels of assurance for traditional applications, and traffic allocation and management capabilities for network operators"* (Wroclawski, 1999, p.32)

Wroclawski points to one specific differences between IntServ and DiffServ:

> *"In traditional QoS architectures...each RSVP router or ATM switch in a net is expected to independently admit, classify, and schedule the packets in a single session or connection request. In contrast, the DiffServ model distributes these functions across a domain, which is typically a single provider's network"* (Wroclawski, 1999, p.32).

3.1.3 MultiProtocol Label Switching (MPLS)

MPLS avoids the situation where a packet header needs to be examined by every router along the route. With MPlS, only the routers at the entry point to an MPLS-enabled network (the edge routers) need to read the packet headers. The routers at the core of an MPLS-enabled network (the core routers) only read a 32-bit label attached to the packet.

A core router reads a label and decides on the next hop. It then creates a new label for the packet (switches the label). The edge routers, on the other hand, either attach a label when the packet enters an MPLS-enabled network, or permanently remove the label when the packet exits.

Apart from saving time by reading a short label instead of a header, the used labels can be stored to provide routing information for the packets that follow (Dutta-Roy, 2000).

3.1.4 Summary

IntServ identifies what needs to be done. DiffServ and MPLS aim to improve on IntServ by reducing the packet processing time. DiffServ aims to simplify a router's task by absolving it from having to store "per-flow" data. MPLS saves processing time by making it possible for routers to read (short) labels rather than (long) headers.

3.2 The OSI Model

The Open Systems Interconnection Reference Model (OSI model) was developed by the International Standards Organisation (ISO) as a reference frame for developing and integrating protocols. The OSI model has not been defined at the higher levels (Layers 5 and 6) to any great detail. It appears that the model requires further development at the Session and Presentation Layers. This contributes to it not having been implemented at these levels to any significant extent.

An alternative model is the United States Department of Defence Model (DoD Model), also known as the TCP/IP Reference Model. This model has been widely used by Internet developers, and all of the popular Internet protocols, such as TCP, IP, ftp, etc. were defined within the framework of the DoD model.

For the purposes of this thesis, however, the DoD model is not a suitable framework The main reason is that it does not distinguish between Data Link functions and physical functions, that is, it does not have layers equivalent to the OSI Data Link Layer and OSI Physical Layer. In addition, the DoD model does not have any protocols defined in its lowest layer, the Host-to-Network Layer.

In this thesis, the OSI model is used as theoretical framework for investigating bandwidth management methods, standards, protocols and devices. All protocols, including the DoD-based Internet protocols, are viewed, and all analysis is carried out in the context of the OSI model. For example, bandwidth management methods are categorised as Layer 1 Bandwidth Management, Layer 2 Bandwidth Management, etc, as appropriate.

4.0 CASE STUDY

4.1 Introduction

4.1.1 Why the Case Study?

The case study of the ECU data network illustrates current practices in network management, including bandwidth management. The ECU data network is a large, diverse network spread over geographically separated campuses. It provides actual examples of some of the methods, problems and solutions identified in the literary review.

In addition, the case study helped this researcher to –

- gain an overview of the ECU data network from both technical and administrative points of view.
- acquire some practical knowledge required to realistically evaluate the bandwidth management methods identified in the Review of Literature
- develop an awareness of network managers' current priorities and concerns
- find out which technologies are currently in use.

4.1.2 Procedure

Data for the case study was gleaned from interviews of IT staff involved in the planning, management and maintenance of the network. Staff that were interviewed included the –

- Campus IT Coordinator, Operations and Systems Programming Branch. This branch administrates the University's networks. The IT Coordinator is responsible for IT Operations at three of the campuses.
- Manager, Communications Services. This branch is responsible for the University's WAN and backbone networks.

- Technical Consultant, Communications Services – provides specialist input into the planning, upgrading and maintenance of the WAN and backbone networks.
- Manager, School of Computer and Information Science (SCIS) support group. This group is responsible for the local area networks at SCIS.
- IT Officers at SCIS - maintain LAN hardware and software.
- User Support Officers - provide desktop support at the Mt Lawley campus.

Permission to interview IT staff had previously been obtained from the Head of School, Computer and Information Science, and from the Manager, Communications Services. No ethics committee clearance was needed.

The interviews were in the form of a series of open-ended questions. No tape recordings of interviews were made, although notes were taken. In some cases the interviewing process was iterative, that is, follow-up questions were put to some interviewees.

4.2 Case Study Details

The study of the Edith Cowan University (ECU) data network was carried out in April 1999. During the remainder of that year the network was subject to extensive upgrading and the Case Study details were updated as the changes became known. However, as upgrading is a continual process, the documentation that follows may not reflect all the changes.

4.2.1 Perth Academic Research Network

The Perth Academic Research Network (PARNet) caters for Perth's four publicly funded universities and the Commonwealth Scientific Investigations and Research Organisation (CSIRO). A major upgrade of PARNet began in 1995. As a result, the four universities and the CSIRO are now connected to PARNet by means of 34 Mbits/sec microwave links radiating from the PARNet state hub at the BankWest tower, which is located in the central business district.

4.2.2 Edith Cowan University Network

In Perth's central business district, and in the built-up areas of Perth, microwave links are the most practical communications medium because they avoid excavation of cable trenches. Microwave links are therefore used to interconnect WANs in the Perth metropolis.

Microwave links require an unobstructed line of sight line, as can be provided at the top of the BankWest tower, one of Perth's tallest buildings. The PARNet State hub is located at the BankWest tower, and so is ECU's main hub. In addition, ECU's internal traffic, that is, the traffic between its main campuses, is also propagated via 34 Mbits/sec microwave links radiating from the BankWest tower. Thus the BankWest tower not only serves as the PARNet State hub, but also as a relay station between ECU's main campuses in the suburbs of Churchlands, Mt Lawley and Joondalup.

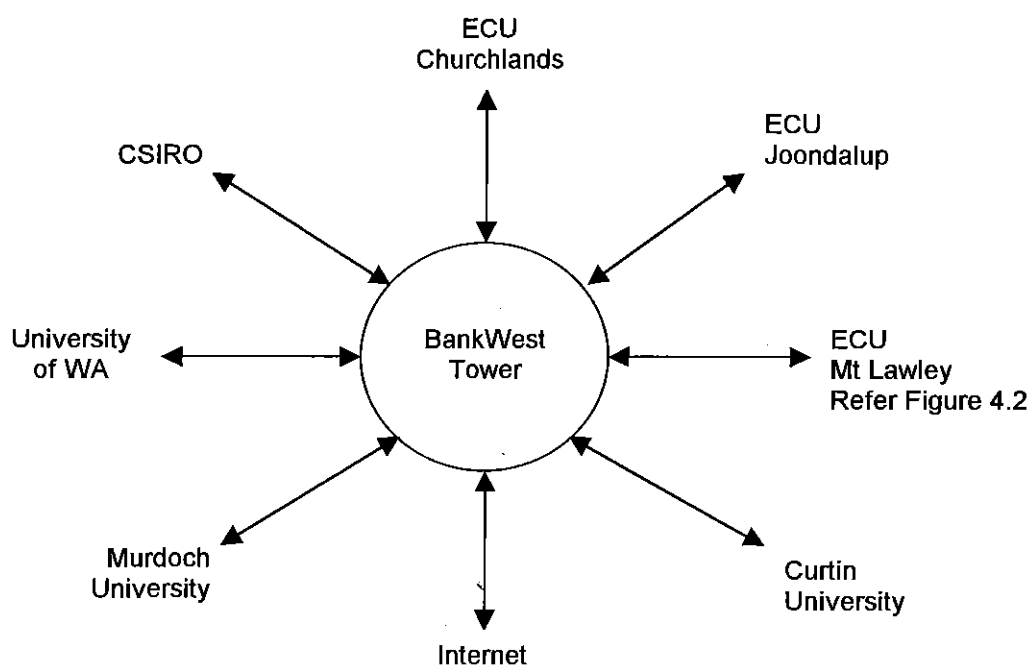

Figure 4.1 - BankWest Tower's Dual Role

State hub and relay station between ECU campuses

The links between ECU's main campuses and the BankWest tower comprise the following channels –

- 1 x 34Mbits/sec data channel using ATM
- 16 x 2Mbits/sec channels carrying various traffic including telephones. Some of the 16 channels are not used and function as spares.

At the time of data collection, plans existed for increasing the microwave link capacity from 34 Mbits/sec to 1 Gbit/sec, to make provision for video and audio streaming.

4.2.3 Extent of ATM

The microwave links use ATM technology. However, ATM is limited to the PARNet backbone and to the WAN links between ECU's main campuses. The campus networks themselves are Ethernet/IP networks.

The interfaces between the ATM backbone and the Ethernet/IP campus networks are ATM switches (Centillion C100) located at the BankWest tower and at each of the main campuses. The ATM switches, including those at the BankWest tower, were installed and are managed, by the ECU Communications Branch.

4.2.4 Campus Network Structure

An ECU campus network consists of many Ethernet LANs. A typical LAN will connect the computers in a building, or the computers on a particular floor, or in a particular wing, of the larger buildings. The LANs themselves are interconnected by fibre cable installed between buildings and between the floors and/or wings of the same building.

Within the campuses, conventional Ethernet hubs have been progressively replaced with Ethernet switches. One significant effect of Ethernet switches is to collapse the backbones and to flatten the networks. At Mount Lawley, for example, the LANs on the campus are now directly connected to the Compatible Systems 4000R router – refer Figure 4.2. That is, the campus backbone has been collapsed to the extent where it can be viewed as consisting of the Compatible Systems 4000R router, a Bay Networks Backbone Link Node (BLN) router and an Accelar 1000 routing switch.

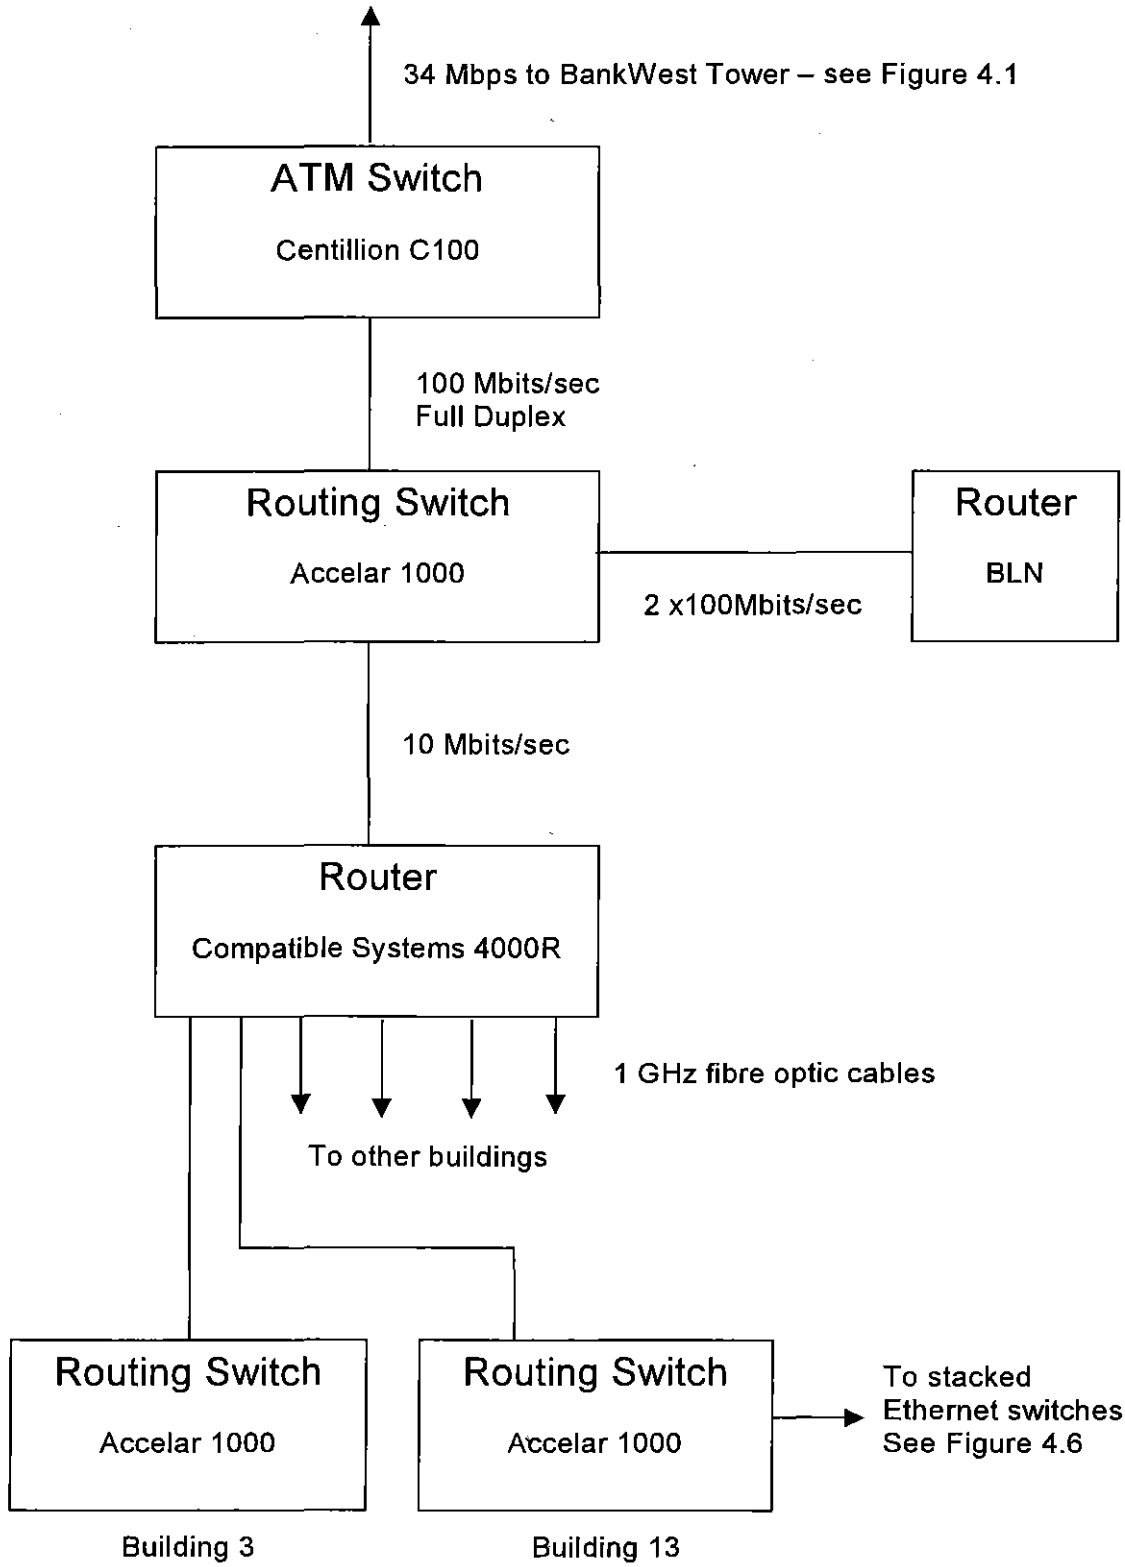

Figure 4.2 - Campus Backbone Network

Mt Lawley Campus, Edith Cowan University

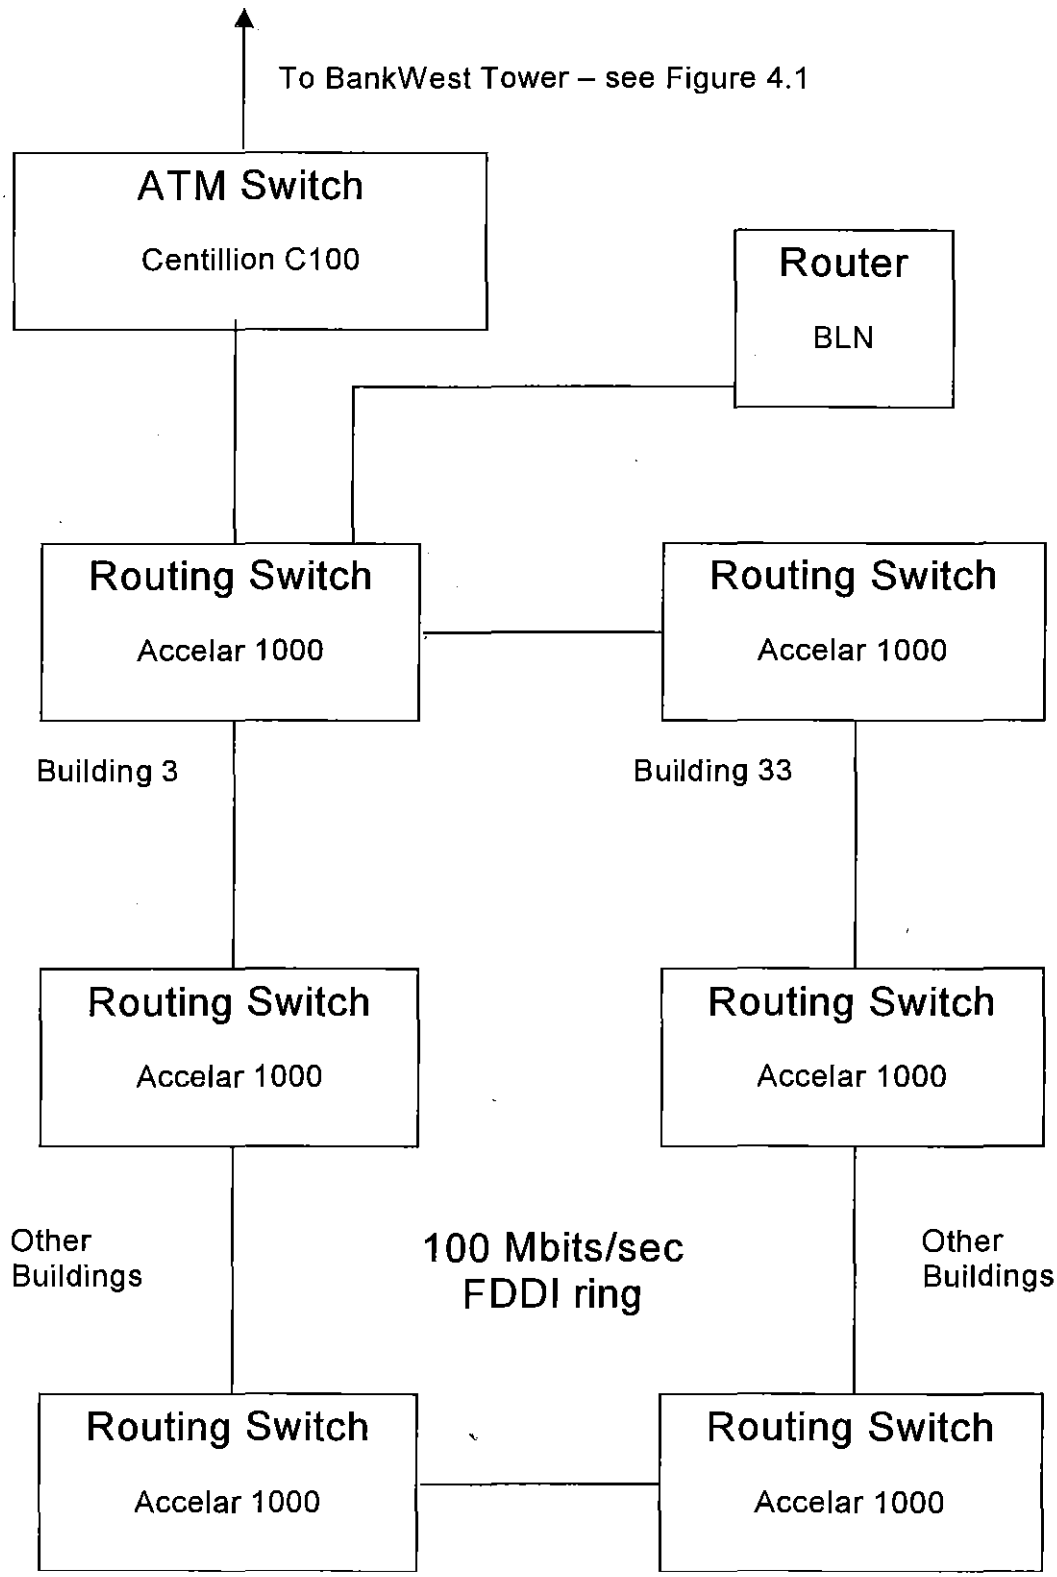

Figure 4.3 - Planned Campus Backbone Network

Mt Lawley Campus, Edith Cowan University

With a collapsed backbone all inter-LAN traffic within a campus is channelled to the routers on the backbone, making the backbone routers a single point of failure. The planned Fibre Distributed Data Interface (FDDI) will take the pressure of the backbone routers and provide higher reliability - refer Figure 4.3.

Security Domains

Traditionally, the ECU data network has been administered as three separate security domains. Hosts attached to these domains have increasing access rights in the following order:

- Student
- Academic Staff
- Administrative Staff

Apart from providing differing levels of security, these sub-networks facilitate broadcasting to selected groups of users. In the past, the requirement for separate security domains resulted in three physically separate sub-networks within each campus.

Implementing Security Domains on a Collapsed-backbone Network

As mentioned above, the introduction of switches has led to collapsed backbones and hence, physically flat networks. Although the campus backbone networks are now flat, the ECU network can still be administered as three logically separate security domains. This is due to the availability of switches with VLAN capability, which allow different types of users to be segregated into different logical LANs.

The sub-networks, which started out as separate physical networks, with staff and students connected to different hubs, are in the process of being converted to Virtual LANs. For example, at Mt Lawley campus, the Student sub-network now consists of a combination of seven VLANs each with a Class C address such as 139.230.35.0 and 139.230.42.0. This means that staff and student computers can be connected to the same Ethernet switch and still be on separate sub-networks.

At the time of writing, the security sub-networks are implemented as VLANS at campus backbone level and as physical sub-networks at the lower levels. Figures 4.5 and 4.6 show the lower levels where the Academic and Student sub-networks are still physically separate.

4.2.5 Operating Environment

The LANs on ECU campuses are required to cater for a variety of computers, including IBM-compatible PCs, Macintosh PCs and UNIX machines. This diversity necessitates the use of multiple protocols and frame types. Thus -

- UNIX machines, and IP applications generally, are using the TCP/IP protocol stack and the Ethernet II frame.
- IBM PCs communicating under NetWare 4.11 (the most widely used Network Operating System (NOS) on ECU campuses) will default to the SPX/IPX protocol stack and the IEEE802.2 frame.
- IBM PCs communicating under legacy NetWare 3.x are using the SPX/IPX protocol stack and the IEEE802.3 frame.
- Macintosh PCs under Appletalk are using the IEEE 802.3 frame with the SNAP extension.

Plans are under way to simplify network operation by limiting future Layer 3 protocols to IP and the frame type to Ethernet II - refer 4.3.7 "Future Trends".

4.2.6 Routing

The Accelar 1000 routing switch and the BLN router (refer Figure 4.2) together perform the routing on the backbone. The routing switch, being hardware-based, can route packets at "wire speed". It has been configured to route the IP packets arriving at the backbone.

Non-IP packets are routed by the software-based, multi-protocol BLN router. The routing of non-IP packets is therefore a slower process. The alternative method of

encapsulating non-IP packets ("tunnelling") entails additional overhead and is also not optimum. The elimination of non-IP traffic along the backbones is network management's stated aim.

The BLN router provides support for LAN emulation (LANE), which allows Ethernet/IP networks on different campuses to communicate across the PARNet ATM backbone.

Apart from their use in backbone networks to complement conventional routers, Accelar routing switches are also being added to LANs, where they can be configured to independently route (at Layer 3) and switch (at Layer 2).

4.3 Discussion

4.3.1 Bandwidth Management

Methods of bandwidth management used at ECU include network segmentation using multiple network interface cards (NICs) and network segmentation using Ethernet switches (micro-segmentation). The two methods are used in the Academic and Student sub-networks at the School of Computer and Information Systems (SCIS) – refer Figures 4.4 and 4.6 respectively.

Effect of Micro-segmentation

(a) Academic Network at SCIS

At the time of data collection, the Academic network had not been upgraded and was still using Ethernet hubs.

At the Main Rack, four stacked hubs were needed to accommodate all the users on the Academic sub-network – see Figure 4.5.

Bus bandwidth = 10 Mbits/sec

Maximum practicable utilisation = 70%, estimated

Maximum number of users = 64

User's share of the bandwidth

= 10 Mbits/sec x 0.7 / 64

= 109 Kbits/sec

(b) Student Network at SCIS

The Student network at SCIS had been upgraded and the original Ethernet hubs had been replaced with Ethernet switches – refer Figure 4.6.

Port capacity = 100 Mbits/sec

Maximum practicable utilisation = 100%

Maximum number of users = 96

Backplane capacity = 2.56 Gbits/sec

Provided the aggregate bandwidth of all 96 ports does not exceed 2.56 Gbits/sec (which is the likely scenario), a student's share of bandwidth is equal to the port capacity, that is –

User's share of the bandwidth = 100 Mbits/sec.

A comparison of the bandwidth available to Academics and Students shows the marked improvement that micro-segmentation has on bandwidth availability. Providing additional bandwidth to all students, without differentiating between students or applications, corresponds to the "bigger pipe" approach – refer Section 2.4.2 'The "Bigger pipe" Approach'.

