AFIT/GCS/ENG/96J-02

Performance Study of a Brazilian Air Force Data Communications Network

THESIS

Oswaldo Zanelli
1st Lieutenant, Brazilian Air Force

AFIT/GCS/ENG/96J-02

Approved for public release, distribution unlimited.

AFIT/GCS/ENG/96J-02

Performance Study of a Brazilian Air Force Data Communications Network

THESIS

Presented to the Faculty of the Graduate School of Engineering

of the Air Force Institute of Technology

Air University

In Partial Fulfillment of the

Requirements for the Degree of

Master of Science (Computer Systems)

Oswaldo Zanelli, B.S.

1st Lieutenant, Brazilian Air Force

June, 1996

Acknowledgments

It took quite a long time to make this thesis. Without the help and encouragement of several individuals, it could never been done. The first of these is Captain Richard A. Raines, my advisor, who patiently and continuously guided me in this study. I wish to express my sincere thanks to him for his advice and guidance not only during this study, but throughout the whole period I have been in AFIT. I also wish to thank my committee members Major Gerald Gerace and Major Robert F. Mills for their time and cooperation in this study. A word of thanks is also due to Engineer Luis Fernando Barreto who provided me with essential information about the X.25 network.

Finally, I would like to dedicate this study to my lovely daughter, Gabriela, born during this process, who was my main source of inspiration, and add a little more than the customary expression of appreciation to my wife Luciana for her encouragement, understanding, love, and secretarial support throughout this effort.

Table of Contents

List of Figures

List of Tables

Abstract

This thesis consists of a performance analysis of two network models using different protocols, the X.25 and Asynchronous Transfer Mode (ATM). X.25 is currently used in the data communications network infrastructure in the Brazilian Air Force. ATM is a high speed network technology. Simulations of both models were performed to compare the two technologies.

This study examines the traffic load that each network can handle while still providing a good quality of service (QoS). The X.25 network model proved to have a sufficient capacity to handle the projected workload, but it presented a low QoS when a more intense traffic load was introduced. The results obtained from the ATM simulation proved that the ATM model can support much higher traffic loads with no degradation of QoS.

This comparison is important since no study of this kind has been realized before and because this network was designed and implemented for a different type of traffic than the one generated by today's applications. The implementation of a high-speed technology will bring many benefits, such as: increased throughput, higher data rates and multimedia traffic.

Performance Study of a Brazilian Air Force Data Communications Network

1 Introduction

By the end of the last decade, the Brazilian Air Force (BAF) designed a private data communications network to interconnect its information systems. This provided a more secure, inexpensive and reliable way of sharing data, and resources. The technology used was the X.25 network protocol. X.25 is a well established packet-switching service usually used to connect remote terminals to host systems. It provides any-to-any connections for simultaneous users. It has extensive overhead for error checking, turning to be a very reliable way to setup network links with unreliable telephone systems. The X.25 interface supports line speeds up to 64 Kbps although much of the throughput is overhead from error checking. This service has been used extensively throughout the world. Chapter 2 explains the X.25 protocol in a detailed manner.

More than 10 years have passed since the X.25 implementation within the BAF. Technological advances in microprocessor designs have allowed for increased system performance on the order of hundreds of million-instructions-per-second (MIPS) rates. As a result, communication rates across interconnection mediums such as LANs and WANs have increased to support computer system applications. The Asynchronous Transfer Mode (ATM) technology is, among the high-speed network technologies, the one which is attracting the most interest. ATM is a broadband technology used for transmission of voice, video, and data over LANs or WANs. It is a cell relay technology where the data packets have a fixed size, reducing processing overhead. It provides minimal overhead for error control, depending on reliable transmission systems. ATM is designed to operate in very high speeds. ATM is also described in Chapter 2.

1.1 Objective

This thesis presents a performance study on two network models, one using X.25 and the other using ATM. The objective of this study is to analyze the reliability of both models. This analysis will indicate the benefits of implementing a high-speed network technology to the data communication network of a corporation.

1.2 Scope

This study presents general background information on X.25 and the ATM protocols in Chapter 2. The model developed for the X.25 network which represents the Brazilian Air Force Data Communication Network is presented in Chapter 3. In the following chapter, a similar model using the ATM protocol is described. Simulations parameters and results are analyzed in Chapter 5. Finally, conclusions and recommendations for future research are discussed in Chapter 6.

1.3 Approach and Methodology

The approach chosen for this study was to use modeling and simulation results to evaluate the performance of two different network models. The simulation technique was chosen among the techniques for performance evaluation according to its characteristics listed in Table 1.1. An important factor on this decision was the availability of a software tool, used for communications network modeling and simulation.

The first step in this study was to create a model of the current data communication network. The next step was to implement a similar model to the above using the ATM technology. Both models were accomplished using BONeS Designer, a commercial software modeling tool, used for communication network modeling and simulation. Evaluation of these models and simulation results were then analyzed.

Table 1.1 - Criteria for selecting an evaluation technique [Jai91].

Criterion	Analytical Modeling	Simulation	Measurement
1.Stage	Any	Any	Postprototype
2.Time required	Small	Medium	Varies
3.Tools	Analysts	Computer languages	Instrumentation
4.Accuracy	Low	Moderate	Varies
5.Trade-off evaluation	Easy	Moderate	Difficult
6.Cost	Small	Medium	High
7.Saleability	Low	Medium	High

1.4 Summary

In Chapter 1, the objective of this study was presented, the scope was defined, and the approach and methodology used throughout this study was described. Chapter 2 provides the background information necessary to understand both the X.25 and the ATM technologies. Chapters 3 and 4 describe the X.25 and ATM models, respectively, developed for this study. The simulations performed with both models are presented in Chapter 5 where their results are also discussed.

2 *Background*

2.1 Introduction

Communication networks have long proved their importance. From the first telegraph message sent, to today's mobile communications, communication networks gone through a series of evolutions. The X.25 protocol is a major contributor to the evolution of communication network. It was developed to offer compatibility between computers and packet switched networks. X.25 was established in 1976 and is widely used throughout the world until today. Another emerging technology is the asynchronous transfer mode (ATM). Although ATM is still on its developing and experimental phase, it has been chosen as the switching and multiplexing technique for Broadband Integrated Service Digital Network (B-ISDN), which will offer a wide diversity of new communication services that employ high bit rates.

Section 2.2 presents the X.25 protocol architecture and the services offered by it. Section 2.3 describes the ATM reference model and some pending issues of the ATM protocol.

2.2 X.25 Protocol

The rapid advances in computer technology over the past decades have made feasible dynamic allocation communication, were no scheduling of bandwidth is done before a message is received. This type of communication technology is known as packet switching. Packet switching divides the input of information into small segments, called packets, of data which move through the network.

In 1969, with the cost of dynamic allocation switching falling below that of transmission lines, many private and public packet networks around the world were being developed. This cost reduction drove the need for a standard on the host-network interface. In 1976, the International Telegraph and Telephone Consultative Committee (CCITT) adopted Recommendation X.25 as the standard interface between packet network equipment and the user devices operating in the packet mode. These standards were to facilitate the connection of the several types of data terminal equipment (DTE) to the various public packet switched networks being developed by the time [Ryb80]. X.25 standards define the procedures necessary for a packet mode terminal to access the services provided by a packet-switched public data network. A revised version of X.25 was approved in 1980, adding new capabilities to the X.25 interface and to end-to-end service. Two other revisions were made in 1984 and 1988.

The architecture of the X.25 interface is very similar to the Open System Interconnection (OSI) architecture. The interface between the DTE and the data terminating equipment (DCE) consists of three levels of control procedures:

1. The physical level;

2. The frame level; and

3. The packet level.

The physical level protocol is the lowest level of the X.25 architecture. This level specifies the requirements of CCITT Recommendation X.21, which was developed for packet data networks. It defines functional, mechanical, and electrical interface to the modem, communication facility, or its equivalent DCE. Recommendation X.21 bis which is compatible with RS-232C is also specified.

The frame level specifies the use of data link control procedure compatible with the High Data Link Control (HDLC) of the International Standards Organization (ISO). In X.25, these procedures are defined as the balanced link access procedure (LAPB). LAPB controls the interchange of data across the link between the DTE and the DCE.

The packet level protocol defines procedures for establishment and clearing of virtual calls and permanent virtual circuits, describes the packet formats, and delineates the procedures for data transfer, flow control, and error recovery. X.25 packets are divided into two types: data packets and control packets, which can be distinguished by a bit in the packet header. Control packets can be subdivided into six groups: call setup, flow control, supervisory, confirmation, diagnostic, and interrupt. Within each group, packet types have been defined to perform various functions. Packet types are identified in pairs carrying the same identifier.

2.2.1 Network Services Available to X.25 DTEs

The following services may be provided on public data networks:

1. Switched virtual circuits (SVC) ;
2. Permanent virtual circuits; and
3. Datagrams.

Some networks may offer both datagram and virtual circuit services. A switched virtual circuit is a temporary association between two DTEs. It is established when a call request issued by a DTE is accepted by the called DTE. Some networks implement the facility known as fast-select where a setup request is sent together with the first data packet. A permanent virtual circuit is a permanent association between two DTEs. This way it does not requires call setup or call clearing action by the DTE. A datagram is a self contained packet which has enough information to allow it to be routed to its destination address without the need of any earlier exchange. Datagrams offer no guarantee of delivery but they have a high probability of being successfully delivered.

A DTE may establish concurrent virtual circuits with several others DTEs over a single physical access circuit. This is done through statistical multiplexing, which is a technique similar to synchronous time division multiplexing, but instead of allocating time slots on a fixed per call basis it does on a dynamic basis [SaA94]. Each packet contains a logical channel (LC) number which identifies the packet with a switched or permanent virtual circuit, or a datagram.

Remote asynchronous devices such as printers, and terminals do not implement the three levels of the X.25 protocol. In order to provide protocol conversion and packet assembly and disassembly functions, standards were developed for these devices. These devices are connected to a network through a device called packet assembly and disassembly (PAD). CCITT recommendations X.28, X.3, and X.29 define the operation of asynchronous devices to PAD interface, the services offered by the PAD, and the interaction between the PAD and the host system.

2.2.2 X.25 Broadcast Service

Standard X.25 packet switching networks provide point-to-point communications. Broadcast services are desired for many applications and provide benefits such as simultaneous data to multiple destinations, and reduction of transmission bandwidth. Pant [Pan89] describes how X.25 packet switching networks can be upgraded to provide point-to-multipoint (broadcast) service.

The OSI model defines broadcasting as a higher layer function [CC184]. By this definition, a host broadcasting to several destinations over a standard point-to-point network would require several simultaneous sessions. Data packets would need to be repeated on each one of the sessions, consuming host resources and transmission bandwidth.

The broadcast service described by Pant operates at the OSI network layer of a standard packet network. The host indicates a set of destinations to the network and then transmits its packets. The copy and distribution of data is a function of the network. This implementation provides reduction in the usage of transmission facilities, and reduction of processing in higher layers since only one session is involved. Compared to common broadcast systems such as satellites, it provides higher reliability because of built-in error detection and recovery mechanisms.

This implementation has a distributed network design where selected nodes are linked by operator defined internal connections in a manner that ensures no closed paths. The connections defined are logical connections and do not reflect the physical topology. This means that not all nodes need to participate in the broadcast service. This enables the use of this implementation in any existing packet switching network. The broadcast service consist of a hardware and software architecture that allows both broadcast and point-to-point lines.

2.2.3 LAN Interconnection

The de facto solution to LAN interconnection has been to provide interconnectivity through bridges and routers via private lines [BaW91]. This solution has proved to be costly and complicated. Some advantages of using X.25 services are support for switched virtual circuits, a common backbone for connecting multiple technologies, support for dial-up access, and access to services such as electronic mail.

Barret and Wunderlich [BaW91] describe ways to interconnect LANs that carry medium-speed applications (less than 56 Kb/s), that are appropriately matched for X.25 services. According to the authors, what most influences in the type of LAN interconnection is the end-user application. The traffic patterns associated with the end-user application must be fully understood. Most of the traffic generated in LANs are generated by file servers. File servers send entire files to be processed on the client computer, generating large amounts of traffic. This is suitable for local environments where bandwidth is not a problem. Since data is the major resource shared in WANs, database server applications provide a more efficient solution than file servers. Database servers are accessed through commands generated by a client process that submits or request information from the database server. The server performs database manipulations and transmits only what was requested by the client, using bandwidth in a more effective way in WAN environments.

These applications are supported transparently over WANs by using bridges and routers to establish connectivity. Routers are more appropriate when dealing with X.25. Some advantages of using routers are more efficient routing protocols, and easier administration of the internetwork by segmenting each of the connected networks into subnets. X.25 routers establish switched virtual circuits to others X.25 routers by mapping X.25 addresses to network layer addressing information. As an example, consider the Internet protocol (IP). In this case, there will be at least one X.25 address mapped to an IP router. When receiving an IP packet from a LAN, the router will check IP address and see if any SVC is established to the remote router. If not, it sends a call setup packet. If an SVC exists, the router sends the data through an X.25 data packet. After an SVC is established to a remote router, all subsequent data from the local router to this remote router is transmitted through that SVC.

2.3 Asynchronous Transfer Mode

In 1988 the asynchronous transfer mode (ATM) was chosen as the switching and multiplexing technique for Broadband Integrated Service Digital Network (B-ISDN). B-ISDN services require high-speed digital networks using ATM. It is based on optical fiber transmission, new services such as multimedia services, and ATM that is used to transport and route information through the network.

ATM combines advantages of both circuit switching and packet switching. Similar to packet-switching, it uses information on the cell header to identify and define each communication, providing a connection with a rate according to the specific need of the service. As in circuit-switching, ATM sets up a path from the sender to the receiver, but instead of identifying a connection by the slot number (synchronous transfer mode), it carries the identifier (cell header) along with the data in every slot. Contrary to synchronous transfer mode, the cells are statistically multiplexed on the link, allowing the total bandwidth available to be distributed on a demand basis among several applications.

Applications that involve images, such as medical applications [ChH92] that require rendering and transmission of x-ray, consume gigabits of information. On this kinds of applications real time access is needed most of the times in order to allow transmission of images to remote sites, and permit simultaneous view of images on both sites. For this amount of information to be transmitted in a short period of time, a high speed network is needed. Medical applications using ATM network in the gigabit-per-second range are already available [Bru94, Ran95]. One of these applications is VISTAnet. VISTAnet allows physicians to access a supercomputer in order to optimize treatment plans. It also enables medical records to be accessed remotely facilitating remote expert consultations. The network backbone operates at the 2.4 Gb/s rate and the user connections operate at the 622 Mb/s rate. Another application for an ATM network is in the area of distributed network computing. In this scenario, geographically distributed computers can be used to cooperatively solve problems that were previously solved by costly supercomputers. Distance learning service, where remote classroom sessions are available, is another application where high speed networks are needed to transport integrated voice, data, and image traffic [Ran95].

Applications such as the ones described above will generate a heterogeneous mix of traffic. Today's networks do not support this heterogeneous mix of traffic in an efficient way. Contrary to other networking techniques, ATM will provide a single network to all types of traffic. ATM can be used for both Local Area Networks (LANs) and Wide Area Networks (WANs). In the following subsections the ATM cell, and the ATM protocol reference model are explained. Some issues of the ATM technology are also presented.

2.3.1 ATM Cell

ATM uses short and fixed length cells of 53 bytes. Each cell has a five byte header and a 48 byte information field. The header is transmitted first and contains the addressing information. In B-ISDN all types of traffic presented by the users are formatted into ATM cells.

Two parts of the header are of particular interest in determining which cells belong to any particular connection: the virtual path identifier (VPI) and the virtual channel identifier (VCI). The values of these fields are used by the routing protocol to determine the path and channel a cell will traverse. There are two types of connections based on the VPI and VCI fields: virtual path connection (VPC) and virtual channel connection (VCC).

A virtual channel connection is made up of the concatenation of virtual channel links (VCLs). A VCL exists between two switching nodes and is defined by the routing tables at switching points that examine both VCI and VPI fields on the header of an ATM cell. A VPC is defined in a way similar to as a virtual channel connection except that the only field used for switching is the VPI. A virtual path connection contains several VCCs that are switched together as one unit. As Le Boudec explains [Bou92], another way of viewing this is to consider that there are two layers of ATM connections. The top layer being the virtual connection layer and the bottom layer being the virtual path layer. The switching elements of the lower layer (virtual path switches) examine only the VPI part of the header, while the switching elements of the top layer (virtual channel switches) examine both VPI and VCI fields. That makes all virtual channel connections with the same VPI to switch together at a VP switch. The advantages of using VCC and VPC over a single virtual path are such that two hosts can multiplex many application streams using the VCI fields on the header of the cells to identify each stream, and several virtual paths can share the physical transmission path with each other. This makes possible a single user-network interface to carry several VP connections.

2.3.2 ATM Protocol Reference Model

The ATM protocol reference model is based on standards developed by the International Telecommunications Union (ITU). It is divided into three layers: the Physical Layer, the ATM layer and the Adaptation Layer (AAL). These three layers are now described.

The Physical Layer defines a transport method for ATM cells between two ATM entities. It can transfer ATM cells at the user-network interface in two ways. One is an externally framed synchronous transmission structure. Using this method, cells are written on the byte stream provided by an underlying transmission system, that can be the North American transmission format Synchronous Optical Network (SONET), or its European derivation, the Synchronous Digital Hierarchy (SDH). The second transmission node is a cell based asynchronous transmission structure, were the frames used by the transmission structure match exactly the ATM cells.

SONET specifies how information data is framed and transported synchronously through optical fiber transmission links without requiring all the nodes and links to have the same synchronized clock for data transmission. The basic data rate supported by SONET is the 51.84 Mb/s synchronous transport signal-level 1 (STS-1) frame. Multiples of this bit rate constitute higher bit rate streams.

The STS-1 frame is composed of a rectangular structure, 9 rows by 90 columns, totaling 810 bytes (Figure 2.1). This format is transmitted row by

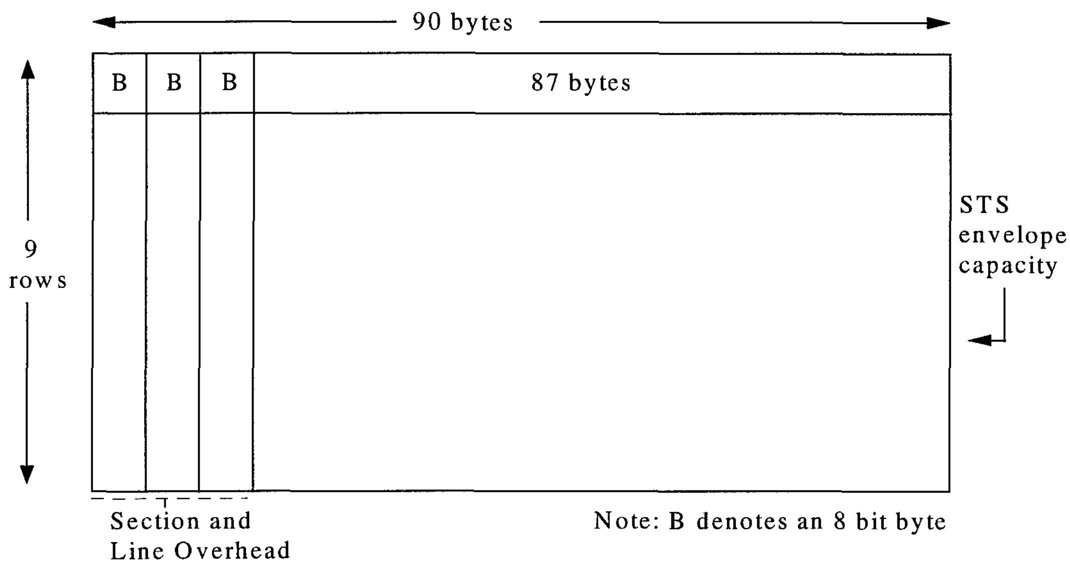

Figure 2.1 - STS-1 frame.

row, from left to right. The STS-1 frame has a frequency of 125 microseconds. It is divided in two areas, the transport overhead (TOH) and the synchronous payload envelope (SPE). The first three columns (27 bytes) of an STS-1 frame are assigned to the TOH, which is subdivided into section overhead (9 bytes) and line overhead (18 bytes). The section and line overhead are responsible for error monitoring, system maintenance functions, synchronization, and identification of payload type. The synchronous payload envelope is the remaining 9 rows by 87 columns. From these 783 bytes, nine bytes in the first column are the path overhead (POH), that support the transport of the payload from the point it is assembled to the point it is disassembled. The remaining 774 bytes form the actual payload. An important point is that SONET can transport both ATM and STM cells in its payload. That is, the physical layer is independent of the payload type. For B-ISDN, the contents of the payload will always be 53 bytes cells. Mapping of ATM cells into the payload is done by aligning by row, the byte structure of every cell with the byte structure of the SONET payload. Since an integer number of ATM cells will not fit in the payload, a cell may cross the payload boundary [ATM93].

Higher rate SONET signals are formed by synchronously multiplexing a number (N) of STS-1 frames. The higher rate signals are exactly N times the basic rate. A STS-12 is 12 times 51.840 Mb/s, or 622.080 Mb/s. The size of a STS-N signal is N x 810 bytes while the transport overhead is N x 27 bytes. The POH is not part of the transport overhead [Dav94]. Only one path overhead is shown, the remainder is allocated to the payload. Thus in this case, the payload is (N x 783 - 9) bytes.

