

Air Force Institute of Technology

AFIT Scholar

Theses and Dissertations

Student Graduate Works

3-1999

Analysis of USAF Battlelabs' Communication Infrastructure Over Various Protocols

Rocky A. Favorito

David M. McIllece

Follow this and additional works at: <https://scholar.afit.edu/etd>



Part of the [Computer Sciences Commons](#)

Recommended Citation

Favorito, Rocky A. and McIllece, David M., "Analysis of USAF Battlelabs' Communication Infrastructure Over Various Protocols" (1999). *Theses and Dissertations*. 5151.

<https://scholar.afit.edu/etd/5151>

This Thesis is brought to you for free and open access by the Student Graduate Works at AFIT Scholar. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of AFIT Scholar. For more information, please contact AFIT.ENWL.Repository@us.af.mil.

AFIT/GCS/ENG/99M-04

ANALYSIS OF USAF BATTLELABS'
COMMUNICATION INFRASTRUCTURE
OVER VARIOUS PROTOCOLS

THESIS

ROCKY A. FAVORITO
CAPTAIN, USAF

DAVID M. MCILLECE
1ST LT, USAF

Approved for public release; distribution unlimited

The views expressed in this document are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

AFIT/GCS/ENG/99M-04

Analysis of USAF Battlelabs'
Communication Infrastructure
Over Various Protocols

THESIS

Presented to the faculty of the Graduate School of Engineering
of the Air Force Institute of Technology
Air University
In Partial Fulfillment of the
Requirements for the Degree of
Master of Science (Electrical Engineering)

Rocky A. Favorito
Captain, USAF

David M. McIllece
1st LT, USAF

March, 1999

Approved for public release; distribution unlimited

This page intentionally left blank.

AFIT/GCS/ENG/99M-04


Analysis of USAF Battlelabs'
Communication Infrastructure
Over Various Protocols

THESIS

Presented to the faculty of the Graduate School of Engineering
of the Air Force Institute of Technology
Air University
In Partial Fulfillment of the
Requirements for the Degree of
Master of Science in Electrical Engineering

Rocky A. Favorito
Captain, USAF

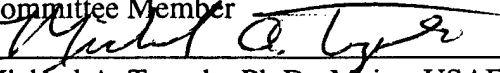
David M. McIllece
1st LT, USAF



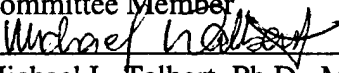
Richard A. Raines, Ph.D., Major, USAF
Committee Chairman



Stuart C. Kramer, Ph.D., Lieutenant Colonel, USAF
Committee Member



Michael A. Temple, Ph.D., Major, USAF
Committee Member



Michael L. Talbert, Ph.D., Major, USAF
Committee Member

Approved for public release; distribution unlimited

ACKNOWLEDGEMENTS

First of all, I would like to acknowledge the big guy upstairs for making this all possible. I would also like to thank my family and friends for being there for me during these times of distress and confusion. Also, I would like to thank my thesis partner and friend, Dave McIllece for his support, insight, cooperation, and hard work on this effort. Next, I would like to thank my advisor and mentor, Major Raines for his guidance, support, mentorship, empathy, and sense of humor throughout this lengthy, and taxing process. I would also like to thank Major Talbert for his support and encouragement. Finally, I am indebted to our priceless system administrator, Dave Doak, for his never-ending systems support.

Capt. Rocky A. Favorito

My heartfelt appreciation and gratitude goes out to those without whom I most certainly could not have completed this thesis. First I thank my Father in Heaven and His Son, my Savior, for their loving guidance and answers to many prayers. Next must come my lovely wife, Angela, for her unfailing love and unwavering support. Without her, this accomplishment would mean very little. I also express my thanks to my advisor, Maj. Richard Raines, for without his guidance and direction this work would have never materialized. I also want to thank Capt. Rocky Favorito for his insights, his commitment to rigor, and most of all his patience! I have learned much more with him than had I gone it alone. I am also indebted to Maj. Michael Temple, who taught me the finer points of Communication Theory, and Lt. Col. Stuart Kramer, who brought a sense of order to the increasingly complex world of engineering through which I wander.

Lt. David M. McIllece

TABLE OF CONTENTS

Acknowledgements	vi
Table of Contents	vii
List of Figures	ix
List of Tables.....	xi
Abstract	xiv
Chapter 1 Introduction	1
1.1 Overview	1
1.2 Research Goal	1
1.3 Summary	2
Chapter 2 Literature Review	3
2.1 Introduction	3
2.2 Air Force Battlelabs	3
2.3 Defense Information System Network.....	4
2.4 OSI Reference Model.....	7
2.5 Virtual Private Networks.....	9
2.6 Multimedia	12
2.7 Data Compression	17
2.8 Traffic Modeling	20
2.9 ISDN and ATM.....	29
2.10 ATM Implementations	38
2.11 Summary	47
Chapter 3 Methodology.....	49

3.1 Methodology Overview.....	49
3.2 Problem Definition.....	49
3.3 Value System Definition	53
3.4 System Synthesis.....	55
3.5 Model Definition	57
3.6 Summary	71
Chapter 4 Analysis	72
4.1 Analysis Overview	72
4.2 System Evaluation.....	72
4.3 Decision Analysis.....	88
4.4 Summary	101
Chapter 5 Conclusions	103
5.1 Goal Restatement	103
5.2 Conclusions	103
5.3 Contributions.....	104
5.4 Suggestions for Future Work	104
Appendix A: Initial Confidence/Error Analysis.....	103
Appendix B: Raw Data	111
Appendix C: Final Confidence/Error Analysis	119
Bibliography.....	120
Vita.....	123

LIST OF FIGURES

Figure 2-1 DISN CONUS Architecture [DIS98b]	6
Figure 2-2 OSI Reference Model	8
Figure 2-3 OSI Encapsulation	9
Figure 2-4 Peer Network Layer VPN	11
Figure 2-5 Overlay Network Layer VPN	11
Figure 2-6 Link Layer VPN	12
Figure 2-7 IFP Source	22
Figure 2-8 Two-State Markov Chain	23
Figure 2-9 Pictorial "Proof": Self-Similarity	24
Figure 2-10 Group of Pictures	28
Figure 2-11 B-ISDN Protocol Architecture [RIL97]	30
Figure 2-12 STS-N Basic Frame Format [McD95]	33
Figure 2-13 VPC and VCC Operation	35
Figure 2-14 ATM Protocol Stack	37
Figure 2-15 Native vs. Conventional ATM Use	38
Figure 2-16 Native ATM Protocol Structure	39
Figure 2-17 IP over ATM Protocol Structure	42
Figure 2-18 LANE vs. Native ATM	43
Figure 2-19 LANE Protocol Structure	44
Figure 2-20 LANE Operation	46
Figure 3-1 System vs. Environment	53
Figure 3-2 Value System Hierarchy	55

Figure 3-3 DISN ATM Backbone Network.....	58
Figure 3-4 Virtual Battlelab Environment	59
Figure 3-5 Common Site Configuration.....	60
Figure 3-6 ESPDU Probability Density	62
Figure 3-7 VBR Video Packet Size PDF	64
Figure 3-8 Native ATM in OPNET.....	65
Figure 3-9 IP over ATM in OPNET.....	67
Figure 3-10 LANE in OPNET	68
Figure 4-1 Average Pareto Interarrival Time	75
Figure 4-2 Pictorial “Proof” of Self-Similarity (Time Traces)	76
Figure 4-3 Self Similar Distributions (Histograms).....	76
Figure 4-4 DIS Source: Average Packets/Second.....	77
Figure 4-5 VBR Video Source: Average Bytes/Second	78
Figure 4-6 LANE Required Runs.....	80
Figure 4-7 Average ETE Delay Performance Comparison.....	84
Figure 4-8 CDV Performance Comparison.....	85
Figure 4-9 Tightly-Couple Delay Comparison	86
Figure 4-10 Bandwidth Utilized.....	87
Figure 4-11 Performance-Weighted Final Scores	99
Figure 4-12 Evenly Weighted Final Scores	100
Figure 4-13 Cost-Weighted Final Scores	101
Figure A-1 Confidence/Error Combinations for Native ATM.....	107
Figure A-2 Confidence/Error Combinations for IP over ATM.....	108
Figure A-3 Confidence/Error Combinations for LANE	110

LIST OF TABLES

Table 2-1 Air Force Battlelabs Descriptions.....	4
Table 2-2 Multimedia Data Types	15
Table 2-3 ESPDU Data Averages	26
Table 2-4 MPEG Frame Sizes.....	29
Table 2-5 SONET STS-N/OC-N Rates [McS95]	32
Table 3-1 DIS Traffic Properties.....	62
Table 3-2 VBR Video Traffic Properties	63
Table 3-3 System Utility Function	69
Table 4-1 Pareto Interarrivals Sample Set Data	75
Table 4-2 DIS Source Sample Set Data	78
Table 4-3 Pareto Interarrivals Sample Set Data	79
Table 4-4 Simulation Resource Requirements	80
Table 4-5 95%/10% Requirements	82
Table 4-6 ETE Delay Constraint.....	83
Table 4-7 MOP Results Summary	88
Table 4-8 System Utility Function	89
Table 4-9 Raw Score Matrix	89
Table 4-10 Confidence Matrix	91
Table 4-11 Weight Matrix.....	92
Table 4-12 Performance-Weighted Undiscounted Score Matrix	96
Table 4-13 Performance-Weighted Discounted Score Matrix	96

Table 4-14 Evenly-Weighted Undiscounted Score Matrix	96
Table 4-15 Evenly-Weighted Discounted Score Matrix	97
Table 4-16 Cost-Weighted Undiscounted Score Matrix	97
Table 4-17 Cost-Weighted Discounted Score Matrix	97
Table A-1 Results for Native ATM Sample Runs	106
Table A-2 Results of Native ATM Confidence Interval Analysis	106
Table A-3 Results for IP over ATM Sample Runs	107
Table A-4 Results of IP over ATM Confidence Interval Analysis	108
Table A-5 Results for LANE Sample Runs	109
Table A-6 Results of LANE Confidence Interval Analysis	109
Table B-1 ETE Delay (secs) Data for Native ATM.....	111
Table B-2 Cell Delay Variation (secs) Data for Native ATM	111
Table B-3 Cell Loss Ratios for Native ATM	112
Table B-4 Utilization (%) for Native ATM Links	112
Table B-5 ETE Delay (secs) Data for IP over ATM.....	113
Table B-6 Cell Delay Variation (secs) Data for IP over ATM	113
Table B-7 Cell Loss Ratios for IP over ATM	114
Table B-8 Utilization (%) for IP over ATM Links	114
Table B-9 ETE Delay (secs) Data for LANE.....	115
Table B-10 Cell Delay Variation (secs) Data for LANE	116
Table B-11 Cell Loss Ratios for LANE.....	117
Table B-12 Utilization (%) for LANE Links	118
Table C-1 Average ETE Delay Error Tolerance.....	119

Table C-2 Average CDV Error Tolerance	119
Table C-3 Prob(Delay < 100 msec) Error Tolerance	119
Table C-4 Bandwidth Utilization Error Tolerance.....	119

ABSTRACT

As a result of the doctrines exposed in the Air Force's "Global Engagement: A Vision for the 21st Century Air Force," the concept of a "Battlelab" was devised to support and develop the key warfighting capabilities needed in the next century. Although geographically separated, these labs must have the capability to exchange information with each other in a "Virtual Battlelab Environment" (VBE). Although many types of data will be exchanged, only variable-bit-rate (VBR) video and distributed interactive simulation (DIS) traffic are modeled in the VBE. The research described in this thesis utilizes a systems engineering approach to investigate the performance of wide-area networking technologies and recommend a connectivity solution. Three networking protocols were considered: Native ATM, IP over ATM, and LAN Emulation. In addition, three analyses were performed: performance-weighted, evenly-weighted, and cost-weighted. Native ATM was the recommended solution for the performance-weighted analysis. IP over ATM was recommended for the cost-weighted and evenly-weighted analyses, although by only slight margins over LANE. Finally, the merit of a systems engineering approach to assist in tradeoff and decision making analyses was also demonstrated.

CHAPTER 1

INTRODUCTION

1.1 Overview

As a result of the doctrines exposed in the Air Force's "Global Engagement: A Vision for the 21st Century Air Force," [Glo98] the concept of a "Battlelab" [Bat98] was devised to support and develop the key warfighting capabilities needed in the next century. The primary mission of these Battlelabs is to identify and evaluate innovative operational concepts that will advance the core competencies and warfighting capabilities of the Air Force.

Although geographically separated, these labs must have the capability to exchange information with each other at desired data rates. It is projected that all types of data will be exchanged, from voice communications and data transfer to real-time video and graphics-intensive distributed interactive simulations. The effectiveness and efficiency of this data communications capability will determine in large part the ability of the Battlelabs to achieve their mission.

1.2 Research Goal

The Battlelabs are currently investigating alternatives for network connectivity. Consequently, AFIT has been tasked by Headquarters United States Air Force, Directorate of Operational Requirements, Battlelab Integration Division (HQ USAF/XORB) to assist in this process by using modeling and simulation to evaluate the effectiveness of available networking technologies. The goal of this research is to

investigate the capabilities of wide-area networking technologies for handling a defined set of traffic loads between these geographically separated Battlelabs and recommend a solution based on the needs and desires of HQ USAF/XORB.

1.3 Summary

This chapter introduced the goal of this research, including a brief introduction of the Air Force's "Battlelab" concept and their computer communication needs. The remainder of this thesis is organized as follows. Chapter 2 presents a review of the concepts necessary to address the complex field of high-speed data communications, including key technologies such as Asynchronous Transfer Mode (ATM). Chapter 3 lays out a systems engineering methodology which is used to define the problem at hand, devise the value system by which proposed systems will be evaluated, synthesize possible solutions, and model those proposed solutions. Chapter 4 details the analysis portions of this process by first evaluating the proposed systems' outputs, which are used in the decision making process to arrive a recommended solution. Chapter 5 concludes this report by summarizing the results, and also by providing recommendations for follow-on work to this research.

CHAPTER 2

LITERATURE REVIEW

2.1 Introduction

This chapter gives the necessary theoretical background required to devise an effective connectivity architecture for high-speed wide-area network traffic. To do this, Section 2.2 begins with further information on the current Air Force Battlelabs. Section 0 contains background information regarding the Defense Information System Network (DISN). Then Section 2.4 gives a brief synopsis of the Open Systems Interconnection (OSI) reference model. In Section 2.5, concepts surrounding virtual private networking (VPN) are discussed. Multimedia issues are covered in Section 2.6, and Data Compression is discussed in Section 2.7. Section 2.8 introduces traffic modeling, followed by Section 2.9, which covers Integrated Services Digital Networks (ISDN) and asynchronous transfer mode (ATM). Options for implementing ATM are presented in Section 2.10, followed by a summary of the chapter in Section 2.11.

2.2 Air Force Battlelabs

In July of 1997, six Air Force Battlelabs were created. The primary mission of these Battlelabs is to readily identify innovative operational and logistical concepts to advance the Air Force's core competencies. In addition, they leverage existing Air Force resources and expertise to measure the potential worth of these concepts. Each Battlelab consists of a permanent staff of approximately 15-25 personnel [Bat98]. Table 2-1 contains a brief description of each of the six Air Force Battlelabs.

Table 2-1 Air Force Battlelabs Descriptions

Name	Location	Commanding Unit	Mission
Air Expeditionary Force Battlelab (AEFB)	Mountain Home AFB, ID	Air Combat Command	Identify innovative operational and logistics concepts for rapidly deploying, sustaining, and employing Air Expeditionary Forces, and to measure their potential for advancing the Air Force's core competencies and joint warfighting.
Command and Control Battlelab (C2B)	Hurlburt Field, FL	Air Combat Command	Identify innovative command and control and battle management operations and logistics concepts and measure their potential to advance the Air Force's core competencies and joint warfighting.
Force Protection Battlelab (FPB)	Lackland AFB, TX	Air Force Protection Group	Identify innovative concepts for protecting Air Force personnel, facilities, and weapons systems and to rapidly measure their potential worth using field ingenuity, modeling and simulation, and actual employment of exploratory capabilities in operational environments.
Information Warfare Battlelab (IWB)	Kelly AFB, TX	Air Intelligence Agency	Identify and rapidly measure the worth of innovative concepts which advance the Air Force's core competencies.
Space Battlelab (SB)	Schriever AFB, CO	Air Force Space Command	Identify innovative space operations and logistics concepts and measure their potential for advancing the core competencies of the Air Force as an integral component of joint warfighting.
Unmanned Air Vehicle Battlelab (UAVB)	Eglin AFB, FL	Air Combat Command	Identify innovative UAV operations and logistics concepts and measure their potential for advancing the core competencies of the Air Force.

2.3 Defense Information System Network

The Joint Chiefs of Staff has identified the need for an integrated information transmission system to support the 21st Century warfighter [JRO95]. This system must

replace stovepipe legacy communication systems with a seamless transport capability that can maintain pace with emerging technologies and meet the changing Command, Control, Communications, Computers and Intelligence (C4I) demands of the warfighter. This information transfer infrastructure is the Defense Information Systems Network (DISN).

The DISN provides the warfighter with a cost-effective, integrated services platform with built-in reliability (redundancy). Required services include: data, voice, imagery, and video. These services need to be provided in low, medium, and high bandwidth environments. Imagery services have been projected to dominate all other services by several orders of magnitude by 2010 [DIS96].

2.3.1 DISN Objectives

As we move into the 21st Century, the costs associated with maintaining legacy systems become harder to justify. The warfighter needs a global, scalable network that is capable of accepting technology insertions [JRO95]. The following are objectives cited by the Defense Information Systems Agency (DISA) for DISN [DIS98b]:

- Satisfy Warfighter Requirements
- Seamless Connections Between Deployed Forces/Home Bases
- Capitalize on Emerging Technologies
- Interoperable with All DOD Services, Government Agencies, and Allies
- Provide Needed Capacity/Connectivity for Warfighter
- Meet Surge Requirements Anytime/Anywhere
- Rapidly Reconfigurable to Meet Warfighter Needs
- Cost Effective Affordable Services
- Real-Time Management Capabilities (Peace and War)
- Cost Recovered through Effective Usage-Based Billing.

Another area of concern is that of Information Warfare (IW) against the United States Government, both from home and abroad. Specifically, protection is needed against attacks such as denial of service, unauthorized monitoring and disclosure of sensitive information, and unauthorized modification of networked databases and services. The importance of IW increases with the technological advances in the information realm.

2.3.2 DISN CONUS Configuration

To meet the needs of the stateside bases, DISA has developed the architecture shown in Figure 2-1 below for the DISN CONUS.




Figure 2-1 DISN CONUS Architecture [DIS98b]

The Synchronous Optical Network (SONET) backbone will serve as the physical layer medium, while future ATM installations plan to integrate value-added services, such as circuit-switched and data services, dedicated point-to-point services, and video teleconferencing [DIS98b].

2.4 OSI Reference Model

After the development of many incompatible data communication systems in the 1970s, the International Standards Organization (ISO) and the International Telecommunications Union – Telecommunications Sector (ITU-T) defined the concept of an “open system,” which is compatible with other open systems regardless of vendor type or technology [Sta96]. Although only a set of recommendations, the OSI model can be used to manage the complexity of a system under design, as well as provide a designer the ability to provide communications solutions at any given level of the model. It should be remembered that the OSI model only recommends standards concerning the transfer of information between layers, and not the internal workings of an open system.

2.4.1 OSI Layers

The OSI model consists of seven layers, as shown in Figure 2-2. The physical layer is concerned with the media used in transporting data between network nodes. Frequencies, voltage levels, and connector specifications are examples of physical layer characteristics. The data link layer is concerned with the flow control, bit framing, and error correction of information on a link between a transmitter and receiver node pair. The network layer provides communication between end applications through a network, and as such it is responsible for routing and relaying data over nodes and subnetworks if

necessary. The transport layer is added on top of the network layer to provide ETE error correction for data transport. Since information may have differing flow needs (one-way, alternating two way, simultaneous two-way), the session layer is placed above the transport layer to manage information flows. In addition, information often differs in its representation or format, and to make open system interoperable, data conversion is performed by the presentation layer. Finally, the application layer represents the user of the lower layer services.




Figure 2-2 OSI Reference Model

2.4.2 Layering and Encapsulation

The power and flexibility of the OSI model becomes apparent when its concept of operation is explained. First, network entities at a given layer must communicate according to the protocol of their layer. Also, communication only takes place between adjacent layers in the OSI model. Data from higher layers is encapsulated with protocol control information (PCI) and passed to lower layers as protocol data units (PDUs).

These lower layers treat the higher layer PDUs as service data unit (SDU) “payloads” (see Figure 2-3)[Sta96]. This allows entities to communicate independent of lower layer implementations. As long as a layer presents data to a lower or higher layer in the format it requires, the entities can communicate. Efficiency may vary depending on the implementation, but the open system allows different implementations to interoperate.

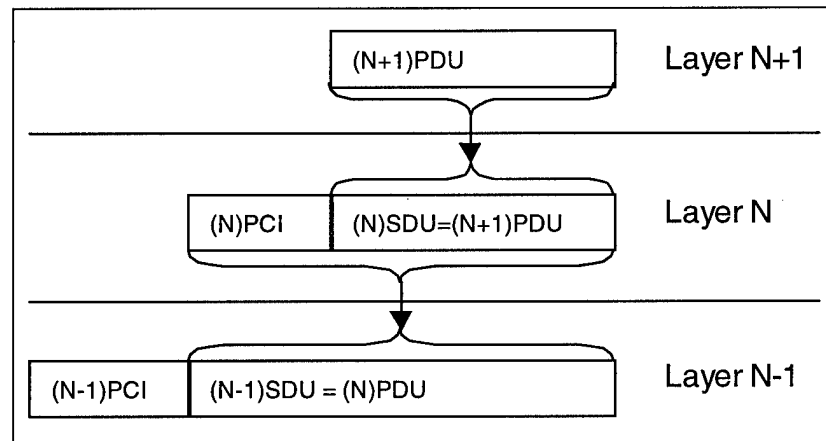


Figure 2-3 OSI Encapsulation

2.5 Virtual Private Networks

2.5.1 Private and Virtual Private Networking

Modern computer networking has its roots in telephony-based telecommunications, where bandwidth is typically allocated in fixed amounts using a form of time or frequency division multiplexing. It is perhaps due to this history that high speed, private wide area networking has traditionally taken the form of dedicated circuits between end stations. These end stations perform any switching necessary, and are responsible to maximally utilize the dedicated bandwidth resource. Such an approach

offers simple, reliable, and often secure transmission, but is not easily scaled and is costly--particularly if the links are underutilized [McS95].