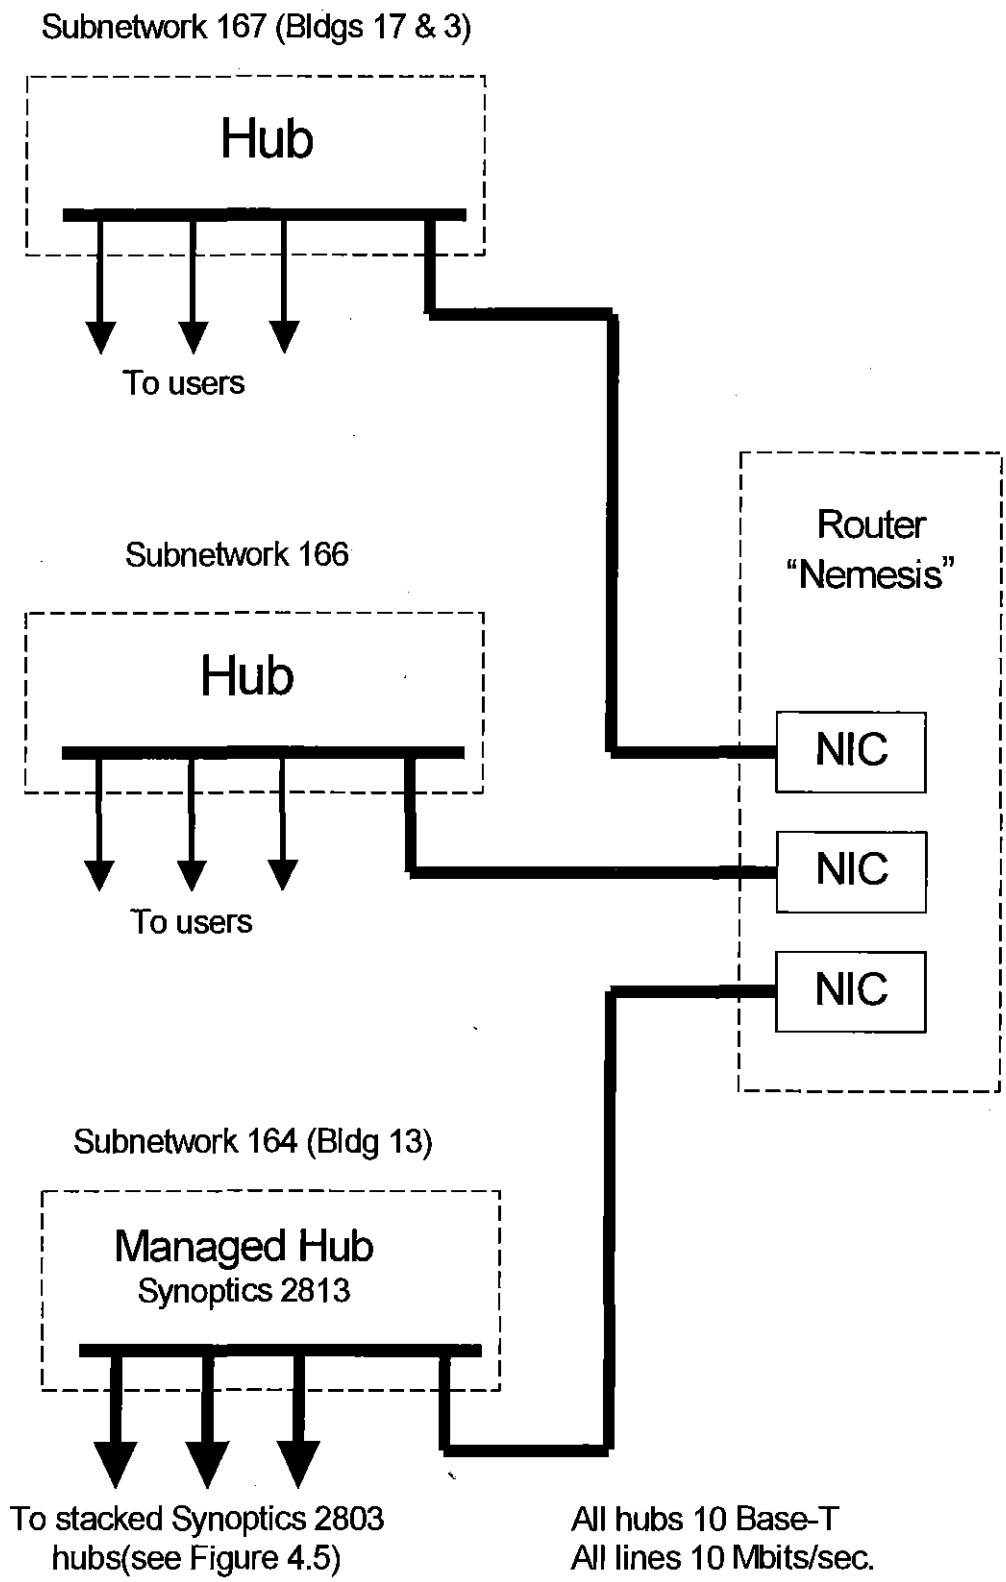

Figure 4.4 - Network segmentation using multiple NICs
Academic Sub-network at SCIS, Main Rack, Building 13

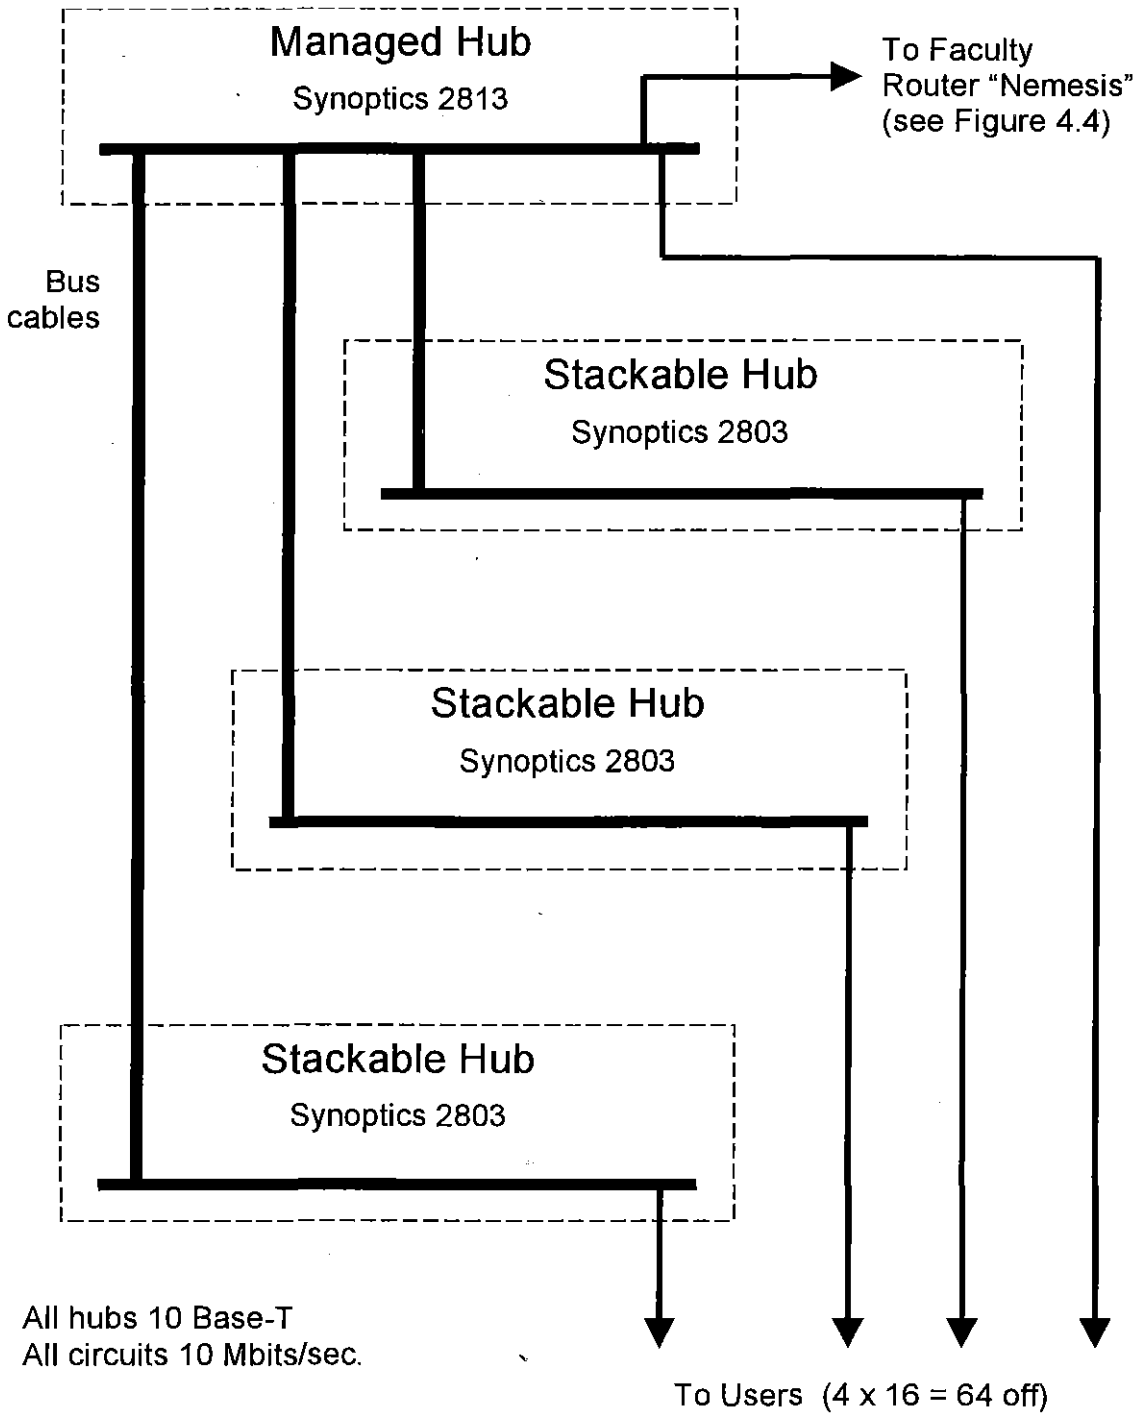

Figure 4.5 - Stacked Ethernet Hubs

Academic Sub-network at SCIS, Main Rack, Building 13

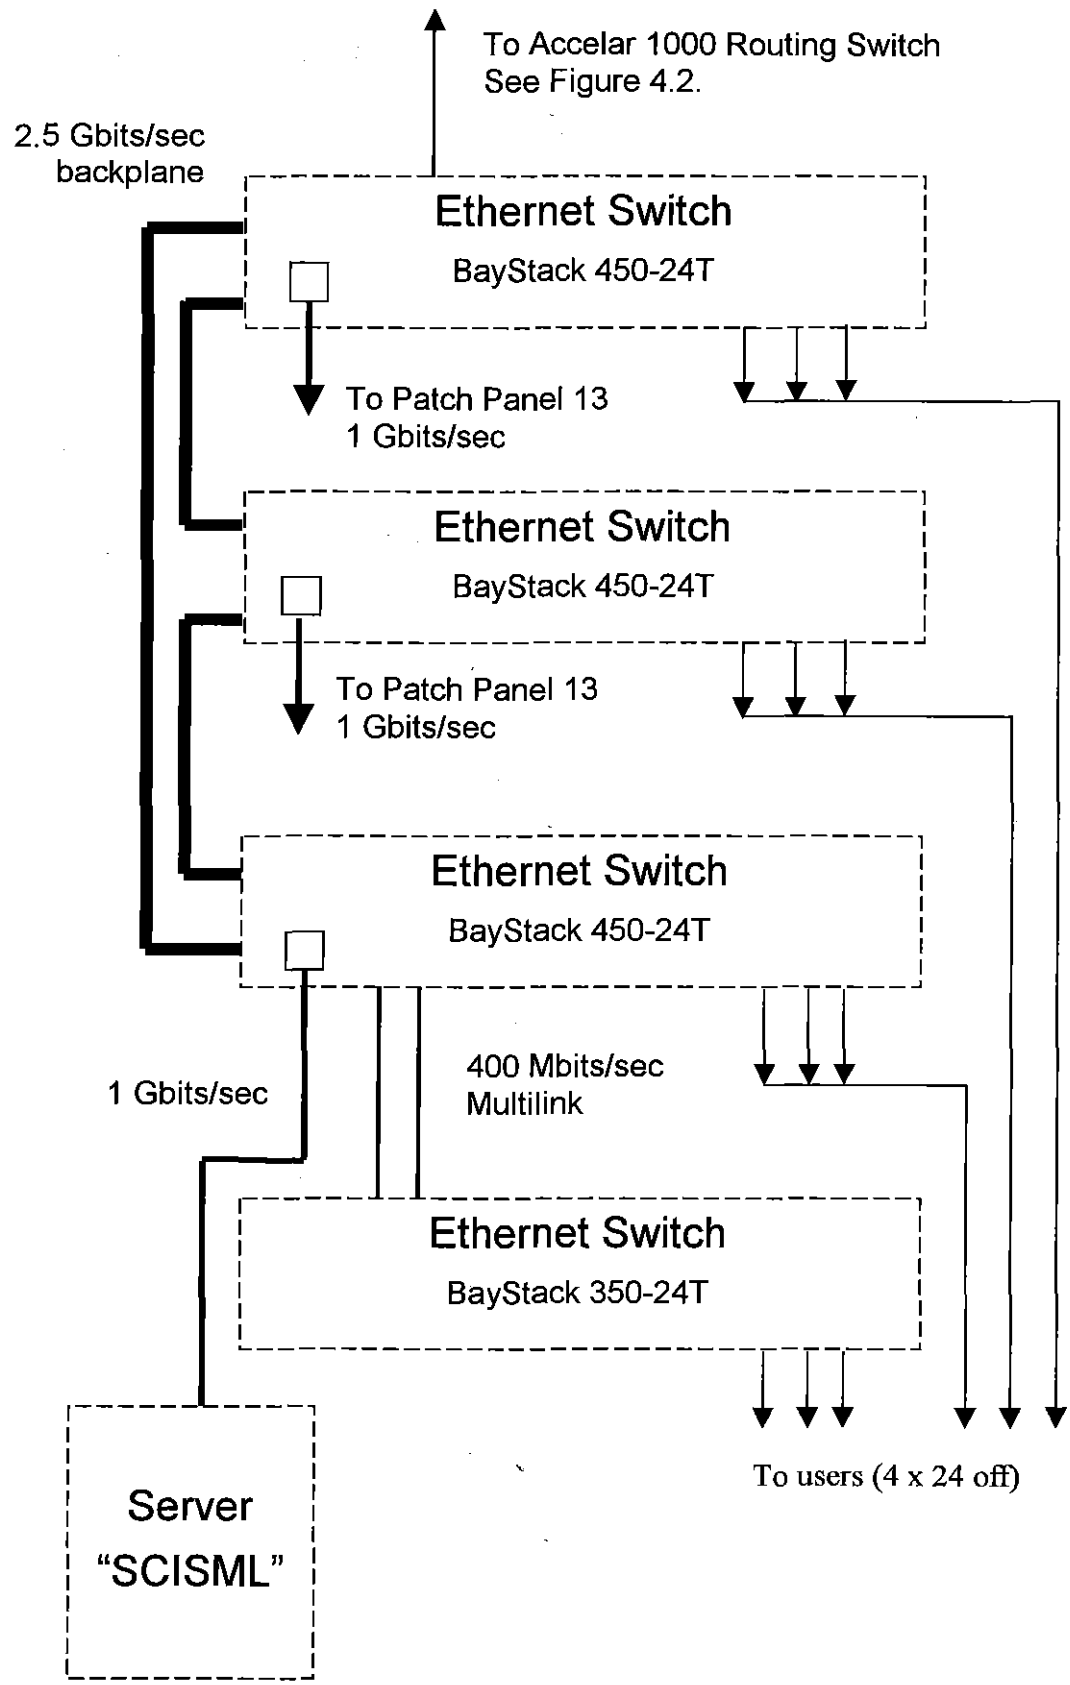

Figure 4.6 - Stacked Ethernet Switches
Student Sub-network at SCIS, Main Rack, Building 13

4.3.2 QoS vs High-bandwidth Approach

Section 4.3.1 demonstrates the effectiveness of Ethernet switches in providing additional bandwidth to LAN users. Ethernet switches increase the bandwidth available at a port. This means more bandwidth for all users, in keeping with the recommendations of Odlyzko (1998). For the short term at least, the use of Ethernet switches avoids the need for QoS at LAN level. Ethernet switches are providing LAN managers with a breathing space during which they can avoid the extra complications and overheads of QoS.

It is likely that the pressure for QoS will be felt most at WAN level, where the cost of bandwidth is high. The WAN links between ECU's main campuses carry both voice and data on single, shared channels. At WAN level too, there is no separate charging for different applications.

Within a campus, the traffic is split into voice and data, and separate telephone (analogue) and data (digital) networks are maintained. However, even after voice traffic is separated out, the data networks still carry data from diverse applications, including multimedia, web browsers and file transfers, over single, shared channels. It is the scenario envisaged by Huitema (1997), a scenario that lends itself to buying bandwidth in bulk, using simple Service Level Agreements (SLAs), if any.

However, this situation is not likely to last. The pressure for QoS is mounting. A recent ECU "Position Paper" on Communications (ProjectPlanningTeam, 1999) states: "*The future ECU network must provide a number of levels of QoS to enable a number of services to be delivered. Data is well catered for but video must be enable (sic) on the network at the earliest opportunity. Voice integration will need be done, but in a later timeframe.*"

4.3.3 Virtual LANs

The advent of switches at ECU has resulted in some originally physical sub-networks being implemented as VLANs. While the transition from Ethernet hubs to Ethernet switches is incomplete, VLANs are implemented at the higher levels, only, of the campus network, because legacy Ethernet hubs are still in use at the lower levels.

At ECU, VLANS are used mainly for security reasons. However, the VLANS also help to conserve bandwidth by reducing the bandwidth used by broadcasts.

4.3.4 Prioritisation

Prioritisation is still at the planning stage at ECU. The preference is for prioritisation based on the type of application rather on type of user (ProjectPlanningTeam, 1999). This would mean that the security-based VLANs already implemented, would not be used to distinguish between student, academic and administrative users.

Application-based prioritisation would entail a more extensive use of VLANs, particularly protocol-based VLANs. Since particular applications often use a specific protocol (Breyer & Riley, 1999, p.174), protocol-based VLANs could be used to distinguish between different types of applications. The throughput on a VLAN handling high-priority applications could then be enhanced by, say, using dedicated servers.

It should also be noted that the BayStack 450 switches used predominantly at SCIS only support port-based VLANs, not protocol-based VLANs. Protocol-based prioritisation could entail the cost of replacing switches.

4.3.5 Local Area Network Emulation (LANE)

The ECU data network is a hybrid network that illustrates how ATM and Ethernet/IP can coexist. As described under 4.2.1 "Perth Academic Research Network" above, ECU's WAN links operate under ATM, whereas the LANs at the campuses use Ethernet and the IP protocol. ECU's ATM network provides IP support in the form of LANE, which allows the Ethernet/IP networks on either side of the ATM network to communicate with each other.

As pointed out in the Review of Literature, LANE by itself does not support QoS. To implement an end-to-end QoS scheme across ECU's campuses would require interaction between QoS architectures as discussed by (Wolf, 1999) – refer 2.3.5 "Types of QoS".

4.3.6 Networking Components

As explained in Section 3.2 "The OSI Model", the OSI model is used as the theoretical framework for this investigation. Table 4.1 lists some of the networking components featured in the Case Study, and in this thesis generally, and positions them according to their functions within the OSI framework.

Table 4.1 - Networking Components and their Functions

OSI Layer	Networking Component
4	RSVP TCP
3	IP, ICMP routers routing switches VLANs (IP-address based)
2	VLANs (port-based) 802.1p prioritisation FDDI Ethernet switches Ethernet hubs NICs (framing)
1	NICs (encoding, eg Manchester) UTP & fibre cables

Networking Component	ATM Layer
LANE ATM switches (assemble packets from cells & disassemble)	ATM Adaptation
ATM switches (assemble cells from bits & disassemble)	ATM
microwave links	Physical

4.3.7 Future Trends

Network planners at ECU can consider three possibilities:

1. migrate to ATM (ATM to the desktop)
2. retain ECU's ATM/Ethernet hybrid network
3. end-to-end Ethernet/IP QoS (DiffServ, MPLS).

At the LAN level, Ethernet/IP dominates and is becoming even more entrenched: *"In the near term (next 18 months), IP and IPX will need to be support (sic) across all networks, but using only the Ethernet II frame type. In the longer term, once all legacy equipment has been upgraded or decommissioned, only IP will be used on the network."* (ProjectPlanningTeam, 1999). It is therefore unlikely that ATM will replace Ethernet/IP at the LAN level in the foreseeable future, ruling out a complete migration to ATM.

At the WAN level, on the other hand, ATM is well established and proven. In the first instance, therefore, retaining the hybrid system would be the preferred option for ECU. The chances of ECU being able to retain its hybrid network will depend on how effectively, and how timely, Wolf's type of interaction (Wolf, 1999) can be established between ATM QoS and Ethernet/IP QoS. If the two QoS systems cannot be integrated, then recent developments in DiffServ and MPLS might lead to solutions that exclude ATM altogether.

5.0 BANDWIDTH MANAGEMENT

This chapter identifies, categorises and discusses methods of bandwidth management[3] with a view to addressing the first research objective: Investigate bandwidth management methods.

Bandwidth management methods can be categorised in accordance with the OSI model, as illustrated in Table 5.1.

Table 5.1 - Summary of Bandwidth Management Methods

OSI Layer	Bandwidth Management Method
4	Admission Control based on Circuits TCP Rate Control RSVP
3	Load-shedding Admission Control based on Packets Buffer Allocation Flow Control using Choke Packets Network Segmentation using Routers Network Segmentation using VLANs Queuing of Packets
2	Network Segmentation using Bridges Network Segmentation using multiple NICs Network Segmentation using Switches Full-duplex Ethernet 802.1p Prioritisation
1	Upgrade NICs Upgrade cable

[3] For the purposes of this thesis "bandwidth management" is taken to include all hardware and software methods of controlling and allocating bandwidth.

5.1 Layer 1 Bandwidth Management

At OSI Layer 1, the Physical Layer, bandwidth is managed by choosing appropriate transmission media and networking equipment. Higher bandwidth is achieved by installing additional hardware capacity, as illustrated in the following two examples.

5.1.1 Replacing copper with fibre cable

Category 5, Unshielded Twisted Pair (UTP) copper cable has been used extensively to connect PCs to wiring closets in commercial buildings. Cable installed between wiring closets and telecommunications outlets is denoted "horizontal cable" in AS 3080 (AS/NZS3080, 1996). AS 3080 limits the length of a "horizontal cable" to a maximum of 100m. Category 5, UTP cable is officially rated at 100 MHz, which makes it suitable for Fast Ethernet.

The UTP cable could be replaced with 62.5/125 μm multimode fibre optic cable. With the current generation of optical couplers, fibre cables can be operated economically at 1GHz. Therefore where copper cable is replaced with fibre cable, the available bandwidth immediately increases from 100 MHz to 1GHz.

Furthermore, 1GHz is low compared to fibre's much higher inherent bandwidth – 30,000 GHz is currently considered achievable (Tanenbaum, 1996, p.87). Thus fibre cable has the added advantage of potential operation at even higher frequencies once economic optical couplers are available.

Note, however, that the length of the fibre cable, when used as horizontal cable, would still be limited to 100m, because AS 3080 stipulates the same 100m limit for all types of horizontal cables.

Also note that, although Category 5e UTP cable, which is suitable for Gigabit Ethernet, is now available (Spurgeon, 2000, p.213), the Case Study suggests that 1GHz fibre is still the most cost-effective 1GHz medium.

5.1.2 Upgrading Ethernet hubs and network interface cards

Replacing 10 Mbits/sec Ethernet hubs and network interface cards (NICs) with 100 Mbits/sec Ethernet hubs and NICs provides a tenfold increase in bandwidth.

Practical considerations: Installing additional capacity can entail substantial capital works such as the purchase and installation of cabling inside buildings and across the campus. Apart from the cost, the issues that need to be considered when evaluating its effectiveness include the -

- **Managerial overhead**. It is unlikely that this type of work can be done in-house. The engagement and supervision of external contractors requires managerial time and effort.

- **Loss of control**. The installation of external cabling can entail interfacing with other LANs not under the control of the local network manager. The network manager having chosen to upgrade network capacity may find himself moving into areas where he no longer has full control.

Summary:

1. Bandwidth management at Layer 1 amounts to facilitating the flow of bits along cable segments and across the transceiver circuits on the NICs.

2. Bandwidth management at Layer 1 is akin to providing an undifferentiated service. It corresponds to the "Bigger pipe" Approach described in the Review of Literature (refer 2.4.2 'The "Bigger pipe" Approach').

5.2 Layer 2 Bandwidth Management

Bandwidth can be managed at OSI Layer 2 where frames from different sources are contending for access to the shared Ethernet bus. By reducing contention in the collision domain, each device connected to the Ethernet bus receives a larger share of the available bandwidth. Contention can be reduced by -

1. network segmentation
2. implementing full-duplex links
3. prioritising frames.

5.2.1 Network Segmentation

As described by Fitzgerald (Fitzgerald & Denis, 1996, p274), network segmentation entails the breaking up of the network into smaller parts: "*By carefully identifying how much each computer contributes to the demand on the server, and carefully spreading those computers to different network segments, the network bottleneck can often be broken*".

1. Traditional Methods of Network Segmentation

Traditional methods of increasing the bandwidth available to users rely on network segmentation using bridges and multiple network interface cards (NICs). From a technical point of view these two methods are relatively simple, but they do require additional hardware and modifications to existing networking equipment and cabling.

- **Network Segmentation using Bridges**: The Ethernet bus is divided into multiple segments joined by bridges. Because of the increased number of segments, new hubs would be required in addition to the bridges.
- **Network segmentation using multiple Network Interface Cards**: The Ethernet bus is again divided into multiple segments. Each segment is connected to separate network interface card (NIC) on the server. This means the server will be required to route between the NICs.

2. Network segmentation using Switches

This is a relatively new method of network segmentation, in which Ethernet hubs are replaced with Ethernet switches.

Since each port on a switch represents a collision domain, adding a switch to a network is equivalent to breaking up the network into as many smaller parts as there are ports. The creation of such large numbers of small segments has been referred to as micro-segmentation.

The dramatic effect of micro-segmentation is illustrated by the Edith Cowan University Case Study: the Academic sub-network uses hubs and can provide an average of 109 Kbits/sec per users, whereas the Student sub-network, which uses switches, can normally provide 100 Mbits/sec per user (refer 4.3.1 "Bandwidth Management").

Ethernet switches, being current-technology devices, have other advantages over Ethernet hubs, such as low latency and support for VLANs and prioritisation, so that micro-segmentation is usually associated with these other benefits.

Practical considerations: Network segmentation, although requiring additional hardware (hubs, NICs and jumper cables), does not entail major outlays. The installation of networking hardware is not as labour-intensive as the installation of permanent cabling. The cost is likely to be less than the cost of installing long cable runs and more likely to be within the network manager's operating budget. Managerial overhead and loss of control are unlikely to be issues (as they are for Layer 1 bandwidth management).

5.2.2 Full-duplex Ethernet

Full-duplex Ethernet provides almost twice the bandwidth of traditional Ethernet networks, and a dedicated, full-duplex connection eliminates contention altogether.

Practical considerations: The NICs must support full-duplex mode. Likely to entail NIC replacement costs. Although the cost of NICs has dropped markedly, the labour costs due to re-configuring client and server software can be substantial.

5.2.3 IEEE 802.1p Prioritisation

802.1p prioritisation is a method of queuing. It regulates and shapes traffic and can therefore be considered a bandwidth management method. In particular, it is classified as a Layer 2 bandwidth management method because it involves the queuing of frames.

As explained in the Review of Literature (refer 2.8 "Priority Queuing implemented in hardware"), 802.1p prioritisation is traditional priority queuing implemented in hardware. The details of 802.1p prioritisation are explored in Chapter 6.

Summary

1. Bandwidth management at Layer 2 amounts to facilitating the flow of frames between hubs, switches, clients and servers.

2. The advent of 802.1p-capable switches means that it is now possible to differentiate between types of frames. Chapter 6 describes tests aimed at evaluating this capability.

5.3 Layer 3 Bandwidth Management

At Layer 3, bandwidth is managed by avoiding congestion. Congestion is caused by too many packets traversing a network at a given time. As explained in the Review of Literature (refer 2.5 "Congestion Control"), congestion has a snow-balling effect which wastes whatever bandwidth is still available.

The number of packets in a network can be controlled at Layer 3 by using routers. Routers control the rate at which packets are injected into the network. At Layer 3, congestion can be reduced by the following methods:

1. Load-shedding
2. Admission control
3. Buffer Allocation
4. Flow Control using Choke Packets
5. Network segmentation using Routers
6. Network segmentation using Virtual LANs
7. Queuing

5.3.1 Load-shedding

When a router cannot handle all arriving packets, it discards the excessive packets. Packets may be simply discarded, or, they may be discarded in accordance with a priority scheme. Priority dropping schemes are difficult to implement because they require cooperation from all users (Tanenbaum, 1996, p.392). In addition, Bajaj et al. (1998b) found that priority dropping was not as effective as may be expected.