When a STS-N signal is converted into an optical signal, it is called an optical carrier level-N. During conversion, no additional bandwidth is added to the signal, so a STS-N and its corresponding OC-N have the same rate. The lowest optical signal used by SONET is the OC-1. Rates up to OC-255 or 13.2192 Gb/s are possible [Dav94], although N is currently defined as 1, 3, 9, 12, 18, 24, 36, and 48. Table 2.1 shows the relationship of line rates, optical carrier signal, and the corresponding STS levels.

Table 2.1 - SONET rates and bandwidths.

Optical Carrier	STS levels	Line rates (Mb/s)
OC-1	STS-1	51.840
OC-3	STS-3	155.520
OC-9	STS-9	466.560
OC-12	STS-12	622.080
OC-18	STS-18	933.120
OC-24	STS-24	1244.160
OC-36	STS-36	1866.240
OC-48	STS-48	2488.320
OC-96	STS-96	4976.640
OC-192	STS-192	9953.280

SONET framing supports signals with rates lower than 51.84 Mb/s. The STS-1 payload can be subdivided into smaller payload structures called Virtual Tributaries (VT) that transport signal at less than DS3. There are VT mappings for DS1, DS1C, DS2, and DS3 [Boe88]. DS3s are carried directly in the payload, while DS1s, DS1Cs, and DS2s traffic are carried in a STS-1 where the payload has been optimized with Virtual Tributaries. Ethernet and 16 Mb/s Token-ring local area networks will eventually have mappings defined.

The cell is the basic element of the ATM layer. Its functionality is defined by the fields present in the ATM cell header. The cell header at the B-ISDN UNI is different from the B-ISDN NNI header in the use of bits 5 to 8 of the first octet. At the B-ISDN NNI, these bits are part of the VPI, while in the UNI, they make up the Generic Flow Control (GFC) field. Figure 2.2 illustrates the cell header for both UNI and NNI. The GFC field is used by the UNI to limit the amount of traffic entering the network during congestion periods. The VCI/VPI fields are used for channel identification and multiplexing. The Payload Type (PT) field indicates whether the cell contains user or network control information. The Cell Loss Priority (CLP) field is used to determine whether a cell may be discarded during congestion periods. It is assigned a higher or lower priority value. The Header Error Control (HEC) uses an 8 bit error code to correct errors on the header, caused during transmission.

The function of the adaptation layer is to provide a link between the services required by the higher layers and the ATM cells used by the ATM layer. The AAL must distinguish between all types of traffic, because of their different transmission requirements. Four classes of ATM services are defined according to time relation between the source and destination, bit rate, and connection mode (Table 2.2).

<table>
<tr><td colspan="2">Generic Flow Control</td><td colspan="3">Virtual Path Identifier</td></tr>
<tr><td colspan="2">Virtual Path Identifier</td><td colspan="3">Virtual Channel Identifier</td></tr>
<tr><td colspan="5">Virtual Channel Identifier</td></tr>
<tr><td colspan="2">Virtual Channel Identifier</td><td>Payload Type</td><td>Res</td><td>Cell Loss Priority</td></tr>
<tr><td colspan="5">Header Error Control</td></tr>
</table>

Figure 2.2 - ATM header format at the Use Network Interface (UNI).

Table 2.2 - Classes of ATM service.

Feature	Class 1	Class 2	Class 3	Class 4
Time relation	required	required	not required	not required
Bit rate	constant	variable	variable	variable
Connection mode	connect-oriented	connect-oriented	connect-oriented	connectionless

Initially there were four types of AAL protocols to support the four service classes. They were called AAL type 1, 2, 3, and 4. Since the difference between types 3 and 4 were minor, they were merged into a single type called AAL type 3/4. After this association, a simpler AAL for class 3 was defined - AAL5. This new adaptation layer for connection-oriented service was more efficient in transmission and processing because unnecessary functions from AAL3 were eliminated.

The AAL is subdivided into sublayers: The segmentation and reassembly sublayer (SAR) processes data units so that they can fit into ATM cells and reconstructs data units from ATM cells. The convergence sublayer (CS) performs functions such as multiplexing, cell loss detection, and flow control.

2.3.3 Flow Control for Available Bit Rate Service

ATM is a networking protocol developed to support applications with distinct quality of service (QoS) parameters. In order to support the large spectrum of applications expected in B-ISDN, a family of service classes was defined in terms of QoS provided to the users. The ATM Forum, has already published definitions for traffic management for services classes with fixed traffic profile [ATM93].

Constant Bit Rate (CBR), and Variable Bit Rate (VBR), are two service classes which have performance parameters already specified. CBR and VBR can be classified as guaranteed traffic, where a fixed traffic profile can be obtained, and an explicit guarantee of service is given by the network. This guarantee can be viewed as a contract between the traffic source and the network. Before connection setup, the traffic source describes its traffic characteristics and request a certain QoS to the network. During transmission the network ensures that the traffic conforms to the one described before setup. Examples of traffic that may require guaranteed service are real-time traffic such as circuit emulation, voice, and video.

Available Bit Rate (ABR) is another service category. ABR is defined for services where the traffic characteristics are unknown, or incapable of being predicted. This way there is no explicit contract between the network and the user specifying the traffic profile and the QoS required. The network divides the available bandwidth between the users. ABR service is potentially useful for several applications, but it is particularly useful for data traffic, where no firm guarantee of bandwidth is usually required. In the case of data traffic, where each packet of data is segmented into ATM cells, cell loss may lead to retransmission of the entire packet by a higher protocol layer [BoF95]. After Floyd and Romanov demonstrated in a study that cell loss under congestion could lead to the collapse of throughput for packet data applications [FlR94], the control mechanisms for ABR traffic were limited to feedback mechanisms. The two main feedback mechanisms proposed to the ATM Forum were credit-based flow control, and rate-based flow control. The latter was chosen by the ATM Forum.

Rate-based schemes use feedback information from the network to control the rate at which the user should transmit. Feedback from the network to the end system gives users the necessary information to adjust their transmission rate according to available bandwidth, so that congestion is controlled or avoided. Bonomi and Fendick present the rate-based framework that is being developed for the support of ABR services, and its basic operation [BoF95]. The rate-based framework supports both end-to-end flow control and segmentation of the control loop (virtual sources and destinations), that is done in the intermediate switching elements. This can improve performance by reducing the length of the control loop. The framework also defines mechanisms and control information formats to allow switches implementing different types of feedback control to coexist in the same loop.

The basic operation of congestion control schemes encompassed by the rate-based framework is as follows: the source creates a connection with a call setup request, where parameters such as peak cell rate (PCR), minimum cell rate (MCR), and rate decrease factor (RDF) are specified either by the source or by the network. After permission is granted, the source begins transmission. The ABR source is allowed a rate to transmit, called allowed cell rate (ACR), which is initially set to the initial cell rate (ICR). A resource management (RM) cell is sent before data cells in the beginning of transmission, and as the first cell after every idle period. The RM cell is a non-user cell format, not yet standardized, and is reserved for control functions in ATM networks [New94a]. The source rate is controlled by the return of these RM cells. The source places the rate at which it may transmit (ACR), and the rate at which it wishes to transmit (usually peak cell rate) in the explicit rate field in the RM cell. The RM cell is transmitted through the network. Switches in the RM path use this information to allocate bandwidth among ABR connections, and may also reduce the explicit rate. When the destination receives the RM cell, it changes the direction of the RM cell to the source, and adjust parameters when the explicit rate is not supported. During the backward travel, the switches may examine the RM cell and do any necessary adjustment regarding bandwidth. When the RM cell arrives at the source, the source adjusts its rate (ACR) based on the information carried by the RM cell.

2.3.4 Connectionless Service on ATM

The existing customer networks send most of its data traffic over local area networks, such as Ethernets and Token Rings. These LANs offer connectionless service, where no call setup procedure is required. In order to interconnect these LANs with public ATM networks, connectionless service must be provided on top of connection-oriented technology. Currently, ITU recommends two general approaches that offer connectionless service in an ATM network: the indirect approach and the direct approach.

In the direct approach presented by [New94b], a connectionless server function (CLSF) is realized using connectionless servers. Basically, a connectionless server (CLS) is a packet switch attached to an ATM switch. All connectionless data that is sent by internetworking units (e.g. bridges and routers) to the ATM network are directed by the ATM switch to the CLS. The server is responsible for routing the message. The connectionless servers are connected by virtual paths through the ATM switches that each CLS is attached (Figure 2.3).

One of the advantages of this approach is that each connectionless LAN needs only one connection, called Internetworking Unit (INU), at the edge of the network to send its data. This way INUs are not responsible for routing decisions, the ATM network is. Another advantage is that in comparison with the indirect approach, the direct approach reduces the number of connections required to realize connectionless service; since only the CLS are connected. Some disadvantages are: the performance of the ATM switch is not fully utilized; and the connectionless servers will have to perform at full speed to keep up with the ATM switches, which can become a performance problem.

The indirect approach is another form of implementing a connectionless service on top of a connection-oriented network. It is implemented by interconnecting ATM internetworking units through the virtual connections established between each INU. The difference from the direct approach is that connectionless servers are removed and the INUs form the interface between the LAN and the ATM wide area network (Figure 2.4).

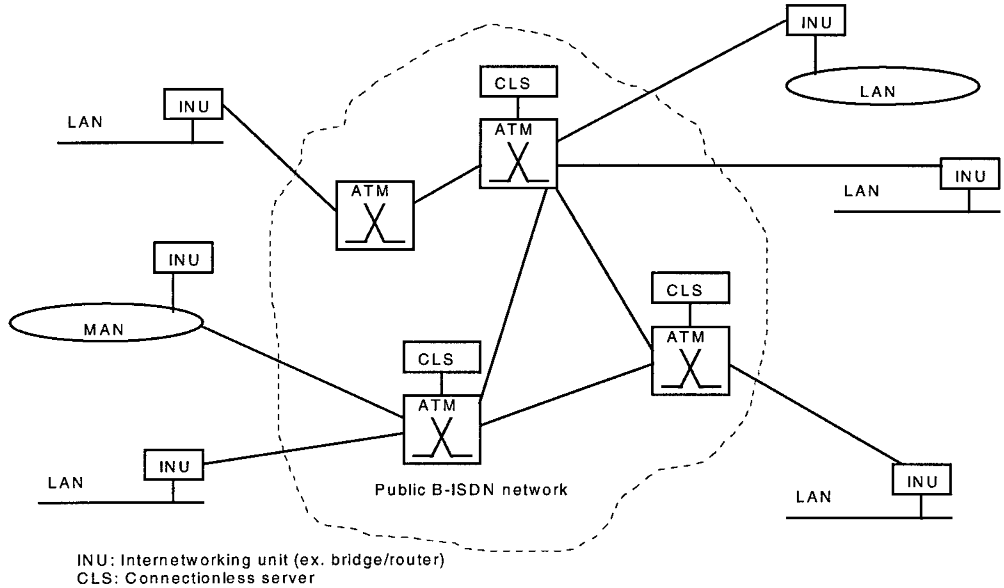

Figure 2.3 - The direct approach to connectionless service [New94b].

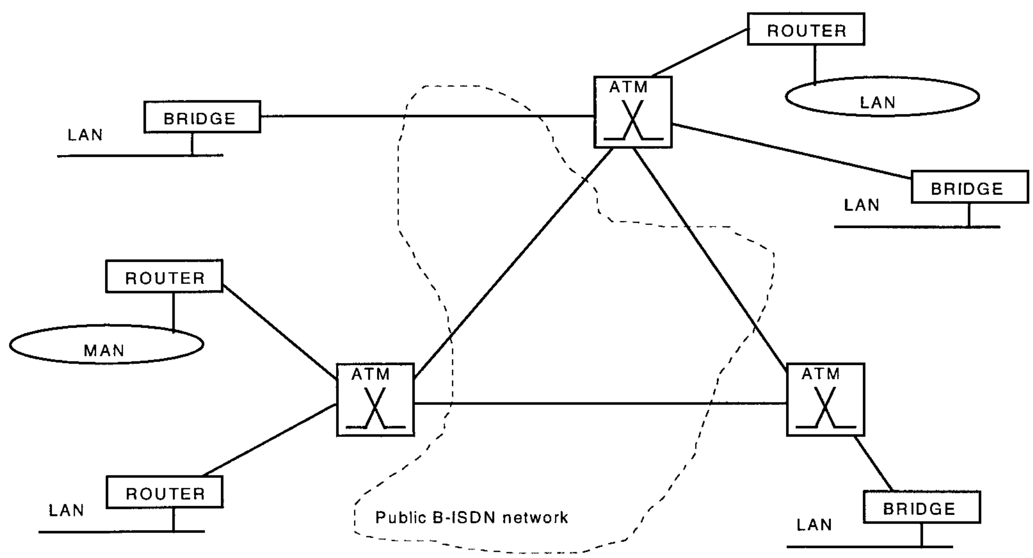

Figure 2.4 - The indirect approach to connectionless service [New94b].

The indirect approach advantages are that cost and management are reduced, due to less networking devices. A single switching mechanism and full ATM switching capacity is used. One disadvantage is that it requires guaranteed allocation of network resources such as bandwidth.

2.3.5 LAN Emulation

Not long ago, the bandwidth offered by LANs such as Ethernet and Token-ring were seen as an infinite amount of bandwidth. With today's high technology workstations and traffic created by multimedia applications, technologies that offered what once seemed an infinite amount of bandwidth, are no more as efficient as they used to be. According to [New94b], for ATM technology to be successful in LANs it must offer the same service for data traffic offered by present LANs. Almost all installed LANs conform with the IEEE 802 family of protocols, and uses a medium access sublayer (MAC) address.

LAN emulation is the process of supporting an IEEE 802 connectionless packet transfer service at the MAC sublayer on top of ATM. Although the network layer protocols can be implemented on top of ATM, the major idea in implementing a specific ATM MAC sublayer is that compatibility with the large existing data communications protocols, applications, and equipments must exist. In the IEEE 802 standards, the MAC is a sublayer of the data link layer and is specific to a particular LAN (ex. Token-ring). This implementation would include ATM to the IEEE 802 family of protocols together with Ethernet, Token-ring, Token-bus, and Distributed Queue Dual Bus (DQDB). The solution is to design an ATM MAC sublayer above the adaptation layer that emulates the services offered by an IEEE 802 LAN at the MAC sublayer [New94b]. This approach enables ATM end-systems to interconnect with other networks through bridges and routers, the same way Ethernet and Token-ring do today.

2.4 Summary

This chapter presented an overview of X.25 and ATM technology. The information presented above does not cover the two subjects completely, but goes through important aspects pertinent to the study that will be presented in the following chapters.

The architecture of the X.25 interface was reviewed, followed by a presentation of the services available to an X.25 DTE. A description on how to broadcast service on an X.25 network works was also presented. The solution to LANs interconnection described using X.25 services, has proved to be cheaper than the commonly used interconnection through bridges and routers using private lines.

Section 2.3 described the ATM protocol reference model. Bonomi and Fendick present a rate-based framework for flow control and also demonstrate the basic operation of congestion control schemes encompassed by the framework. Two approaches to offer connectionless services in an ATM network were described.

3 X.25 Model

3.1 Introduction

This chapter presents the methodology used in the modeling and simulation of the X.25 protocol network. Section 3.2 discusses the portions of the protocol modeled. An overview of the model, explaining the most important modules and functions is presented in Section 3.3. Verification and validation of this model are described in Section 3.4. Section 3.5 summarizes this chapter.

3.2 Modeled Aspects of the X.25 Protocol

This model is based on the model described in [Cad94b]. It emphasizes the packet layer call processing phases and the DTE/DCE interface. Layer 2, which provides data link services to layer 3 was modeled as a First In First Out (FIFO) server instead of the specified LAP or LAPB. This was done to improve the simulation performance by reducing execution times. This model includes the following aspects of the X.25 protocol:

- Virtual Call Service - Different packet types flowing through the DTE and DCE model the three phases of a virtual call (call setup, data transfer, and call clearing). The packet sequences specified for the X.25 protocol are modeled in detail.
- Logical Channels between the DTE and DCE - As specified by the standard. The range of channel numbers are implemented as global parameters of the model, as

described in Section 3.3.2.2. The logical channels may be used for both incoming and outgoing calls.

- Multiplexing - The model uses a number of logical channels for simultaneous virtual calls. By this, it allows multiplexing in the same way the X.25 protocol does. Parameters set by the user specify the number of logical channels between a DTE and a DCE, and the number of virtual circuits in the network. The combination of these parameters allows the number of simultaneous multiplexed channels to be varied. Up to 4095 virtual calls may be in progress at a given DTE at the same time.
- Flow Control - The model does not allow rejection of packets, only positive acknowledgments are produced. Flow control was implemented using window size and the Delivery confirmation bit (D bit). The window size for data packets can be any positive integer. The D bit indicates who acknowledges a received packet. If set to 1, the receiving DTE generates the acknowledgement, confirming the delivery. If set to 0, the acknowledgement is generated at the local DCE. The D bit must be set for each X.25 station.

In order to improve the model's efficiency, the model was simplified by not modeling some parts of the protocol. Error recovery, and sequence numbers for flow control were not modeled since the network connecting the X.25 stations is designed to deliver packets to the correct destination, in the order they are received (all packets associated to a VC follow the same path through the network), and without loss. Only one way data transfer is modeled.

3.3 Model Overview

3.3.1 The Data Structures in the X.25 Model

A total of ten data structures (DS's) were used on this model. Data structure inheritance was applied in the X.25 packet DS. Their names and relation to other data structures appear on Table 3.1. Each user defined data structure is also described below.

Table 3.1 - Data Structures in the X.25 Model.

DATA STRUCTURE	RELATION
X.25 External DS	COMPOSITE
X.25 Network DS	COMPOSITE
X.25 Packet	COMPOSITE
Call Accepted & Call Connected	X.25 Packet
Call Request & Incoming Call	X.25 Packet
Clear Request & Clear Indication	X.25 Packet
Data Request & Data Indication	X.25 Packet
Receiver Ready	X.25 Packet
MAC-Data.req	COMPOSITE
Router DS	COMPOSITE

3.3.1.1 X.25 External DS

The X.25 External data structure is used to model requests issued by users of a network to their local DTE. After the completion of the request, copies of this DS are also an output at the destination DTE, representing a serviced request, and at the sending node, representing an acknowledgement. This acknowledge can be either a successful acknowledge or failure of the file transfer. This will be indicated by the value of the Request Type field. The information conveyed by a request is reflected on the fields of the data structure. These are described as follows:

a) Time Created - This field holds the current simulation time when the actual data structure is created by the requesting user.

b) Source DTE - This field is set to the address of the requesting station.

c) Destination DTE - This field is set to the address of the remote station to which the request should be carried out.

d) Request Type - A set type indicating the meaning of the data structure. A "0" indicates the DTE that the data structure is a request. A request that arrived at the destination DTE is specified by a "1" (Indication). A "2" (Acknowledgement) indicates the status of a request sent to a DTE. This status is specified by the Status of Request field explained below.

e) Status of Request - This field only has meaning when this DS is coming out of the sender DTE as an Acknowledgement. This field is set by the local DTE according to the status of a request. The status codes and their meanings are shown on Table 3.2. Rejections by the local DTE occur when it cannot allocate an logical channel. Rejections by the DCE occur either when the DCE cannot allocate a virtual circuit or the remote DCE cannot allocate a logical channel to the remote DTE.

f) Message Length - The length in bytes of the file to be transferred. The sending DTE uses this number to determine how many packets are sent to transmit this file.

g) Data to Deliver - This field is used for encapsulation of another DS. This encapsulated DS is passed out to the destination DTE when all the request has been transferred.

Table 3.2 - Status Codes.

CODE	MEANING
0	Request has completed successfully.
1	Request has been rejected by the local DTE.
2	Request has been rejected by the DCE.

3.3.1.2 X.25 Packet DS

Communication between a DTE and a DCE of a same X.25 station is made through several X.25 Packets. The X.25 Packet is the "parent" of six others data structures that inherit all of its fields. The "children" data structures (described below) flow between a DTE and a DCE causing the different phases of a virtual call to progress. The fields of the X.25 Packet data structure are described below:

a) Time Created - This field is set to the current simulation time whenever a newly created X.25 data structure is transmitted.

b) Type - This field is set to a type code whenever a "children" packet from the X.25 Packet is created. The type codes are shown in Table 3.3.

c) LC # - This field identifies to which logical channel number a packet belongs. Packets from a given call have the same logical channel number.

Table 3.3 - X.25 Packet type codes.

CODE	PACKET TYPE
1	Call Request & Incoming Call
2	Call Accepted & Call Connected
3	Clear Request & Clear Indication
4	DTE & DCE Clear Confirmation
5	Data Request & Data Confirmation
6	Receiver Ready

3.3.1.2.1 Call Request & Incoming Call DS

This data structure is issued by the sender's DTE to request that a virtual circuit be initiated. The fields Sender and Receiver are set to the source and destination addresses of the corresponding call. The field DS to Transfer encapsulates a copy of the file transfer request that is sent out of the destination DTE when the call is completed.

3.3.1.2.2 Call Accepted & Call Connected DS

This is a control packet that contains no fields. When received by the local DTE, it causes the data transfer phase of a call to initiate.

3.3.1.2.3 Clear Request & Clear Indication

This packet contains a single field indicating why the call is being cleared by either DTE. The Cause field contains a code representing why the call is being cleared.

3.3.1.2.4 Data Request & Data Indication

In a real system, a data packet contains a segment of the message being transmitted. On this model, the message segment is replaced by an integer stored in the field Length. This integer represents the length, in bytes, of such segment. The D bit field is used to determine where the acknowledgement received for each data packet is generated.

