

NAVAL POSTGRADUATE SCHOOL MONTEREY, CALIFORNIA



THESIS

ANALYSIS AND DESIGN OF A UNIVERSAL TRAFFIC NETWORK

by

John B. Hunter
Holger Großmann

September 2000

Thesis Advisor:
Second Reader:

Gilbert M. Lundy
Wolfgang Baer

Approved for public release; distribution is unlimited

THIS QUANTITY EMPLOYED 4

20001117 032

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 2000	3. REPORT TYPE AND DATES COVERED Master's Thesis	
TITLE AND SUBTITLE: Analysis and Design of a Universal Traffic Network			5. FUNDING NUMBERS	
6. AUTHOR(S) Hunter, John B. and Großmann, Holger				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) As the field of computer networking has evolved, so too has the use of these networks. Modern networks must be capable of performing more than simple data transfer. To be of value, a network must be able to handle the convergence of different types of traffic – voice, video, and data – and the Quality of Service requirements associated with each type. This thesis performs a detailed analysis of the different types of traffic, the two primary transmission media, fiber optical and copper based connections, and the connection-orientation technology to route the traffic. Presented in this thesis is a fiber-based hybrid network consisting of Asynchronous Transfer Mode at the backbone layer and Frame Relay and Passive Optical Networking at the local access layer. The proposed Universal Traffic Network, based on present-day technology, is a viable solution to the challenge imposed by the convergence of different traffic types.				
14. SUBJECT TERMS Computer Networks, Asynchronous Transfer Mode, Frame Relay, SONET, Fiber Optics, Passive Optical Networks, Gigabit Ethernet, Multimedia Communication, Transmission Media			15. NUMBER OF PAGES 138	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

ANALYSIS AND DESIGN OF A UNIVERSAL TRAFFIC NETWORK

John B. Hunter
Lieutenant, United States Navy
B.A., Virginia Polytechnic Institute & State University, 1993

Holger Großmann
Lieutenant Commander, German Navy
M.S., University of the German Armed Forces Munich, 1990


Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE


from the

NAVAL POSTGRADUATE SCHOOL
September 2000

Authors:




John B. Hunter



Holger Großmann


Approved by:



Gilbert M. Lundy, Thesis Advisor



Wolfgang Baer, Second Reader



Dan Boger, Chair, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

As the field of computer networking has evolved, so too has the use of these networks. Modern networks must be capable of performing more than simple data transfer. To be of value, a network must be able to handle the convergence of different types of traffic - voice, video, and data - and the Quality of Service requirements associated with each type.

This thesis performs a detailed analysis of the different types of traffic, the two primary transmission media, fiber optical and copper based connections, and the connection-orientation technology to route the traffic. Presented in this thesis is a fiber-based hybrid network consisting of Asynchronous Transfer Mode at the backbone layer and Frame Relay and Passive Optical Networking at the local access layer. The proposed Universal Traffic Network, based on present-day technology, is a viable solution to the challenge imposed by the convergence of different traffic types.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
	A. ENVIRONMENT.....	1
	B. MOTIVATION.....	2
	C. FOCUS AND EXCLUSIONS.....	3
	D. THESIS ORGANIZATION.....	5
II.	NETWORK TRAFFIC REQUIREMENTS.....	7
	A. VOICE (TELEPHONY).....	8
	B. VIDEO.....	13
	C. DATA.....	17
III.	TECHNOLOGY ANALYSIS.....	21
	A. COPPER VERSUS FIBER.....	22
	1. Segment Length.....	22
	2. Cost.....	25
	3. Bandwidth.....	30
	B. FIBER FUNDAMENTALS.....	32
	1. Multi-mode & Single-mode Fiber.....	34
	2. Transmission Impairments.....	37
	C. SONET.....	42
	1. Supported Speeds.....	44
	2. SONET Multiplexing.....	45
	3. SONET Frame Format.....	46
	D. CONNECTION ORIENTATION.....	50
	1. ATM - Asynchronous Transfer Mode.....	53
	2. FR - Frame Relay.....	61
IV.	NETWORK DESIGN & SPECIFICATION.....	69
	A. ATM BACKBONE.....	71
	B. FEEDER PORTION.....	77
	1. FUNI - Frame-Based User-to-Network Interface.....	78
	2. PON - Passive Optical Network.....	82
	C. LOCAL ACCESS.....	85
V.	ALTERNATIVE TECHNOLOGIES.....	93
	A. IP OVER FIBER.....	93
	B. PHOTONIC ROUTING.....	98
	C. TERABIT LINKS.....	102
	D. GIGABIT ETHERNET.....	104
VI.	CONCLUSION.....	111
	LIST OF REFERENCES.....	113
	INITIAL DISTRIBUTION LIST.....	117

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 2.1 Audio and Video Codecs Under H.323.....	10
Figure 2.2 Speech Compression Algorithms.....	11
Figure 2.3 Video Compression Algorithms.....	15
Figure 2.4 QoS Requirements for the UTN.....	20
Figure 3.1 Office Building Network Design.....	29
Figure 3.2 Typical Optical Fiber.....	34
Figure 3.3 Multi-Mode Graded-Index Optical Fiber.....	36
Figure 3.4 Single-Mode Optical Fiber.....	37
Figure 3.5 Fiber Attenuation and Wavelength.....	38
Figure 3.6 Dispersion for Different Fibers.....	41
Figure 3.7 Supported SONET Speeds and SDH Equivalents.....	45
Figure 3.8 Basic SONET STS-1 Frame Format.....	47
Figure 3.9 Frame Overlap.....	48
Figure 3.10 Traditional (non-ATM) Statistical Multiplexer.....	55
Figure 3.11 ATM Statistical Multiplexer.....	56
Figure 3.12 ATM UNI cell format in bits.....	56
Figure 3.13 Virtual Path & Circuits.....	57
Figure 3.14 Frame Relay Frame Format.....	63
Figure 3.15 Use of Subframes in Frame Relay.....	65
Figure 4.1 Conceptual Design of the UTN.....	69
Figure 4.2 UTN Level Two - ATM Backbone.....	72
Figure 4.3 ATM Point-to-Multipoint Connection.....	76
Figure 4.4 UTN Level One - Feeder Portion.....	78
Figure 4.5 FUNI Reference Model.....	79
Figure 4.6 AAL5 Frame carrying a Frame Relay Frame.....	80
Figure 4.7 Multiplexing FR Circuits across an ATM Circuit.....	81
Figure 4.8 Functionality of Passive Optical Networking.....	83
Figure 4.9 UTN Level Zero - Local Access.....	85
Figure 4.10 Node Model of a Frame Relay Workstation.....	87
Figure 4.11 Intelligent Distribution of a Video Signal.....	89
Figure 5.1 IP to Fiber Layering Methods.....	95
Figure 5.2 Optical Routing Mechanism.....	100
Figure 5.3 802.3 Frame Format.....	104
Figure 5.4 Gigabit Ethernet Maximum Distance Supported.....	105

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS

AAL	ATM Adaptation Layer
ADX1	ATM Data Exchange Interface
ANSI	American National Standards Institute
ATM	Asynchronous Transfer Mode
AWG	Arrayed Waveguide Grating
BECN	Backward Explicit Congestion Notification
B-ISDN	Broadband-Integrated Service Digital Network
C/R	Command/Response
CBT	Core Based Tree
CIR	Committed Information Rate
CPC	Common Part Conversion
CRC	Cyclic Redundancy Check
CLP	Cell Loss Priority
Codec	Coder/Decoder
dB	decibels
DE	Discard Eligibility
DLCI	Data Link Connection Identifier
DVMRP	Distance Vector Multicast Routing Protocol
DWDM	Dense Wave Division Multiplexing
EA	Extended Address
EDFA	Erbium-Doped Fiber Amplifier
EMI	Electromagnetic Interference
FCS	Frame Check Sequence
FDM	Frequency Division Multiplexing
FECN	Forward Explicit Congestion Notification
FR	Frame Relay
FRAD	Frame Relay Access Device
FRIPIF	Frame Relay to IP Interface
FTP	File Transfer Protocol
FUNI	Frame Based User-to-Network Interface
GBE	Gigabit Ethernet
Gbps	Gigabits per second

GFC	Generic Flow Control
HDLCL	High-level Data Link Control
HEC	Header Error Control
HO	Head Office
HTTP	Hyper Text Transfer Protocol
IEX	Internet Exchange
IGMP	Internet Group Management Protocol
IP	Internet Protocol
ITU	International Telecommunications Union
KB	Kilobyte
Kbps	Kilobits per second
LAN	Local Area Network
LOH	Line Overhead
Mbps	Megabits per second
MBps	Megabytes per second
MHz-km	Megahertz per kilometer
M-JPEG	Motion-Joint Picture Experts Group
MOSPF	Multicast Open Shortest Path First
MPEG	Moving Picture Experts Group
ms	milliseconds
nm	nanometer
NNI	Network-to-Network Interface
NTP	Network Time Protocol
OADM	Optical Add Drop Multiplexer
OC	Optical Carrier
OLT	Optical Line Terminal
ONT	Optical Network Unit
OSI	Open Systems Interconnection
PBX	Private Branch Exchange
PCM	Pulse Code Modulation
PDU	Protocol Data Unit
PIM	Protocol Independent Multicast
POH	Path Overhead

PON	Passive Optical Networks
POS	Packet Over SONET
PPP	Point-to-Point Protocol
PSTN	Public Switched Telephone Network
PVC	Permanent Virtual Circuit
QoS	Quality of Service
RSVP	Resource Reservation Protocol
SAR	Segmentation And Reassembly
SDH	Synchronous Digital Hierarchy
SDU	Service Data Unit
SMDS	Switched Multimegabit Data Service
SOH	Section Overhead
SONET	Synchronous Optical Network
SPE	Synchronous Payload Envelope
SSCS	Service Specific Convergence Sublayers
STM	Synchronous Transport Module
STS	Synchronous Transport Signal
SVC	Switched Virtual Circuit
TCP	Transmission Control Protocol
TDM	Time Division Multiplexing
TDMA	Time Division Multiple Access
TOH	Total Overhead
UDP	User Datagram Protocol
UNI	User-Network Interface
UTN	Universal Traffic Network
VC	Virtual Circuit
VCC	Video Channel Connection
VCI	Virtual Channel Indicator
VDVCC	Video Virtual Channel Connection
VoIP	Voice-over-Internet Protocol
VP	Virtual Path
VPI	Virtual Path Indicator
VT	Virtual Tributary

VTG	Virtual Tributary Group
WAN	Wide Area Network
WDM	Wave Division Multiplexing

ACKNOWLEDGMENTS

The authors would like to express their sincere gratitude to Professors Bert Lundy and Wolfgang Baer whose guidance, wisdom and enthusiasm made this work possible. In addition, they would like to thank their families and friends for their support. Specifically, Jutta and Katja for putting up with the late nights and long hours. To Schatzie and Ki, for helping to remember the more important things in life. Finally, to our parents, for instilling in us the desire and wherewithal to achieve our goals.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. ENVIRONMENT

Without question, computer networks are fast becoming an integral part to the everyday lives of people throughout the world. What began over 20 years ago with experimental systems like the ARPANET - connecting mainframe computers over long-distance telephone lines - has resulted into a booming business. News of businesses utilizing computer networks to perform functions, which until recently had to be completed manually, abound in the media. Even in the home there are computer networks that resemble those of a small corporate office. From the commercial world to the private sector, this array of networks caters to delivering different forms of traffic, such as voice, video and data. The next evolutionary step in these networks is the convergence of the various types of traffic on the same network. For example, simultaneously making a telephone call, watching television, and surfing the Internet, all from one personal computer. However, this convergence of traffic requires a new type of network model. Present day models - whether designed to carry voice or small amounts of data - simply are not adequate to handle the merging of information.

B. MOTIVATION

In order to achieve the ultimate goal of bringing a converged computer network into every household and business, a careful analysis needs to be conducted. The end result of the analysis will be a Universal Traffic Network (UTN). The analysis will not be too intricate, however a number of technical aspects will be evaluated.

To begin, a breakdown of the network requirements will be addressed. Specifically, what types of traffic must the UTN be capable of handling and what are the associated bandwidth requirements for each? Additionally, understanding the various Quality of Services (QoS) necessary to supporting each type of traffic is significant. After having defined the requirements, the authors proceed to the next logical phase in the design of the UTN.

The second step considered in developing the UTN is the technology necessary to satisfy the requirements. In particular, which types of physical transmission mediums (i.e., copper or fiber) and connection orientations (e.g., circuit-switched, packet-switched, cell-switched, or combinations of each) should the UTN employ? Through a sensible analysis, a sound argument is made for which technology best suits the UTN model.

The final step in the analysis is a hypothetical model of the UTN. The model will describe where the various technologies are to be implemented. Specifically, it focuses on the local backbone of the network, access points to the local backbone, and the interface to end devices or nodes in the network.

C. FOCUS AND EXCLUSIONS

First, wherever appropriate, a concerted effort has been made to detail and explain the various technologies involved in the design of the UTN. However, the authors assume the reader has a basic understanding of certain networking principles such as, elementary electromagnetic fundamentals, knowledge of network components (e.g., routers, switches, hubs), circuit switching versus packet switching, the Open Systems Interconnection (OSI) Reference Model, etc.

Specifically, this thesis intends to concisely describe a computer network capable of handling the merger of the three major forms of traffic - voice, video and data - in a single network. The analysis is based on existing or present-day technologies - both physical connectivity and connection-orientation. However, there are limitations to the analysis.

The computer network industry is a huge conglomerate made up of various corporations. Consequently, there are network devices that perform the same function, however each component built by different vendors. For example, there are over fifty companies that specialize in producing Frame Relay Access Devices (FRADs). While some vendors offer better products than others, the authors do not attempt to recommend one specific brand of network hardware over another. For this reason, the thesis is limited in terms of the explicit hardware devices utilized in the design of the UTN. Instead, the focus is on the principles behind the technology vice the actual device itself.