On the other hand, some type of discrimination between certain types of packets, e.g. data packets and acknowledgment packets, would seem almost essential. For example,

discarding TCP acknowledgment packets would be counterproductive and cause more congestion because it would further delay the release of buffers (Farrell, 1996).

5.3.2 Admission Control based on Packets

In this scheme, a router needs to acquire a permit before it can transmit a packet. Permits are granted only if the number of packets in the network is below a pre-set limit. This method requires continual monitoring of network statistics.

5.3.3 Buffer Allocation

In a connection-oriented network buffers can be allocated at each intermediate router, once a virtual circuit is set up. Should congestion occur while the packets are in transit, the packets can then be stored in the pre-allocated buffers. Congestion is relieved because the stored packets are no longer in transit and have in effect been taken out of the network.

5.3.4 Flow Control using Choke Packets

When a router receives too many packets, it sends a choke packet to the source. The source then reduces the flow of packets.

The choke packet can be send all the way to the source (end-to-end), in which case the response may be too slow. Alternatively, the choke packet may also be send to the intermediate routers (hop-by-hop), in which case the response will be faster, provided the intermediate routers have sufficient buffer capacity to store the delayed packets.

The Internet Control Message Protocol (ICMP) can be used to send choke packets. ICMP is defined in Layer 3.

Flow Control versus Congestion Control: Flow control differs from congestion control in that it regulates traffic between a sender and a receiver, whereas congestion control is aimed at regulating the number of packets traversing a network at a given time.

Flow control is not always an effective measure against congestion, because it reduces throughput and may increase a packet's transit time. Flow control is only effective in reducing congestion if it manages to reduce the number of packets in the network at the critical points and at the critical time.

5.3.5 Network segmentation using Routers

The network is divided into multiple segments joined by routers. The routers are used to route packets to destination networks. Each router maintains a routing table consisting of destination addresses and the corresponding interfaces.

For a given number of destination addresses, adding more routers means fewer destination addresses per router. Hence smaller routing tables, less time taken to search a table, quicker forwarding and less congestion.

5.3.6 Network segmentation using Virtual LANs

One of the new concepts facilitating the management of bandwidth is Virtual LANs (VLANs) defined in IEEE Standard 802.1Q. VLANs can be based on ports, MAC (Multiple Access Control) addresses, Layer 3 protocols, IP network addresses or network subnet addresses. The concept allows network designers to group together particular users, independent of their physical locations.

VLANs provide smaller broadcasting domains, reducing the amount of bandwidth used up by broadcasts. In this way VLANs increase the bandwidth available to users.

Creating VLANs is really another form of network segmentation, where the segments are virtual rather than physical.

VLANs are supported by the new Ethernet switches currently on the market.

5.3.7 Queuing

Queuing is a method of traffic shaping (refer 2.5 "Congestion Control" in the Review of Literature). Traffic shaping removes peaks and provides a more even traffic flow that allows better use of the available bandwidth.

The various types of queuing methods are discussed in the Review of Literature (2.6 "Queuing").

5.3.8 Summary

1. Layer 3 bandwidth management comprises the traditional, software-based methods of congestion control. These methods aim to control the flow and number of packets in a network.

2. The use of VLANs is a new method of reducing congestion. The effect of VLANs on prioritised flows of Ethernet frames is investigated in Test 6 of Chapter 6.

5.4 Layer 4 Bandwidth Management

At Layer 4, bandwidth is managed by controlling the end-to-end connections between hosts. Layer 4 protocols, such as TCP, can be used not only to control the number of connections, but also to control the data flow between the hosts. This capability alone allows Layer 4 bandwidth management to be tackled in two ways, namely by -

1. limiting the number of end-to-end connections (Admission Control)
2. controlling the flow of data between two specific hosts (TCP Rate Control).

5.4.1 Admission Control based on Circuits

Where a connection-oriented protocol such as TCP sets up virtual circuits prior to transmission, congestion can be reduced by limiting the number of new circuits. Thus, once congestion reaches a pre-set level, no more circuits are set up.

5.4.2 TCP Rate Control

TCP, being a point-to-point protocol, can be used to control data flows between specific sender/receiver pairs on the network. TCP uses a sliding window to determine the number of packets a sender is allowed to send before receiving an acknowledgement.

This type of flow control also limits the number of packets traversing the network at a given time. Applying TCP flow control on a large scale, that is, by controlling the flow between many sender/receiver pairs, will benefit the network as a whole because fewer packets will be injected into it. TCP Rate Control is therefore also a means of congestion control.

Congestion can be controlled by dynamically by adjusting the TCP window in response to transmission timeouts on the network. According to (Tanenbaum, 1996, p.536), this is a practicable method because "most transmission timeouts on the Internet are due to congestion".

TCP Rate Control shapes the traffic, making it smoother and more predictable. It reduces the users' bandwidth requirements. This is in contrast to other bandwidth

management methods, such as network segmentation, which increase the bandwidth available to users.

Together with queuing, TCP Rate Control is used in commercially available "IP Traffic Shapers" (Reardon, 1998).

5.4.3 RSVP

In addition to Admission Control and TCP Rate Control, a third method of managing bandwidth at Layer 4 is RSVP. As explained in 2.8.1 "The limits of 802.1p prioritisation", RSVP is designed to monitor and control bandwidth along each link between sender and receiver. It can therefore be used to optimise bandwidth utilisation.

5.4.4 Summary

Layer 4 bandwidth management uses connection-oriented point-to-point protocols to provide end-to-end control over bandwidth.

5.5 Conclusions on Bandwidth Management

1. Layer 1 bandwidth management can entail high capital costs.
2. Layer 2 Ethernet switching is currently the most effective method of increasing bandwidth.
3. Bandwidth management at Layers 3 and 4 entails router configuration and programming. The design and implementation of these software solutions may require considerable effort and expertise from the network manager. The need to periodically monitor the bandwidth management scheme would further add to labour costs.

6.0 EXPERIMENTAL WORK

6.1 Introduction

Small, experimental LANs were set up to evaluate the effectiveness of 802.1p prioritisation.

6.1.1 Reason for conducting prioritisation tests

The significance of 802.1p prioritisation as a component of QoS was pointed out by Breyer and Riley – refer 2.8.1 "The limits of 802.1p prioritisation". Breyer and Riley (Breyer & Riley, 1999) see 802.1p prioritisation as a significant step toward implementing QoS in Ethernet/IP networks and one that would advance the development of Ethernet/IP QoS to a level comparable to that of ATM QoS).

IETF' s efforts at developing QoS have been directed mainly at the IP protocol and the work has been done in the context of WANs. However, Internet traffic usually needs to traverse one or more LANs as it leaves one host for another. Prioritisation, as defined in the IEEE 802.1p standard, is the first standardised method of providing a differentiated service for LANs.

The significance of 802.1p prioritisation is recognised by suppliers who are promoting it as part of their proprietary QoS schemes, e.g. Xylan's Switched Network Services – refer (Determan, 1999).

802.1p-capable equipment is now available. In fact, the Bay Networks BayStack 450 Ethernet switch, which is widely used in the School of Computer and Information Science (SCIS), is 802.1p-capable.

Nevertheless, the case study of the Edith Cowan University campuses revealed that prioritisation was not used to any significant extent within SCIS nor within ECU as a whole. There appears to be some reluctance on the part of network managers to implement 802.1p prioritisation. This experiment is aimed at discovering possible reasons.

6.1.2 Theoretical vs Empirical Testing

The theoretical performance of priority queuing has been extensively investigated by authors including (Bajaj et al., 1998a) and (Berger & Whitt, 1998). These authors used mathematical modelling and analysis to predict performance. However, theoretical evaluations may not reveal practical problems. To determine possible practical reasons for the slow acceptance of 802.1p prioritisation, empirical testing is required. In contrast to theoretical evaluations, in this experiment real (though experimental) networks are tested.

The advantage of having an experimental LAN is that the traffic, unlike the traffic in the operational ECU network, can be controlled. One disadvantage is that hardware availability limits the extent to which Internet conditions can be simulated. The equipment available for this experiment comprised a fast, current-technology, Ethernet switch and a limited number of relatively slow PCs.

6.1.3 Overview of Tests

The experimental work comprises two series of tests using different test beds. In both series of tests the strategy was to first validify the test bed in the familiar environment of a shared Ethernet.

First Series of Tests

Test 1 - Ethernet hub with one client at a time to measure individual throughputs.

Test 2 - Ethernet hub with combinations of clients to determine the effect of adding more users to a shared Ethernet.

Test 3 - Ethernet switch with one client at a time to measure individual throughputs.

Test 4 - Ethernet switch with combinations of clients to determine the effect of adding more users to a switched Ethernet.

Test 5 - Ethernet switch with prioritised ports to test prioritisation.

Test 6 - Ethernet switch with prioritised ports on separate VLANs to determine the effect of VLANs on throughput.

Second Series of Tests

In the Second Test Series, more and faster machines were used to generate heavier traffic. In addition, two Ethernet switches were used so that more pressure could be put on one particular switch port.

Test 7 - Ethernet hub with one client/server pair at a time to establish baselines for hub operation. Repeatability examined.

Test 8 - Ethernet hub with combinations of client/server pairs to determine the effect of adding more users to a shared Ethernet. LANalyzer consistency checked. Throughput normalised.

Test 9 - Ethernet switch with one client/server pair at a time to establish baselines for switch operation. Repeatability examined.

Test 10 - Ethernet switch with combinations of client/server pairs to determine the effect of adding more users to a switched Ethernet. Throughput normalised.

Test 11 - Ethernet switch with one prioritised port to test prioritisation.

Test 12 - Ethernet switch with all ports but one prioritised.

Test 13 - Ethernet switch with prioritised input and output ports to investigate reverse data flow.

6.2 First Series of Tests

6.2.1 Test Bed

Five identical PCs (C1 to C5) were configured as Netware 4.11 clients and arranged to simultaneously transfer files to a sixth PC (S) configured as a Netware 4.11 file server.

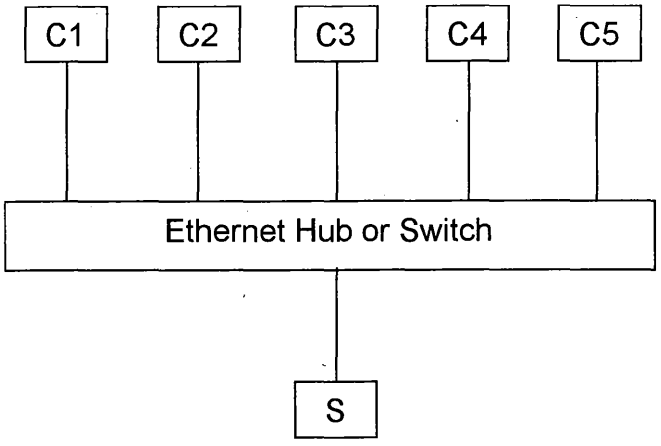

Figure 6.1 - Test Bed for First Test Series
General Case

All PCs (server and clients) were Intel 486 SX, 33 Mhz fitted with 10 Mbits/sec Genius network interface cards and running Windows 95. The networking device was either a 10BaseT, Netgear EN516 hub[4] (10 Mbits/sec) or a BayStack 450 switch (10/100 Mbits/sec), depending on the type of test.

Table 6.1 - Clients and Server used in the First Test Series

Processor	Intel 486 SX 33 Mhz
Network Interface Cards	Genius, 10 Mbits/sec
Cables	Category 5, UTP, maximum length: 12 m
Local Operating System	Windows 95
Network Operating System	Netware 4.11

[4] In this thesis the term "hub" refers to a repeater, whereas "switch" refers to a multi-port bridge.

The software on the clients was identical. It consisted of Windows 95 and Netware 4.11 client software. The software had been installed as a single image downloaded from the same source.

6.2.2 Generating Network Traffic

The traffic was generated by means of DOS batch files containing re-entrant code that continuously copied a set of test files from the clients to the server. The set of test files consisted of six files ranging between 2 to 20 MB. The files were created with the archiving utility WinZip, which was used to compress a large number of program and data files into a set of six large files.

This traffic pattern generated in this way is not typical of Internet links, as it does not contain any real-time traffic: a more realistic traffic pattern would have consisted of a combination of file transfers and (prioritised) video/audio transmissions. Equipment and expertise for setting up, controlling and monitoring real-time traffic was not available.

The importance of the traffic pattern is explained under 2.8.2 "Effect of Load". The generation of realistic traffic patterns can be considered for future experimental work.

Data transfer rates between each client and the server were measured using Novell's LANalyzer for Windows, Version 2.2, which was installed on each client. It monitors the traffic on the LAN segment a client is connected to.

The arrangement made use of Windows' multi-tasking capability, running LANalyzer while the file transfer was taking place. In Tests 1 to 6, the file transfer was taking place in the foreground, with LANalyzer running in the background. The implications of running LANalyzer and the file transfer on the same client are discussed in 6.11.5 "Conclusions from Test 7".

6.3 Test 1: Ethernet hub with one client at a time

6.3.1 Setup

The five clients C1, C2, C3, C4 and C5 were connected to an Ethernet hub one at a time, and programmed to continuously copy test files to the server, S.

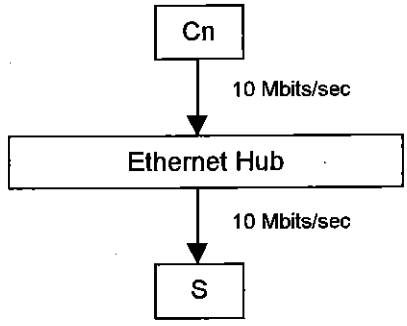

Figure 6.2 - Test Bed for Test 1

6.3.2 Purpose

This test was conducted to –

1. determine each client's throughput capability
2. note and quantify significant variations in performance between identical clients, if any
3. check LANalyzer for correct operation.

6.3.3 Procedure

All readings taken during this test (and all following tests) were as measured and tabulated by LANalyzer. LANalyzer provides graphs and tables in real time - see Figure 6.3 below. The values in the following tables, e.g. Table 6.2, Table 6.3, etc, have been extracted from the real-time tables provided by LANalyzer.

All average and peak values in this experiment represent overall averages for a particular test period. LANalyzer was re-started (and its counters re-set) at the beginning of each test.

Run-up Period

In this test, and in Tests 2, 3 and 4 below, readings were taken only after variables such as Kbytes/sec, Utilisation, etc had had time to converge to steady state values. Readings were taken approximately 15 mins after a file transfer commenced. During this time, referred to as the "run-up time", the LANalyzer graphs and tables were monitored to ensure that the readings were converging to steady state values.

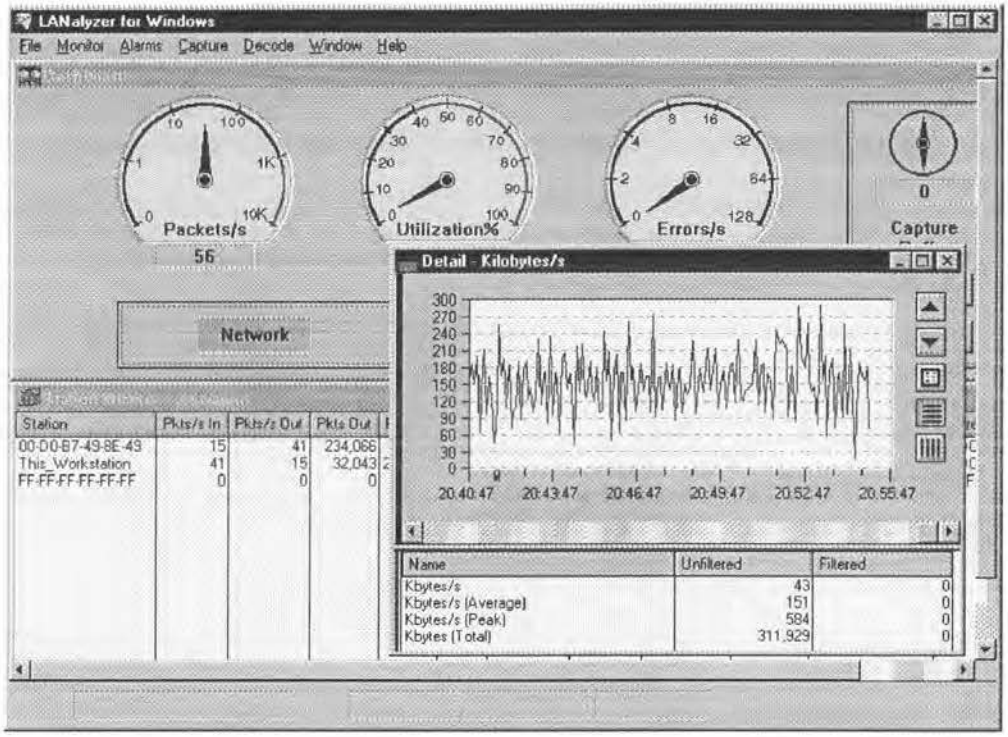

Figure 6.3 - Typical LANalyzer Screen

6.3.4 Results

The individual performances of the five clients are summarised below. Table 6.2 is interpreted as follows:

Over a run-up period of approximately 15 mins, Client C1 was transferring the test files to the server at an average rate of 220 Kbytes/sec. Similarly, over a similar run-up period, also lasting approximately 15 mins, Client C5 was transferring the same set of test files at an average rate of 224 Kbytes/sec.

Table 6.2 - Individual Client Performance (shared Ethernet)

	Client				
	C1	C2	C3	C4	C5
Kbytes/s (Avg)	220	261	223	219	224
No.Packets/s (Avg)	185	214	186	185	186
Utilisation (%)	18	21	18	17	18

The average throughputs for Clients C1 to C5 varied between 220 Kbytes/sec and 261 Kbytes/sec. The average of the average throughputs[5] for the five clients was (220 + 261 + 223 + 219 + 224)/5 = 229 Kbytes/sec. The percentage variation from this average is +14% (Client C2, the fastest) and – 4% (Client C4, the slowest).

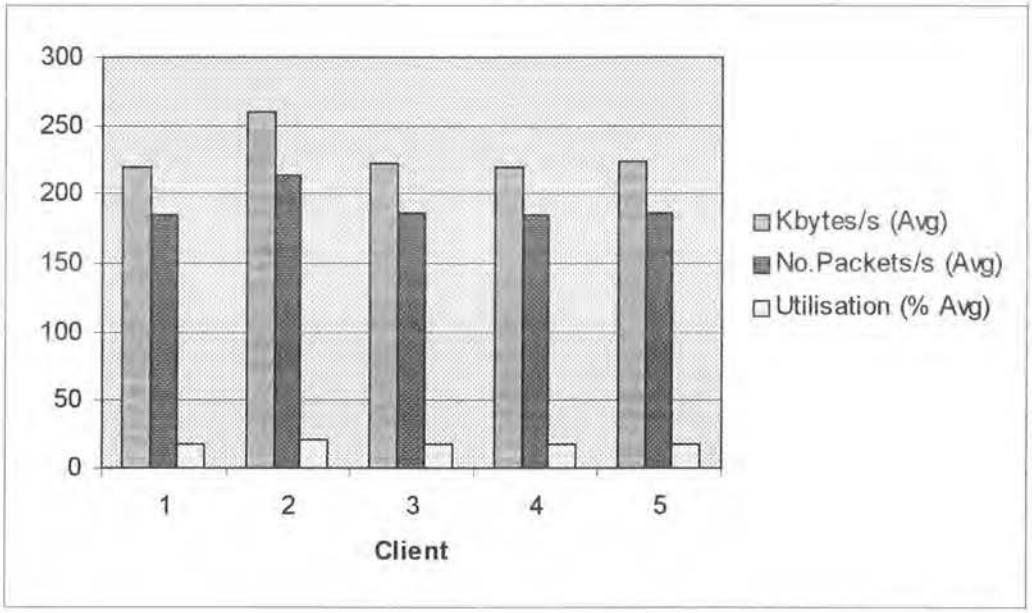

Figure 6.4 - Variations in Client Performance (shared Ethernet)

Throughput Level

An average throughput of 229 Kbytes/sec appears low. Keeping in mind that for the purposes of this experiment it is the clients' relative throughputs that are of importance,

[5] All LANalyzer readings are averaged with respect to time. In addition, in this thesis, many variables are averages for a set of clients. To avoid having to continually refer to averages of averages, the "average throughput", when it refers to time, will from now on be referred to as the "throughput" When averages are mentioned they will refer to averages calculated for a set of clients.

rather than the actual values, it is nevertheless appropriate to examine the actual values, if only to ascertain that they represent a reasonable result.

Let us compare the average throughput of 229 Kbytes/sec with the bandwidths of the components through which the data passes. These components are as follows:

1. Hard disk drive

The maximum transfer rate of the hard disk drive is determined by the rotational speed and the track capacity. For a typical 500 Mbyte hard drive (6000 rpm, 63 sectors/track, 512 bytes/sector, 16 heads, 1651 cylinders) as used in this experiment the

maximum sustainable transfer rate to a first approximation

= Track capacity x number of revolutions/sec

= 63 sectors x 512 bytes/sector x (6000/60) bytes

= 3,225,600/1024 Kbytes/sec

= 3150 Kbytes/sec

2. Interface between the PC and the Hard Disk

According to Rosch (1999, p.435) for an AT Attachment interface with a 16-bit connection, the peak transfer rate

= 8 Mbytes/sec

= 8 x 1024 Kbytes/sec

= 8192 Kbytes/sec.

3. Industry Standard Architecture bus

For a 16-bit ISA bus, the maximum transfer rate is 8Mbytes/sec (Rosch, 1999, p.557), i.e. 8192 Kbytes/sec.

4. Network Interface Card

The NICs are rated at 10 Mbits/sec. They should be able to handle 10,000,000 bits/sec.

Maximum theoretical transfer rate

= 10,000,000 / (8 x 1024) Kbytes/sec

= 1220 Kbytes/sec

Clearly, none of the above four components can be the limiting factor that would reduce the average throughput to 229 Kbytes/sec.

It appears then that in Test 1 (where only one client was active at a time) the relatively low throughput was due to overheads incurred when data transfers are negotiated between the disk drive, memory, DOS and other programs. According to (Rosch, 1997):

> *"The throughput between your drive and controller is higher than between drive and memory. And the actual throughput to your programs – which must be managed by your operating system – is slower still. Throughput to DOS on (sic) the order of a few kilobytes/sec is not unusual for hard disk drives that have quoted transfer rates in excess of twenty megabytes per second."*

This explains the gulf between the relatively high bandwidths of the above components and the actual throughput.

6.3.5 Conclusions from Test 1

1. The main purpose of Test 1 was to show relative machine performance. Analysis of the results shows that the clients, although identical with respect to both hardware and software, vary significantly in their ability to transfer files, as is shown by their throughput, which varies by 18% (14 % + 4%). The results from subsequent tests need to be interpreted in the light of these variations.

2. Another aim of Test 1 was to establish a degree of confidence in the test bed and method. Figure 6.4 above demonstrates a correlation between Kbytes/sec (Avg), No.packets/sec (Avg) and Utilisation (%Avg), suggesting that LANalyzer is functioning correctly in the shared Ethernet environment of Test 1.

3. The relatively low throughput is to be expected when the overheads of the interfaces are taken into account.

6.4 Test 2: Ethernet hub with combinations of clients

6.4.1 Setup

As for Test 1, but with different combinations of clients transferring files simultaneously.

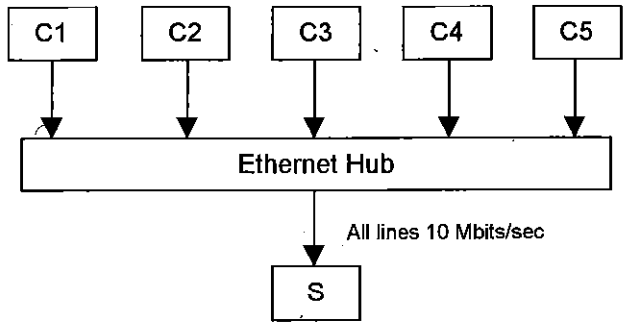

Figure 6.5 - Test Bed for Test 2

6.4.2 Purpose

Test 2 aims to –

1. To ascertain that the experimental network is behaving like a standard Ethernet by determining the effect of adding more clients
2. provide a basis for comparison with the Ethernet switch
3. check LANalyzer utilisation figures.

6.4.3 Procedure

Four test runs were conducted, each lasting about 15 mins. One set of readings was taken at the end of each test run. Beginning with client combination C1 + C2, the test runs were then repeated with C1 + C2 + C3, etc, each test run presenting a heavier load to the server and to the network.

6.4.4 Results

The performances of several combinations of clients are summarised in Table 6.3, which is interpreted as follows:

Over a run-up period of approximately 15 mins, Clients C1 and C2 together transferred the files at an average rate of 356 Kbytes/sec. On a machine basis, the average transfer rate for C1 and C2 was 356/2 = 178 Kbytes/sec per machine. This compares with an average for C1 and C2 of (220 + 261)/2 = 240 Kbytes/sec in Test 1 (where C1 and C2 were running as single units).

Similarly, at the end of another run-up period also lasting approximately 15 mins, Clients C1, C2 and C3 together transferred the same set of files at an average rate of 484 Kbytes/sec (Avg), and the average rate per machine was 484/3 = 161 Kbytes/sec.

Table 6.3 - Combinations of Clients in a shared Ethernet

Client combination	C1	C1 + C2	C1+C2+C3	C1+C2+C3+C4	C1+C2+C3+C4+C5
Number of clients	1	2	3	4	5
LAN throughput	220	356	484	530	547
Throughput/Client	220	178	161	132	109
No Packets/sec	185	304	425	464	480
Utilisation (%)	18	29	40	43	44
Calculated Utilisation	18.0	29.2	39.6	43.4	44.8

Checking Utilisation

For a 10BaseT network the maximum theoretical data transfer rate is 10 Mbits/sec, that is, 10,000,000/(8 x1024) = 1221 Kbytes/sec. The utilisation of the Ethernet bus when all five clients are active is therefore 547/1221 = 44.8%. This calculated utilisation figure matches the measured utilisation figure of 44%. In fact, Table 6.3 shows that all the calculated utilisation figures closely match the figures provided by LANalyzer.