The concept of a virtual private network (VPN) is to share resources in a public network on an on-demand basis, improving the utilization of existing resources and decreasing the costs of usage of those resources. Stated another way, a VPN is a private network constructed within a public network infrastructure [FeH98]. The “privacy” portion of a VPN depends upon encryption since the data is transmitted through the public network.

2.5.2 Types of VPNs

Depending on the needs of the user, a virtual private network can be implemented at a number of layers in the OSI reference model. In particular, VPNs can be implemented at the application, transport, network, or link layer.

2.5.2.1 Network Layer VPNs

If implemented at the network layer, VPNs are created by using the routing function of the network layer. They can take the form of either “peer” or “overlay” VPNs. A “peer” VPN is one in which the network layer routing path is computed on a hop-by-hop basis [FeH98]. Each node in the transit path is a peer with adjacent next-hop nodes, and at each node the data unit is raised to the network layer to determine the next hop in the VPN path. This is shown in Figure 2-4, where the dotted line represents the flow of data units in the VPN.

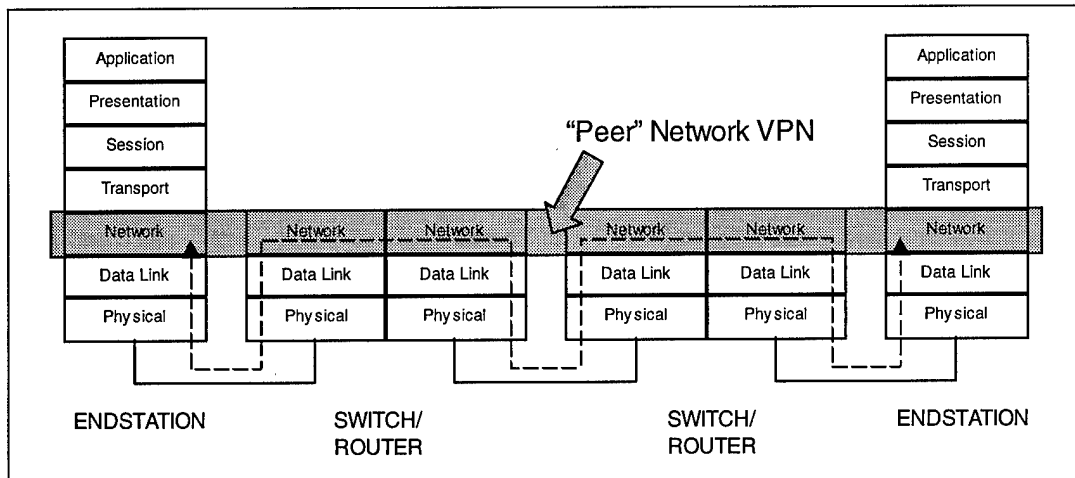


Figure 2-4 Peer Network Layer VPN

An “overlay” VPN uses the underlying link layer to “cut through” the network to another edge node, where operation returns to the network layer. This makes the nodes one hop away from each other at the network layer, as shown in Figure 2-5. Examples of this overlay version include Frame Relay, ATM, and tunneling implementations [FeH98]. VPNs can then be created as a collection of tunnels or virtual paths through the host network.

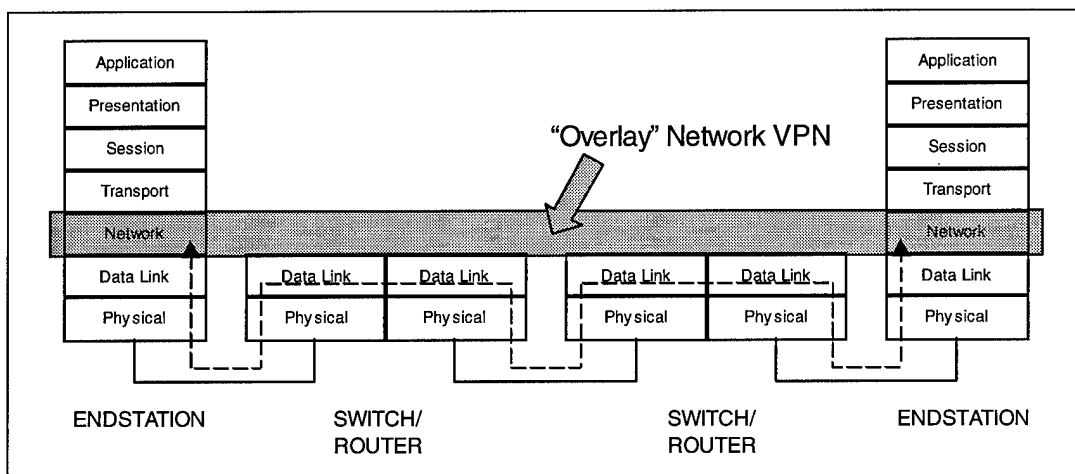


Figure 2-5 Overlay Network Layer VPN

2.5.2.2 Link Layer VPNs

Link layer VPNs are created by implementing a bridging as opposed to routing solution at the link level to provide connectivity. In this sense they are very similar to traditional private networks, which connect nodes with dedicated physical links as shown in Figure 2-6. Link layer VPNs however, have neither synchronized data clocks nor dedicated transmission paths like private networks. ATM and Frame Relay networks are often used as the link layer over which higher level protocols form VPNs.

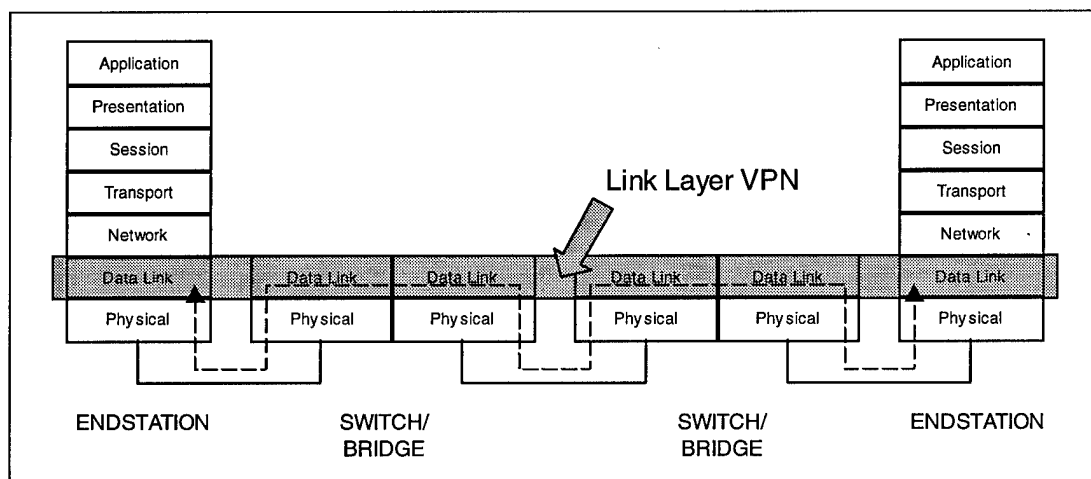


Figure 2-6 Link Layer VPN

2.5.2.3 Application and Transport Layer VPNs

Although not very common, VPNs can be formed at the application or transport layer by using encryption. Examples could include encrypted voice or data sessions between two or more end users.

2.6 Multimedia

So, what is multimedia? Unfortunately, there is no singularly agreed upon definition for multimedia. Most organizations define multimedia from their own

perspective to satisfy their respective interests. While this lack of standardization may raise immediate concerns (e.g. lack of interoperability), there is some good news. Most definitions for multimedia generally contain the same semantics. Furthermore, standards are in development (i.e. MPEG-7) to better define multimedia data as well as its content [WuI98a]. In general, *multimedia* (as its name suggests) is the combination of two or more data types into a single integrated delivery system. This system is typically interactive in nature.

The widespread use of multimedia data can be attributed to five events [Gib98]:

- Data compression research
- Human perception-based studies
- Standards
- Advances in computer processors
- Advances in communication networks.

2.6.1 Data Compression Research

Multimedia applications are reaping the benefits of approximately a quarter of a century of data compression research. Because of this research, it is possible to represent data more efficiently.

2.6.2 Human Perception-Based Studies

Many studies have been conducted to determine the inherent limitations of the human senses, namely the eyes and ears. The results of these studies have fueled several data compression techniques to obtain higher data compression rates. Therefore, one can represent multimedia data in considerably fewer bits with no detectable loss in quality.

2.6.3 Standards

In general, standards provide low-cost, high-quality products through competition, and enable interoperability in a multi-vendor, multi-platform, multi-protocol environment [WuI98b]. This definition also holds true for data compression standards such as MPEG and JPEG.

2.6.4 Advances in Computer Processors

Computer processor technology advances have brought multimedia technology to the Personal Computer (PC). These advances include higher frequencies, instruction set extensions to accommodate multimedia operations (i.e., MMX), and data buses with greater bandwidth. In addition, many PCs now come with graphic accelerators.

2.6.5 Advances in Communication Networks

The advances in PCs have pushed the bandwidth bottleneck to the networks. However, the combination of fiber optics and high-speed switching technologies promises to alleviate this bandwidth problem. Fiber optic cables provide high-capacity bandwidth with extremely low error-rates. This efficiency allows for more data to be sent with less overhead. However, to truly realize these increased data rates, faster switching devices are needed.

2.6.6 Multimedia Data Types

Multimedia data types fall into one of two categories with respect to time: *discrete* or *continuous*. Table 2-2 lists examples of each. These classifications, however, are not mutually exclusive. That is, one can convert a discrete type to make it continuous, and visa versa. For example, it is possible to create continuous data type

(animation) by time-sequencing closely related images at a rate of at least 16 images/second [Rag98]. Similarly, one can create an image (discrete type) by freezing a video frame (continuous type).

Table 2-2 Multimedia Data Types

Classification	Data Type
Discrete	Text Graphics Image
Continuous	Audio Video (image and audio) Generated Media (animation and music) Speech

2.6.7 Multimedia Data Characteristics

Multimedia data have the following characteristics:

- huge file sizes
- abstract data types
- spatial properties
- temporal properties
- inter-media synchronization.

2.6.7.1 Huge File Sizes

Modern advances in hardware have given computers the ability to display many colors. However, this ability comes with a significant cost. For example, a true-color Red Green Blue (RGB) pixel typically consumes 24 bits. Each individual color component requires eight bits to represent 256 possible shades. This allows for the creation of 2^{24} or 16,777,216 possible colors [Sal98]. Therefore, an image with a resolution of 640 x 480 pixels would require 921,600 bytes of storage. Now, extend this

to a 2-hour video with an associated data rate of 30 frames per second. This video would require approximately 199 Gigabytes of secondary storage or approximately 306,255 CD-ROMs. Even with the vast cost reductions for storage media, this requirement is still overwhelming, especially when this scenario includes a network connection, where bandwidth is a limited resource.

2.6.7.2 Abstract Data Types

Multimedia data such as audio and video are complex data types. They possess attributes that require non-standard data types. That is, they do not fall into the realm of applications satisfied by traditional alphanumeric data types (i.e. integer, string, etc). For instance, video is an aggregate object made up of image, audio, temporal properties, spatial properties, and synchronization properties, to name a few. It is very difficult to model and implement these attributes using relational modeling techniques. Therefore, one should employ object-oriented analysis and design techniques when modeling multimedia. This effort should also benefit from the emerging MPEG-7 standard, which specifies multimedia content necessary for the efficient search and access of the objects.

2.6.7.3 Spatial Properties

Image data has inherent spatial properties that must be modeled. Specifically, image data contain pixels, which have the following spatial dimensions:

- Vertical location
- Horizontal location
- Magnitude (defines color).

These values are stored to recreate or properly display image data. Certain compression techniques (i.e., JPEG and MPEG) eliminate the spatial redundancies in images and conserve storage and bandwidth resources.

2.6.7.4 Temporal Properties

Multimedia data types may also have temporal properties. Video frames, for example, not only have spatial dimensions, but also time components. In the United States, video applications typically require a frame rate of 30 frames/second, whereas voice is sampled at a rate of 8000 samples/second [Rag98]. These media are time-dependent, and thus continuous with respect to time.

2.6.7.5 Inter-media Synchronization

Consider video, which consists of audio and image data streams. To provide quality video, we need to synchronize these two streams. This means that we have to model, store, and process these inter-media temporal relationships. Synchronization across networks, where all users share and compete for bandwidth, may prove difficult to achieve. In fact, perfect synchronization is rarely achievable. However, a predefined Quality of Service (QoS) still needs to be satisfied.

The media considered in this thesis are Distributed Interactive Simulation and Variable Bit Rate (VBR) Video. More specifically, VBR MPEG video will be modeled and analyzed.

2.7 Data Compression

Data compression is the elimination of redundant information in data files. Put another way, it is the representation of a source in digital form with as few bits as

possible while maintaining an acceptable loss in quality [Gib98]. So, why is compression needed? As mentioned earlier, multimedia data files tend to be very large, requiring inordinate amounts of storage and bandwidth. Data compression allows for reduction of these requirements. In addition, data compression reduces the time and bandwidth needed to transmit and receive space-intensive files such as images, audio, and video over Wide Area Networks (WANs) (e.g., the Internet). Finally, storage and bandwidth resources are limited. As soon as their amounts are increased, new applications come along that drive even higher requirements.

The benefits of data compression are obvious in terms of increased storage requirements and bandwidth utilization. Nonetheless, these benefits come with associated costs (i.e., no free lunches here). A case in point, popular error-detection and correction codes such as Cyclic Redundancy Codes (CRC) require redundancy in data files to carry out their functions. Compression effectively removes this added protection, leaving them more susceptible to errors. This may be tolerable in a fiber optic environment, but this may not be acceptable in an already error-prone wireless environment. Therefore, it is important to exercise caution when compressing data. That is, one should scale the degree of compression to meet the needs of the particular application.

2.7.1 Intraframe Compression vs. Interframe Compression

Intraframe Compression reduces the amount of spatial redundancy in a frame. This method is employed in the JPEG and MPEG standards. This method allows for cleaner editing. On the other hand, *interframe compression* not only eliminates spatial redundancy within a given frame, but also the temporal redundancies between

consecutive frames. Only differences and motion vectors are sent to the decoder, allowing better compression ratios.

2.7.2 Lossless vs. Lossy Compression

When using lossless compression, the size of the file is reduced without losing any information. This type of compression is necessary for applications such as computer programs and medical x-rays, where the loss of a single bit may render the data worthless. However, this type of compression is not optimal, achieving typical compression ratios of approximately 2:1 [WuI98a].

Lossy compression methods, on the other hand, take advantage of the inherent limitation in the human eye to detect certain levels of intensity. They also capitalize on the shortcomings of the human ear to detect certain frequencies. Consequently, these techniques (if properly used) allow significantly more compression than lossless methods, with little or no detectable loss in quality.

2.7.3 Symmetrical vs. Asymmetrical Compression

Symmetrical compression involves a process in which the decoder experiences the same amount of work as the encoder. Usually, the decoder executes the same algorithm as the encoder, but in reverse. This type of compression lends itself well to real-time applications such as video conferencing. With *asymmetrical compression*, the coder performs significantly more work than the decoder. CD-ROM applications use this form of compression to expedite playback.

2.8 Traffic Modeling

Before discussing protocols used to implement a solution, traffic sources and source modeling should also be addressed. To obtain a useful evaluation of any particular protocol architecture, source traffic should be generated in as realistic a manner as possible.

2.8.1 Characterizing an Arrival Process

In traditional telecommunications, the models used to describe the arrival process of calls or connections were based on a Poisson distribution. Here, the arrivals are independent of one another and the time between arrivals is exponentially distributed. In the case of large public switched systems where the traffic sources were all of the same type (constant bit rate voice), the Poisson assumption was a good model of actual system behavior.

However, as telecommunications networks began to diversify and carry packetized information, the Bernoulli distribution was introduced to characterize the slotted, discrete nature of packet switching. The Bernoulli distribution is essentially a discrete-time version of the Poisson distribution, where the time axis of arrivals is divided into discrete slots. A packet or cell arrives in a slot with probability $p(<1)$, and is empty with probability $1-p$. Instead of being exponentially distributed, the time between arrivals is characterized by the geometric distribution and the number of arrivals in a unit of time is Bernoulli distributed. Again, as in the Poisson model, successive arrivals are independent of one another.

2.8.2 Early Burst Modeling

Early application of Poisson and Bernoulli models to emerging high-speed packet switching networks (such as ATM) did not capture the inherent bursty nature of modern traffic sources [PeE96]. Instead of constant or continuously varying data rates, most traffic sources (data, video, etc.) generate bits for an interval, and then sit idle. This burst and idle process varies widely according to the nature of the traffic source and its intended function.

To capture the alternating nature of bursts, an *on/off* distribution was first overlaid on the Poisson and Bernoulli distributions. This on/off or "interrupted" behavior is modeled in the Poisson case by generating arrivals in Poisson fashion during an on time, and then sitting idle during an off period. These periods are alternated continuously and are exponentially distributed. This method is referred to as an *interrupted Poisson process* (IPP) [PeE96]. Similarly, the *interrupted Bernoulli process* (IBP) alternates with Bernoulli distributed arrivals in the *on* periods, and alternates between on and off periods in a geometrically distributed fashion.

2.8.3 Burstiness and Correlation

These models improved in their modeling of bursty behavior, but lacked the notion of correlation in a burst. As the nature of bursty sources was further studied, the notion that successive arrivals were independent (uncorrelated) was not true to observed behavior. The nature of this inter- and intra-burst behavior is still an object of much research, but a number of ways have been devised to model such sources.

One such approach is to use a fluid model, where arrivals occur in continuous fashion during a burst. This model is referred to as an *on/off fluid source* or as an

interrupted fluid process (IFP), and is depicted in Figure 2-7. In addition, this model can take one of two refinements, the first being that of *independent burst periods*. The IFP captures the dependence of the arrivals within a burst, but the burst periods themselves may not be dependent. In this case, the burst periods could be Uniform, Poisson, or Bernoulli distributed, or any distribution that accurately models the independent, random nature of the inter-burst occurrence.

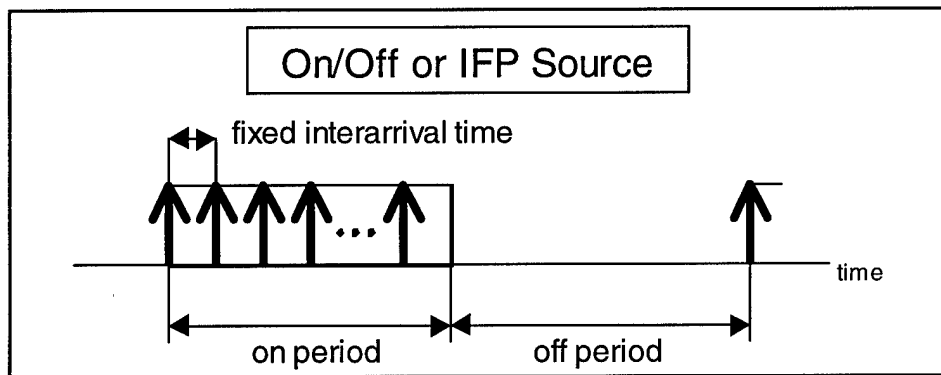


Figure 2-7 IFP Source

The second form is that of the *dependent*, generally distributed on/off source [ZhA94]. In this version, the source still generates data at a constant rate during the on period, but the duration of the *on* period is bounded by the length of a fixed frame. A frame length is one burst period plus its associated idle period. The off period then lasts for the duration of the fixed frame, after which the on period once again starts. In this manner, not only are the individual cells or packets in a burst dependent, but the on (burst) and off (idle) periods are also dependent. The on and off periods are related in that they sum to the length of the declared fixed frame. The length of any given burst (and hence the length of the off period) may follow any general distribution, depending on the behavior of the particular source.

More advanced models exist which attempt to capture the elements of correlation in bursty sources. Many of them use Markov chains where the rate of arrivals depends on which state in the Markov chains a source finds itself in. This is referred to as *Markov modulation* (MM), and many variants exist. They include a *Markov modulated Poisson process* (MMPP), a *Markov modulated Bernoulli process* (MMBP), or *Markov modulated Fluid process* (MMFP). In each case, the process can find itself in one of many different states, and in each, arrivals occur at a state-dependent rate according to their respective distribution [PeE96]. IPP/IBP/IFP are actually special cases of these MMPP/MMBP/MMFP models, where the Markov chain alternates between two states, on and off (see Figure 2-8) [ZhA94]. As the complexity of the distribution increases, so does the difficulty of incorporating it into the source model.