A final limitation to this thesis is actually simulating the UTN. While it is true a number of powerful network simulation and analysis tools are available, the results they produce are only as accurate as the input. To draw accurate conclusions from using a network analysis tool, one must have provided exact input parameters (e.g., buffer sizes, cycles per second, etc.) As alluded to in the previous paragraph, this thesis does not specify the actual hardware devices, thus simulation based on assumptions is not valuable. Although a detailed simulation is not conducted, a limited analysis with figures depicting the network implementation is provided.

In summary, designing and implementing a network of any size requires a great deal of time, money and research. The authors of this thesis understand that developing a network on the scale of the UTN is a colossal endeavor. As a result, an attempt is not made to address every aspect of the UTN, as it is not within the confines of this thesis. However, the authors are convinced this thesis does provide a solid foundation to implementing the UTN.

D. THESIS ORGANIZATION

Including this introductory chapter, this thesis is organized in the following chapters:

Chapter II: Network Traffic Requirements Analysis. This chapter describes the different types of traffic that will be transmitted across the UTN. The types of traffic are categorized into three separate sections: voice, video and data. Each section provides a cursory overview of the specific traffic type and concludes with the requirements the UTN must be capable of supporting.

Chapter III: Technology and Services Analysis. In this chapter, the various types of technology and services are identified for the UTN in order to satisfy the requirements defined in Chapter II. The first half of the chapter is devoted to arguing in favor of fiber over copper as the

transmission medium for the UTN and continues with a perfunctory synopsis of the physical fundamentals of fiber optic technology. The second half of the chapter delineates the types of connection-orientation the UTN will employ, Asynchronous Transfer Mode (ATM) and Frame Relay (FR).

Chapter IV: Network Implementation. Described in this chapter is a high-level analysis of the framework behind the UTN. Three levels of the network are explored, the local backbone, the feeder portion with access points, and the distribution level. Various schematics are presented utilizing a commercial network performance analysis tool, OPNET Modeler.

Chapter V: Alternative Technologies. The UTN is not the sole solution to converging networks. There are areas of ongoing research, which could be implemented in a future UTN. Presented in this chapter are four key areas; IP over fiber, Photonic Routing, Terabit Links, and Gigabit Ethernet.

Chapter VI: Conclusion. An overall assessment of the UTN and possible drawbacks to the design are presented in the final chapter. Based on present-day technology, the UTN is a practical solution to the problem of converging networks.

II. NETWORK TRAFFIC REQUIREMENTS

Before entering into a discussion of the UTN requirements, the authors wish to make clear that the underlying protocol for each type of traffic is the Internet Protocol (IP). It is assumed the reader has a basic understanding of IP and therefore it will not be discussed in-depth. It is important to note because in addition to the various types of traffic, each of the mediums, services and connection-orientation outlined in this thesis are capable of handling IP. This means that any type of traffic encapsulated in IP packets will be capable of being transported across the network.

The first logical step in designing any network is determining what types of requirements or services the network will support. The analysis need not be too complex, however it must encompass all the services. Otherwise, the design will be flawed or not feasible at all. Furthermore, a requirements analysis sets the stage for formulating what types of mediums, protocols and connect-orientation are required to successfully support the services. It is for this reason, the analysis begins by outlining the types of requirements the UTN must be capable of supporting. The authors envision a network capable of supporting three

fundamental types of services, Voice (Telephony), Video, and Data.

A. VOICE (TELEPHONY)

Traditionally, a typical telephone call is made by accessing the Public Switched Telephone Network (PSTN) via an analog signal on a pair of copper wires connected to the home or business. However, with the advent of packet-switching technology, the digitization of data, and alternative transmission mediums (e.g., wireless, cable and fiber optics), a new paradigm is rapidly replacing the old model. Whereas before, there were separate voice and data networks - each with their own techniques - the two are now converging. Initially, combining voice and data required data to look like voice and be transmitted via circuit-switched technology. However, as the amount of data traffic surpassed voice, the convergence took a 180-degree turn. Nowadays, voice is capable of resembling data and the two can reside in a packet form on the same network using packet-switching technology.

In order to send voice traffic over a data network, speech must first be digitized by sampling the analog signal at various levels and encoded into bits. On the receiving end of the conversation, the bits must be decoded to

recreate the voice traffic to its original form. This process is accomplished through the use of codecs (coders/decoders). The quality of the conversation depends chiefly on the physical number of bits used to encode the voice. In an ideal world, an infinite amount of bits would be available for sampling so that when recreating the signal, it would appear virtually the same as the original signal. However, this would also require an unnecessary amount of bandwidth and robust voice sampling is not necessary to reproducing a quality signal. Additionally, the type of codec used must be extremely efficient; otherwise a substantial delay will result. In short, the key to designing a network that supports voice is based on three QoS criteria - low latency, low jitter and minimal packet loss.

There are a number of things which can impact guaranteed QoS for voice traffic such as, type of transmission medium, network architecture (i.e., ATM versus Ethernet), etc. For now, the focus is on methods to minimize the processing delay of the signal itself. Later, the types of technology the UTN will employ, to further guarantee the QoS necessary for voice traffic, such as packet loss and jitter will be discussed.

H.323, defined by the International Telecommunications Union-Telecommunications Standardization Sector (ITU-T), is the international standard for transmitting multimedia across a packet-switching network. Under the umbrella of H.323, are a set of protocols (see figure 2.1) for the encoding of voice and video. H.323 is independent of the packet network and the protocols over which it runs are not specified.

H.323 Standard	
Audio Codecs	Video Codecs
G.711	H.261
G.721	H.263
G.728	
G.729	
G.723.1	

Figure 2.1 Audio and Video Codecs Under H.323

For purposes of the UTN, the only codecs of interest are those that pertain to Voice-over-Internet Protocol (VoIP). In addition to the codecs specified under H.323, there are a number of speech compression algorithms for handling VoIP. Figure 2.2 lists the major codecs, the data rate, and compression delay associated for each algorithm.

Codec	Data Rate (Kbps)	Delay (ms)
G.711	64	0.125
G.721	32	0.125
G.728	16	22.5
G.729	4 - 8	15
G.723.1	5.3 and 6.3	37.5
FS-1016	4.8	15
FS-1015	2.4 - 4.8	22.5
GSM	13	20

Figure 2.2 Speech Compression Algorithms

The question remains, from this set of codecs, which one best suits the UTN model and provides the highest QoS? The answer is G.711.

G.711 describes the requirements for a codec using Pulse Code Modulation (PCM) - a standard method by which analog waveforms are digitized. The process entails taking a voice sample every 1/8000-second, with only eight bits being sent to encode each sample. As a result, only 256 different levels may be encoded and a rate of 64 Kbps is attained. Not coincidentally, this is the same rate supported by the PSTN and because of this, G.711 is often referred to as an

uncompressed codec. However, this does not fully answer the question of why G.711 is best suited for the UTN. Why not simply select a codec with a higher compression ratio to minimize the use of bandwidth?

In general, codecs with larger levels of compression introduce more delay. For example, as illustrated in figure 2.2, G.729 encodes speech at four to eight Kbps, but adds a one-way delay of approximately 15 milliseconds (ms). As a rule of thumb, voice traffic should not exceed a round-trip delay of more than 400 ms [Ref. 1]. Anything above this threshold is intolerable to carrying on a normal conversation. When using G.711, the delay is negligible. In addition, the UTN will provide the user with ample bandwidth for voice traffic and G.711 was specifically designed for networks where bandwidth is not limited. A final reason for selecting G.711 is its interoperability. While there are many vendors for various VoIP applications, H.323 stipulates G.711 as the mandatory speech codec standard. As a result, users from the UTN could anticipate (theoretically) being able to communicate with users from another network.

In summary, G.711 is highly effective in minimizing the processing delay for digitizing a voice signal. In addition, it allows for sending as much data about the voice sample with as few bits as possible thus maintaining high-quality

speech, which is effectively indistinguishable from using the conventional PSTN. While G.711 does not play as important a role in the QoS requirements for voice traffic as the actual network implementation, by using this codec it does specify the necessary bandwidth for voice in the UTN, 64Kbps per channel.

B. VIDEO

Akin to voice, the digitizing of an analog video signal involves sampling it at various points and converting it into a digital video stream. However in contrast to voice, a digital video signal has somewhat different QoS requirements. While latency and jitter are still important QoS requirements, some packet loss is tolerable in a video session. The biggest difference between the two types of traffic is bandwidth utilization. Therefore, the focus is not exclusive to minimizing processing delay. A video signal requires considerably more bandwidth because of its crucial components, such as resolution, color depth and frame rate. Resolution refers to the horizontal and vertical dimensions of a video signal, color depth the number of bits used to convey color, and frame rate, the number of frames displayed per second. A simple example will aid in understanding the

profound impact the three criteria can have on bandwidth utilization.

In this example, the intent is to transmit an uncompressed digital video signal across a network. The video session will have a resolution (R) of 800 horizontal pixels by 600 vertical pixels, a color depth (C) using 24 bits (3 bytes) and standard frame rate (F) of 30 frames per second. To determine the required bandwidth (B), a straightforward mathematical formula is used, $B = R * C * F$. Therefore, in order to send this digital video to the network, a bandwidth of approximately 345.6 Mbps (or 43.2 MBps) would be required. Fortunately, there are established methods to reducing the amount of bandwidth utilization in sending digital video, namely compression techniques.

There are essentially two types of compression techniques, lossless and lossy. A lossless compression algorithm is a technique whereby the data restored during decompression is exactly the same as the original data. This type of technique is useful when sending data that cannot tolerate errors, such as executables, because any changes in the digital makeup of the file will render it useless. In contrast, a lossy compression algorithm does not guarantee the decompressed data will appear as the original. In fact, a lossy algorithm removes data that cannot be restored

later. Additionally, lossy compression algorithms typically achieve much higher compression ratios than lossless algorithms and work well on video images files because the changes in the data are not detectable to the human eye. For these reasons, the center of discussion is on the use of lossy compression.

There is an assortment of lossy algorithms for encoding or compressing a video signal. Figure 2.3 lists the foremost compression algorithms for video.

Standard	Data Rate
H.261	$P * 64 \text{ Kbps}$
H.263	$P * 64 \text{ Kbps}$
MPEG	384 Kbps - 3 Mbps
MPEG2	4 - 9 Mbps

Figure 2.3 Video Compression Algorithms

The ITU-T has defined H.261 and H.263, as the first and second-generation video encoding standards. The baseline data rate is 64 Kbps because they were intended for use in an Integrated Service Digital Network (ISDN). Integral multiples ($P * 64 \text{ Kbps}$) are also possible with H.261 and H.263. While these two perform well, the preferred standard for video compression is MPEG. Named after its designers,

the Motion Pictures Expert Group, MPEG is a compression standard that delivers high quality video. For purposes of the UTN, the compression algorithm of choice is the MPEG2 standard.

A great deal can be said about the process through which MPEG2 encodes and decodes a video signal; it is sufficient to state that the algorithm is highly complex. More important are the effectiveness and performance of using MPEG2 in the UTN. For example, in comparison to H.261 and H.263, MPEG2 is more aptly suited to the high bandwidth envisioned in the UTN. While the actual encoding algorithms behind H.261 and H.263 are similar to MPEG2, again it is important to note they were targeted for ISDN speeds, where bandwidth is available in 64 Kbps increments. In contrast, MPEG2 was intended for high-speed networks and operates at data rates between four and nine Mbps [Ref. 2]. MPEG2 is able to achieve this data rate because of its impressive compression ratio. While the ratio varies according to parameters, such as resolution, a typical compression ratio for MPEG2 is 90-1. Previously, computer processors were not fast enough to decompress MPEG video streams. However, with today's high-speed processors decompressing video streams with high resolution is possible [Ref. 3].

Undoubtedly, a great deal more can be said about providing video services over a network. For example, what about supporting a general application like video teleconferencing? While this is an important issue, it opens discussion to variety of topics (e.g., whether to use UDP or TCP, M-JPEG or MPEG, etc.) and exceeds the scope of this thesis. Instead, the focus has been on minimizing the decompression delay and bandwidth for the type of video the UTN will support, broadcast television or video-on-demand at a data rate of four to nine Mbps. Designed with high speed networks in mind, MPEG2 significantly minimizes compression and decompression delay and bandwidth utilization. It is therefore the compression algorithm which best supports the type of video running across the UTN. While these factors are not directly related to the QoS for video, they do have an impact on setting the requirement for bandwidth in the UTN and aid in minimizing the overall delay associated with transmitting video.

C. DATA

Conceptually, all of the requirements described herein could be considered simply as binary streams of information or data. However, for purposes of the UTN, data is quantified as services other than voice or video. It is

important to differentiate between the three because of the various QoS each require. As will be discussed later, the UTN will employ various methods (e.g., virtual circuits, various frame sizes, etc.) to handle each separately. So instead of analyzing a specific type of data application, it is more important to focus on the requirements for data.

Of the three traffic types and with the exception of packet loss, data is perhaps the lowest in terms of QoS requirements. This does not mean data is not important. In fact, the lower QoS requirements can be an advantage, such as more flexibility. For example, if during a telephone conversation packets are lost, there is not ample time for applications on either end to request and retransmit the lost packets and the session could be considered bunk. In contrast, if packets are lost during a data transfer session, there is time for recovery. Therefore, the key QoS requirement for data is to minimize the delay associated with "waiting" for the data.

Fortunately, the combined advancements in high-level protocols and high-speed processors have made great strides in minimizing the delay in data transfer. The hand-off between higher layers of the OSI model occurs much faster and delay is often negligible. In fact, many of today's computers are able to process information faster than the

medium can deliver (e.g., a 56Kbps dial-up connection). This does not mean computing power has out-paced the medium (think fiber optics), but it does illustrate that a bottleneck is occurring with regards to bandwidth.