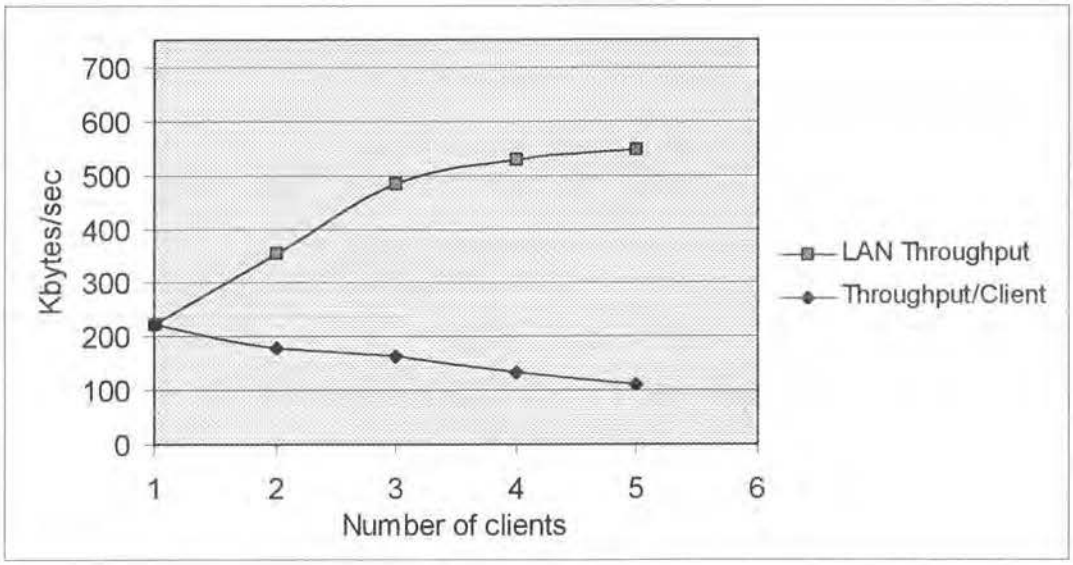

Figure 6.6 - Adding more Clients to a shared Ethernet

6.4.5 Conclusions from Test 2

1. As more clients access the Ethernet bus, the total transfer rate (avg Kbytes/sec for the bus) increases but the individual transfer rates (avg Kbytes/sec per machine) decrease. This shows the effect of increasing contention and collisions.

2. The results from Test 2 were as expected for an Ethernet: Figure 6.6 shows similarities to standard Ethernet performance graphs provided by Boggs, Mogul and Kent (1988) and confirms that the experimental Ethernet is working correctly.

3. LANalyzer utilisation figures have been shown to be correct.

4. As in Test 1, the LANalyzer readings for Utilisation and No.Packets/sec correlated with the readings for Kbytes/sec, providing further evidence that LANalyzer is functioning correctly

6.5 Test 3: Ethernet switch with one client at a time

6.5.1 Setup

Each PC was connected to a separate port on a BayStack 450 Ethernet switch. This switch has a 2.56 Gbits/sec switch capacity. The switch's in-built autosensing function set all switch ports to 10 Mbits/sec.

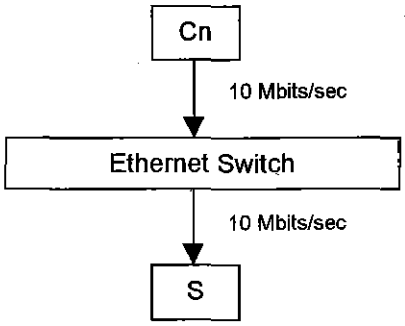

Figure 6.7 - Test Bed for Test 3

As each port on the switch represents a different collision domain, a copy of LANalyzer had to be run on every client in order to measure the throughput for each port. This necessitated having to use the same client to both generate and monitor traffic. The possibly adverse effect of multi-tasking the clients is discussed under Test 7 results.

6.5.2 Purpose

As for Test 1 but for a switched Ethernet environment.

6.5.3 Procedure

Same as for Test 1.

6.5.4 Results

The results from Test 3 are summarised in Table 6.4, which is interpreted as follows:

Over a run-up period of approximately 15 mins, Client C1 transferred the files at an average rate of 210 Kbytes/sec.

Table 6.4 - Individual Client Performance (switched Ethernet)

	Client				
	C1	C2	C3	C4	C5
Kbytes/s (Avg)	210	203	209	204	218
No.Packets/s (Avg)	176	169	173	169	182
Utilisation (%)	17	16	17	16	18

Individual throughputs for the five clients varied between 203 Kbytes/sec and 218 Kbytes/sec. The average throughput is (210 + 203 + 209 + 204 + 218)/5 = 209 Kbytes/sec. The percentage variation from this average is +4% (C5, the fastest) and –3% (C2, the slowest).

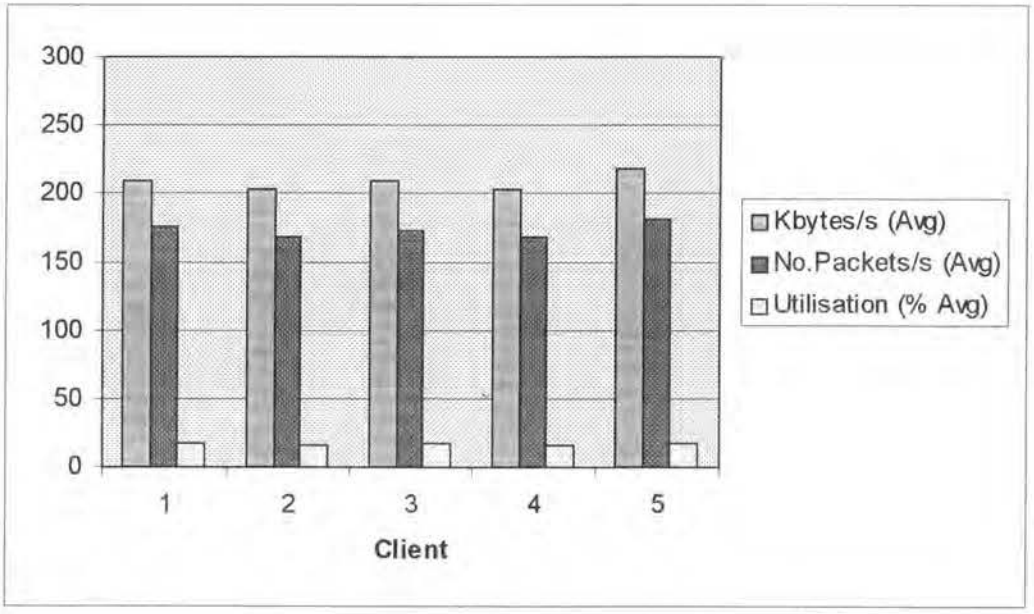

Figure 6.8 - Variations in Client Performance (switched Ethernet)

6.5.5 Conclusions from Test 3

1. In the switched Ethernet environment of Test 3, the clients showed significant variations (7 %) in throughput capability. The variations, although smaller than the variations measured in the shared Ethernet of Test 1, need to be taken account of in the prioritisation tests that follow.

2. The calculated average throughput rate of 209 Kbytes/sec compares to 229 Kbytes/sec for the hub in Test 1, showing that under light loads, as when only one client is accessing the network, the hub is 10% faster than the switch.

3. There is a correlation between Kbytes/sec (Avg), No.Packets/sec (Avg) and Utilisation (%Avg), suggesting that LANalyzer also works correctly in a switched Ethernet network.

4. The readings for Utilisation, 16% (Avg) to 18% (Avg) match the corresponding readings obtained for the hub in Test 1.

6.6 Test 4: Ethernet switch with combinations of clients

6.6.1 Setup

As for Test 2, but using an Ethernet switch instead of the Ethernet hub.

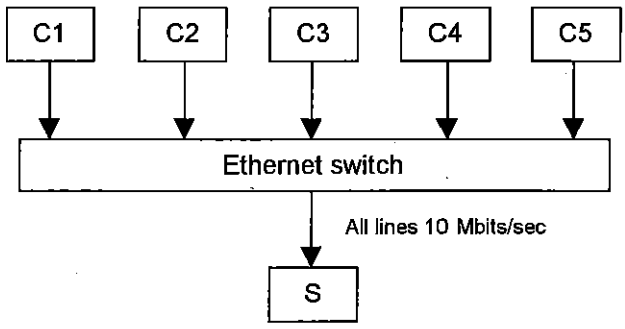

Figure 6.9 - Test Bed for Test 4

6.6.2 Purpose

This test is preparatory to the prioritisation tests that follow. It is intended to ascertain that the switched Ethernet is operating correctly.

6.6.3 Procedure

Same as for Test 2, that is four test runs each presenting a heavier load to the network.

6.6.4 Results:

The results for Test 4 are summarised in Table 6.5, which is interpreted as follows:

During a run-up period of approximately 15 mins, Client C1 was transferring the files at 171 Kbytes/sec. At the same time, Client C2 (connected to a different network segment by virtue of being on a different switch port) was transferring the same set of files at 173 Kbytes/sec. The total throughput, or LAN throughput, was (171 + 173) = 344 Kbytes/sec. The average throughput per client was (171 + 173)/2 = 172 Kbytes/sec.

The maximum load on the server was (721 Kbytes/sec x 1024 x 8) / 1,000,000 = 5.9 Mbits/sec on average. This represents a substantial load for the 10 Mbits/sec network, but a very small load for the 2.56 Gbits/sec BayStack 450 switch.

Table 6.5 - Combinations of Clients in a switched Ethernet

Client Combination	C1	C1+ C2	C1+C2+C3	C1+C2+C3+C4	C1+C2+C3+C4+C5
Number of Clients	1	2	3	4	5
Individual Throughputs (Kbytes/sec)	210	171,173	160,150 152	127,129 133,134	161,125,146 139,150
LAN Throughput (Kbytes/sec)	210	344	462	523	721
Avg Throughput/Client (Kbytes/sec)	210	172	154	131	144
Throughput/Client (No.packets/sec)	176	147,148	138,130 131	110,112 114,117	144,121,130 124,134
Individual Line Utilisation (%)	17	14,14	13,12,12	10,10 10,11	13,10,11 11,12

The values for "LAN Throughput" and "Avg Throughput/Client" were calculated. All other values were measured.

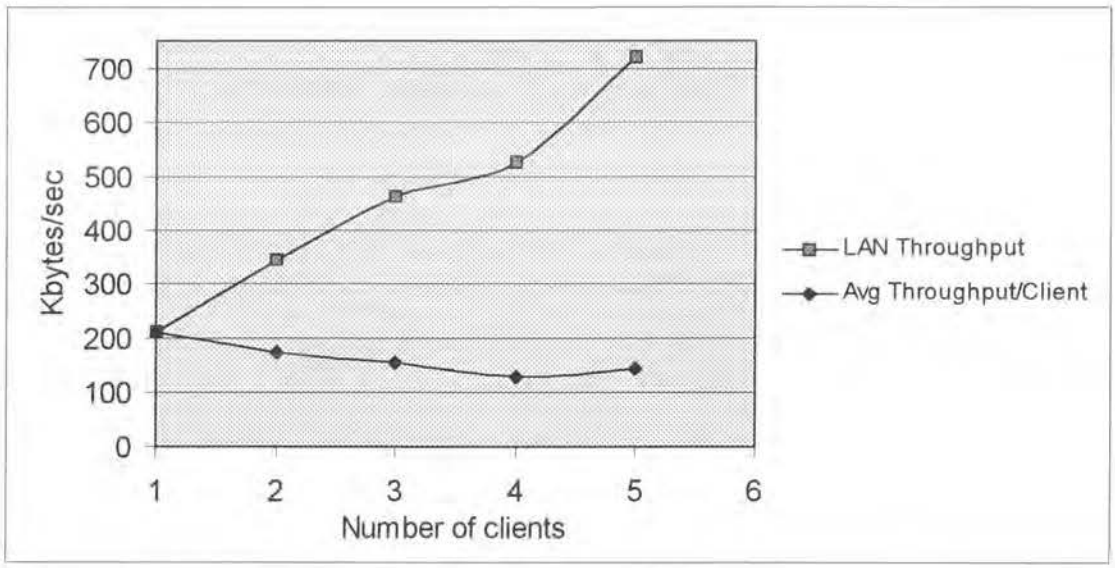

Figure 6.10 - Adding more Clients to a switched Ethernet

6.6.5 Conclusions from Test 4

1. As more clients are added to the switched network, the LAN throughput increases (up to 721 Kbytes/sec, see Table 6.5). The throughput per client decreases steadily from the one-client to four-client combination, but increases for the five-client combination.

If we discard the five-client result as an experimental aberration, then a trend emerges that is similar to the one obtained for the shared Ethernet in Test 2. The similarity is shown in Figure 6.11 which uses the results from Tests 2 and 4 to compare the average throughputs from clients in shared and switched Ethernet networks.

In Test 2 the decrease in individual throughput could be attributed to an increasing number of collisions due to increasing utilisation (up to 44%, refer Table 6.4). However, in the switched environment of Test 4 collisions cannot occur. So, what caused the decrease in individual throughputs observed for the switched network (if it is not coincidental)?

- **Line speed?** The maximum load that was put on any line during Test 4 was 5.9 Mbits/sec (721 Kbytes/sec, Table 6.5) on the server server. This maximum is well below the nominal 10 Mbits/sec that a 10BaseT Ethernet is capable of handling. Line speed was therefore not a factor that caused the decrease in individual throughput.

  This conclusion is also borne out by extrapolation of the "LAN Throughput" curve in Figure 6.10 - the upward slope of the curve suggests that the LAN can easily handle more than five clients.

- **Switch Latency?** There is a finite time interval between the last bit of an Ethernet frame entering a switch and the first bit exiting. The latency of a switch can throttle the throughput when the number of active ports multiplied by their port speed, exceeds the backplane capacity of the switch.

  The BayStack 450 switch has a rated backplane capacity of 2.56 Gbits/sec. This has been confirmed, in essence, by independent tests (TheTollyGroup, 1998) which demonstrated a backplane capacity between 2.0 and 2.4 Gbits/sec, depending on frame size. Against that, in Test 4, the maximum load that could have been presented by the network was only –

  5 ports x 10 Mbits/sec per port = 50 Mbits/sec.

The switch was therefore forwarding frames at full line speed and was not blocking any of the ports. Switch latency was not a factor that caused the decrease in individual throughput.

It appears that the general decrease in individual throughput as more clients are added to the switched network, is due to limitations of the server rather than the limitations of the network. It is possible that the server's hard disk is struggling to handle all the files simultaneously arriving from five clients. The server is likely to be a bottleneck.

Alternatively, the five readings obtained for the switched Ethernet may simply be unrelated fluctuations. To clarify this situation, the Second Test Series removes the suspected bottleneck by providing a dedicated server for each client. In addition, in the Second Test Series, a 6th client is added to test the unexpected upward trend of the Switched-Ethernet curve at the five-client mark.

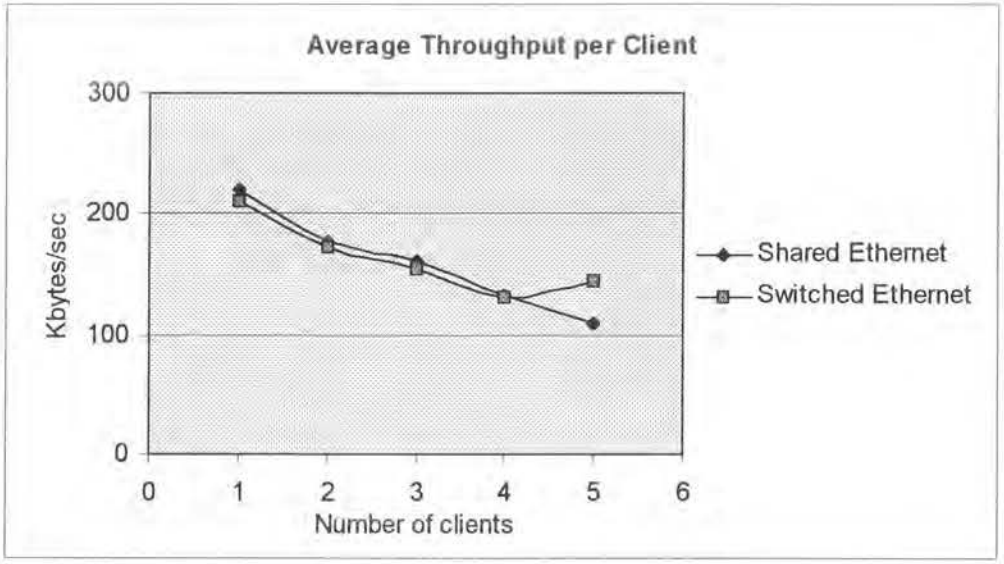

Figure 6.11 - Adding more Clients to shared and switched Ethernets

2. Under light load, such as when 2, 3 or 4 clients were accessing the network simultaneously, the switch was slightly slower than the hub - compare the LAN throughputs of 344, 462 and 523 Kbytes/sec in Table 6.5 with the Ethernet bus

throughputs of 356, 484 and 530 Kbytes/sec (Avg) respectively in Table 6.3 - Combinations of Clients in a shared Ethernet.

This can possibly be attributed to the switch having a higher overhead – it has to examine each received frame and read the destination address, whereas the hub simply forwards a received frame to all other ports on the hub.

3. Under relatively heavy load, such as when 5 clients were accessing the network simultaneously, the switch was faster than the hub - compare the calculated total throughput of 721 Kbytes/sec in Table 6.5 with the measured Ethernet bus throughput of 547 Kbytes /sec (Avg) in Table 6.3.

   This can be attributed to the delays due to increasing number of collisions at heavy loads out weighing the hub's advantage of smaller overheads.

4. In this switched network, the LANalyzer readings for Utilisation and No.Packets/sec correlated with the readings for Kbytes/sec.

6.7 Test 5: Ethernet switch with prioritised ports

6.7.1 Setup

The five clients were connected to the switch and arranged to simultaneously transfer files to the server. This was the same setup as in the last part of Test 4 (Client Combination C1+C2+C3+C4+C5), except that in Test 5, Ports 4 and 5 were configured as high-priority ports.

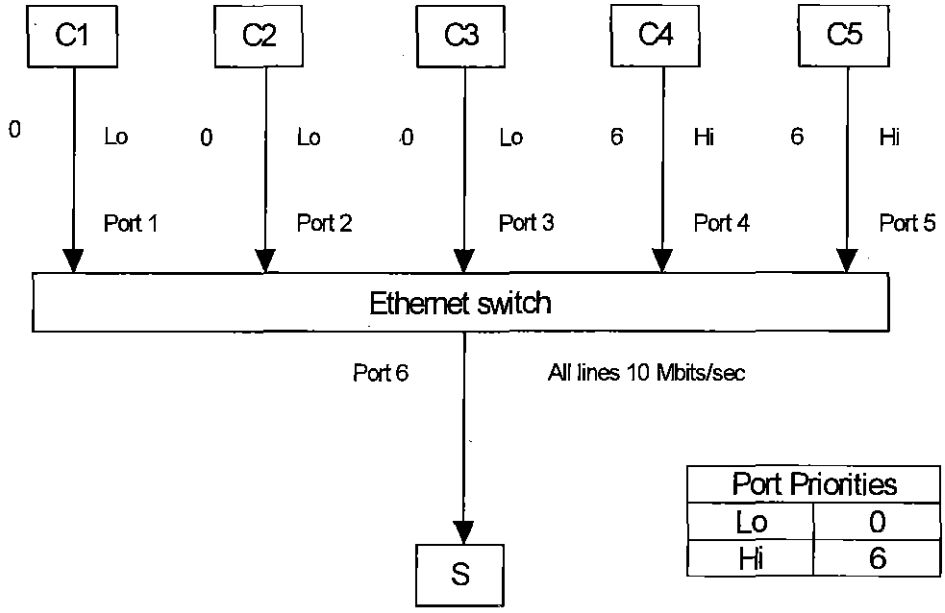

Figure 6.12 - Test Bed for Test 5

Port Priority was set to "0" for the Ports 1, 2 and 3, and to "6" for the Ports 4 and 5. Port Priority 0 was mapped to "Low Traffic Class" and Port Priority 6 to "High Traffic Class". In accordance with the IEEE 802.1Q tagging rules, the untagged frames arriving from the clients are assigned user priorities equal to the port priorities, which makes traffic from Clients C4 and C5 "High Traffic Class".

6.7.2 Purpose

To note the effect of prioritising the ports on an 802.1p-capable Ethernet switch.

6.7.3 Procedure

Two test runs were conducted. During the first test run the simultaneous transfer of files was allowed to proceed for an hour approximately. A set of readings (Reading 1) was taken.

The clients were then re-started. The file transfers were recommenced and LANalyzer was re-started on each client. Sets of readings (Reading 2, 3 and 4) were taken after 10mins, 20mins and 3hrs35mins respectively.

6.7.4 Results

The results from Test 5 are tabulated in Table 6.6, which is interpreted as follows:

Over a run-up period of approximately one hour, the low-priority clients (C1, C2 and C3) were transferring data at 112, 117 and 121 Kbytes/sec, while the high-priority clients (C4 and C5) were transferring the same data at 119 and 110 Kbytes/sec. The average throughputs for low and high-priority clients were 117 and 115 Kbytes/sec respectively.

Table 6.6 - Ethernet switch with prioritised ports

Reading	Run-up Period	Throughput (Kbytes/sec)						
		High-priority Clients			Low-priority Clients			
		C4	C5	Avg	C1	C2	C3	Avg
1	1 hr	119	110	115	112	117	121	117
2	10 mins	108	116	112	107	112	116	112
3	20 mins	109	113	111	105	111	113	110
4	3hrs 35mins	104	116	110	108	111	108	109

Similarly, over a run-up period of approximately 10 mins, the low-priority clients were transferring data at 107, 112 and 116 Kbytes/sec. The high-priority clients were transferring the same data at 108 and 116 Kbytes/sec. The average throughput at this stage was 112 Kbytes/sec for both low and high-priority clients.

Then, after extending the 10min run-up period to 3hrs 35mins, the low-priority clients were transferring data at 108, 111 and 108 Kbytes/sec. The high-priority clients were transferring the same data at 104 and 116 Kbytes/sec. The average throughputs for low and high-priority clients were 109 and 110 Kbytes/sec respectively.

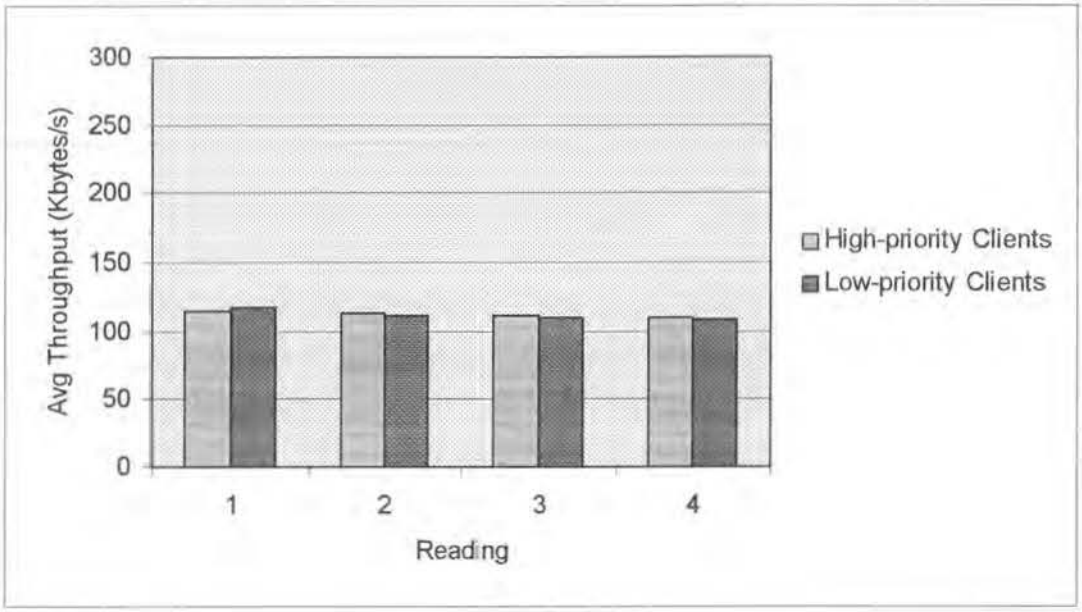

Figure 6.13 - Throughput from High and Low-priority Clients

6.7.5 Conclusions from Test 5

1. Prioritisation

The results from Test 5 do not demonstrate higher throughput for high-priority ports.

Even over a test period exceeding 3 hours, the low-priority ports were never held up by the high-priority ports, at least, not for any significant amount of time. It appears that the switch was able to consistently clear the low-priority queues before the arrival of the next high-priority frame. In other words, the switch, being very fast (2.56 Gbits/sec backplane capacity), may not have been put under sufficient pressure for prioritisation to come into its own. Test 11 addresses this issue.

Although the above may be the reason for the unexpected test results, the following two issues need to be examined.

Do the NICs need to be 802.1p compliant?

Based on the IEEE 802.1Q tagging rules described in the BayStack 450 User Manual, the frame processing mechanism for Test 5 would have been as follows:

The frames arriving from the clients were IEEE 802.2 frames, the default frame type for NetWare 4.11. The frames were untagged, that is, they did not contain any priority information.

In Test 5, all frames entering the switch at Ports 1, 2, 3, 4 and 5 were sent to Port 6, the port to which the server was connected. Frames entering Ports 1, 2 and 3 were assigned low priority and sent to the low-transmit queue on Port 6. Similarly, frames entering Ports 4 and 5 were assigned high priority and sent to the High-transmit queue on Port 6.

Port 6 was untagged. According to the 802.1Q Tagging Rules, tagged frames exiting an untagged port will have their tags removed. Hence the frames exiting the BayStack switch remained untagged, as they were before they entered the switch. Thus the frames arriving at the server's NIC were still IEEE 802.2 frames. This means that the clients' NICs and the server's NIC do not need to handle tagged frames. It should also mean that the NICs used in this experiment do not need to be 802.1p compliant.

Is it necessary to implement VLANs?

One possible explanation for the unexpected results could be that the throughput was affected by the switch's configuration. As noted under "VLANs" in 2.10.1 "ISO/IEC 15802-3", 802.1p prioritisation uses the VLAN frame to carry priority information. The VLAN frame carries an additional 4-byte field - the "tag". In addition to a 3-bit User Priority field, the "tag" also includes a 12-bit VLAN Identifier (VID) field. If, as was the case in Test 5, the VID field is not specified, then the BayStack 450 switch will insert a default value of "1", so that all frames belong to VLAN1. Every tagged frame, it appears, must be associated with a VLAN.

In view of the above, it may also be necessary to implement two or more specific VLANs in order to use the VLAN frame. Test 6 is aimed at testing this hypothesis.

2. Prioritisation Overhead

In Test 5 the average throughputs for both low and high-priority clients are consistently lower than the average throughput per client obtained in Test 4 – compare, for example, the average throughputs of 117 and 115 Kbytes/sec in Table 6.6 - Ethernet switch with prioritised ports, with the average throughput of 144 Kbytes/sec obtained for the 5-client combination in Table 6.5 - Combinations of Clients in a switched Ethernet.

Test conditions in Test 4 were the same as in Test 5, except that in Test 4 the switch was not configured for prioritisation. The results therefore suggest that prioritisation introduces an overhead that reduces the average throughput from clients.

6.8 Test 6: Ethernet switch with prioritised ports on separate VLANs

6.8.1 Setup

The previous test, Test 5 above, was carried out with the BayStack 450 switch in its default VLAN configuration, that is, with all ports belonging to the same VLAN (VLAN 1).

Test 6 is a repetition of Test 5, except that the low-priority and high-priority ports are assigned to separate VLANs – VLANs 1 and 2 respectively.