3.3.1.2.5 DTE & DCE Clear Confirmation DS

This data structure is used as an indication to the local DTE that a given call has successfully terminated. Similarly to the Call Accepted & Call Connected DS, there are no information fields in this DS.

3.3.1.2.6 Receiver Ready

This packet acknowledges the receipt of data packets. The receipt of a Receiver Ready packet represents an acknowledgment of the longest outstanding data packet.

3.3.1.3 X.25 Network DS

This data structure flows through the network, between two DCEs. It encapsulates X.25 packets that are transferred from one DCE to another. The information contained in this data structure enables the network to deliver it to the correct DCE without looking into the encapsulated X.25 packet. The X.25 Network DS fields are described below:

a) Call Setup Flag - This field may have two values. Its value is "1" if the encapsulated field is a call request, and "0" otherwise. This information is used to keep track of routing information.

b) Virtual Circuit - All the packets in a virtual call will have the same virtual circuit number. The network uses this field and the Call Setup Flag to setup routing tables used for routing packets through the network.

c) Calling DTE - The identification number of the initiator of the call.

d) Called DTE - The identification number of the DTE which received the call.

e) Packet Source - The identification number of the DTE which generated this packet.

f) Packet Destination - The identification number of the DTE to which this packet should be delivered.

g) Packet Length - The length in bytes of the packet. This field is used to calculate any transmission delay.

h) X.25 Packet - The encapsulated X.25 Packet that will be extracted at the destination DCE.

3.3.1.4 MAC-Data.req DS

The Carrier Sense Multiple Access with Collision Detection (CSMA/CD) is a protocol commonly used for local area networks. The MAC-Data.req DS implements the request made by the Link Layer to the MAC Layer that a frame (packet) be sent. This request is encapsulated in the Router DS information field. Since this model is not concerned with the CSMS/CD protocol, its functions are not explained in detail. It is presented here so the reader can better understand the Ethernet model. The fields of this DS are: source, destination, data, and length.

3.3.1.5 Router DS

This data structure was used to maintain compatibility between the traffic generator used in this model with the one used in the ATM model described in the following Chapter. It is composed of an Information field, that holds the information to be delivered, a Source and Destination field that are the addresses of LANs, and an Information Length field that holds the size of the information.

3.3.2 Model Description

Two similar models were developed for the X.25 Network. The first models the actual X.25 Network. Its traffic source generates a traffic similar to the one generated by the current network. The second model has a traffic source that models the traffic generated by an Ethernet LAN. Besides this, the second model uses an Enhanced X.25 Station module that simulates the reuse of virtual circuits. The models are scaled down to ten X.25 stations. Both are described below using a top-down approach.

3.3.2.1 Current X.25 Network

The system described below is a reduced scale model of the current data communications network used in the Brazilian Air Force. It attempts to capture the performance and topology of the real system. The system level modules for the X.25 network model are shown on Figure 3.1. This diagram is composed of two system blocks denoted X.25-System-A, and X.25-System-B, which are explained below. Both systems are interconnected by a bi-directional link. The other three blocks are sinks used to collect system's statistics. The parameters shown on Figure 3.1 are explained below.

Both X.25-System-A (Figure 3.2) and X.25-System-B (Figure 3.3) are made up of five X.25 Stations, and a router (Router-A, or Router-B). Each X.25 Station is connected to the central router through bi-directional links. A File Transfer Request (FTR) block is connected to each X.25 Station. Each one of these modules are explained in the following sections.

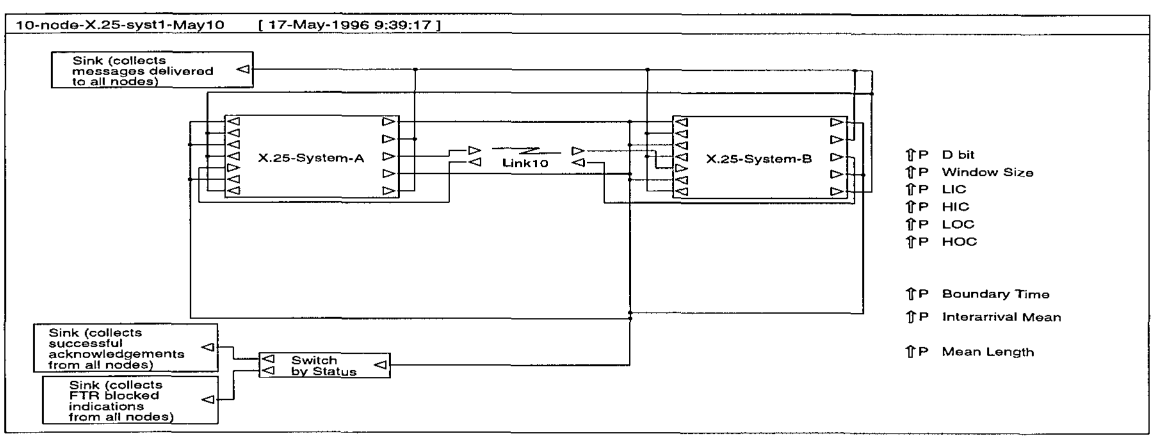

Figure 3.1 - System level modules.

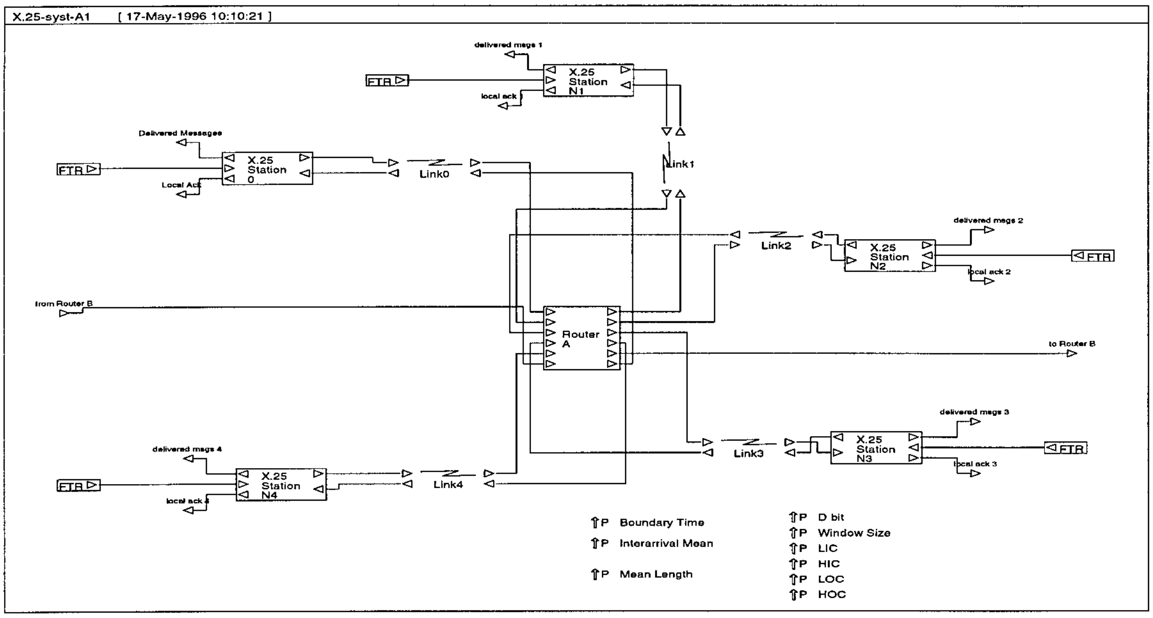

Figure 3.2 - X.25-System-A.

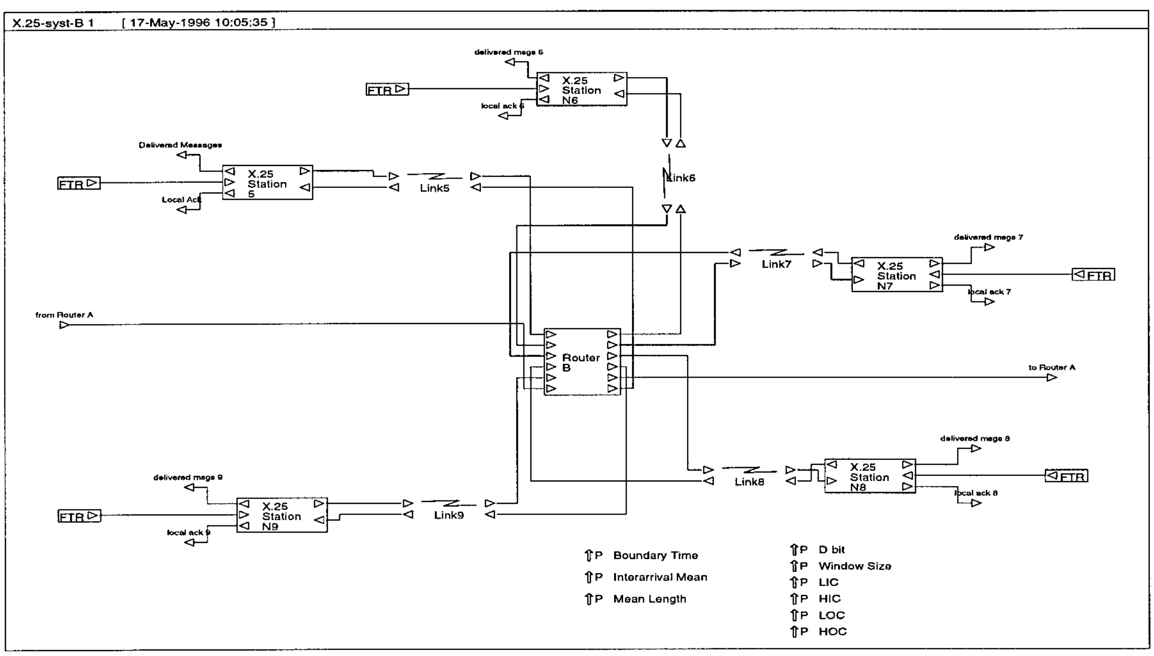

Figure 3.3 - X.25-System-B.

3.3.2.1.1 X.25 Station

The X.25 Station module is shown in Figure 3.4. It receives new file transfer requests through the New Request input port. All file transfer acknowledgements exit the station via the Local Ack port. File transfer indications exit via the Delivered Messages port at the destination DTE. The two external ports on the right of the diagram provide the interface between the DCE and the internal network.

The blocks between the DTE and the DCE model the link layer. They select the packet length and calculate a service time, according to the link rate between the DTE/DCE pair. The service time is used as a parameter to the FIFO with Servers block. Each packet is delayed for the specified service time before exiting the queue.

The DTE block receives requests and looks for an outgoing logical channel in the Ready state (the LC state is kept in internal memory). If it finds one, it changes the state of the logical channel to Waiting and a Call Request packet is sent to the DCE. A copy of the request is stored and leaves the DTE through the Local Ack port when the call is completed. If there are no logical channels to the DCE in the Ready state, the request is rejected, and sent out of the DTE through the Local Ack port with its Status of Request field set. These same operations happen when the call is blocked by the remote DCE.

Embedded modules on the DTE implement a state machine for the sequences in a call. Incoming packets from the DCE are routed to one of the four state processors (Waiting, Clearing, DTE transmitting, and DTE receiving) according to the state of the LC. The check for an available position in the window flow control is also done inside the DTE by a simple check on the value of this memory vector.

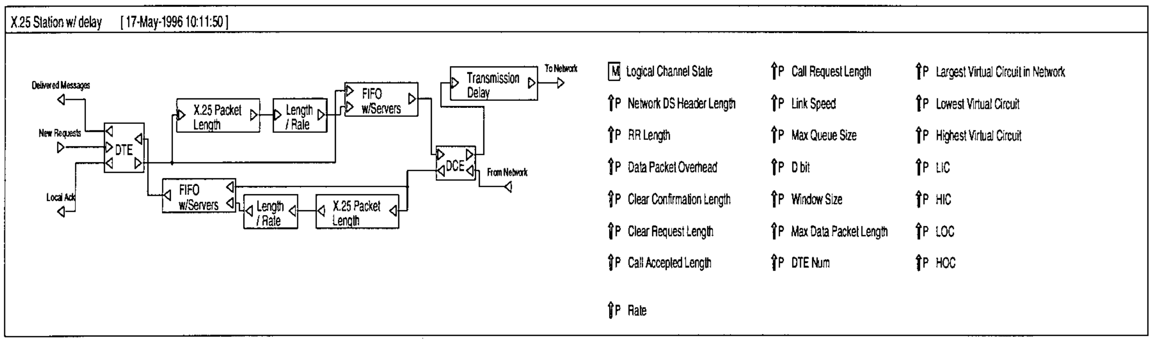

Figure 3.4 - X.25 Station.

The DCE block provides the interface between the DTE and the network. It also transmits Receiver Ready packets for data packets that do not have their D bit set. When Call Request packets arrive from the DTE, a virtual circuit number for the network is allocated. The DCE then builds a X.25 Network DS that is used to encapsulate all X.25 packets going to the network. If the encapsulated packet is a Clear Confirmation packet, the DCE sets the LC used for the call back to the Ready state and frees the network virtual circuit. After leaving the DCE, the Network DS packet passes through a Transmission Delay block before entering the network.

All packets coming from the network have their encapsulated X.25 packet extracted by the DCE. The DCE also performs allocation of the incoming logical channels to the DTE for packets coming from the network, by searching the LC state memory for a logical channel in the Ready state. If no channels are available it rejects the call request by returning a Clear Request to the source DTE.

The parameters Lowest Incoming Channel (LIC), Highest Incoming Channel (HIC), Lowest Outgoing Channel (LOC), and Highest Outgoing Channel (HOC) are exported to higher modules up to the system level. They are used to keep track of the state of the channels between each DTE/DCE pair, and identify packets of a same call at the local or remote station. The LIC, and HIC parameters are assigned by the DCE, while the DTE assigns the LOC, and HOC. Up to 4095 calls may be in progress at a given DTE. The D bit and the Window Size are also exported up to the system level.

Two others parameters that are set by the DCE are the Lowest and Highest Virtual Circuit in the network. The virtual circuit (VC) number identify packets from the same call in the network. A DCE converts a logical channel number to a virtual circuit number for outgoing packets, and does the opposite conversion for incoming packets. To disable the chance of more than one DCE assigning the same VC number for a call, there are non-overlapping ranges of virtual circuit numbers that each DCE can use.

3.3.2.1.2 FTR

A FTR module (Figure 3.5) creates file transfer requests in the form of X.25 External data structures. The interarrival times are exponentially distributed. Although in the real system the traffic destination is not uniformly distributed the model assumes it is because the actual distribution was not available. The length of each file transfer is also exponentially distributed.

3.3.2.1.3 Routers

Both routers, Router-A, and Router-B, shown in Figure 3.6 and Figure 3.7 respectively, share the same functions. The 6-Way-Switch routes each arriving packet to its proper destination through its six output ports. Five of the output ports are directed to the four external X.25 Stations directly connected to that switch. The remaining port is connected to the other router. The 6-Way-Switch module for Router-A is shown in Figure 3.8.

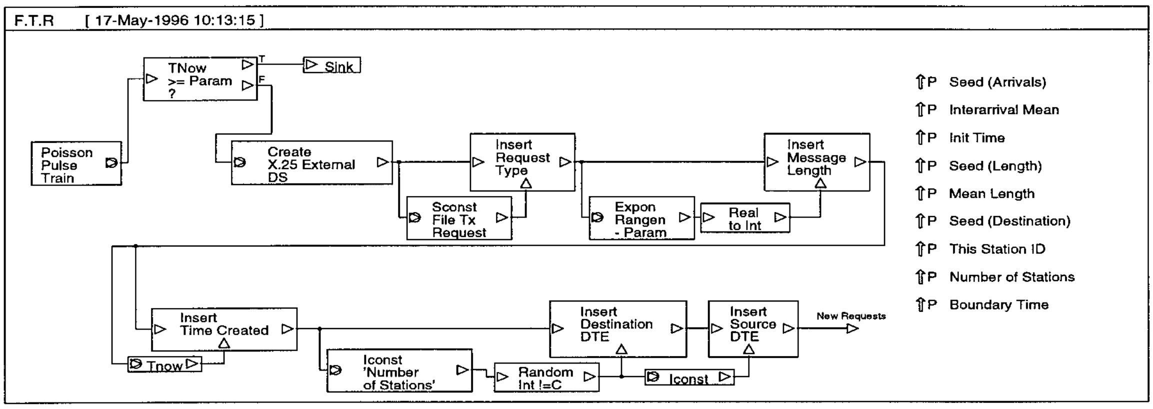

Figure 3.5 - FTR module.

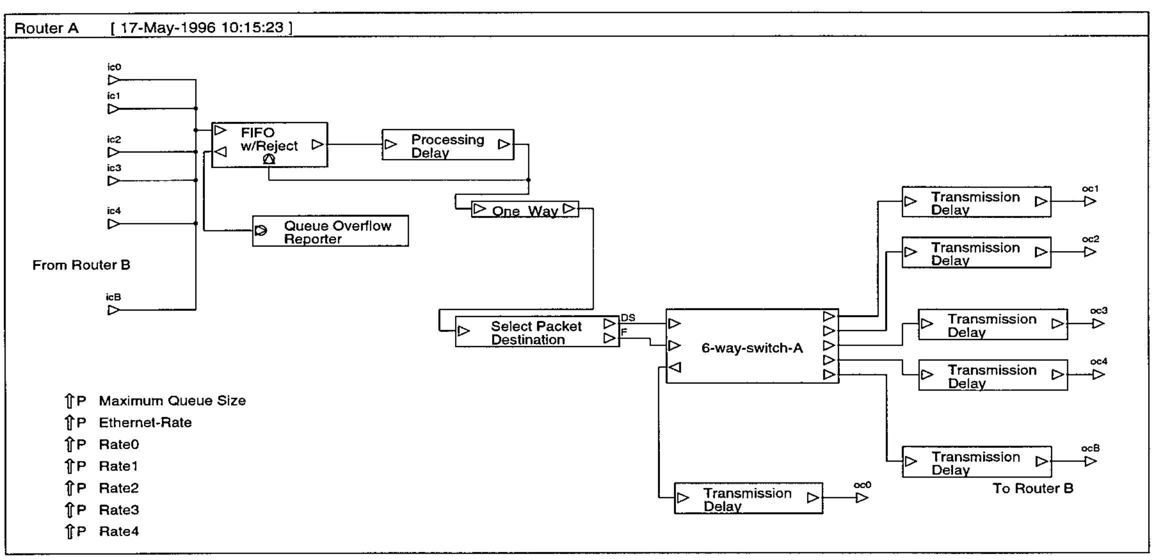

Figure 3.6 - Router-A module.

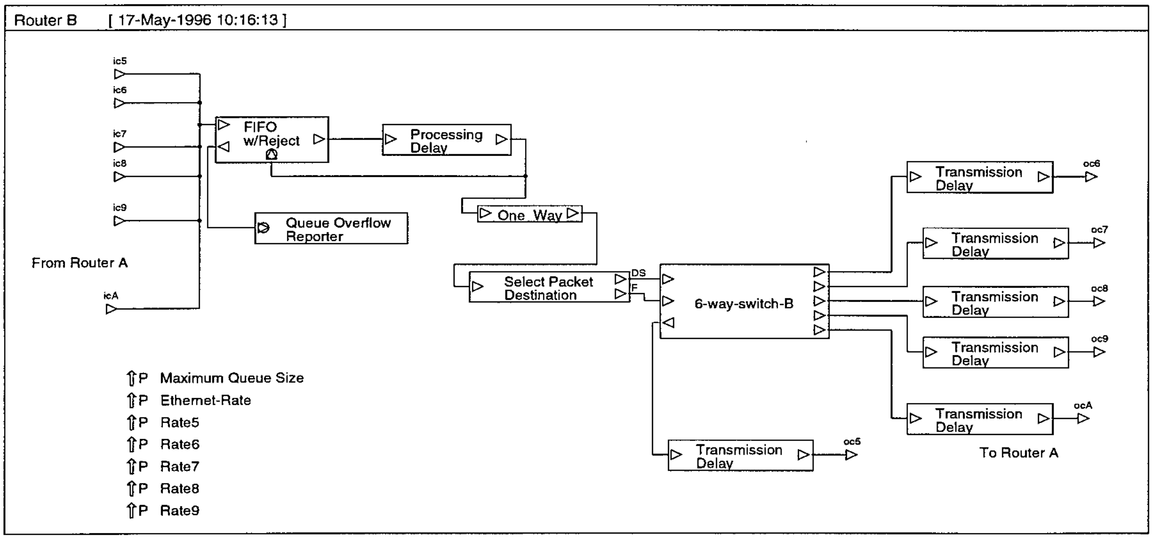

Figure 3.7 - Router-B module.