For the UTN, the key to minimizing the amount of time a user spends waiting on data transfer is providing sufficient bandwidth, in this case 1.5Mbps. The yardstick of 1.5Mbps is based solely on today's standard data applications (e.g., HTTP, FTP, NTP, etc.) For example, when surfing the Internet, loading of web pages should be similar to that of turning the pages of a book, or no longer than one second. Based on an average web page file size of generally 100KB, 1.5Mbps is more than enough bandwidth. Granted the user will experience a longer delay for the transfer of larger files, such as MP3s, but this is obvious, expected by the user and not the norm. To be sure, bandwidth demand will continue to increase over time. However, providing the equivalent of a T-1 line to each user is far superior to today's standard data networks and serves as a sound baseline to supporting the data QoS requirements.

The matrix depicted in figure 2.4 encapsulates the various QoS and bandwidth requirements for each of the three UTN traffic types, voice, video and data. The numeric values in the matrix fields define the threshold for each of the

major transmission impairments to quality service and the bandwidth required for each type of service. For example, voice must have a latency no greater than 400ms, a jitter buffer delay no greater than 20 ms, and a packet loss less than five percent [Ref. 4]. If a field in the matrix does not have a specific value (e.g., low), this means the value is relative. For example, the goal of the UTN is to minimize or keep low the jitter time for transmitting data.

Traffic Type	Latency	Jitter	Packet Loss	Bandwidth
Voice	400ms <	20 ms <	5 % <	64 Kbps
Video	400ms <	20 ms <	7 % <	4-9 Mbps
Data	~ 1 s	300ms <	Zero	1.5 Mbps

Figure 2.4 QoS Requirements for the UTN

In summary, this section has focused on stating the QoS requirements for the various types of traffic the UTN must be capable of supporting. In addition, some of the methods (i.e., G.711 and MPEG2) that indirectly impact the QoS requirements have been looked at that specifically minimizing processing delay and bandwidth utilization. The following chapters explore the technology and implementation of the UTN to further minimize the adverse affects on QoS.

III. TECHNOLOGY ANALYSIS

With the network traffic requirements defined, the focus shifts to the second logical step to designing the UTN, an analysis of the technology to meet the requirements. It is important point out that an analysis of wireless technology is not included. Without question, major advancements are being made in wireless technology and could play a future role in the UTN. However, in its current state, wireless technology does not provide the necessary bandwidth to meet the requirements of the UTN so the discussion is dedicated to wireless technologies.

The analysis of the technology behind the UTN follows a layered approach similar to the OSI model. The analysis begins at the physical level by analyzing the two primary transmission mediums, copper and fiber. Valid arguments are presented for fiber as the medium of choice and the discussion continues with a detailed description of the fundamentals to fiber optic technology. The next tier in the analysis looks at the vehicle for carrying data across the network, Synchronous Optical Network (SONET). The basic principles, like frame format, transmission speeds and multiplexing are outlined. The chapter concludes by describing the two connection-orientation technologies

(i.e., Asynchronous Transfer Mode and Frame Relay) that will best fulfill the QoS requirements outlined in Chapter II.

A. COPPER VERSUS FIBER

From the early days of the telegraph, to modern day Local Area Networks (LAN), copper has been the primary medium for relaying information. However, with networks expanding and applications demanding more and more bandwidth, the use of copper, especially for the deployment of new networks, is no longer the primary transmission medium of choice. For a network such as the UTN, copper will simply not meet the requirements. There are a number of reasons why fiber should be the transmission-medium of choice versus copper. While there are different types of fiber (i.e., single-mode and multi-mode) and different types of copper (i.e., twisted pair and coaxial), the arguments presented below apply in general to all types. The three most important reasons for selecting fiber over copper are Segment Length, Cost and Bandwidth.

1. Segment Length

Copper is restricted in terms of its maximum length (distance) between network segments (e.g., routers, switches, end nodes, etc.) For example, when using copper as the transmission medium, the maximum segment length is

approximately 2500 meters in an Ethernet [Ref. 5] and roughly 5 kilometers in a WAN [Ref. 6]. Beyond these distances, the signal deteriorates to the point where the information on the carrier signal is indistinguishable, meaning the information is lost. These physical segmentation length limits of copper are attributed primarily to three factors: attenuation, distortion and noise.

Attenuation is the loss of energy as a signal propagates through guided media, such as copper wires or fiber. The level of attenuation depends on the frequency. A signal loses energy logarithmically with distance. This loss is expressed in decibels (dB) per unit of distance. For copper, attenuation is more detrimental than in fiber because each signal on a copper wire is not simply one waveform, but composed of a series of Fourier components and this leads to distortion.

The different levels of attenuation for different frequencies cause distortion. If all Fourier components of a decomposed fundamental frequency were equally affected by attenuation, the resulting signal would be reduced in amplitude but not distorted. However, in a copper medium, each component is attenuated by a different amount, which results in a different Fourier spectrum at the receiver and consequently in a different signal. At too high levels of

attenuation, the signal may become so distorted that the receiver cannot even detect it. Another possibility is caused by the different propagation speeds of the different Fourier components. For digital data, fast components from one bit may overtake slow components of the preceding bit, thus interfering with the representing signal and increasing the probability of incorrect reception.

There are methods to circumventing the problems associated with attenuation and distortion, usually by using what is commonly referred to as structured cabling (e.g., amplifiers or repeaters). However, even with structured cabling, it is often not possible to restore enough of the signal to its correct original form. In addition, using structured cabling introduces more delay and cost for the active components of the network.

A final impairment to the segmentation length of copper, and perhaps the largest, is noise. Noise can be defined as the effect of unwanted energy on a signal from sources other than the transmitter. Copper is susceptible to thermal noise and Electromagnetic Interference (EMI), induced by the environment. Thermal noise cannot be avoided since the random motion of the electrons in a wire causes it. EMI could affect a copper wire if it were in close proximity to a large electric current, for example power

lines. Energy surges on power lines can create impulse noise, which - for digital data - can literally wipe out bits. Another affect of EMI, known as cross talk, can occur between two wires that are close to each other due to the inductive coupling. In contrast to copper, fiber is not prone to thermal noise or EMI. In a fiber medium, the signal is not carried by way of electricity, rather a beam of light. In a beam of light, the information travels via photons, which are free of electrical charges. As a result, no electrons are present to affect each other or be influenced by electrons outside the fiber wire.

While fiber does have physical length limitations, they are nowhere near as restrictive as copper. In fact, distances of up to 9000 miles for transoceanic routes are now possible [Ref. 7]. In general, because of the inherent low attenuation of signals transmitted via fiber, much longer distances are attainable than when using copper.

2. Cost

Perhaps the biggest misunderstanding regarding implementing a network with fiber is cost. To begin, there is a misconception that in order to use fiber, a large amount of technical training is required. People interested in designing or implementing a new network often steer clear of using fiber in fear of the high cost associated with

hiring engineers familiar with fiber optic technology. It is unfortunate, because using fiber is not as involved, nor as costly as it may appear. The misconception largely stems from the knowledge required to connect fiber lines.

In the early days of fiber optic technology, the techniques used to connect fiber optic lines were complex. The mechanical splicing of fibers required a high degree of precision. Even more effort was necessary if two pieces of fiber had to be fused (melted) to form a solid connection. However, the technology has advanced far enough that connecting fiber lines is relatively easy. While it still requires a professional to install, the time and consequently the cost related to this is relatively low.

A second myth about the cost of fiber is the price tag for the actual medium itself. Per unit of distance (e.g., meter or foot) fiber is cheaper than copper. For example, a single-mode, 12-strand optical fiber costs approximately \$0.60 per foot, whereas a shielded twisted pair copper wire costs approximately \$0.75 per foot. In general, fiber is more inexpensive to purchase in bulk quantity than copper. Clearly a cost analysis for implementing a network based on the simple metric of price per unit of distance for the medium would be impractical.

In order to perform a thorough cost analysis of implementing a fiber based network, one must also consider factors that have a more profound effect on the overall price tag, like cost of installation and ancillary equipment. First the cost associated with the installation of fiber can be enormous. For example, the labor cost for installing fiber-optic lines is approximately \$100/meter for underground and \$30/meter for above ground. While these figures are certainly substantial, one must recall the added benefits that will be realized by installing the fiber. In addition, once the fiber is in place, the cost related to physical maintenance of the line is less than that for copper.

The second largest factor to the cost of implementing a fiber based network concerns the supplementary equipment required to use fiber, such as the light emitting devices, repeaters, amplifiers, etc. While the price tag associated with these pieces of physical hardware are often expensive, they are decreasing. The industry associated with fiber optics is not as mature as copper. As the industry matures standards are being put in place, which is one of the primary reasons for the decline in prices for fiber optic equipment. Take for example the adaptors used in a fiber network. Initially there were numerous types, in limited

production, thus increasing the cost. However, as the industry settles, it appears that only two solutions will most likely become standards: Volition VF-45 and MT-RJ.

At this juncture it may appear that copper is still cheaper than fiber in terms of overall cost. This is simply not true. The authors do not claim that the cost of implementing a fiber-based network is inexpensive; the point is that the cost is cheaper or equal to implementing one with a traditional copper architecture. To further illustrate the overall cost advantages, figure 3.1 details a simple corporate intranet utilizing a copper versus fiber architecture.




Figure 3.1 Office Building Network Design

The left side of the figure depicts a traditional copper Ethernet. The right side is the same corporate network, only implemented with fiber. Notice that in the Ethernet model, a separate active component (e.g., router or switch) is located on each floor. This is required due to the length restrictions imposed on each segment. Additionally, air conditioning, as well as the noise of the active components, must be considered. It is readily apparent how rapid cost could accumulate. In contrast, the fiber-based network does not require an active component on

each floor. In this example, the fiber network needs only one central router office for the entire building complex. As a result, no added costs are incurred for additional active components or office space to house the equipment, nor is there as large a requirement for air conditioning, and noise is minimized.

In summary, the overall costs associated with implementing a network with fiber (as the transmission medium) continue to decline. The technical expertise and cost of physical hardware are much cheaper than a decade ago. In fact, the cost savings are 30-50% over that of initial fiber optic network solutions [Ref. 8]. More importantly, the cost of implementing a fiber network is comparable to that of a traditional copper Ethernet due to the elimination of peripheral equipment and overhead. Different installations will yield different numbers, however in a realistic comparison between copper and fiber, an all-fiber network is equal or cheaper in cost.

3. Bandwidth

The foremost advantage of using fiber versus copper is bandwidth. While remarkable strides have been made to prolong the use of copper by squeezing even higher data rates from the medium, these methods are reaching their physical limits. Emerging technologies, such as Gigabit

Ethernet, are not considering copper as the medium for high speed networks (i.e., 10 Gigabit Ethernet) in engineering design [Ref. 9]. Simply stated, the bandwidths realized by using fiber dwarf those in comparison to copper. For example, through the use of Dense Wave Division Multiplexing (DWDM), bandwidths ranging from 160 Gbps to 400 Gbps [Ref. 10] are now achievable.

While copper might still be a viable medium - perhaps within the confines of a home - it should be obvious from the arguments presented in this section that fiber is the medium of choice for the UTN. Fiber optics is the solution to supporting an infrastructure, such as the UTN, capable of carrying voice, video and data [Ref. 11].

Although only three of the major advantages of using fiber over copper have been addressed, it is important to briefly point out there are a number of additional advantages to using fiber [Ref. 12].

- BER: Fiber has lower bit error rates (BER) than any other medium.
- Security: Fiber is very difficult to tap. By monitoring signal strength it is easy to determine a possible wiretap.
- Compact: Fiber is very strong, small and lightweight.

- Maintenance: Fiber optic components require much less upkeep than copper.

In summary, fiber offers more pros than cons for any transmission medium.

B. FIBER FUNDAMENTALS

To better understand the benefits and tremendous bandwidths realized by using fiber, this section focuses primarily on the physical aspects of fiber optics. First a brief introduction to how data is transmitted via light is introduced, next an analysis of the two specific types of fiber (i.e., multi-mode and single-mode) and concluding with a few of the physical transmission characteristics associated with the medium.

The basic principle behind data transmission using fiber optics is to "hook" the data on light impulses instead of electromagnetic signals. In order to have light refracted inside fiber, it is essential the light energy be completely contained. This containment is achieved by making use of a medium's refractive index, which is defined as the ratio of the velocity of light in the material to the velocity of light in a vacuum. When light encounters a boundary between materials of differing refractive indices (e.g., air and water), a portion of the light is bent while the remainder

is reflected. Light rays can be inserted into a fiber at different angles, which defines the modes (e.g., single-mode or multi-mode.) As long as the angle of light stays below a pre-defined critical angle, all of the light energy will be reflected. If the light exceeds this critical angle or limit, energy will pass through the boundary and the goal of total containment is not achieved.

As illustrated in figure 3.2, a typical fiber-optic cable consists of three parts: core, cladding, and coating. To send data through the cable, light is inserted at the core. The light cannot escape from the core because of the higher refraction index of the core in comparison to the cladding. This is possible because the boundary between the cladding and the core, which reflects all light inserted at less than the critical angle back into the core. If complete containment is achieved (i.e., all light inserted has been reflected), this is called total internal reflection or "fiber-optic effect". The coating is simply an environmental protection agent.




Figure 3.2 Typical Optical Fiber

Data can be inserted to the fiber using different encoding methods. Regardless of the method, the light must be modulated, either directly or externally. For analog encoding the transmitter changes amplitude, phase, or frequency of a light wave in a smooth continuous fashion, while digital transmitters shift the same attributes between distinct states. Basically, a logical one is represented by a light impulse, a logical zero by the absence of light.