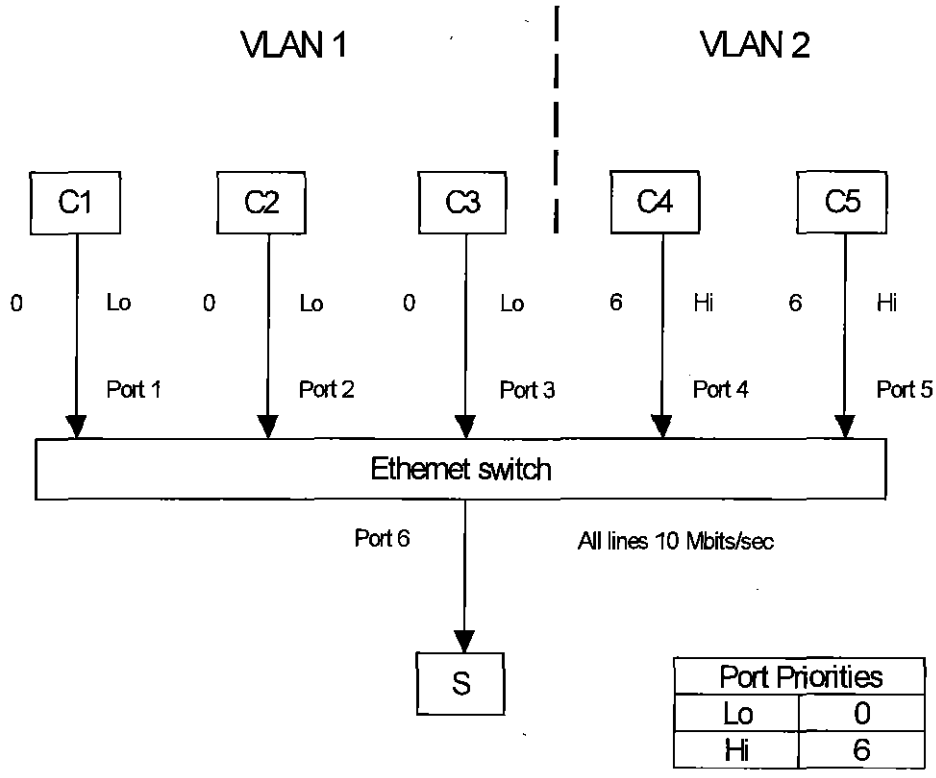

Figure 6.14 - Test Bed for Test 5

6.8.2 Purpose

To determine whether the switch's VLAN configuration affected the results from Test 5.

6.8.3 Procedure

Same as for Test 5, except that one test run only was conducted.

6.8.4 Results:

The results from Test 6 are summarised in Table 6.7.

Table 6.7 - Ethernet switch with prioritised ports on separate VLANs

Reading Set	Run-up Period	Throughput (Kbytes/sec)						
		High-priority Clients VLAN1			Low-priority Clients VLAN2			
		C4	C5	Avg	C1	C2	C3	Avg
1	15 mins	107	116	112	111	108	105	108
2	25 mins	112	115	114	109	104	109	107
3	40 mins	110	114	112	108	106	112	109

Firstly, comparing the throughputs from Tests 5 (no VLANs, Figure 6.13) and Test 6 (VLANs implemented, Figure 6.15), no significant differences can be detected for both the prioritised and non-prioritised ports. This suggests that VLAN configuration has no effect on throughput.

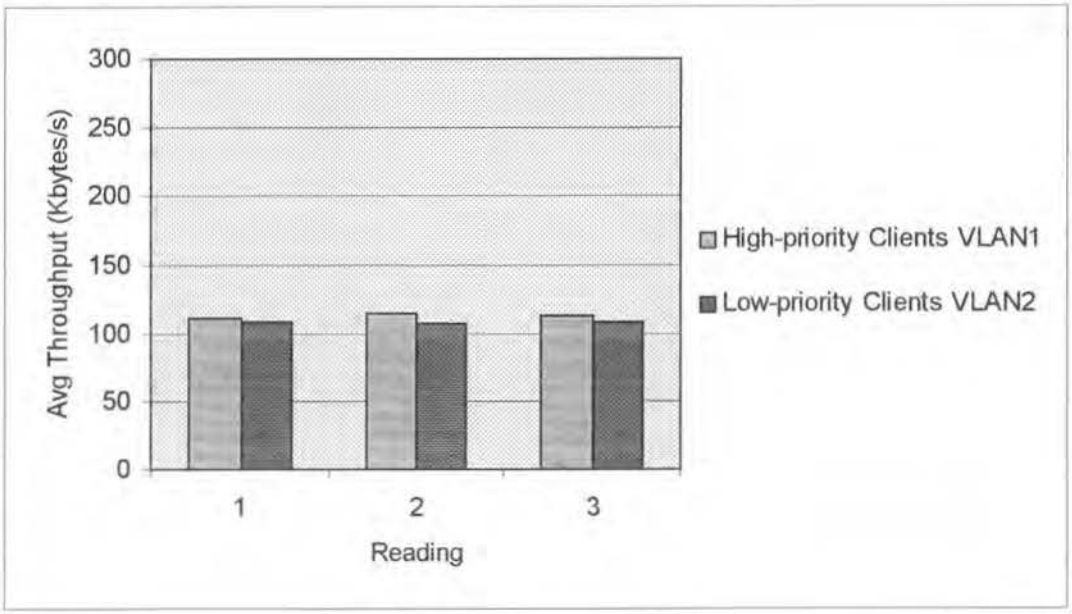

Figure 6.15 - Throughput from High and Low-priority Clients on separate VLANs.

Secondly, comparing the prioritised and non-prioritised throughputs in Test 6, Figure 6.15 shows that, on average, the high-priority ports achieved higher throughputs than the low-priority ports. However, the increases were not consistent – they were not achieved for every high-priority port and not within every test period. In addition, the increases are small and, in percentage terms, are comparable to the variations in individual machine performance, as measured in Test 3.

6.8.5 Conclusion from Test 6

The prioritisation results from Test 6 are essentially the same as the prioritisation results from Test 5. The implementation of two specific VLANs in Test 6 had no effect on the prioritisation mechanism. The VLAN configuration of an Ethernet switch can be disregarded in the context of prioritisation. This was done in the second series of tests.

6.9 Discussion (First Series of Tests)

Tests 1 to 6 did not yield the expected result, which was that with a large sample (and noting that thousands of frames would pass through the switch during a 15 second test period) the laws of statistics would ensure that the high-priority ports passed more data than the low-priority ports.

The First Test Series was summarised by Engel and Maj (1999) as follows: *"There were no significant differences in the data transfer rates recorded for prioritised and non-prioritised ports. This suggests that prioritisation provides no significant advantage when Ethernet switches are lightly loaded."*

Additional Tests: Tests 5 and 6 were repeated on a 3-machine network using similar hardware and the same software. Again, the advantage of prioritising ports could not be demonstrated. These additional tests, though not relied upon for primary data, nevertheless strengthen the above conclusion that prioritisation provides no significant advantage when Ethernet switches are lightly loaded.

Three factors that may have contributed to the unexpected result, are discussed below.

1. Burstiness of the traffic

From Bajaj, Breslau and Shenker (1998a), some of whose findings were outlined in the Review of Literature (refer 2.8.2 "Effect of Load"), we know that the "*benefits of priority service*" are less marked for smooth loads and that these benefits "*only occur at higher levels of utilization*".

In Tests 1 to 7 the continuous transfer of large files gave rise to a relatively smooth load with few bursts - see Figure 6.3 - Typical LANalyzer Screen.

In addition, the BayStack switch, rated at 2.56 Gbits/sec, was forwarding frames at average rates of approximately 5 Mbits/sec respectively. This represents a very low utilisation of the switch.

Thus we have the conditions under which, according to Bajaj et al, the benefits of prioritisation are either small or non-existent.

2. Interactions between concurrent file transfers

According to Bajaj et al (1998a, p.76) best-effort applications tend to absorb delays.

The file transfers in Tests 5 and 6 are examples of best-effort applications. Any delays, in forwarding frames from one of the three low-priority ports could have been absorbed by the other two low-priority file transfer operations, thus rendering the low-priority ports as fast as the high-priority ports.

3. The presence of buffers

The dramatic effect of buffers on throughput is demonstrated by Guerin, Kamat, Peris and Rajan (1998), who propose the use of buffers as a QoS mechanism - refer 2.6.1 "Priority Queuing".

The file transfer in Tests 1 to 7 could have been expedited by the PCs' several data buffers. In particular, the presence of buffers could have made the file transfers less dependent on potential bottlenecks such as the low-priority ports. In other words, the buffers could have assisted the file transfer operations to "absorb delays".

6.10 Second Series of Tests

The tests in the Second Series are follow-up tests to the First Series. The Second Series tests were conducted to –

1. establish a broader statistical baseline for each machine
2. determine the extent of repeatability
3. determine the effect of time on data transfer rates
4. eliminate the bottleneck due to a single server having to process files arriving simultaneously from several clients – refer 6.6.5 "Conclusions from Test 4".

1. Baselines

The results from Test 3 showed that under switch operation the individual performance of clients varied by 7 %. Tests 5 and 6 showed that the effects of prioritisation, if at all measurable with the equipment available, did not substantially exceed this variation in individual performance. In other words, the variations in individual performance could be comparable in magnitude to the differentiation expected from prioritisation. There is a possibility that variations in machine performance are masking the effects of prioritisation. To reduce the chances of this happening in the Second Test Series, a baseline with a much broader statistical base was established for each client.

2. Repeatability

In addition to variations in individual performance, there was also the possibility that the machines were not performing consistently. Such lack of repeatability could also mask the effects of prioritisation. Thus in the Second Test Series, each test was repeated three times.

3. Effect of Time

It was noted during preparations for Test 1 (refer 6.3.3 "Procedure") that throughput depended to some extent on the time it was measured and that the averages measured by LANalyzer required about 15 mins to converge to steady state values. To ensure that

measurements were not unduly affected by time, all measurements in the first series of tests were only taken after a 15 min "run-up period". In addition, of course, measurements were taken as close together as possible - typically, a set of readings would be taken within 30 secs. While these precautions minimised the effect of time on throughput measurements, they may not have entirely eliminated time as a factor affecting measurements of throughput.

Consequently, in the second series of tests, the effect of time was investigated in detail: throughput was measured at different times and Throughput vs Time graphs were plotted for each client and for various client combinations.

4. Bottleneck

To eliminate the server bottleneck, additional servers were used to share the demands from multiple clients. Eliminating this possible bottleneck, allowed a greater load to be placed on the Ethernet switch under test.

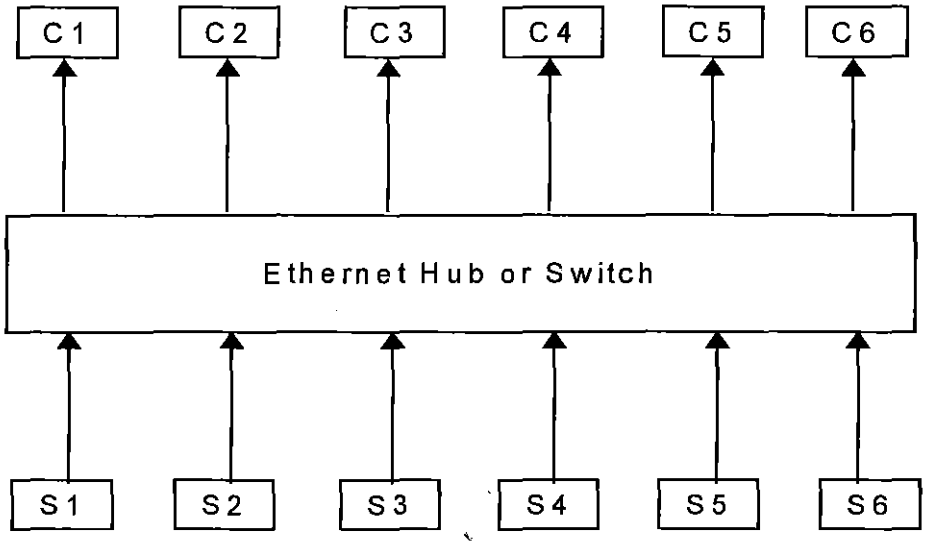

Figure 6.16 - Test Bed for Second Test Series

General Case

6.10.2 Test Bed

The basic test bed for the second series of tests is shown in Figure 6.16. Six identical PCs were configured as Netware 4.11 file servers and arranged to simultaneously

transfer files to another set of six identical PCs configured as Netware 4.11 clients. As in the first series of tests, the traffic is generated by means of DOS batch files containing re-entrant code that produce simultaneous and continuous file transfers.

The salient feature of this revised test bed, compared with the test bed for the first series of tests, is that each client is receiving files from a dedicated server. Each client logs on to its respective server so that the files are copied from Servers S1, S2, S3, S4, S5, S6, S7 and S8 to Clients C1, C2, C3, C4, C5, C6, C7 and C8 respectively.

It should be noted that the second series clients designated C1 to C5 were different from the five clients (also designated C1 to C5) used in the Preliminary Tests. The clients used in the second series of tests were Intel 486DX, 66 Mhz and the servers were Intel Pentiums, 166 Mhz. Clients and servers were fitted with Intel EtherExpress PRO/100+ PCI network interface cards, capable of operating at either 10 or 100 Mhz and at either full or half duplex.

Table 6.8 - Clients and Servers used in the Second Test Series

Processor	Clients: Intel 486DX 66 Mhz Servers: Intel Pentium 166 Mhz
Network Interface Cards	Intel EtherExpress PRO/100+ PCI 10/100 Mbits/sec
Cables	Category 5, UTP, maximum length: 12 m
Local Operating System	Windows 95
Network Operating System	Netware 4.11

6.10.3 LANalyzer in the Foreground

During the first series of tests, LANalyzer was run in the background, with the file transfers taking place in the foreground. In the second series, LANalyzer was run in the foreground and the file transfers took place in the background. For a discussion of this issue refer 6.11.5 "Conclusions from Test 7".

6.11 Test 7: Ethernet hub with one client/server pair at a time

6.11.1 Setup

Seven client/server pairs designated CS1, CS2, CS3, CS4. CS5, CS6 and CS7 were connected to an Ethernet hub, one at a time. Each client (Cn) was programmed to continuously copy test files from its respective server (Sn).

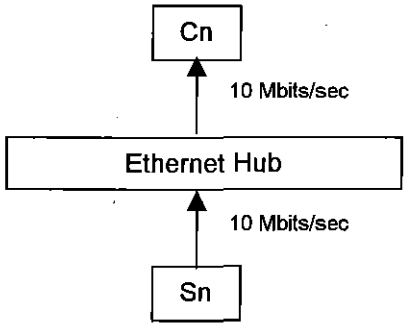

Figure 6.17 - Test Bed for Test 7

6.11.2 Purpose

To establish a baseline for each client/server pair under hub operation, taking into account the effect of time.

6.11.3 Procedure

Each test run commenced with the activation of a DOS batch file that initiated the file transfer from server to client. LANalyzer (and hence the data capture) was started only after the file transfer had commenced.

Client/server pairs were run singly and the rate at which files were transferred was monitored and recorded. Three test runs were conducted for each client/server pair[6].

[6] Note that, in the context of the second series of tests, references to a "client" or to a "server" may also refer to a client/server pair.

6.11.4 Results

The results are tabulated and plotted in the tables and graphs of Figure 6.18. For example, Part (a) of Figure 6.18 is interpreted as follows:

Five minutes into "Test Run a" (abbreviated "Test a" in the tables and graphs), Server 1 was transferring files to Client 1 at an average transfer rate of 149 Kbytes/sec. After 45 mins, at the completion of "Test Run a", the transfer rate was down slightly to 147 Kbytes/sec. Similarly, during two subsequent test runs, "b" and "c", the transfer rates dropped slightly from 146 and 152 Kbytes/sec to 143 and 147 Kbytes/sec, respectively.

Effect of Time

In all cases LANalyzer would show a sharp peak immediately after it was launched. After 30 seconds, the transfer rate would begin to settle down, decreasing slowly, so that between 5 and 15 mins, there would typically be a decrease of about 2%. After the first 15 mins the transfer rate would then remain, indefinitely, within a 1% band.

Repeatability

Each graph in Figure 6.18 shows three curves corresponding to three test runs conducted for each client/server pair. The test runs were performed under identical conditions but separated in time. The time intervals between test runs varied between one hour and several days. The degree of repeatability is shown by the closeness of the three curves to each other.

The graphs show that, generally, after 15 mins operation, the curves converge to within a 5% band. An exception is CS5, which took 30 mins to converge.

Figure 6.18 - Single Client/Server Pairs on an Ethernet Hub

Time	Kbytes/sec		
(mins)	Test a	Test b	Test c
5	149	146	152
10	149	150	149
15	146	147	147
20	144	146	
25	146	144	
30	147	143	
35	147	143	
40	146	143	
45	147	142	
50		143	

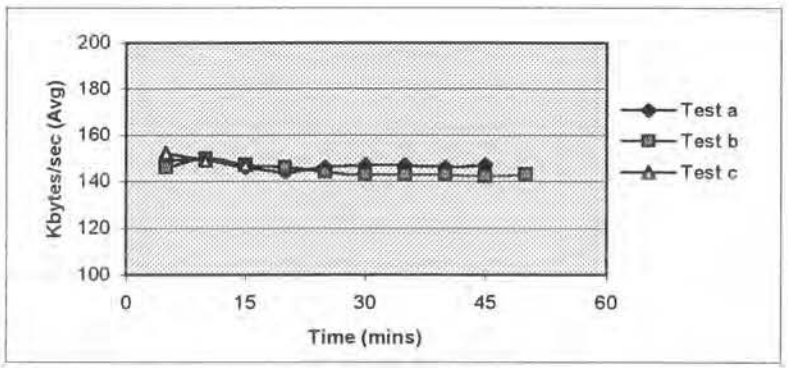

(a) Client/Server Pair CS1

Time	Kbytes/sec		
(mins)	Test a	Test b	Test c
5	161	151	157
10	159	151	155
15	158	149	154
20	154	149	153
25	152	147	153
30	152	147	153
55			150
70			150
85			150
2hrs			150

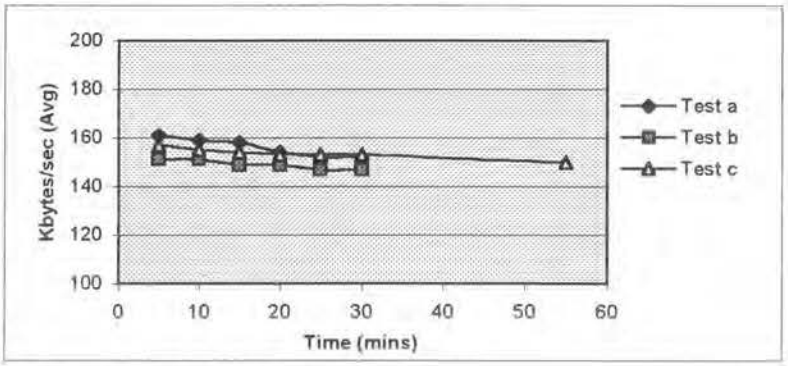

(b) Client/Server Pair CS2

Time	(Kbytes/sec)		
(mins)	Test a	Test b	Test c
5	149	151	151
10	146	150	145
15	146	149	144
20	147	149	146
25	147	150	147
30	149	150	147
35	151	149	146
40			146
45			147
15hrs	149		

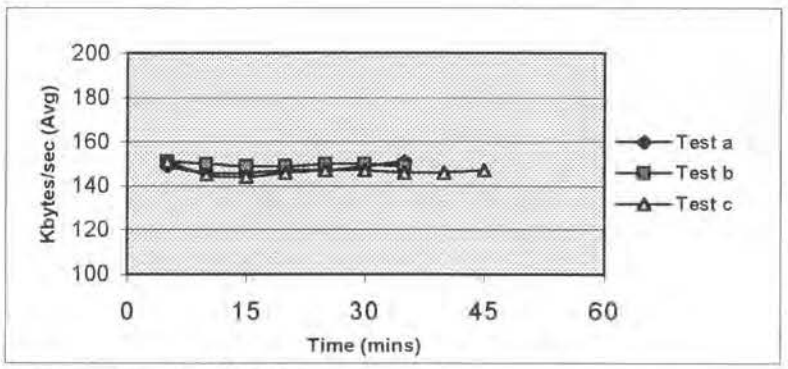

(c) Client/Server Pair CS3

Time	(Kbytes/sec)		
(mins)	Test a	Test b	Test c
5	136	134	134
10	135	130	131
15	132	128	130
20	132	128	131
25	132	129	132
30	132	128	133
35	132	129	132
40		129	132
45		129	
15hrs			130

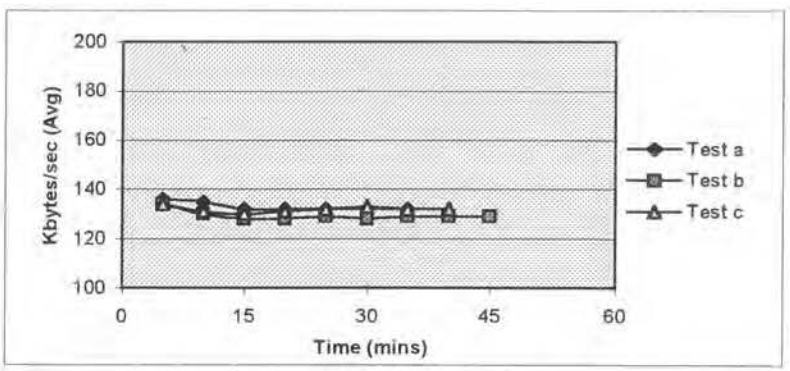

(d) Client/Server Pair CS4

Time	Kbytes/sec		
(mins)	Test a	Test b	Test c
5	196	151	165
10	176	150	161
15	167	149	155
20	164	150	156
25	162	151	155
30	161		155
35	159		155
40	158		154
45			154
50			154
55			154
2 hrs		152	
18 hrs	152		

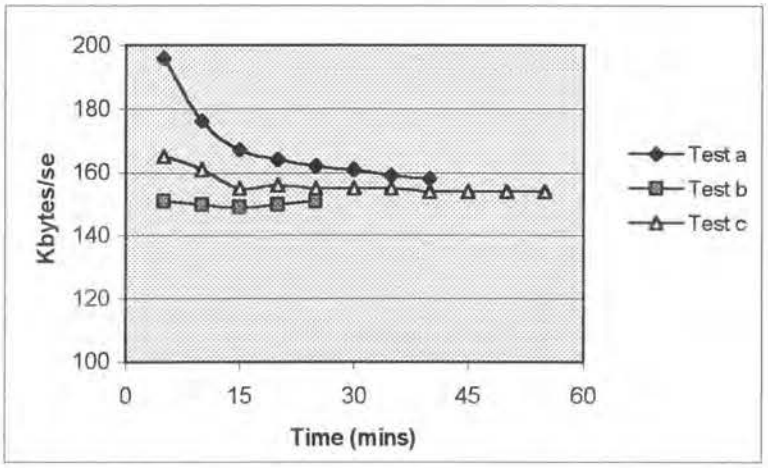

(e) Client/Server Pair CS5

Time	Kbytes/sec		
(mins)	Test a	Test b	Test c
5	157	149	148
10	155	150	152
15	152	150	150
20	150	148	149
25	151	148	149
30	150	149	148
35	150	149	148
40	150	149	148
45	150	149	147
50	149		149
55	149		
21 hrs	148		

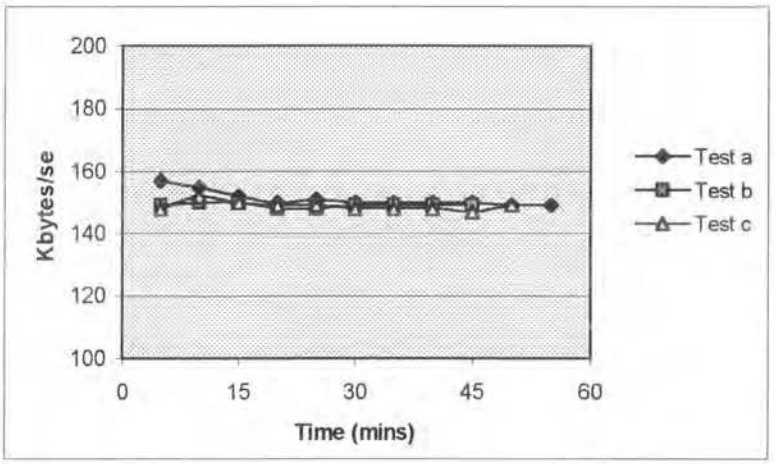

(f) Client/Server Pair CS6

Time	Kbytes/sec		
(mins)	Set 1	Set 2	Set 3
5	175	154	156
10	160	153	154
15	157	151	155
20	154	151	154
25	153	152	154
30	150		154
35	150		

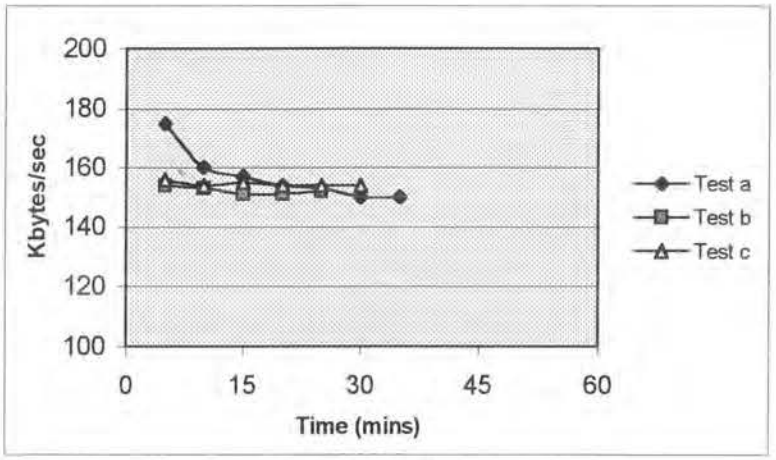

(g) Client/Server Pair CS7

Baseline Parameters

By associating a set of performance parameters with each client/server pair, the baseline can provide a basis for quantitatively comparing the machines.

The characteristic performance of each client/server pair under hub operation is shown in Figure 6.18 - Single Client/Server Pairs on an Ethernet Hub. After an examination of the graphs and the tables in Figure 6.18 the most suitable baseline parameters were deemed to be-

- The maximum transfer rate measured at the 15-minute mark. For example, the table for CS3 in Figure 6.18 shows that after 15 mins operation, the throughputs for client/server pair CS3 when connected to a hub were 146, 149 and 144 Kbytes/sec for Test Runs 1, 2 and 3 respectively. The 15-min Maximum (Max) for CS3 under hub operation is therefore 149 Kbytes/sec.

- The minimum transfer rate measured during the three test runs. For example, Table CS3 in Figure 6.18 (c) shows that the minimum throughput (Min) recorded during the three test runs is 144 Kbytes/sec. Therefore Min for CS3 under hub operation is 144 Kbytes/sec.

The steady state value to which the transfer rate converges during repeated test runs. The steady state values were determined by examining the tables and graphs in Figure 6.18. For example, in the graph for CS3 in Figure 6.18, the curves corresponding to Test Sets 1, 2 and 3 converge to 149 Kbytes/sec. The Steady State value for CS3 under hub operation in Figure 6.19 is therefore 149 Kbytes/sec.

The values of the baseline parameters for each client/server pair under hub operation are tabulated and plotted in Figure 6.19. For comparison purposes, the baseline parameters arrived at under switch operation in Test 9 are shown in Figure 6.20. It can be seen that client/server pair CS4 is consistently slower than the other pairs. As the difference is significant, about 10%, it was deemed necessary to normalise the results from subsequent tests. The Steady State values, which have a high degree of time independence, were used as a basis for normalisation – refer Test 8.