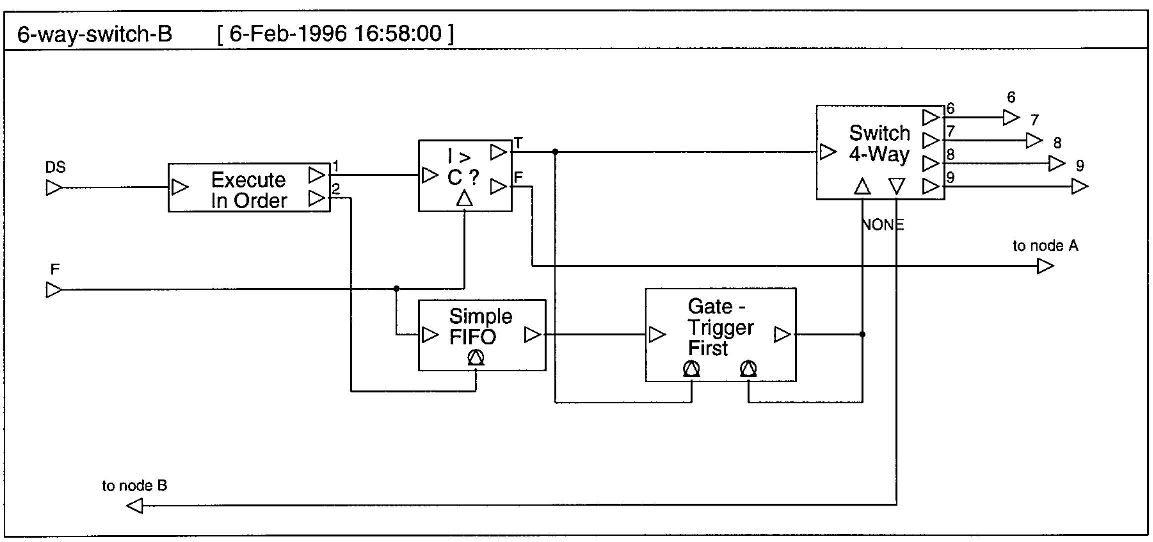

Figure 3.8 - 6-Way-Switch module.

3.3.2.2 X.25 - LAN Interconnection Network

This model was built to analyze the behavior of the X.25 network when used to interconnect local area networks. The system level modules are shown in Figure 3.9. Similar to the previous model, it is composed of two system blocks denoted X.25-LAN-

System-A (Figure 3.10) and X.25-LAN-System-B. The two sink blocks are used to collect system's statistics. This model differs very little from the one just presented. The differences are the traffic source and the X.25 station that are described below.

3.3.2.2.1 LAN module

The LAN module (Figure 3.11) creates X.25 External data structures that are sent to the X.25 Stations as new requests. It is composed of a traffic generator that models the traffic from an Ethernet LAN (Figure 3.12). This

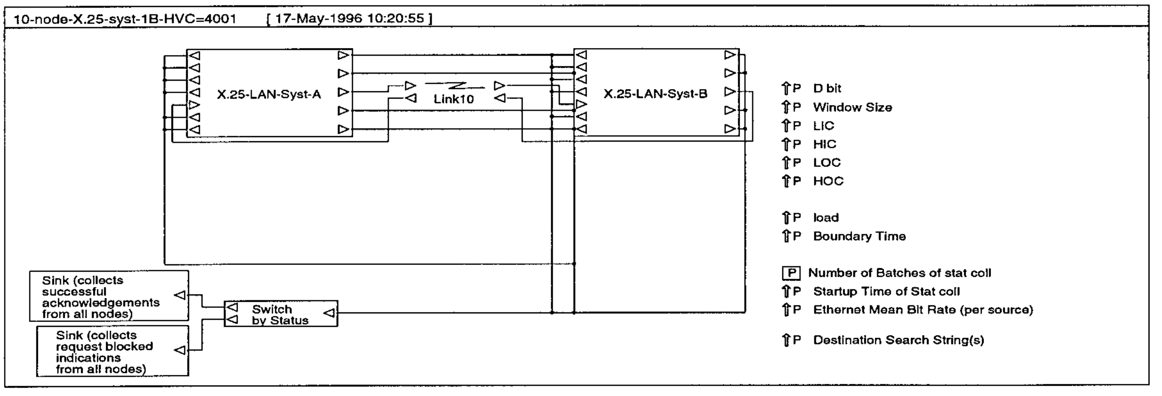

Figure 3.9 - X.25-LAN Interconnection system model.

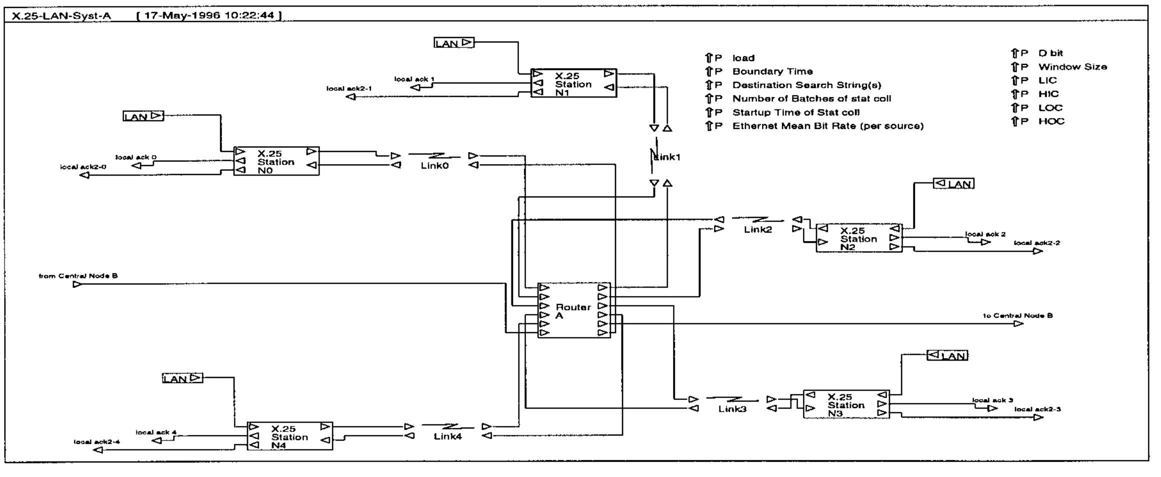

Figure 3.10 - X.25-LAN-System-A module.

model is a statistical model taken from [Nei91]. The study performed by Neir indicates that the packet distribution of a LAN traffic is basically bimodal, consisting of large data packets, and small acknowledgments packets. The model randomly sends a small or large burst. The probability of a small burst being generated is 0.4. After a burst is sent there is an exponential silence time before another burst is generated. Each burst passes through a switch that controls what percentage of the traffic being generated is going into the network. This percentage is specified by the Load parameter. The bursts that go into the network generate a MAC-Data.req DS that implements a request made by an Ethernet LAN. This data structure is then encapsulated into a Router DS. The parameters from the LAN module that are passed to higher levels of the system are explained below:

• Load - percentage of traffic generated that goes into the network;

• Number of Ethernets - number of Ethernets LANs a source is modeling;

• Boundary Time - controls the generation of new bursts (new packets);

• Name of Local Node - string identification for the LAN;

• Destination Search String(s) - specifies the destination the source is sending to. The traffic generated is evenly distributed between the destinations that match the string;

• Ethernet Mean Bit Rate (per source) - aggregate mean rate of traffic from all LANs this source is modeling. Must be less than "Number of Ethernets * 10 Mb/s";

• Seed - Seed value for the random number generator.

The statistics for the Ethernet traffic are shown on Table 3.4.

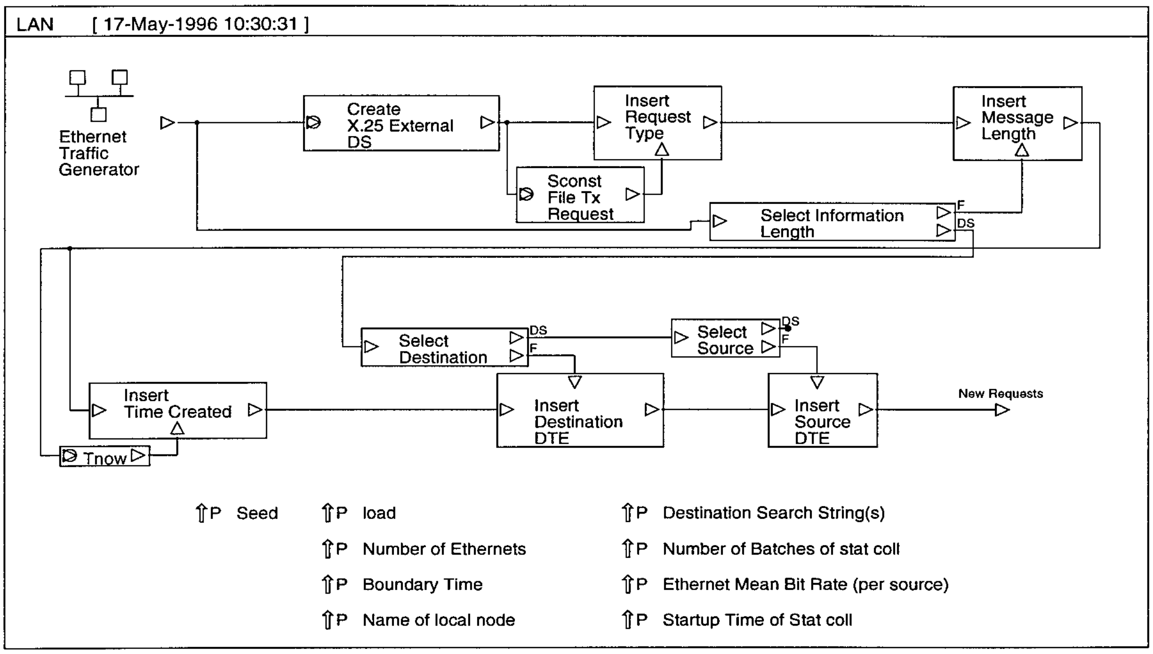

Figure 3.11 - LAN module.

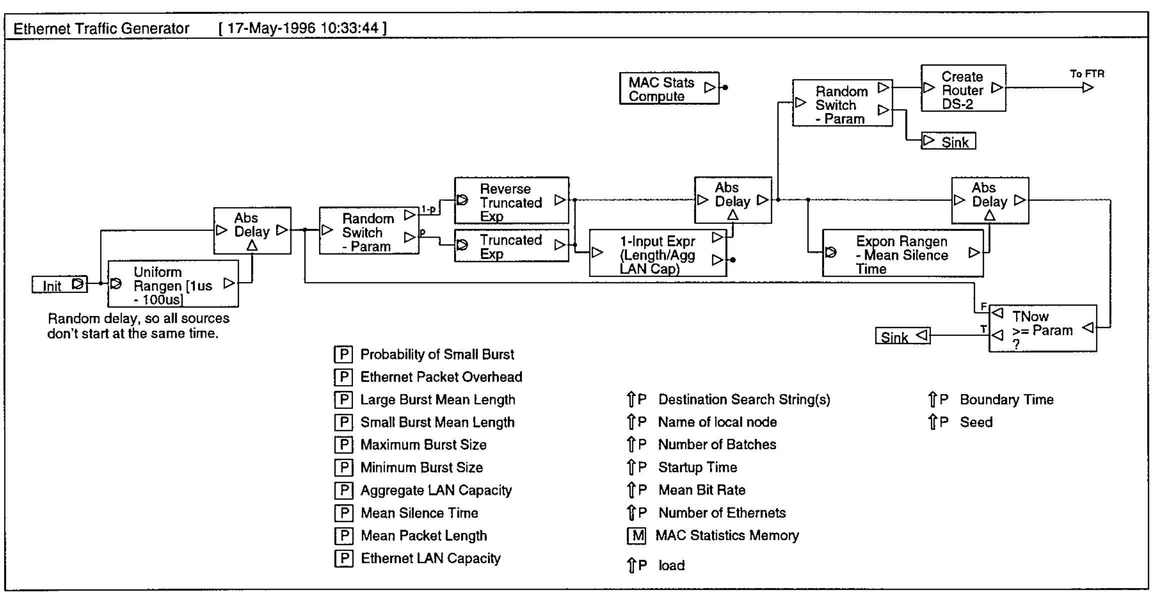

Figure 3.12 - Ethernet model.

Table 3.4 - Ethernet traffic statistics.

mean packet length (burst size)	5386 bits
maximum burst size (max packet length)	12000 bits
minimum burst size (min packet length)	368 bits
small burst mean length	24 bits
large burst mean length	3936 bits
probability of sending a small burst (ack)	0.4

3.3.2.2.2 Enhanced X.25 Station

The internals of the Enhanced X.25 Station are depicted in Figure 3.13. It contains a DTE block, a DCE block, and additional structures. The blocks above the dotted line simulate the reuse of virtual circuits, suggested in [Jom89, BaW91] for the interconnection of LANs, which are connectionless networks, with the connection-oriented service offered by X.25 networks. Whenever a packet arrives through the New Requests port, the X.25 Station examines the destination address to determine whether a virtual circuit exists or not. If it exists, and it is on the data transfer phase (DT Out), the length of the incoming packet is added to the remaining number of bytes yet to send on this VC. Since there are no actual data bytes transmitted on the Data Request DS, only the Length field is extracted. If no VC exists to the destination the new request is directed to its DTE and a regular X.25 call is initiated. The same happens if the established VC is not on the DT Out phase.

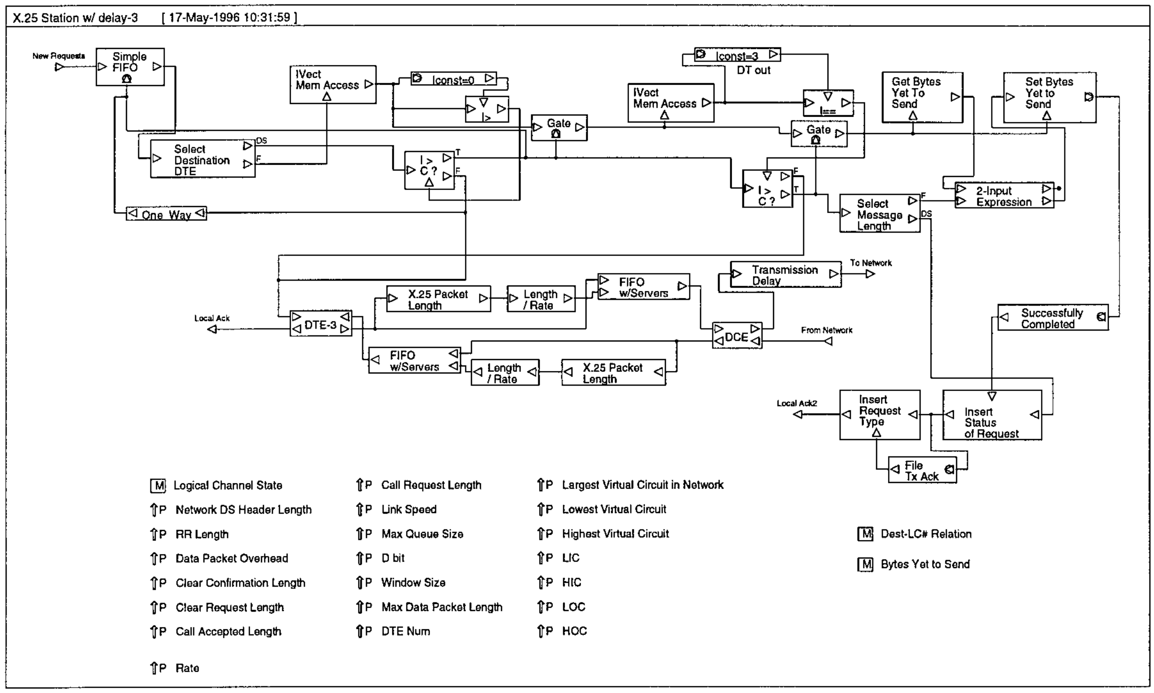

Figure 3.13 - Enhanced X.25 Station.

After a request has been serviced, an acknowledgement exits the station through one of the Local Ack ports. The DCE ports denoted To Network, and From Network are the interface between the station and the network. The function of the modules below the dotted line are the same as explained in Section 3.3.2.1.1.

3.4 Verification and Validation

Model verification consisted of three types of tests: the verification that each possible route in the system was working as expected, verification of all blocks involving delays, and service times to ensure that these were being added correctly, and packet trace verification.

3.4.1 Individual Routes Verification

Every possible route that a packet can take in the system was evaluated to check if the switches were working as expected. These tests were made by choosing one node and replacing the Init block from its traffic generator by an One Pulse block. The other nodes had their Startup Time parameters set to a number greater than the simulation time (TStop). This way a single packet would transverse the network. The destination of each packet was user specified by setting the parameter Destination Search String. Generic probes were placed throughout the system to monitor the packet's route. The results showed that all routes are working as expected when a single packet transverse the system.

3.4.2 Delay Verification

In order to verify that the delays added to the system were implemented correctly, tests similar to the one explained in the previous section were performed. Probes were placed in input, and output ports of queue, and delay blocks. With ten instances of this tests every delay block in the system was verified. Final results were obtained by comparing the difference between the time when the packet gets to its destination, and the time when the packet was generated with the sum of all delays that were added to this packet. The result of this sum was exactly the same as the difference, on every 10 instances, proving that the delays are being correctly inserted in the system.

3.4.3 Packet Trace Verification

The phases of a virtual call were checked for proper function. With the Type field of the X.25 Packet DS, every type of packet can be determined. Figure 3.14 shows all phases that a call goes through. The Y axis is the one logical channel number multiplied by 10 plus the Type field. This computation is done to help visualize all phases of a single logical channel.

On Figure 3.14, the first point on the plot (2, 201) is a Call Request packet, followed by a Call Accepted packet at (3, 202). At TNOW = 3, what seems like one packet is really two Data Packets that were sent. The two Receiver Ready packets at (x_1, 206), and (x_2, 206) are related to the two Data Packets. A Clear Request packet (7, 203) is sent at TNOW = 7, and a Clear Confirmation is received at TNOW = 12.

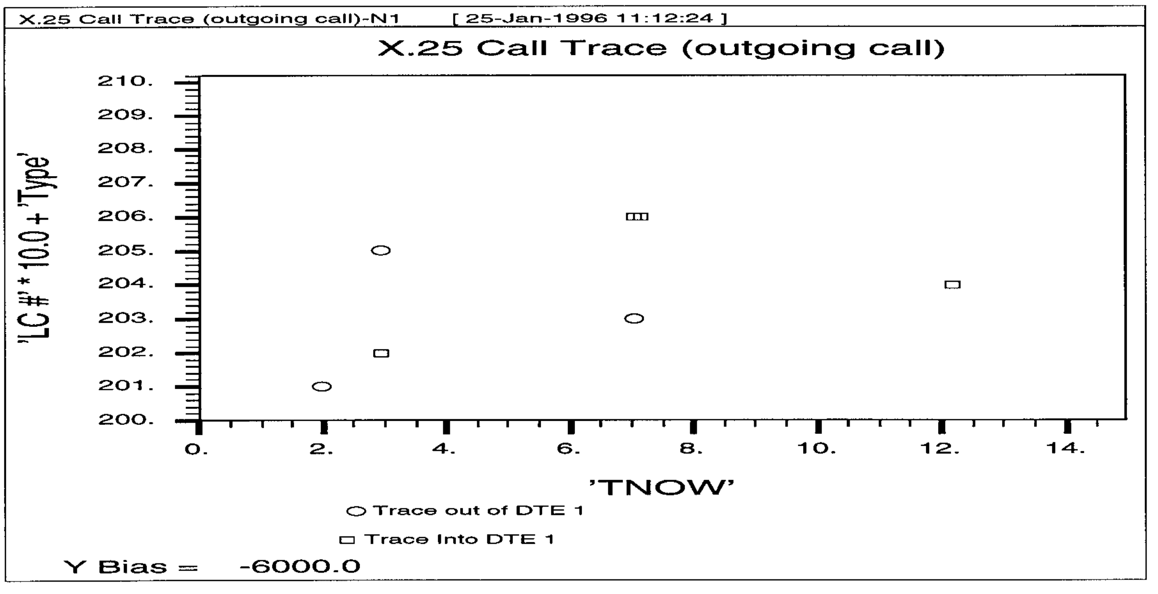

Figure 3.14 - Packet trace.

3.4.4 Model Validation

From the three key aspects for model validation listed by [Jai91] only two are validated in this section. The third key aspect, output values, and conclusions, is validated in Chapter 5. The validity tests of assumptions, and parameter values were made by comparison with previous works, and real systems.

The modeled aspects of the protocol were similar to the ones used on previous works on X.25 systems [Cad94b, ReF90]. Since this model is a scaled down representation of a real system, its representativeness was validated through an interview with [Bar96], who has a considerable knowledge of the real system, and provided reference material for this effort. Input parameters were validated against previous works, already referenced, and were within the limits defined by the X.25 protocol recommendations. Real system measurements were not available.

3.5 Summary

This chapter covered the X.25 model. The portions of the protocol that are being implemented in the model were explained. An overview of the model implementation was given. This overview did not covered the low level modules of the model, however it gave an explanation on the major functions of the model. Finally, verification and validation of the model was discussed in Section 3.4.

4 ATM Model

4.1 Introduction

The ATM network model presented in this chapter is implemented to compare its performance against the X.25 network model. To facilitate this comparison, some aspects of the model, such as topology, and number of switches, are implemented similarly to the X.25 model. Section 4.2 discusses what portions of the ATM protocol are modeled. The data structures used on this model are described on Section 4.3. A description of the main modules used on the model, together with a model description is presented in Section 4.4. The model verification and validation is described in Section 4.5. Section 4.6 summarizes this chapter.

4.2 Modeled Aspects

In a comparison study between the OSI reference Model (RM) and the B-ISDN Protocol Reference Model (PRM) presented in [Sta96], the author examines the data communication services of both reference models. The functions performed by the network layer of the OSI RM, as presented, are basically the same as the ATM layer on the B-ISDN PRM. Therefore, this model concentrates on functions similar to those of Layer 3 of the X.25 protocol presented in the previous chapter. These functions are located in the ATM layer of the B-ISDN PRM, which is concerned with operations on ATM cells.