Before discussing the problems related to transmission impairments with fiber, it is necessary to first analyze the two primary types of fiber: multi-mode and single-mode.

1. Multi-mode & Single-mode Fiber

As the name implies, a multi-mode fiber can contain several different modes. Recall that the modes of a fiber

refer to the different angles of the light inserted. In addition to the name differentiating it from a single-mode fiber, multi-mode fiber can be further divided into step-index and graded-index types.

In a step-index fiber, the refraction index of the core is the same throughout its diameter and the corresponding refraction index of the cladding is one step lower. A typical size designation for this type of optical fiber is 200/230, which means that the core has a diameter of 200 microns and the cladding a diameter of 230 microns. In terms of throughput, the step-index fiber is slower. The maximum bandwidth usually tops out at about 20 megahertz per kilometer (MHz-km), but the usable bandwidth is even lower, approximately five MHz-km (MHz-km is the typical unit of measurement for bandwidth when discussing fiber optics. For example, a fiber with a maximum bandwidth of 500 MHz-km is capable of lossless transmission at a frequency of 500 MHz over a distance of 1 km, or 1 GHz over a distance of 500 m).

An alternative to a step-index fiber is a graded-index optical fiber, where the difference in core to cladding refractive indices is gradual. The core has a high index in the center, which becomes gradually lower as it approaches the outer diameter of the cladding. The core dimensions of this optical fiber type are normally in the range of 50 to

62.5 microns with a cladding of 125 microns in diameter, as depicted in figure 3.3. Usual bandwidths range between 100 and 800 MHz-km.




Figure 3.3 Multi-Mode Graded-Index Optical Fiber

The second major type of optical fiber is single-mode. In terms of throughput, a single-mode step-index fiber can achieve the highest transmission rates. With its relatively small core of 8 - 9 microns and a cladding size of 125 microns (figure 3.4), single-mode fiber allows only one mode of transmission. In other words, only one ray of light is inserted. Due to its excellent bandwidth-length product of more than 25 GHz-km and low attenuation, single-mode fiber is ideal for long distance communications. It should be noted that one drawback to having such a small core is aligning or splicing single-mode fiber for connections is more difficult.




Figure 3.4 Single-Mode Optical Fiber

2. Transmission Impairments

Fiber is not a "pure" transmission medium and impervious to errors. As is the case with copper, fiber is also subject to some transmission impairments. Fortunately the impairments are not as debilitating and there are methods to minimizing their effects. Therefore, a brief discussion of the two primary transmission impairments to fiber, specifically attenuation and dispersion, and solutions to these problems are outlined below.

Perhaps the biggest impairment is the decrease in signal quality due to absorption (loss of energy) in relation to distance, commonly known as attenuation. Loss of signal strength is dependent on two factors, the wavelength of the ray of light and the quality of the glass in the core. The level of absorption differs for all wavelengths,

but there are three wavelength "windows" that have the least absorption levels (illustrated in figure 3.5).




Figure 3.5 Fiber Attenuation and Wavelength

The first window appears at 850 nanometers (nm). Fiber that operates on 850 nm is typically used in low-cost solutions due to the less complex manufacturing process for transmitter and receiver. The second window at 1300 nm has better absorption characteristics and less material dispersion (discussed in the next paragraph). However, equipment operating in this window is more expensive than that designed for 850 nm. The smallest absorption rates are

attained in the final window, at wavelengths of about 1500 nm, which make it ideal for long distance connections.

Fortunately, the effects due to attenuation can be compensated by use of signal regenerators. Unfortunately, these devices need to perform an electro-optical conversion in order to analyze the signal, recreate and reconvert it before placing it back onto the fiber strand. This dual conversion takes time and adds additional delay to the network.

An alternative to traditional signal regenerators are in-fiber amplifiers like the Erbium-Doped Fiber Amplifier (EDFA), which is created by doping an otherwise ordinary single-mode fiber with erbium during fabrication. The principle operation behind an EDFA is relatively simple: a high-powered laser source at 980 nm is mixed with the arriving (and relatively weak) input signal on a special section of erbium-doped fiber, typically about 10 m of fiber coiled to reduce the form factor of the device. This pump signal excites the erbium ions to a higher energy state. The input signal photons striking those excited ions cause them to transfer energy to the signal and consequently the erbium ions return to a lower energy state until pumped again. The input signal leaves the EDFA with a higher energy level.

The second most common transmission impairment to fiber is dispersion, often divided into two categories, modal and chromatic dispersion.

Modal dispersion occurs due to the different distances each mode has to travel based on its insertion angle. The steeper the angle, the longer the distance a light ray has to travel, which changes the shape of the signal. Modal dispersion is a primary concern in multimode fibers, where more than one mode is present. The largest effect of modal dispersion can be observed in multimode-step-index fiber (refer to figure 3.6). To overcome this deficiency, graded-index fiber was developed. In a graded-index fiber, a light ray that moves from the core outwards is refracted at each layer until in the outmost layer it runs parallel to the fiber axis. Since the refraction index is gradually decreasing outwards, light propagates faster in the outer areas, which eliminates most of the effect of modal dispersion.




Figure 3.6 Dispersion for Different Fibers

Chromatic dispersion is the sum of wave-guide dispersion and material dispersion. Waveguide dispersion depends on the refraction index profile and geometry of the fiber. In some instances, waveguide dispersion can cause light to cluster and eventually disappear. However, the larger impact of chromatic dispersion is due to material dispersion. Material dispersion comes from the fact that each light impulse, even if transmitted from one source, is composed of several components of the optical spectrum and the propagation speed of each of these components vary.

There are a number of light sources available, such as light emitting diodes and lasers. However, a perfect monochromatic light source does not exist and there will always be variations in frequency within one ray of light. There are methods of reducing the effects of chromatic dispersion, such as waveguide dispersion counteracting the effects of material dispersion, but the effect will always be present. Instead, it is more important to determine what function the type of fiber will serve in a network. For example, in a step-index multimode fiber, the effects of chromatic dispersion are high. However if the fiber is simply being used to span a short distance (e.g., a campus) the effects are negligible.

Summing up, fiber is not immune to the laws of physics. There are physical impairments to transmitting data via fiber optics. Fortunately, these impairments can be diminished through the methods such as EDFA or sound reasoning as to which type of fiber is necessary to complete the task.

C. SONET

Without question, one of the major points to using fiber is to transport gigabits or even terabits of in a shorter time frame than it is possible when using

traditional copper media. However, simply having a fiber optic line in place does not answer the question of "how" to place the information onto the carrier. As mentioned earlier, it is anticipated that the majority of applications in the UTN will be based on IP. However, the question still remains on how to place an IP packet, or for that matter an ATM cell, FR frame, or any container of bits representing data, onto a fiber line? There are methods such as DWDM, yet it is still under development and the direct placement of Ipv4 packets onto fiber is not feasible because of the missing start and end delimiters of those packets [Ref. 13]. Luckily, a very mature and capable technology already exists to placing data onto a fiber strand, Synchronous Optical Network (SONET).

Before beginning a discussion of the inter-workings behind SONET, a point regarding SONET and Synchronous Digital Hierarchy (SDH) needs to be addressed. Prior to SONET, fiber-based systems in the United States (U.S.) were all vastly proprietary and there was very little interoperability. Frustrated, telecommunication carriers approached the American National Standards Institute (ANSI) asking for a standard. As a result, ANSI defined SONET as the standard for optical telecommunications transport in the U.S. However, this did not solve the problem of

incompatibilities with networks outside the U.S. Consequently, the ITU-T began an effort of its own to standardize digital transmissions over fiber known as SDH. The SDH specifications are defined in standards G.707, G.708, and G.709 and the rules of SDH take precedence on international links. Despite some minor differences in implementation, all SONET specifications are in compliance with the ITU's SDH recommendations. Therefore in the following sections, SONET and SDH are used interchangeably.

1. Supported Speeds

SONET currently defines ten different speeds, but only four speeds are supported by equipment. Each speed level has an associated Optical Carrier (OC) level and an electrical level transmission frame called the Synchronous Transport Signal (STS). The equivalent SDH designations, referred to as Synchronous Transport Modules (STMs), are also defined for the four supported speeds. Figure 3.7 details the physical line bit rate, the payload bit rate, and the overhead bit rate for each supported SONET level and lists the SDH equivalents for these SONET levels.

Optical Level	Electrical Level	Line Rate (Mbps)	Payload Rate (Mbps)	Overhead Rate (Mbps)	SDH Equivalent
OC - 3	STS - 3	155.520	150.336	5.184	STM - 1
OC - 12	STS - 12	622.080	601.344	20.736	STM - 4
OC - 48	STS - 48	2488.320	2405.376	82.944	STM - 16
OC - 192	STS - 192	9953.280	9621.504	331.776	STM - 64

Figure 3.7 Supported SONET Speeds and SDH Equivalents

It is important to note that the relation between overhead and payload remains "fixed" for SONET/SDH at a rate of 3.45%. This is a definite advantage in comparison to other asynchronous multiplexing schemes, which increase exponentially with higher line rates [Ref. 14]. The reason for this constant rate can be explained by examining how SONET performs multiplexing.

2. SONET Multiplexing

Unlike the complicated multiplexing structures used in the old digital hierarchy, multiplexing STS signals is fairly simple: an STS-N is equal to N byte-interleaving STS-1's. Basically, to combine multiple STS streams, every stream gets allocated a slot in a round robin fashion. The allocation uses the same schemata as conventional Time Division Multiplexing (TDM), with the exception that throughout the network-link the time must be synchronized (as the name SONET implies). Having an STS-3 signal, which is composed of three STS-1 signals, the device receiving the

STS-3 signal knows exactly which time slots contain data of which enclosed STS-1 stream. The same applies for an STS-12 signal composed of four different STS-3 signals, which is composed of three STS-1 signals each, and so on.

The reason for basing everything off the STS-1 stream and with a line rate of 51.840 Mbps is relatively simple. It is derived directly from the basic STS-1 SONET frame format.

3. SONET Frame Format

A basic STS-1 SONET frame consists of 810 bytes and is transmitted 8000 times per second, which leads to a line rate of 51.840 Mbps. As illustrated in figure 3.9, the 810 bytes are divided into nine rows with 90 bytes each, where the first three columns constitute the overhead. As depicted in figure 3.8, the Total Overhead (TOH) is further divided into Section Overhead (SOH) and Line Overhead (LOH). The SONET payload is carried in the Synchronous Payload Envelope (SPE) and its rate equals 50.112 Mbps.

		Columns									
		1	2	3	4	5	6		88	89	90
R O W S	1	SOH	POH						88	89	90
	2	91	92	93	94	95	96		178	179	180
	3	181	182	183	184	185	186		268	269	270
	4	271	272	273	274	275	276		358	359	360
	5	361	362	363	364	365	366		448	449	450
	6	451	452	453	454	455	456		538	539	540
	7	541	542	543	544	545	546		628	629	630
	8	631	632	633	634	635	636		718	719	720
	9	721	722	723	724	725	726		808	809	810



Transport
Overhead
(TOH)



Payload

Figure 3.8 Basic SONET STS-1 Frame Format

The SPE contains more overhead, called the Path Overhead (POH), but this overhead is considered as part of the user data. POH can be found in the first column of the SPE and is needed for additional flexibility. Since the signal arriving at the terminal multiplexer is very often an electrical one, it is probably not as tightly synchronized as a regular SONET signal. However, instead of waiting until a new frame is created, SONET places the information wherever it is convenient. Thus, an SPE will in most cases not be aligned with the first row of the TOH. A pointer (depicted as an arrow in figure 3.11) within the TOH points to the actual beginning of the SPE to ensure that the receiving end is able to detect the beginning of a SONET information payload section. Even more flexibility is gained because the first byte of an SPE does not necessarily start

in row four of the frame, hence an overlap will occur. The flexibility of placing payload envelopes whenever they arrive at the terminal multiplexer is shown in figure 3.9.




Figure 3.9 Frame Overlap

As mentioned earlier, basic STS-1 SONET frames are combined to form STS-N frames. In terms of frame structure the procedure is very simple. Instead of having one TOH, an STS-N frame contains N TOH's and the same applies for the SPE. An STS-3 frame is 270 columns wide (3-90) with the first 9 columns being the TOH (3-3) and the remaining 261 columns (3-87) are payload capacity. The number of rows does not change with higher-level frames but remains constant at

nine. Therefore, the overhead is just N times larger than for a basic STS-1 frame. Again, this is one of the distinguishing features of SONET over other technologies, where the overhead increases exponentially with speed.

While the preceding paragraphs have focused primarily on the basic principles behind SONET, the focus now shifts to some of the additional advantages and key features of SONET. First, SONET does not have to be channelized. Essentially this means, that although higher SONET rates are gained by adding basic STS-1 streams, it is possible to offer the unchannelized full resulting bandwidth. For example, a subscriber could receive the total resulting 155.52 Mbps of an OC-3 fiber link. This process is known as concatenation and the designation for a concatenated STS-N is to add a lower case "c" to the notation. For example, a concatenated STS-3 becomes an STS-3c.

A second added feature of SONET is the concept of a Virtual Tributary (VT). For instance, having an STS-3 capable line it would be nice to divide the available bandwidth into smaller portions so as to offer subscribers lower rate payloads like 1.5 Mbps. With SONET, this is achieved by mapping the lower rate signals into sections of an STS-1 frame. These sections are called VTs, carrying independent and probably different types of sub-rate

payloads. VTs are then grouped into Virtual Tributary Groups (VTGs). Each STS-1 frame is either divided into exactly seven VTGs or not divided at all. Currently there are four VT sizes defined: VT 1.5 can transport the equivalent of an DS-1 signal (1.5 Mbps), VT 2 matches an E-1 (2.048 Mbps) signal, VT 3 takes an unchannelized DS-1c signal, and VT 6 is capable of delivering an DS-2 signal [Ref. 15].