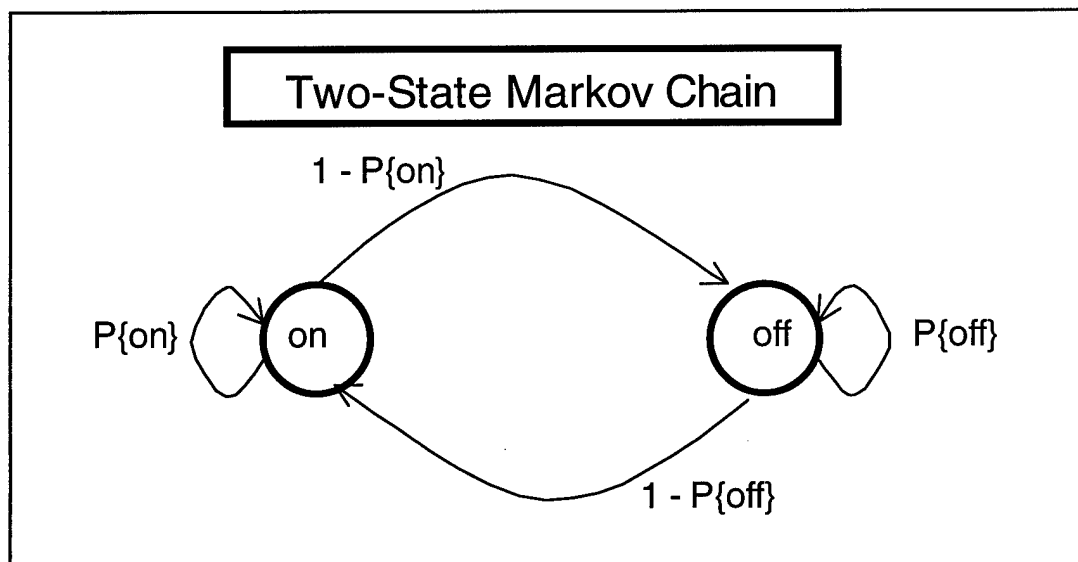


Figure 2-8 Two-State Markov Chain

2.8.4 Self-Similarity

Extending the notions of burstiness one step further, it has been observed that several forms of traffic exhibit *self-similar* qualities, simply defined as burstiness over a wide range of time scales [LeT95]. Also known as “fractal” processes, self-similar processes have no natural burst length, and when examined at short, intermediate, or long time scales, similar bursty qualities are manifest. Pictorial “proofs” can be used to verify self-similar processes by showing traffic patterns that maintain bursty characteristics over wide-ranging time scales and “are intuitively ‘similar’ to one another (in a distributional sense)” [LeT95]. An example of one such proof taken from [LeT95] is shown below in Figure 2-9. Poisson-type processes may exhibit bursty characteristics over certain time scales (i.e., have a natural burst length), but when observed over longer time periods, they become smooth. See [LeT95] for a more complete development of self-similar processes.




Figure 2-9 Pictorial “Proof”: Self-Similarity

Development methods for generating fractal processes are ongoing, but one way to generate such a process is to aggregate many simple Poisson arrivals with service

times that have infinite variances (“heavy tail” distributions) [LeT95]. The pareto distribution,

$$f(x) = ax^{-(a+1)}, a > 0, 1 \leq x \leq \infty \quad 2-1$$

with $1 < a < 2$ (stable pareto), fits this criteria [LeT95]. The mean of this distribution is:

$$Mean = \frac{a}{a-1} \quad 2-2$$

provided $a > 1$ [Jai91]. By this definition, a stable pareto distribution has infinite variance.

2.8.5 DIS Traffic

To be of use, a traffic source model must strike a balance between accuracy/complexity and practicality/simplicity. Depending on the simulation overhead and design complexity incurred, the most accurate traffic model may not always be used. Different types of traffic may also call for different source models. One of the most resource-intensive types of traffic to be considered in the Virtual Battlelab Environment (VBE) is that of distributed interactive simulation (DIS).

2.8.5.1 DIS Defined

DIS has evolved from early DoD simulation programs like SIMNET, and consists of simulations of virtual battle spaces conducted with multiple manned and computer generated forces. In [HoL95] it states, “The primary mission of DIS is to define an infrastructure for linking simulations of various types at multiple locations to create realistic, complex, virtual ‘worlds’ for the simulation of highly interactive activities.” The connectivity of the simulation is based on the premise that each node maintains its

own version of the simulated space, that it is responsible for certain entities in this space, and that it passes only those changes to the space for which its entities are responsible. This form of operation requires the network protocol in use to have a multicasting/broadcasting capability.

2.8.5.2 Entity State Protocol Data Units

Exchange of data is accomplished by a standardized set of protocol data units (PDUs) of various types and classes. PDUs are generated as events in the simulation occur, such as firing of weapons, detonations, or energy emissions. By far the most common PDU type generated is referred to as an entity state PDU (ESPDU). One recent study found that ESPDUs account for approximately 96% of PDUs generated in DIS simulations [PuW95]. ESPDUs are generated at specified intervals whenever an entity's articulated parts (position, orientation, configuration, etc.) are changed.

Given the entity class of an object (aircraft, tank, etc.), formulas can be used to estimate the size of ESPDUs. As a result, DIS demonstration tests have provided average and peak ESPDU generation rates (in PDUs per second) for use in traffic estimation. See Table 2-3 for a summary of these empirical estimates [PuW95].

Table 2-3 ESPDU Data Averages

Entity State PDU Average Data Values		
	Air Entities	Land Entities
Average ESPDU Size (Bytes)	464	224
Average Issue Rate	1 ESPDU / second	.17 ESPDU / second

2.8.5.3 STOW DIS Data

Initiated by the Advanced Research Projects Agency (ARPA), the Synthetic Theater of War (STOW) program is an effort to expand the current DIS capability for large-scale exercises [NgB96]. STOW-Europe (STOW-E) was conducted in November 1994, and its data traffic was gathered and analyzed. Among the lessons learned from this analysis, it was shown that DIS traffic exhibits both Poisson-type and self-similar characteristics [NgB96].

The scale of STOW-E was approximately 2000 entities (DIS simulations have typically been on the order of thousands of entities), but DARPA's objective was to implement a virtual battle space with 100,000 entities by the year 2000 [PuW95]. This could result in aggregate traffic of approximately 50,000 PDUs/second, with up to 1,000 multicast groups. In addition to DIS traffic, support traffic such as pre-simulation file transfer and concurrent video must be considered [PuW95]. These applications' traffic characteristics differ from those of DIS, but should be considered to fully account for network requirements.

2.8.6 VBR Video: MPEG-2

Video data traffic most often accompanies DIS traffic to aid in set up and operation of DIS exercises [PuW95]. Video typically comes in two forms: VBR and constant bit-rate (CBR) video. Forcing video to be CBR results in added delay, wasted bandwidth, and variable quality [Org93]. On the other hand, VBR video has constant quality and takes advantage of the efficient statistical multiplexing of VBR connections in ATM. This thesis looks specifically at MPEG-2 VBR Video, which incorporates both intraframe and interframe compression to generate a VBR source.

MPEG-2 is a standard used for the compression of video. Target applications for MPEG-2 include: High Definition TV, interlaced digital video, cable, and satellite TV, Digital Versatile Disk, and Video on Demand over computer networks [Wu98a]. MPEG-2 is an asymmetrical, lossy compression scheme that utilizes both interframe and intraframe compression. MPEG video frames come in three formats: Intra (I) frames, Forward Predicted (P) frames, and Bi-directionally Predicted (B) frames.

- I Frames - coded using intraframe compression.
- P Frames - predicted via motion compensation from previously encoded I frame. Differences are encoded and forwarded along with displacement (motion vectors) to decoder.
- B Frames - motion compensated prediction and interpolation is performed using either I frames, P frames or both. These are never used for prediction, as error propagation will result.

MPEG-2 frames are typically transmitted at 30 frames/second. In addition, they are sent in a fixed sequence known as a group of pictures (GOP). A group of pictures consists of all frames between two consecutive I frames (including the first I frame). A typical GOP is shown in Figure 2-10.




Figure 2-10 Group of Pictures

The distance between I frames is denoted as N , where N must be less than or equal to 16. When N is larger, errors will propagate ($N = 15$ in Figure 2-10). The distance between consecutive I and P, P and P, and P and I frames is denoted as M ($M = 3$ in Figure 2-10), and there are $(M - 1)$ B frames between these pairings. Table 2-4 lists average sizes (in bits) for I, P, and B frames taken from a standard test sequence coded by the MPEG Test Model method with $N = 15$ and $M = 3$ [Fog98].

Table 2-4 MPEG Frame Sizes

Level	I	P	B	Average
30Hz @ 4 Mbit/sec	400,000	200,000	80,000	130,000

2.9 B-ISDN and ATM

2.9.1 B-ISDN

Today, separate networks are used to carry voice, data, and video. This is largely because each media type has different characteristics. Data tends to be “bursty” in nature with significant “dead time,” resulting in low utilization of network bandwidth. In addition, data is rather insensitive to delay. Voice and video, on the other hand, require steady data rates, and are very sensitive to delay. Regardless of these differences, separate networks are neither cost-effective nor justifiable given modern object-oriented technologies. There is a need for a common network that supports all types of media.

Broadband Integrated Services Digital Network (B-ISDN) is the technology that promises to integrate all communication services (voice, video, data, telephony, etc.) over

a single network to create the "Information Superhighway" envisioned by users worldwide [WuI98b]. The primary transport mechanism for B-ISDN is the Asynchronous Transfer Mode (ATM), which will run on the Synchronous Optical Network (SONET). SONET is an international standard that provides the physical layer medium for ATM via high-speed, low-error fiber optic cables. The B-ISDN protocol architecture is shown in Figure 2-11.




Figure 2-11 B-ISDN Protocol Architecture [RIL97]

2.9.2 B-ISDN vs. N-ISDN

B-ISDN technology stems from the original ISDN standards published in the 1980s, which is now commonly referred to as Narrowband ISDN (N-ISDN). N-ISDN utilizes existing telephone networks (copper pairs) to carry information traffic. The two service configurations supplied by N-ISDN are [McS95]:

- *Basic Rate Interface (BRI)*- two 64 kbps (B) channels for user data, and one 16 kbps (D) channel for network management and control. The BRI is primarily intended for voice, data, and videophone.

- *Primary Rate Interface (PRI)*- twenty-three 64 kbps (B) channels and one 64 kbps (D) channel. The PRI is primarily intended for high-bandwidth applications such as LANs and Private Branch Exchanges (PBXs).

These circuit-switched configurations supply predefined channel rates, reducing flexibility and scalability to the end user. Furthermore, loosely defined standards further hampered the success of N-ISDN.

B-ISDN brought forward the separate user, control, and management aspects of ISDN, and improved upon them. High-speed fiber optic cables supply high data rates with a bit error-rate of less than one in a billion [Cad94]. This efficiency allows the transfer of error correction and detection functions from the network to the end systems. Also, B-ISDN is cell-switched over virtual channels, allowing dynamic allocation of bandwidth on demand via statistical multiplexing. This type of multiplexing allows B-ISDN to handle “bursty” data much more efficiently than time-division multiplexed ISDN circuits. Finally, strong international support for this technology should allow for the realization of applications such as videophones, video-on-demand, and distributed interactive simulations.

2.9.3 SONET

Synchronous Optical Network (SONET) is a standard for the physical layer framing structure for ATM [Cad94]. This standard is necessary to provide multi-vendor interoperability among network connection equipment. SONET links have three major advantages over traditional Plesiochronous Digital Hierarchies (PDH) links in use by phone companies [McS95]:

- Higher transmission rates

- Direct multiplexing (no intermediate multiplexing stages)
- Low bit error-rates.

SONET signal rates known as Synchronous Transfer Signals (STS-M), along with their associated Optical Carriers (OC-N) are depicted in Table 2-5. In addition, the equivalent number of digital subscriber (ISDN) lines is listed for comparison purposes.

Note: a given STS-M can be carried on any OC-N so long as $N \geq M$.

Table 2-5 SONET STS-N/OC-N Rates [McS95]

STS-N or OC-N level	Bit Rate (Mbps)	Number of DS0s	Number of DS1s	Number of DS3s
1	51.84	672	28	1
3	155.52	2016	84	3
6	311.04	4032	168	6
9	466.56	6048	252	9
12	622.08	8064	336	12
18	933.12	12096	504	18
24	1244.16	16128	672	24
36	1866.24	24192	1008	36
48	2488.32	32256	1344	48
96	4976.00	64512	2688	96
192	9952.00	129024	5376	192

As evidenced in Table 2-5, above, SONET equipment handles a wide variety of traffic. It transfers DS1, DS3, and ATM using point-to-point and ring-based systems. For additional survivability, SONET is commonly deployed in dual, self-healing rings. Survivability could be further increased if carriers would start running fibers down different cables, protecting them from construction backhoes.

The basic SONET STS-N frame format is shown in Figure 2-12. The frame consists of section overhead, line overhead, path overhead, and a synchronous payload envelope [WuI98b]. Transport overhead is composed of both section and line overhead, which is processed by all equipment that STS frame encounters. Path overhead, on the

other hand, is only processed by destination equipment. Each byte of payload corresponds to one 64 kbps channel. In addition, STS frames are sent every 125 μ s—the period for voice signals [WuI98b].

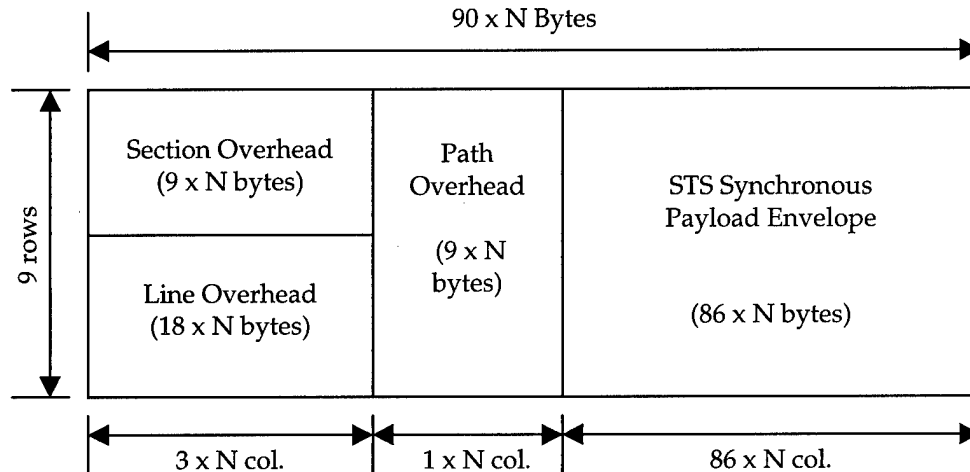


Figure 2-12 STS-N Basic Frame Format [McD95]

2.9.4 Why ATM?

In the world of high-speed wide area networking, Asynchronous Transfer Mode (ATM) has emerged as the "most promising technology for supporting future broadband communications." [SiJ95] As communications traffic becomes more diverse in nature (voice, video, data, etc.) and requires more and more bandwidth, existing networks do not have the capability to efficiently support this heterogeneous mix of traffic. ATM promises the capability to support many different kinds of traffic in an integrated, efficient manner. Much of the research of efficient allocation of resources in wide area network (WAN) environments has been addressed using an ATM approach. Since the problems that will be addressed in this work are in a high-speed WAN environment,

ATM will be the underlying protocol used for data transfer at the lowest layers (in an OSI model sense).

2.9.5 ATM Structure and Operation

ATM consists of packaging data in fixed sized packets, called cells. Each cell is 53 bytes in length, comprised of 5 bytes of header and 48 bytes of payload data. Fixed cells allow switching, routing and multiplexing functions to be simplified and streamlined. As a result, ATM switches can operate at extremely high bit rates.

In addition, ATM is primarily connection-oriented (somewhat like traditional telephone networks) in that resources are allocated along a path from source to destination. ATM connections are referred to as "virtual" connections, since resources are not strictly dedicated, but rather shared with other connections along the path. This sharing allows greater utilization of links through asynchronous statistical time division multiplexing (asynchronous because no slots or framing of cells is used in transmission)[Sta96].

End-to-end ATM connections are accomplished via virtual channel connections (VCCs) and virtual path connections (VPCs). VCCs are individual ATM connections, and are set up from a source to a destination. As shown in Figure 2-13a, a number of VCCs may be bundled over a single VPC. VPCs consist of two VPC termination points, a physical route, and an optionally assigned capacity [FrH96]. A VCC can use a single VPC if it connects the intended source and destination nodes (such as VPC 3 in Figure 2-13b). In the case that a single VPC does not connect an intended source and destination, a VCC can use several VPCs en route (such as VPC 1 and VPC 2 in Figure 2-13b).




Figure 2-13 VPC and VCC Operation

2.9.6 ATM Service Classes

ATM divides the types of traffic source services it supports into four classes, each with a corresponding ATM adaptation layer (AAL) type [LiP97, [SiJ95]:

2.9.6.1 Class A - Constant Bit-Rate (CBR) Service:

AAL1 supports this connection-oriented service in which the bit-rate is a constant. Examples include 64kbit/sec voice, uncompressed fixed-rate video streams, and dedicated leased lines use by private networks.

2.9.6.2 Class B - Variable Bit-Rate (VBR) Service:

AAL2 supports this service, which is also connection-oriented, but has bit-rates that are variable. These services require bounded end-to-end (ETE) delays, but may or may not have a "real-time" constraint. If the service operates real-time, it is referred to as VBR-RT (for real time), and the variance of the ETE delay must meet specified bounds. Examples include interactive packetized voice and compressed packetized video, such as

video conferencing. VBR-NRT (non-real time) still must meet overall delay requirements, but variance in delay is not tracked since there is no real-time requirement. An example of a VBR-NRT service is multimedia email.

2.9.6.3 Class C - Available Bit-Rate (ABR) Service:

This service class consists of connection-oriented data traffic, such as file transfer and email. There are no required bounds on ETE delay, and bit rates are variable. Bit rates are typically dictated by network loading, and are only guaranteed over the VC at a minimum cell rate (MCR). Since guaranteed bandwidth may not be available after CBR and VBR services' needs are met, the MCR is most often set to zero and ABR operates on an availability basis. Non-zero MCRs may result in a denied connection. AAL3 and AAL4 were merged into AAL3/4 for Class C, and AAL5 is also often used to support this class.

2.9.6.4 Class D - Unspecified Bit-Rate (UBR) Service:

Any left over capacity can be used by the UBR class, which is the "best effort" connectionless data service in ATM. Much like ABR, email and file transfer often use UBR, but no resources are allocated and UBR cells are dropped first when congestion occurs. Cell loss recovery and ETE performance must be accomplished by whatever applications use this service. Both AAL5 and AAL3/4 may be used to support Class D.

Although intended for use with their associated classes, AAL types may be used to support any class of service. In actuality, the majority of vendors make ATM products that support all classes with AAL5, and the ATM Forum has focused most of its efforts on AAL5 [SiJ95].

2.9.7 ATM and the OSI Model

Does ATM operate at the data link layer or network layer? Can it operate at the transport layer as well? Answers to these questions depend on your point of view. In practice, ATM is used in a variety of ways. Higher layer protocols may be used “on top of” the ATM layers, but since ATM has many inherent capabilities, higher layer protocols may not be necessary in certain situations.

Figure 2-14 shows the ATM (or B-ISDN) protocol structure, consisting of the ATM adaptation layer (AAL), the ATM Layer, and the physical layer. The AAL interfaces higher layer protocols or applications to the ATM layer by performing segmentation and reassembly between packets of data and ATM cells. The ATM layer performs traffic management and flow control, as well as the VP and VC routing functions described above in Section 2.9.5. Framing and bit timing at the physical layer is most often implemented using SONET, as described in Section 2.9.3. Other physical layer standards have also been proposed, such as ATM over twisted pair [SiJ95].




Figure 2-14 ATM Protocol Stack

2.10 ATM Implementations

When considering the use of ATM, it should be remembered that ATM can either be used directly with applications or as a lower layer circuit over which higher protocols operate. How ATM is used can be defined by where the application program interface (API) resides [TrE95]. Figure 2-15 shows a conventional versus a “native” ATM implementation. The conventional approach is to operate a higher layer protocol, such as TCP/IP, over ATM, and use ATM as an underlying circuit network. “Native ATM” solutions are being devised to take advantage of the multimedia support capabilities of ATM by interfacing applications directly to ATM.

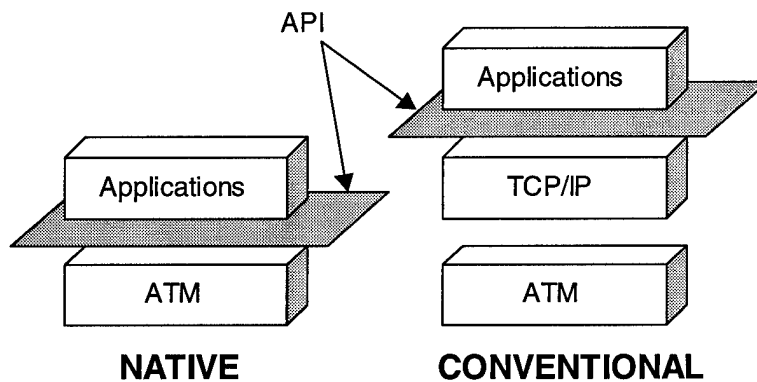


Figure 2-15 Native vs. Conventional ATM Use

The conventional approach has been widely used as IP traffic began to utilize the high-speed improvements of ATM as an underlying network technology. However, there exists a disconnect between IP's inherent connectionless service and ATM's connection-oriented service. As such, ATM's inherent strengths (QoS and high data rates) are not used to their fullest extent. As ATM technology becomes more prevalent, though, ATM to the desktop may become more feasible.

2.10.1 Native ATM

The Native ATM version operates as shown below in Figure 2-16, with all link, network, and higher layer responsibilities delegated to the ATM protocol. A Native ATM solution is essentially a network-layer VPN since the functions of the network layer are performed by the ATM protocol. This can further be defined as a “peer” VPN model, since at each “hop” in the ATM network, the packet stream returns to the network layer (ATM layer) [FeH98].

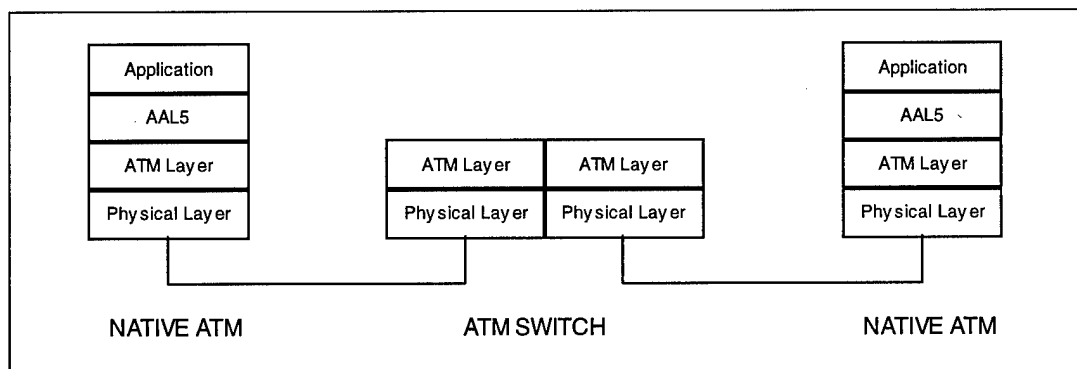


Figure 2-16 Native ATM Protocol Structure

The interfaces between ATM end-stations and ATM switches are defined by the user-to-network interface (UNI) and the network-to-network interface (NNI). Although UNI is an evolving specification (currently UNI 3.0/3.1) [ATM98], it is defined, together with the NNI, to perform routing via VPI/VCI connections. The UNI specification works with ATM Address Resolution Protocol servers (ATMARPs) to provide registries of ATM addresses so that hosts (both ATM hosts and non ATM hosts such as IP) may register their ATM addresses [Bob97][All95]. If an ATM client knows the address of a host the connection is made, but if it is unknown the ATMARP is queried for the address.

Multicast operation is yet to be finalized by the ATM Forum, but one form it could take is that of overlaid point-to-multipoint connections [All95], where every node

maintains a connection to every other node. If only point-to-point connections are supported, this solution is called an " N^2 mesh." Although resistant to node failures, this approach requires an excessive number of connections to be maintained. In addition, this approach is not very scalable, especially since most ATM ports have a limit of 1024 connections due to memory constraints [Bob97]. This approach could be used if the number of nodes is relatively small.

Another approach is to use a multicast server, where a designated server acts as single destination for all clients. Unidirectional point-to-point connections are set up for each client to the multicast server, and the server sets up a single point-to-multipoint connection to the clients for each multicast group [All95]. Such an approach is much more scaleable, requiring fewer connections to be maintained and simplifying multicast group subscription and de-subscription. The multicast server, however, can act as a bottleneck and/or a single point of failure.

Both of these approaches can be improved if a multicasting tree concept is utilized. To minimize the number of VPCs and VCCs that must be maintained, "root" point-to-multipoint connections can be established, with additional connections added on as "leaves" from the nearest point on the path [Bob97]. Subscription and de-subscription require complex switching procedures, but the amount of redundant data is greatly reduced in point-to-multipoint connections.

2.10.2 IP over ATM

Classical IP over ATM offers a way of running network protocols (e.g., TCP and UDP), and various applications (e.g., FTP, WWW, etc.) directly over ATM. This approach treats ATM as a LAN and partitions an ATM network in several logical

subnets known as Logical IP Subnetworks (LIS) [Xu97]. Current implementations of ATM provide a replacement for [RFC98a]:

- Local area networks (LANs)
- Local area backbones between existing LANs
- Links between IP routers.

All end systems within a given subnet have the same IP prefix and address mask. In addition, all members in a given LIS are interconnected via the underlying ATM network. Hence, these systems have both an IP address and an ATM address. Since the two are not logically coupled, some means is necessary to resolve IP addresses into ATM addresses and vice versa. This means is the ATMARP service previously mentioned in Section 2.10.1. ATMARP service only has local significance within each LIS. Inter-LIS communications must be accomplished via an IP router.

In the Classical IP over ATM approach, IP packets are carried in ATM AAL5 and have a default maximum transmission unit (MTU) length of 9180 bytes. However, the alternate MTU sizes may be negotiated during call setup to better suit the needs of the cooperating IP members [RFC98a]. This default MTU is large enough to hold Ethernet, token ring, FDDI, and SMDS. This allows for greater interoperability and better performance since most router overhead is due more to number of packets handled versus number of bits handled [RFC98a]. Since AAL5 is not reliable, it is up to higher layer protocols to provide retransmissions in the event of failures.

The classical approach is not very optimal since it still requires data traffic between end systems of different logical subnets to traverse IP routers even though both end systems are attached to the same underlying ATM network. This is undesirable

because each router has to reassemble ATM cells into IP packets for routing purposes. This process introduces high latencies and bandwidth bottlenecks. As a result, the Next Hop Resolution Protocol (NHRP) was developed to address these problems [RFC98b].

NHRP allows an end system in one LIS to resolve (via an IP address) the ATM address of a system in a foreign LIS and set up an ETE ATM connection, known as a “shortcut” or “cut through,” between them. In this approach, all IP workstations are only one “hop” away from each other at the IP level, even though several ATM switches are traversed between the end stations.

The IP over ATM protocol stack is shown in Figure 2-17. Like Native ATM, IP over ATM can also be considered a network-layer VPN, but this time the network layer is implemented using IP.

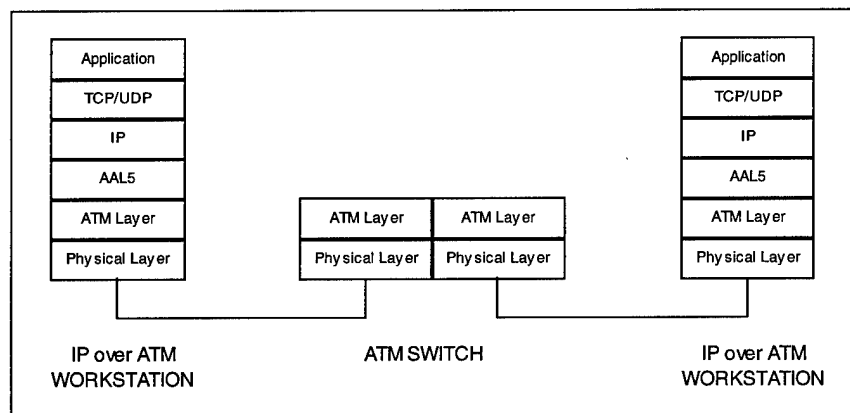


Figure 2-17 IP over ATM Protocol Structure

In addition, this would be considered an “overlay” network-layer VPN since the IP layer cuts through the underlying ATM network [FeH98].

2.10.3 LAN Emulation

In an effort to encourage the interoperation of ATM with legacy LAN systems, the ATM forum developed the LAN Emulation (LANE) signaling scheme. LANE was devised as a way in which ATM could operate below the legacy LAN medium access control (MAC) protocols to bridge their transmitted traffic and make the LAN applications function as if all were located on a single LAN. Figure 2-18 shows how ATM resides below the MAC and IP layers, and implies that in LANE operation, it must emulate those characteristics of the MAC that it does not inherently possess. Some of these characteristics negate ATM strengths, whereas others provide capabilities that ATM lacks [TrE95].

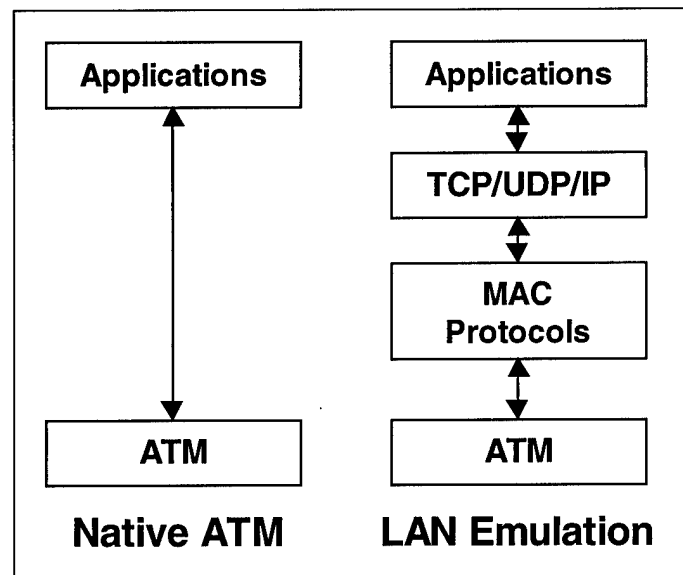


Figure 2-18 LANE vs. Native ATM

The LANE protocol, shown in Figure 2-19 LANE Protocol Structure

, is designed to operate such that all layers above the MAC/LANE layer perceive no change from standard shared-access LAN operation. This can be considered a link-

layer VPN since the LANE layer is designed to emulate the MAC layer, which is analogous to the data link layer for Ethernet and Token Ring protocols.




Figure 2-19 LANE Protocol Structure

2.10.3.1 Segmentation and Reassembly

Since an Ethernet LAN frame (maximum size of 1518 Bytes) does not fit into a 48 byte ATM payload, frames must be segmented to fit in ATM cells. As discussed in Section 2.9.7, the AAL layer provides this function. Due to the unnecessary complexity and increased overhead of AAL3/4 (9 bytes versus 5 bytes), AAL5 was chosen by the ATM Forum as the adaptation layer for LANE.

2.10.3.2 Providing Connectionless Service

User access to legacy LAN networks is connectionless and therefore fairly simple. User applications simply send out information (typically on a shared medium) whenever access is permitted. ATM emulates this connectionless characteristic by resolving MAC addresses to ATM addresses, and then setting up connections over which data can flow. In this sense it hides all connection-oriented issues from the higher layers, and, as long as

connection setup can be performed fast enough, the ATM network is in turn hidden from the connectionless characteristics of the higher layers [TrE95].

In a WAN environment, resources may be wasted if utilization is not high, but if traffic generated by the higher protocols is heavy, this effect is lessened. Also, since the ATM interconnections are primarily emulating MAC broadcasting behavior, the number of connections can become unwieldy if the number of emulated LANs (ELANs) is large.

2.10.3.3 Providing Broadcast/Multicast Service

A shared-medium LAN can provide multicasting services very easily since the shared medium operates in a broadcasting fashion. A user simply listens to the medium, and if the frame is intended for or desired by the user, it is taken. If not, it is ignored. Although only functional over a local area, this quality of shared access makes multicasting a strength of legacy LANs.

ATM (and IP for that matter) require complex routines to provide multicasting services. LANE, in addition to providing for reuse of LAN equipment and applications, also inherits the multicasting strengths of shared-access LANs. ATM provides paths for multicast traffic by using point-to-multipoint connections, governed by a multicast server similar in concept to that explained in Section 2.10.1.

2.10.3.4 LANE Architecture

The current LANE architecture, detailed in the LANE UNI (LUNI) phase 1 specification [All95], and explained in [TrE95] consists of four elements: LANE clients (LECs), the Broadcast/Unknown Server (BUS), the LANE Server (LES), and the LANE Configuration Server (LECS). Figure 2-20 shows the LANE concept of operations.

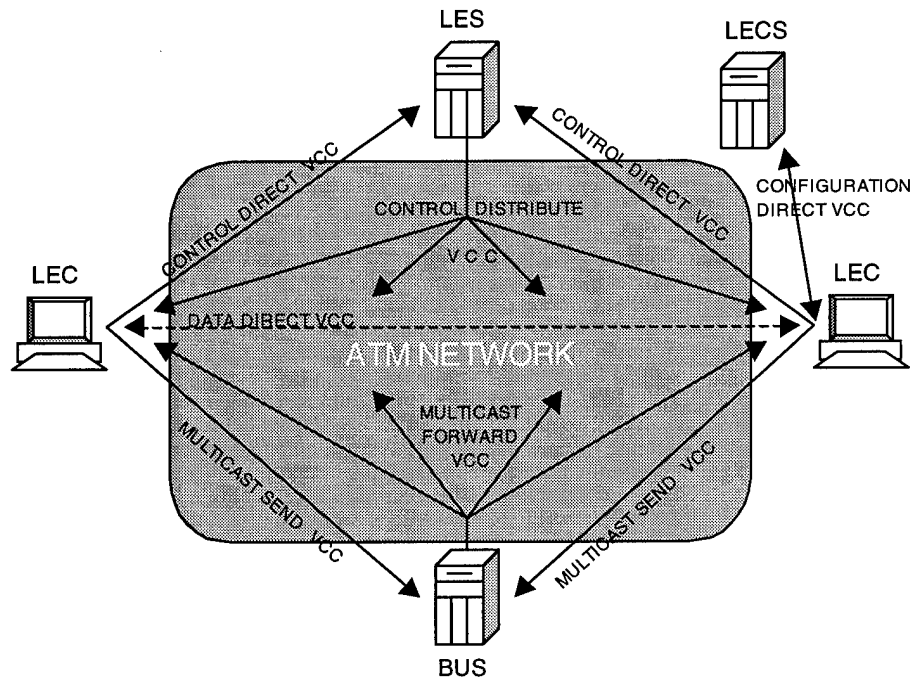


Figure 2-20 LANE Operation

LANE Client (LEC): A LEC is any station (workstation, server, bridge, router, etc.) connected to the ATM network that participates in the LANE service. It has both a MAC IEEE 48-bit address and an ATM address, which it resolves directly with the LES. When a LEC wishes to communicate with another LEC, it transmits directly over an ATM path (Data Direct VCC) as long as it knows the MAC-ATM resolved address of the destination. If unknown, it queries the LES (via Control Direct VCC) and the LES passes the address (via Control Distribute VCC) to aid the LEC in setting up the direct connection.

LANE Server (LES): The LES's primary function is to perform address resolution protocols (ARPs), and it can either broadcast ARPs to all LECs (via Control Distribute VCC), or maintain a resolution table and pass ARPs only to the requesting and requested nodes (via Control Direct VCCs). Optionally, it may act as a pure multicast

server that routes multicast traffic itself, rather than just forward all multicast data back to requesting nodes in replied ARPs.

Broadcast/Unknown Server (BUS): The BUS has two primary functions. The first is to take broadcast messages from the LECs (on Multicast Send VCCs) and forward them to all LECs (via Multicast Forward VCC). The second is to locate LECs whose addresses are unknown to the LES. LECs send broadcast or unknown requests directly to the BUS, which in turn broadcasts to all LECs using its broadcast point-to-multipoint connection.

LANE Configuration Server (LECS): The LECS is used when a LEC initiates LANE activity. The LECS resides at a preset ATM address and points the LEC to the LES to begin activity (via Configuration Direct VCC). If more than one ELAN exists on a network, the LECS can direct the LEC to the appropriate LES for a particular emulated LAN.

In summary LANE provides for economical reuse of legacy LANs, as well as multicasting and broadcasting capability over an ATM network. Problems with scalability, as well as weaknesses associated with LANE server failures or bottlenecks make LANE unsuitable for environments containing large numbers of nodes.

2.11 Summary

This chapter has presented the theoretical background necessary to devise an effective connectivity architecture for a high-speed wide area ATM network. It began with an overview of the Air Force Battlelabs, followed by a description of the Defense Information System Network. Next, a brief review of the OSI Reference Model was presented as well as a discussion of Virtual Private Networking. Multimedia traffic and

data compression were both addressed, followed by a development of traffic source modeling. Finally, the basics of ISDN and Asynchronous Transfer Mode were discussed, followed by an explanation of signaling options involving ATM.

Chapter 3 introduces a systems engineering approach that can be used to assess and analyze the VBE interconnectivity problem with respect to the background information and concepts discussed in this chapter. This systems approach relies heavily on the user's requirements, and is tailorable to any problem domain.

CHAPTER 3

METHODOLOGY

3.1 Methodology Overview

The remainder of this thesis follows a systems engineering approach as described by [Hal69], with the first four steps of that process included in this chapter. This process is not free from subjectivity, and is heavily dependent on user requirements. When these requirements are not available or undefined, assumptions should be made based upon the “best” information available to ensure optimal results. This process, when correctly implemented, is a valuable tool in aiding both the decision-maker and the design engineer. As a result, a systems engineering approach is used in this study, and reasonable assumptions are made where necessary.

Section 3.2 begins with a description of the problem and general assumptions, followed by Section 3.3, which lays out the objectives of the study in terms of a value system definition. Section 3.4 details the system synthesis, listing the possible options for solving the defined problem. Finally, Section 3.5 describes the model to be used to evaluate the devised solutions and Section 3.6 gives a brief summary. The remaining steps in this process are treated in Chapter 4.

3.2 Problem Definition

3.2.1 Scope of Problem

The Air Force Battlelabs must have the capability to meet their high-speed data communication needs, ranging from standard voice traffic to video to distributed

interactive simulation. The collection of these communication needs has been termed the "Virtual Battlelab Environment"(VBE) [SOW98]. This VBE consists of nine geographically separated nodes all generating some form of worst-case traffic to all other nodes. Three competing solutions will be devised, and their performance evaluated to determine which best meets stated constraints and measures of performance.

3.2.2 Players

The chief decision making authority resides in HQ USAF/XORB, with Concurrent Technologies Corporation (CTC) on contract to perform systems design to determine an optimum solution. CTC is considered the chief decision maker (CDM) for this particular study, since the Air Force Institute of Technology (AFIT) supports CTC by performing research and technical studies [SOW98]. Attempts to define HQ USAF/XORB's needs are ongoing, and, since many of these requirements are as yet undefined, reasonable assumptions have been made wherever stated requirements were unavailable. In addition to HQ USAF/XORB, CTC, and AFIT, the remaining six players in the VBE are the Air Force Battlelabs.

3.2.3 Assumptions

The primary assumptions for this problem can be divided into two categories: *constraints* and *alterables*. Constraints are stated needs that must be met by any proposed solution. Alterables (also known as factors) are unconstrained parameters that can be varied or traded to determine their relative impact or importance. The constraints and alterables of this study are listed below.

3.2.3.1 Constraints

The following constraints have either been determined from HQ USAF/XORB, or have otherwise been assumed:

- As implied in [SOW98], ATM technology will be the low-level means of providing high speed connectivity for the VBE. However the specific implementation of ATM or any overriding protocols is not fixed, as explained below in Section 3.2.3.2.
- No dedicated transmission services will be considered due to cost considerations and Department of Defense (DoD) ATM policy[DIS98a]. Instead, some form of Virtual Private Networking must be utilized.
- Defense Information System Network (DISN) ATM Services [DIS98a] will be the backbone network over which the VBE forms a Virtual Private Network (VPN).
- Several configurations can be used to connect a given Battlelab to DISN ATM Services as explained in DISA's ATM system specification [DIS98a]. It is assumed that each node has its own ATM site switch that connects to the nearest DISN SDN ATM switch. It is further assumed that both these switches are located on-site.
- Only Distributed Interactive Simulation (DIS) traffic with Variable Bit Rate video traffic is modeled. All other sources of traffic are assumed minor in comparison with these two traffic sources, and are not modeled to decrease model complexity.
- A single distributed simulation (with video support) is simulated for a fixed period of constant activity (5 seconds).
- A worst-case scenario is modeled requiring each node to maintain 1000 DIS entities along with VBR video traffic in a simulation, which is broadcast to all other nodes.

- All simulated DIS and video traffic is assumed to be unclassified in nature.
- End-to-end delay from the application layer of one node to the application layer of any other node cannot exceed a specified amount (see Section 3.3 for values).
- A bandwidth-on-demand concept of operations is assumed, with bandwidth costs calculated on a per-bit transmission basis.

3.2.3.2 Alterables

The following parameters are alterable in that they can be varied to determine their feasibility and potential for solving the stated problem.

- The utilization of the bandwidth of the transmission lines that connects each Battlelab to the DISN ATM network can vary, but should be minimized to reduce costs (since costs are calculated on a per-usage basis).
- The protocol architecture used to deliver DIS and video application data over the ATM backbone is not constrained.
- Cell delay variation should be minimized to reduce jitter, which results from applications that are not synchronized.
- Cell losses should be kept to a minimum to reduce the impact to DIS and video output quality.

3.2.4 System and Environment

To better understand the relationship between constraints and alterables, Figure 3-1 below shows the boundary between the system (what can be varied or altered) and the environment (what is fixed or constrained).

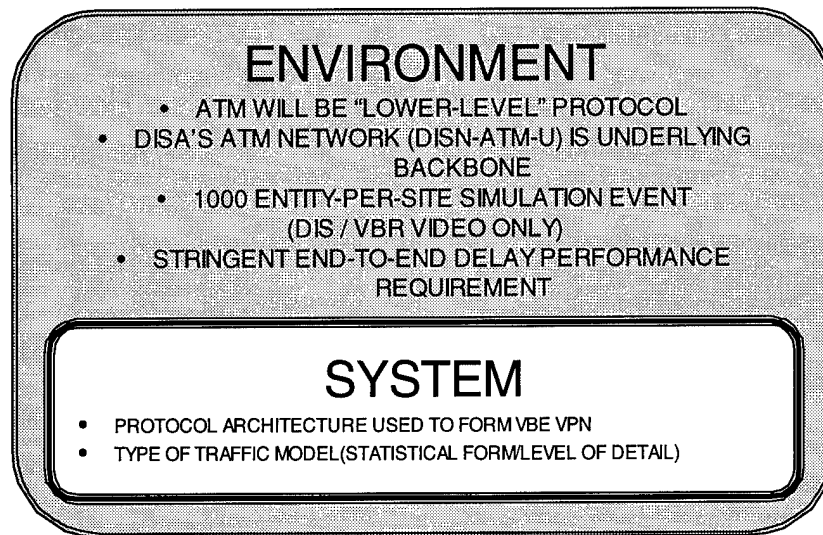


Figure 3-1 System vs. Environment

3.3 Value System Definition

This section defines what determines an “acceptable” solution (one that meets stated constraints), as well as what further defines a “better” solution (the one that performs best among alterable parameters). The definition of a value system allows the outcomes of different solutions to be compared and rated. Performance requirements are “pass-fail” criteria, and measures of performance (MOPs) are scaled criteria that are used to determine the relative performance among proposed solutions. Both requirements and MOPs are listed below.

3.3.1 Performance Requirements

For this research, only one requirement must be met: the ETE delay from any application layer client to any other application layer server must not exceed 300 msec [PuW95]. Although not specified by the CDM, it is further assumed that this requirement must be met with a probability of 99%. If not met, this requirement causes a proposed

solution to be eliminated. This is the only fixed requirement that each proposed solution must meet to be considered viable for this study.