Client/Server Pair	Throughput (Kbytes/s)		
	Max	Min	Steady State
CS1	147	142	144
CS2	158	147	150
CS3	149	144	149
CS4	132	128	130
CS5	167	149	154
CS6	152	147	149
CS7	157	150	152

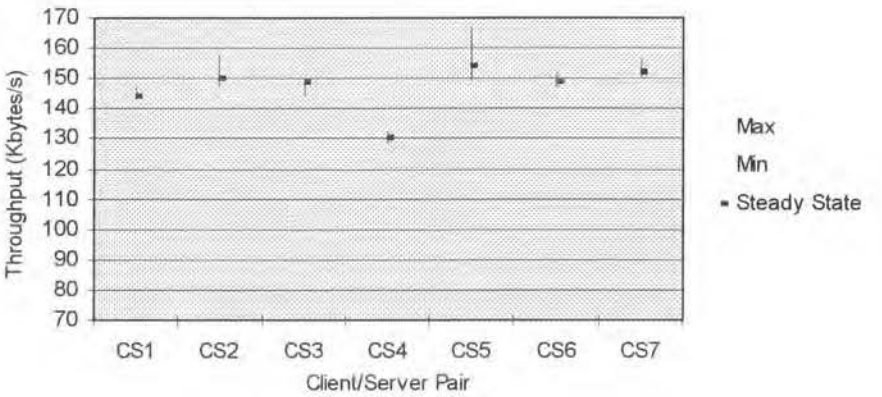

Figure 6.19 - Baseline Parameter Limits (shared Ethernet)

Client/Server Pair	Throughput (Kbytes/s)		
	Max	Min	Steady State
CS1	87	84	85
CS2	90	81	86
CS3	88	83	85
CS4	77	70	73
CS5	89	87	88
CS6	93	84	85
CS7	91	84	87

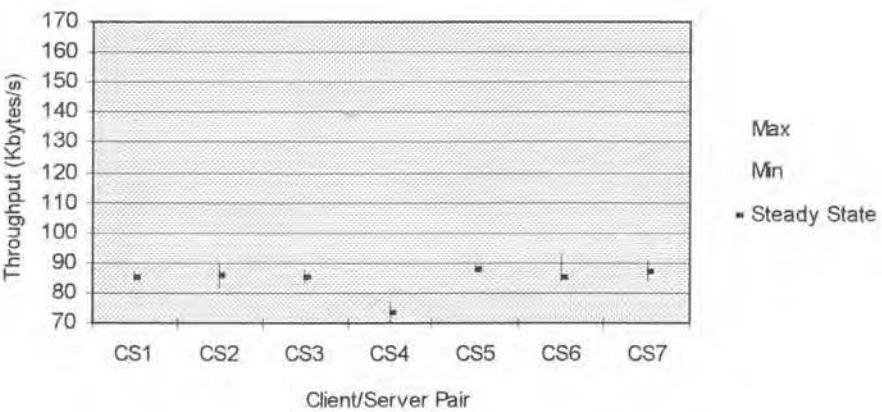

Figure 6.20 - Baseline Parameter Limits (switched Ethernet)

6.11.5 Conclusions from Test 7

1. Effect of Time

LANalyzer readings are reliable after 15 mins operation. Readings taken during the first 5 mins should be ignored.

2. Effect of running LANalyzer in the Foreground

As previously stated, LANalyzer was run in the background during the first series of tests and in the foreground during the second series.

Running LANalyzer in the foreground enabled it to be run with fewer interruptions, thus removing an experimental variable. For example, when during the Second Test Series, the threshold limits for the LANalyzer's alarms were checked, LANalyzer did not have to be brought into the foreground because it was already in the foreground.

It is likely can that running LANalyzer in the foreground contributed to the high degree of repeatability achieved in the second series of tests. This assessment is subjective because a quantitative assessment of repeatability would require many more test runs and is outside the scope of the experiment.

On the other hand, with file transfers taking place in the background, throughput was significantly reduced throughout the Second Test Series, even though attempts were made (more machines, faster machines) to increase the throughput and to put more pressure on the relatively fast Ethernet switch.

As an example, during Test 1 when file transfers were taking place in the foreground, Client C1 achieved a throughput of 220 Kbytes/sec (refer Table 6.2 - Individual Client Performance (shared Ethernet)). During Test 7, during which file transfers were taking place in the background, client/server pair CS1 only achieved 146 Kbytes/sec (refer Figure 6.18 - Single Client/Server Pairs on an Ethernet Hub, CS1, Test a, 15 min mark), even though the clients had been upgraded for the second series of tests.

3. Effect of using the same client to both generate and monitor traffic

Regardless of whether it is the better to run the file transfers in the foreground or background, the fact that the file transfer rate decreased so markedly when switched to the background (by as much as 64%, depending on initial load) shows that the Intel 486s are labouring heavily under the combined load of LANalyzer and file transfer. This is also brought out by the throughputs between single client/server pairs being low, e.g. 149 Kbytes/sec for CS1 at the beginning of Test Run 1 (refer Figure 6.18 - Single Client/Server Pairs on an Ethernet Hub) compared to 220 Kbytes/sec for C1 in Test 1 (refer Table 6.2 - Individual Client Performance).

Although the repeatability and consistency checks carried out during the Second Test Series suggest that the LANalyzer readings are not affected, it may be argued, that just as the LANalyzer is affecting the file transfer, so the file transfer may be affecting the LANalyzer.

This is an area to be considered for further experimental work – refer 6.19 "Suggestions for future Experimental Work".

6.12 Test 8: Ethernet hub with combinations of client/server pairs

6.12.1 Setup

Five different combinations of client/server pairs were connected to the hub. The combinations consisted of 2, 3, 4, 5 and 6 client/server pairs. The pairs in a combination were arranged to simultaneously transfer files from a server to its respective client.

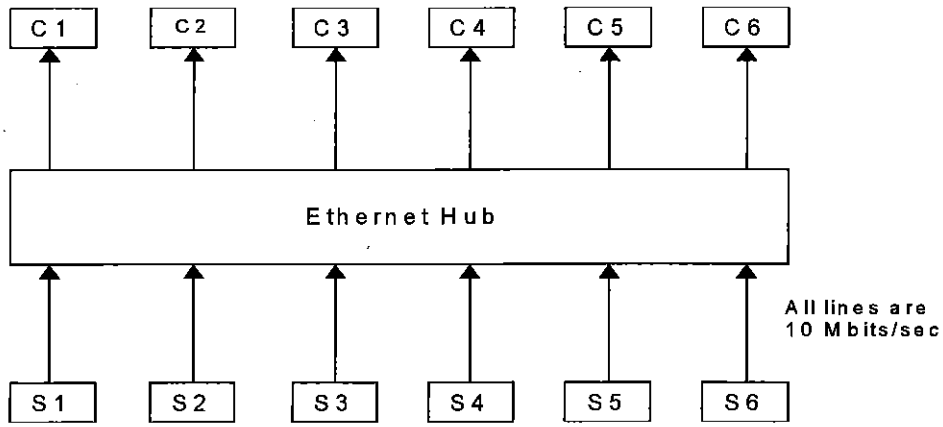

Figure 6.21 - Test Bed for Test 8

Monitoring Network Traffic

In this setup there is only one collision domain, so that for the purposes of this test only one LANalyzer would be required to monitor the network traffic. Nevertheless, in this test, a copy of LANalyzer is run on every client for consistency with Test 10. In Test 10, where a switch is used instead of a hub, each client is in a different collision domain, so that LANalyzer does need to be run on each client. Running LANalyzer on every client in both Test 8 and Test 9 allows the results from the tests to be legitimately compared.

6.12.2 Purpose

1. To progressively load an Ethernet hub, that is, the same purpose as for Test 2, but also taking account of the effect of time.

2. To test LANalyzer for consistency.

6.12.3 Procedure

File transfers were started between each client and its respective server. Then the LANalyzers were started simultaneously (within seconds of each other).

6.12.4 Results

The results from Test 8 are summarised in Figure 6.22 to Figure 6.26 below. The table in Figure 6.22 is interpreted as follows: Five minutes after the LANalyzers on Clients C3 and C5 were started, the traffic on the network as measured by the two LANalyzers was 325 and 322 Kbytes/sec. The average bus throughput was therefore (325 + 322)/2 = 324 Kbytes/sec. The average throughput per client was 324/2 = 162 Kbytes/sec.

The adjacent graph shows how the average throughput per client (Avg/Client) varied over the 50-minute test period.

Time (mins)	Kbytes/sec			
	CS3	CS5	Avg	Avg/Client
5	325	322	324	162
10	311	312	312	156
15	307	308	308	154
20	305	305	305	153
25	305	305	305	153
30	304	303	304	152
35	302	302	302	151
40	303	303	303	152
45	301	301	301	151
50	300	300	300	150

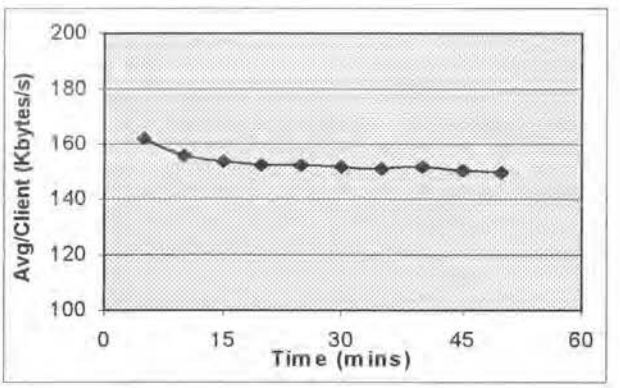

Figure 6.22 - Two Clients on a Hub

Time (mins)	Kbytes/sec				
	CS1	CS2	CS3	Avg	Avg/Client
5	402	408	412	407	136
10	399	403	406	403	134
15	399	402	403	401	134
20	403	403	406	404	135
25	402	403	404	403	134
30	403	405	404	404	135
35	402	403	404	403	134

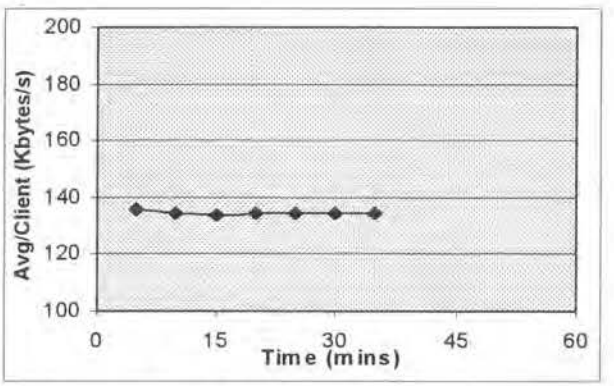

Figure 6.23 - Three Clients on a Hub

Time	Kbytes/sec					
(mins)	CS1	CS2	CS3	CS4	Avg	Avg/Client
5	518	509	503	507	509	127
10	502	499	495	498	499	125
15	498	493	491	493	494	123
20	493	492	490	491	492	123
25	492	490	488	489	490	122
30	491	489	489	490	490	122
35	492	491	490	490	491	123
40	491	490	489	489	490	122
45	490	490	489	488	489	122

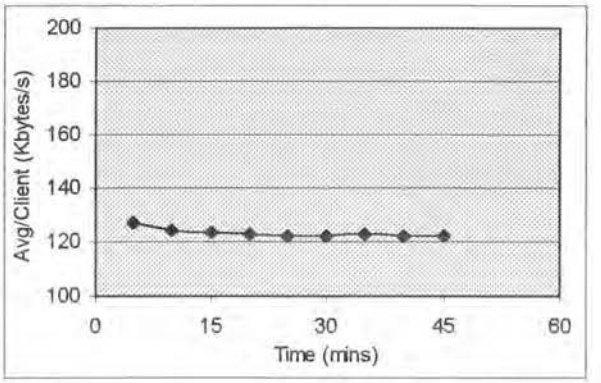

Figure 6.24 - Four Clients on a Hub

Time	Kbytes/sec						
(mins)	CS1	CS3	CS4	CS5	CS6	Avg	Avg/Client
5	616	608	601	596	592	603	121
10	601	597	595	594	590	597	119
15	598	594	591	591	591	594	119
20	592	589	588	590	588	590	118
25	591	590	589	589	588	590	118
30	590	589	587	587	587	588	118
35	589	588	587	587	586	588	117
40	589	588	586	587	586	588	117

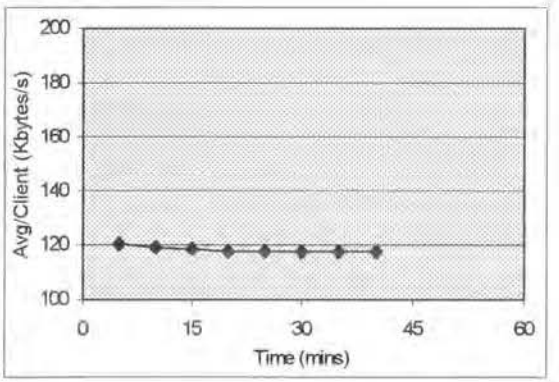

Figure 6.25 - Five Clients on a Hub

Time	Kbytes/sec							
(mins)	CS1	CS3	CS4	CS5	CS6	CS7	Avg	Avg/Client
5	713	700	690	687	679	677	691	115
10	683	677	672	673	669	669	674	112
15	671	669	665	664	662	662	666	111
20	668	666	663	662	660	659	663	111
25	665	663	662	662	660	660	662	110
30	665	663	662	662	660	660	662	110
35	664	662	660	661	659	660	661	110
40	664	662	660	661	660	660	661	110

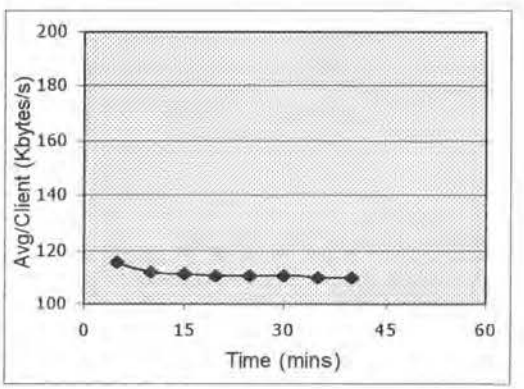

Figure 6.26 - Six Clients on a Hub

Effect of Time

Looking at the graph in Figure 6.22, it can be seen that as time progressed, the average throughput for CS3 and CS5 decreased slightly, settling down to a steady state value of 150 Kbytes/sec. This steady state value, and the steady state values for the other 4 client/server combinations tested, are listed in . Both the raw and the normalised values are listed.

Effect of loading an Ethernet Hub

The effect of increasing the number of active clients (in this case client/server pairs) on a shared Ethernet is illustrated by the family of curves in Figure 6.27, which combines the curves from the preceding graphs.

Test 8 only deals with combinations of two or more client/server pairs. To allow a comparison to be made with a typical single client/server pair, a curve from the Test 7 results was included in Figure 6.27 (the "1-client" curve). The Test 7 results include 21 curves (7 client/server pairs, 3 test runs each) to choose from. The chosen curve shows CS5 performance as measured during test run 3 (refer Figure 6.18(e)). CS5 was chosen because because the "2-clients" combination also includes CS5 (see Figure 6.22).

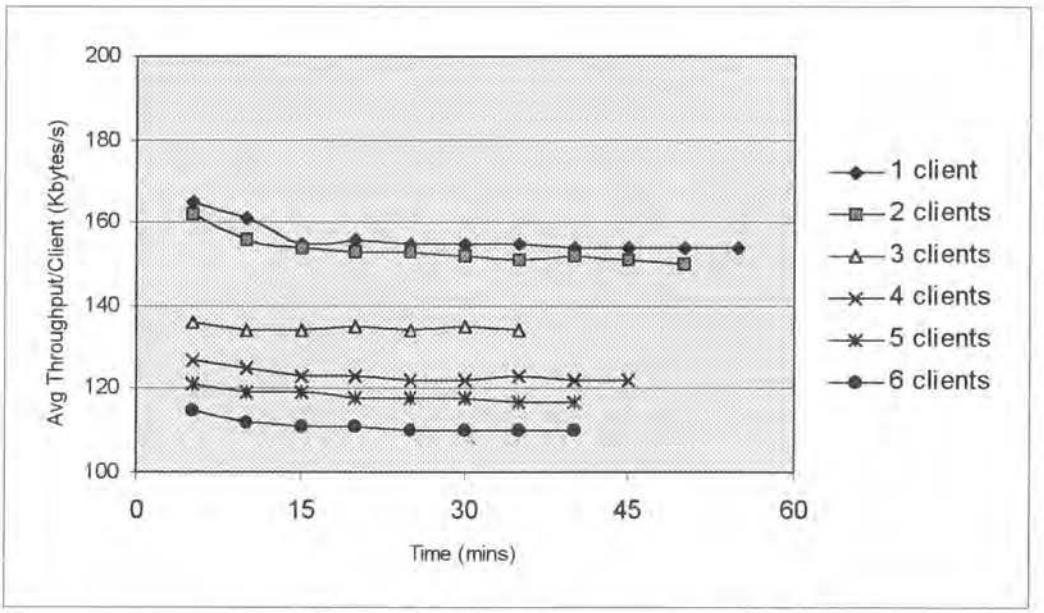

Figure 6.27 - Adding more Clients to a Hub

As was the case in Test 2, the throughput per client decreased as the number of clients increased. For example, with two clients the steady state throughput per client was

150 Kbytes/sec, with three clients it was down to 134 Kbytes/sec. The relationship between steady state throughput per client and the number of clients is plotted in Figure 6.28.

The shape of the "Hub normalised" curve (steady throughput around one and two clients, then a slow decrease as the number of clients increases) is similar to curves obtained by Boggs, Mogul and Kent (1988, Figure 3-3, p.229) when they measured the throughput on a standard Ethernet and drew a graph of "Mbits/sec" versus "Number of Hosts". This result reinforces are similar result obtained in Test 2.

The values shown in Table 6.9 are steady state values derived from the family of throughput-versus-time graphs in Figure 6.27. These steady state values are designated "Raw" in Table 6.9 because they do not take account of the throughput capacity of individual machines. The "Raw" figures are therefore normalised as explained below.

Table 6.9 - Steady State Throughput for Combinations of Clients

Number of Clients	Throughput / Client (Kbytes/sec)			
	Hub		Switch	
	Raw	Normalised	Raw	Normalised
1	147	147	84	84
2	150	145	87	86
3	134	133	80	83
4	122	125	82	83
5	117	118	84	85
6	110	110	83	83

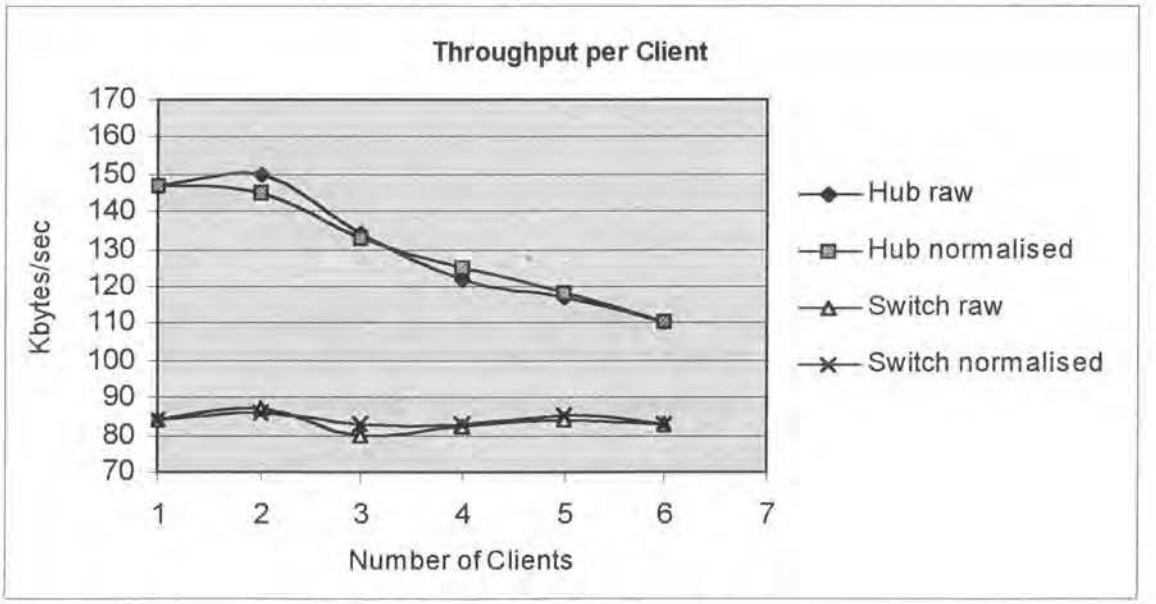

Figure 6.28 - Adding more Clients to a Network

For comparison purposes, the above table and graph also include the results obtained under switch operation in Test 10.

Normalising the Throughput

Because of significant variations in throughput capacity of individual machines (refer Tests 7 and 9) the average throughput per client for each combination of client/server pairs was normalised using the steady state throughput of individual client/server pairs. The Steady State throughputs of individual client/server pairs are listed in Figure 6.19.

From Figure 6.19 the average Steady State throughput for the 7 client/server pairs

= (144 + 150 +149 +130 + 154 + 149 + 152) / 7 = 147 Kbytes/sec.

The normalised, average throughputs per client for each client/server combination are as follows –

For CS3/CS5:
150 x (2 x 147) / (149 + 155) = 150 x 0.97 = 145 Kbytes/sec.

For CS1/CS2/CS3:
134 x (3 x 147) / (144 + 150 + 149) = 134 x 0.99 = 133 Kbytes/sec.

For CS1/CS2/CS3/CS4:
122 x (4 x 147) / (144 +150 + 149 + 130) = 122 x 1.02 = 125 Kbytes/sec

For CS1/CS3/CS4/CS5/CS6:
117 x (5 x 147) / (144 + 149 + 130 + 154 + 149) = 117 x 1.01 = 115 Kbytes/sec.

For CS1/CS3/CS4/CS5/CS6/CS7:
110 x (6 x 147) / (144 + 149 + 130 + 154 + 149 + 152) = 110 x 1.00
= 110 Kbytes/sec (i.e. unchanged by normalisation).

The normalised values are tabulated and plotted in Figure 6.28 - Adding more Clients to a Network. It can be seen the effect of normalisation is to push up the throughput of those combinations that include client/server pair CS4, and to push down the throughput of those combinations that do not include CS4, thus compensating for the relatively low throughput from CS4.

LANalyzer Consistency

Since the procedure calls for LANalyzer to be run on every client (see "Monitoring Network Traffic" in Section 6.12.1) it was also possible to check whether the LANalyzer readings are consistent. The redundant LANalyzers provided a cross-check on each other.

Spot tests indicated that generally, LANalyzer readings are consistent within 1%. For example, Figure 6.26 shows that after 30 mins operation, the six LANalyzer readings varied between 660 and 665 Kbytes/sec. This is a 0.8% variation, which is considered adequate.

6.12.5 Conclusions from Test 8

1. As in Test 2, the MAC network behaved predictably with the data transfers for each client slowing down, as bus contention and the number of collisions increased. The results for hub operation once again match the results obtained by Boggs, Mogul and Kent (1988, Figure 3-3, p.229).

2. Readings of Kbytes/sec (Avg) provided by LANalyzer are generally consistent and do not vary by more than 1%, once the average has had time to build up.

6.13 Test 9: Ethernet switch with one client/server pair at a time

6.13.1 Setup

Seven client/server pairs connected to a switch one at a time, i.e. same as Test 7 except that a switch is used instead of the hub.

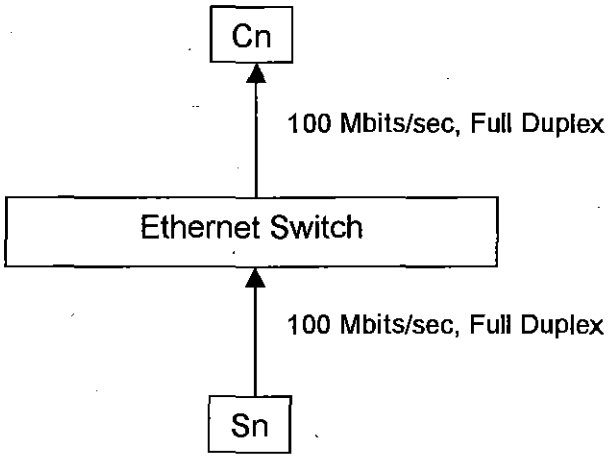

Figure 6.29 - Test Bed for Test 9

6.13.2 Purpose

To establish baselines for the client/server pairs under switch operation.

6.13.3 Procedure

As in Test 7, client/server pairs were run singly, with three test runs for each pair.

6.13.4 Results

The results from Test 9 are tabulated and plotted in Figure 6.30.

Figure 6.30 - Single Client/Server Pairs on a Switch

Time	Kbytes/sec		
(mins)	Test a	Test b	Test c
5	90	86	90
10	87	85	88
15	85	84	87
20	85	85	86
25		84	86
30	85	85	85
35			85
40			84
55	85		

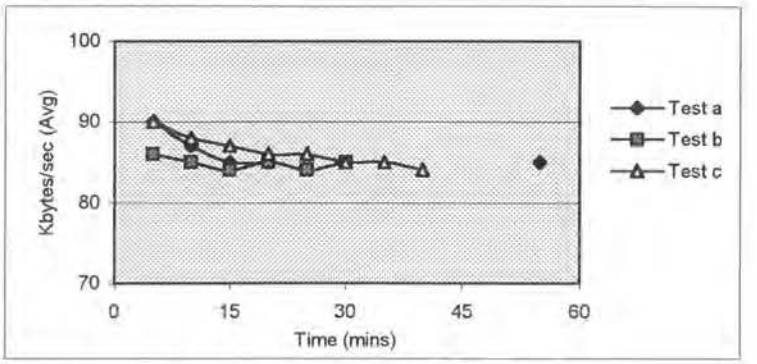

(a) - Client/Server Pair CS1

Time	Kbytes/sec		
(mins)	Test a	Test b	Test c
5	90	93	84
10	89	88	81
15	89	87	81
20	88	87	83
25	88	86	83
30	88		83
35	88		84
40	88		84
45	88		
55	87		

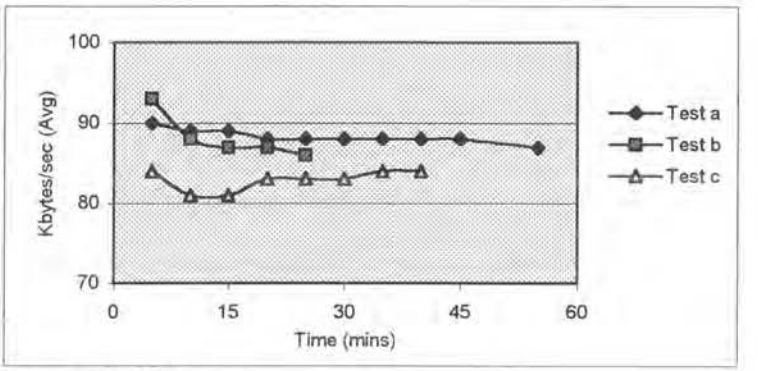

(b) Client/Server Pair CS2

Time	Kbytes/sec		
(mins)	Test a	Test b	Test c
5	86	92	86
10	83	89	85
15	84	88	86
20	85	87	86
25	85	86	86
30	84	86	87
35	84	86	86
40	84	86	86
45		86	86
16 hrs	86		

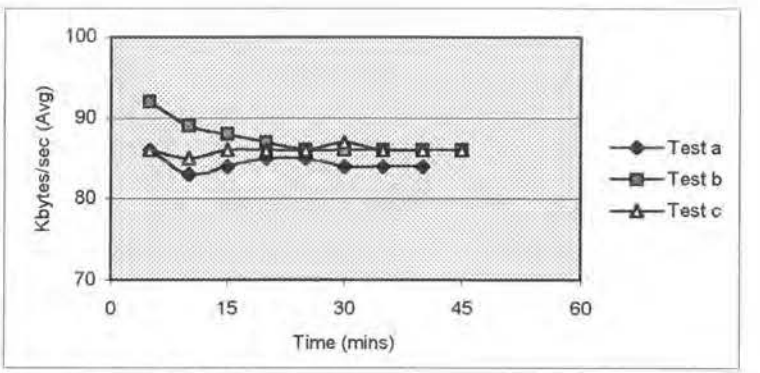

(c) Client/Server Pair CS3

Time	Kbytes/sec		
(mins)	Test a	Test b	Test c
5	75	74	70
10	76	72	
15	76	72	72
20	77	72	71
25	77	72	
30	76		
35	76		
40	76		
2 hrs		72	
19 hrs	73		

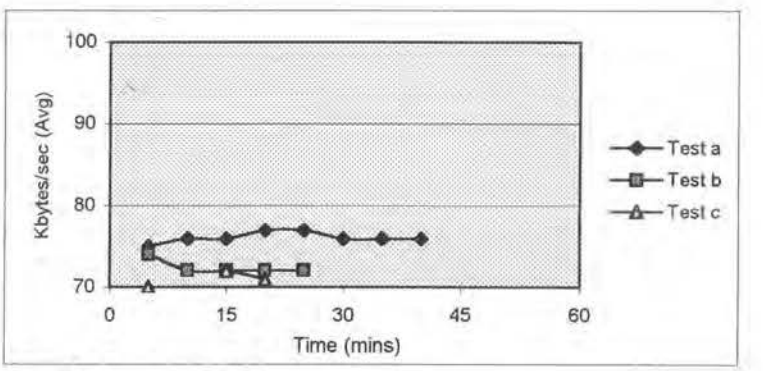

(d) Client/Server Pair CS4

Page 115 does not exist in the original thesis

Repeatability

Again, as in Test 7, the three curves in each family are close to each other and after 15 mins operation they converge to within a 10% band.

Baseline Parameters

The baseline parameters for switch operation are tabulated and plotted in Figure 6.20. They were derived in the same way as the baseline parameters for hub operation, viz.:

- **The maximum transfer rate measured at the 15-minute mark**. For example, the table for CS3 in Figure 6.30(c) shows that after 15 minutes of operation, the throughputs for client/server pair CS3 connected to a switch were 84, 88 and 86 Kbytes/sec for Test Runs a, b and c respectively. The 15-min Maximum (Max) for CS3 under switch operation is therefore 88 Kbytes/sec.

- **The minimum transfer rate measured during the three test runs**. For example, table in Figure 6.30(c) shows that the minimum throughput (Min) recorded during the three test runs is 83 Kbytes/sec. Therefore Min for CS3 under switch operation is 83 Kbytes/sec.

- **The steady state value to which the transfer rate converges during repeated test runs**. The steady state values were picked by examining the tables and graphs in Figure 6.20. For example, in the graph for CS3 in Figure 6.30(c), the curves corresponding to Test Runs a, b and c converge to 85 Kbytes/sec. The Steady State value for CS3 under switch operation is therefore 85 Kbytes/sec.

As was the case in the shared Ethernet environment of Test 7, client/server pair CS4 was considerably slower than the other pairs. This means that the results for the switched Ethernet will also need to be normalised.

Variations in individual Throughput

For the shared Ethernet, refer Figure 6.19 - Baseline Parameter Limits (shared Ethernet), the variation in individual throughput = (154-130)/147 = 16%.

For the switched Ethernet, refer Figure 6.20 - Baseline Parameter Limits (switched Ethernet), the variation in individual throughput = (87-73)/84 = 17%.

Where 147 Kbytes/sec and 84 Kbytes/sec are the average steady state throughputs for the shared and switched Ethernets respectively.

6.13.5 Conclusions from Test 9

1. Effect of Time

Figure 6.30 shows the variations in throughput over time. The shape of the throughput-versus-time curves is similar to the ones obtained for hub operation in Test 7, i.e. an initial peak followed by a slow decrease down to a steady state value. However, the pattern is not as predictable as it was under hub operation. Indeed for four (out of 21) test runs, the throughput actually increases over time.

The results suggest that the shared channel provided by an Ethernet hub exercises a steadying influence on data transfer rates.

2. Throughput

Once again, with this type of experimental setup, the throughput under switch operation is considerably less than under hub operation. This may be attributed to –

- the additional overhead of a switch – a switch needs to unpack a frame and read the destination address, whereas a hub simply forwards a frame to all of its ports.
- with only one client on a shared Ethernet, frames are not delayed when accessing the Ethernet bus.

The variations in individual throughput, this time based on the steady state figures, were approximately the same for shared and switched Ethernets (16% and 17% respectively). This contrast with the Test 3 results, which were based on a 15-minute reading. The result suggests that the effect of time should be taken into account when measuring throughput.

6.14 Test 10: Ethernet switch with combinations of client/server pairs

6.14.1 Setup

Combinations of 2, 3, 4, 5 and 6 client/server pairs connected to switch, i.e. same as for Test 8, except that a switch is used instead of a hub.[7]

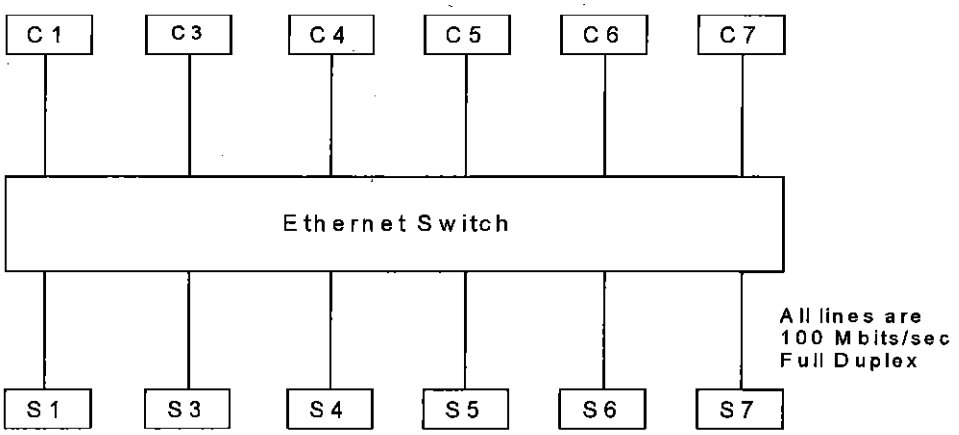

Figure 6.31 - Test Bed for Test 10

6.14.2 Purpose

To determine the effect, if any, of increasing the number of clients on a switched Ethernet. This is a repeat of Test 4, but with more stringent test conditions.

In Test 4 the throughput per client decreased steadily as the number of clients increased from one to four, leading to speculation that this decrease was due to the server bottleneck (refer 6.6.5 "Conclusions from Test 4").

In this test, as in all other Series Two Tests, the server bottleneck has been removed.

6.14.3 Procedure

Same as for Test 8, that is, five test runs each presenting a heavier load to the network.

[7] Client C2 developed a fault while tests were in progress. Client/server pair CS2 was consequently replaced by CS7.

6.14.4 Results

The results from Test 10 are tabulated and plotted below. The table in Figure 6.34 - Four Clients on a Switch, for example, is interpreted as follows: Five minutes after the LANalyzers on Clients C1, C3, C4 and C5 were started, the traffic on the four network segments to which C1, C3, C4 and C5 were connected was 85, 93, 76 and 87 Kbytes/sec. The average throughput per client was therefore (85 + 93 + 76 + 87)/4 = 85 Kbytes/sec.

Looking at the graph in Figure 6.34, it can be seen that as time progressed the average throughput per client decreased slightly, settling down to a steady state throughput of 82 Kbytes/sec.

Time (mins)	Kbytes/sec		
	CS1	CS3	Avg/Client
5	84	89	87
10	85	88	87
15	85	87	86
20	85	88	87
25	85	88	87
30	85	88	87
35	85	87	86

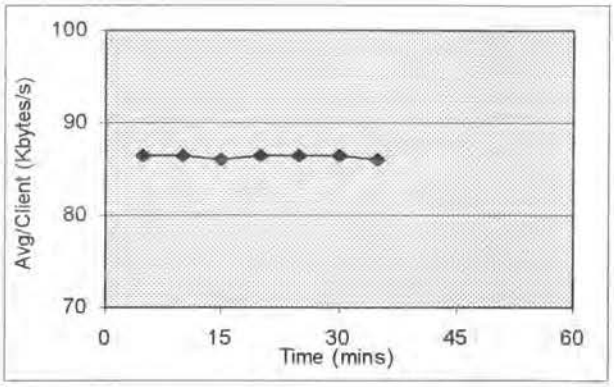

Figure 6.32 - Two Clients on a Switch

Time (mins)	Kbytes/sec			
	CS1	CS3	CS4	Avg/Client
5	86	82	70	79
10	85	83	73	80
15	84	83	73	80
20	84	83	73	80
25	84	83	73	80
30	85	83	72	80
35	84	83	73	80
40	85	83	73	80

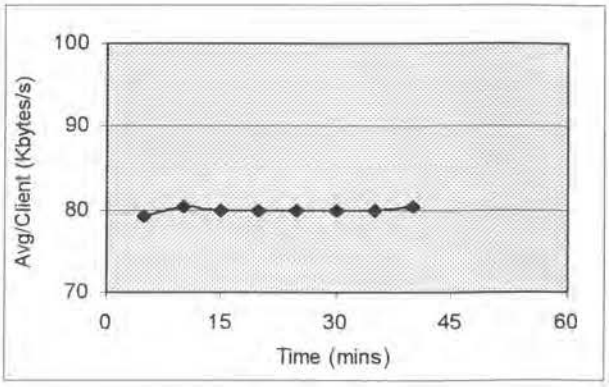

Figure 6.33 - Three Clients on a Switch

Time	Kbytes/sec				
(mins)	CS1	CS3	CS4	CS5	Avg/Client
5	85	93	76	87	85
10	83	90	73	87	83
15	84	88	73	86	83
20	82	89	73	86	83
25	82	88	73	86	82
30	82	89	73	86	83
35	82	88	73	87	83
40	82	88	73	86	82
45	82	88	72	86	82
20 hrs	84	87	73	87	83

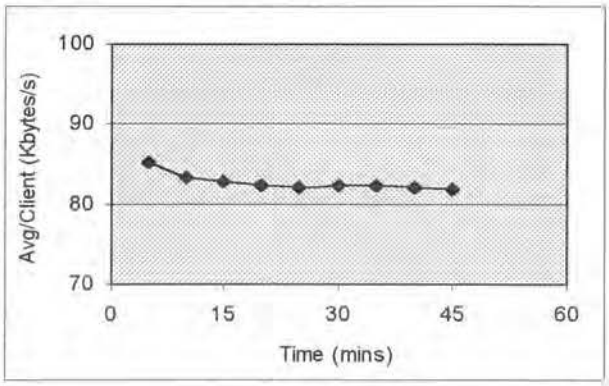

Figure 6.34 - Four Clients on a Switch

Time	Kbytes/sec					
(mins)	CS1	CS3	CS4	CS5	CS7	Avg/Client
5	83	84	76	92	90	85
10	85	84	72	88	92	84
15	84	85	73	89	92	85
20	85	85	73	89	91	85
25	85	85	73	88	90	84
30	85	85	73	88	89	84
35	85	85	72	88	89	84
40	85	85	73	88	89	84
45	85	86	73	87	89	84
50	85	86	72	87	89	84
55	85	86	73	87	89	84
60	85	86	73	87	89	84

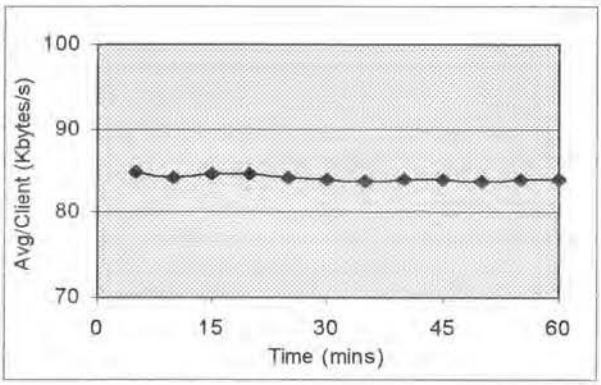

Figure 6.35 - Five Clients on a Switch

Time	Kbytes/sec						
(mins)	CS1	CS3	CS4	CS5	CS6	CS7	Avg/Client
5	81	88	70	86	88	82	83
10	84	83	73	83	87	83	82
15	85	83	73	84	87	84	83
20	85	82	73	85	87	84	83
25	85	82	73	85	87	84	83
30	85	83	73	84	87	84	83
35	84	83	73	84	87	85	83
60	84	84	74	86	87	85	83

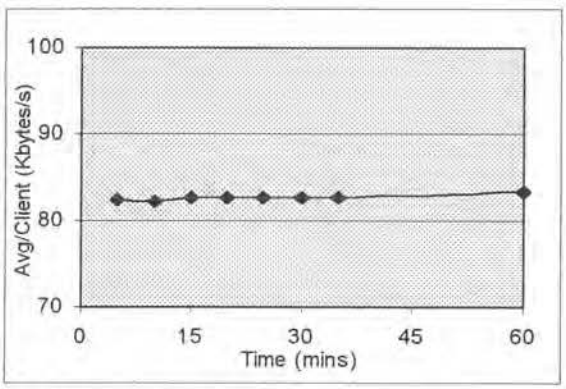

Figure 6.36 - Six Clients on a Switch

This steady state throughput, together with the steady state throughputs for the other four combinations of client/server pairs is listed in Table 6.9 - Steady State Throughput for Combinations of Clients, where it represents the "Raw" (pre-normalisation) figure for a four-client combination.

Effect of Time

The effect of time on the throughput of combinations of clients is similar to its effect on the throughput of individual clients – again we are getting an initial peak followed by convergence to a steady state band. However, the curves are flatter – compare, for example, Figure 6.35 - Five Clients on a Switch, with Figure 6.30 - Single Client/Server Pairs on a Switch. The flatness of the curve in Figure 6.35 is due to the averaging effect of multiple clients.

Similarly, the averaging effect makes the curves for combinations of clients look the same, for both shared and switched Ethernets – compare, for example, Figure 6.25 - Five Clients on a Hub, with Figure 6.35 - Five Clients on a Switch.

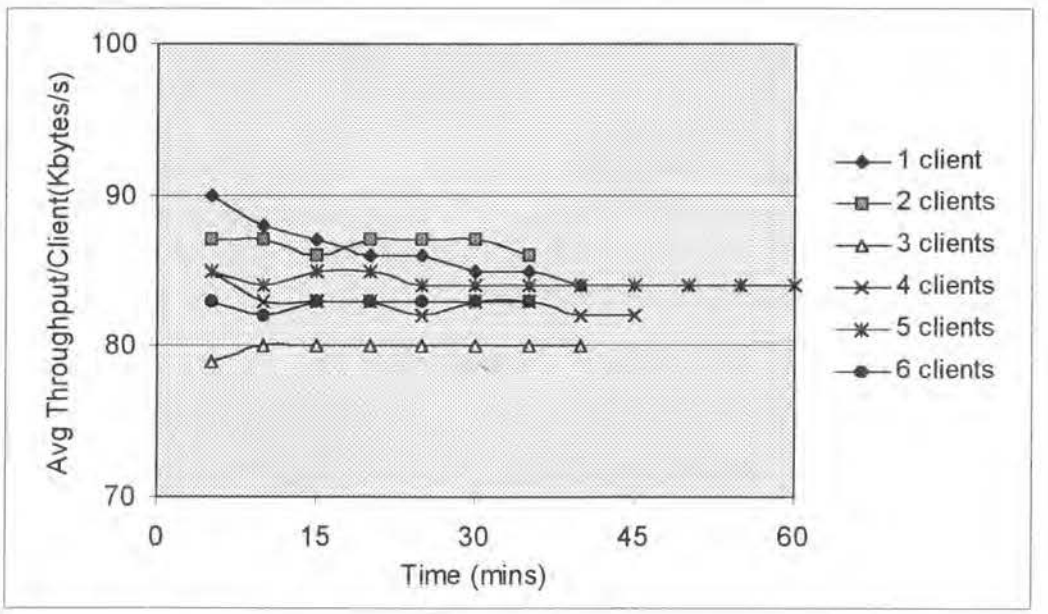

Figure 6.37 - Adding more Clients to a Switch

Effect of loading an Ethernet Switch

As Figure 6.37 shows, no clear pattern emerges as more client/server pairs are added. There is no consistent decrease in throughput per client as the load on the switch increases. The switch shows no strain under the offered load.

Normalising the Throughput

As was done for the Test 8 results, the combined throughput for each combination of client/server pairs was normalised using the Steady State throughput of individual client/server pairs. The Steady State throughputs under switch operation are listed in Figure 6.20 - Baseline Parameter Limits (switched Ethernet).

From Figure 6.20 the average Steady State value for the seven client/server pairs
= (85 + 86 +85 +73 + 88 + 85 + 87)/7 = 84 Kbytes/sec.

The normalised, average throughputs per client for each client/server combination are as as follows –

For CS1/CS3:
87 x (2 x 84)/(85 + 85) = 87 x 0.99 = 86 Kbytes/sec.

For CS1/CS3/CS4:
80 x (3 x 84)/(85 + 85 + 73) = 80 x 1.04 = 83 Kbytes/sec.

For CS1/CS3/CS4/CS5:
82 x (4 x 84)/(85 +85 + 73 + 88) = 82 x 1.02 = 83 Kbytes/sec

For CS1/CS3/CS4/CS5/CS6:
84 x (5 x 84)/(85 +85 + 73 + 88 + 85) = 84 x 1.01 = 85 Kbytes/sec.

For CS1/CS3/CS4/CS5/CS6/CS7:
83 x (6 x 84)/(85 +85 + 73 + 88 + 85 +87) = 83 x 1.00
= 83 Kbytes/sec (i.e. unchanged by normalisation).

The normalised values are tabulated and plotted in Figure 6.28 - Adding more Clients to a Network. Figure 6.28 shows that normalisation compensates for the slowness of CS4 – it lifts the throughput of slow combinations such as CS1/CS3/CS4 and lowers the throughput of fast combinations such as C1/C3. The normalised curve is flatter.

6.14.5 Conclusions from Test 10

The flatness of the normalised curve (see the "Switch normalised" curve in Figure 6.28) illustrates that in this switched Ethernet network, the throughput per client remains essentially constant, independent of the number of clients (as one would expect in the absence of contention and collisions).

The results from Test 10 shed light on the results from Test 4. It can now be seen that the throughput of 144 Kbytes/sec obtained for the five-client combination in Test 4 (refer Table 6.5 - Combinations of Clients in a switched Ethernet) was not an experimental aberration, but rather part of the irregular upward and downward movements that could be expected from a non-normalised set of results.

This means that the Test 4 results, like the Test 10 results, show that we have a lightly loaded switch and clients of varying throughput capabilities. There is no pattern, no general decrease as there was in the shared Ethernets of Tests 2 and 8. The hypothesis that the Test 4 results are demonstrating the effect of a server bottleneck (as was put forward in Section 6.6.5, "Conclusions from Test 4") must also be discounted.

6.15 Test 11: Ethernet switch with one port prioritised

6.15.1 Setup

The six servers S1, S3, S4, S5, S6 and S7 were arranged to simultaneously copy files to Clients C1, C3, C4, C5, C6 and C7 respectively.

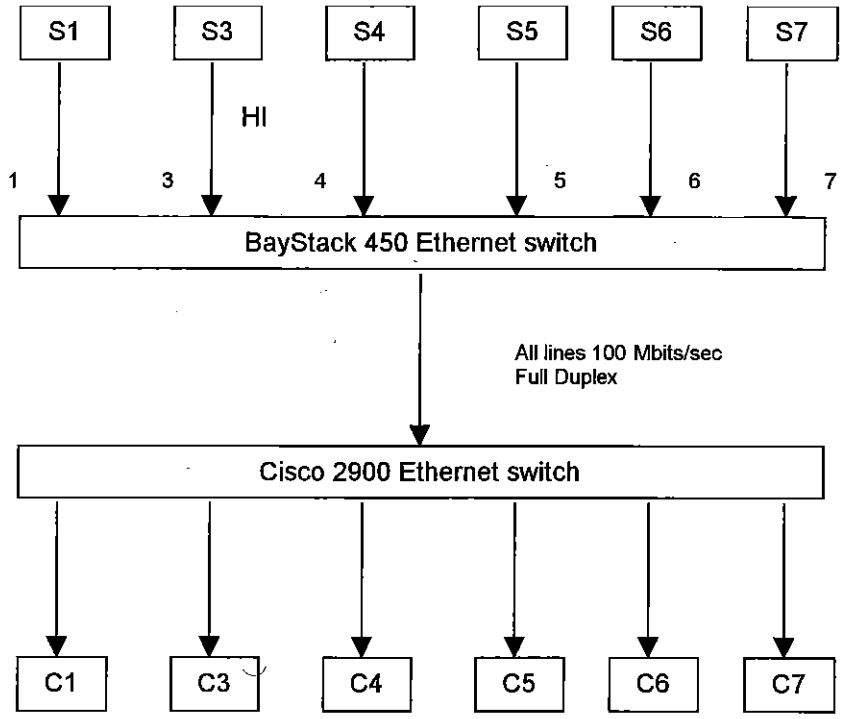

Figure 6.38 - Test Bed for Test 11 (a)

All traffic was channelled through one output port (Port 8) on the BayStack switch, in an attempt to stress the port to a point where prioritisation can come into its own (that is, where its contribution to throughput outweighs the additional overhead).

Port Priority on selected input ports on the BayStack switch was set to "6". "6" was then mapped to "High Traffic Class" (as was done in Test 5). For example, in Test 11 (a), the Port Priority for Port 3 was set to "6", which made the traffic between S3 and C3 "High Class Traffic", as indicated in the diagram below.

The output port (Port 8) on the BayStack switch was not prioritised, nor were any ports on the Cisco switch.[8]

[8] The Cisco 2900 switch does not support prioritisation.

6.15.2 Purpose

To note the effect of prioritising the ports on an 802.1p-capable Ethernet switch. This is a repeat of Test 5 but with more stringent test conditions.

The "Conclusions from Test 5" noted that the load put on the switched network in Test 5 may have been insufficient to demonstrate the effect of prioritising some ports. The revised setup adopted for Test 11 includes two Ethernet switches and additional servers (one server per client) to put maximum pressure on a single port on the switch under test.

6.15.3 Procedure

As in the previous tests, file transfers were initiated first. Once all 6 clients were copying files from their respective clients, the LANalyzers programs were started simultaneously (within seconds of each other). Throughput was again recorded at 5 min intervals.

The test was repeated with a different port, Port 5, set to high priority.

6.15.4 Results

The results from the two tests, Tests 11(a) and 11 (b), are summarised in the tables and graphs below.

Test 11 (a) – Server 3 on Priority Port

Table 6.10 is interpreted as follows: Five minutes after the LANalyzers were started, the throughput on the high-priority server (S3) was 86 Kbytes/sec and the throughputs on the other servers (all low-priority) were 87, 74, 79, 87 and 87 Kbytes/sec. The average low-priority throughput was therefore 83 Kbytes/sec.

Table 6.10 - Server S3 on Priority Port

Time (mins)	Throughput (Kbytes/sec)							
	High Priority Server (S3)		Low Priority Servers					
	Raw	Normalised	S1	S4	S5	S6	S7	Avg
5	86	87	87	74	79	87	87	83
10	86	87	86	73	79	87	84	82
15	86	87	87	73	80	87	85	82
20	86	87	87	73	81	87	85	83

Figure 6.39 compares S3's normalised throughput with the average throughput from the low-priority servers.

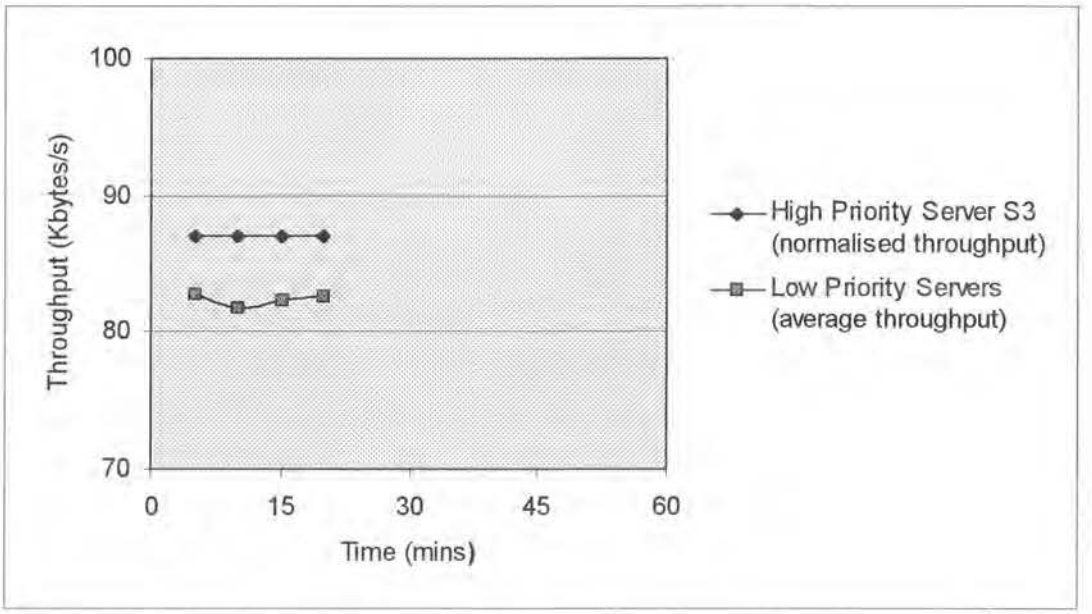

Figure 6.39 - Server 3 on Priority Port

Normalising the Throughput

Because CS3 is a relatively slow client/server pair (refer Figure 6.20 - Baseline Parameter Limits (switched Ethernet)) normalisation will push up its throughput.

Normalised throughput for S3

= 86 Kbytes/sec x S3 baseline throughput / Average baseline throughput

= 86 Kbytes/sec x 85/84

= 87 Kbytes/sec.

Test 11 (b) – Server S5 on Priority Port

Table 6.11 and Figure 6.40 summarise the results from the second test run in which the port connecting S5 was prioritised.