This model reuses modules from [Cad94a]. The model assumes that all connections are set up prior to the start of the simulation. This assumption is done because of two reasons. First because the simulations are short due to the high volume of traffic generated. The second reason is that in real systems few calls are set up or cleared in this short period of time. Since end-to-end delay is not analyzed in the comparison between both models in this study, the call set up or clear overhead is not considered.

4.3 Data Structures

Four data structures are used in the ATM network model. They are the Cell, Converge Sublayer/Segmentation & Reassemble (CS/SAR), MAC-Data.req and Router data structures. The Cell DS represents the ATM layer, while the CS/SAR and Router DS represent the ATM Adaptation Layer. The MAC-Data.req DS, as explained in the previous chapter, represents a request made by a LAN. The others data structures are explained below.

4.3.1 Cell DS

Since cells are what flow through the ATM layer, the Cell DS has almost the same fields as the ATM cell. The VPI and VCI fields have a different meaning on the ATM model. They do not implement the concept of a virtual path and virtual connection as explained on Chapter 2. Here, these two fields carry the source and destination node address. These addresses are assigned automatically at the start of the simulation to LAN nodes. The PT field identifies the type of cells, as a normal or a Hello Cell. The Cell Loss Priority (CLP) field when set, indicates that a cell is more likely to be discarded. The HEC field is not used in this simulation. The Info fields carries the payload that can encapsulate any data structure. The Time Created field stores the time this DS is created. The Path field is a vector that stores the path that a cell takes to gets to its destination. This vector is composed of identification numbers of modules that a cell passes through. Each Internetworking Unit, B-NT2, and Cell Switch module has an identification number. The IWUs are numbered first, starting from zero. The numbering continues with B-NT2 modules and finish with the Cell Switches. Each index of the vector indicates a hop of the path. Index zero indicates where the cell entered the ATM network. The Next Hop field indicates the index in the Path Vector of the next hop. This field helps the system to ensure that the cell is going through the correct path by comparing the value of the path (Next Hop) and the identification number of that node.

4.3.2 CS/SAR DS

This data structure models the activity in the ATM Adaptation Layer 5. The Info field encapsulates a Router DS. The Length field carries the length of the packet being delivered. This length is added to the CS overhead and any padding necessary to make the length of the packet a multiple of 48 bytes.

4.3.3 Router DS

This data structures encapsulates packets coming from a LAN. The Information field carries a MAC-DATA.req DS (Chapter 3) that is delivered on the destination node. The Source and Destination fields carry addresses of LAN nodes. The Information Length carries the length of the Information field.

4.4 UNI and NNI Modules

The main modules can be divided in two groups. The first, UNI Modules, is composed of the Ethernet Traffic Source Model, the IWU, the B-NT1 and B-NT2 modules. The second group, consists of the memory initialization module (Init Mems), and the Cell Switch module.

4.4.1 Ethernet Traffic Model

This model is a specialization of the traffic generator from the LAN module explained on Section 3.3.2.4. The Ethernet model is shown in Figure 4.1. Its additional parameters are explained below:

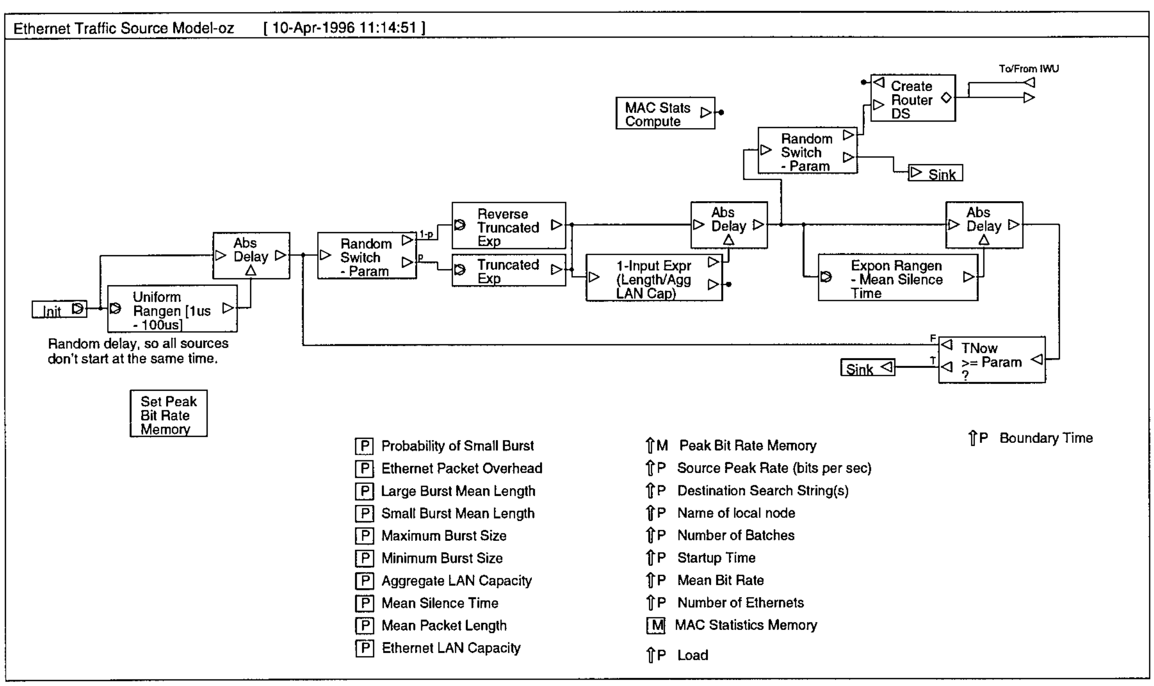

Figure 4.1 - Ethernet model.

- Peak Bit Rate Memory - This memory is exported up to the system level and is initialized by the memory initialization module (Init Mems), explained below. It is represented by a matrix were the row index is the source address, and the column index is the destination address. A source, destination pair is seen as a VC, where the content of this pair holds the peak bit rate allowed for this VC. The value of peak bit rate for each VC from a particular source is obtained by dividing the Source Peak Rate parameter by the number of destinations to where it transmits. If a cell violates this rate it has its CLP field set as low priority.

- Source Peak Rate (Bits per Sec) - Peak rate allowed from this source.
- Destination Search String - Indicates to where this source is sending. Traffic is evenly distributed over all sources that match this string. More than one string can be specified.
- Mean Bit Rate - The mean bits per second from this traffic source.
- Load - This parameter controls the amount of traffic that leaves this traffic source.

4.4.2 IWU

This device is a Terminal Adapter (TA) for a LAN. Its function is to interface non B-ISDN equipments to B-ISDN compatible equipment. It performs its function by segmenting into ATM cells the packets received from the Ethernet Traffic Model, and reassembling packets from cells. Assembly, and reassembly are performed by the AAL5 internal modules. Assembly is done by adding an overhead to the incoming packet (router DS) and padding its length to be a multiple of 48 bytes. The last cell in a sequence has a payload type of one. When such a cell is received, the AAL5 start reassembling the received cells into a packet. If any cell is missing or added, the packet is not delivered to the destination.

At the beginning of the simulation, an IWU sends cells with a payload type equals 2 (Hello Cells) throughout the network, to establish the network topology. When a Hello Cell passes over a link, its Information field is filled with the length of the link interconnecting both points that are exchanging Hello Cells. When a Hello Cell is received, the content of the information field is read and placed in the appropriate location in the Cost Matrix parameter.

The IWU parameters are explained below. The Peak Bit Rate Memory parameter was explained in Section 4.4.1. The IWU is shown in Figure 4.2.

- Network Address List - This memory is exported to the system level, and is used by routing modules to determine were to send cells to.
- Cost Matrix - Also exported to the system level. This memory is initialized at the start of the simulation with Hello Cells obtaining the length between every networking device (IWU, B-NT2, and Cell Switches). With the ID Number of two modules, the Cost Matrix knows if two modules are directly connected.
- Router Address List - A string that determine the address of the network sitting behind the IWU.
- Maximum Queue Size - The maximum number of cells allowed in the queue that holds cells that are leaving the IWU.
- Link capacity - Capacity of the link leaving the IWU to the ATM network. This parameter is used to calculate the time cells will wait in the queue before being sent out of the IWU. The capacity used is the rate available to ATM cells. These rates are obtained after all appropriate overheads have been removed. Table 4.1 lists the available rates.

- ID Number - A unique non-negative integer that must be different for every module that has this parameter.

Table 4.1 - Rates available to ATM cells [Cad94a].

Transmission Type	Total Bit rate (MB/s)	Rate Available to ATM cells (MB/s)
DS1	1.544	1.536
DS3	44.636	40.536
STS-1 (SONET)	51.840	49.536
TAXI	100.000	98.148
STS-3 (SONET)	155.520	149.760
STS-12 (SONET)	622.080	599.040

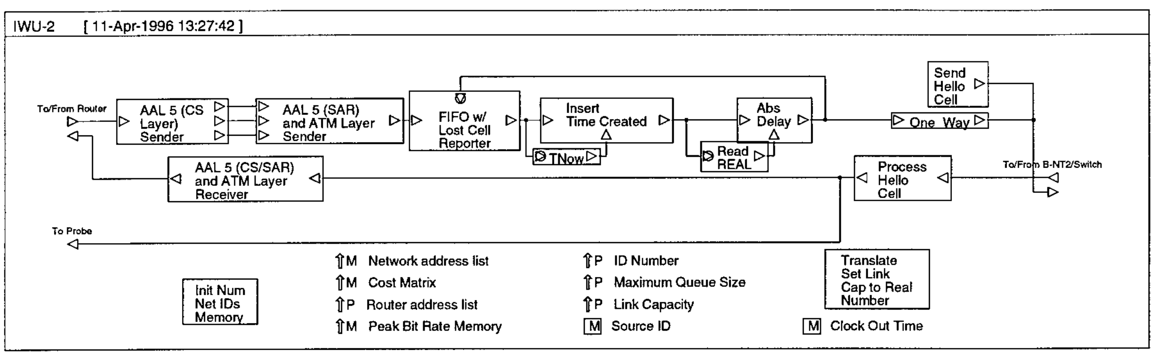

Figure 4.2 - IWU model.

4.4.3 B-NT2

The B-NT2 module, shown in Figure 4.3, multiplexes cells streams together and functions like a private branch exchange (PBX). The architecture of the B-NT2 module is of a Bus-Type Switch with output queuing [HaH93]. For this type of switch, an access algorithm allocates the bus to each input controller at constant intervals. There are no buffers in the input controller since each of these are able to service a cell before the arrival of another cell. This is achieved by a high-speed time division multiplexing bus, where conflict free transmission is guaranteed if the total capacity of the bus is at least the sum of the capacities of all input links [PrS87]. Queuing only occurs on the output port of the switch.

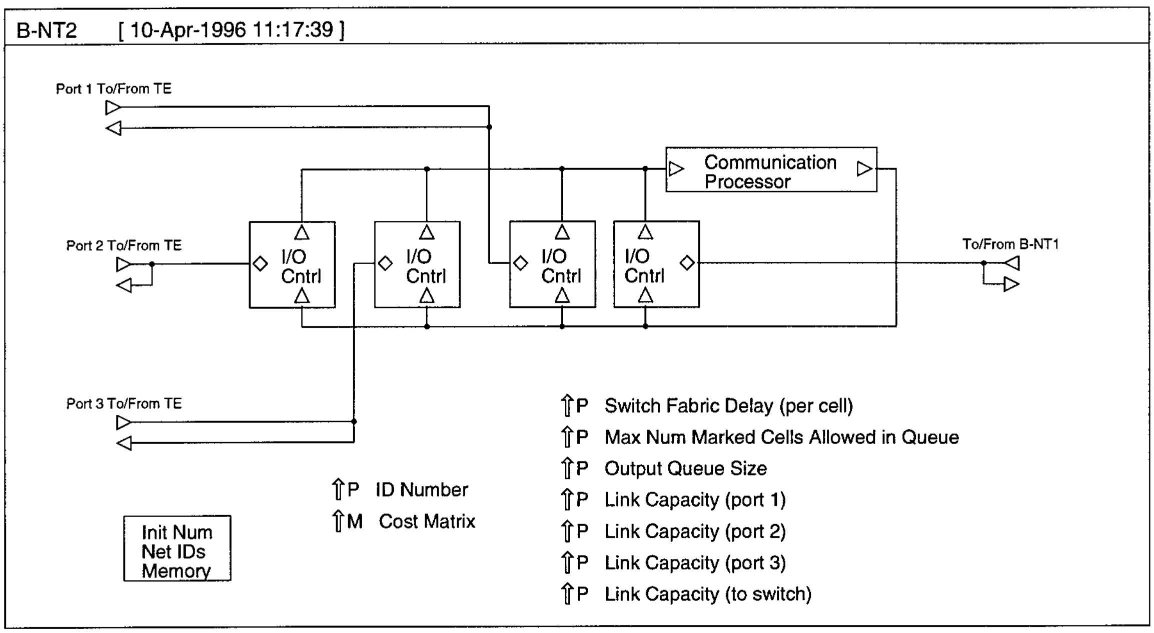

Figure 4.3 - B-NT2 module.

Incoming cells from IWUs are delayed by the amount of time specified in the Switch Fabric Delay parameter set at the system level. This parameter corresponds to the mean bus propagation time of a Bus-Type Switch, that is the average amount of time for a cell to move from the input of the switch to the output buffer. Cells are then routed to a FIFO with partial buffer sharing queue. For this type of queue, low priority cells (CLP bit set) only occupy a portion of the queue. This limited space is determined by the Max Num Marked Cells Allowed in Queue parameter, set at simulation time. When this limit is reached, marked cells are then discarded.

Parameters ID Number, Cost Matrix, and Link Capacity have the same meaning as in the previous modules. Switch Fabric Delay (per cell) is the mean amount of time a cell takes from the input of the switch to the output queue. The Output Queue Size parameter specifies the size of the queue in cells.

4.4.4 B-NT1

The B-NT1 module (Figure 4.4) performs access control. A monitor for the cell peak rate controls every VC from the source before the cells enters the ATM network. Any cell that violates the agreed upon rate has its CLP bit set and are more likely to be discarded. The monitoring algorithm was obtained from [ATM93] and is demonstrated in Figure 4.5. This algorithm is applied to every source, destination pair in a B-NT1 module. The parameters for this model are defined below, followed by an explanation of the algorithm. The Peak Bit Rate Memory parameter has the same meaning as in previous modules.

- Cell Delay Variation Tolerance (τ) - Represents any distortion that can be introduced between the time the cell is generate and the time the cell is checked for violations.
- TAT Memory - Matrix memory that holds the theoretical arrival time for each connection.

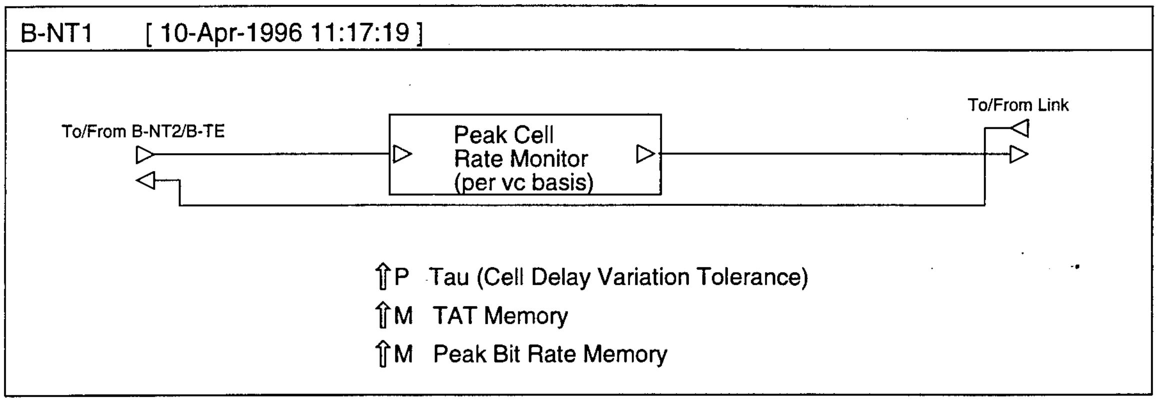

Figure 4.4 - B-NT1 module.

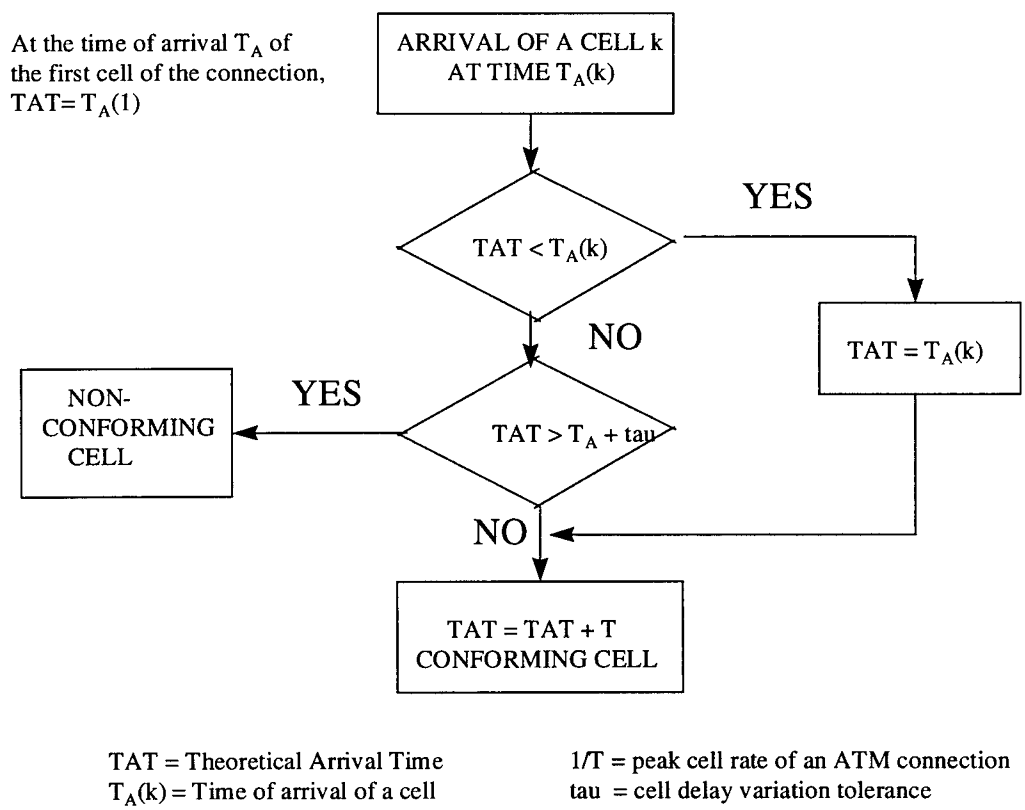

Figure 4.5 - Cell rate monitor algorithm [ATM93].

At the arrival of a cell at time Ta(1), the theoretical arrival time (TAT) is initialized to the current time Ta(1). For cells that follow, if the arrival time of the k^{th} cell, Ta(k), is after the current value of the TAT, then the cell is conforming and TAT is updated to the current time Ta(k), plus T (T = 1 / Peak Cell Rate of an ATM connection). If the arrival time of the k^{th} cell is greater than or equal to Ta(k) $-\tau$, then the cell is also conforming, and the TAT is increased by T. If the k^{th} cell is less than the TAT $-\tau$, then the cell is non-conforming and the TAT is not modified.

4.4.5 Init Mems

This module (Figure 4.6) initializes all system memories at time zero. Three of these memories were not discussed yet. They are the Num Net Ids, SAR Mem, and Num Cells Lost. The Num Net Ids memory holds the total number of IWUs, B-NT2s, and Cell Switches. It is used to dimension the Cost Matrix Memory. The SAR Mem memory is a square matrix with an entry for every source/destination pair. This entry holds the number of cells received since the last cell with a payload type of one (last cell in a sequence). This memory is used in the AAL5 (IWUs) to reassemble packet. After a packet is reassembled the length of the packet is checked against the value in the SAR Mem memory. The Num of Cells Lost memory is similar to the SAR Mem memory except that its entries hold the number of lost cells for each virtual circuit (source/destination pair).