A final and biggest feature of SONET is its ability to support a variety of payload types. For the UTN, the primary packet form will be IP. However, as alluded to earlier, it is not feasible to place IP packets directly onto fiber. Instead, IP packets must be encapsulated in a separate container, which leads directly to the next section in the analysis, the connection orientation of the UTN.

D. CONNECTION ORIENTATION

When designing a network capable of handling voice, video and data integration, there are numerous options from which to choose a connection orientation. Some examples include Switched Multimegabit Data Service (SMDS), X.25, ATM, FR, etc. Each orientation has its own pros and cons; the key is selecting the correct tool for the problem. The authors propose a "Hybrid Network", utilizing both ATM and FR technology, which best suits the connection orientation

for the UTN. For this reason, the remainder of the chapter focuses on the two technologies.

Both technologies make use of the concept of virtual circuits (VCs). Additionally, ATM introduces virtual paths (VPs), which will be discussed together with the concept of VCs in the ATM section. Traditional telephone communication is based on the existence of physically switched circuits to connect the users on both ends through physical lines. The physical lines usually offer much more bandwidth than is required to transmit the communication data, resulting in wasted bandwidth. VCs were designed to enable maximum utilization of physical lines and allow the same line to act as a medium for more than one transmission simultaneously.

In most cases, data packets or frames must travel through a series of switches or routers. In a packet-switched network, the route for each packet may vary from its successor or predecessor. Consequently, packets may not arrive in the correct order and the application or a higher-level protocol is needed to ensure orderly and complete reception. In contrast, while a VC has the same characteristics as a physical switched circuit - each packet will arrive in the correct order using exactly the same route. In order to enable bi-directional circuits, a route in the reverse direction must be established, which in most

cases will differ from the route in the opposite direction. Both routes constitute virtual channels; a virtual circuit consists of two virtual channels.

The two primary types of VCs supported by FR and ATM are Permanent Virtual Circuits (PVCs) and Switched Virtual Circuits (SVCs). PVCs are similar to a dedicated telephone line; the paths are fixed and once set, not available on demand or on a call-by-call basis. While it is possible the actual route through a network may change from time to time, the endpoint devices on a PVC will not change. Of note, unused bandwidth for a PVC can be used for use by other PVCs or SVCs. SVCs are comparable to a telephone dial-up connection. A client must specify the destination address, similar to a telephone number. While implementing SVCs is more complex than using PVCs, the major advantage is that SVCs permit any-to-any connectivity. As a result, SVCs are highly economical, in terms of bandwidth, for supporting endpoint devices that infrequently communicate with each other.

Why not design a pure IP network using fiber? The question returns one to the requirements of the UTN for the answer. While a pure IP network works quite well for simple data (non-voice & video) transfer, IP is not connection oriented, which is necessary to fulfill the requirements

stated for the UTN. This is the fundamental reason why analyzing two of the predominant connection-oriented technologies available, ATM and FR, is important to designing the UTN.

1. ATM - Asynchronous Transfer Mode

ATM, a connection-oriented, fast packet switching technology, was initially designed as the technology to implement Broadband Integrated Services Digital Network (B-ISDN). B-ISDN replaced the analog telephone local loop with a digital service. This became necessary in order to support the higher bandwidths being offered by improvements in fiber optic technology [Ref. 16]. However, ATM went above and beyond simply supporting higher bandwidths. In addition, B-ISDN services were intended to support scalability, so that a call would be able to obtain the bandwidth and QoS that meets the specific needs of that call. ATM fully met the requirement of scalability. In fact, the QoS capabilities offered by ATM are one of its greatest advantages.

ATM departs from the classic time division multiplexed telecommunications network. In contrast to a time division multiplexed system, where the bandwidth is fixed for the duration of usage, ATM-type multiplexing makes resources that are currently not in use available to other users. ATM provides for the flexible distribution of resources through

advanced statistical multiplexing. This guarantees an average data transmission rate based on pre-agreed bandwidth specifications.

ATM also differs from traditional statistical multiplexing, where no time slots are reserved, by its use of fixed-length packets, which is what makes ATM so distinct. ATM's fixed-length packets are 53 bytes - 5 header bytes and 48 payload bytes. In order to distinguish these fixed-length packets from variable-length packets used in traditional networks, they are called cells.

Why use cells? There are several reasons. The first advantage is related to the physical hardware. It is much easier to design and build hardware for processing packets if one already knows the size of each packet. Second, since all packets are of the same size, it is possible to have numerous switching elements all doing the same thing in parallel. Third, in a traditional multiplexed system, large frames introduce unacceptable delays for time-sensitive data.

Figure 3.10 illustrates how an unacceptable delay may occur in a traditional multiplexed system. In this example, C has a large data-frame to send and A has a smaller, but higher priority, data-frame to send. Once the traditional multiplexer has started to process C's data, its entire

frame must be processed before any other data can be served, thus introducing a delay to A's traffic.




Figure 3.10 Traditional (non-ATM) Statistical Multiplexer

Figure 3.11 shows the same scenario in an ATM network. In this case, the large data-frame from C is broken into cells and interleaved with other cells, thus significantly reducing the delay for A's traffic. Figure 3.11 also displays the inherent scalability of ATM. Because ATM uses a flexible cell mechanism, rather than a strict synchronous time slot, it is called asynchronous. The ability to send data at a variable rate is an advantage for the processing of bursty traffic. Moreover, ATM offers the opportunity of prioritizing time-sensitive traffic (e.g., a voice transmission) over other traffic (e.g., normal binary file transfer).




Figure 3.11 ATM Statistical Multiplexer

Depending on where one is located in an ATM network, there are two types of cell formats, a User-Network Interface (UNI) and Network-to-Network (switch-to-switch) Interface (NNI). The sole difference between the two cell formats concerns the Generic Flow Control (GFC). In an NNI cell, the GFC field is replaced with four extra bits for the Virtual Path Indicator (VPI). Figure 3.12 details a generic UNI cell and a brief description of each field follows.

4	8	16	3	1	8	384
GFC	VPI	VCI	Type	CLP	HEC	Payload

Figure 3.12 ATM UNI cell format in bits

The GFC field is intended for local use only and not widely used. It must be set to 0 when a cell crosses a UNI.

The next two fields, the VPI and Virtual Channel Indicator (VCI), reflect the two-tiered ATM cell address scheme. As illustrated in figure 3.13, cells flow along virtual channels that reside inside of a virtual path. The VCI specifies to which channel a cell belongs. VCI numbers can be any number between 0 and 65,535. Similarly, the VPI number determines the virtual path that contains the channel, and can range from 0 to 255. These numbers may change at each switch the path or channel passes since their scope is local only.




Figure 3.13 Virtual Path & Circuits

The fourth field in an ATM cell is the type. As the name suggests, this field identifies the information contained in the particular cell - either user data or operations, administration, and management data.

The next bit, the Cell Loss Priority (CLP), describes the importance of that particular cell. In other words, it is used as a priority indicator for cells that may be discarded in case of a congestion situation.

The last byte of the header is the Header Error Control (HEC) field. The HEC provides a means to check the validity of a cell header. It contains the result of a Cyclic Redundancy Check (CRC) and may be used for the correction of single-bit errors. Generally, cells with corrupted headers will be discarded.

The last point to address about ATM concerns the various services it was intended to support. ATM was designed to support all sorts of services - voice, video, and data. To accomplish this goal, a series of ATM Adaptation Layers (AALs) were defined. There are a total of five AALs, AAL0 through AAL5 (AAL3 and AAL4 were merged into one layer, AAL3/4). AAL0 is best described as a null data link layer. AAL1 and AAL2 were designed to support applications like voice, which require guaranteed bit rates. AAL3/4 was intended to support both connection and

connectionless oriented services, however due to its shortcomings, this layer is rarely used. Due to the inefficiencies of AAL3/4, a fifth AAL was proposed, called AAL5. This final layer is the choice for most circuits that support data transfer and is more efficient in its use of bandwidth than AAL3/4.

In summary, there are a number of advantages to selecting ATM as the connection orientation of a network:

- Fixed-length cells enable fast, hardware-based switching at speeds far beyond the range of a conventional telecommunications network. Furthermore, the switching equipment does not have to perform complex logical decisions, thus being less expensive.

- The bandwidth available for an application can easily be adjusted to match the actual requirements of that application, hence less bandwidth is wasted during idle transmission times.

- Administration of an ATM network is relatively simple due to its use of virtual paths containing virtual channels. Although a path can have literally thousands or even millions of virtual connections, it can be controlled as only one connection. Therefore the switches need to store much less connection-state information than in a conventional network.

- Cell headers contain all necessary information that is needed to identify their channel. Therefore cells from many different channels can be mixed on one path.

- ATM supports voice, data, and video traffic.

After looking at this brief description of ATM and its benefits, one could pose the question; why not design the UTN using just ATM?

ATM is a technology that is well suited for high bandwidths in the backbone of a network. However, a major drawback to ATM is the procedure for physically accessing the network. In an ATM network all devices - including endpoint devices - must have an ATM address that is used when signaling to establish a VC. Furthermore, emulating a LAN-environment with ATM requires additional hardware devices and protocols, such as LAN Emulation servers and LAN Emulation client software. Combined, accessing an ATM network and LAN emulation, add to the complexity of the endpoint devices and substantially increases the cost of the network. In the UTN, there could be countless devices requiring access to the network, a simple example being a neighborhood of homes. So how will the UTN support cost-efficient, high-speed bandwidths to the home? Fortunately, there already exists a well-defined technology that compensates for the disadvantages of ATM, Frame Relay.

2. FR - Frame Relay

Similar to ATM, Frame Relay is a way of transmitting data by segmenting information into frames or packets. Slightly predating ATM, FR was conceived as part of the ISDN specification, but soon evolved into its own network service. As networks transitioned from analog to digital lines, not only did available bandwidth increase, but also the error rate in transmissions decreased significantly. Moreover, as noise-free fiber-optic lines came into widespread use, transmission errors became almost negligible.

During this time frame, the most popular method for multi-protocol data communications was X.25. A robust error-correcting protocol, X.25 was well suited for overcoming the limitations of analog lines. However, with the digitization and use of fiber in transmission lines, its disadvantages strongly outweighed its advantages.

In a X.25 network, the error-checking capability actually hampers a transmission. For example, during a transmission, error checking is performed at each node (switch) along the route. If a node receives erroneous data, it requests the transmitter to resend the frame. This extensive node-processing time and lengthy packet queuing at switch locations, introduces a considerable delay. In

contrast, in a FR network, all error checking takes place at the originating and receiving device. An intermediate node that receives an erroneous frame simply discards that frame relying on the intelligence of the receiving node to correct for this frame loss. Given that digital lines are relatively free of errors, FR simply eliminates the need for internodal error-correction, thus allowing for remarkably improved and faster throughput of data.

Like ATM, FR is a connection-oriented, fast packet-switching technology. The basic difference between FR and ATM is the variable length versus the fixed length packets (cells). In a FR network, frame sizes are not limited to 53 bytes, but unique to the network or vendor device - upwards of 4,096 bytes in some instances [Ref. 17].

FR is highly efficient in terms of bandwidth utilization because of statistical multiplexing. However, one should be aware that unlike ATM, in which cells are relatively small, in a FR network frames are usually much larger. Referring back to figure 3.12, the delay a high priority frame may encounter depends on the size of a lower priority frame currently being processed at the time the frame with higher priority arrives at a transmission device. In other words, the longest possible delay of high-priority traffic is directly proportional to the maximum frame size.

As mentioned earlier, when FR was being developed the designers were focused on minimizing the overhead associated with older technology (such as X.25). A frame was filled with information as necessary. However, the frame handled speed of transmission and throughput via high-speed communications and with a lower overhead. In figure 3.14, a standard frame format is shown. As one can see, there is very little overhead to FR.




Figure 3.14 Frame Relay Frame Format

Each frame begins and ends with flag bytes. The method used in FR was borrowed from the High-level Data Link Control (HDLC) layer two transmission protocol. The first flag is used to signal the arrival of a new frame to a switching device in the network. The closing flag on a frame indicates the frame end of frame and the switching device can begin processing the entire frame.

The next field is the frame header, which contains a Data Link Connection Identifier (DLCI), the largest field in

the header. The DLCI is the FR virtual circuit number, which has local significance corresponding to a particular destination. Although older implementations used various sized DLCIs, these are obsolete, and a 10-bit DLCI is the norm. The DLCI allows data coming into a frame relay switch to be sent across the network.

Also contained within the frame header are a series of bits known as DLCI Subfields. FR does not use the Command/Response (C/R) Bit. It can be used by higher layer protocols and is not modified by the FR network. The Extended Address (EA) Bit serves as a flag. It is set to zero in all bytes except for the last one. Hence, an EA bit can be used to indicate if the length of an address field is two, three, or four bytes. A Discard Eligibility (DE) Bit can be used to discard frames that were sent in excess of the pre-assigned rate. Switching devices to reroute traffic if the network is congested can use the remaining two DLCI Subfields, the Forward Explicit Congestion Notification (FECN) and Backward Explicit Congestion Notification (BECN) bits.

Following the DLCI is the payload. To promote interoperability, vendors and service providers are required to support payloads with a minimum of 1600 bytes. FR can be expanded to accommodate up to 4096 bytes, however not all

suppliers support this large a frame. Another significant point about the payload is the use of subframes.

Each switching device in a FR network has a predetermined limit to the number of frames it can switch per second. If a large number of short length frames arrive at a switching node, the potential exists for a bottleneck to occur. This could be detrimental if the user were to send voice traffic - which is time sensitive - across a FR network. However, this delay can be minimized through the use of subframes. The basic idea behind subframing is multiplexing voice and/or data on a single FR DLCI. Subframes are contained within a single frame and prevent short frames from occurring, avoid limits on frames per second, and improve efficiency with reduced overhead. A simple example, illustrated in figure 3.15, demonstrates how this goal is achieved by inserting three voice samples into one frame.




Figure 3.15 Use of Subframes in Frame Relay

The last byte of the frame, before the closing flag signaling the end of the frame, is the Frame Check Sequence (FCS) field. The FCS is used to detect errors in a frame. If an error is detected, the node detecting the error doesn't request a retransmittal, it simply discards the frame and readies itself to receive the next frame. While this method could have a catastrophic impact on network efficiency if the transmission medium were noisy, in the UTN all lines are fiber optic, thus the probability for errors is very low. In essence, the FCS is used only to check for corruption. However, it is because of FR's discard policy that the FCS plays a more important role than in traditional networks.

Concluding this discussion of FR, a lot can be said of the benefits it provides over alternative connection-orientation technologies. With the UTN in mind, listed below are some of these advantages:

- Well-established and widely adopted standards that enable interoperability with other network technologies, such as ATM [Ref. 18].

- Supports high bandwidth speeds and is very flexible. Bandwidth can be easily added to support increasing demand or lowered for underutilized links. For example, a customer may only require telephony and data, but no video service.

- Any-to-Any connectivity. Any node in a FR network can communicate with any other node via a PVC or dynamically via a SVC.

- There are numerous FR equipment vendors, thus driving down the cost for customer premise equipment. Additionally, because of FR's intrinsic simplicity, not only is the equipment inexpensive, but easy to configure and operate.

- Perhaps most important is FR's flexible packet size, enabling it to support voice, data and video.

Concluding the discussion on the connection orientation of the network, the authors wish to point out that the proposed "hybrid-network" of ATM and FR is not the sole solution to converging networks. The selection of these two technologies is based on their maturity and ease of interoperability. There are a number of possible alternatives, which are explored and discussed in detail in Chapter V. However, based on "present-day" technology, a combination of ATM and FR best suits the model of the UTN.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. NETWORK DESIGN & SPECIFICATION

Having defined the traffic requirements and technology to be used, the final step in the analysis of the UTN is the actual network design and specification. Since the model is built from the ground up, the authors assume the absence of any existing infrastructure, which provides a great deal of flexibility in the design. The network described herein is designed in such a manner as to be deployed to a large geographical area with thousands of end terminals. The methodology to the design of the UTN follows a tiered approach as depicted in figure 4.1.




Figure 4.1 Conceptual Design of the UTN

The network is split into three distinct levels. The analysis begins with Level 2, detailing the traffic flow over the ATM backbone. This level is contained in one city and provides interconnection between the neighborhoods via OC-192 links thus serving as a local city backbone. Next is Level 1, referred to as the "Feeder Portion" or the multi-neighborhood backbone of the network. This part of the UTN provides a Frame Relay based backbone, connecting several smaller neighborhoods via OC-48 links. In this level, traffic is fed from the backbone to the local access points by introducing two concepts; Frame based User-to-Network Interface (FUNI) and Passive Optical Networks (PON). In Level 0, the last tier of the network, is the local access portion of the UTN. This section describes a neighborhood or an average street, where two FR methods for granting access to users at the end nodes of the UTN are described.

It is important to note, that while no specific analysis of the layer above Level 2 (in this case the proposed DWDM layer) is provided, a SONET mapping to the layer is considered and discussed in the next chapter. By omitting an analysis of this layer, the authors admit the problem of interconnecting alternate networks to the UTN - both heterogeneity and scale - needs to be addressed in

follow-on work. Again, the focus is on implementing the layers contained by the UTN.

The overall design of the UTN adheres to two basic principles: all traffic is handled by circuits and no circuit shall contain duplicate information. The latter is especially important for video traffic, which could easily consume a colossal and unpredictable amount of bandwidth.

A. ATM BACKBONE

The local city backbone makes use of conventional ATM over SONET technology. Typically in network designs, one sees a picture of an ATM cloud with no description to what resides "inside" the cloud. With this in mind, it is important to at least detail the connections and switches within the ATM backbone itself. Internal to the ATM cloud, each switch is connected via multiple single mode OC-192 fiber links, supporting bandwidths of up to 10 Gbps. As to the switches themselves, they will be capable of switching speeds of at least of 2.4 Gbps. This benchmark is based on the ATM Forum specification for 2.4 Gbps at the physical layer [Ref. 19]. The specification is founded on SONET and SDH standards and details how ATM cells are mapped onto the SONET/SDH frame. The ATM cell payload is scrambled and the resulting cell stream is placed into a SPE. ATM cell

extraction operates in similar fashion, only in reverse order. The standard is not only proof that 2.4 Gbps is achievable using ATM in conjunction with SONET, but an indicator that higher data rates will be realized in the near future.

External to the ATM cloud are two types of point-to-point links connecting to various nodes in the network (refer to figure 4.2, created with OPNET Modeler v7.0B). The first are OC-192 links connected to Head Offices (HOs) which provide web access outside the UTN by performing regular Internet Exchange (IEX) tasks. The second type are OC-48 connected to various neighborhoods.




Figure 4.2 UTN Level Two - ATM Backbone

As to the HOs, there will be in place IEX web servers, dedicated to providing SVCs to end users with predetermined bandwidth. Recall for the UTN, at most 1.5 Mbps per connection will be provided. Depending on the number of possible connections, a careful calculation must be performed to determine the amount of web servers that will be required to provide the necessary bandwidth and management of circuits. This is especially important to avoid link congestion. In addition, not only must the number of web servers be adjusted to accommodate the number of possible web users, but also their physical location is an important design factor. Fortunately, with SVC this task is much easier accomplished than in a standard regular unstructured IP network. If a user requires web access, an SVC will be established providing a Committed Information Rate (CIR) of 1.5 Mbps, thus the bandwidth needed for n users is equivalent to n times 1.5 Mbps. Since all users are not requesting web access at all times, a rough estimate of 2000 users per single link appears to be acceptable. Depending on the amount of fiber strands within a point-to-point link and the respective switch capacity, the number of locations for each HO can be determined. Furthermore, the capacity of the web servers involved determines the number of servers per HO.

To support voice services, a CIR of only 64 Kbps is required for each SVC as explained in the requirements analysis. As a result, a calculation for the number of web servers to handle voice traffic is not necessary. In fact, to avoid additional load on the web servers, independent voice servers will perform the task of phone switching, either within the local network, via long haul connections for long distance service, or through access to Private Branch Exchanges (PBXs) to connect to the PSTN.

Clearly the video portion is the most challenging aspect to the design of the UTN. The video signal envisioned will contain all available television stations (channels) and will be inserted to the backbone from video servers located in the HO. As this signal constitutes the largest demand on the backbone, it will be broadcast only once across the backbone. In order to achieve a broadcast, the pool of "well-known" ATM Virtual Channel Connections (VCCs) will be increased throughout the local backbone and its descendents, so that specific Video Virtual Channel Connections (VDVCCs) can be identified. Currently, all well-known VCCs are using VPI zero and VCIs 1-6, 16, and 17. General broadcast signaling is performed on VCI number two [Ref. 20].

In the UTN, the VDVCCs are permanent and assigned a predefined VCI. The "broadcast" video signal transported

through a VDVCC will terminate at the routers within each neighborhood (discussed further in the Local Access section). How will this additional forwarding functionality be achieved? The authors purpose introducing a slight modification to the ATM switches, more specifically, incorporating into the switch a means to recognize a video cell and being able to multiplex the signal. While it would appear that multiplexing would have a negative effect on overall throughput, this is only true if a switch treats a video cell like any other cell. To distinguish a video cell there are several possible solutions. Perhaps the most effective method is to modify the GFC field, which is normally only used by equipment within the end user device [Ref. 21]. By using a unique GFC for video cells, the ATM switch can instantly turn around and multiplex the incoming cell to all outgoing VCIs with the well-known video value; hence hardware switching is not replaced by software methods. This proposed method of broadcasting differs from the current multicasting schemes because it assumes that every network end device must be provided with all available video channels.

As an alternative, the native ATM support for point-to-multipoint connections could be utilized. In point-to-multipoint (broadcast), the connection is unidirectional in

nature and initiated by one "root" node and received by all the other nodes, which are called leaves [Ref. 22]. These connections cannot be used for communication in the opposite direction, but this is not required for a broadcast video stream, as depicted in the figure 4.3. It is important to note that the ATM UNI Version 4 does not support point-to-multipoint connections, so that this version of an UNI will not be used within the UTN.




Figure 4.3 ATM Point-to-Multipoint Connection

Summing up the ATM backbone video portion, it is assumed that through slight modifications in the switching process and VCI handling, a video signal containing all available channels is present only once throughout the backbone. The bandwidth needed depends on the number of

channels and the bandwidth one channel requires. As outlined in the requirements for the UTN, a video channel requires approximately four to nine Mbps, thus 50 channels would occupy approximately 200 to 450 Mbps of the backbone capacity.

B. FEEDER PORTION

Stepping down one level, the design now enters into the feeder portion of the UTN. At this level a conversion from ATM traffic to FR traffic and a subsequent passive traffic splitting takes place. The feeder portion of the UTN is possibly the most complex part of the overall network (figure 4.4, created with OPNET Modeler v7.0B) because it incorporates PON principles as well as the conversion from frame based traffic to ATM. To map ATM cells into FR frames, and vice versa, a Frame-Based User-to-Network Interface is required.




Figure 4.4 UTN Level One - Feeder Portion

1. FUNI - Frame-Based User-to-Network Interface

The concept of a FUNI is based on the ATM Data Exchange Interface (ADXII) and is capable of transparently supporting Service Specific Convergence Sublayers (SSCS) and other higher layers [Ref. 23]. Additionally, the FUNI supports basic ATM UNI tasks such as VPI/VCI multiplexing, signaling, network management, traffic policing, and, optionally, operations, administration, and maintenance functions. Figure 4.5 is a reference model that illustrates the traffic flow between a UTN feeder portion and the ATM based HO.




Figure 4.5 FUNI Reference Model

Without delving too deep into the mapping function, it is important to highlight various steps in the process. At the Segmentation And Reassembly (SAR) sublayer payloads are stuffed into cells at the sending end and extracted at the receiving end. At the Common Part Conversion (CPC) sublayer, additional administrative functions are performed, such as packaging outgoing data into AALx frames and processing incoming AALx frames. The use of "AALx" indicates that the FUNI version used (version 2.0) in the UTN supports AAL3/4 and AAL5. To be consistent with the adaptation layer, the

maximum CPC sublayer length specified in ITU-T recommendation I.363.5, FUNI Service Data Units (SDUs) of sizes up to 65535 octets will be supported. An aside, the FR-ATM inter-work function need not be performed by devices that are part of the networks they join. In fact, interworking can be performed by either a FR network switch with an ATM interface or an ATM network switch with a frame relay interface [Ref. 24]. However, for purposes of the UTN, the conversion is conducted by a separate FUNI device.

The FR side to the interface is a relatively simple process. When a FR device transmits a frame across a VC, the FR service specific convergence sublayer examines the incoming DLCI and looks up the VPI and VCI of the ATM VC on which the frame needs to be transmitted. It then places the FR address and payload fields into the payload of an AAL5 frame by simple encapsulation (illustrated in figure 4.6).




Figure 4.6 AAL5 Frame carrying a Frame Relay Frame

Incoming AAL5 frames are reassembled, their CRC is validated, and the padding is removed. Thereafter, a new FCS is appended to the address and payload fields and the FR

frame is forwarded onto the VC identified by the DLCI in the address field. Appropriate mapping must be performed between FR discard eligibility flags and ATM cell loss priority flags as well as between FR and ATM congestion notification flags. It is interesting to note that traffic for several FR circuits can be bundled across one ATM circuit that connects the two interface devices, as illustrated in figure 4.7.




Figure 4.7 Multiplexing FR Circuits across an ATM Circuit

This "bundling" of multiple DLCIs will be used as often as possible in the UTN. The outgoing FR frames will be translated into the same VPI-VCI combination. Incoming AAL5

frames will start with a FR address field, from which the DLCI can be extracted so that the frames can be inserted into the correct VC. An obvious result from employing this method is a much more efficient use of bandwidth.

2. PON - Passive Optical Network

The next major element in layer 1 of the UTN is the PON device. In a typical hybrid network that employs both ATM and FR, the device immediately following the conversion of ATM to FR is a router. However, by adding a routing device one incurs a processing delay. Therefore an alternative solution was required for the UTN. Enter the concept of a PON.

Fundamentally, on the downstream portion of a transmission into a PON device, all incoming traffic is captured and multiplexed to each outgoing port. The upstream portion uses a Time Division Multiple Access (TDMA) protocol to place incoming packets or frames onto high capacity fiber lines (refer to figure 4.8). More formally, these two directions are defined as follows [Ref. 25]:

- Downstream: direction for signals traveling from the Optical Line Terminals (OLT) to the Optical Network Unit (ONT).
- Upstream: direction for signals traveling from the ONT to the OLT.




Figure 4.8 Functionality of Passive Optical Networking

Transmission in both downstream and upstream directions can take place on the same fiber and components (duplex working) or on separate fibers and components (simplex working). For purposes of clear distinction between up and downstream traffic, the UTN utilizes simplex working. Some PON vendors are using sophisticated encryption algorithms to address data security issues, but this is not necessary for the UTN because the OLTs are FR switches with added intelligence, hence they are still part of the provider's network.

As to the actual capacity of the PON device, according to ITU-T recommendation G.983, the transmission line rate of a PON device should be a multiple of 8 kHz and have one of the following nominal line rates:

- Option 1: Symmetric 155.52 Mbps upstream/downstream
- Option 2: Asymmetric 155.52 Mbps upstream / 622.08 Mbps downstream

For the UTN the asymmetric solution is most feasible because video traffic is broadcast downstream only.