Table 6.11 - Server S5 on Priority Port

Time (mins)	Throughput (Kbytes/sec)							
	High Priority Server (S5)		Low Priority Servers					
	Raw	Normalised	S1	S3	S4	S6	S7	Avg
5	80	76	85	90	74	89	84	84
10	82	78	85	87	74	86	84	83
15	82	78	84	87	73	86	84	83
20	82	78	86	87	73	87	84	83

CS5 is a relatively fast client server pair (refer Figure 6.20). Hence its raw throughput needs to be "normalised downwards" (multiplied by the ratio 84/88) in order to allow for a valid comparison with the low-priority servers.

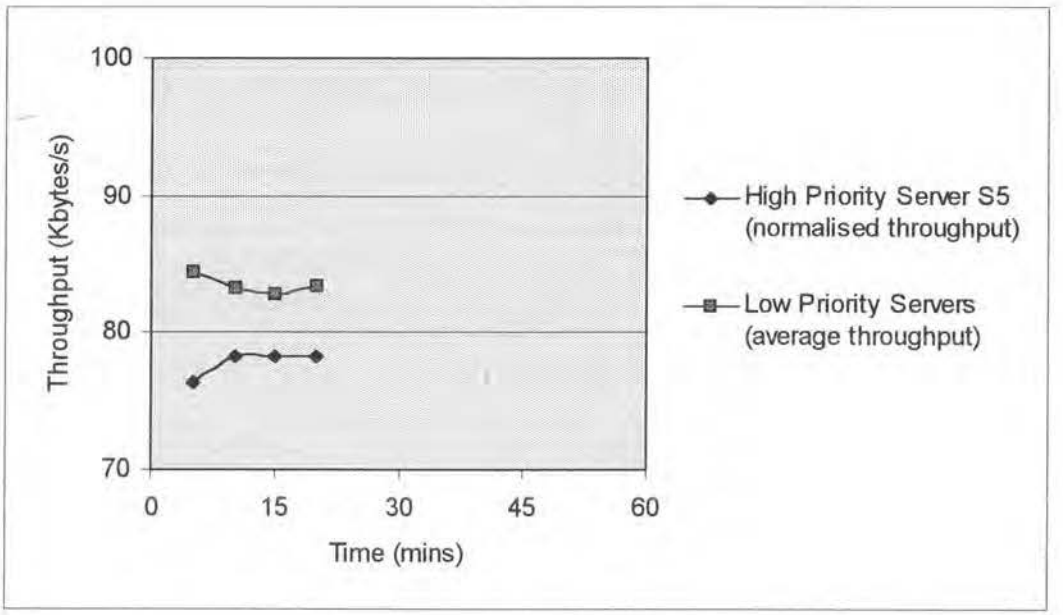

Figure 6.40 - Server S5 on Priority Port

6.15.5 Conclusions from Test 11

Once again, lack of consistent test results limits any conclusions that might be drawn regarding the effectiveness of prioritisation. As was the case in the previous prioritisation tests (Tests 5 and 6), the results from Tests 11(a) and 11(b) are ambiguous: Tests 11(a) shows a marginal increase in throughput for the prioritised port, whereas Test 11(b) reverses the result.

It appears that, even though frames from 5 servers (the low-priority Servers S1, S3, S4, S5 and S6) all had to pass through a single queue (the low-priority queue on Port 8), the switch was able to clear this queue without any significant delays.

In Test 11, the low-priority queue on Port 8 was processing most of the traffic. To investigate what happens when the high-priority queue gets most of the traffic, an additional test, Test 12, was devised.

6.16 Test 12: Ethernet switch with all but one port prioritised

6.16.1 Setup

Similar to the setup in Test 11, the difference being that all input ports, except Port 5, were configured as high-priority ports.

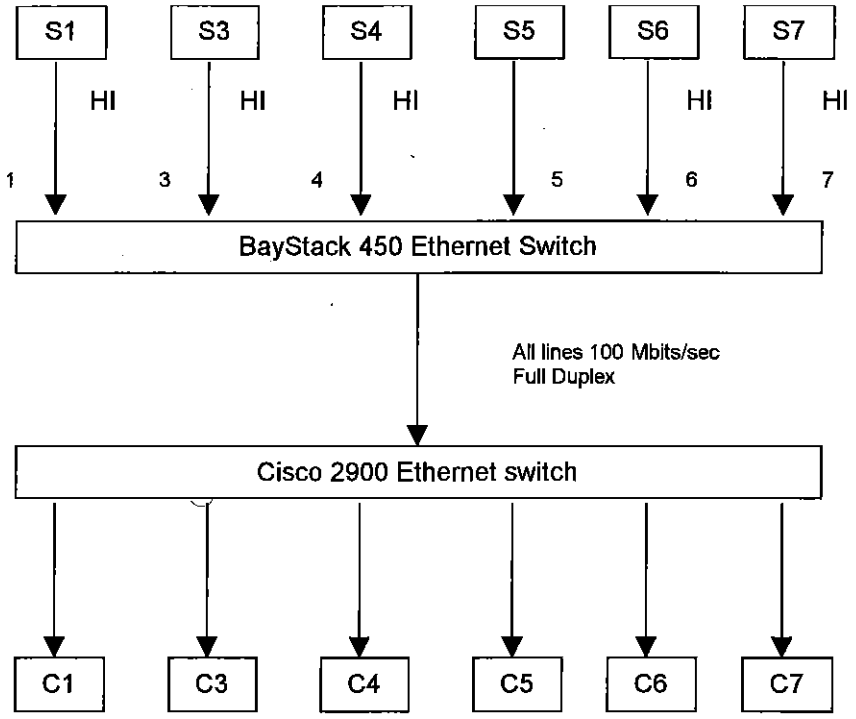

Figure 6.41 - Test Bed for Test 12

6.16.2 Purpose

Test 12 is a follow-up from Test 11. Again, the purpose was to note the effect of prioritisation. However, in Test 12 a relatively large number of ports are prioritised so that greater stress is put on the priority queue on the output port (Port 8).

6.16.3 Procedure

As for Test 11, except that the test was not repeated.

6.16.4 Results

The results for Test 12 are summarised in Table 6.12 and Figure 6.42. The "Raw" throughput from Server 5 was normalised, as it was in Test 11, by multiplying by the ratio 84/88.

Table 6.12 - Servers S1, S3, S4, S6 and S7 on Priority Ports

Time (mins)	Throughput (Kbytes/sec)							
	High Priority Servers						Low Priority Server (S5)	
	S1	S3	S4	S6	S7	Avg	Raw	Normalised
5	87	88	73	86	85	84	85	81
10	86	86	73	87	82	83	85	81
15	87	86	73	85	82	83	83	79
20	85	85	73	85	84	82	83	79
25	85	85	74	85	83	82	83	79
30	85	85	74	85	83	82	83	79
40	85	85	74	85	84	83	83	79

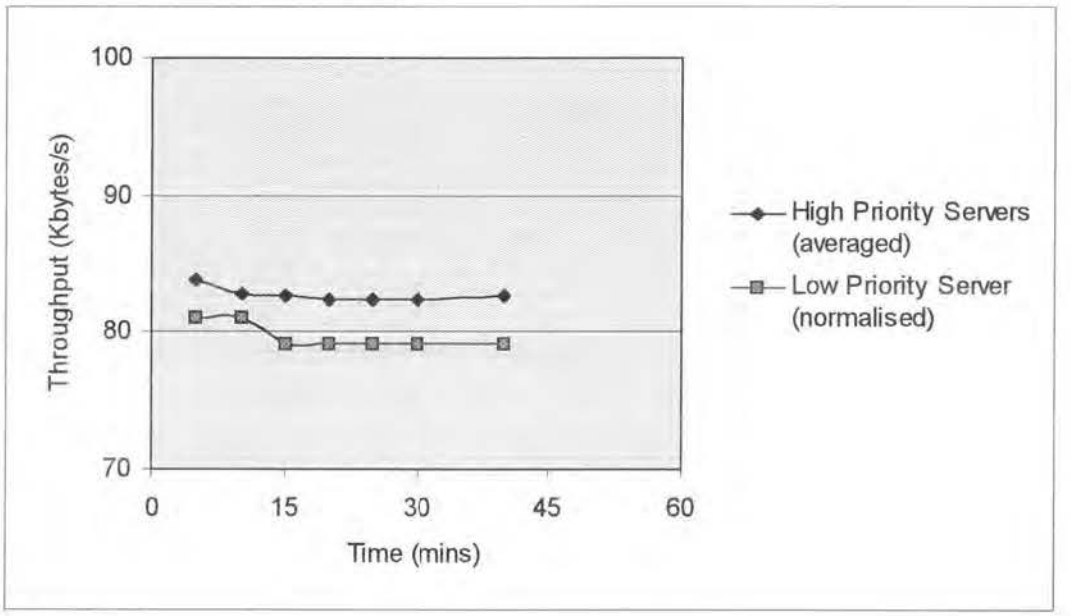

Figure 6.42 - Servers S1, S3, S4, S6 and S7 on Priority Ports

6.16.5 Conclusions from Tests 12

Test 12 shows a marginal increase in throughput for the prioritised ports. However, the difference in throughputs between prioritised and non-prioritised ports were small, allowing the conclusion, that in the tested environment, the effect of prioritisation is not marked.

One possible reason could be that in all prioritisation experiments thus far, the forward streams, only, were prioritised. The reverse flows, consisting of acknowledgement frames were not prioritised. Test 13 addresses this issue.

6.17 Test 13: Ethernet Switch with Input and Output Ports prioritised

6.17.1 Setup

In this test, the traffic flow between high-priority server/client pairs was prioritised in both directions. Both input and the output ports connecting priority clients and servers were set to high-priority.

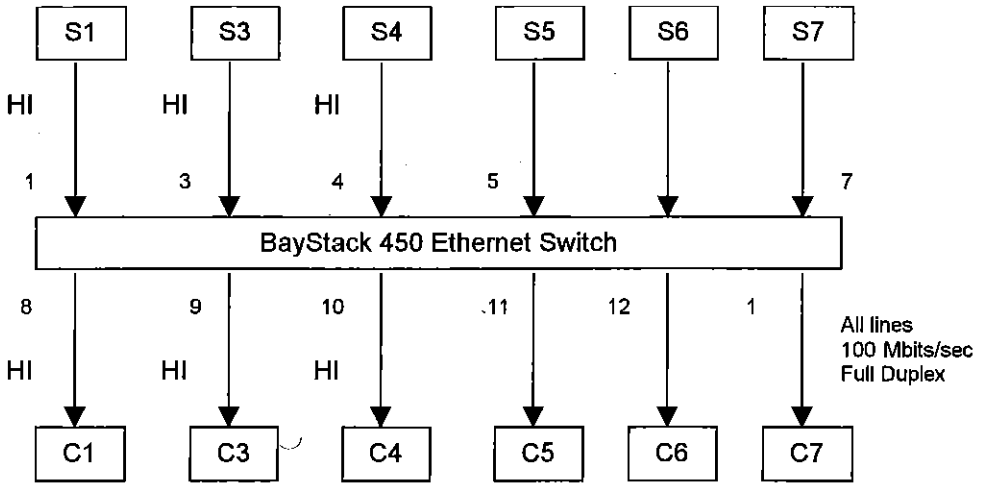

Figure 6.43 - Test Bed for Test 13

The test bed comprises one Ethernet switch and is similar to Test 10. A preferred test bed would have been the one used for Tests 11 and 12, in which all traffic was channelled through one port. In Test 13 this was not possible because it would have required two 802.1p-capable Ethernet switches when only one was available.

6.17.2 Purpose

To investigate the effect of prioritising the reverse flow.

The reverse flow (from receiver to sender) consists of acknowledgment packets. Measurements taken during the previous tests showed it typically is about 1.2 % of the forward stream.

In the preceding prioritisation tests, only the input ports were set to high priority, the assumption being that only the forward stream needs to prioritised.

Test 13 is intended to test the hypothesis that delays in the arrivals of acknowledgment packets could slow down the rate at which the sender (in this case the servers) forwards frames to the receiver (the clients), even when –

1. the forward stream is prioritised with respect to the other forward streams
2. the lines are full duplex.

6.17.3 Procedure

As in previous tests, the file transfers were started first, followed by the activation of LANalyzer on each client. A set of readings was taken every 5 minutes for 35 minutes. Readings were taken as close together as possible and in the same order each time.

6.17.4 Results

The results for Test 13 are summarised in Table 6.13, which is interpreted as follows: Five minutes after the LANalyzers were started, the throughput on the high-priority servers (S1, S3 & S4) was 88, 89 and 72 Kbytes/sec respectively. The average throughput for the high-priority servers was therefore 83 Kbytes/sec.

The throughputs on the low-priority servers (S5, S6 & S7) were 85, 88 and 79 Kbytes/sec. The average low-priority throughput was therefore 84 Kbytes/sec.

Table 6.13 - Servers S1, S3 & S4 on Priority Ports

Time (mins)	Throughput (Kbytes/sec)								
	High Priority Servers					Low Priority Servers			
	S1	S3	S4	Avg	Normalised Av	S5	S6	S7	Avg
5	88	89	72	83	90	85	88	79	84
10	88	86	73	82	90	86	88	80	85
15	87	86	72	82	89	85	88	79	84
20	87	86	72	82	89	86	89	80	85
25	86	86	72	81	88	86	89	80	85
30	86	86	72	81	88	87	89	80	85
35	85	86	71	81	88	87	88	82	86

Normalising the Throughput

Because CS1/CS3/CS4 is a slower combination than CS5/CS6/CS7 (refer Figure 6.20 - Baseline Parameter Limits (switched Ethernet)) normalisation will push up its throughput.

Normalised average throughput for CS1/CS3/CS5

= Avg throughput x CS5/CS6CS7 base throughput / (CS1/CS3/CS4) base throughput

= 83 Kbytes/sec x 87/80

= 90 Kbytes/sec.

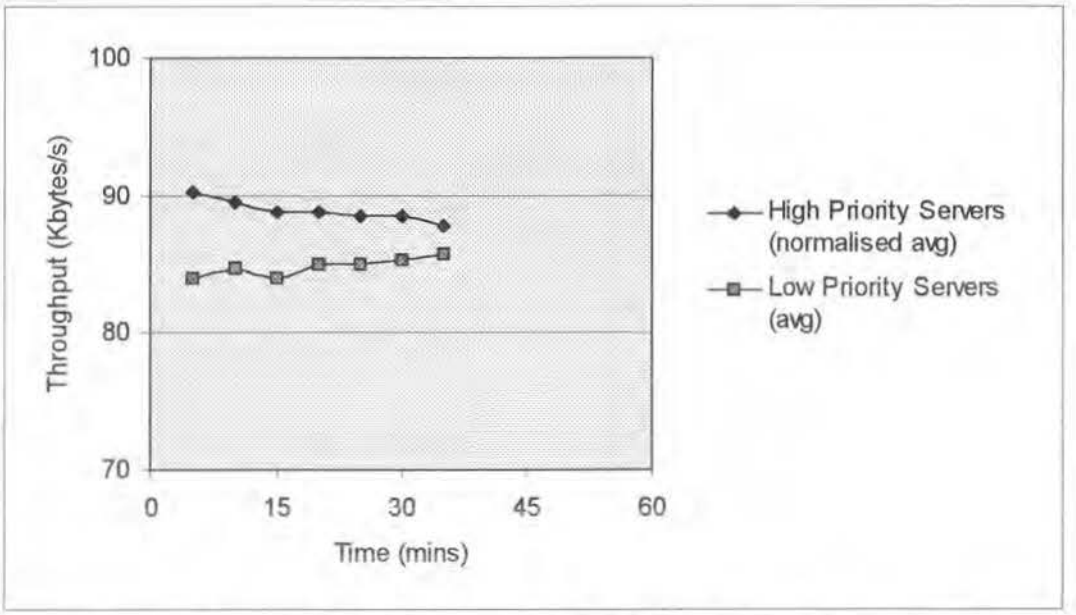

Figure 6.44 - Servers S1, S3 & S4 on Priority Ports

Figure 6.44 compares the average throughput from the low-priority servers with the normalised average throughput from the high-priority servers.

6.17.5 Conclusions from Test 13

The normalised throughput from the prioritised servers exceeds the throughput from the non-prioritised servers by about 5%. Although the difference is small, Test 13 does show higher throughput for prioritised servers, in contrast to the ambiguous results obtained from Tests 11 and 12. The results from Test 13 suggest that reverse flows need to be considered when data streams are prioritised.

6.18 Summary of Test Results

The main results, and the tests in which the results were established, are as follows:

1. Throughput is low compared to the theoretical capabilities of the hard disks and NICs (Test 1).

2. There are significant variations in the throughput capabilities of individual machines, although hardware and software are identical (Tests 1, 3, 7 and 9).

3. In shared Ethernets the throughput was generally higher than in switched Ethernets (Tests 3 and 9).

4. The VLAN configuration of an Ethernet switch has no effect on throughput or prioritisation (Test 6).

5. Throughput from a multi-tasked client is greatly reduced when the file transfer takes place in the background (Test 7).

6. In shared Ethernets the throughput per client decreased as more clients were added to the network (Tests 2 and 8).

7. In switched Ethernets the throughput per client was not affected by the number of clients on the network (Tests 4 and 10).

8. Throughput from high-priority ports is not consistently, and not significantly, higher than the throughput from low-priority ports (Tests 5, 6, 11, 12 and 13).

The above results are specific to the test beds used – the scope for generalising is limited. The experimental factors that need to be kept in mind are fast Ethernet switch, slow PCs, multi-tasking and relatively steady, non-bursting traffic.

6.19 Suggestions for future Experimental Work

1. Concerns were raised in 6.11.5 "Conclusions from Test 7", regarding a possible need to separate traffic generation and network monitoring. In the preceding tests, a client would simultaneously carry out both these functions.

   Since LANalyzer will not run on a NetWare file server, and since an instance of LANalyzer needs to be run on every port of an Ethernet switch (because each port corresponds to a separate LAN segment), the equipment available for this project does not allow the traffic generating and monitoring functions to be separated. This is an area to be considered for further experimental work. For example, a commercial traffic generator could be used to generate the traffic, leaving the client PC with the sole function of monitoring the traffic.

2. All prioritisation tests in this project were done using only one 802.1p-capable networking device, the BayStack 450 Ethernet switch. However, 802.1p prioritisation and the associated 802.1Q tagging is intended to operate over many networks, with tagged packets entering and exiting multiple networking devices and carrying priority information along the way. It may be inappropriate to judge 802.1p prioritisation on the results obtained from a single-device network.

   It is suggested that future experiments should involve at least two Ethernet switches that are fully 802.1p compliant. Network cards and any other intervening networking devices would need to support 802.1p prioritisation.

7.0 CONCLUSIONS

Bandwidth management provides the key to meeting future demands for bandwidth. Current bandwidth management methods are identified and classified in Chapter 5. In Chapter 5, a new system of classifying bandwidth management methods is presented. The classification is based on the OSI layers.

The most effective method of managing bandwidth is the use of Ethernet switches to segment LANs (micro-segmentation). The case study of the Edith Cowan University networks underscores the popularity and effectiveness of switches. The popularity is indicated by the large number of switches installed at ECU in recent years. The effectiveness is demonstrated in Chapter 3 where simple calculations show a 1000-fold increase in the bandwidth of the Student sub-network

A bandwidth management method with the potential to be incorporated in future QoS architectures is 802.1p prioritisation. 802.1p prioritisation is priority queuing implemented in hardware. The queuing of packets is a congestion control method normally performed by routers in Layer 3. 802.1p prioritisation, on the other hand, involves the queuing of frames, which makes it a Layer 2 function.

802.1p prioritisation could form an essential link in end-to-end QoS, providing QoS in LANs to complement other QoS mechanisms, such as RSVP, which are designed to provide QoS across WANs. As such it needs to be integrated into an overall QoS architecture and its effectiveness needs to be evaluated.

The experiments detailed in Chapter 6 aim to test the effectiveness of 802.1p prioritisation within the limits of the experiment. The experimental results are summarised in Section 6.18. The most significant finding is that the throughput from high-priority ports is not consistently, and not significantly, higher than the throughput from low-priority ports.

An important fact brought out by the case study was that 802.1p prioritisation is not used to any significant extent at Edith Cowan University. This is, at least in part, due to the introduction of switches temporarily meeting the demands for more LAN bandwidth and obviating the necessity of introducing bandwidth saving measures. On the other

hand, it may due to busy LAN managers not wanting to tackle an as yet immature technology. If this is the case, then the Case Study supports the experimental results, which suggest that 802.1p prioritisation is not yet ready for implementation on a large scale.

As it is, the driving force for QoS is not LAN bandwidth, but the more expensive WAN bandwidth. While the IETF is still busy developing QoS for WANs, LAN managers and hardware vendors have a breathing space in which to refine 802.1p prioritisation.

8.0 REFERENCES

Adams, S. (1998, Aug 1998). A Measure of Success. *Telecommunications International Edition,* 34-40.

Angin, O., Campbell, A. T., Cheok, L. T., Liao, R., Lim, K. S., & Nahrstedt, K. (1997, 1997). *Report on the 5th IFIP International Workshop on Quality of Service.* Paper presented at the International Workshop on Quality of Service, Center for Telecommunications Research, Columbia University.

Apostolopoulos, G., Guerin, R., Kamat, S., & Tripathi, S. (1998). *Quality od Service Based Routing: A Performance Perspective.* Paper presented at the ACM SIGCOMM'98.

AS/NZS3080. (1996). Telecommunications installations - Integrated telecommunications cabling systems for commercial premises .

Bajaj, S., Breslau, L., & Shenker, S. (1998a). *Is Service Priority Useful in Networks?* Paper presented at the Joint International Conference on Measurement and Modeling of Computer Systems.

Bajaj, S., Breslau, L., & Shenker, S. (1998b). *Uniform versus Priority Dropping for Layered Video.* Paper presented at the Proceedings of the ACM SigComm' 98 Conference on Applications, Technologies, Architectures & Protocols for Computer Communication.

BayNetworks. (1998). *Using the BayStack 450 switch.* Santa Clara, USA: Author.

Berger, A. W., & Whitt, W. (1998). Effective Bandwidths with Priorities. *IEEE/ACM Transactions on Networking, 6*(4), 447-460.

Boggs, D. R., Mogul, J. C., & Kent, C. A. (1988). *Measured Capacity of an Ethernet: Myths and Reality.* Paper presented at the SIGCOMM '88 Symposium on Communications Architecture and Protocols.

Bolot, J, & Turletti, T. (1994). *A rate control mechanism for packet video in the internet.* Paper presented at the Computer Communications (IEEE Infocom), Toronto.

Breslau, L., & Shenker, S. (1998). Best-Effort versus Reservations: A Simple Comparative Analysis. *ACM SigComm.*

Breyer, R., & Riley, S. (1999). *Switched, Fast, and Gigabit Ethernet.* (3rd ed.): Macmillan.

Carpenter, D., & Kandlur, D. D. (1999). Diversifying Internet Delivery. *IEEE Spectrum, November 1999,* pp. 57-61.

Clark, D. (1996). Strategic Directions in Networks and Telecommunications. *ACM Computing Surveys, 28*(4), 679-690.

Coffman, K. G., & Odlyzko, A. M. (1998, October 2). *The size and growth rate of the Internet.* Available: www.firstmonday.dk/.

Determan, S. (1999). *Switched Network Services: Technology overview* (White Paper): Xylan.

Dutta-Roy, A. (2000). The cost of quality in Internet-style networks. *IEEE Spectrum, September 2000*, 57-62.

Engel, A., & Maj, S. P. (1999, 3-5 September 1999). *Towards Quality of Service on the Internet: an Educational Study.* Paper presented at the 3rd Baltic Region Seminar on Engineering Education, Goteborg, Sweden.

Faloutsos, M., Banerjea, A., & Pankaj, R. (1998). *QoSMIC: Quality of Service sensitive Multicast Internet protoCol.* Paper presented at the ACVM SIGCOMM'98.

Farrell, C. (1996). Computer Communications 352 Lecture Notes . Curtin University.

Fitzgerald, J., & Denis, A. (1996). *Business Communications and Networking.* (5th ed.).

Foo, S., Hui, S. C., & Yip, S. W. (1999). Enhancing the quality of low bit-rate real-time Internet communications services. *Internet Research: Electronic Networking Applications and Policy, 9*(3), 212-224.

Guerin, R., Kamat, S., Peris, V., & Rajan, R. (1998). Scalable QoS Provision Through Buffer Management. *ACM SigComm.*

Halsall, F. (1993). *Data Communications, Computer Networks and Open Systems.* (3rd ed.). New York: Addison-Wesley.

Held, G. (1997). *High-Speed Networking with LAN Switches.* New York: John Wiley &Sons.

Huitema, C. (1997). *IPv6 The new Internet Protocol.* (2nd ed.). Upper Saddle River: Prentice Hall.

ISO/IEC Final DIS 15802-3. (1998). *Information Technology - Telecommunications and information exchange between systems - Local and metropolitan networks - Common specifications - Part 3: Media Access Control (MAC) Bridges (Incorporating IEEE P802.1p: Traffic Class Expediting and Dynamic Multicast Filtering).*

Kilkki, K. (1999). *Differentiated Services for the Internet.* Indianapolis: Macmillan Technical Publishing.

Maamria, K. (1998, Aug 98). Quantifying Internet Quality. *Telecommunications International Edition.*

Minoli, D., & Alles, A. (1996). *LAN, ATM and LAN Emulation Technologies*. Boston, London: Artech House.

Minoli, D., & Schmidt, A. (1999). *Internet Architecture*. New York: Wiley.

Odlyzko, A. (1998, September 12, 1998). *The Internet and other networks: Utilization rates and their implications*. Available: www.research.att.com/~amo.

ProjectPlanningTeam. (1999). *ECU Communications Upgrade Project: Position Paper* . Perth: Edith Cowan University.

Reardon, M. (1998). *Traffic Shapers: IP in Cruise Control*, [Internet]. Available: www.data.com.

Rosch, W. L. (1997). *Hardware Bible Premier Edition*. Indianapolis: SAM Publishing.

Rosch, W. L. (1999). *Hardware Bible*. (5th ed.). Indianapolis: Que.

Spurgeon, C. E. (2000). *Ethernet: The Definitive Guide*. Sebastopol: O'Reilly & Associates.

Stoica, I., & Zhang, H. (1999). *Providing Guaranteed Services Without Per Flow Management*. Paper presented at the ACM SIGCOMM'99.

Tanenbaum, A. S. (1996). *Computer Networks*. (3rd ed.). Upper Saddle River, New Jersey: Prentice-Hall.

Tebbutt, D. (1998, November). Quality and equality. *Australian Personal Computer*, 67-68.

Teitelbaum, B. (1999). *Quality of Service for Internet2*. Paper presented at the First Internet2 Joint Applications / Engineering QoS Workshop.

TheTollyGroup. (1998). Test Report No.8286: BayStack 450-24T Switch, Fast Ethernet Switch, Competitive Evaluation, Backplane Capacity : Author.

VanHouweling, D. (1999). *Preface to Proceedings*. Paper presented at the First Internet2 Joint Applications / Engineering QoS Workshop.

VanJacobson. (1998). *Differentiated Services for the Internet*. Paper presented at the First Internet2 Joint Applications / Engineering QoS Workshop.

Wolf, L. C. (1999). Multimedia applications in heterogeneous Internet/ATM environments. *Internet Research, 9*(1), pp.49-57.

Wroclawski, J. (1999). *Evolution of End-to-End QoS: Design Philosophy*. Paper presented at the First Internet2 Joint Applications / Engineering QoS Workshop.