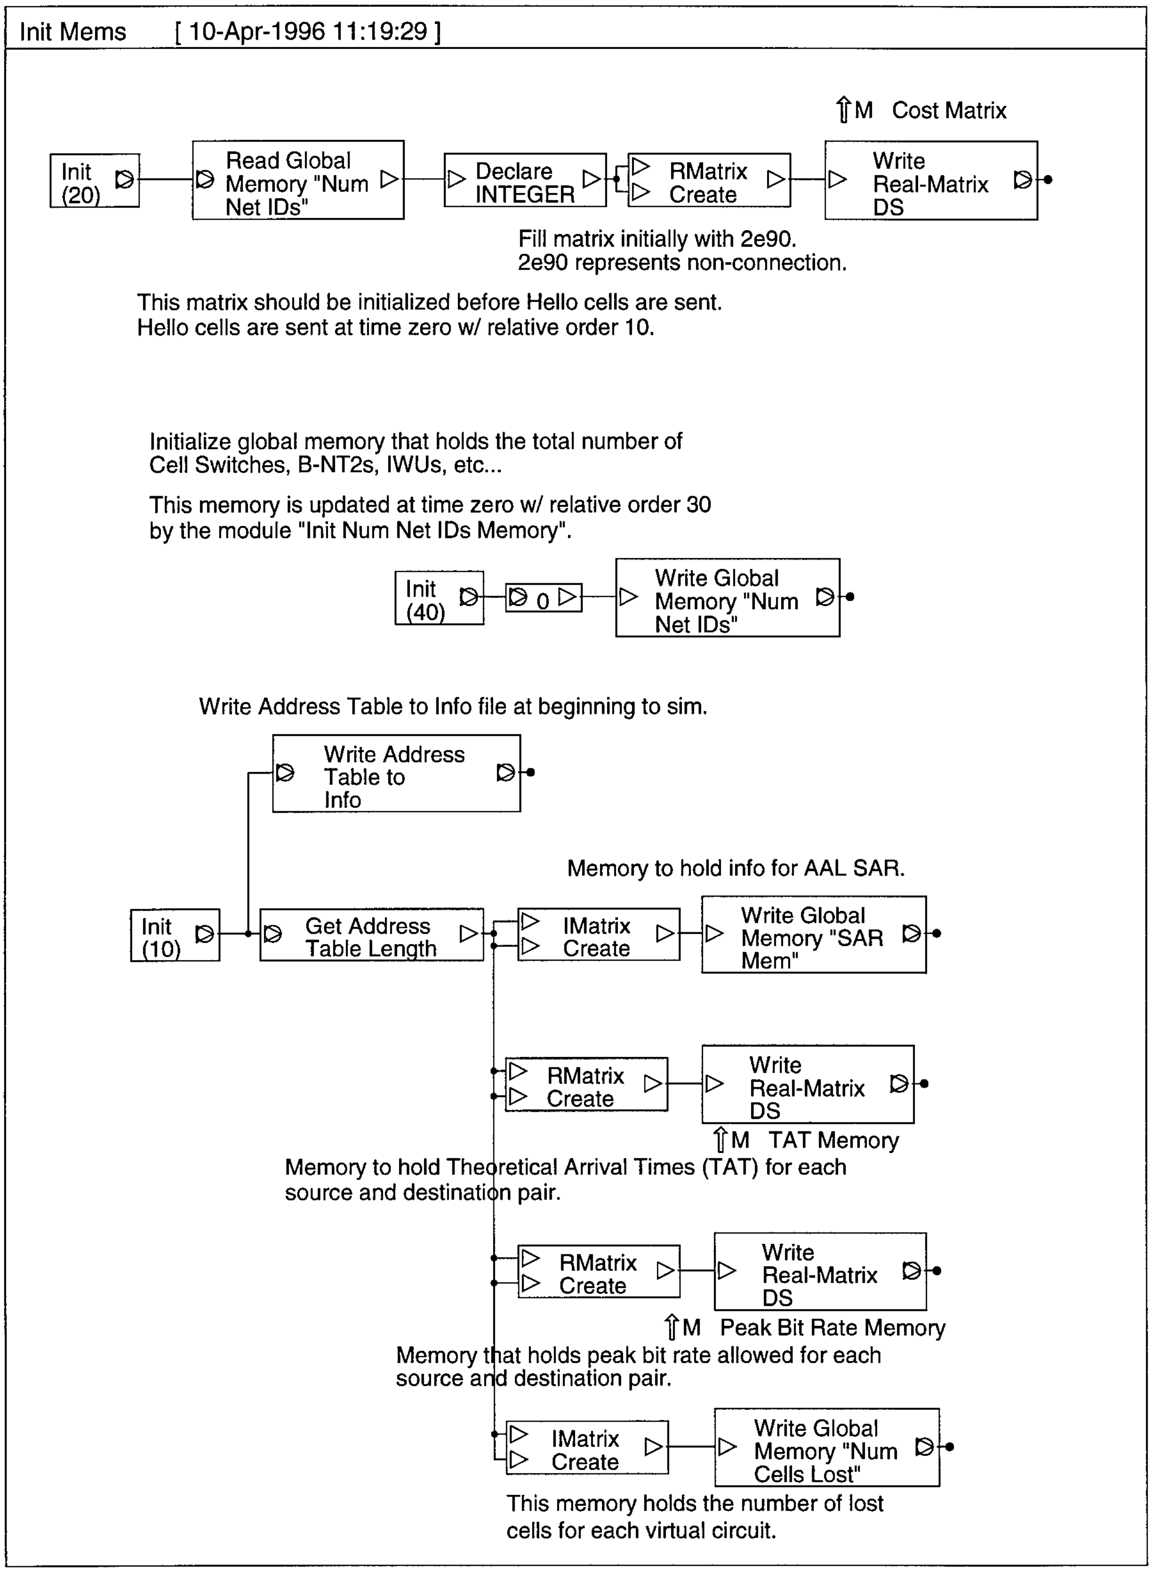

Figure 4.6 - Init Mem module.

4.4.6 Link

This module adds propagation delay to cells that pass through it. The Link Propagation Delay per Unit Length parameter indicates the amount of time it takes a bit to move a unit length. For fiber optic cable this value is approximately 5μs per kilometer [Cad94a]. The Link Length parameter is the length of the link connecting two networking devices.

Hello Cells passing through a link experience no queuing or propagation delays. The length of the link is put in the payload of the Hello Cell, and this information is used to create the Cost Matrix, which carries the length between every networking device.

4.4.7 Cell Switch

The ATM switch basically consists of two components: a switching element, and a buffer memory. The switching element transfers ATM cells from an input port to an output port. The buffer memory, located on the output port, stores cells for queuing. This module simulates a shared buffer memory switch [End93], where the cell buffer memory for the output queues is shared among all the switch output ports. The parameter Max Num Marked Cells Allowed in Switch, set at a higher level, indicates how many buffer spaces can be occupied by cells with the CLP bit set. The Number Buffers in Switch, set at simulation time, gives the total buffer space to be shared among the output queues. Parameters ID Number, Cost Matrix, Switch Fabric Delay, and Link Capacity were explained in previous sections. The Cell Switch module is shown in Figure 4.7.

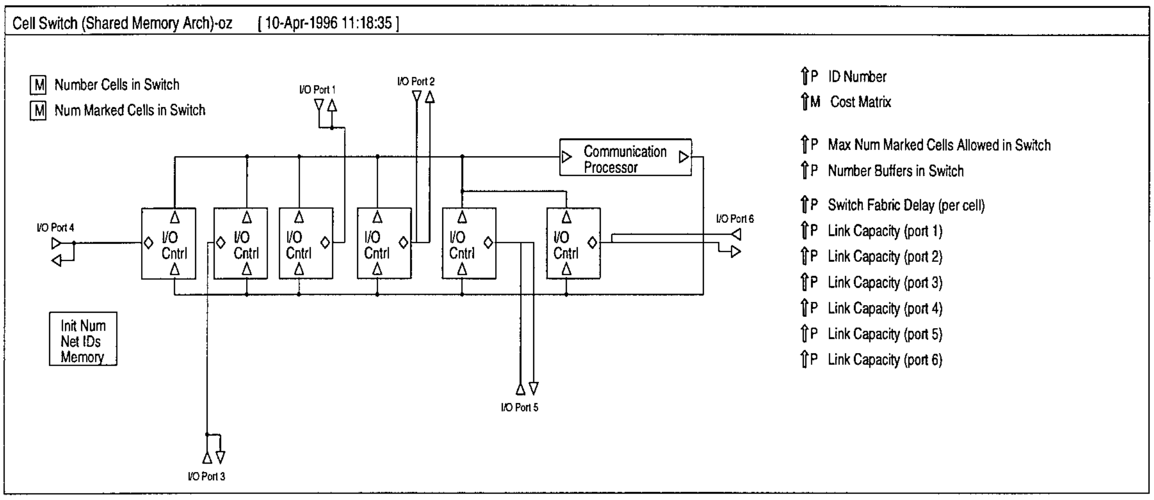

Figure 4.7 - Cell Switch module.

4.4.8 UNI

The UNI modules were built with the intent to reduce the number of blocks in higher levels. Each UNI represents a site. As shown in Figure 4.8, it is made up of three traffic sources, each connected with an IWU. These three traffic sources, give the flexibility to vary the traffic going into the ATM network. The B-NT2 module receives cells from the IWUs and passes them to the B-NT1. The ID Number parameters for the IWUs and B-NT2s are set inside the UNI blocks.

4.4.9 ATM Model Description

The ATM model shown in Figure 4.9 is based on the X.25 model described in the previous chapter. It has the same topology as the X.25 model. Network devices specific to an ATM environment were added. Furthermore, the number of Ethernet modules at each site has tripled. This is due to the fact that ATM is to carry greater traffic loads than

X.25. The system level module is shown in Figure 4.10. The Sink block collects cell statistics from traffic arriving on every IWU.

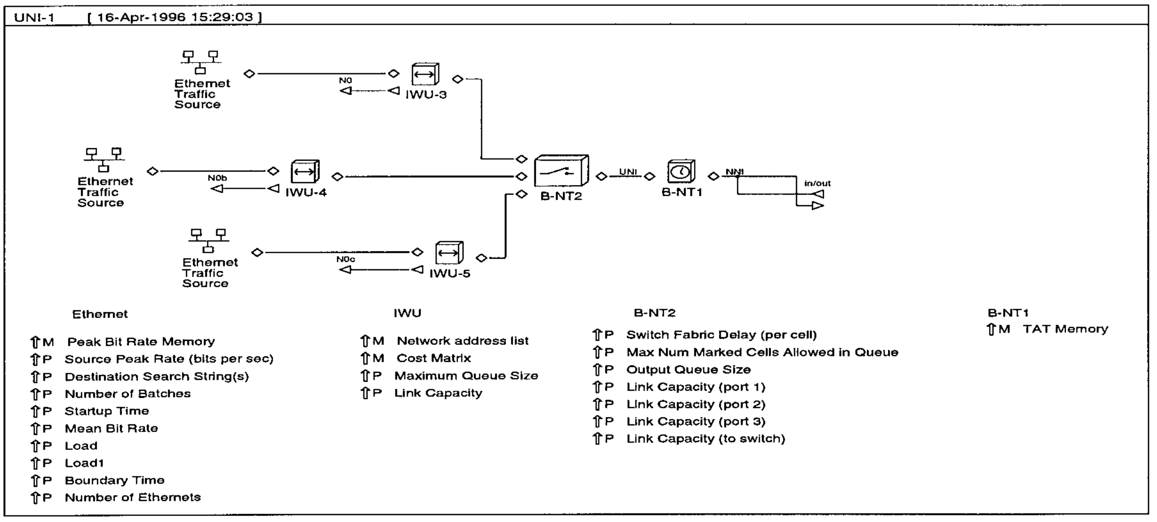

Figure 4.8 - UNI module.

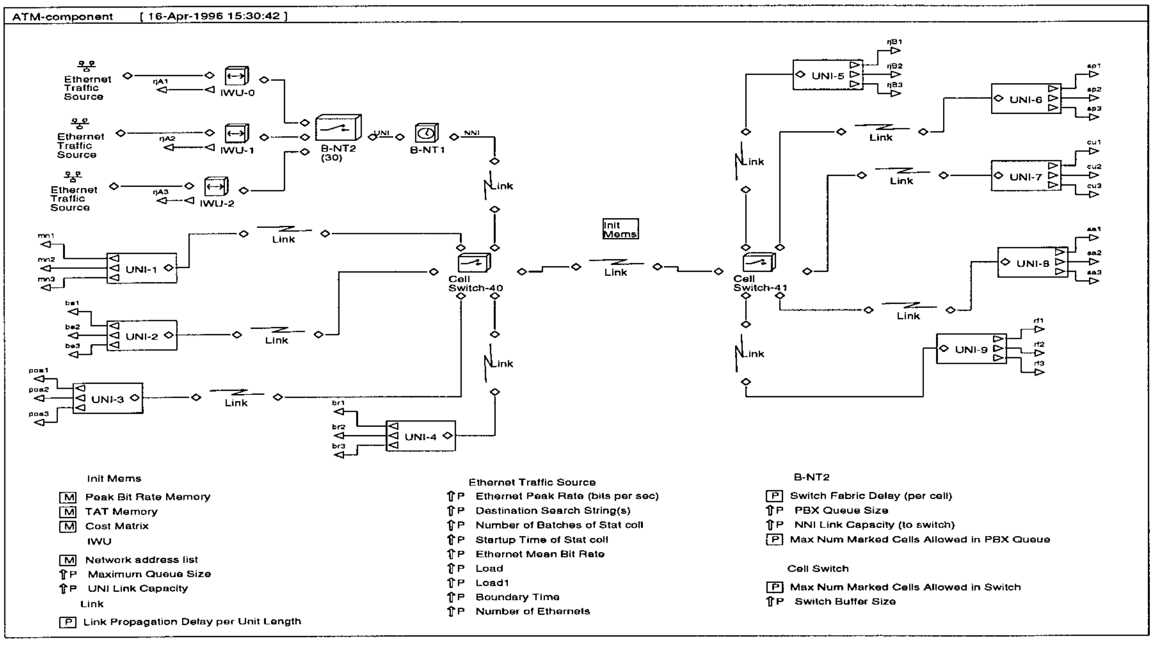

Figure 4.9 - ATM model.

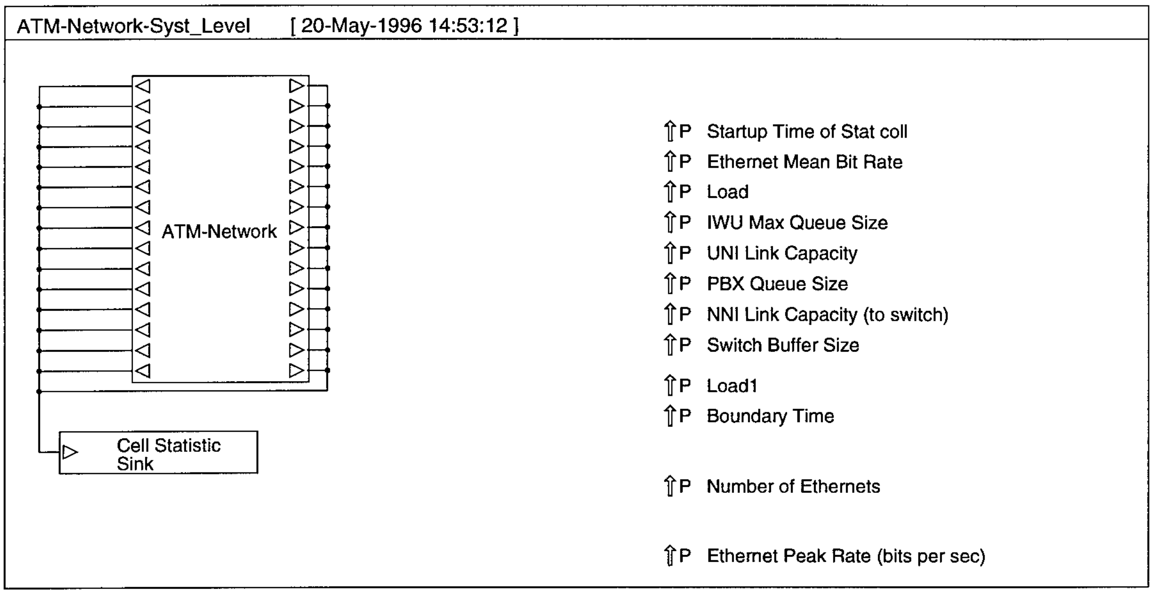

Figure 4.10 - System level model.

4.5 Verification and Validation

The ATM model verification consisted of the following tests:

1. Verification of ATM Library modules that were modified;
2. Verification of routes;
3. Delay verification;

Each of the above tests, and how the model was validated is explained in the following sections.

4.5.1 Modules Verification

The ATM Library modules provided by Designer were assumed to be free of any errors. This is due to the fact that other systems were already built using these same modules, and relying on the assumption that the provider would not distribute an unverified product. This way, only the modules that were altered were tested. These modules are the Ethernet Traffic Source and the Cell Switch. The first was tested to verify if the blocks added were performing the functions of controlling the traffic load out of the module, and controlling the time when to stop generating bursts that results in new packets. This test was performed using generic probes to verify the outputs of each block. The Cell Switch module had additional I/O Control blocks added. It was verified for correct routing by generating a single packet/cell and checking its destination using generic probes on the output ports of the switch. The results showed that both modules were performing as expected.

4.5.2 Routes Verification

Since both modules that perform routing functions (Cell Switch and B-NT2) were verified for proper functioning it was expected that there would be no routing problems. This belief was confirmed after the routes were verified. The test was performed by creating a single burst in a traffic source, and checking if the packet generated arrived at the proper destination (set by the Destination Search String parameter). Two traffic sources, each connected with a different Cell Switch were chosen as test sites. Having each test site sending packets to every one of the 29 possible Ethernet LANs, every possible route of an B-NT2 or Cell Switch was tested.

4.5.3 Delay Verification

The delay encountered by a packet was a function of the following factors: synchronization delays, transmission delays, queuing delays and propagation delays. This delay was verified by generating a single packet and recording the delays added to this packet/cell with the use of generic probes. The difference between the time the packet was generated and the receipt of the packet by the receiver matched the sum of all delays added to the packet. This test was performed with a source/destination pair to represent the system.

4.5.4 Model Validation

Similar to the X.25 model, the validity tests of assumptions, and parameter values were performed by comparison with other ATM models. The modeled aspects of the protocol were similar to the ones described on other ATM models [Cad94a]. The model was also compared with the X.25 model to make sure both had common features that enabled a comparison between both.

4.6 Summary

This chapter presented the ATM model. The modeled aspects of the protocol were described. The ATM data structures were explained, followed by a presentation of the main modules used in the model. Section 4.5 described how the model was verified and validated.

5 *Simulations*

5.1 Introduction

This chapter explains the results of simulations performed with the X.25 model and the ATM model. For the X.25 network model, two types of traffic models are examined. The first approximates the actual traffic that runs in the Brazilian Air Force data communications network. The second uses traffic sources that simulate an Ethernet LAN traffic. Section 5.2 describes the X.25 simulations and discusses their results. Section 5.3 explains how the ATM simulation was performed and explains its results. A summary of this chapter is presented in Section 5.4.

5.2 X.25 Simulations

The major difference between the X.25 simulations is the traffic source used. The first simulation uses the FTR block as the traffic source and the X.25 station module, both introduced in Chapter 3. This model represents the current X.25 network running in the Brazilian Air Force data communications infrastructure. The traffic characteristics of this model are explained below.

The traffic that runs in the X.25 data communications network is operational data that is carried out through remote host connections. It is made up of transactions with a mean size of 500 bytes, file transfers averaging 1.5 Mbytes and electronic mail (email) messages with a mean size of 512 bytes. The node with the highest traffic generates a daily traffic of about 6960 transactions, 60 file transfers, and 567 email messages. This totals about 3.5 Mbytes of transactions, 60 Mbytes of file transfer and about 300 Kbytes of email messages. These statistics are from the current X.25 network. Although the system provides three types of services, according to [Jai91], a metric should reflect the performance of services at the system level. This way the overall traffic composed of file transfers, transactions, and email, is modeled as file transfer requests in order to have a single metric. The sum of the daily traffic is approximately 94 Mbytes. If the average of an FTR is 512 bytes, approximately 2 FTRs per seconds are generated.

Four simulations were performed with this model. Each one had different values for the number of outgoing channels. Additionally, all simulations used the same design parameters, which are shown in Table 5.1. The plot in Figure 5.1 displays the resulting reject ratio obtained from the four simulations. The reject ratio is obtained by dividing the number of unsuccessful requests by the number of requests issued. The plot suggests that the number of outgoing channels is the limiting factor, since increasing them increases the request rate. As expected, as more traffic is input into the system the reject ratio increases. Assuming the system's quality of service is 1%, arrival rates of almost 10 FTRs per second could be processed for a 100 outgoing channels system. For this same quality of service, arrival rates of 20, 29, and 39 FTRs per second could be accommodated by 200, 300, and 400 outgoing channels systems respectively. The 300 and 400 outgoing channels curve presents more data points than the other two curves. This is because the simulations for these systems had to have the arrival rate parameter with more iterations in order to measure when the reject ratio becomes greater than 0%.

Table 5.1 - X.25 simulation parameters.