What about bi-directional transmissions in a PON? As the UTN will require signaling (especially for video) this is important. Bi-directional transmissions can be accomplished through wavelength division multiplexing (WDM) on a single fiber, however as stated before, the UTN utilizes SONET. In addition, there is an unintended benefit of using SONET with a PON device, the elimination of complex optical branching splitters like Arrayed Waveguide Grating (AWG) or Dielectric Fiber, which would be necessary when using DWDM. The PON device in the UTN will work at wavelengths in the range of 1260 to 1360 nm, which meets the recommendation of G.983.

The final aspect of the to the feeder portion of the UTN is the FR segment. Recall from figure 4.4 there are a number of FR switches connected to the PON device and

neighborhoods via OC-48 fibers. These are standard FR relay switches and there is not much to describe about their function within the UTN. The only point to emphasis is the "compatibility" of this portion of the network with the rest. As outlined in ATM Forum's Physical Layer Interface Implementation Agreement [Ref. 26], FR supports SONET/SDH. Therefore, the FR portion of level 1 in the UTN is interoperable with the remainder of the network.

C. LOCAL ACCESS

The final level to the UTN is the local loop and the point where users gain access to the network (figure 4.9, created with OPNET Modeler 7.0B).




Figure 4.9 UTN Level Zero - Local Access

Users gain access through FR equipment. While any device with a serial port and appropriate software can act as a FRAD [Ref. 27], two options for implementing the local access are possible. The first option installs the required software at the end users' device, such that the device resides in the home itself. In this situation, the FRAD in the home performs all necessary functions in the creation and decapsulation of frames. These types of devices are often referred to as FR workstations, and have client-server applications running over TCP/UDP. As illustrated in figure 4.10 (created with OPNET Modeler 7.0B), the FR access functionality is an internal process to the device.




Figure 4.10 Node Model of a Frame Relay Workstation

A short analysis of how frames will be handled within the FRAD is important. When a packet arrives at the FRAD from one of the applications running on top of the FRAD module, it will be forwarded to the FR network after checking for valid VC information. When a frame arrives from the network to the FRAD, the DLCI value is analyzed in order to determine if the incoming frame is destined for this device. The Frame Relay to IP Interface (FRIPIF) uses tables created by the IP layer to perform a lookup between IP addresses and other unique FRAD devices. Depending on the

origin of the IP datagram - either from the IP layer or from the FRAD layer (after it has been decapsulated from the FR frame) - the datagram is either forwarded to the FRAD to be sent to the network or decapsulated and handed over to the higher layers. Regarding the Transport Adaptation Layer (TPAL) in figure 4.10, The TPAL presents a basic, uniform interface between applications and transport layer models. This step in the process is only included for simulation purposes and would not exist in a real world device.

External to the FRAD, the end users will connect to a FR router. Since the video signal will be terminated at this router, it will utilize a second routing table which is flexible in terms of being interactive with the end user. An end user "signals" the router to which video channel they require and this will be inserted into the routing table. As a result, the router will not be required to multiplex all video channels to all outgoing ports, rather it performs an intelligent multiplex. Since each video frame contains multiple video channels in a predetermined pattern the FRAD can demultiplex using simple timing schemes. Figure 4.11 illustrates how this will be accomplished through a simple example of six end users and five channels.




Figure 4.11 Intelligent Distribution of a Video Signal

This approach appears to be the most feasible in terms of throughput, while the task of preparing the data for use within an IP-based application is left to the end users' device. If it were left up to the FR router to perform both routing and multiplexing, in addition to frame-to-IP decapsulation of the video signal, it would quickly become a potential bottleneck and delay source. Given today's high processing power of home devices and that the end user is only interested in a small percentage of the data running over the FR router (one video channel vice up to 50 channels

at the router), the onus is placed on the home device and circumvents the potential backlog at the FR router. This concept can easily be proven when using sophisticated voice compression algorithms without noting any delay introduced by the decompression at the end users site.

An alternative, and to add more flexibility for the end user to choose devices, one could implement a FRAD externally to the home. In this case, users of one group (e.g., neighborhood street) connect to a FRAD in a similar to fashion to an Ethernet. However, depending on the volume of the video traffic, this might cause a significant delay at the FRAD. A second argument against placing the FRAD external to the home concerns the use of SVCs. Recall that when a SVC is established, the user specifies the destination to which he wishes to connect and a SVC is established only for the duration of the connection. The added benefit of SVCs is that no network resources (except an address) are used when there is "silence" on the line. Having an external FRAD responsible for establishing SVCs would impose even more load to those devices, consequently increasing the chances of bottlenecks right at access to the network.

Summarizing the local access portion of the UTN, utilizing end user devices that serve as the FRAD is the

most appropriate solution since it delegates processing tasks, such as frame-to-IP conversions, to the devices with the least amount of networking tasks. In addition, the FRAD is capable of responding to heavy traffic loads by varying the queue depths before network congestion even occurs [Ref. 28]. By placing the FRAD internal to the home or end user, it also falls in line with one of the basic principles to FR, pushing the data-conversion to a higher level, in turn reducing the load on the network nodes like switches or routers.

THIS PAGE INTENTIONALLY LEFT BLANK

V. ALTERNATIVE TECHNOLOGIES

To state that the UTN is the ultimate solution to converging networks would be a misnomer. There are a number of areas of research currently being conducted that also address the problem. Some of these were kept in mind in the design of the UTN so as to facilitate an easy migration to them. In any event, it is important to address a few of the major areas where ongoing research is being conducted, specifically; IP directly over fiber, Photonic Routing, Terabit Links, and Gigabit Ethernet.

A. IP OVER FIBER

It is universally accepted that IP packets have become the Layer 3 Protocol Data Unit (PDU) of choice, thus eliminating the need for multi-protocol routers and multi-protocol networks. The Internet, as it presents itself today, is a single-protocol network, and the protocol used is IP. It is also true that fiber links are the physical connections of choice. Consequently, one of the main questions that must be addressed when designing a network is what is the best way, in terms of cost efficiency, for fiber networks to be used to connect IP routers? There is not a simple answer. Currently it is not feasible to place IP-

packets directly onto a fiber link, simply because of the absence of a "start of packet" indicator in the IP packet header and the lack of standards that define what happens on the line between IP packets.

IP does not provide delineation or framing tools, hence it is impossible for a receiver or intermediate node to determine whether there is a sufficient buffer available to process an incoming packet. This problem could be overcome by adding start-end indicators to the IP-header, but the problem addressed above is of more significance. In contrast to an ATM cell, which is filled with an idle pattern when no information is carried, an IP packet always contains information since it is only created when information needs to be transmitted. When nothing has to be sent, nothing happens on the carrier. This idling time could easily be confused with a line failure. One can solve this issue by adding idle patterns onto the carrier in order to keep the line busy, but since these idle patterns are also just a series of zeros and ones, a bit pattern equal to the idle pattern could very well appear in the middle of an IP-payload. IP simply does not provide the necessary transparency to inform the receiver whether the pattern constitutes an idle pattern or useful information.

The issues of packet length, idle pattern, and transparency have all been solved at the data link layer (layer 2 of the OSI model), but not at the IP packet layer (layer 3). Consequently, the layer 2 frames must be used in order to avoid these issues, since they always have a length code and always define a transparent inter-frame fill pattern.

So it appears in the meantime, that in order to use IP packets they must be placed into some form of frame before transmission over fiber. The main question is, what kind of frames and how many frame levels are needed? Figure 5.1 illustrates the different possibilities of layering between IP and fiber:




Figure 5.1 IP to Fiber Layering Methods

Since in the very early stages of SONET/SDH IP was not recognized as a standard protocol, the only way to carry IP packets inside a SONET/SDH frame was to encapsulate them in either a VT for low speeds, or ATM cells for high speeds. This method is commonly known as the legacy method. When Packet over SONET (POS) was defined it became possible to use Point-to-Point Protocol (PPP) frames to hold IP packets, thus eliminating the need for ATM cells. Recently, Gigabit Ethernet (GBE) arrived on the scene and gathered much attention. GBE is best envisioned as a series of GBE switches connected by fiber links. GBE usually utilizes wavelengths of 850 nm, while 1300 nm is possible as well. GBE frames can either be placed directly onto a fiber, or shifted onto DWDM Systems with Optical Add Drop Multiplexers (OADM) and optical cross-connects. However, if GBE frames are placed directly onto fiber, only one stream can be on the line at one time. By using DWDM and OADM optical networking, multiple GBE streams can be transmitted simultaneously on the same physical fiber.

However, now the question becomes why not simply implement the Optical Network method, thus eliminating GBE and SONET/SDH completely? There are quite a few arguments pointing toward a solution without SONET/SDH, subsequently without ATM. Once GBE is combined with DWDM and OADM, a GBE

switch (switching GBE frames) can send and receive multiple IP packet streams on separate wavelengths over the same fiber, which is sometimes called photonic networking or IP over lambda. Without a doubt, this will more than likely be the future of networking and SONET/SDH will give way to these newer methods. However, the current standards for GBE in conjunction with DWDM/OADM, are still being defined, whereas SONET/SDH standards are well defined and have been proven to be functional. Additionally, GBE is progressing towards the 10 Gigabit Ethernet and therefore the efforts are migrating to the new technology. This is one of the primary reasons why GBE was not selected for the UTN. Deploying a network based on technology, which might be replaced by a successor without ever reaching maturity, is not a viable solution.

While the authors agree that the future belongs to photonic networking, SONET/SDH will continue having a place as a legacy, or even consumer networking method. Implementing a SONET/SDH-based network (like the UTN) "now", will not be a financial risk, since the bottom line remains that the fiber and equipment can be reused - at least partially - in a pure photonic network. In contrast, implementing a pure photonic network at this time would constitute a high financial uncertainty. A technology still

in the developmental stages, such as GBE, changes so rapidly that it is impossible to include those improvements in a large network like the UTN.

As alluded to in the previous paragraphs, the future of large networks appears to be one based on all-optical technology, perhaps utilizing DWDM. While DWDM will not provide an infinite amount of bandwidth, it will certainly provide enough bandwidth so that bandwidth will no longer be considered the limiting factor to network design. To be sure, there will be applications in the future that consume enormous amounts of bandwidth, however for the foreseeable future it is safe to assume that DWDM offers nearly unlimited bandwidth. Therefore, to develop networks and network devices with merely unlimited bandwidth, the focus of attention must shift from queuing theory and techniques to administer the lack of throughput. Instead, attention must be placed on the devices (e.g., switches and routers) and medium (fiber) themselves. Consequently, the next two sections describe two major research areas that are currently addressing this problem.

B. PHOTONIC ROUTING

If the assumption is that future networks will run IP over DWDM over fiber, this raises the question of how to

eliminate the requirement for first changing the incoming photons to an electrical signal before they can be routed and reverted back into an optical signal? To perform this function an optical cross-connect will be required. An optical cross-connect shuffles the wavelengths and fibers by configuration, not packet-by-packet or cell-by-cell based on header content as true routing or switching does.

Since all data is held within a light beam, then routing the beam means switching it from one physical fiber to another. Consequently, the light beam has to be absorbed from one line and inserted into another line, or in other words, the light beam has to change direction at some time. Analyzing the physics of light, an intuitive way to redirect light is to take advantage of its reflection characteristics. Since the beam must not lose any photons, the reflection must be total, thus mirrors would be a possibility. Light entering an input port could be reflected by mirrors so as to leave on a different output port. The mirrors will be positioned depending on the magnitude of the direction change.

Once it is determined how the direction can be changed, the question remains on how the cross-connect knows where to route the light and consequently how to position the mirrors? This information is contained in the header, so the

header needs to be analyzed in order to find the intended destination. Currently the only way to analyze the data contained in the header is to interpret them bit by bit. However, since computers only work with electrical signals, at a minimum the header must be changed into electricity. Figure 5.2 shows the basic principle of an optical router.




Figure 5.2 Optical Routing Mechanism

In order to achieve this optical routing, the following steps must be performed successfully. Since the incoming fiber may host many different wavelengths, they need to be isolated first before they can be analyzed. The analysis is very simple. First a header needs to be detected. The header could contain special optical patterns, such as a delimiter, which makes it distinct and easily recognized. The contents

of the header must be copied into an electrical signal by means of a standard optical receiver in order to make it readable for a computer. The optical packet is fed into a delay loop, since the header analysis and subsequent mirror positioning takes up more time than the optical signal needs to travel through the cross-connect. The analysis of the header follows the same rules as in conventional routers or switches. Basically, a routing table look-up needs to be performed to determine the appropriate output port for the optical packet and the mirrors need to be adjusted. The mirrors are etched onto a silicon substrate and the adjustment is done electrically to direct the optical signal from the delay loop to the output port.

While this method of routing still needs to convert parts of the optical packet (namely the header) into electricity, the routing itself is photonic routing. However, to perform true photonic routing without any conversion of light into electricity would require true optical computers - a technique that is not yet available. There are routers that exist today which are already capable of routing Terabits per second, however they are not pure optical routers and are usually formed by "clusters" of routers - typically between 14 and 128 routers appearing as one operating unit [Ref. 29]. This leads to the second major

area of research being conducted to handle enormous amounts of data.

C. TERABIT LINKS

Although DWDM seems to be the technology of the future, it is not an end all solution. At present, the maximum carrying capacity of a commercially available DWDM system is 1.6 Tbps, the rate touted by Nortel Networks for its OPTera system. Current-generation DWDM systems use optical amplification to extend the distance between points of optoelectronic regeneration of the signal, usually maxing out at 124 to 248 miles. A new technology - Soliton - is currently being developed to enhance the capacity envelope of DWDM. In a lab demonstration a French company called Algety showed that - using Soliton - it is possible to deliver 51 wavelengths at speeds of 20 Gbps each, achieving an aggregate throughput of 1 Terabit per second over a distance of 620 miles [Ref. 30].

Solitons are a form of unusually long-lasting standing waves formed by the combination of two or more separate waves under certain conditions. As discussed in chapter 3.3, light pulses sent down a fiber spread out in length as they propagate. This spreading is called dispersion, and its amount is wavelength dependent. By creating pulses in a

special shape related to the reciprocal of the hyperbolic cosine, dispersion effects can be minimized. Essentially this is done by inserting laser pulses onto a fiber with intentional nonlinearities. The soliton can be broken down into its original constituent waves by an interfering wave at the receiving end, thereby allowing the originating signals to be recaptured. However, the precise replication of the original signals is quite difficult. Solitons keep their pulse shape over very long distances without dispersion and minimal loss in coherency.

In summary, bandwidth and segmentation length are for all intensive purposes, no longer a limiting factor to network design. With the advent of photonic routing and Solitons, networks based on today's technology will change from best administering the lack of throughput, to choosing the physical media and the active components that best fits the needs for the actual purpose of the network. Design will become much easier in terms of bandwidth concerns, but on the other hand more difficult with respect to over sizing networks. This will have a profound affect on the requirements analysis, as the design phase must be done backwards. Soon the question will be how much bandwidth can the network consume at a maximum, followed immediately by how to limit costs in order not to not oversize the network

by magnitudes of three or more? The issue will no longer be on whether to use ATM or FR.

D. GIGABIT ETHERNET

If the UTN was not designed from the ground up, but built on an existing Ethernet architecture, the technology of choice would most likely be Gigabit Ethernet (GBE). GBE follows the same form, fit and function as its 10 Mbps and 100 Mbps Ethernet predecessors, thus a migration to higher-speed networking would be possible without disrupting network operation. All three Ethernet speeds use the same IEEE 802.3 frame format (figure 5.3), full-duplex operation and flow control methods. In half-duplex mode, GBE facilitates the same fundamental CSMA/CD access method to resolve contention for a shared media.




Figure 5.3 802.3 Frame Format

There are only two single-LAN GBE physical topologies, either two stations connected by a single point-to-point link utilizing half or full-duplex mode, or a group of stations (in half-duplex mode) connected to a single repeater through individual point-to-point links [Ref. 31].

GBE does not support a shared-bus medium, such as coaxial cables in full-duplex mode. However, since one of the key concepts for the UTN is the use of fiber optics, this discussion does not include shared media issues, neither is half-duplex an option for the UTN.

GBE can run over fiber optic lines and achieves substantial data rates, at least when compared to conventional Ethernet and Fast Ethernet. However, GBE was chiefly developed to replace existing Ethernets, which as a maximum are deployed in campus sizes. Essentially, the fundamental focus of GBE is to provide one Gbps bandwidth for campus networks with the simplicity of Ethernet [Ref. 32]. The length restrictions imposed on GBE, as shown in figure 5.4 [Ref. 33], was one of the primary reasons why GBE is not used within the UTN.

NAME	Wavelength	Mode	Diameter [μ m]	Distance
1000 BASE-SX	770-860 nm	Multi	62.6	≤ 275 m
1000 BASE-SX	770-860 nm	Multi	50	≤ 550 m
1000 BASE-LX	1270-1355 nm	Multi	50 or 62.5	≤ 550 m
1000 BASE-LX	1270-1355 nm	Single	10	≤ 5000 m

Figure 5.4 Gigabit Ethernet Maximum Distance Supported

GBE within the UTN would greatly affect the design flexibility in terms of maximum link length. It would

introduce restrictions and minimize the advantages of using fiber as the medium. Moreover, GBE runs at a maximum speed of one Gbps, a performance that wastes the true capabilities of fiber optical lines. Currently, the 10 Gigabit Ethernet (10GBE) standard is being developed, however the specifications are not mature enough to attempt to design a network based on this new technology. As a result, it is also too early to prove or disprove, if an easy migration from GBE to 10GBE is even possible. In addition, given the fact that 10GBE will soon be a reality constitutes a threat to GBE. One can expect that vendors will soon focus their attention on the new technology and cease supporting the old.

Of more importance is an issue not even addressed in the original GBE specifications, QoS. For example the Resource Reservation Protocol (RSVP). RSVP is a method by which to request and provide quality of network connections. In order to use RSVP, each network component in the link between client and server must support RSVP before meaningful results can be achieved. In a GBE network, RSVP would have to be added to each individual component. In contrast, the technology chosen for the UTN natively supports QoS issues, without the need to add another protocol like RSVP.

Another issue regarding GBE concerns IP multicasting. As addressed in Chapter IV, one of the main design principles of the UTN is that no two circuits shall contain duplicate information. If GBE were used in the UTN, it would be very difficult to prevent duplicate information from being transmitted. This would have a profound effect, especially in the case of video data. The problem is GBE (or any IP routed network), lacks the functionality to perform IP multicasting.

The routing of multicast packets within a large IP network is difficult. Conceptually, it would appear to be a simple task. A transmitter sends data with a dedicated multicast IP (the IP addresses 224.0.0.0 to 239.255.255.255), while the receiver notifies the transmitter as whether or not it is interested in receiving this multicast IP. However, the problem is whom must the receiver notify?

There are two options; a client could notify the original transmitter (e.g., a video server) or its local router. For the first option, the transmitter must know how many and which receivers are listening for the IP multicast. However, the problem is that the receiver must know the exact address of the transmitter. This goes contrary to the concept of IP multicasting, because the receiver should only

be required to choose the appropriate multicast group (e.g., the desired TV channel), not the transmitter that broadcasts this channel. Moreover, by requiring the receiver to select a specific server, a large scalability problem arises for the server itself. For example, if there were just a thousand TV viewers changing the channel frequently, the server would be required to use a large portion of its available bandwidth and processing power simply for administrative tasks. The other solution, where the client notifies its local router to join the appropriate multicast group, simply shifts the problem to an alternate computer. In this case the router adjacent to the transmitter must keep track of all IP multicast data.

It is important to note that there exist some solutions to the problem of IP multicasting, such as the Distance Vector Multicast Routing Protocol (DVMRP) and Multicast Open Shortest Path First (MOSPF). However, these solutions depend on specific unicast routing protocols. To avoid this dependency, the Protocol Independent Multicast (PIM) has been developed, which basically operates similar to the Core Based Tree (CBT) protocol. CBT implements a core router, where all traffic runs through, whereas PIM uses rendezvous points with similar functionality. Without addressing further details of these protocols, it is a fact that none

of the currently proposed solutions overcome the imminent scalability problems at the core router or the rendezvous points [Ref. 34]. A variation to PIM, PIM Source Only, which requires notification directly at the server, might be able to better support video over IP, but it works only with version 3.0 of the Internet Group Management Protocol (IGMP 3), and has not yet been implemented.

In summary, GBE is an alternative to the technology chosen for the UTN. However, due to its length restrictions it could only be deployed at Level 1 or Level 2 of the UTN and would add significant complexity in the transformation from Ethernet frames to ATM cells. Also, as currently proposed, GBE does not offer a migration path to DWDM, like SONET/SDH. In addition, once applications running over the network overtake the capacity of GBE, a significant financial burden would be required to increase the available throughput. Last, until the issue of IP multicasting is solved, little can be said about the effectiveness of video over IP running over GBE.

THIS PAGE INTENTIONALLY LEFT BLANK

VI. CONCLUSION

The aim of this thesis was to describe a network capable of handling three types of traffic, namely voice, video, and data, and their associated QoS requirements. Through careful analysis and based on specifications of the various technologies, a mixture of ATM and FR equipment operating over SONET, meets this goal. The authors acknowledge a more detailed analysis (e.g., analytical model and simulation) of the network implementation is in order, however this was not the intention. The approach has been from a conceptual level and serves as an underpinning from which the UTN could be built. In spite of not having a simulated model, there are two areas where possible bottlenecks could occur in the UTN. The first is in the FR switches at Level Zero. If these switches are incapable of discarding DLCIs not destined for their table fast enough, congestion will most certainly occur in these nodes. A second area where the network could have faults is in the PON device. A generous assumption is being made that vendors of this equipment will indeed follow the specifications outlined for their manufacturing. If the multiplexing or regenerating of all incoming traffic (composed of voice, video and data) is not performed rapidly and accurately, the

network could become bogged down with error congestion traffic.

In conclusion, the UTN proposed in this thesis is a viable solution to the QoS requirements outlined in Chapter II. Based on present day technology, a hybrid-network consisting of ATM and FR connected via fiber strands, best suits the abstract model. In addition, with the foundation of SONET as the transmission facility - rather than solely on WDM or DWDM - the UTN provides the necessary bandwidth to meet all requirements. Granted emerging technologies, similar to those discussed in Chapter V, may quickly make the UTN obsolete, this is beside the point. The UTN would easily position designers for future implementations, regardless of the protocol, application or technology selected.

LIST OF REFERENCES

1. Cisco Systems, Inc., *CISCO CCIE Fundamentals: Network Design and Case Studies*, p. 484, Macmillan Technical Publishing, 1998.
2. Cisco Systems, Inc., *CISCO CCIE Fundamentals: Network Design and Case Studies*, p. 475, Macmillan Technical Publishing, 1998.
3. Davie, L. & Peterson, D., *Computer Networks*, p. 552, Morgan Kaufmann, 2000.
4. Rowe, M., *Measure of VoIP Networks for Jitter and Loss*, Test & Measurement World, December 1999, Cahners Business Information.
5. Rech, J., *Dehnbare Netze*, c't 2000, Issue 3 January 2000, p. 242, Heise Verlag.
6. Tanenbaum, A. S., *Computer Networks*, p. 92, Prentice Hall PTR, 1996.
7. Muller, N. J., *Desktop Encyclopedia of Voice and Data Networking*, p. 189, McGraw-Hill, 2000.
8. Rech, J., *Glas statt Kupfer*, c't 2000, Issue 10 May 2000, p. 255, Heise Verlag.
9. Gigabit Ethernet Alliance, *Gigabit Ethernet: Accelerating the Standard for Speed*, 1998.
10. Bates, R. J., *Broadband Telecommunications Handbook*, p. 478, McGraw-Hill, 2000.
11. Russell, T., *Telecommunications Protocols*, p. 70, McGraw-Hill, 2000.

12. Goralski, W. J., *SONET*, pp. 46-47, McGraw-Hill, 2000.
13. Goralski, W. J., *SONET*, p. 536, McGraw-Hill, 2000.
14. Goralski, W. J., *SONET*, pp. 166-167, McGraw-Hill, 2000.
15. Goralski, W. J., *SONET*, p. 212, McGraw-Hill, 2000.
16. Feit, S., *Wide Area High Speed Networks*, pp.278-279, Macmillan Technical Publishing, 1999.
17. Bates, R. J., *Broadband Telecommunications Handbook*, p. 155, McGraw-Hill, 2000.
18. Frame Relay Forum, *The Basic Guide to Frame Relay Networking*, p. 12, Frame Relay Forum, 1998.
19. The ATM Forum Technical Committee, *2.4 Gbps Physical Layer Specification*, AF-PHY-0133.000, ATM Forum, 1999.
20. Kercheval, B., *TCP/IP Over ATM*, p. 171, Prentice Hall PTR, 1998.
21. Feit, S., *Wide Area High Speed Networks*, p. 310, Macmillan Technical Publishing, 1999.
22. Feit, S., *Wide Area High Speed Networks*, pp. 294-295, Macmillan Technical Publishing, 1999.
23. The ATM Forum Technical Committee, *Frame Based User-to-Network Interface (FUNI) Specification v2.0*, AF-SAA-0088.000, ATM Forum, 1997.
24. Bates, R. J., *Broadband Telecommunications Handbook*, p. 208, McGraw-Hill, 2000.

25. ITU-T, *G.983 High Speed Optical Access Systems Based on Passive Optical Network (PON) Techniques*, ITU-Telecommunications Standard Sector, 1998.
26. The Frame Relay Forum, *Frame Relay/ATM PVC Network Interworking Implementation Agreement*, The Frame Relay Forum Document Number FRF.5, 1994.
27. Feit, S., *Wide Area High Speed Networks*, p. 151, Macmillan Technical Publishing, 1999.
28. Muller, N. J., *Desktop Encyclopedia of Voice and Data Networking*, p. 668, McGraw-Hill, 2000.
29. Goralski, W. J., *SONET*, pp. 540-541, McGraw-Hill, 2000
30. Dawson, F., *Light Moves Generate New Optical Heat*, Inter@ctive Week, September 1999.
31. Seifer, R., *Gigabit Ethernet*, p. 223, Addison Wesley, 1998.
32. Gigabit Ethernet Alliance, *Gigabit Ethernet: Accelerating the Standard for Speed*, 1998.
33. Seifer, R., *Gigabit Ethernet*, p. 246, Addison Wesley, 1998.
34. Leitner, F., *Bilder für Alle*, c't 2000, Issue 12 May 2000, p. 219, Heise Verlag.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center.....2
8725 John J. Kingman Rd., Suite 0944
Ft. Belvoir, VA 22060-6218
2. Dudley Knox Library.....2
Naval Postgraduate School
411 Dyer Rd.
Monterey, CA 93943-5101
3. Prof. Dan C. Boger, Code CS.....1
Chairman, Department of Computer Science
Naval Postgraduate School
Monterey, CA 93943-5106
4. Prof. Gilbert M. Lundy, Code CS/Ln.....1
Department of Computer Science
Naval Postgraduate School
Monterey, CA 93943-5100
5. Prof. Wolfgang Baer, Code CS/Ba.....1
Department of Computer Science
Naval Postgraduate School
Monterey, CA 93943-5118
6. John Hunter.....3
1632 Montmorency Drive
Vienna, VA 22182
7. Holger Großmann.....3
Treeneblick 7
24852 Sollerup
Germany