3.3.2 Measures of Performance

The measures of performance that are considered in this study are listed below and summarized in Figure 3-2. The MOPs are divided into two categories: those that measure effectiveness (performance), and those that measure efficiency (cost).

3.3.2.1 Performance MOPs

Average End-to-end Delay: The average delay of packets from the application layer of all hosts to that of all clients, measured in milliseconds.

Average Cell Delay Variance (CDV): The average variance of delay of ATM cells, measured in milliseconds.

Cell Loss Ratio: The ratio of dropped ATM cells to the total number of ATM cells sent (unitless).

Tightly Coupled Delay Measure: The probability that ETE delay of ATM cells does not exceed 100 milliseconds. This is considered a measure of performance since the delay from any application layer client to any other application layer server should not exceed 100 msec for tightly coupled simulation [PuW95].

3.3.2.2 Cost MOPs

Cost of Implementation: Estimated level of cost incurred by a solution (i.e., hardware, software, infrastructure impacts), measured as either low, medium, or high.

Cost of Operation (Bandwidth): Assumed to be limited to and dominated by the cost of demanded bandwidth on all access lines, measured as a total percentage of bit-rates utilized.

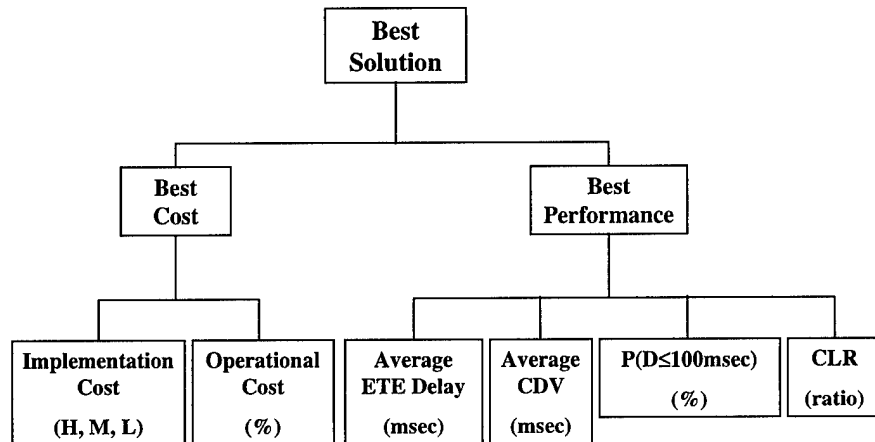


Figure 3-2 Value System Hierarchy

3.4 System Synthesis

Many approaches could be taken to assess the performance of a solution. These range from installing hardware and software and testing its operation to conceptual “pen and paper” analysis. This study develops a computer model of the VBE and the connectivity options using a networking simulation software package. Many such simulation tools exist (such as MATLAB, COMNET, OPNET and BONES Designer) and each has its own strengths for modeling computer network behavior. Since this study is primarily concerned with comparing the performance of existing protocols, the chosen tool should possess the capability to quickly model currently available protocols. It should also be interoperable with outside sources of data, to maximize the reuse of previous work.

3.4.1 Modeling Tool

After a review of available simulation packages, OPNET (by MIL3) was selected as the tool to be used for this analysis. This decision was based on a number of factors. First, OPNET has an extensive library of commonly used protocols (ATM, TCP/IP, Ethernet, etc.). It also has a helpful graphical user interface. Another important consideration was OPNET's role in the Department of Defense (DoD).

In January 1997, a large technical working group met to address the DoD's concerns regarding the potential burden that combat may place on its communications infrastructure. One of the primary objectives of this meeting was to assess the communications modeling tools available for the development of the Networks and Warfare Simulation (NETWARS) Model, which was initiated to address the DoD's concerns. This working group overwhelmingly chose the OPNET modeling tool, and, since this research will lay the foundation for ongoing VBE support, OPNET was also chosen for this study.

Some related DoD WAN simulation models have already been developed in OPNET, and whenever possible, this work attempts to leverage these previously developed models. The greatest example of these attempts is this study's reuse of the DISN ATM backbone model developed in [Gla98]. Lastly, using OPNET will promote reuse for future VBE efforts as well as interoperability with other DoD models.

3.4.2 Possible Network Protocols

Several protocols could be used in the VBE VPN including the internet protocol (IP) over ATM and LAN Emulation (LANE). As mentioned in Chapter 2, ATM itself has the capability to operate directly with the application layer, making a "Native ATM"

solution another option. Any proposed solution must perform with minimum delays and delay variance, and should offer a multicast capability.

3.5 Model Definition

After considering the protocol configurations that could be implemented, Native ATM, IP over ATM, and LANE were selected as the protocol solutions to be evaluated and compared in this thesis. Although the specifics of Native ATM's operation have yet to be defined, preventing its wide-spread use, it was still selected as a baseline solution. This "ATM to the desktop" approach has the promise of fully utilizing the high-speed switching and QoS capabilities that ATM was designed to offer. At least one other DoD WAN connectivity study used a native ATM protocol suite as a baseline, so this is not without precedent [Gla98]. IP over ATM was selected since it is the most likely solution to be implemented. The IP addressing scheme has seen wide-spread use as evidenced in the explosive growth of the internet. A "tunneling" or "cut-through" use of IP over an underlying ATM network is typical. LANE was selected because of its inherent broadcast/multicast capability, and because it allows maximum reuse of legacy LAN equipment and applications [TrE95]. Although it suffers from scalability problems, LANE's use in the VBE can be justified since the number of end stations in this model is currently small in number.

3.5.1 Common Model Components

Although they are distinct models, Native ATM, IP over ATM, and LANE all share many common components. This was done to improve the comparability of results

when the solutions are evaluated. All three protocol suites are well supported in the OPNET simulation tool.

3.5.1.1 Common Backbone Model

As stated in Section 3.2.3.1, the DISN ATM network is used as the backbone over which the VBE will operate. Figure 3-3 shows the OPNET model of the unclassified DISN ATM network, which was obtained from [Gla98]. Each ATM switch, referred to as a bandwidth manager (BWM), is modeled as an 8-port switch with output buffers 65,536 cells in size. This is taken from a FORE systems switch specification, and was verified as a realistic value by FORE systems [FOR99]. Each switch is assumed to have negligible fabric delay, and performs usage parameter control (UPC) functions on non-compliant cells by “tagging” them for selective discard when congestion occurs. All ATM links are modeled as STS-3 links (155 Mbps). These parameters were validated using DISA supplied information [DIS98a].

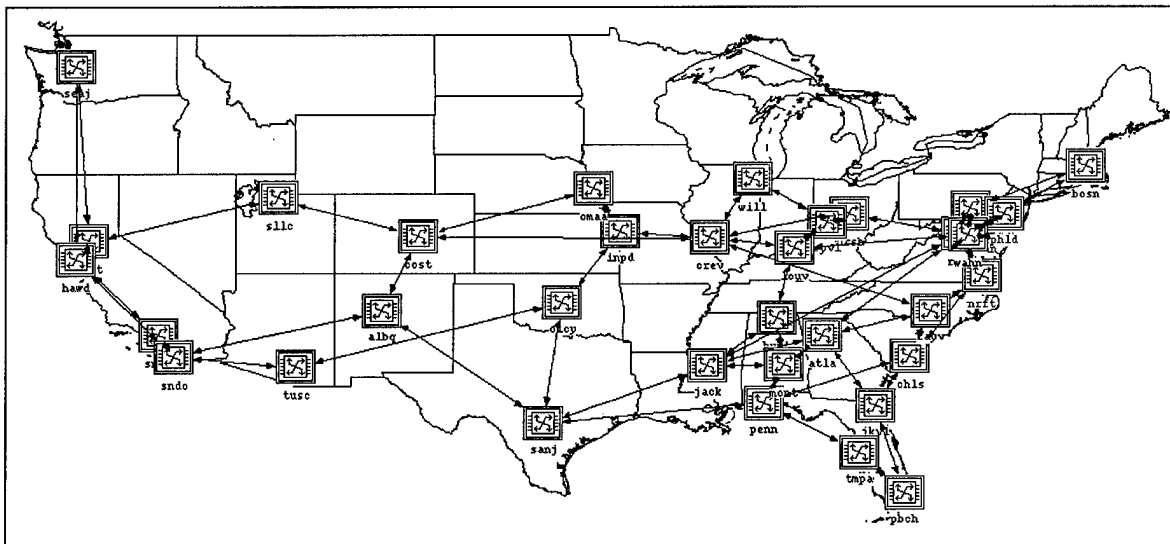


Figure 3-3 DISN ATM Backbone Network

Figure 3-4 shows the VBE geographic node locations, with their connections to the DISN ATM “cloud” (backbone) shown as lines to the ATM cloud icon. Connection into the ATM backbone is actually made at the ATM BWM switch closest in proximity to each VBE node.




Figure 3-4 Virtual Battlelab Environment

3.5.1.2 Common Site Model

In accordance with the assumption stated in Section 3.2.3.1 that all nodes are built in an identical, worst-case configuration, each has been created using the same site model shown in Figure 3-5. The service delivery node (SDN) ATM switches and the user controlled site switches are modeled in an identical fashion to the DISN BWM switches mentioned above. ATM links between clients and servers and ATM switches are assumed to be at STS-3 rates (155 Mbps). The DIS and Video workstation shown generates broadcast traffic to all other remote DIS or Video servers. This client-server traffic is modeled as one-way traffic (DIS and Video) generated at each workstation

destined for an assigned server (DIS or Video, respectively). Each of the client's connections are constructed as point-to-point ATM connections to all other nodes, intended to simulate point-to-multipoint connections (OPNET only supports point-to-point connections).




Figure 3-5 Common Site Configuration

3.5.1.3 Common Traffic Characteristics – DIS

Although generated randomly, the statistical characteristics of each node's generated traffic is assumed to be identical. The traffic model to be used for all DIS traffic attempts to capture the characteristics observed in the STOW-E traffic analyses of [NgB96] and [Sch95]. Specifically, traffic is generated continuously using an *on-off* source, with the *on* and *off* periods uncorrelated. Furthermore, this source follows a Poisson-type process (exponentially distributed period lengths). This Poisson variation is then modulated by a self-similar process by using pareto distributed interarrival times

[LeT95] during the on periods. This pareto PDF is set up in accordance with the observation of [NgB96] that DIS traffic appears to exhibit Poisson-type characteristics modulated by self-similar processes. Estimates for mean *on* and *off* periods were made from actual trace data in [Sch95] and [NgB96]. Schlorff observed that more than 84% of DIS traffic measurement samples (10 ms sample period) contained no data [Sch95]. Accordingly, the burst duration was selected to have a mean of 15 ms, and an *off* time with a mean of 85 ms.

The distribution of packet (ESPDU) sizes and mean packet interarrivals was determined from [PuW95]. Of 1000 entities assumed to be maintained per site, 100 were assumed to be air entities and 900 to be land entities. With rates of 1.0 air ESPDU/sec and 0.17 land ESPDUs/sec, 100 Air ESPDUs and 153 Land ESPDUs are generated per second, on average. The combined rate of packet generation is then 253 packets/sec. Since packet generation occurs only during bursts, and bursts occur on average only 15% of the time, the rate of packet generation during a burst is $253/0.15 = 1686.7$ packets/sec. Inverting this gives the mean interarrival time during a burst to be 0.592 msec/packet. Although the ESPDU types have associated mean packet sizes (224 Bytes for land ESPDUs, 464 Bytes for air), if the simplifying assumption is made that the probability of a given packet type is distributed according to its portion of the total packet generation rate, then the distribution of generated packets can be described as:

$$\begin{aligned}
 P(\text{Land ESPDU}) &= \frac{153 \text{ packets/sec}}{253 \text{ packets/sec}} = 0.6047 \\
 P(\text{Air ESPDU}) &= \frac{100 \text{ packets/sec}}{253 \text{ packets/sec}} = 0.3953
 \end{aligned}
 \tag{3-1}$$

The probability density function for packet sizes would then be as shown below in

Figure 3-6. Table 3-1 summarizes the statistical properties of the generated DIS traffic.




Figure 3-6 ESPDU Probability Density

Table 3-1 DIS Traffic Properties

DIS Traffic Statistical Properties	
Mean Burst Duration	15 milliseconds
Burst Duration PDF	Exponential
Mean Lull (Off) Duration	85 milliseconds
Lull Duration PDF	Exponential
Mean Packet Interarrival Time (During a Burst)	0.592 milliseconds
Packet Interarrival PDF	Pareto (stable)
Packet Size	224 Bytes or 464 Bytes
Packet Size PDF	Discrete: P(224B) = 0.6047 P(464B) = 0.3953

3.5.1.4 Common Traffic Characteristics – Video

As mentioned in Chapter 2, the video model design used in this study is based on a standard test sequence from the MPEG test model method [Fog98]. MPEG video is typically transmitted at a rate of 30 frames/sec. This translates to an interframe time of 33.33 msec. To model this application in OPNET, a constant packet interarrival PDF was constructed with a mean packet interarrival time of 33.33 msec.

Although MPEG frames can be transmitted in any order, most implementations send them in some fixed pattern. This study assumes a GOP of 15, containing one I frame, four P frames, and ten B frames. The associated sizes for these frames are 50,000 bytes, 25,000 bytes, and 10,000 bytes, respectively. Table 3-2 depicts the statistical properties of this MPEG video.

Table 3-2 VBR Video Traffic Properties

VBR Video Traffic Statistical Properties	
Mean Packet Interarrival Time	33.333 milliseconds (30 packets/sec)
Packet Interarrival PDF	constant
Packet Size	10,000, 25,000, or 50,000 Bytes
Packet Size PDF	Discrete: $P(10,000 \text{ B}) = 0.6667$ $P(25,000 \text{ B}) = 0.2667$ $P(50,000 \text{ B}) = 0.0667$

To simulate these frame sizes along with their probability of occurrence in OPNET, the PDF shown in Figure 3-7 was created. This PDF creates frames of variable size with the occurrence probabilities shown, and at a fixed interarrival time of 33.33 msec.




Figure 3-7 VBR Video Packet Size PDF

As mentioned above, MPEG frames are typically generated in a fixed pattern. However, this is not an absolute requirement with respect to the MPEG standard. In fact, more advanced encoders attempt to optimize the placement of these frames according to local sequence characteristics in the context of more global characteristics [Fog98]. Therefore, this OPNET model generates I, P, and B frames randomly, but with the same probability of occurrence as depicted in Figure 3-7. This enables the use of OPNET's built-in traffic generator and its associated statistics. Finally, this model creates an average data rate of 4 Mbps, which is consistent with the MPEG test model's sequence [Fog98].

3.5.2 Native ATM Specific Components

After organizing the model architecture such that all non-protocol specific items were common, three models were created where the only difference in construction was the implementation of either a Native ATM, an IP over ATM, or a LANE protocol. The protocol stacks operate by passing data down from higher layers for transmission to other nodes, and receive data up from lower layers as information arrives. If not destined for a

given node, the switching modules at the lower layers pass the data until its arrival at its destination.

The Native ATM model was implemented in OPNET as shown in Figure 3-8. The application layer is implemented by the client module, which generates DIS and Video traffic as described above. The ATM adaptation layer and its application program interface (API) are implemented by the tpal, tpal_if, and AAL OPNET modules. The user and management planes of the ATM layer are implemented by the ATM_mgmt and ATM_layer modules, and the transmission and reception of ATM cells is accomplished in the ATM_trans module via point-to-point receivers and transmitters (pr_* and pt_* modules, respectively).

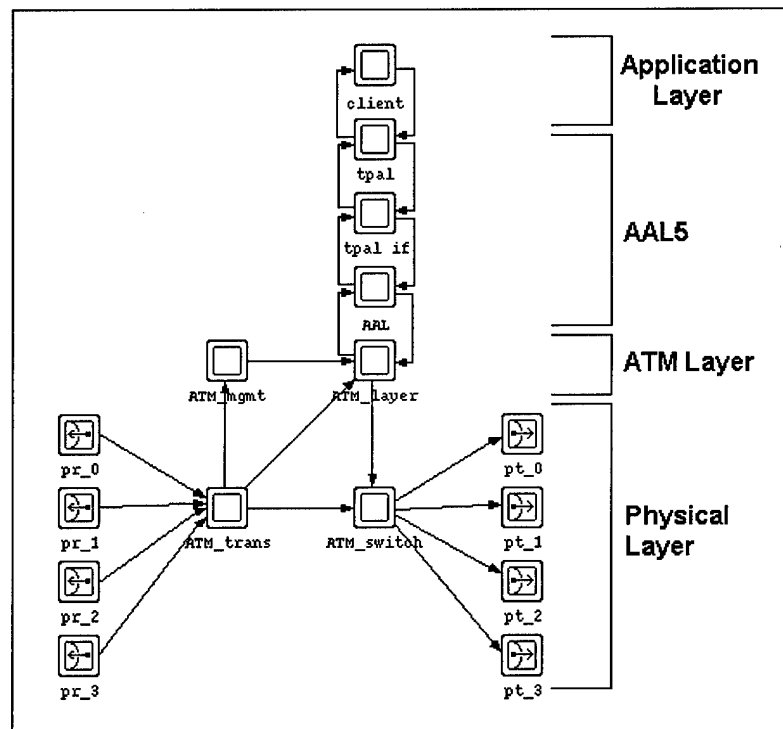


Figure 3-8 Native ATM in OPNET

The connections between DIS and Video Clients and their respective destination servers were accomplished using an N^2 mesh of point-to-point connections. This mesh

was implemented because the current version of OPNET does not support point-to-multipoint connections. For this study, an N^2 mesh is a reasonable solution since the VBE network only consists of 9 end nodes. An N^2 approach, however, limits the growth and scalability of a solution, since the number of connections required to support a large number of nodes would rapidly exhaust system resources and addressing capability. It should be noted that, since all proposed solutions operate on top of this mesh of ATM virtual circuits, this scalability issue is common to all solutions. It is anticipated that future versions of OPNET will provide the capability to model multicasting in ATM to mitigate this weakness.

3.5.3 IP over ATM Specific Components

The OPNET implementation of this stack is shown in Figure 3-9. The AAL, ATM, and physical layers are implemented as before in Native ATM, but now the IP (network) layer is implemented with the `ip_encap`, `ip`, and `ipal` OPNET modules. Although the TCP and OSPF (open shortest path first) modules are shown in the OPNET stack, they are not used. Since low delay variance is more important than reliability in real time applications like DIS and Video [PuW95], UDP is used instead of TCP. This is because TCP's reliable service incurs unacceptable delays for retransmissions, whereas UDP's connectionless operation is more suitable for achieving low delay variance. RIP (routing information protocol), a form of Bellman-Ford routing, is shown above the UDP module, and is used to implement the routing function for the UDP/IP stack. Finally, the `tpal` module is used to implement the API as in Native ATM.




Figure 3-9 IP over ATM in OPNET

3.5.4 LANE Specific Components

The OPNET implementation of LAN Emulation is shown in Figure 3-10. Again the lower layers (AAL5, ATM, and physical layers) are the same as in the previous protocols, but now the LANE layer is implemented with the LANE, LANE_if, and arp OPNET modules. These layers emulate the MAC layer of the Ethernet protocol, allowing the higher layers to operate as if all remote stations were located on a single LAN. The remainder of the upper level modules (UDP, RIP, etc.) operate in the same

manner as described in IP over ATM above. UDP is used for LANE instead of TCP for the same reasons as IP over ATM in Section 3.5.3. It is further assumed that the workstation clients and servers are emulating an IEEE 802.3 Ethernet LAN (not a token ring LAN).

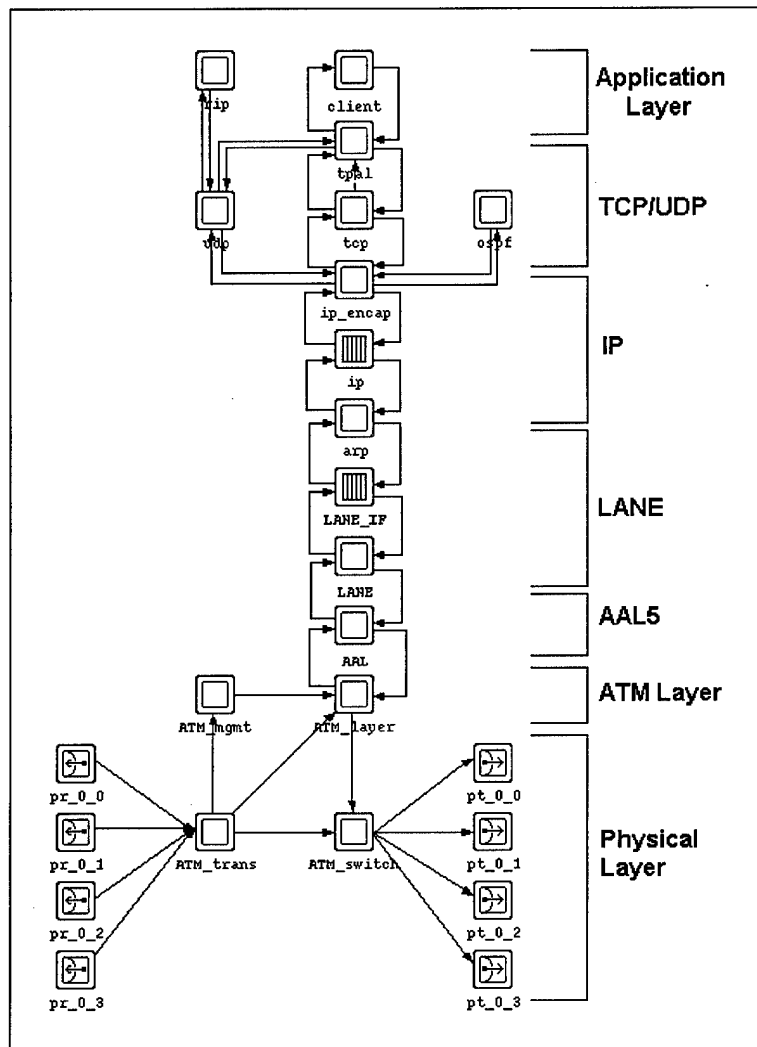


Figure 3-10 LANE in OPNET

Although not pictured, one other difference exists in OPNET between the LANE model and the other two. The LES and BUS servers that LANE requires were added to

the site model of the Space Battlelab. Although somewhat arbitrary, this location was selected since it resides at a central geographic location in the VBE.

3.5.5 System Utility Function

After defining the proposed models and the measures of performance by which they are compared, the relative value to the CDM of a particular measure's outcome must be defined. This information is captured in a *system utility function*. The system utility function is an explicit representation of the relative ratings assigned by the CDM to each level of performance in each measure. In addition, this utility function is used to translate the raw performance data values into relative scores for comparison purposes. A 0-10 scale is used (10 representing the best relative value). The system utility function used for this study is shown below in Table 3-3.

Table 3-3 System Utility Function

	Measures of Performance	MEASURE / SCORE											UNITS
		0	1	2	3	4	5	6	7	8	9	10	
PERFORMANCE	MOP #1: Avg ETE Delay	300	270	240	210	180	150	120	90	60	30	0	msec
	MOP #2: Avg CDV	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0	msec
	MOP #3: CLR	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	10^{-10}	$<10^{-11}$	ratio
	MOP #4: P(Delay < 100ms)	0	10	20	30	40	50	60	70	80	90	100	%
COST	MOP #5: Implement Cost		H				M				L		L,M,H
	MOP #6: BW Utilization Cost	100	90	80	70	60	50	40	30	20	10	0	%

The rationale behind these selections is based upon the useful bounds of these quantities and the best judgement of the authors (since input from the CDM was unavailable). For average ETE delay, only values below 300 msec have value, since this is the upper bound on acceptable delay. The worst-case CDV was selected to be 1 msec, which is the CDV Tolerance (CDVT) for highly compressed MPEG Video [Koe94]. Target values for CLR for most services in ATM networks is on the order of 10^{-6} [YeY94], so this is selected as the median value on the utility scale. The utility increments increase and decrease logarithmically to the bound shown. Cell losses should remain low in the proposed models since the only source of cell losses would be from buffer overflows (transmission bit errors are assumed to be zero). CLR values follow an inverse logarithmic scale with CLR values below 10^{-10} rated with a value of 10. A direct percentage scale is used to map probability scores for the 100 msec delay bound MOP into utility scores. Implementation cost estimates of low, medium, and high are mapped into the values of 1, 5, and 9 respectively, to evenly distribute the scoring of these estimated costs. Finally, bandwidth utilization costs are inversely mapped (by percentage) into the utility scores, because with a bandwidth-on-demand pricing scheme costs increase with rising utilization levels.

The accuracy of these score mappings may be enhanced by using more accurate mapping models, to include both linear and non-linear approaches. In addition, the CDM's inputs (if available) would likely alter the above utility function, yielding potentially different analysis results. Therefore, every attempt should be made to obtain the CDM's inputs for this analysis.

3.6 Summary

This chapter has followed the first steps of a systems engineering approach to define the problem to be solved, detail the criterion by which solutions will be judged, synthesize possible approaches, and define a model which will be used to evaluate the chosen solutions. The proposed solutions consist of using Native ATM, IP over ATM, and LANE protocol architectures to provide connectivity for the Air Force Battlelabs' VBE. A common model environment was set up with only the proposed solution protocols varying between three models. The performance of these models is evaluated and compared in the Chapter 4 by utilizing the remaining systems engineering steps.

CHAPTER 4

ANALYSIS

4.1 Analysis Overview

The remaining two steps of the systems engineering process described in Chapter 3 are implemented in this chapter to evaluate and analyze the outputs of the proposed model solutions. Section 4.2, System Evaluation, begins with an explanation of the procedures used to achieve a specified level of statistical significance. The verification of model performance within specified confidence and error bounds is then addressed, followed by the raw performance data output by those models. Section 4.3 then presents the Decision Analysis process. The proposed models are then scored and recommendations are made based upon these scores.

4.2 System Evaluation

4.2.1 Statistical Significance

Since the random numbers generated by computers are not truly random, simulation results are somewhat dependent on the initial seeding of the random number generator. In addition, the choice of a particular seed could yield uncharacteristic model results. Therefore, it is important to determine the number of runs necessary to obtain a certain level of confidence in simulation results.

4.2.1.1 Central Limit Theorem

To determine the number of runs necessary to obtain a certain level of confidence, it is important to understand the central limit theorem. This theorem states the following:

“if random samples of n observations, y_1, y_2, \dots, y_n are drawn from a population with finite mean μ and variance σ^2 , then, when n is sufficiently large, the sampling distribution of the sample mean can be approximated by a normal density function with mean $\mu_{\bar{y}} = \mu$ and $\sigma_{\bar{y}} = \sigma / \sqrt{n}$ ” [MeS92].

However, running n simulations, where n is sufficiently large ($n \geq 30$) may not be feasible. Therefore, it is sometimes necessary to use the *Student's t* distribution when our sample size (number of simulation runs) is small.

4.2.1.2 Student's-t Approximation

The *Student's t* distribution resembles the normal distribution in its bell-shaped curve. However, this distribution is based on the sample variance as opposed to the known or assumed population variance. In addition, the *Student's t* is dependent on the number of degrees of freedom ν , where $\nu = (n-1)$, and n is the sample size or number of runs. It is described by the following equation [Gla98]:

$$n = \left(\frac{t_{\frac{\alpha}{2}} \cdot s}{error} \right)^2 \quad 4-1$$

where:

- α = level of confidence index
- $t_{\frac{\alpha}{2}}$ = *t*-distribution value (based on α and number of degrees of freedom)
- s = sample standard deviation
- *error* = error tolerance.

4.2.1.3 Determination of Confidence Intervals

The number of runs (samples) required to attain a certain level of confidence, along with an associated error tolerance, can be found by solving the *Student's t* equation iteratively. First a small set of n samples is run. The sample standard deviation s of these n runs is then used with the desired level of confidence and error tolerance to determine the required sample size n' . If n' is less than or equal to n , then the number of runs n is sufficient. If not, then $n' - n$ additional runs must be made to attain required sample size.

4.2.2 Model Verification

4.2.2.1 Pareto Interarrivals

To verify that the models under test are operating properly, their main components' functions should be verified. The first component tested was the random generator that outputs self-similar traffic. Since OPNET's built-in probability densities were not adequate to generate self-similar traffic, a stable pareto distribution had to be built using the PDF editor. This pareto distribution was used to generate self-similar interarrivals with a desired mean of .592 seconds.

Using the development of Section 4.2.1, $n = 5$ runs were executed for a simple traffic source with the interarrivals following the constructed pareto PDF. The statistic of interest, mean interarrival time, was gathered over the runs and produced the results summarized in Table 4-1.

For reasons to be explained in Section 4.2.2.4, a 90% confidence level was selected for $n = 5$ runs, making $t_{\alpha/2} = 2.132$. Solving for the error tolerance, these five

runs attain an error of 4.19%. Since the target mean (0.592 msec) is within 3.7% of the sample mean, the pareto distributed interarrivals are assumed to be correct with a 90% confidence level. One sample's average interarrival time is shown below in Figure 4-1 plotted over a 50 second time interval.

Table 4-1 Pareto Interarrivals Sample Set Data

Run	Mean ia time (msec)	Sample Mean
1	0.5421	0.5697 msec
2	0.5986	
3	0.5730	Sample Std Dev (s)
4	0.5463	2.5063e-5
5	0.5886	




Figure 4-1 Average Pareto Interarrival Time

In addition to mean interarrival time, the self-similarity of the source should be verified. Although more rigorous statistical tests exist for verifying self similarity, a simple pictorial method can be used and is adequate to demonstrate self-similar behavior. This was accomplished by using a "pictorial proof" [LeT95] to show bursty

characteristics when viewed over wide-ranging time scales. For the pareto traffic source, output interarrivals were gathered at random intervals over four time scales, and are shown below in Figure 4-2.

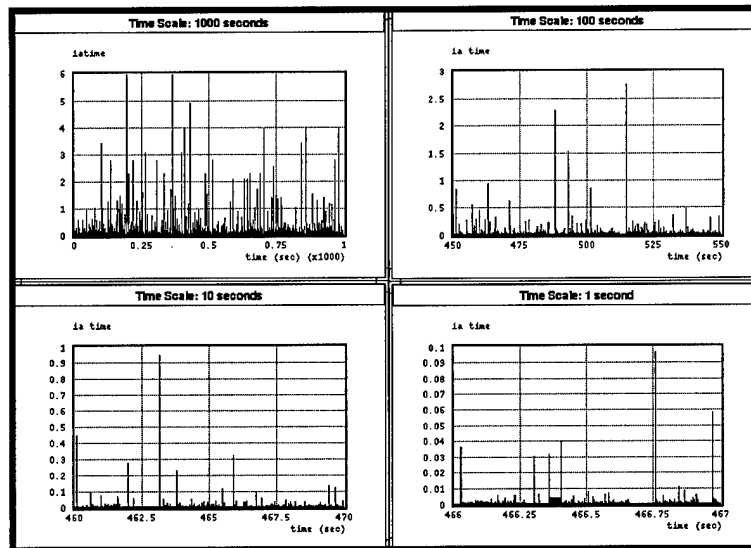


Figure 4-2 Pictorial "Proof" of Self-Similarity (Time Traces)

Figure 4-3 depicts the histogram of each respective time trace and is included to show that these traces are similar in a distributional sense.

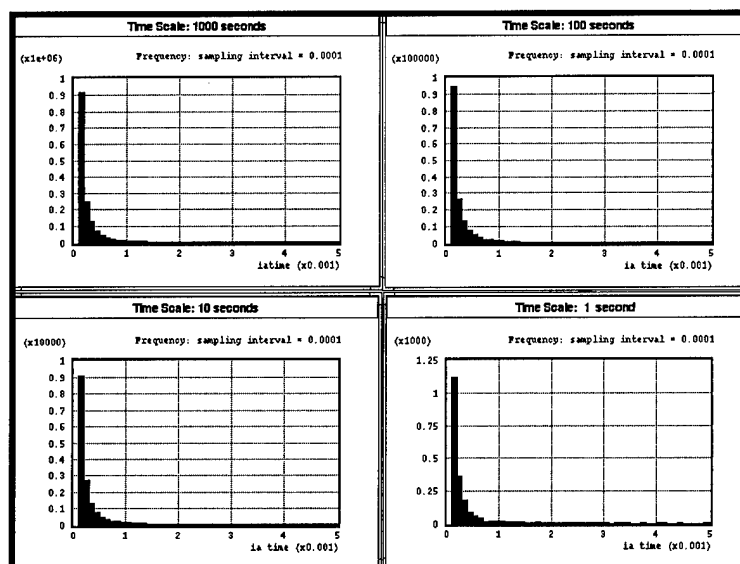


Figure 4-3 Self Similar Distributions (Histograms)

4.2.2.2 DIS Source

Following on the verification of self-similar generation of interarrival times, the generation of packets by the DIS source was also verified. This was accomplished by running an initial sample of $n = 8$ runs, where the output of a single traffic source generating DIS traffic was probed to collect the number of packets produced over time. This output was then differentiated to give the rate of packets generated per second, and this was then averaged over time. An example output of one of the runs is shown below in Figure 4-4.




Figure 4-4 DIS Source: Average Packets/Second

The statistic of interest here is mean packet generation rate, and the sample set results are summarized in Table 4-2.

For $n = 8$ runs, a 90% confidence level yields $t_{\alpha/2} = 1.895$. This gives an error tolerance of 8.48% for these eight runs. Since the desired average packet generation rate of 253 packets/sec is within 6.87% of the sample mean, the DIS Source is assumed to be generating packets at the intended rate with a 90% confidence level.

Table 4-2 DIS Source Sample Set Data

Run	Mean Pkts/sec	Sample Mean
1	277.6	271.68 pkts/sec
2	276.0	
3	293.3	
4	226.0	
5	225.5	
6	272.5	
7	271.4	
8	331.2	
		Sample Std Dev (s)
		34.4 pkts/sec

4.2.2.3 VBR Video Source

In addition to DIS Sources, VBR Video sources are also used to generate traffic in the network. The verification of these sources follows the same procedure as above. First, a sample set of $n = 5$ runs was made to gather sample statistics for average bytes/second output by a single VBR source. As before, a single sample's trace of average bytes/second over time is shown below in Figure 4-5.




Figure 4-5 VBR Video Source: Average Bytes/Second

The desired outcome for average bytes/second output from the VBR source is 500 kBytes/second. The sample set data for all VBR Video sample runs is summarized in Table 4-3.

Table 4-3 Pareto Interarrivals Sample Set Data

Run	Mean Bytes/Sec	Sample Mean
1	526,300	504,452 bytes/sec
2	486,400	
3	495,960	Sample Std Dev (s)
4	492,900	17,838 bytes/sec
5	520,700	

As in Section 4.2.2.1, a 90% confidence level and $n = 5$ runs gives $t_{\alpha/2} = 2.132$. The error tolerance in this case is determined to be 3.32%. As before, since the target average generation rate of 500 kBytes/second is within 1% of the sample mean, the VBR video source is assumed to be generating bytes at the intended rate within 10% error with a 90% confidence margin.

4.2.2.4 System Verification

After verifying the individual components that comprised the common system components, each solution (Native ATM, IP over ATM, and LANE) had to be assembled, and their outputs verified to determine statistical significance and confidence levels. This was accomplished using the same procedure as the verification of the source components. The average ETE delay was selected as a representative measure of system performance, and was used to obtain an initial estimate on the number of runs needed to achieve a 90% confidence interval. This decision to use this statistic was based on the fact that ETE delay was the only elimination criteria statistic.

First, an initial estimate on the feasible number of runs (each of five second durations) was determined to be $n = 5$. This number was based on the test simulation times and output vector file sizes shown in Table 4-4.

Table 4-4 Simulation Resource Requirements

Model	Average Time Required to Simulate 1 second (hours)	Average Output Vector Size (Mbytes)
Native ATM	2	156
IP over ATM	3	156
LANE	3	156

These runs were executed for each model. Next, the required number of runs, n , was calculated over a range of α and *error* values and plotted for each model. A complete report of these run outcomes can be found in Appendix A. The results for the LANE (worst-case) model are depicted in Figure 4-6. Based on these results, a 90% confidence level was chosen.




Figure 4-6 LANE Required Runs

This confidence level was chosen based on the:

- number of runs already completed
- time to complete a given run
- number of reruns due to errors/losses
- number of runs necessary to obtain a given level of confidence.

Since LANE displayed worse-case behavior in both simulation time and number of runs required for a given confidence level, this model was used as the critical path in determining a confidence/error combination. This 90% confidence level required two runs for Native ATM (all five are still used in analysis), seven for IP over ATM, and nineteen for LANE.

To have obtained a 95% confidence level, approximately three more IP over ATM runs and nine more LANE runs would have been needed. However, these numbers are only hypothetical. That is, these numbers would only hold if the following were true:

- All runs were error-free (no lost data files or faulty seeds)
- Runs could be farmed out to multiple machines
- At least 50% of the CPU was available for a given run
- Adequate disk space was available to accommodate huge output files
- No power outages were experienced.

In obtaining a 90% confidence level, though, several of the required runs were lost or contained faulty seeds. More specifically, approximately 10% of all simulation runs just disappeared for no apparent reason, while another 20% had seed errors (both problems are still open with MIL3's technical support personnel). In addition, another twenty runs were lost due to inadequate disk space. However, the error file did not give sufficient details as to why the output file was missing or corrupted. This problem has

been noted and can be controlled in the future. Finally, another fifteen runs were lost due to an unexpected power outage. Table 4-5 forecasts the number of runs and resources that would have been required to obtain a 95%/10% confidence/error combination after adjusting for the losses and faulty seeds previously experienced.

Another item that needed to be taken into account was the time required to collect and process the results for a given simulation run. This variable time would have been in addition to the results from Table 4-5. The best-case scenario would have required approximately two more weeks of work, which was not feasible due to project time constraints.

Table 4-5 95%/10% Requirements

Model	Number of Error-free runs	Re-runs due to Losses/ Faulty Seeds	Total Runs Required	Total Duration to Run 5 second Simulations	Total Disk Space Needed
IP over ATM	3	1	4	60 hrs	600 MB
LANE	9	2	11	165 hrs	1.65 GB

4.2.3 Raw Performance Data

After verifying the number of runs required to attain a 90% confidence level, these simulations were run and the measures of performance were gathered for each. This section reports the output of this data collection beginning with the delay requirement in Section 4.2.3.1, followed by that for each MOP in Sections 4.2.3.2 through 4.2.3.7. A brief summary of this raw output data is then presented in Section 4.2.3.8. Where applicable, error tolerances are shown with their respective mean values using confidence interval bars. Appendix B contains a complete record of the simulation

data gathered and analyzed in this section, and Appendix C contains the actual means and error values for each MOP.

4.2.3.1 ETE Delay Requirement

For a given solution to be considered viable, it must only pass the single requirement of this study: the ETE delay of packets from any client's application layer to that of any server (destination client) must be less than 300 milliseconds with a probability of 99%. As shown in Table 4-6, every proposed solution met this delay constraint.

Table 4-6 ETE Delay Constraint

Solution	Prob(ETE Delay < 300 ms)	%
Native ATM	1	100
IP over ATM	0.993428	99.34
LANE	0.994641	99.46

4.2.3.2 Average ETE Delay

The results of this first MOP are shown below in Figure 4-7, showing the Native ATM performed with average delays about two orders of magnitude lower than IP over ATM or LANE ($0.748 \text{ msec} \pm 5.85 \text{ } \mu\text{sec}$ versus $65.14 \text{ msec} \pm 6.47 \text{ msec}$ and $75.78 \text{ msec} \pm 4.49 \text{ msec}$, respectively).

This is primarily due to fact that IP over ATM and LANE impose larger overhead requirements when encapsulating and segmenting data. Traffic generated at each source is segmented into IP frames in the IP over ATM protocol, with a max IP transfer unit of 9180 bytes, and then segmented again at the AAL layer for ATM transport. In LANE, the source data is first segmented into IP frames with a max size of 1500 bytes (IEEE

802.3 Ethernet), then again at the LANE layer to add the LAN Emulation header (max frame size of 1516 bytes), and finally at the AAL5 layer for ATM transmission. Native ATM avoids this overhead by accepting data in frames up to 64 kBytes and directly segmenting them for ATM transport.




Figure 4-7 Average ETE Delay Performance Comparison

4.2.3.3 Average CDV

The average cell delay variation, or measure of cell “clumping,” is shown for each protocol solution in Figure 4-8. Although it appears that IP over ATM operates with the lowest CDV, all CDV values for the proposed solutions are extremely small (on the order of microseconds). The maximum observed CDV for any simulation was only measured at 0.13 msec, well below the most stringent CDV tolerance of 1 msec for highly compressed video traffic. This is most probably due to the fact that no contention traffic exists in the network, allowing cells to be serviced much quicker than if there had been background traffic.




Figure 4-8 CDV Performance Comparison

4.2.3.4 CLR

There were no cell losses in any of the simulations performed, regardless of the protocol. This resulted in CLRs equal to zero for each solution. This can be attributed to two main factors: lack of contention traffic, and large buffer sizes in all ATM switches. The scenarios simulated in this study do not account for the quantity or distribution of contentious traffic that exists in real-world operations. Background traffic was omitted to determine the “best case” performance of each tested implementation. Since only VBE traffic is modeled, the buffer sizes of the ATM switches in the network were more than adequate to accommodate bursts of cells during transmission. The actual buffer sizes were assumed to be 65536 cells, which was validated as realistic from [DIS98a] and [FOR99].

4.2.3.5 Tightly-Coupled Delay Measure

The results of how well a solution met the tightly-coupled delay bound of 100 milliseconds are shown in Figure 4-9.




Figure 4-9 Tightly-Couple Delay Comparison

Expressed as the probability that the delay of any given cell is less than 100 ms, Native ATM clearly performs best among the proposed protocols. Like average ETE delay, these probabilities are likely dependent on the overhead requirements of each protocol.

4.2.3.6 Implementation Cost

Although not a direct product of the simulations themselves, the estimated implementation costs must be determined to accurately discern the overall cost of a given solution. The relative level of costs are determined by the rationale that the emerging and largely unimplemented Native ATM technology would incur proportionally large costs when compared to its competitors. IP over ATM has seen wide use in carrying internet traffic over high-capacity backbones, and LANE emulation by definition was created to maximize the reuse of legacy LAN infrastructure and operating systems [ATM95]. Therefore the implementation cost score are Native ATM = H, IP over ATM = L, and LANE L.

4.2.3.7 Bandwidth Usage Cost

The time-averaged bandwidth utilization of all links that connect each VBE site to the ATM backbone was collected for each protocol, and the results averaged over all simulations is shown below in Figure 4-10. It is evident that Native ATM utilizes the least amount of bandwidth, with an average usage of 31.79% for all access links. IP over ATM averages 44.19% and LANE averages 46.02% for utilization over all VBE access lines. The additional overhead associated with the IP over ATM and LANE protocols are most likely the cause of their higher utilization.




Figure 4-10 Bandwidth Utilized

4.2.3.8 Performance Summary

Table 4-7 summarizes the results of the raw performance data for each solution model. This data in its current tabular format, is used to further determine the overall utility of each proposed protocol.

Table 4-7 MOP Results Summary

	Native ATM	IP over ATM	LANE
Avg ETE Delay	7.48E-04	6.51E-02	7.58E-02
Avg CDV	4.04E-06	2.11E-06	3.89E-06
CLR	0.0	0.0	0.0
Prob (D < 100)	1.00	0.78	0.72
Implement Cost	H	L	L
BW Cost	31.79	44.19	46.02

4.3 Decision Analysis

After the performance of each model is determined and analyzed with respect to the individual MOPs, the overall performance of the solutions can be compared by using an evaluation method given by [Ath82]. This method builds upon the value system defined in Chapter 3, and utilizes the System Utility Function of Section 0. Again, the involvement of the CDM is important in the remainder of this analysis. Since the user's requirements were not fully defined at the time of this thesis, many of the inputs to this analysis were justified assumptions based upon research and standards, where available.

It is important to note again, however, that this process is not free from subjectivity. That is, differing user inputs will drive different analysis results. Therefore, it is important that the inputs be as concise as possible to achieve optimal results. However, even if the applicability of the analysis outputs is not optimal (based on less than perfect inputs), this systems analysis process still provides a valuable decision-making framework from which systems can be evaluated and compared.

4.3.1 Raw Score Matrix

In the first step of the evaluation process, the system utility function defined in Chapter 3 is applied to the raw scores output for each MOP in Section 4.2.3.8. This

utility function is shown again in Table 4-8 below for convenience. Applying this function to the raw scores produces a *Raw Score Matrix, R*. The application of the utility function to the raw MOP scores is fairly straightforward, with values that fall between defined scoring bounds interpolated in an appropriate manner (linearly or logarithmically). MOP #5, Implementation Costs, is only assigned a score of 1, 5, or 9, since the raw score assigned to this MOP was only determined in the broad terms of a low, medium, or high rating. The Raw Score Matrix is shown below in Table 4-9.

Table 4-8 System Utility Function

	Measures of Performance	MEASURE / SCORE											UNITS
		0	1	2	3	4	5	6	7	8	9	10	
PERFORMANCE	MOP #1: Avg ETE Delay	300	270	240	210	180	150	120	90	60	30	0	msec
	MOP #2: Avg CDV	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0	msec
	MOP #3: CLR	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	10^{-10}	$<10^{-11}$	ratio
	MOP #4: P(Delay < 100ms)	0	10	20	30	40	50	60	70	80	90	100	%
COST	MOP #5: Implement Cost		H				M				L		L,M,H
	MOP #6: BW Utilization Cost	100	90	80	70	60	50	40	30	20	10	0	%

Table 4-9 Raw Score Matrix

	Native ATM	IP over ATM	LANE
Avg ETE Delay	9.98	7.83	7.47
Avg CDV	9.96	9.98	9.96
CLR	10.00	10.00	10.00
Prob (D < 100)	10.00	7.82	7.24
Implement Cost	1.00	9.00	9.00
BW Cost	6.82	5.58	5.40

4.3.2 Confidence Matrix

After determining the raw performance scores, a measure of fidelity of each of these points should be determined. The assessment of the quality of each data point is required to judge the overall confidence of each given solution. A combination of judgement about each model's fidelity and the confidence interval for each MOP is used to assign an estimated confidence score to each element. A low (0.3), medium (0.6), and high (0.9) measure of relative confidence was used in the scoring of the protocol models since the resolution of this assessment was not very fine. Scores for each MOP are selected using the relative confidence interval achieved over all simulation runs, providing a little better granularity (values of 0, 0.1, 0.2, . . . , 1.0). These interval values can be found in Appendix C.

Since Native ATM is the only protocol that has yet to be defined in detail, all scores in the Native ATM column were weighted with a 60% (0.6) confidence value. The IP over ATM and LANE model suites in OPNET closely follow defined real-world specifications, and hence all confidence scores in those columns are weighted with a 90% (0.9) score.

Average ETE delay confidence scores were selected to be 0.9, 0.6, and 0.8 for Native ATM, IP over ATM, and LANE, respectively. This was determined based on the relative size of the error margins for these three scores. CDV scores were determined to be 1.0, 0.9, and 0.8, respectively, using similar judgement as applied for average ETE delay. The cell loss MOPs were all assigned values of .6 since the model used did not include contention traffic, the primary source of buffer overflows. The tightly-coupled delay bound MOP for Native ATM was assigned a score of 1, since it achieved this delay

bound with a probability of 100%. The IP over ATM and LANE MOPs were both given scores of 0.8 since they achieved very comparable scores for this MOP. The implementation cost MOPs were scored based upon the level of confidence in their raw cost scores. Since these raw scores were only specified in rough terms (low, medium, and high), it was assumed that a confidence of only 0.6 is achievable by these estimates. Finally, the bandwidth utilization cost MOPs were all assigned scores of 0.9 due to their tight error bounds. The *confidence matrix*, *C*, is shown below in Table 4-10, with each element's confidence score computed by multiplying its respective model confidence score (one value for each column) by its MOP confidence value.

Table 4-10 Confidence Matrix

	Native ATM	IP over ATM	LANE
Avg ETE Delay	0.54	0.54	0.72
Avg CDV	0.6	0.81	0.72
CLR	0.36	0.54	0.54
Prob (D < 100)	0.6	0.72	0.72
Implement Cost	0.36	0.54	0.54
BW Cost	0.54	0.81	0.81

4.3.3 Weighting Vectors

At this stage, the CDM must determine the relative weight of each measure of performance. These weights are a determination of the relative importance of each measure according to the needs and desires of the CDM. Since the MOPs are hierarchical in nature, a worth analysis is used instead of a direct comparison to determine their values. Relative weights for any given level in the hierarchy sum to one. After the CDM determines the weights of each element in the hierarchy, the weight of

each MOP is determined by multiplying the decimal scores down the hierarchical chain (or across if tabular).

Since the CDM's actual weights are not available, three different versions of the weighting vector are used for this analysis. These versions are for a CDM who is a) primarily performance-driven, b) evenly-driven (cost and performance), and c) primarily cost-driven. These are reflected in the weights of the first tier of the weighting matrix, and these values produce three different weighting column vectors; W_p , the *performance-weighted vector*, W_e , the *evenly-weighted vector*, and W_c , the *cost-weighted vector*. These weight vectors are shown below in Table 4-11.

Table 4-11 Weight Matrix

WEIGHT MATRIX				W_p	W_e	W_c
TIER 1		TIER 2				
PERFORMANCE	0.7 0.5 0.3	MOP #1: Avg ETE Delay	0.4	0.28	0.2	0.12
	0.7 0.5 0.3	MOP #2: Avg CDV	0.3	0.21	0.15	0.09
	0.7 0.5 0.3	MOP #3: CLR	0.1	0.07	0.05	0.03
	0.7 0.5 0.3	MOP #4: P(Delay < 100 ms)	0.2	0.14	0.1	0.06
COST	0.3 0.5 0.7	MOP #5: Implementation Cost	0.4	0.12	0.2	0.28
	0.3 0.5 0.7	MOP #6: BW Utilization Cost	0.6	0.18	0.30	0.42

The first tier weights of 0.3, 0.5, or 0.7 were selected according to the CDM type as explained above. In the second tier of the weighting matrix, the weights that correspond

to each MOP appear. Although these are also to be determined by the CDM, they were selected on a "best information" basis since this input was unavailable.

4.3.3.1 Performance Weight Rationale

Average ETE delay was deemed the most important performance MOP, since delay is the critical statistic used to eliminate proposed solutions. Cell delay variation was given the next highest weight since it closely follows delay as a critical performance characteristic in both DIS and VBR video transmission. Cell losses are weighted least of all MOPs since real-time applications are more tolerant of isolated cell losses than of delay or CDV. Although its desirability to the CDM is unknown, the measure of a solution's ability to support tightly-coupled simulation was weighted between CDV and CLR since it is assumed that the Battelabs will have at least some interest in this enhanced capability.

4.3.3.2 Cost Weight Rationale

The cost of a complex system must consider more than just the first time cost of purchase. The life cycle cost of a system is often determined more by the recurring costs of operation and maintenance (O&M) than by the initial purchase price. In fact, O&M costs can account for well above 60 percent of the total life cycle cost [Ker98]. For this reason, the one-time implementation cost was selected to be a proportionally small percentage of the overall cost, with the O&M costs dominating. As stated in Chapter 3, the cost of bandwidth usage is assumed to dominate these O&M costs, and this category is referred to simply as the "bandwidth utilization cost."

4.3.4 Scoring of Solutions

With the completion of the matrix of raw performance scores R , the matrix of confidence scores C , and the weighting vectors W_p , W_e , and W_c , a systematic approach can now be applied to produce a set of final scores for each alternative. This approach is explained below in Section 4.3.4.1, and then executed in Section 4.3.4.2.

4.3.4.1 Scoring Procedure

Let the weight of a particular MOP (an element of a weight vector W) be represented as w . Let an element of the raw score matrix be denoted r , and its corresponding confidence score be denoted c . For each data element, we can now define u , an *undiscounted* score which does not take confidence into account. An undiscounted score is a performance measure, which is undiscounted in the sense that a measure of its validity has not yet been factored in. This u is determined by:

$$u = r \cdot w \quad 4-2$$

It should be noted that the multiplication performed here and in equations 4-3 through 4-6 is done on an element-by-element basis, not in a true “matrix multiplication” manner. After u has been computed for each data element in the matrix, a total undiscounted score U can be determined for a given solution by summing over the MOPs for that solution

$$U = \sum_{MOPs} u. \quad 4-3$$

The total undiscounted score represents an optimistic estimate of the value of a solution, or a measure of the potential value that solution offers. A *discounted* score d , which takes confidence into account, can be computed for each matrix element. These d values are determined by

$$d = r \cdot w \cdot c = u \cdot c, \quad 4-4$$

and a total discounted score for a given solution, D , is therefore given by

$$D = \sum_{MOPs} d. \quad 4-5$$

A discounted score represents a pessimistic estimate of the relative value of a solution.

Finally, an overall confidence measure C_o can be defined for a given solution by

$$C_o = \frac{D}{U}, \quad 4-6$$

where this confidence measure is typically expressed as a percentage. This is an overall measure of accuracy in rating each alternative, and relates to the range of value scores that can be expected.

By applying the methodology explained above, undiscounted scores, discounted scores, and measures of confidence can be computed for each solution to produce a summary of the relative performances of each option. These three measures can be used to compare the range of potential scores and determine the relative value of solutions.

4.3.4.2 Final Scoring

The results of final scoring are shown below according to the CDM types explained earlier in Section 4.3.3. In each case, the undiscounted score matrix is followed by the discounted score matrix, generated as explained in the previous section. The scores shown have been normalized by the process so that their numerical value shows increasing utility, with 0 the lowest or worst score, and 10 the highest, or best. The performance-weighted undiscounted and discounted analyses are shown in Table 4-12 and Table 4-13, respectively.

Table 4-12 Performance-Weighted Undiscounted Score Matrix

	Native ATM	IP over ATM	LANE
Avg ETE Delay	2.79	2.19	2.09
Avg CDV	2.09	2.10	2.09
CLR	0.70	0.70	0.70
Prob (D < 100)	1.40	1.09	1.01
Implement Cost	0.12	1.08	1.08
BW Cost	1.23	1.00	0.97
Total	8.33	8.17	7.95

Table 4-13 Performance-Weighted Discounted Score Matrix

	Native ATM	IP over ATM	LANE
Avg ETE Delay	1.51	1.18	1.51
Avg CDV	1.25	1.70	1.51
CLR	0.25	0.38	0.38
Prob (D < 100)	0.84	0.79	0.73
Implement Cost	0.04	0.58	0.58
BW Cost	0.66	0.81	0.79
Total	4.56	5.44	5.49

Table 4-14 and Table 4-15 show the evenly-weighted undiscounted and discounted analysis results, respectively.

Table 4-14 Evenly-Weighted Undiscounted Score Matrix

	Native ATM	IP over ATM	LANE
Avg ETE Delay	2.00	1.57	1.49
Avg CDV	1.49	1.50	1.49
CLR	0.50	0.50	0.50
Prob (D < 100)	1.00	0.78	0.72
Implement Cost	0.20	1.80	1.80
BW Cost	2.05	1.67	1.62
Total	7.24	7.82	7.63

Table 4-15 Evenly-Weighted Discounted Score Matrix

	Native ATM	IP over ATM	LANE
Avg ETE Delay	1.08	0.85	1.08
Avg CDV	0.90	1.21	1.08
CLR	0.18	0.27	0.27
Prob (D < 100)	0.60	0.56	0.52
Implement Cost	0.07	0.97	0.97
BW Cost	1.10	1.36	1.31
Total	3.93	5.22	5.23

Finally, the results of the cost-weighted undiscounted and discounted analyses are shown below in Table 4-16 and Table 4-17, respectively.

Table 4-16 Cost-Weighted Undiscounted Score Matrix

	Native ATM	IP over ATM	LANE
Avg ETE Delay	1.20	0.94	0.90
Avg CDV	0.90	0.90	0.90
CLR	0.30	0.30	0.30
Prob (D < 100)	0.60	0.47	0.43
Implement Cost	0.28	2.52	2.52
BW Cost	2.86	2.34	2.27
Total	6.14	7.47	7.32

Table 4-17 Cost-Weighted Discounted Score Matrix

	Native ATM	IP over ATM	LANE
Avg ETE Delay	0.65	0.51	0.65
Avg CDV	0.54	0.73	0.65
CLR	0.11	0.16	0.16
Prob (D < 100)	0.36	0.34	0.31
Implement Cost	0.10	1.36	1.36
BW Cost	1.55	1.90	1.84
Total	3.30	4.99	4.96

4.3.5 Decision Recommendations

After producing the undiscounted and discounted scores for each solution, decision recommendations can be made based upon these scores, the overall confidence level of those scores, and the value system of each type of decision-maker. The decisions recommended for the performance-oriented, evenly-oriented, and cost-oriented CDMs follow below. It is important to note that these scores only reflect relative performance, and are only intended to aid in the decision making process. Consequently, the decision-maker should not make a decision based solely on the numerical output.

4.3.5.1 Performance-Weighted Recommendation

A graphical representation of the results of the performance-weighted analysis is shown below in Figure 4-11. The results indicate the Native ATM solution has a higher potential value over the other protocols, although its unimplemented status drives a slightly lower discounted score and overall confidence. Although IP over ATM and LANE appear to assure slightly better minimal value, future definition of ATM standards will likely raise the overall value of a Native ATM solution and make its performance strengths stand out. Although all solutions are very comparable, the slight potential advantage of Native ATM would make it the recommended choice for a performance-oriented decision-maker.




Figure 4-11 Performance-Weighted Final Scores

4.3.5.2 Evenly-Weighted Recommendation

If the CDM considers the importance of cost and performance to be equal, the analysis of results produces scores as shown in Figure 4-12. These scores show relatively equal performance for LANE and IP over ATM solutions, with the IP over ATM protocol achieving a slight edge in potential gains. Native ATM would probably be eliminated in this analysis, as its undiscounted, discounted, and overall confidence scores all lag those of the other two. IP over ATM would likely be the recommended solution unless other factors not considered in this evaluation, such as policy requirements or time required for implementation, dictated otherwise.




Figure 4-12 Evenly Weighted Final Scores

4.3.5.3 Cost-Weighted Recommendation

Finally, if the CDM is primarily concerned with cost issues, the output of the analysis is as presented in Figure 4-13. The decision most evident from this analysis is that Native ATM should definitely be eliminated from consideration, due to its relatively low scores. The IP over ATM and LANE scores are very close in this analysis, with IP over ATM achieving a slight edge. Other factors not withstanding, IP would likely be the recommended choice for the cost-oriented CDM.




Figure 4-13 Cost-Weighted Final Scores

4.4 Summary

The evaluation and analysis of the solutions proposed in Chapter 3 was accomplished in this chapter. The method of achieving the desired level of statistical significance was discussed, and the verification of the models with a 90% confidence level was explained. Following this, the model output data was analyzed according to the desires of three different types of CDM, and recommendations were made. The output of each analysis indicates that Native ATM would be recommended for a performance-oriented CDM, but not for an evenly-oriented or cost-oriented CDM. IP over ATM would likely be the recommended choice for an evenly-weighted or cost-weighted CDM, with LANE a close contender in each case. Although this study may indicate a winner for a given category, it is important to note that a combination of these approaches may

be necessary to support both modern bandwidth-intensive and legacy data traffic. Finally, this analysis process demonstrates its merit in aiding decision-makers in determining solutions which best meet a defined set of criteria. With the active participation of a CDM, this process can be iterated, allowing for tradeoff analysis studies. These studies can then be used to assist decision-makers in their quest to find optimal solutions.

CHAPTER 5

CONCLUSIONS

5.1 Goal Restatement

The goal of this thesis was to evaluate the performance of wide-area networking technologies and recommend a solution to the problem of successfully supporting a specified traffic load between geographically separated Battlelabs. Assumptions were made where user requirements were not available.

5.2 Conclusions

A systems engineering approach was used to analyze the ability of three proposed systems (Native ATM, IP over ATM, and LANE) to carry specified DIS and VBR Video traffic loads. These systems were evaluated in terms of ETE cell delays, CDV, CLR, and bandwidth utilization costs. In addition, these systems were analyzed from three different perspectives (performance-oriented, evenly oriented, and cost-oriented). The Native ATM solution was recommended on a performance-weighted basis, but eliminated in the evenly-oriented and cost-oriented analyses. The IP over ATM solution was recommended for the cost-weighted and evenly-weighted analyses, with LANE a close contender in each case. It is important to keep in mind, however, that some hybrid approach may be necessary to support both modern bandwidth-intensive data and legacy traffic. Finally, the merit of the systems engineering approach in aiding decision-makers in the comparison and analysis of multiple solutions was demonstrated.

5.3 Contributions

As a result of [SOW98], AFIT is now an integral part of the Air Force's Advanced Battlelab Collaboration Integrated Product Team (ABC-IPT), which was established to design and sustain the VBE. This research lays the foundation and framework for all future AFIT work regarding the VBE. More importantly, this study gives the Battlelabs a baseline from which to work regarding their ongoing efforts to design a communications infrastructure. Finally, this thesis has reinforced AFIT's partnership with the DoD by once again demonstrating AFIT's ability to solve complex operational problems.

5.4 Suggestions for Future Work

This research effort only scratched the surface of possibilities with respect to the analysis of ATM related protocols. As more standards become available, ATM based solutions will be more prevalent. In addition, many simplifying assumptions were necessary with respect to the Battlelabs' requirements (e.g. hardware and software specifications, decision weights, etc.) to conduct this study. As these requirements become better defined, a more realistic model and a more applicable analysis will be possible. In the meantime, there are three primary areas in which this research can be expanded: traffic modeling, multicasting, and other ATM protocols.

5.4.1 Traffic Modeling

As more empirical data becomes available for the modeling of self-similar traffic, the DIS and VBR models can be improved. Also, OPNET Version 5.1 has a new feature that allows for the inclusion of background traffic. This background feature can be used

to model contention traffic from other Battlelab sources (email, ftp, telnet, etc.), as well as traffic from external sources riding on the public portion of the network. This will allow more accurate analysis in terms of bandwidth utilization and cell losses.

5.4.2 Multicasting in ATM

Currently, OPNET does not support true multicasting. Instead, an N^2 mesh had to be developed to simulate point-to-multipoint communications. While this model was feasible with the current Battlelabs' configuration, it is neither optimal nor scalable. Multicasting could be used to improve network efficiency by decreasing network traffic.

5.4.3 Other ATM Protocols

The implementation of several additional ATM protocols such as Multi-Protocol over ATM and NHRP opens the door for additional research opportunities. Both protocols are improvements over Classical IP over ATM, and LANE. These standards are still in their infancies, however, they should be available for implementation soon. Finally, a hybrid approach that combines multiple protocols could be examined to meet the existing and future needs of the Battlelabs.

APPENDIX A:

INITIAL CONFIDENCE/ERROR ANALYSIS

This Appendix contains the tabulated data from the five sample runs for each model. The mean of the means and standard deviation of the means of the ETE delay of each sample run are then used in conjunction with appropriate t-values to calculate the number of runs required to attain various confidence/ error combinations for each model. Finally, these combinations are plotted for each model.

Table A-1 Results for Native ATM Sample Runs

run	mean	std dev	mean (means)	std dev (means)
1	0.000744	0.000173	0.0007482	6.14003E-06
2	0.000749	0.000173		
3	0.00075	0.00172		
4	0.000757	0.000175		
5	0.000741	0.000173		

Table A-2 Results of Native ATM Confidence Interval Analysis

Confidence Interval	t-val	% Error	n
99%	4.604	5	1
		10	1
		15	1
		20	1
		25	1
95%	2.776	5	1
		10	1
		15	1
		20	1
		25	1
90%	2.132	5	1
		10	1
		15	1
		20	1
		25	1




Figure A-1 Confidence/Error Combinations for Native ATM

Although results in Figure A-1 have no practical application, they are shown for consistency in displaying results of the analysis.

Table A-3 Results for IP over ATM Sample Runs

run	mean	std dev	mean (means)	std dev (means)
1	0.0629	0.0516	0.07282	0.009322929
2	0.0838	0.0745		
3	0.0807	0.0776		
4	0.0647	0.0487		
5	0.072	0.07643		

Table A-4 Results of IP over ATM Confidence Interval Analysis

Confidence Interval	t-val	% Error	n
99%	4.604	5	139
		10	35
		15	16
		20	9
		25	6
95%	2.776	5	51
		10	13
		15	6
		20	4
		25	3
90%	2.132	5	30
		10	8
		15	4
		20	2
		25	2




Figure A-2 Confidence/Error Combinations for IP over ATM

Table A-5 Results for LANE Sample Runs

run	mean	std dev	mean (means)	std dev (means)
1	0.0782	0.055	0.07372	0.015235879
2	0.054	0.0463		
3	0.075	0.0601		
4	0.0952	0.0713		
5	0.0662	0.05		

Table A-6 Results of LANE Confidence Interval Analysis

Confidence Interval	t-val	% Error	n
99%	4.604	5	363
		10	91
		15	41
		20	23
		25	15
95%	2.776	5	132
		10	33
		15	15
		20	9
		25	6
90%	2.132	5	78
		10	20
		15	9
		20	5
		25	4




Figure A-3 Confidence/Error Combinations for LANE

APPENDIX B:

RAW DATA

This Appendix contains the pertinent raw data for all runs for each model. Primary statistics are ETE delay, cell delay variation, cell loss ratio, and utilization.

Table B-1 ETE Delay (secs) Data for Native ATM

run	mean	peak	std dev	Total Cells	Prob(D < 300 msec)	# of Busts	Prob(D < 100 msec)	# of Busts	mean (means)	std dev (means)
1	0.000744	0.01453	0.000173	206402	1	0	1	0	0.0007482	6.14E-06
2	0.000749	0.01423	0.000173	210885	1	0	1	0		
3	0.00075	0.01444	0.00172	204033	1	0	1	0		
4	0.000757	0.01442	0.000175	199775	1	0	1	0		
5	0.000741	0.01428	0.000173	200634	1	0	1	0		
Totals					1	0	1	0		

Table B-2 Cell Delay Variation (secs) Data for Native ATM

run	mean	peak	mean (means)	std dev (means)
1	4.01E-06	1.47E-05	4.04E-06	2.0736E-08
2	4.05E-06	2.63E-05		
3	4.06E-06	1.35E-05		
4	4.06E-06	1.53E-05		
5	4.04E-06	1.86E-05		

Table B-3 Cell Loss Ratios for Native ATM

run	CLR
1	0
2	0
3	0
4	0
5	0

Table B-4 Utilization (%) for Native ATM Links

Run	AEF	Space	IW	FP	UAV	C2	AFIT	CTC	XOR
1	31.47	32.29	31.85	32.08	32.01	32.31	32.02	31.63	32.33
2	31.92	31.51	31.92	32.14	31.98	32.97	32.13	32.38	31.66
3	31.58	31.09	32.17	31.45	31.42	31.79	31.86	31.61	31.69
4	32.38	31.53	31.89	31.54	31.55	31.41	31.09	32.69	31.92
5	32.18	30.62	30.5	31.31	31.1	32.4	31.35	31.49	32.53
Average	31.906	31.408	31.666	31.704	31.612	32.176	31.69	31.96	32.026

Average of Averages	31.79422222
---------------------	-------------

Table B-5 ETE Delay (secs) Data for IP over ATM

run	mean	peak	std dev	Total Cells	Prob(D < 300 msec)	# of Busts	Prob(D < 100 msec)	# of Busts	mean (means)	std dev (means)
1	0.0629	0.3021	0.0516	201745	0.99998017	4	0.7807	44242	0.065144	0.01116444
2	0.0838	0.41222	0.0745	204288	0.98480087	3105	0.6769	66005		
3	0.0807	0.46194	0.0776	206825	0.96283815	7686	0.7339	55036		
4	0.0647	0.285	0.0487	204145	1	0	0.7678	47402		
5	0.072	0.41242	0.07643	199650	0.98783872	2428	0.7579	48335		
6	0.05612	0.30491	0.05121	195798	0.99991318	17	0.8427	30799		
7	0.06025	0.26478	0.04548	203859	1	0	0.8105	38631		
8	0.04758	0.26089	0.0411	197063	1	0	0.8819	23273		
9	0.06627	0.29234	0.05305	201536	1	0	0.7506	50263		
10	0.05712	0.26795	0.04375	199879	1	0	0.8215	35678		
Totals		2014788	0.99342859	13240	0.781781508	439664				

Table B-6 Cell Delay Variation (secs) Data for IP over ATM

run	mean	peak	mean (means)	std dev (means)
1	2.21E-06	1.94E-05	2.11E-06	8.95731E-08
2	2.08E-06	9.96E-06		
3	2.26E-06	1.19E-05		
4	2.02E-06	6.69E-06		
5	2.04E-06	1.31E-05		
6	2.08E-06	1.33E-05		
7	2.01E-06	7.52E-06		
8	2.16E-06	1.10E-05		
9	2.06E-06	1.44E-05		
10	2.21E-06	1.52E-05		

Table B-7 Cell Loss Ratios for IP over ATM

run	CLR
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	0
9	0
10	0

Table B-8 Utilization (%) for IP over ATM Links

Run	AEF	Space	IW	FP	UAV	C2	AFIT	GTC	XOR
1	43.95	44.84	44.88	45.21	44.32	43.63	45.55	44.34	43.62
2	45.12	44.49	44.89	44.01	44.14	44.16	42.97	44.01	45.07
3	44.08	44.34	44.04	43.74	43.91	44.31	43.9	45.5	43.36
4	43.72	43.88	43.75	45.55	45.03	44.38	44.7	44.88	43.1
5	44.1	43.6	44.19	43.9	44.18	42.88	43.36	43.16	43.67
6	44.3	43.51	43.95	42.47	43.59	43.45	43.36	42.77	43.9
7	44.54	44.2	42.97	45.33	44.16	43.88	43.94	43.31	44.05
8	44.38	42.62	43.42	44.22	42.77	42.86	44.85	45.31	41.86
9	44.12	44.36	44.23	44.43	43.37	44.07	43.56	44	43.1
10	44.67	43.29	43.09	43.07	42.85	44.03	43.89	44.56	43.46
Average	44.194	44.23	44.35	44.482	44.316	43.872	44.096	44.378	43.764

Average of Averages	44.18688889
---------------------	-------------

Table B-9 ETE Delay (secs) Data for LANE

run	mean	peak	std dev	Total Cells	Prob(D < 300 msec)	# of Busts	Prob(D < 100 msec)	# of Busts	mean (means)	std dev (means)
1	0.0782	0.31363	0.055	204014	0.99941671	119	0.6826	64754	0.07577625	0.012824
2	0.054	0.33243	0.0463	193358	0.99947765	101	0.8605	26973		
3	0.075	0.34029	0.0601	197940	0.99807517	381	0.7103	57343		
4	0.0952	0.40927	0.0713	205501	0.99005358	2044	0.6042	81337		
5	0.0662	0.29374	0.05	197529	1	0	0.7698	45471		
6	0.09731	0.45852	0.07678	205209	0.97830992	4451	0.6401	73854		
7	0.09402	0.37696	0.07685	201794	0.99203643	1607	0.6102	78659		
8	0.07864	0.40379	0.06333	202817	0.99344729	1329	0.7063	59567		
9	0.06934	0.3344	0.05521	199663	0.99931885	136	0.7416	51592		
10	0.08574	0.4189	0.07351	203258	0.97907585	4253	0.6978	61424		
11	0.08757	0.34523	0.07232	201241	0.99613399	778	0.6451	71420		
12	0.0627	0.29184	0.04695	198398	1	0	0.7882	42020		
13	0.06047	0.31705	0.04666	194099	0.99990211	19	0.8149	35927		
14	0.08515	0.32531	0.06525	203147	0.99903518	196	0.6566	69760		
15	0.06421	0.34841	0.05068	199615	0.99860231	279	0.7945	41020		
16	0.05967	0.33822	0.04764	196015	0.99993368	13	0.8175	35772		
17	0.09674	0.41155	0.07063	203000	0.98693103	2653	0.5843	84387		
18	0.06892	0.33074	0.05396	200694	0.9994818	104	0.7686	46440		
19	0.07907	0.3906	0.07007	199863	0.98929267	2140	0.7104	57880		
20	0.07984	0.38766	0.06711	203672	0.98688087	2672	0.7144	58168		
21	0.08094	0.36885	0.06472	205157	0.99491609	1043	0.7285	55700		
22	0.07292	0.32884	0.05521	204272	0.99943213	116	0.7268	55807		
23	0.06515	0.35016	0.05716	198564	0.99574948	844	0.8073	38263		
24	0.06163	0.36927	0.05315	199341	0.99728606	541	0.8269	34505		
Totals					4818161	0.99464132	25819	0.724367243	1328043	

Table B-10 Cell Delay Variation (secs) Data for LANE

run	mean	peak	mean (means)	std dev (means)
1	2.78E-06	2.91E-05	3.89E-06	7.5507E-07
2	5.32E-06	8.42E-06		
3	4.15E-06	7.19E-05		
4	3.55E-06	8.00E-05		
5	3.47E-06	4.51E-05		
6	3.17E-06	4.17E-05		
7	3.64E-06	4.98E-05		
8	3.84E-06	3.80E-05		
9	3.50E-06	5.36E-05		
10	3.10E-06	4.68E-05		
11	3.84E-06	6.09E-05		
12	3.38E-06	5.16E-05		
13	3.32E-06	4.56E-05		
14	3.48E-06	4.14E-05		
15	5.09E-06	1.34E-04		
16	5.06E-06	9.33E-05		
17	4.25E-06	1.30E-04		
18	3.93E-06	1.17E-04		
19	3.14E-06	3.00E-05		
20	3.73E-06	1.05E-04		
21	3.54E-06	5.85E-05		
22	5.73E-06	1.08E-04		
23	4.46E-06	9.01E-05		
24	3.90E-06	7.90E-05		

Table B-11 Cell Loss Ratios for LANE

run	CLR
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	0
9	0
10	0
11	0
12	0
13	0
14	0
15	0
16	0
17	0
18	0
19	0
20	0
21	0
22	0
23	0
24	0

Table B-12 Utilization (%) for LANE Links

Run	AEF	Space	IW	FP	UAV	C2	AFT	CTC	XOR
1	46.65	46.22	45.99	45.71	45.75	46.78	45.52	45.75	45.73
2	45.95	44.54	46.25	44.94	44.29	45.79	46.66	45.85	46.2
3	46.78	44.98	45.04	46.94	46.16	46.69	47	45.52	46.37
4	46.49	46.77	47.58	46.52	45.89	46.71	46.57	46.22	45.82
5	46.15	45.45	45.15	46.11	45.89	46.18	45.01	46.43	46.25
6	45.8	46.07	46.07	45.49	46.92	46.86	47.21	46.75	46.34
7	47.44	47.13	46.72	46.51	44.81	46.6	46.01	46.02	46.5
8	45.52	44.97	46.5	45.15	47.15	46.64	46.1	46.84	46.06
9	46.32	45.92	45.43	44.74	45.02	45.88	46.45	46.53	45.9
10	46.9	45.15	46.15	46.48	46.34	46.36	46.42	46.46	47.93
11	46.38	45.79	46.77	45.49	45.27	46.94	47.05	46.59	46.9
12	46.42	45.6	45.67	45.09	46.08	45.36	45.51	46.14	45.85
13	45.46	45.28	44.81	44.24	44.93	43.35	46.2	46.39	45.47
14	45.9	44.81	46.23	45.04	46.88	47.1	47.03	45.28	45.65
15	46.86	46.22	46.63	45.53	45.96	45.8	46.03	46.53	46.17
16	46.61	45.44	45.13	46.36	44.8	45.56	47.13	47.24	45.4
17	45.26	45.86	45.14	45.79	46.27	46.17	46.28	46.58	45.05
18	46.63	46.43	46.34	44.91	45.68	46.95	46.23	46.32	46.35
19	46.38	45.07	45.84	46.61	45.95	46.56	45.13	45.64	46.93
20	46.2	46.13	45.16	45.35	46.49	46.08	45.66	45.87	46.06
21	46.45	46.83	45.1	47.57	46.83	47	45.96	45.78	46.08
22	47.05	46.08	46.76	46.27	45.09	47.43	44.55	46.08	46.37
23	45.47	46.34	44.23	44.38	46.06	44.84	45.02	46.22	45.88
24	45.67	44.32	46.3	45.69	45.95	45.79	45.64	46.15	45.8
Average	46.28083333	45.725	45.87458333	45.7045833	45.8525	46.22583333	46.09875	46.2158333	46.2158333

Average of Averages 46.02152778

APPENDIX C:

FINAL CONFIDENCE/ERROR ANALYSIS

The following tables present the results of the final confidence/error analysis for each statistic after completing the predicted number of runs needed to achieve a 90% confidence level.

Table C-1 Average ETE Delay Error Tolerance

	mean	error	% error
Native ATM	7.48E-04	5.85E-06	1
IP over ATM	6.51E-02	0.006471	10
LANE	7.58E-02	0.004487	6

Table C-2 Average CDV Error Tolerance

	mean	error	% error
Native ATM	4.04E-06	1.98E-08	0
IP over ATM	2.11E-06	5.19E-08	2
LANE	3.89E-06	2.64E-07	7

Table C-3 Prob(Delay < 100 msec) Error Tolerance

	mean	error	% error
Native ATM	1.0	0.0	0
IP over ATM	7.82E-01	0.034109	4
LANE	7.25E-01	0.026749	4

Table C-4 Bandwidth Utilization Error Tolerance

	mean	error	% error
Native ATM	31.8	0.229725	0.72
IP over ATM	44.0	0.204828	0.47
LANE	46.0	0.121975	0.27

BIBLIOGRAPHY

- [All95] A. Alles, "ATM Internetworking," Cisco Systems White Paper, <http://www.cisco.com/warp/public/614/12.html>, May 1995.
- [Ath82] T. H. Athey, Systematic Systems Approach: An Integrated Method for Solving Systems Problems, Prentice Hall, Englewood Cliffs, New Jersey, 1982, pp. 204-234.
- [ATM95] The ATM Forum Technical Committee, "LAN Emulation Over ATM, Version 1.0, af-lane-0021.000" <http://www.atmforum.com>, January 1995.
- [ATM98] Approved ATM Forum Specifications, <http://www.atmforum.com/atmforum/specs/approved.html>, October 1998.
- [Bat98] "Air Force Battlelabs Homepage," 1998, <http://www.hq.af.mil/xo/afbattlelab/>, 15 Jul 98.
- [Bob97] T.A. Bobo, "Message-Bundle Converting in Internet Protocol Multicast Based High Level Architecture Exercises," MS Thesis, AFIT/GE/ENG/97D-20. School of Engineering, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, December 1997.
- [Cad94] Cadence Design Systems, Inc, "ATM Library Reference," Cadence, Foster City, CA, 1994.
- [DIS96] DISA, "DISN Architecture," 1996, <http://www.disa.mil/DISN/disnarch.html>, 22 Jun 98.
- [DIS98a] Department of Defense DISN ATM System Specification, ver. 1.2c, 17 April 1998.
- [DIS98b] DISA, "DISN CONUS Configuration," 1998, <http://www.disa.mil/org/disnconf.html>, 23 Jul 98.
- [FeH98] P. Ferguson and G. Huston, "What is a VPN," <http://www.employees.org:80/~ferguson/>, rev. 1, April 1998.
- [Fog98] C. Fogg, "What is MPEG?," http://bmrc.berkeley.edu/projects/mpeg/faq/mpeg2-v38/faq_v38.html, 13 July 1998.
- [FOR99] FORE Systems Specification: *ForeRunner ASX-1000 Switch*, http://www.fore.com/products/switch/dsbackbone_asx_1000_oi.html, Feb. 1999.
- [FrH96] V. Friesen, J. Harms, and J. Wong, "Resource management with virtual paths in ATM networks," *IEEE Network*, Sep./Oct. 1996, pp. 10-19.
- [Gib98] J. Gibson, et. al., Digital Compression for Multimedia, Morgan Kaufmann, San Francisco, 1998.
- [Gla98] S. R. Glasgow, "Use of OPNET in the U.S. Army TRADOC Network Requirements Study," *Technical Memorandum TRAC TM-0198*, September 1998.
- [Glo98] "Global Engagement: A Vision for the 21st Century Air Force," <http://www.xp.hq.af.mil/xpx/21/nuvis.htm>, 11 Aug 98.
- [Hal69] A. D. Hall, III, "Three-Dimensional Morphology of Systems Engineering," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-5, no. 2, Apr. 1969, pp. 156-160.
- [HoL95] R.C. Hofer and M.L. Loper, "DIS Today," *Proc. IEEE*, vol. 83, no. 8, August 1995.

-
- [Jai91] R. Jain, The Art of Computer Systems Performance Analysis, John Wiley and Sons, New York, 1991.
- [JRO95] JROCM, "DISN Mission Need Statement," 1995, <http://www.disa.mil/DISN/docs/mns.html>, 22 Jun 98.
- [Ker98] H. Kerzner, Project Management: A Systems Approach to Planning, Scheduling, and Controlling, pg. 757, John Wiley and Sons, INC., New York, 1998.
- [Koe94] J. F. Koegel, Multimedia Systems, ACM Press, New York, NY, 1994.
- [LeT95] W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson, "On the self-similar nature of ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, Feb 94.
- [LiP97] K. Liu, D. Petr, V. Frost, H. Zhu, C. Braun, and W. Edwards, "Design and analysis of a bandwidth management framework for ATM-based Broadband ISDN," *IEEE Communications Magazine*, May 1997, pp. 138-145.
- [McS95] D. McDysan and D. Spohn, ATM Theory and Application, McGraw-Hill, New York, 1995.
- [MeS92] W. Mendenhall and T. Sincich, Statistics for Engineering and the Sciences, Dellen MacMillan, San Francisco, 1992.
- [NgB96] G.D. Nguyen and S.G. Batsell, "Statistical Characteristics of DIS Traffic," *MILCOM 96*, vol. 2, pp. 637-641.
- [Org93] A. Ortega, M.W. garret and M. Vetterli, "Toward Joint Optimization of VBR Video Coding and Packet Network Traffic Control," *Proc. Fifth Int'l Workshop on Packet Video*, Berlin, Germany, March 1993.
- [PeE96] H. Perros and K. Elsayed, "Call admission control schemes: a review," *IEEE Communications Magazine*, Nov. 1996, pp. 82-91.
- [PuW95] J.M. Pullen and D.C. Wood, "Networking Technology and DIS," *Proc. IEEE*, vol. 83, no. 8, August 1995.
- [Rag98] S. V. Raghavan, Networked Multimedia Systems: Concepts, Architecture & Design, Prentice Hall, Upper Saddle River, N.J., 1998.
- [RFC98a] RFC 2225, "Classical IP and ARP over ATM," <http://info.internet.isi.edu/in-notes/rfc/files/rfc2225.txt>, April 1998.
- [RFC98b] RFC 2336, "Classical IP and ARP over ATM to NHRP Transition," <http://info.internet.isi.edu/in-notes/rfc/files/rfc2336.txt>, July 1998.
- [Sal98] D. Salomon, Data Compression: The Complete Reference, Springer, New York, 1998.
- [Sch95] J. S. Schlorff, "Traffic Shaping in ATM Networks with Application to Distributed Interactive Simulation Traffic," MS Thesis, University of Maryland at College Park, 1995.
- [SiJ95] K. Siu and R. Jain, "A Brief Overview of ATM: Protocol Layers, LAN Emulation, and Traffic Management," *Computer Communications Review*, vol. 25, no. 2, Apr. 1995, pp. 6-20.
- [SOW98] Statement of Work for the Air Force Institute of Technology to Support the Air Force Battlelabs' Virtual Battlelab Environment, Draft, 1998.

-
- [Sta96] L. Staalhagen, "A Comparison Between the OSI Reference Model and the B-ISDN Protocol Reference Model," *IEEE Network*, Jan/Feb 1996, pp. 24-33.
- [TrE95] H.L. Truong, W.W. Ellington Jr., J.Y. LeBoudec, A.X. Meier, J.W. Pace, "LAN Emulation on an ATM Network," *IEEE Communications Magazine*, May 1995, pp. 70-85.
- [Wu198a] C. Wu and J. Irwin, "Multimedia and Multimedia Communication: A Tutorial," *IEEE Transactions on Industrial Electronics*, vol. 45, no. 1, February 1998, pp. 4-14.
- [Wu198b] C. Wu and J. Irwin, Emerging Multimedia Computer Communication Technologies, Prentice Hall, Upper Saddle River, NJ, 1998.
- [Xu97] J. Xu, "IP over ATM: Classical IP, NHRP, LANE, MPOA, PAR and I-PNNI," http://www.cis.ohio-state.edu/~jain/cis788-97/ip_over_atm/index.htm, 14 August 1997.
- [YeY94] J. Yei and T. Yang, "A State-Space Reduction Method for Computing the Cell Loss Probability in ATM Networks," *Proc. IEEE ICC'94*, New Orleans, USA, May 1994.
- [ZhA94] Z. Zhang and A. Acampora, "Equivalent bandwidth for heterogeneous sources in ATM networks," *Proc IEEE ICC*, 1994, pp. 1025-1031.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, 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 March 1999		3. REPORT TYPE AND DATES COVERED Master's Thesis
4. TITLE AND SUBTITLE ANALYSIS OF USAF BATTLELABS' COMMUNICATION INFRASTRUCTURE OVER VARIOUS PROTOCOLS			5. FUNDING NUMBERS	
6. AUTHOR(S) Rocky A. Favorito, Captain, USAF David M. McIllece, First Lieutenant, USAF				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology 2950 P Street WPAFB, OH 45433-7765			8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GCS/ENG/99M-04	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ USAF/XORB Major Ross Walton (703) 696-0697 1911 N. Fort Myers Dr., Ste 1200 Rosslyn, VA 22209			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES Major Richard A. Raines (937) 257-3636 x4715 rraines@afit.af.mil				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) As a result of the doctrines exposed in the Air Force's "Global Engagement: A Vision for the 21st Century Air Force," the concept of a "Battlelab" was devised to support and develop the key warfighting capabilities needed in the next century. Although geographically separated, these labs must have the capability to exchange information with each other in a "Virtual Battlelab Environment" (VBE). Although many types of data will be exchanged, only variable-bit-rate (VBR) video and distributed interactive simulation (DIS) traffic are modeled in the VBE. The research described in this thesis utilizes a systems engineering approach to investigate the performance of wide-area networking technologies and recommend a connectivity solution. Three networking protocols were considered: Native ATM, IP over ATM, and LAN Emulation. In addition, three analyses were performed: performance-weighted, evenly-weighted, and cost-weighted. Native ATM was the recommended solution for the performance-weighted analysis. IP over ATM was recommended for the cost-weighted and evenly-weighted analyses, although by only slight margins over LANE. Finally, the merit of a systems engineering approach to assist in tradeoff and decision making analyses was also demonstrated.				
14. SUBJECT TERMS Battlelabs, Asynchronous Transfer Mode, Native Mode, Internet Protocol, IP over ATM, LAN Emulation, LANE, Distributed Interactive Simulation, Variable Bit Rate Video, Systems Engineering, Virtual Private Networking, Wide Area Networks, WAN Protocols			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	