Arrival Rate	(0, 60, 15)
D bit	1
Window Size	2
LIC	1 - 1 - 1 - 1
HIC	900-800-2700-3600
LOC	901-801-2701-3601
HOC	1000-2000-3000-4000
Interarrival Mean	1 / 'Arrival Rate'
Mean Length	512
TSTOP	100
Global Seed	(67, 35, 21)

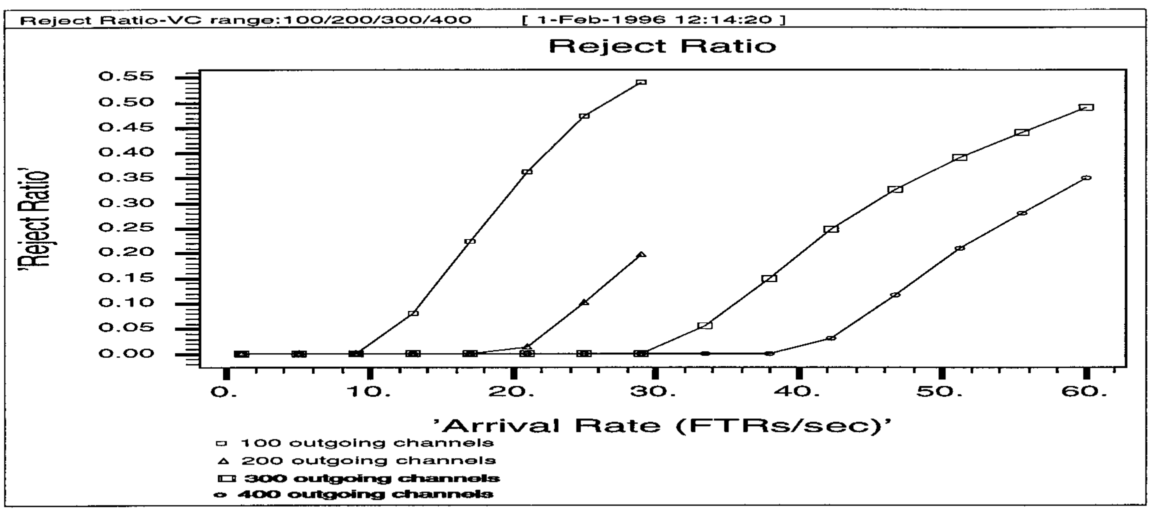

Figure 5.1 - X.25 Reject Ratio

The second simulation (X.25 - LAN Interconnection) uses the LAN traffic source which produces a much higher arrival rate than the one generated in the first simulation. Figure 5.2 shows how many packets (bursts) a LAN traffic source generates per second. For the lowest bit rate (2 Mbps), more than 350 packets/sec were generated. The simulation's parameters are shown in Table 5.2. The Boundary Time parameter controls the time the traffic source stops sending packets into the network. The Ethernet Mean Bit

Rate (per node) parameter was set to simulate a low, a medium, and a high mean bit rate, all below the maximum bit rate of 10 Mbps for an Ethernet. The Load parameter is set to 100% and will allow all the traffic generated by the ten traffic sources into the X.25 network. The number of outgoing channels is set to the highest value (400 channels) used in the first simulation in order to accommodate the increased traffic due to the LAN traffic generator. The simulations results are shown in Figure 5.3. A probe placed on the input port of the Switch by Status module in the system level diagram (Figure 3.9) collects all file transfer acknowledgments (including negative acknowledgments) to compute the reject ratio. The high arrival rates produced by the LAN traffic source increased the reject ratio to unacceptable values much higher than the 1% reject ratio assumed as good QoS parameter. The reject ratio reached values higher than 50% up to more than 90%. This simulation demonstrates the low quality of service offered by the X.25 network when connecting LANs.

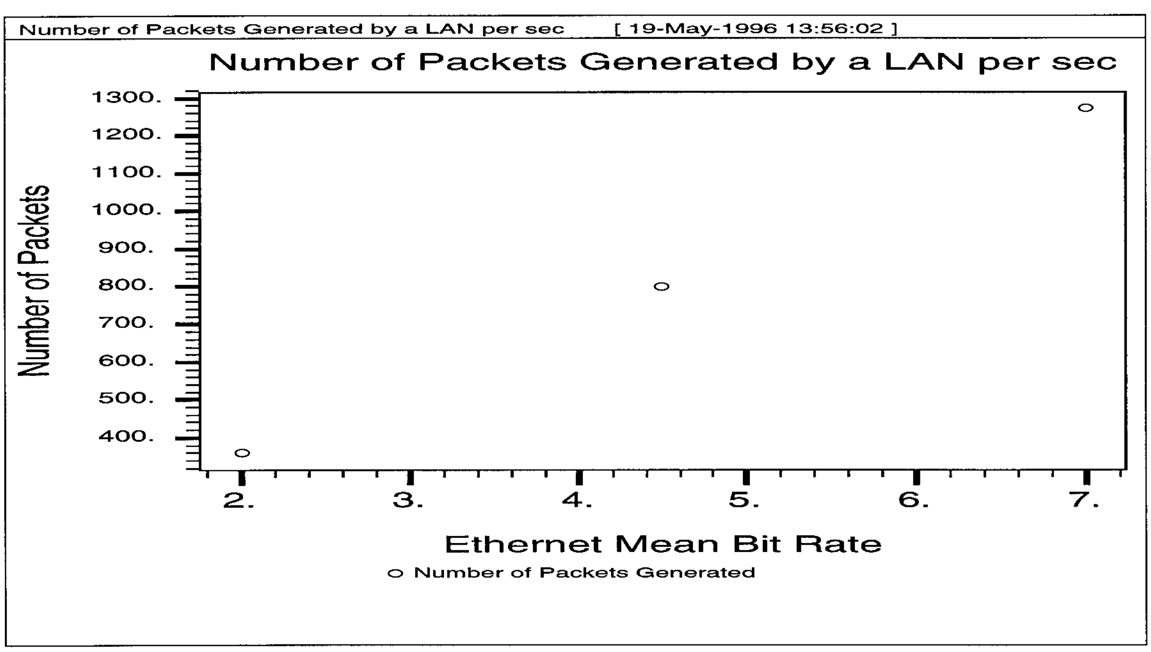

Figure 5.2 - Number of packets/sec generated by a LAN traffic source.

Table 5.2 - X.25 - LAN Interconnection simulation parameters.

D bit	1
Window Size	2
LIC	1
HIC	3600
LOC	3601
HOC	4000
Load	1
Boundary Time	2
Startup Time of Stat coll	0.1 * TSTOP
Ethernet Mean Bit Rate (per source)	(2000000, 3500000, 7000000)
Destination Search String(s)	'ether'
TSTOP	300
Global Seed	(71, 89, 13)

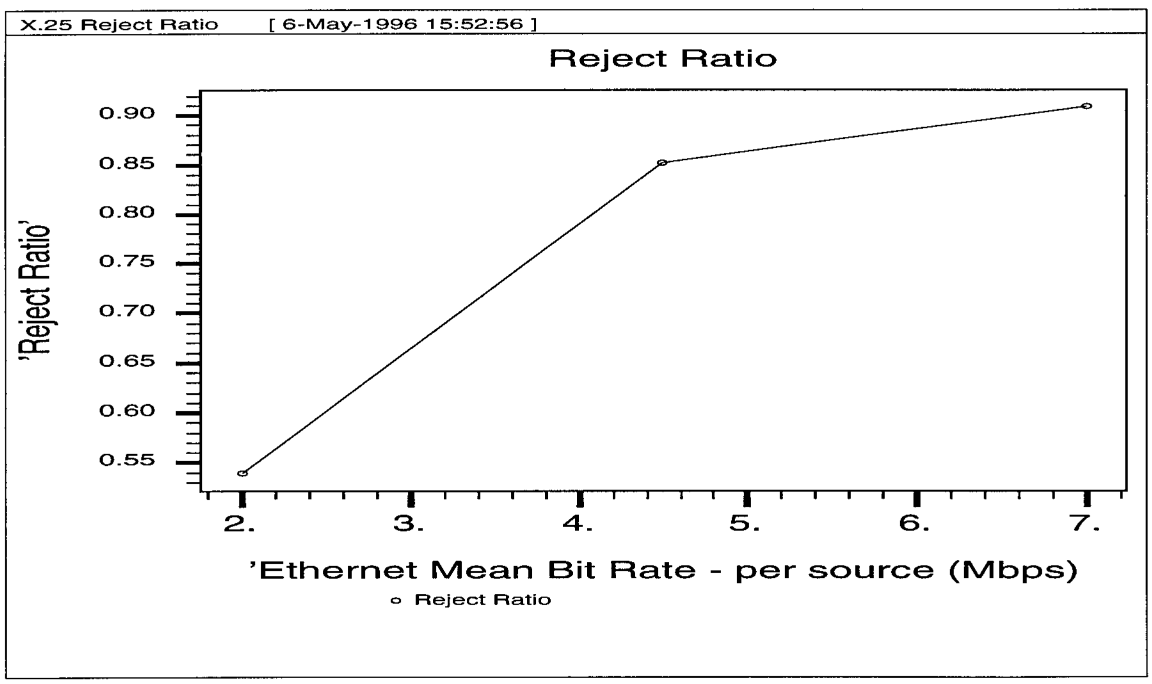

Figure 5.3 - X.25 - LAN Interconnection reject ratio.

5.3 ATM Simulations

The ATM model has the same topology as the X.25 model. Network devices specific to an ATM environment were added. Furthermore, the number of LAN traffic sources increased by a factor of three at each site, since ATM is to carry greater traffic loads than X.25.

Since the behavior of the ATM model was unknown, a simulation was performed iterating the Ethernet Mean Bit Rate parameter and setting the value of the Switch Queue Size to 8000 cells. This value for the queue size was estimated by analyzing the number of traffic sources from others ATM simulations [Cad94a]. A probe was placed in the input port of the Sink block in the system level diagram (Figure 4.10). This sink collects cells statistics from the traffic arriving on every IWU, and this data is used to compute the cell loss ratio. Figure 5.4 displays the results from this simulation. The plot shows that with an Ethernet Mean Bit Rate of 9 Mbps or greater, the cell loss ratio (CLR), which measures the ratio of lost cells and the total number of cells transmitted, is unacceptable, assuming an acceptable value of 0.001. This implies that when maintaining the same network topology, either the queue size needs to be adjusted for higher loads or the traffic has to stay less than 9 Mbps. The queue sizes are investigated below.

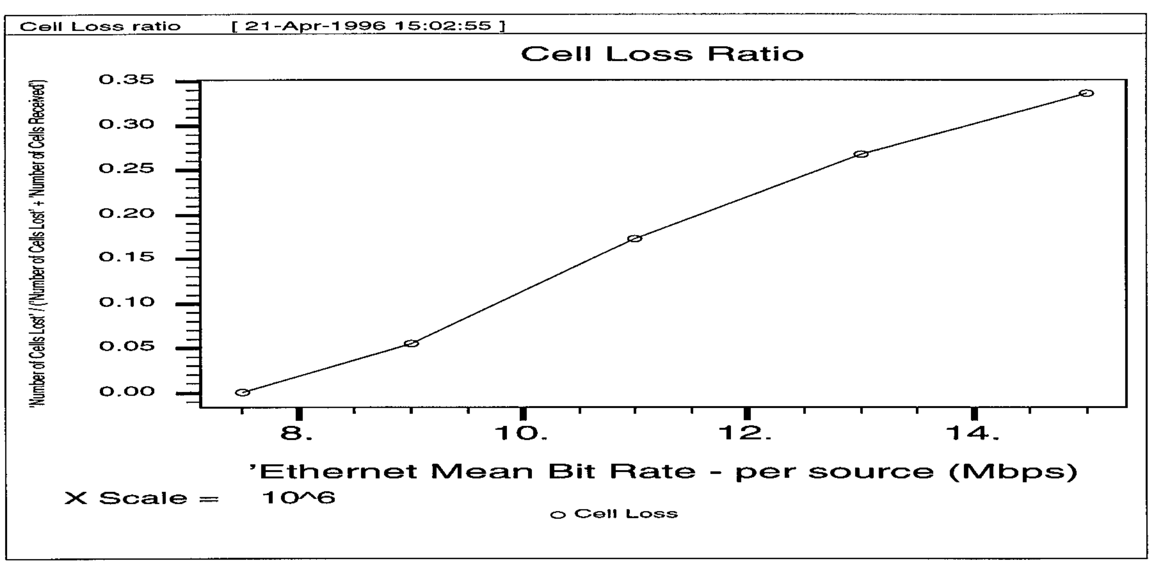

Figure 5.4 - Cell loss.

It is known that queues inside switches may be a bottleneck if not properly set. This occurs when the input traffic saturates a queue. When this happens, cells are discarded, and cell loss increases. A simulation iterating the Switch Queue Size parameter was performed to verify the system's behavior as the switch queue size changes. In this simulation, the Ethernet Mean Bit Rate was kept constant at 9 Mbps, and the switch queue size varied from 4,000 to 40,000 cells. If the expected mean bit rate from a LAN is 4.5 Mbps, then these results give an idea of how to size switch queues. It is important to notice that a traffic source in the ATM model represents two Ethernet LANs, making the Ethernet Mean Bit Rate (per node) for this case equal to 9 Mbps. The result of this simulation is displayed in Figure 5.5. The Y axis shows the cell loss ratio as the switch queue size increases. Queue sizes less than 40,000 cells have a cell loss ratio greater than 9%. A 40,000 cell queue size for the switch yields a cell loss ratio at a 5% level.

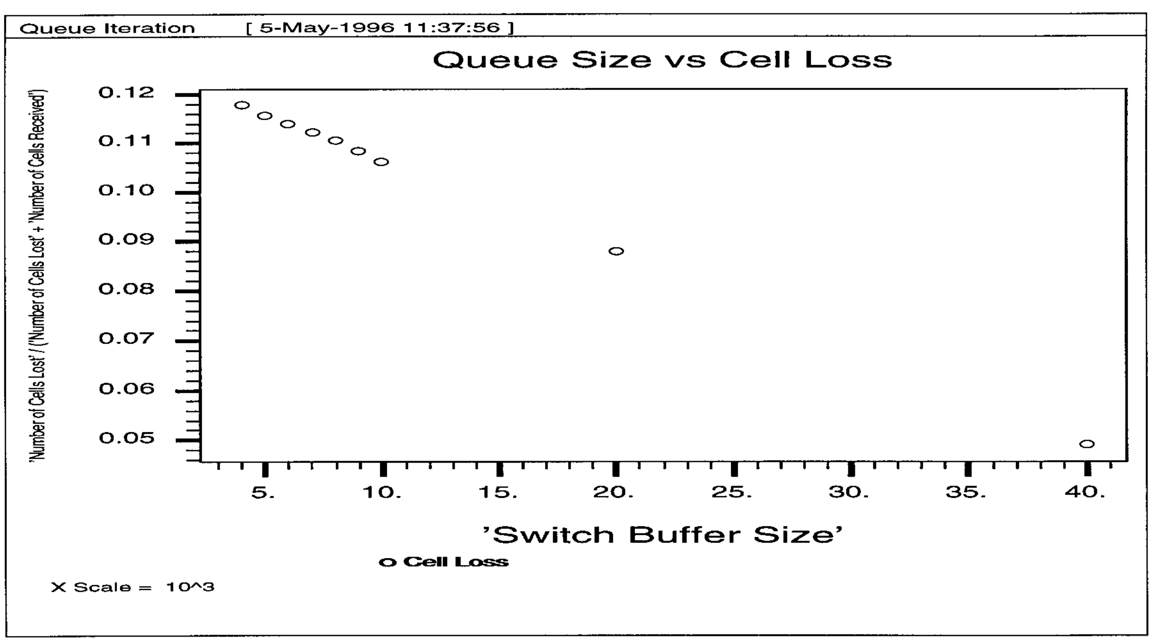

Figure 5.5 - Queue vs. Cell Loss plot.

The cell delay was also investigated to analyze its behavior as the queue size increases. Figure 5.6 displays the Queue Size versus Delay plot. As expected, the delay increases as does the queue size. This is due to the fact that as more cells get queued, the time spent in the queue increases, increasing the end-to-end delay. Cell transfer delay is affected by propagation delay, queuing delay, routing, and switching delays. Propagation and queuing delays are responsible for most of it, due to long distances (up to more than 4,000 Km) between end-to-end nodes and the network's topology. The delays plotted in Figure 5.5 turned out to be acceptable when compared with a permissible delay of 400 ms for voice traffic. The highest delay obtained in the simulation (69 ms) is almost six times less than what is considered acceptable for a delay sensitive traffic such as voice.

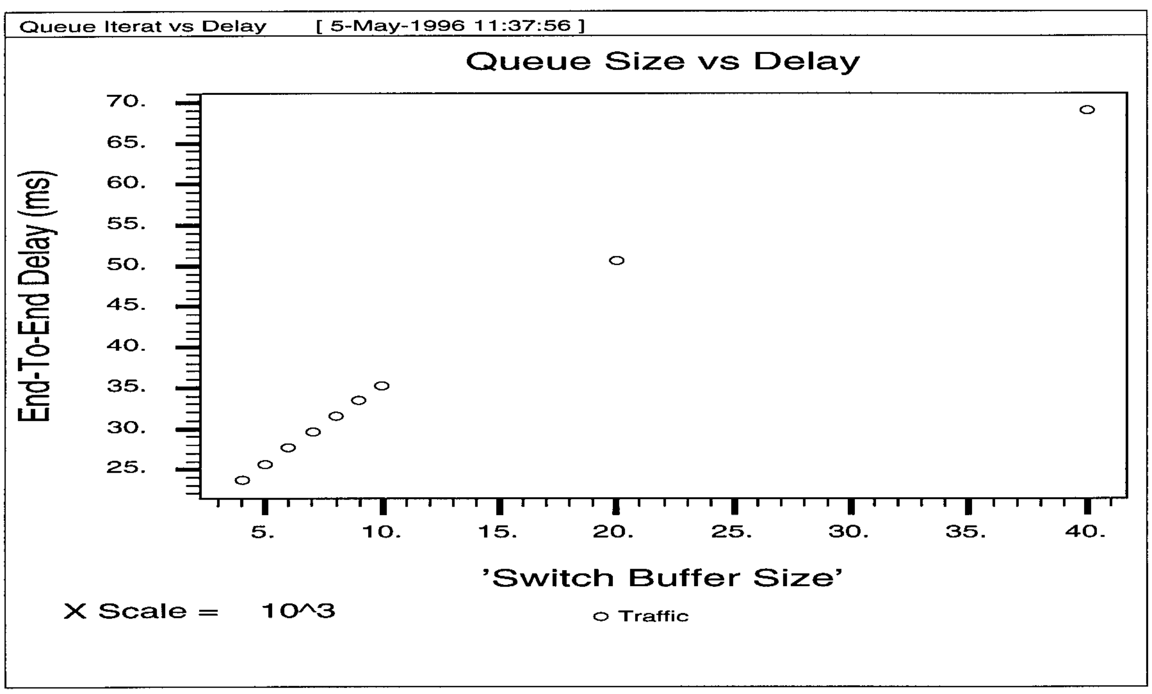

Figure 5.6 - Queue Size vs. Delay plot.

To keep the cell loss ratio in acceptable levels, the Switch Buffer Size must be increased, increasing the number of queued cells and therefore the queuing delay. An alternative to keeping the CLR at reasonable values would be to modify the physical layout of the network by introducing more cell switches and distributing the input LAN traffic through these switches. In this case the switches' buffer size could also be reduced.

Since increasing the Switch Queue Size not only increases the end-to-end delay but also decreases the cell loss ratio, a trade-off has to be considered. The choice made was to use the largest queue size (40,000 cells) considered for the following simulation. Although the end-to-end delay suffered a substantial increase when the Switch Buffer Size was increased by a factor of 8, it is still considered acceptable since a more sensitive delay traffic such as voice, permits delays of up to 400 ms while still maintaining an acceptable QoS.

The following discusses the performance of an ATM - LAN Interconnection Network. In B-ISDN, out-of-band signaling is performed. A user may have multiple signaling entities connected to the network call control management via virtual channel connections and the bit rates allocated to them can be chosen in a way that optimally satisfies the user's needs. The network parameter used to measure performance is the cell loss ratio, which is also a reliability parameter.

The first simulation performed had only one traffic source per site active. This was achieved by setting the Load parameter of the two other traffic sources to zero. The active traffic source on each site simulated a single LAN, and the Ethernet Mean Bit Rate and Load parameters were the same as those used in the X.25 - LAN Interconnection simulation. This was done to compare the behavior of both models when the traffic sources in both generate the same amount of traffic. Although different performance measures are used because of differences in standards, both measure the reliability of a network.

The plot in Figure 5.7 shows that the ATM network had no cell loss when exposed to the same traffic source used in the X.25 - LAN Interconnection model. As demonstrated in Figure 5.3, in the X.25 simulation the reject ratio was much higher than the acceptable 1% level commonly used. This means that the ATM model should be able to accommodate more traffic before it produces unacceptable values.

The second simulation had all the traffic sources enabled and simulating two LANs. The simulation parameters are shown in Table 5.3. The Ethernet Mean Bit rates represent the aggregate mean rate of traffic from the two LANs each source was modeling. The Switch Queue Size was set to 40,000 cells. The Load parameter was set to maintain the network aggregate mean rate of traffic into the network to a value that would keep the CLR at acceptable values below 0.001 when a medium traffic rate (4.5 Mbps) was generated. The network aggregate mean bit rate of traffic input into the network is obtained by Equation 1. This equation represents how much traffic enters the network. This calculation was performed with the medium mean bit rate (9 Mbps) used in the simulation and the result obtained was 135 Mbps.

Network Mean Aggregate Bit Rate = Ethernet Mean Bit Rate * Load * Number of LANs (1)

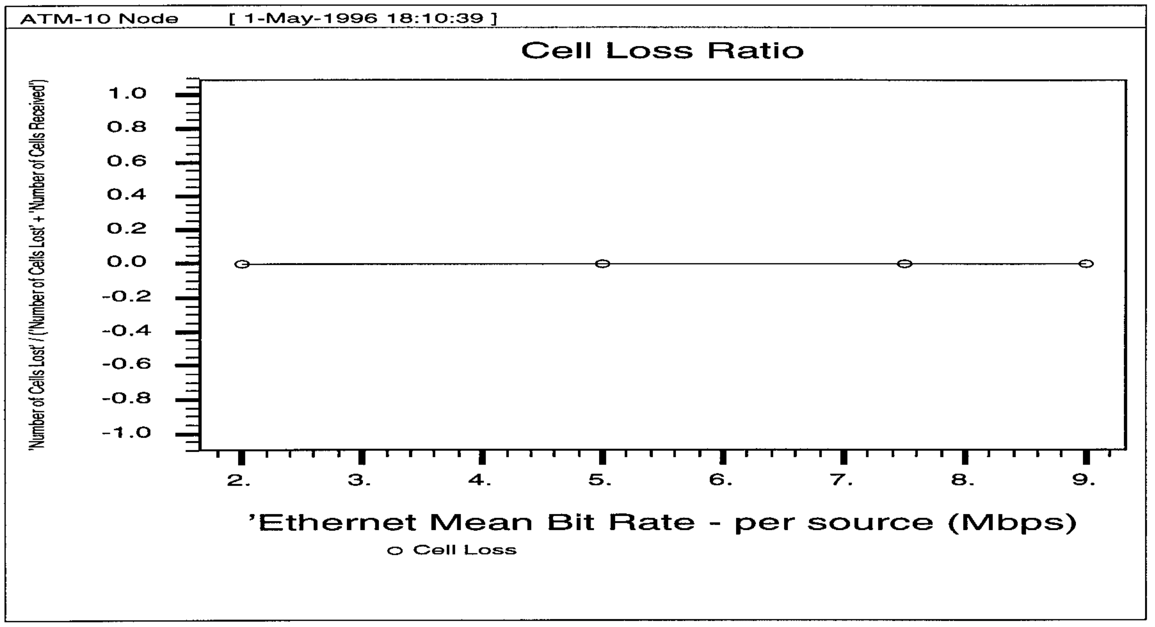

Figure 5.7 - Cell loss with same traffic as in the X.25-LAN Interconnection simulation.

Table 5.3 - ATM - LAN Interconnection simulation parameters.

Startup Time of Stat coll	0.1 * TSTOP
Ethernet Mean Bit Rate (per source)	(4000000, 9000000, 14000000)
Load	0.5
IWU Max Queue Size	500
UNI Link Capacity	STS-1
PBX Queue Size	2000
NNI Link Capacity	STS-1
Switch Queue Size	40000
Boundary Time	2
Number of Ethernets	2
Ethernet Peak Rate (Mbps)	1.25 * Ethernet Mean Bit Rate (per source)
TSTOP	2
Global Seed	(41, 17, 73)

Three plots were generated for this simulation. The plot in Figure 5.8 demonstrates the CLR obtained for the system as the Ethernet Mean Bit Rate increases. The cell loss ratio went above an acceptable value (0.1599) at the highest Ethernet Mean Bit Rate (14

Mbps). For the lower and medium bit rates the CLR remained at zero. If more cells had been delivered, then it is possible that some cell loss might have occurred.

The throughput for the UNI and NNI interfaces was also measured to see how the links were being used. The throughput was obtained by summing the total number of cells that entered the probe and multiplying it by 424 bits (number of bits in an ATM cell). Then, this number was divided by the TSTOP to obtain the average throughput. The plot in Figure 5.9 shows the throughput in the UNI. The maximum throughput measured was in the links between a site and a switch, when the maximum throughput reached 6.35, 14.8, and 23.1 Mbps for the 4, 9, and 14 Mbps data rates, respectively. The STS-1 link can still deliver more traffic than the LANs behind it can submit.

Figure 5.10 displays the throughput in the NNI links (between the Cell Switches). Because of the network topology, the traffic that comes from fifteen traffic sources have to pass through a single switch. This is why the NNI throughput reaches the peak of the STS-1 link (51.84 Mbps).

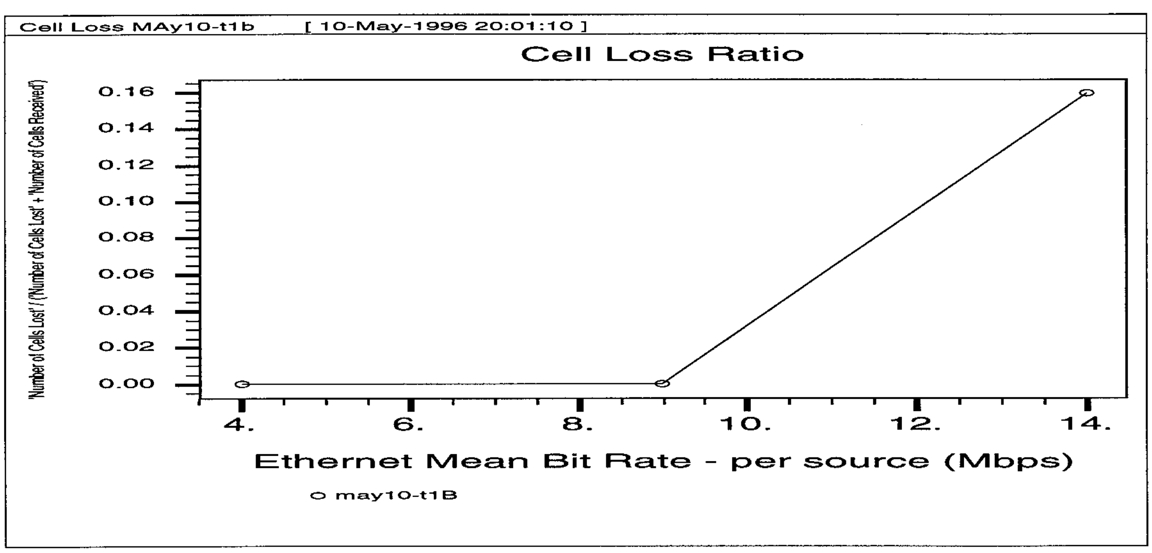

Figure 5.8 - Cell loss from ATM - LAN Interconnection model.

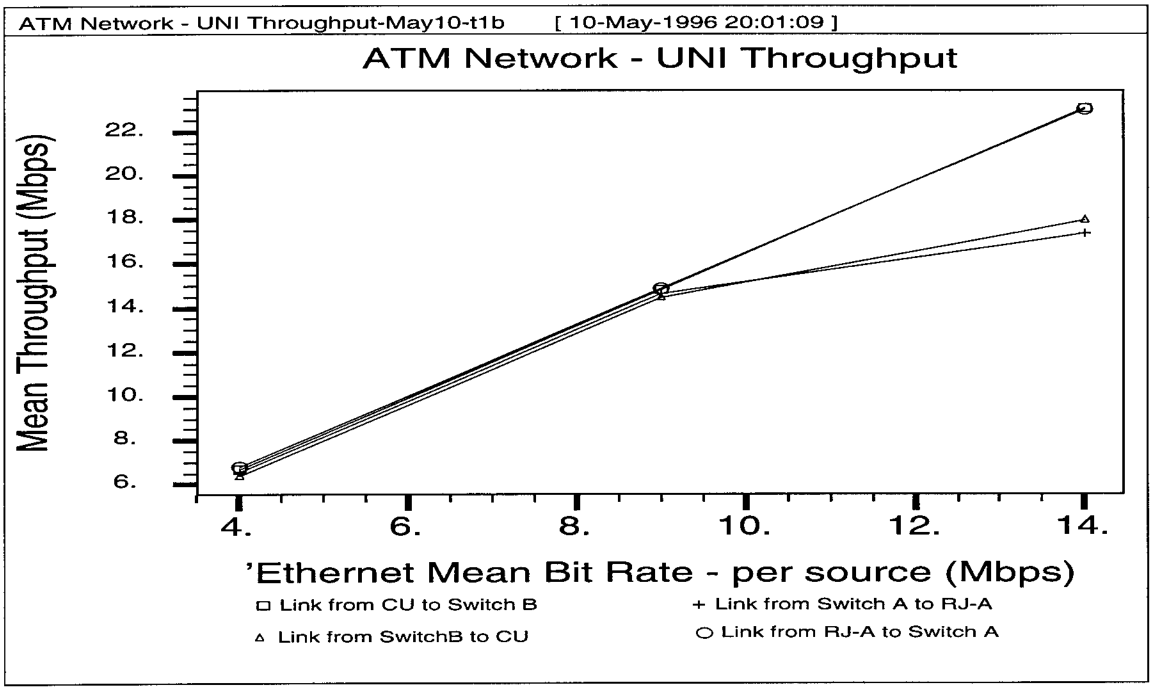

Figure 5.9 - UNI throughput.

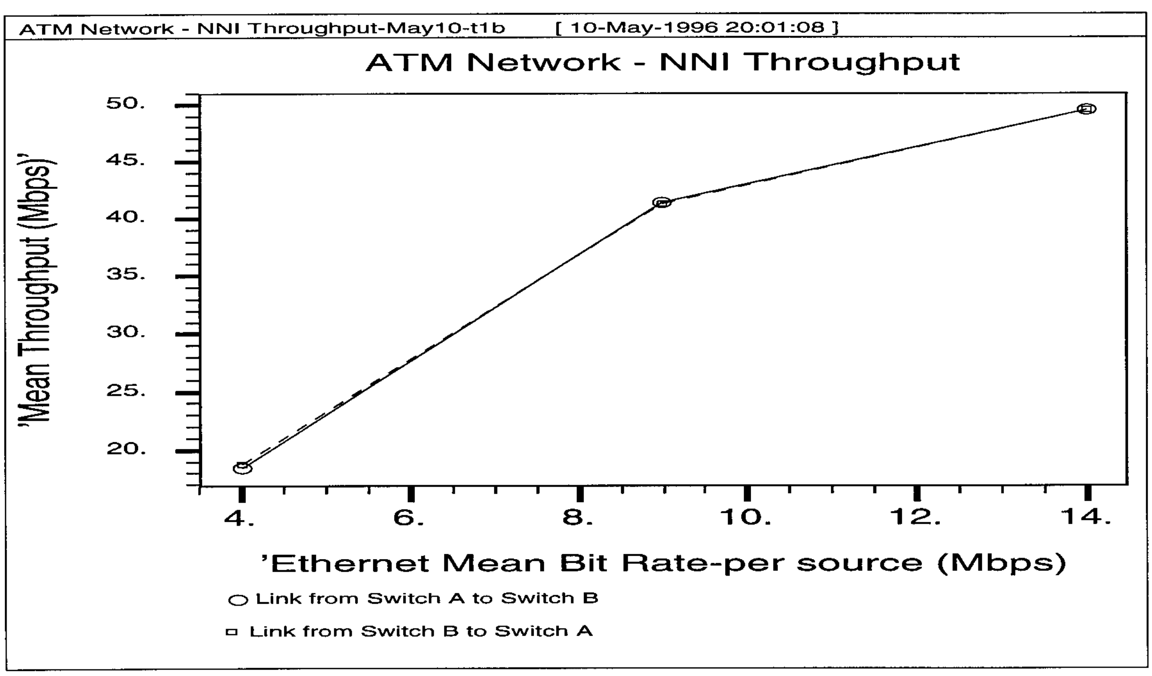

Figure 5.10 - NNI throughput.

Modifying the network topology by increasing the number of cell switches and distributing the traffic through the switches is one way of reducing the load on the NNI. An alternative is use links with higher capacity which are available to ATM.

5.4 Summary

This chapter presented the results of the X.25 and ATM simulations. The model for the Brazilian Air Force X.25 network was shown to have sufficient capacity to handle the projected workload. When the LAN traffic was inserted in the model the quality of service dropped to unacceptable levels. The ATM simulations proved to handle the traffic that took the X.25 model to unacceptable levels of QoS, and still maintain an acceptable value of QoS for ATM networks when the traffic was increased.

6 *Conclusions and Recommendations*

6.1 Introduction

The X.25 data communications network of the Brazilian Air Force was implemented in 1992. Until now, no study has been performed to analyze its performance. This study attempted to illustrate the behavior of the actual system and demonstrate a high-speed networking technology in this same environment. Comparing these two technologies was a way to recognize and emphasize the benefits of the implementation of a high-speed technology. Simulation was the performance evaluation technique used in analyzing these complex networks.

In Chapter 1, the objective of this study was presented, the scope was defined, and the approach and methodology used throughout this study was described. Chapter 2 provided the background information necessary to understand both the X.25 and the ATM technologies. Chapters 3 and 4 described the X.25 and ATM models, respectively, developed for this study. The simulations performed with both models were presented in Chapter 5 where their results were also discussed.

6.2 Conclusion

During the years that have passed since the implementation of the X.25 network, traffic characteristics have changed. Although the model for the current network proved to have sufficient capacity to handle the projected workload, it presented a low QoS when a more intense traffic load was introduced.

The number of LANs currently installed plus the ones currently being installed in the Brazilian Air Force, illustrate the rapid growth of data communications in this Force. More than 100 LANs are in the process of installation this year [Bar96]. The interconnection of these LANs will inevitably increase the load in the WAN environment. As demonstrated in Figure 5.2 and 5.3, the high arrival rates generated by the LANs produce reject ratios much higher than an acceptable level of 1%. The network aggregate mean bit rate in the X.25 - LAN Interconnection simulation, calculated using Equation 1 was 45 Mbps in a medium data rate. When the same traffic, as in the X.25 simulation, was introduced in the ATM network, the simulation results proved that such a traffic load is supported with no degradation of QoS. With an aggregate mean bit rate 3 times greater than the X.25 simulation, the ATM - LAN interconnection simulation presented a CLR lower than the acceptable QoS of 0.001. This is where the demand for a high-speed networking technology, such as ATM, increases. The deployment of higher speed LANs technologies such as 100BaseT, 100VG-AnyLAN and the already used FDDI will increase the aggregate traffic into the WAN, and intensify the demand for a higher speed WAN.

Another important point observed in this study is the adjustment of the buffer size inside the cell switches. In order to obtain an acceptable cell loss, the buffers must be set adequately. The delays suffered by the traffic also depend on the switch buffer size. Depending on the type of traffic the network is carrying, the delay is an important performance metric to take in consideration when adjusting buffer sizes. In the simulation performed, the end-to-end delays obtained turned out to be acceptable delays since they were lower than what is considered an acceptable value for a more delay sensitive traffic such as voice.

An important factor that should be observed in order to improve the QoS is the network topology. For the ATM network, the traffic generated by the LANs is distributed between two switches. If the physical layout of the network was modified by introducing more cell switches and distributing the input traffic through these switches, a better CLR could be obtained. This would also enable a reduction in the buffer sizes in the switches.

The results obtained in Chapter 5 indicate that an ATM network is an adequate network to support LAN interconnections. It also accommodates other services and its capacity can be increased to much higher levels. Introducing the ATM technology to the data communications network brings several benefits. As already mentioned, ATM was designed to bring a wide range of services under a single integrated network. One of these benefits is to unify the different services carried by the Air Force through several separated networks.

Today, there are several different services that rely on different networks. Some of them are listed below:

- Operational data communications network (X.25 data communication);
- Telex network;
- Leased lines voice network;
- Public voice network;
- Messaging systems through public telephone lines;
- Local Area Networks.

Among these networks are different technologies. Some transport redundant information, and all have different costs. Most networks are implemented with old technologies that are limited in data rates and to specific services, which does not mean that they do not provide an adequate service. With ATM, this wide range of services can be carried out under a single integrated network with increased performance and reliability. Services not available in the Brazilian Air Force, such as lifelike video conferencing, distance learning can also be enabled.

6.3 Recommendations

This study has provided a comparison of two different networking technologies under the most similar environment as possible. Due to the diversity, and time restraints of this study, certain topics could not be implemented and/or investigated. These topics, listed below, can be investigated in future research.

1. Modify the network topology by implementing a greater number of cell switches in the ATM model. By doing so, the CLR of the network may change and make this network more reliable, and able to accept a higher traffic.

2. Implement the use of satellite (wireless) communications. This could reduce the cost involved with lengthy terrestrial links.

3. Analysis of the cost of implementing an ATM network with such dimensions and characteristics.

As a final recommendation, it is advisable to pursue a strategy when implementing a new technology. As in the implementation of any new technology, risks exist. Until this date, seamless interoperability between ATM products was not yet obtained [Con95]. The adoption of a proprietary approach may constraint future interoperability and increase costs. Standard-based solutions must be chosen. These considerations, when coupled with the existing lack of knowledge on the ATM technology, mandates that the Brazilian Air Force pursue a cautious approach to the introduction of ATM technology.

A well established methodology and full understanding of the technology will make the implementation of ATM a success, bringing benefits and helping the Air Force with its needs.

Bibliography

[ATM93] ATM Forum, *"ATM User-Network Interface Specification,"* Version 3.0, Prentice-Hall, 1993.

[Bar96] Barreto, Luis Fernando. Engineer in the Brazilian Air Force, DIRINFE, Brazil. Personal Correspondence. 01 November 1995.

[BaW91] J.J.Barret and E.F.Wunderlich, *"LAN Interconnection Using X.25 Network Services,"* IEEE Network Magazine, September 1991, pp.12-16.

[Boe88] R. J. Boehm, *"The SONET Interface and Network Applications"*, IEEE GLOBECOMM, 1988.

[BoF95] F. Bonomi, and K. W. Fendick, *"The Rate-Based Flow Control Framework for the Available Bit Rate ATM Service,"* IEEE Network, March 1995, pp.25-39.

[Bou92] J. Y. Le Boudec, *"The Asynchronous Transfer Mode: a tutorial,"* Computer Networks and ISDN Systems,Vol.24, 1992, pp.279-308.

[Bru94] W. A. Bruwer et al., "*VISTAnet and MICA: Medical Applications Leading to the NCIH*," IEEE Network magazine, November/December 1994.

[CAD94a] Cadence Design Systems Inc. *BONeS Designer ATM Library Reference.* Version 1.0, December 1994.

[Cad94b] Cadence Design Systems Inc. *BONeS Designer Modeling Guide.* December 1994.

[CCI84] CCITT Data Communication Networks: Open System Interconnection (OSI), system description techniques. Recommendations X.200 - X.500, 1984.

[ChH92] K. Chipman, P. Holzworth, and L. Loop, *"Medical applications in a B-ISDN field trial,"* IEEE Journal in Selected Areas on Communication, Vol.10, September 1992, pp.1173-1187

[Con95] ATM Consortium (http://www.iol.unh.edu/general/IOL-General_Newsletter/Spring95-Homepage.html).

[Dav94] R. P. Davidson, Broadband Networking, New York: Wiley-QED, 1994.

[End93] N. Endo, *"A Shared Buffer Memory Switch for an ATM Exchange,"* IEEE Transactions of Communications. Vol.41, No.1, January 1993.

[FlR94] S. Floyd and A. Romanow, *"Dynamics of TCP traffic over ATM networks,"* Proceedings ACM SIGCOM'94, September 1994, pp.79-88.

[HaH93] R. Handel, M.N. Huber ans S. Schroder, ATM Networks-Concepts, Protocols, Applications. 2nd Edition. New York: Addison-Wesley, 1994.

[Jai91] R. Jain, The Art of Computer Systems Performance Analysis. New York: John Wiley & Sons, Inc., 1991.

[Jom89] P. Jomer, *"Connection Caching of Traffic Adaptive Dynamic Virtual Circuits,"* SIGCOMM 89, ACM, New York, 1989, pp.13-24.

[Nei91] L. A. Neir, *"Time to Delivery Queuing for Resource Allocation in Congestion Control In frame Relay Networks,"* Master's Thesis, University of Kansas, 1991.

[New94a] P. Newman, *"ATM Local Area Network,"* IEEE Communications Magazine, March 1994.

[New94b] P. Newman, *"Traffic Management of ATM Local Area Networks,"* IEEE Communications Magazine, Vol.32, No.8, August 1994, pp.44-50

[Pan89] R. Pant, *"X.25 Broadcast Service,"* IEEE Network Magazine, July 1989, pp.20-26.

[PrS87] M. De Prycker and M. De Somer, "Performance of a Service Independent Switching Network with Distributed Control," IEEE Journal on Selected Areas in Communications, vol.5, No.8, October 1987, pp.1293-1302.

[Ran95] M. N. Ransom, *"North Carolina Information Highway: Megabits Driving Gigabits,"* IEEE Journal on Selected Areas in Communications, Vol.13, No.5, June 1995, pp.788-792.

[ReF90] D. Reznik and V. S. Frost, *"An X.25 Network Model using BONeS,"* Telecommunication and Information Sciences Laboratory, University of Kansas Center for Research, 1990.

[Ryb80] A. Rybczynski, "*X.25 Interface and End-to-End Virtual Circuit Service Characteristics*," IEEE Transactions on Communications, Vol.Com-28, No.4, April1980, pp.500-510.

[SaA94] Tarek N. Saadawi and Mostafa H. Ammar with Ahmed El Hakeem, Fundamentals of Telecommunication Networks,Wiley-Interscience, 1994.

[Sta96] L. Staalhagen,"*A Comparison Between the OSI Reference Model and the B-ISDN Protocol Reference Model,*" IEEE Network, January/February 1996, pp.24-33.

Vita

Oswaldo Zanelli was born on August 26, 1966 in Rio de Janeiro, Brazil. He is a First Lieutenant in the Brazilian Air Force. He graduate from the Brazilian Air Force Academy on 1987 and attended the Catholic University of Rio de Janeiro (PUC-RJ), where he received a Bachelor of Systems Analysis in 1990. He worked at the Air Force Computer Science and Statistics Department (DIRINFE) since 1990. Then, in 1994, he was assigned to study at the Air Force Institute of Technology (AFIT), where he received the degree of Master of Science in Computer Systems in June of 1996.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE June 1996

3. REPORT TYPE AND DATES COVERED Master's Thesis

4. TITLE AND SUBTITLE

PERFORMANCE STUDY OF A BRAZILIAN AIR FORCE DATA COMMUNICATIONS NETWORK

5. FUNDING NUMBERS

6. AUTHOR(S)

Oswaldo Zanelli, 1st Lieutenant, Brazilian Air Force

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Air Force Institute of Technology,
WPAFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT/GCS/ENG/96J-02

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

This study consists of a performance analysis of two network models using different protocols, the X.25 and the asysnchronous transfer mode (ATM). The models were built using BONeS Designer, a modeling/simulation tool. Simulation of both models were performed to compare the two technologies.

14. SUBJECT TERMS

X.25, asynchronous transfer mode, ATM, modeling, simulation, network

15. NUMBER OF PAGES 89

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL