


NAVAL POSTGRADUATE SCHOOL MONTEREY, CALIFORNIA



THESIS

**ANALYSIS OF THE NAVAL POSTGRADUATE SCHOOL
COMPUTER NETWORK ARCHITECTURE**

by

Paul Eric Wiedenhoeft

September 1994

Principal Advisor:
Associate Advisor:

Arthur L. Schoenstadt
James C. Emery

Approved for public release; distribution is unlimited.

DTIC QUANTITY INDICATED 6

19950117 028

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 1994	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE: ANALYSIS OF THE NAVAL POSTGRADUATE SCHOOL COMPUTER NETWORK ARCHITECTURE (U)			5. FUNDING NUMBERS	
6. AUTHOR(S) Wiedenhoef, Paul Eric				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE *A	
13. ABSTRACT (maximum 200 words) The computer network on the Naval Postgraduate School campus has become an integral part of the operations of the Naval Postgraduate School organization. An analysis of the network architecture will help formulate strategic plans that will support the network and the Naval Postgraduate School to the end of the century. This study describes the Naval Postgraduate School computer network architecture, driving forces, limitations, and possible measures of network benefits. It considers network alternatives and reasonable transition strategies. The study offers recommendations for improvements to the existing network architecture.				
14. SUBJECT TERMS Computer Network Architecture; Cabling; Local Area Network; Network Topology;			15. NUMBER OF PAGES 181	
			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	

NSN 7540-01-280-5500 Standard Form 298 (Rev. 2-89)

Prescribed by ANSI Std. Z39-18

Approved for public release; distribution is unlimited.

ANALYSIS OF THE NAVAL POSTGRADUATE SCHOOL COMPUTER NETWORK
ARCHITECTURE

Paul Eric Wiedenhoef
Lieutenant Commander, United States Coast Guard
B.S., United States Coast Guard Academy, 1983

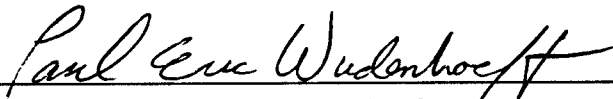
Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

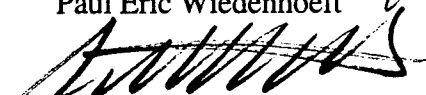
from the

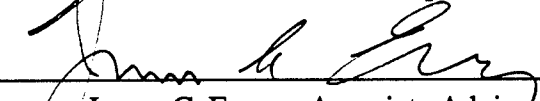
NAVAL POSTGRADUATE SCHOOL
September 1994

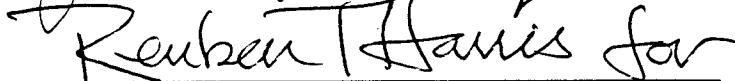
Author:


Paul Eric Wiedenhoef

Approved by:


Arthur L. Schoenstadt, Principal Advisor


James C. Emery, Associate Advisor


David R. Whipple, Chairman
Department of Systems Management

ABSTRACT

The computer network on the Naval Postgraduate School campus has become an integral part of the operations of the Naval Postgraduate School organization. An analysis of the network architecture will help formulate strategic plans that will support the network and the Naval Postgraduate School to the end of the century.

This study describes the Naval Postgraduate School computer network architecture, driving forces, limitations, and possible measures of network benefits. It considers network alternatives and reasonable transition strategies. This study offers recommendations for improvements to the existing network configuration.

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

ACKNOWLEDGMENTS

The author acknowledges the support, encouragement, and prayers of his wife, Kathleen, and his children, Rebekah and Timothy. The author also thanks Raul Romo and Terry Gentry for patient explanations and applications of technical and practical aspects of computer local area networks.

TABLE OF CONTENTS

I. INTRODUCTION	1
A. OBJECTIVES	2
B. RESEARCH QUESTIONS	2
C. METHODOLOGY	3
1. Questionnaire and Interviews	3
2. Literature Review	4
D. CHAPTER SUMMARY	4
II. BACKGROUND	5
A. DEFINITIONS	5
1. Local Area Network	5
2. Internetwork	5
3. Subnetwork	5
4. Segment	5
5. Network Architecture	6
6. Network Infrastructure	6
7. Transmission Medium	6
8. Node	6
9. Distribution Device	6
10. Port	7
11. Channel Capacity	7
12. Data Rate	7
13. Bandwidth	7
14. Throughput	8
B. TRANSMISSION MEDIA	8
1. Twisted Pair	8
a. Unshielded Twisted Pair	9
b. Shielded Twisted Pair	9
c. Twisted Pair Connectors	10
2. Coaxial Cable	10
a. Thick Coaxial Cable	11
b. Thin Coaxial Cable	11
c. Coaxial Cable Connectors	11

3. Optical Fiber	12
a. Optical Fiber Connectors	13
4. Wireless	13
C. DISTRIBUTION DEVICES	13
1. Bridge	13
2. Concentrator	14
3. Gateway	15
4. Hub	15
5. Network Interface Card	15
6. Repeater	16
7. Router	16
8. Switch	17
9. Transceiver	18
D. NETWORK TOPOLOGIES	18
1. Bus	18
2. Ring	18
3. Star	19
4. Mesh	19
E. LOCAL AREA NETWORK STANDARDS	20
1. OSI Reference Model	20
2. IEEE 802 Reference Model	21
3. Medium Access Control Protocols	22
a. IEEE 802.3	23
b. Ethernet vs. IEEE 802.3	25
c. IEEE 802.5	25
d. IEEE 802.12	26
4. Fiber Distributed Data Interface	27
5. Asynchronous Transfer Mode	27
F. CHAPTER SUMMARY	28
III. COMPUTER NETWORK ARCHITECTURE AT NPS	29
A. BRIEF HISTORY	29
B. BACKBONE AND SUBNETWORKS	30
1. Lower Network Layers	31
2. Higher Network Layers	32

3. Devices on the Network	32
4. Data Distribution	33
5. Survivability	33
C. MANAGEMENT AND SUPPORT	34
D. USERS AND APPLICATIONS	37
E. CHAPTER SUMMARY	38
IV. DRIVING FORCES OF NETWORK ARCHITECTURE	39
A. ORGANIZATIONAL ISSUES	39
1. Organizational Vision and Potential	39
a. Vision and Goals	39
b. Potential	41
2. Inertia, Management, and Funding	42
3. Human Resources	43
B. NETWORK PLANNING ISSUES	44
1. Scalability	44
2. Sustainability	45
3. Survivability	45
4. Security	46
C. APPLICATIONS AND HIGHER NETWORK LAYER ISSUES	46
1. Applications as Bandwidth	46
2. Data Compression	51
3. Distributed Systems	52
4. Network Management Applications	52
D. TECHNOLOGY ISSUES	53
1. Practical Limitations and System Bottlenecks	53
a. Channel Capacity	53
b. System Bottlenecks	57
2. Network Technologies	60
E. CHAPTER SUMMARY	62
V. MEASURES OF NETWORK BENEFITS	65
A. EFFICIENCY MEASURES	66
1. Cost per Port on a Distribution Device	66
2. Cost per Network Adapter	66
3. Cost per Node on the Network	67

4. Cost per Megabit of Performance	67
5. Bandwidth per Port	67
6. Packet Throughput	68
7. Network Simulation	69
8. Capacity Assessment	70
9. Network Monitoring	70
B. EFFECTIVENESS MEASURES	71
1. Business-Value Approach	71
2. Quality of Service	73
C. NPS MEASURES	75
VI. REASONABLE ARCHITECTURE ALTERNATIVES AND TRENDS	79
A. ARCHITECTURE ALTERNATIVES	79
1. Cabling	80
a. Cabling Between Buildings and Between Wiring Closets	80
b. Cabling within Wiring Closet	82
c. Cabling from Wiring Closet to Network Nodes	83
2. Network Connectivity	87
a. Network Interface Card	87
b. Repeater	88
c. Concentrator Hub	88
d. Switching Hub	89
3. Internetwork Connectivity	89
a. Shared Backbone	90
b. Collapsed Backbone	91
c. Backbone between Routers	92
B. ARCHITECTURE TRENDS	93
1. Meeting Immediate Needs	93
2. Scalability and Integration	93
3. Backbones	94
4. Virtual LANs	95
C. NPS APPLICABILITY	96
VII. REASONABLE TRANSITION STRATEGIES	99
A. STRATEGIES	100
1. Higher-Bandwidth Switch-Based Architecture	100

2. Evolution	102
3. Parallel Backbones	109
B. EXAMPLES	111
1. Fairfield University	111
2. Johns Hopkins School of Medicine/University of Maryland Medical Center	114
3. Concurrent Technologies Corporation	117
VIII. CONCLUSIONS AND RECOMMENDATIONS	121
A. CONCLUSIONS	121
B. RECOMMENDATIONS	122
1. General	122
2. Network Architecture Specifics	124
a. Subnetwork Cabling	125
b. Distribution and End Devices	125
3. Further Research	126
a. Measures of Benefits	126
b. Organizational Issues	126
c. Network Management Applications	127
C. SUMMARY	127
APPENDIX A. NETWORK ARCHITECTURE QUESTIONNAIRE	129
APPENDIX B. SUBNETWORK SUMMARY	133
APPENDIX C. NETWORK DIAGRAMS	141
APPENDIX D. CAPACITY ASSESSMENT	145
LIST OF REFERENCES	147
BIBLIOGRAPHY	153
INITIAL DISTRIBUTION LIST	163

LIST OF TABLES

1. Unshielded twisted pair EIA/TIA categories.	9
2. Examples of implementations by network layer.	21
3. IEEE 802.3 variants.	24
4. Driving Factors for Higher-Bandwidth LANs.	47
5. Data rate requirements for a number of applications typical of collaborative computing.	50
6. I/O bus throughput.	59
7. LAN upgrade technology tradeoffs, April 1994.	73
8. Parameters for user-oriented measures of quality.	74

LIST OF FIGURES

1. Representative network topologies.	19
2. OSI reference model and IEEE 802 model.	22
3. NPS computer network management organizational structure.	35
4. Different application bandwidth requirements.	48
5. Data rate vs. distance for various transmission media.	54
6. Potential system bottlenecks from end-to-end.	58
7. Cost/Performance Trends of Network Topologies.	61
8. Maturity of Higher-Performance Network Technologies.	62
9. Typical NPS subnetwork architecture.	79
10. Three main categories of cabling in a LAN.	80
11. Shared backbone configuration.	90
12. Collapsed backbone.	91
13. Backbone between routers.	92
14. Higher-bandwidth switch architecture.	101
15. An evolutionary transition to ATM.	106
16. Parallel backbones strategy.	110
17. JHU/UMMC late 1993 configuration.	115
18. JHU/UMMC network configuration (projected for late 1994).	116
19. CTC parallel backbones network.	119

I. INTRODUCTION

The Naval Postgraduate School (NPS) computer network is a strategic asset to the NPS organization. Although it is necessary for the day-to-day conduct of business throughout the campus, current organizational strategic plans do not address it. An analysis of the network architecture will provide insight into the strategic importance of the network and provide input into strategic planning for the network. An analysis of the computer network architecture provides an understanding of the network's role in supporting its users and the organization. An analysis also provides insight into how the network might better support its users and the organization in the future.

A wide variety of customers share the Naval Postgraduate School computer network. An overall network strategy is necessary to meet the needs of these customers.

[Katzan] proposes this strategy must include specification of three main items:

- *Current Position* (Where are we?)
- *Goals* (Where are we going?)
- *Direction* (How do we get there?)

Current position is a determination of installed equipment, network topologies, organization, and use of the network. *Goals* are the future position of the network, dependent on technology and the goals of the organization. *Direction* is the plan to move from the current position to the future position. As the futures of the organization and technology are uncertain, the network strategy must include numerous reasonable alternatives, yet still address the unexpected.

There is no current documented strategic plan for the NPS computer network.

There is little documentation of the present architecture. There are no current documented goals. There are no documented plans how to reach any goals. In the second half of 1994, efforts began to address some of these shortcomings as they relate to computing in general on the NPS campus.

A. OBJECTIVES

This study shall assist network administrators in their continued development of strategic plans that will provide a computer network architecture capable of supporting NPS to the turn of the century and beyond. This study will identify specifics regarding current position of the NPS computer network, consider reasonable goals of the network, and study transition strategies appropriate for obtaining the network goals. The study will offer recommendations for continued strategic improvement to the existing network configuration.

B. RESEARCH QUESTIONS

To obtain the objectives of this thesis and gain insight into requirements for plans and goals of the network, managers and those with influence over the network's role must ask relevant questions. [Whittman] offers questions to ask regarding network architecture in an organization:

- What is the state of the existing network? What do measurements and network data indicate? Is there a map of the network? Is change needed?
- What is important to the organization? (Reliability, bandwidth, and flexibility are possible considerations.) Can the network architecture match the organizational culture? Is it possible to stay a step off the leading edge of network technology?
- What changes in the organization are likely to affect the load on the network in the next few years? Will the number of nodes on the network increase substantially? What about down-sizing? What new applications will be introduced?
- Is there a plan to evaluate the results of implementing a new technology?
- Will a given network architecture make the network management's life easier?

To obtain the objectives of this study, the list of questions simplifies to the following:

- What is the current architecture of the NPS computer network?
- What are the driving forces affecting the NPS compute local area network?
- What are reasonable alternatives to the current NPS computer network architecture?
- What are reasonable transition strategies for the NPS computer network?

C. METHODOLOGY

1. Questionnaire and Interviews

A questionnaire gathered information on the current network architecture. The questionnaire used for this study identified specifics about the physical layout of the network architecture and provided insight into the management and the users of the network. Interviews of network administrators, users, consultants, and vendors provided

additional insight not captured in the questionnaire answered questions raised by responses to the questionnaire. Practical experience with cabling, network management, and hands-on troubleshooting supplemented both the questionnaire and the interviews.

2. Literature Review

Computer network technology is advancing at a startling pace. Current trade magazines and professional journals provided insight into four major areas of computer network management: limiting factors and driving forces behind advances in network architecture, alternatives for the NPS network architecture, ways in which a network architecture can be considered beneficial to its users and the organization, migration strategies and lessons learned from transition experiences of managers of local area networks in other organizations.

D. CHAPTER SUMMARY

This thesis provides information, alternatives and recommendations to network managers useful for strategic network planning. The following chapter reviews terms and concepts that lay the foundation for subsequent chapters. The subsequent chapters address research questions as they apply generically to computer networks and specifically to the NPS computer network.

II. BACKGROUND

This chapter reviews terms and concepts necessary for understanding discussions in subsequent chapters.

A. DEFINITIONS

1. Local Area Network

A local area network (LAN) is a system of computing resources interconnected via common transmission media, data distribution devices, and network interfaces in order to share information. The Institute of Electrical and Electronics Engineers (IEEE) distinguishes LANs from other computer networks in that they are optimized for a moderate-sized geographic area such as a single building or a campus. A LAN is generally owned, used, and operated by a single organization.

2. Internetwork

An internetwork is a group of LANs interconnected by a data distribution scheme such as a backbone or a router. This study refers to internetworks as "networks."

3. Subnetwork

A subnetwork is a portion of a network that by itself meets criteria to be considered a LAN and typically uses only one network topology.

4. Segment

A segment is a group of network resources on a network or subnetwork sharing a single segment of cabling.

5. Network Architecture

Network architecture is the planned structure of a network and a description of data formats and procedures used for communication on a network. Architecture implies orderly arrangement, deliberate design, and organization.

6. Network Infrastructure

Network infrastructure is the structure of the network regardless of planning. It is the combination of cabling, distribution devices, and network topologies that provide services to the network resources and users.

7. Transmission Medium

The path, or channel by which data is distributed between stations on a network. Examples include guided media such as copper-based coaxial cable and twisted pair, light-based optical fiber, and unguided media such as microwave or infrared channels.

8. Node

A node is the network access point on a transmission medium for a computing device that originates and/or is the end recipient of data across a computer network.

9. Distribution Device

Any device in a network that distributes data through a transmission medium to another station. Examples include bridges, concentrators, gateways, hubs, repeaters, routers, and switches.

10. Port

A port is an input/output connection on a network distribution device through which data is received from and/or distributed across a transmission medium.

11. Channel Capacity

Channel capacity is the maximum rate at which data can be transmitted over a given path, or channel under given conditions, typically expressed in bits per second (bps) or megabits per second (Mbps).

12. Data Rate

Data rate is the rate, in bits per second (bps) or megabits per second (Mbps) at which data is communicated. Data rate is a function of signal transmission rate (typically binary for computer LANs) and bandwidth.

13. Bandwidth

The difference between the highest and lowest frequencies of the transmitted signal as restricted by the transmission medium and the transmitter, expressed in Hertz (Hz) or MegaHertz (MHz). Bandwidth is directly proportional to channel capacity. Because of this relationship between bandwidth and channel capacity, bandwidth is often used as a measure of the channel capacity. For instance, all else being equal, a doubling of bandwidth corresponds to a doubling of channel capacity.

14. Throughput

Throughput is a rate at which useful data is communicated between end devices on a network. It is a function of data rate and encoding scheme. Throughput is a measure of the useful data across the network after removing network transport data from the encoding scheme.

B. TRANSMISSION MEDIA

The choice of transmission medium for network connectivity is the foundation on which an entire network is built. Different transmission media are appropriate in different situations. Network managers must weigh cost and performance characteristics of the various choices in order to meet the demands of network management, user needs, and fiscal constraints. In this section, types of networks refer to the transmission media that support them. Subsequent sections discuss the types of networks.

1. Twisted Pair

Twisted copper pair cabling is common in network environments. This is largely due to early local area network concerns of reducing costs by taking advantage of previously installed (and unused) voice-grade telephone twisted pairs. The use of lower grade twisted pair is no longer adequate as higher data rates and electro-magnetic concerns require higher quality transmission media and connections. The outer sheath around the cable is typically plenum grade, used among distribution devices and between distribution devices and wall-mounted jacks, or PVC grade, used between wall-mounted jacks and network interface cards.

a. Unshielded Twisted Pair

Unshielded Twisted Pair (UTP) is widely accepted for network data transmissions because it is cost-effective and easy to install. Ever-increasing demands placed on UTP networks prompted an industry standards organization, Electronics Industry Association/Telecommunications Industry Association (EIA/TIA), to develop specifications for system performance. Table 1 summarizes the EIA/TIA standards. Currently, the highest performance specification is the Category 5 EIA/TIA-568 standard.

Category	Maximum data rate per pair (Mbps)	Attenuation (dB per thousand ft)	Cost-- PVC grade (cents / ft)	Cost -- Plenum grade (cents / ft)
1	(not specified)	(not specified)	5 - 15	40
2	4	8 @ 1 MHz	8 - 20	40
3	10	30 @ 10MHz 40 @ 16MHz	12-25	30 - 45
4	16	22 @ 10MHz 31 @ 20MHz	20 - 45	45 - 85
5	100	32 @ 10MHz 67 @ 100MHz	25 - 45	45 - 60

Table 1. Unshielded twisted pair EIA/TIA categories. After [Marks].

b. Shielded Twisted Pair

Shielded Twisted Pair (STP) has inherent quality and high system performance. There are primarily two types: 100 Ohm and 150 Ohm. IBM introduced 150 Ohm STP in 1984 and has since improved the capabilities of the cable so that it complies with the EIA/TSB-53 proposed standard of up to 300 MHz for data and 600 MHz for video signals. When installed with the proper connectors, the cable meets the

requirements of the FDDI standard for 100 Mbps data transmission at 100 meters. [One Network Place]

c. Twisted Pair Connectors

Registered Jack (RJ) terminations and punch-down blocks, common in the telecommunications field, are appropriate connectors for twisted pair cabling. The most common jacks used for data-grade twisted pair cabling include RJ-11 (2 pair) and RJ-45 (4 pair). Punch-down blocks are practical for terminating and interconnecting a high density of twisted pairs such as in a wiring closet.

2. Coaxial Cable

Coaxial cable (coax) is a two-conductor, metallic electrical cable used for radio frequency (RF) and digital data communications transmission. The cable is constructed with a single solid or a stranded center conductor that is surrounded by a dielectric layer, an insulating material of constant thickness and high resistance. The second conductor is a layer of aluminum foil, metallic braid or a combination of the two encompassing the dielectric and acting both as a shield against interference (to or from the center conductor) and as the return ground for the cable. Finally an overall insulating layer forms the outer jacket of the cable. Coaxial cable is generally superior in high-frequency applications such as networking. However for shorter distances (up to 100 meters), UTP or STP cable is generally just as reliable when using differential modulation techniques (such as with 10Base-T) [Medici]. The most common types of coaxial cable are those 50 Ohm cables

used in IEEE 802.3 CSMA/CD networks. Other coaxial cables (e.g., 75 Ohm RG-62) are also used in LAN environments.

a. *Thick Coaxial Cable*

Thick coaxial cable is typically used in 10Base5 network environments. It is 50 Ohm cable designed for use as backbone cabling. This cable was specified for the original Xerox "Ethernet" CSMA/CD network, so it is often referred to as Ethernet cable. This leads to confusion because several transmission media can now be used. This coaxial cable is described as "thick ethernet" or simply "thicknet" in this study. Thicknet has an outside diameter of 0.375-0.405 inches.

b. *Thin Coaxial Cable*

Thin coaxial cable is 50 Ohm cable that is often used for horizontal cable runs in Thin Ethernet (10Base2) environments. RG-58 and its variants are examples of 50 Ohm coaxial cable with outside diameters of 0.165-0.195 inches. This coaxial cable is described as "thinnet" in this study.

c. *Coaxial Cable Connectors*

Coaxial cables mentioned above are typically terminated with loads that match the impedance of the cable. Segments of cable connect to each other and to terminal equipment using standard connectors that properly align the respective center conductors. Bayonet Neill-Concelman (BNC) connectors are standard for thinnet and allow quick connection and disconnection of segments. A special "T" configuration of a series of cable segments with BNC connectors allows a transceiver to "tap" into the cable

and not disrupt signals on the shared medium. For thicknet cabling, a medium attachment unit (MAU), also known as a transceiver, taps directly into the core of one continuous cable segment. If a transceiver is not directly on the device it is attaching to the network, a segment of attachment unit interface (AUI) cable spans the distance between the transceiver and the device.

3. Optical Fiber

Optical fiber outperforms both twisted pair and coaxial cable in LAN environments. Optical fiber transmission uses a different part of the frequency spectrum than twisted pair and coaxial cable transmissions, and is therefore immune to electromagnetic interference (EMI) and radio frequency interference (RFI). The lowest grade of optical fiber signaling, light emitting diodes over multi-mode optical fiber, has data capacity of roughly five times that of twisted pair and can travel over a distance ten times greater. This makes it useful for longer distances transmissions and high capacity channels such as network backbones. Use of optical fiber is limited because of its costs, especially the cost of connectors. The cost per unit length of multi-mode optical fiber is two to four times greater than UTP. Optical fiber connector costs are roughly ten times more than for UTP connectors. ([Black Box] lists terminated 62.5-micron core plenum fiber optic cable at \$118.40 plus \$1.47 per foot. The same catalog lists terminated Category 5 plenum UTP at \$11.00 plus \$0.42 per foot.)

a. Optical Fiber Connectors

Optical fiber connectors must be fused to the end of the cable. Improper fusion caused by misalignment, incorrect temperature, or incorrect fusing time severely degrades the available channel capacity caused by increased transmission losses.

4. Wireless

Wireless, or unguided, transmission systems are appropriate for areas where use of other transmission media is cumbersome, such as open spaces in libraries, older buildings with solid walls, floors, and ceilings with no place for cable runs, and from building to building when cabling is technically or economically infeasible. Wireless transmissions have channel capacities comparable to lower-capacity guided media transmissions (i.e., less than 20 Mbps). General types of wireless transmission media are infrared, microwave, and spread spectrum.

C. DISTRIBUTION DEVICES

This section describes the capabilities of various distribution devices. The name of a device, as presented by vendor, does not necessarily reflect the capabilities of the device. Managers should look at the functionality of a device rather than its name to determine its capabilities. Subsequent sections on network topologies and network standards refer to the definitions in this section.

1. Bridge

[Newton] defines a bridge as a distribution device that connects LANs using similar or dissimilar media and signaling systems such as ethernet and token ring. When

used for connecting LANs, a bridge connects LANs at the IEEE 802 medium access control (MAC) sub-layer of the data link layer. Bridges forward packets destined for another LAN. Bridges are normally either source routing bridges or transparent bridges. A third bridging method, Source Routing Transparent (SRT), enables a bridge to act as both a transparent and a source routing bridge.

2. Concentrator

[Newton] states, "It makes the network connections." Some wiring concentrators are dumb, making only physical connections between network segments. Others are intelligent, making networking decisions, and providing network diagnostics. A wiring concentrator can have bridges and routers that divide the network into segments. It can have the hardware necessary to change from one transmission medium to another (e.g., twisted pair to optical fiber). It can contain the hardware to change from one network type to another--for example, from ethernet to token ring.

Some LANs use concentrators, or access units, that allow network devices to be interconnected through a central point in a star wiring topology. Attaching devices through a central concentrator typically simplifies the maintenance of a LAN.

"Concentrator" is a very generic term for a distribution device. Bridges, gateways, hubs, multiplexors, routers, and switches all might be concentrators in appropriate configurations.

3. Gateway

[Freedman] defines a gateway as a computer that performs protocol conversion between different types of networks or applications. For example, a gateway can connect a microcomputer LAN to a mainframe network. An electronic mail, or messaging, gateway converts messages between two different messaging protocols.

4. Hub

[Newton] defines a hub as the point on a network where numerous circuits are connected. A hub is also called a switching node, especially in star-topology LANs. Hub hardware can be either passive or active. "Passive hubs" add nothing to the data being transmitted. "Active hubs" regenerate signals and may monitor traffic for network management. "Intelligent hubs" are computers that provide network management and may also include bridging, routing, and gateway capabilities.

Wiring hubs are useful for their centralized management capabilities and for their ability to isolate nodes from disruption. Hubs are becoming so comprehensive that some even offer optional expansion boards that include a file server and network operating system. A hub's star topology improves troubleshooting over bus topology, in which all nodes are connected to a common cable. Active, intelligent hubs incorporate functions found in bridges, routers, and switches.

5. Network Interface Card

A network interface card (NIC), or network adapter, provides an interface from a device at a node to the network. It works at the two lowest layers of the OSI reference

model with the network software and computer operating system to transmit and receive messages on the network. The NIC takes streams of 1s and 0s from the network and formats them into frames. The frames are then passed to higher level protocols for additional processing. NICs are most commonly identified by several characteristics:

- Medium access control topology, such as token ring, ethernet, FDDI
- Bus interface, such as ISA, EISA, MCA, NuBus, PCI.
- Data path width, such as 8, 16, 32, or 64 bits
- Physical media connection, such as coaxial cable through AUI or BNC, UTP through RJ-45, or optical fiber through ST connectors.

6. Repeater

A repeater is a distribution device used to receive a digital signal, recover the pattern of 1's and 0's and retransmit the new signal. A repeater overcomes the attenuation losses of a digital signal over distance on a transmission medium. It can also be used to "segment" a network. If medium access control (MAC) layer or physical layer problems occur on a particular segment, the repeater may isolate that segment from the rest of the network.

7. Router

A router is a computer system that routes messages from one LAN to another. It is used to interconnect similar and dissimilar networks and can select the most expedient route based on traffic load, line speeds, costs, and network failures. Routers maintain address tables for all nodes in the network and work at the network layer of the OSI

reference model. Distributing at the network layer takes more time than IEEE 802 MAC layer devices such as bridges.

Routers break apart the LAN into smaller LANs for improved security, troubleshooting, and performance. For example, an internetwork protocol (IP) router can divide a network into subnetworks so that only traffic destined for particular IP addresses can pass between segments. Routers with high-speed (gigabit) buses may serve as a "collapsed" internetwork backbone, connecting all networks in the enterprise.

8. Switch

A switch is a mechanical, electrical or electronic device that opens or closes circuits, completes or breaks an electrical path, or selects paths or circuits. In LANs, a switch divides a large network into smaller segments by filtering unnecessary traffic from individual segments. Most distribution devices incorporate some type of switching in their logic.

Two types of switches are "cut-through" and "store-and-forward." Cut-through switches read only the a portion of the header of a packet before switching the incoming signal to another segment. Store-and-forward switches read the entire incoming signal before determining where to route it. Cut-through switches are faster than store and forward switches; while store-and-forward switches are better for filtering "bad" packets or unwanted traffic from a segment.

9. Transceiver

A transceiver connects a device on a node to a network. It mediates transmission and receipt of data by a node on the network. A transceiver may be built directly into a NIC, such as in 10Base2 networks or attached to a NIC by an attachment unit interface (AUI) cable such as in 10Base5 networks.

D. NETWORK TOPOLOGIES

Topologies appropriate for LANs include bus, ring, star, and mesh, as shown in Figure 1. These topologies indicate either the physical layout of the cabling that connects the network devices or the logical or electrical connections among network devices.

1. Bus

Bus topology implies a serial connection of network stations to a shared linear medium with two defined ends. On a logical bus, any signal sent from one station is received by all other stations on the network. Tree topology is a generalization of bus topology in that the serial connections to the bus may be buses themselves.

2. Ring

Ring topology is a concentric grouping of network stations on a continuous shared medium. On a logical ring, any signal sent from one station must pass through the next designated station on the ring before continuing around the ring and back to originating station.

3. Star

Star topology implies a central distribution device surrounded by the nodes of the network arranged in a star-like manner. A logical star network separates each segment from all other segments on the network.

4. Mesh

Mesh topology implies that each node on a network is directly connected to more than one other node on the network. It allows alternate path routing of signals.




Figure 1. Representative network topologies

E. LOCAL AREA NETWORK STANDARDS

This section discusses standards for some common local area networks alternatives used or considered at the Naval Postgraduate School. Various communications and computing committees, such as American National Standards Institute (ANSI), Institute of Electrical and Electronics Engineers (IEEE), and International Standards Organization (ISO), adopt standards for local area networks. Some network alternatives are not yet standards because many standards issues are not resolved. These non-standard alternatives deserve mention because they promise significant improvement in performance over some of the adopted standards.

1. OSI Reference Model

International Standards Organization (ISO) developed the open systems interconnection (OSI) model as a reference for computer communications architecture and as a framework for computer communications protocol standards. The model consists of seven layers: physical, data link, network, transport, session, presentation, and application. Table 2 shows examples of implementations at the various network layers.

This study focuses on the first two layers. The first, physical layer, is concerned with transmission of a bit stream over the transmission medium. It deals with the mechanical, electrical, functional, and procedural characteristics of access to the transmission medium. The second, data link layer, is concerned with the reliable transfer of information across the transmission medium. It deals with blocks of data (frames) and the necessary synchronization, flow control, and error control.

Network layer(s)	Sub-layer	Examples
7. Application	application software, support software, operating systems	word processors, graphics, POSIX, ACMS, Oracle tools, VAX, DOS, Unix, NDIS drivers, ASCII, PostScript, SMTP
6. Presentation	data organization; storage, transmission and data format	HPFS, Macintosh HFS, byte stream, AFP, NFS, SMB, NCP
5. Session and 4. Transport		TCP, SPX, UDP, ASP/ATP, MS NetBEUI, TP4(OSI)
3. Network		IP, IPX, DDP, IS-IS (OSI)
2. Data Link and 1. Physical		IEEE 802.3, IEEE 802.5, Localtalk, FDDI

Table 2. Examples of implementations by network layer. After [Cini].

2. IEEE 802 Reference Model

The IEEE committee 802 developed a set of standards, based on the OSI model, that focuses on the lowest communications layers as they specifically apply to LANs. These first two OSI layers are divided into three IEEE layers, as shown in Figure 2. The highest of these three IEEE layers is the logical link control (LLC) described in the IEEE 802.2 standard. It is responsible for addressing and data link control, and is independent of the topology, transmission medium, and medium access control technique. The lowest two layers are the physical layer the medium access control (MAC) layer. The transmission medium and topology are interdependent with the MAC layer; therefore IEEE 802 developed a series of standards organized by the MAC algorithm used in each case. Higher layers in the two models are the same. [Stallings 1]

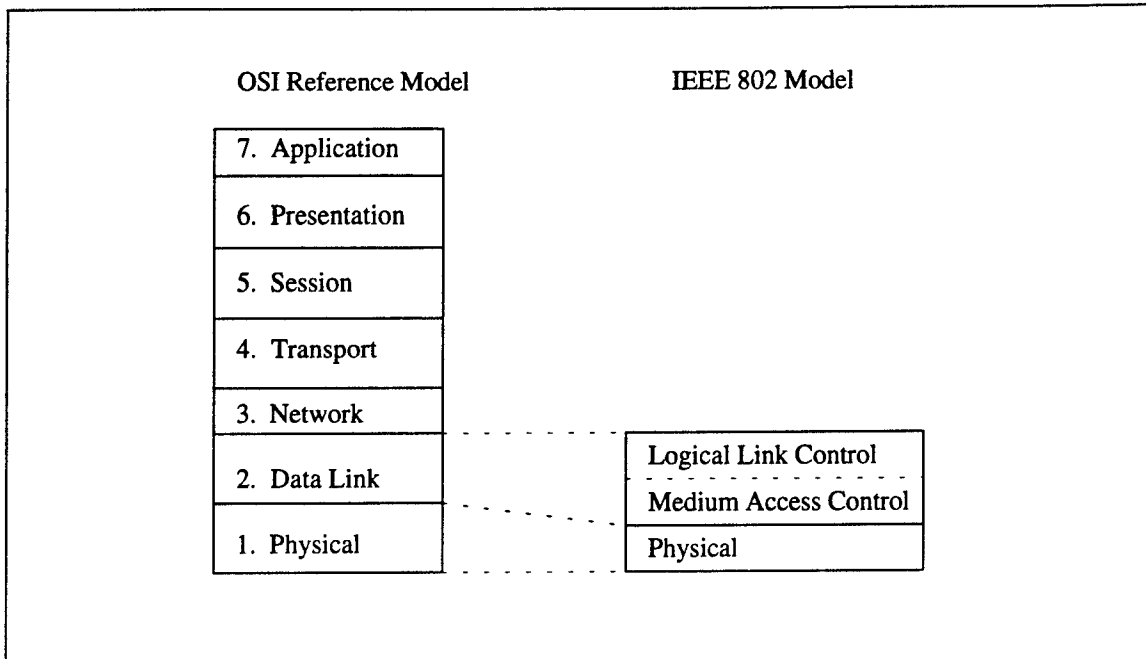


Figure 2. OSI reference model and IEEE 802 model. From [Stallings 2].

3. Medium Access Control Protocols

The medium access control (MAC) layer controls access by devices sharing the transmission medium. The specified MAC technique defines where the access is controlled and how a device gains access to the network.

Central access control provides tight management over channel capacity and simplifies the electronics at the nodes. Distributed access control avoids the single point of failure inevitable in central access control, but increases the complexity of each node.

Access to the network is either synchronous or asynchronous. Synchronous methods, which dedicate channel capacity among nodes, are typically not used in LANs because traffic to and from nodes is generally unpredictable [Stallings 1]. Asynchronous methods include *round robin*, *contention*, and *reservation*.

In *round robin* techniques, each station is given an opportunity, in turn, to transmit. A station may decline to transmit or transmit for a specified length of time before giving up its turn. This is an efficient technique when many stations have data to transmit over long periods of time. 100VG-AnyLAN is an example of a centralized round robin technique. Token ring is an example of distributed round robin.

In *contention* techniques, no access control is exercised over the transmission medium. This technique is efficient when network traffic is "bursty" with light to moderate volumes. All stations contend for time on the network, which must have distributed access control. Carrier sense multiple access with collision detection (CSMA/CD) is an example of a contention technique.

In *reservation* techniques, time on the transmission medium is divided into slots similar to synchronous methods except that stations reserve future slots based on the time needed or type of traffic. Time-division multiplexing (TDM) and Asynchronous Transfer Mode (ATM) are examples of reservation techniques.

a. IEEE 802.3

The IEEE 802.3 standard defines network medium access protocols using carrier sense multiple access with collision detection (CSMA/CD) used on logical bus topology LANs [Stallings 1]. In this technique, a station desiring to send a data frame over the network "listens" to the transmission medium and senses if any other traffic is present. If no traffic is present, the station sends its traffic and continues to listen. If a collision with another station's frame is detected during the transmission of the frame, the

station stops sending, transmits a "jam" signal to all stations, waits a period of time and tries to send again.

The transmission medium used, data rate, and signaling differentiate the options in the series of IEEE 802.3 MAC protocols. Each option has at least one industry nickname or brand name associated with it. Table 3 summarizes the IEEE 802.3 series.

The 100Base-T (proposed) standard is broken down into three classifications based on MAC framing and physical layer specifications for signaling and cabling [Roberts].

IEEE 802.3 designation	Nick name or Brand name	Data rate (Mbps)	Signaling	Transmission medium
10Base5	Thicknet or Ethernet	10	baseband	"Ethernet" 50 Ohm coax
10Base2	Thinnet or Cheapernet	10	baseband	RG-58 series 50 Ohm "thin" coax
10Base-T	10Base-T	10	baseband	Cat 3 UTP (2 pr.)
1Base5	StarLAN	1	baseband	UTP
10Broad36	Broadband ethernet	10	broadband	75 Ohm coax
10Base-F	ethernet over fiber	10	signal-encoded beam of light	multi-mode optical fiber (1 pr.)
100Base-T (proposed)	fast ethernet (fast ethernet alliance)	100	baseband	one of the 100Base-TX, 4T, FX cabling
100Base-FX (proposed)	100 Mbps ethernet over fiber	100	signal-encoded beam of light	multi-mode optical fiber (1 pr.)
100Base-TX (proposed)	Grand Junction proposal	100	baseband	Cat 5 UTP (2 pr.) Type 1 STP (2 pr.)
100Base-4T (proposed)	4T+	100	baseband	Cat 3, 4 or 5 UTP (4 pr.)

Table 3. IEEE 802.3 variants.

b. Ethernet vs. IEEE 802.3

Xerox/PARC Ethernet II and IEEE 802.3 CSMA/CD standards are generally interchangeable. The primary difference between the two is the use of two bytes in the ethernet packet frame. These two bytes are referred to as the "type" or "length" field. For Ethernet II, these two bytes represent the *type* of packet (for instance IP packets are coded 0x800). For IEEE 802.3, these two bytes indicate the *length* of the data field that immediately follows the length field. (The data field can be no less than 46 bytes and no greater than 1500 bytes.) [Medici], [Stallings 2]

This difference is so slight in definition and in application in LANs that I will use the term "ethernet" throughout this study to refer to either of these two CSMA/CD protocols.

c. IEEE 802.5

The IEEE 802.5 standard, called token ring, is adapted from a protocol developed by IBM. In this technique, a particular bit pattern, called a token, is passed from station to station around a logical ring topology. A station desiring to send traffic must wait until the token is available to send its data frame. When the token is available, the station grabs the token and sends its data frame. When the leading edge of the data frame completes its cycle around the ring and after the data frame is completely sent, the transmitting station purges the frame from the network and transmits a new token to the next station on the ring.

d. IEEE 802.12

The IEEE 802.12 proposed standard, called 100VG-AnyLAN, is being developed by Hewlett-Packard Company. 100VG-AnyLAN is an evolution of 10Base-T Ethernet and Token Ring topologies. It can deliver 100 Mbps to every node in a given network with no major software changes from installed 10Base-T Ethernet and Token Ring networks through a type of "polling" scheme in which the hub controls access to the network rather than using the node-based CSMA/CD scheme used by other ethernet standards. This deterministic "Demand Priority" arbitration system enables video, voice, and multi-media applications that require guaranteed bandwidth and predictable, low latencies.

100VG-AnyLAN requires no changes to application software on clients or servers, and is compatible with current network operating systems. 100VG-AnyLAN requires no changes to bridges and routers, supporting the ethernet or token ring framing and network management systems already in use. This topology gives network administrators a seamless way to boost the performance of 10Base-T by a factor of 10. 100VG-AnyLAN gives organizations a networking system that allows them to realize these higher data rates over existing UTP (the "VG" is short for "voice grade"), thus obviating one of the biggest costs of network upgrades, recabling. [Hewlett-Packard 2]

4. Fiber Distributed Data Interface

The ANSI Fiber Distributed Data Interface (FDDI) standard and its copper cable variants are based on the IEEE 802.5 token ring standard but modified for higher speeds. They are robust and reliable network protocols that provide 100 Mbps on a shared medium to a maximum of 1000 repeating nodes. Two counter-rotating rings allow self-healing after cable or equipment malfunctions. Dual homing (attachment of critical equipment via two independent connections to the network) is an option that increases survivability of the network and devices should a fault occur in a network interface or one of the two counter-rotating rings. FDDI concentrators can detect and disconnect faulty nodes. Twisted Pair-Physical Medium Dependent (TP-PMD) is based on an earlier proprietary specification called Copper Distributed Data Interchange (CDDI). Two STP variants of FDDI are Twisted Pair-FDDI (TP-FDDI) and Shielded Distributed Data Interface (SDDI). [Miller 2]

5. Asynchronous Transfer Mode

Asynchronous transfer mode (ATM) is a technique for high speed transfer of data based in high-speed switching and small fixed-length packets drawn from telephony standards. It promises high capacity on star or mesh logical topology networks. ATM has not been completely adopted as a standard by ANSI, IEEE, or ISO. An ATM forum is attempting to resolve such issues as transmission media, signaling, node limitations, routing, security, multi-casting, LAN emulation, and IP encapsulation [Strauss 1].

ATM establishes a virtual circuit or channels with a reserved capacity between the transmitting station and the receiving station. Using small fixed-size (53 bytes) packets called cells, the switches in the circuit can transfer the fixed size cells at a constant data rate in the hardware with minimal software intervention. This allows the switching to be extremely fast, especially when compared to other LAN technologies. In LAN environments, ATM is designed for data transfer rates between 45 Mbps and 2.4 Gbps and possibly higher [Feltman]. Because the data rate can be high and constant, this technique is appropriate for real-time multimedia transmissions (such as live video) and high aggregate bandwidth needs.

F. CHAPTER SUMMARY

This chapter reviewed terms and concepts used throughout the remainder of the analysis. The following chapter analyzes the current computer network architecture and network management at NPS using many of the terms and concepts found in this chapter.

III. COMPUTER NETWORK ARCHITECTURE AT NPS

A. BRIEF HISTORY

The Naval Postgraduate School (NPS) computer network evolved over more than a decade from small separate departmental local area networks (LANs). As these networks grew, departments perceived the benefit of interconnecting LANs to share information, use of electronic mail and other network services, and access to wide area networks such as DDN, NFSNET, DISNET, and BARRNET. Many of these services were already available through the mainframe computer and its distributed "dumb" terminals throughout the campus, but other factors drove the demand to provide these services on the LANs. The emergence of end-user computing, supercomputers, and powerful graphics and engineering workstations contributed towards the interconnection of departmental LANs. [Norman], [Leahy]

The informal goal of network efforts has been, to this point, connectivity, i.e., connecting to the network every desktop and computing resource on the campus that could benefit from being linked to other computing resources. To this end, network managers have installed quick, easy-to-install, and inexpensive cabling, network interfaces and connections to the campus backbone for every device from the lowest-end PC to the mainframe computer and the supercomputer.

According to the NPS Director of Academic and Administrative Computing, Code 51, the basic connectivity goal should be achieved during 1994. After this goal is met, informal goals of the network include providing enhanced services such as increased

bandwidth, improved network management, improved user services, and more and better applications for the users.

B. BACKBONE AND SUBNETWORKS

The computer network on the NPS campus is generally a backbone network connecting numerous subnetworks. The information in this section summarizes data collected during 1994 from a questionnaire (Appendix A), interviews with network managers and technicians, and efforts to map the campus network by tracing physical cabling. Appendix B provides a tabular summary of information collected about subnetworks from questionnaire responses, interviews, and router configuration data. Appendix C provides graphical views of representative subnetworks and cabling maps.

Much of the information collected in these appendices may already be inaccurate as the network continues to change. Subnetwork administrators are continually changing the configuration of their subnetworks to meet demands of the users. The configuration changes include adding cabling, distribution devices, and other network resources; re-routing existing cabling; and relocating network resources. Much of the reconfiguration work depends on availability of funding, supplies, technicians, and time.

The NPS computer network is a heterogeneous amalgamation of network technologies at all layers. The following paragraphs describe, in general, the physical status of the network.

1. Lower Network Layers

The present configuration of the NPS computer network architecture is primarily a collection of shared medium topologies with capacities of 10 Megabits per second (Mbps) or less. Coaxial cable bus ethernet topologies (10Base5 and 10Base2) dominate the subnetworks. Other subnetwork topologies are FDDI, token ring on shielded twisted pair copper wire (STP), Apple Computer's Localtalk, and ethernet over unshielded twisted pair copper wires (10Base-T). In addition to these topologies, others, such as Artisoft's LANtastic, run in some areas of the campus without direct connection to the network. Other LANs remain disconnected from the backbone for (primarily) security reasons.

The campus backbone is "collapsed"--that is, running on the internal backplane of interconnected Cisco AGS+ and CGS routers. A 10 Mbps ethernet backbone connects the routers for transferring information among the subnetworks and for accessing wide area network (WAN) connections.

Many of the subnetworks are backbones themselves. Standard coaxial ethernet cables ("thicknet") connect to the routers via transceivers and attachment unit interface (AUI) cables. The cables extend to the buildings and floors where most of the devices are attached to the subnetwork. Single-port and multi-port transceivers tap into thicknet and connect to distribution devices through AUI cables. Cabling between distribution devices and the nodes is AUI, thicknet, thinnet, or UTP, as appropriate, for the network adapters or network interface cards (NICs) at the nodes.

A few of the thicknet subnetworks connect to the routers via fiber-optic transceivers that interface between coaxial cable and multi-mode optical fiber segments.

2. Higher Network Layers

A variety of session layer and network layer protocols run on the NPS computer network. These protocols are listed below.

- Apple Computer, Inc.'s Appletalk Session Protocol/Appletalk Transaction Protocol (ASP/ATP) for Appletalk networks
- Banyan, Inc.'s Interprocess Communication Protocol/Vine's IP for Banyan Vines networks (ICP/VIP)
- Department of Defense's Transmission Control Protocol and Internet Protocol for DoD networks (TCP/IP)
- Digital Equipment Corp.'s Network Services Protocol (NSP) for DECnet networks
- Novell, Inc.'s Sequenced Packet Exchange protocol/Internet Packet Exchange protocol (SPX/IPX) for Netware networks (and based on XNS)
- Xerox Corporation's Xerox Network Services (XNS).

3. Devices on the Network

Devices on the network include an Amdahl mainframe computer, a Cray YMP supercomputer, high-end workstations such as Silicon Graphics, Incorporated (SGI), Hewlett-Packard (HP), and other RISC processor machines; low-end IBM PC or compatible microcomputers; and Apple Macintosh microcomputers. Other devices on the network include a variety of file servers, printers, terminal servers, and others.

4. Data Distribution

Data distribution devices include routers, ethernet multi-port transceivers on thicknet, fiber-optic transceivers, ethernet multi-port repeaters between thicknet and thinnet ethernet segments, 10Base-T repeating hubs, token ring multi-station access units (MAUs), FDDI hubs, and a variety of LAN-to-mainframe gateways.

Some research initiatives bypass the network as a means of data distribution. For example, some videoteleconferencing applications run on ISDN lines. Also, departments have modems that connect to off-campus organizations without direct DoD Internet access. The oceanography department's (COAC) lab is considering running a dedicated T1 (1.5 Mbps) to another site.

5. Survivability

Some precautions are in place or planned to allow the network to withstand network disasters caused by loss of electrical power. These include battery back-up and emergency generator for the computer center and uninterruptable power supplies on some servers and end nodes. The battery back-up system for the computer center has been out of operation since November 1993 because of wiring and switching problems. Some distribution devices store their configurations on file servers. When both the server and the distribution device "go down," the distribution device configuration cannot be restored until the server comes back up.

The cabling plant consists of cables running from the routers to other distribution devices. There are no redundant paths for subnetwork cabling to recover from loss of

connectivity. Without redundant paths, survivability of the network is reduced. With redundant cabling in place, alternate connections could be made to restore network connectivity in case of a disaster that severs the primary runs among distribution devices.

There is no current formal disaster recovery plan for the network. Network managers have not formally identified critical network resources such as servers, nodes, and applications. There is no documented plan for reducing the effects of a catastrophic event through redundancy of data, data distribution, computers, or data storage devices.

C. MANAGEMENT AND SUPPORT

The organizational structure of the NPS network management, Figure 3, reflects the structure of the organization as a whole. There is a small amount of centralized management of network architecture and a great deal of decentralized computing resource effort that meets the special needs of a small number of users for a short period of time. There is no current, formal, written strategic plan for computing and the computer network. The position of Dean of Computer and Information Services, Code 05, has been vacant or occupied in an "acting" capacity since January 1993.

Many subnetworks have either no assigned administrator in title or in name or have their own subnetwork administrators. Subnetwork administrators have varying levels of training and expertise. There is no overall structured training plan to keep subnetwork administrators current in the technologies and applications used or anticipated.

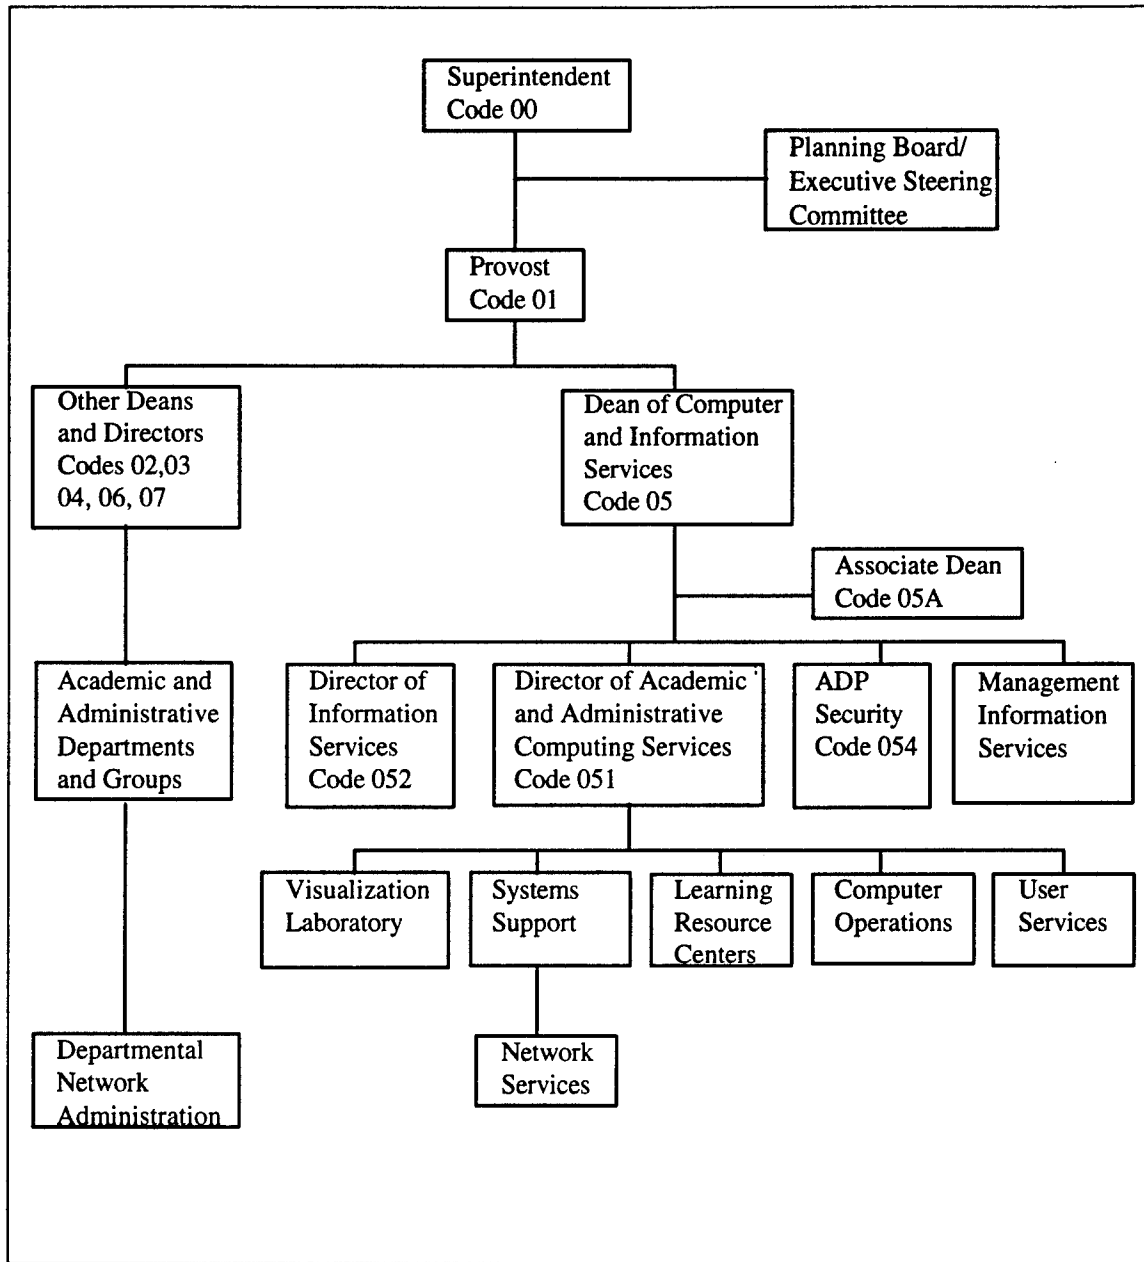


Figure 3. NPS computer network management organizational structure.

There are several advisory and planning committees on computing and networking at NPS. These committees address issues related to current network environment, strategic network planning, computing policies, and computing end-user support services. Their recommendations offer solutions to many network problems.

NPS computer network management is reactionary. [Nolle] observes:

[Network managers] do not proactively manage their networks. [Network managers] have never wanted to manage networks; they'll demand management system availability in case they need it, but they'll try to do without it.

Nolle's observation holds true for NPS as observed during the period of April to September, 1994. While network management applications monitor the network periodically, no person or application consistently monitors network performance. Network technicians devote much of their effort to correcting problems, with little time available for analysis and strategic planning. Network management is considering setting up a network operations center for improved network monitoring and response, but time and funds are devoted elsewhere. Similarly, network administrators are unable or unwilling to devote full time to strategic oversight of network resources (equipment, personnel, and funding) because of other responsibilities.

Organizational support for the network is unclear. Although there are organizational documents implying support of the network (see section on organizational vision in the next chapter), actual support is not evident. There are some indications that this situation may be improving. However, as already noted, the position of Dean of Computer and Information Services, Code 05, remains unfilled by a permanent, full-time,

and qualified person. Funding and billets for the central computing staff have not kept up with the demands placed on central computing [Report].

Without an understanding of the organization's perspective of the network, network managers continue to do their best to provide services to all users. A clear commitment from the organization regarding the criticality of the network would allow network managers to better plan and allocate resources. If the network is a vital service to NPS, the organization must provide resources, including people and funds, to ensure the network can best support its customers. If the network is not a vital service, then the organization should direct the people and funds, presently committed to network support, to areas of the organization considered more vital.

D. USERS AND APPLICATIONS

The users of the network include tenant activities, NPS faculty, staff, and students. The users are located in various buildings throughout the campus. Other users are connected from off-campus sites (e.g., Defense Manpower Data Center). The network supports a wide variety of services and applications summarized below.

- Administration and office management applications
- Research and simulation applications
- WAN connections (e.g., DoD Internet)
- Other internetwork communications (e.g., electronic mail)

E. CHAPTER SUMMARY

The NPS computer network is a heterogeneous collection of network topologies and network resources. This reflects the wide variety of users and demands found in the diverse environment of research, education, and administration. As a vital asset to the organization, the network is not well positioned to best support the organization's changes. The organization has the potential to position itself as a more viable, higher-quality research and academic asset to the Navy and DoD. Without a more viable, higher-quality network, NPS may not achieve that potential. If the network management does not strategically plan in anticipation of the forces acting on it, the school itself becomes a less viable, more vulnerable asset. The following chapter discusses the forces that drive or limit the network architecture's support of the organization.

IV. DRIVING FORCES OF NETWORK ARCHITECTURE

Various forces, ranging from organizational issues to technical issues, significantly impact the role of the campus network at NPS. Individually, these forces may either drive managers to change the NPS network architecture or attempt to keep it from changing.

A. ORGANIZATIONAL ISSUES

1. Organizational Vision and Potential

[Molta], [Sprague and McNurlin], [Wittman], and others emphasize the need for a direct relationship between organizational vision, strategy, and investments in information technologies and network architecture. Business redesign and process improvement must be the driving forces of network architecture if the network is to support organizational improvements. Investments in information technology must clearly be linked to the visions and goals of an organization. These investments must support an organization's strategy and anticipate or quickly respond to changes in an organization's business environment.

a. Vision and Goals

An organization's vision and goals are important driving forces of its network architecture. For NPS, these driving forces can be found high in the Department of the Navy organization. For example, the Office of the Chief of Naval Operations Graduate Education Policy commits the Navy to keeping NPS as the primary source of excellent graduate education. The policy stresses innovation, the highest standards of excellence, and the unique professional needs of the Navy and the Department of Defense.

This clearly signals that NPS must remain on the leading edge of technologies and education to give the Naval Service "a comparative advantage over potential adversaries" [CNO]. The NPS computer network must be able to support this policy.

The NPS Executive Steering Committee is developing an overall strategic plan for NPS. *NPS Mission Statement*, *NPS Vision 2000*, and *NPS Guiding Principles* are parts of the plan. These parts, revised and published in the Spring of 1994, help bring the plan into focus [Wargo]. Four of the ten points in *NPS Vision 2000* are driving forces of network architecture:

It is NPS's vision to be recognized as the graduate school of choice for defense establishment students and as a *premier research university* at home and abroad... [emphasis added]

Our programs will continue to *grow to meet the emerging specific needs* of all services, DoD and the government as consistent with our mission... [emphasis added]

Our research will continue to be recognized throughout the government as providing valuable, responsive and cost-effective products, relevant to current and future defense applications. *We will remain on the leading edge of technology*, management and war fighting improvements... [emphasis added]

Our faculty will be even more sought after as participants in the most *prestigious national and international research activities*, and for high-level DoD positions and consultations... [emphasis added]

Research on the leading edge of most technologies today requires interconnected computers to share data or processing power. The network must be able to support such research by continually offering high availability of network resources and

by rapidly integrating leading-edge technologies with a minimal requirement for configuration change.

NPS Guiding Principles support *NPS Vision 2000*. The communications-enabling properties of the computer network architecture can enhance each of the 13 guiding principles. Two principles directly drive the network architecture:

QUALITY COMES FIRST. As our products and services are viewed, so are we viewed. We will achieve quality through daily emphasis on continuous improvement of our products, services and processes...

INFRASTRUCTURE DEVELOPMENT SHOULD LEAD, NOT LAG PROJECTED GROWTH. Investments in training, technology, and facilities in advance of expected program growth are made when financially possible...

The network architecture, as part of the infrastructure, enhances or detracts from the quality of products and services provided by NPS. The network is part of the underlying technology that must enable growth in programs offered at NPS. Clearly, the vision from the top of the organization down to the guiding principles drive requirements of the NPS computer network. The summary driving force is the need to keep NPS graduate education and faculty research of the highest quality and near the leading edge in many areas.

b. Potential

A more difficult driving force to define involves the potential of the organization as supported by the network. Potential could mean increased prestige, improved organization and management, increased revenues, or decreased costs across the organization. Potential, in this instance, is the increased ability of NPS to recoup costs by

taking on new revenue-generating projects it would not have been able to consider without the enabling network architecture. A potential research sponsor might consider network performance a resource factor in deciding what organization can best meet sponsors' requirements for delivery time and quality. Projects such as integration and interoperability testing, battlefield simulations, three-dimensional modeling, videoteleconferencing, and distance learning all become more "do-able" as the network architecture improves. Users can take on a greater variety or more challenging research projects if the network has a high-level of capability.

2. Inertia, Management, and Funding

Organizational inertia at NPS lead the forces preventing the network from changing to meet the current needs and future demands of the users. Overcoming this inertia is a challenge for any change within the organization. As noted in Chapter III, the network architecture reflects the organizational structure of the school. Each department and lab has its own computing initiatives, sources of funding, and assigned personnel. There is little coordination or pooling of resources among the departments on network initiatives that could be of mutual benefit.

The 1994 Report of the Committee on the Role of Computing at NPS [Report] notes a general dissatisfaction among the users of computing resources. This dissatisfaction has not yet been a great enough motivator to overcome the inertia of the organizational structure that encourages departments to pursue uncoordinated goals. The departments themselves might be frustrated by agonizingly slow bureaucratic procedures

and policies and thus are not sufficiently motivated to pursue coordinated efforts that could benefit NPS as a whole. Government budgeting and procurement processes, personnel actions, billet reassignments, and coordinating mechanisms between and among the academic and military organizational structures all are factors. None of these factors is designed to allow quick coordinated response to program sponsor demands, research initiatives, and internal demands of faculty, staff, and students. Nor are they designed to acquire requisite funding even if a quick, coordinated response could otherwise be mustered. Thus, departments pursue their own goals and manage their own small networks with their own personnel and sources of funding. [Report]

3. Human Resources

Network management's capability to effectively deploy human resources in support of network users is a driving force of network architecture. With the right people in the right place at the right time, managers can actively meet the needs of the users. "The right place" and "the right time" are technical network management issues. "The right people" is a human resources management issue made more challenging by a tight network budget.

Factors involved in maintaining the right people include:

- technical network training and expertise
- application expertise
- credibility with users
- interpersonal skills

Network managers and technicians at NPS have varying skill levels and technical training. No central formalized training program exists to give them the skills and experience they need to better serve their customers, the users of the network, and the applications residing thereon. This results in an inefficient and ineffective maintenance of the network architecture and delivery of service to the users. Quality of network services and user assistance is not consistent across the network [Report]. Users perceive a general lack of competence and coordination among network personnel (even if this perception is based on only one negative interaction with one member of network management). On more than one occasion, technicians from different departments have nullified each others efforts in solving network problems because of lack of communication between the departments, the technicians, and the users. Different interpretations of technical network training between the technicians also caused problems.

B. NETWORK PLANNING ISSUES

The basic tenets of network design are driving forces in planning and implementing changes in the NPS network architecture. Network managers must consider scalability, sustainability, survivability, and security in changes in network architecture.

1. Scalability

Scalability means the ability of the network to grow and change with as little "pain" as possible. Adding more nodes should require only minor configuration changes in the wiring closet and a NIC. Changes in topology or integration of new technologies

should require only changes in a wiring closet and perhaps at affected nodes, while cabling remains intact. Increases in bandwidth demands should require only changes in the distribution devices. Network technicians should only need to work within the wiring closet, changing out or adding a new distribution device, changing an integrated module, changing a port on the distribution device, or making software changes (e.g., a dynamic bandwidth-on-demand algorithm).

2. Sustainability

Sustainability means network managers can maintain the network with a reasonable number of people working a reasonable number of hours. The test equipment, repair tools, and management software effectively isolate and correct problems.

Replacement parts are available at reasonable costs. Maintenance costs and problem histories are maintained and readily accessible. Network monitoring and configuration control can be done remotely. The network configuration is well documented including cable plans and network resource inventory.

3. Survivability

Survivability means that network management has plans and procedures in place that allow the network to recover reasonably quickly and well from interruptions in power and breaks in transmission media. Critical network resources are identified and isolated. Redundancy is built in. Power losses and fluctuations are minimized. Technicians can replace modules in distribution devices without bringing the entire network down.

4. Security

Security means the network resists tampering, theft, and tapping of transmission media and network resources. Network managers must protect resources from malicious or accidental damage to equipment. Physical security helps protect higher layer applications and files from unauthorized access by making it more difficult to physically tap into the network. Network managers should select distribution devices that can be configured to prevent unauthorized access.

C. APPLICATIONS AND HIGHER NETWORK LAYER ISSUES

1. Applications as Bandwidth

"One of the fundamental challenges in this brave new world is grasping an understanding of the bandwidth requirements of new applications [Minoli]." At the network layers that are the focus of this study, applications on a network are considered primarily for the bandwidth required to adequately maintain a data transfer rate acceptable to the user. The demand for this bandwidth by applications is a driving force in the architecture of the NPS computer network.

According to a 1994 survey conducted by Infonetics Research, Incorporated, San Jose, California, 104 network administrators cited the driving factors for higher-bandwidth LANs listed in Table 4. [MacAskill 3]

Driving Factor	percent of respondents
Image Transfer and Imaging	63
Graphics	50
Visualization/CAD/CAE	46
Videoteleconferencing	42
Aggregate Bandwidth	34
Custom Applications	33
Transaction Processing	25
Software Development	24
Groupware	19
Financial Applications	18

Table 4. Driving Factors for Higher-Bandwidth LANs. After [MacAskill 3].

Many of these same concerns are appropriate for NPS. The bandwidth required for certain types of applications are summarized in Figure 4. At NPS, the users demanding the most bandwidth for their applications are faculty, staff, and students involved in research involving simulations, large data file transfers, and real-time applications such as videoteleconferencing. Staff administrative applications demand lower bandwidth at present. With advances in videoteleconferencing, administrative functions might use this application as a means of improving their services.

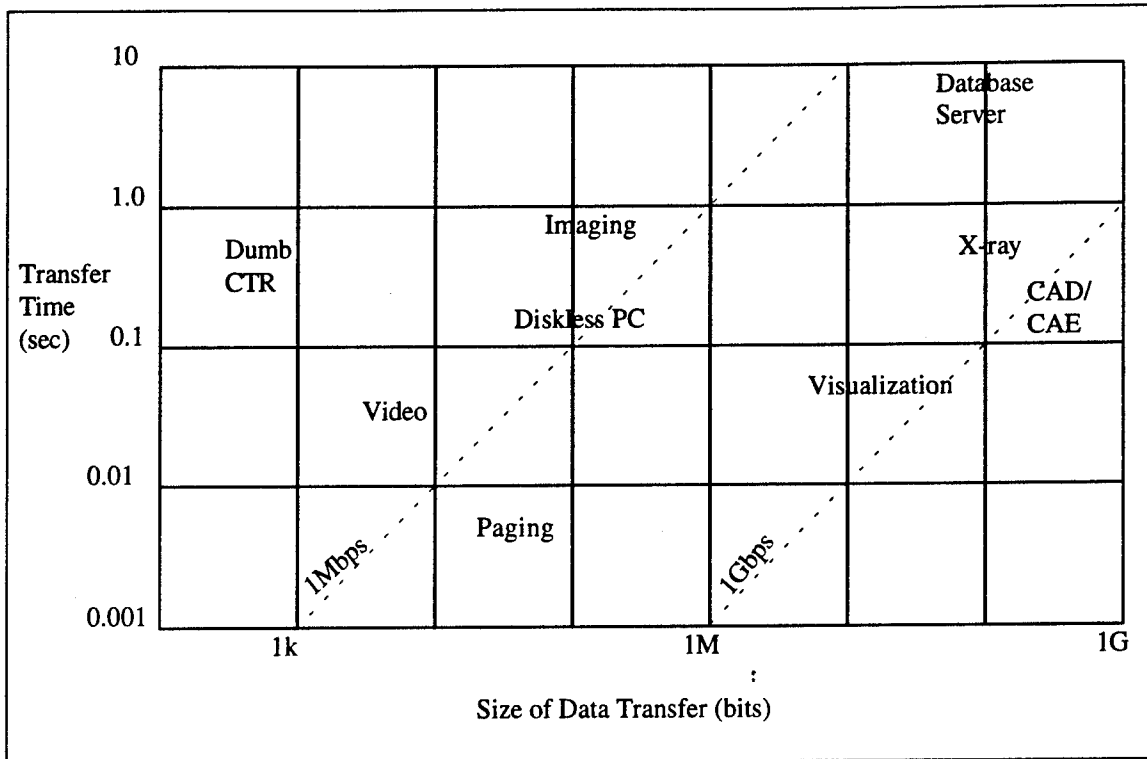


Figure 4. Different application bandwidth requirements. Shown as a function of time and size of data transfers. After [Sirbu].

Electronic mail and file access (e.g., NFS) are presently the most widely used network applications [Metcalf]. According to [Minoli], "As collaborative computing evolves from E-mail and transaction-based data applications to applications that involve person-to-person videoconferencing, video and image data bases, and multi-media, existing LAN technologies will quickly become inadequate." However, as Table 5 indicates, and [McClimans] emphasizes, the real bandwidth requirement for most multi-media applications is not in data rates of tens or hundreds of Mbps that would render NPS 10 Mbps subnetworks inadequate, but a smaller rate that can be considered non-blocking and having a fixed, low-latency characteristic. Today's technology can

deliver high-quality compressed video at less than 1.5 Mbps for full-screen applications [Strauss 1]. The driving force of applications, then, is not any one application (when properly deployed) but the aggregate of many applications employed by many users concurrently across the network.

The challenge to network management is two-fold. First, network managers must more efficiently and effectively use the existing network architecture to its maximum potential. Second, managers must ensure improvements to network architecture are planned, funded, and executed. to meet the future demands of the users.

Application	Transaction length (sec)	Message length (octets)	Throughput (bps/user)
Traditional Database Read	30	1.2k	320
Traditional Database Retrieval	9	1.2k	1.06k
Traditional Database Browse	3	1.2k	3.2k
PC Server (Client/Server)	20	12k	4.8k
Database Retrieval	4	4.3k	8.6k
Image Database Retrieval, Business Imaging System	50	60k	9.6k
Image Database Retrieval, Business Imaging System (evolving applications)	12	120k	80k
Engineering Imaging System	3	36k	96k
Multimedia: Voice Annotated Text			33k
Multimedia: Voice Annotated Image (Business Image Quality)			38k
Multimedia: Voice Annotated Image (high-quality image)			86k
Teleconference			128k
Higher-quality teleconference			768k
Video Distribution (entertainment using MPEG-2 compression)			6M
Visualization: Chemistry	1	80k	640k
Visualization: Genetics	3	1M	2.67M
Visualization: Biology	1	800k	6.4M
Visualization: Fluid Dynamics	1	2M	16M
Visualization: Weather Forecasting	0.2	1M	40M
Visualization: Particle Physics	0.03	3M	800M

Table 5. Data rate requirements for a number of applications typical of collaborative computing. From [Minoli, 1994].

2. Data Compression

Data compression plays an important role in maximizing the use of available network bandwidth. The two major algorithms for data compression are lossy and lossless [Nelson, p. 130]. Lossy compression means some of the data is lost in the compression. This is appropriate for applications such as video, graphics, and sound data where speed of transmission or compression is more important than accuracy and resolution. Examples of lossy compression are Joint Photographic Experts Group (JPEG) and Motion Picture Experts Group (MPEG). JPEG compresses single graphic images. MPEG algorithms compress video and motion pictures. These lossy techniques compress appropriate data to ratios of 10 to 1 with little or no apparent loss of resolution. Ratios of 250 to 1 or greater are possible. These techniques are essential for transmitting real-time video across 10 Mbps LANs. A nominal full-motion real-time video on a screen with 640 by 480 pixel resolution and using 8-bit color at 24 frames per second requires throughput of nearly 60 Mbps. A 10 Mbps LAN can handle a reasonable MPEG lossy compression of 50 to 1.

Lossless compression is slower but allows exact reproduction of compressed data. This is suitable for programs and data files that must not suffer loss of information through compression. It is also suitable for high-quality video and sound files that do not require real-time transmission across a network. Ratios of 2 to 1 are typical of lossless compression algorithms. Examples include Huffman coding, arithmetic coding, the Lempel and Ziv methods, and CCITT V.42bis. Users can send compressed files in

packets across the network at data rates available on the network and reassemble and uncompress them at another node. [Nelson]

Data compression is essential in wide-area connections. Within the NPS LAN environment, there is no visible financial cost for use of the bandwidth. This is not the case for wide area transmissions where common carrier companies charge for access, connection time, and bandwidth usage. Data compression reduces the wide area network connection time and bandwidth usage, thereby reducing telecommunications expenses for the school compared to uncompressed transmissions.

3. Distributed Systems

Phenomenal advances in microprocessor technology and increases in the power of microcomputer processing drive changes in computing away from the centralized "mainframe" paradigm of the past. Distributed systems, such as client/server models, put more capabilities in the hands of end users of computer systems. Computer network architecture must keep pace with this trend to satisfy users who seek to improve their own efficiency and effectiveness at NPS. Distributed high performance processors require high performance connections to maximize the potential of the distributed system paradigm. [Sprague and McNurlin]

4. Network Management Applications

Computer applications assist network management in monitoring and controlling their network. These applications automate many network management tasks such as network traffic monitoring, resource inventory, resource status, and configuration control.

When properly deployed, these applications allow more efficient and effective centralized management of the network with reduced manpower by identifying network problems and identifying possible causes and solutions. [Sprague and McNurlin]

D. TECHNOLOGY ISSUES

1. Practical Limitations and System Bottlenecks

a. Channel Capacity

All transmission media of practical interest in local area networks are of limited bandwidth. The limitations are a consequence of the physical properties of the transmission medium and deliberate limitations at the transmitter. The transmitter limits the bandwidth to prevent interference from other sources or to control costs and complexity of both the transmitter and the receiver. The challenge, then, is maximizing efficiency of a transmission channel with a given bandwidth. Factors such as white noise, impulse noise, attenuation (as a function of frequency and distance), and delay distortion (intersymbol interference) all effect the capacity of any transmission scheme. [Stallings 2, pp. 57-66]

[Shannon] presents a "capacity" formula for the theoretical efficiency of transmission schemes. In this formula, the channel capacity, C , in bps, is related to the bandwidth, W , in Hz, and the signal-to-noise ratio, S/N , measured in dB.

$$C \leq W \log_2(1 + S/N) \quad (\text{Equation 1})$$

There are limitations to the use of this formula beyond the scope of this study but it does give a relative measurement for comparison of transmission schemes. Figure 5

summarizes the practical limitations of transmission media typically used in local area networks.

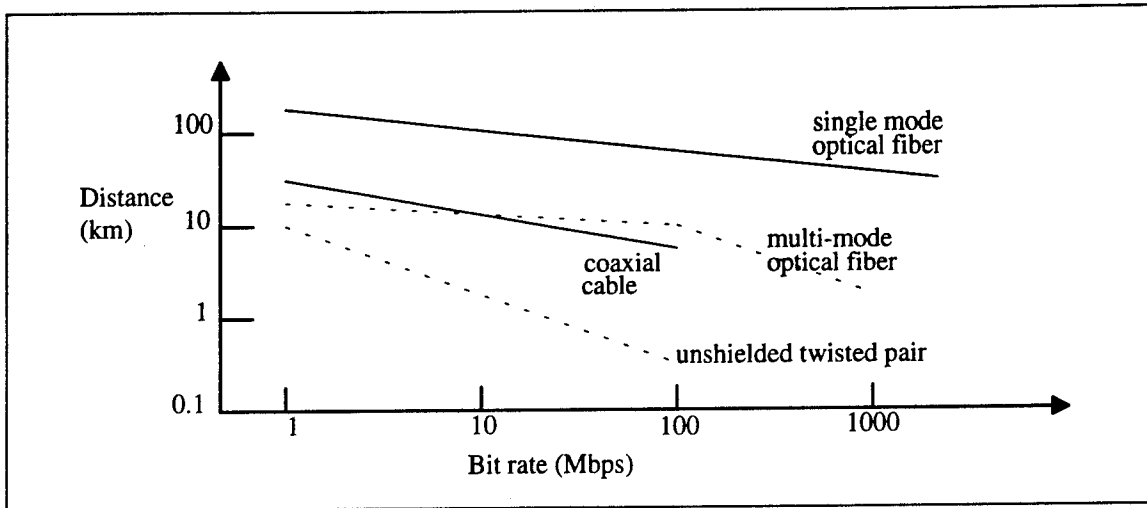


Figure 5. Data rate vs. distance for various transmission media. After [Fraser and Henry].

Practical limitations of a transmission medium in a LAN environment are not in the transmission medium itself but in the engineering trade-offs made to send a baseband digital signal over the transmission medium. There are engineering trade-offs involving data rate, cable length, number of cables installed, number of nodes, electrical characteristics of the cable, encoding scheme, and transmitter/receiver complexity and sensitivity. For example, all else being the same, to achieve a higher the data rate, the length of the transmission medium must be reduced to maintain the same quality transmission. [Stallings 1]

Compared to other transmission media, twisted pair is limited in distance, bandwidth, and data rate. Attenuation for twisted pair is a very strong function of frequency of the signal. In baseband digital signaling systems such as IEEE 802.3

10Base-T, a 10 MHz signal carries a digital signal over 100 meters on one pair of unshielded twisted pair (UTP) [Bryan 1]. The encoding scheme essentially translates one bit of information into one bit of transmitted signal on the wire resulting in a 10 Mbps channel capacity. The new Fast Ethernet Alliance 100Base-TX specification uses a 125 MHz signal and an encoding scheme that allows a data rate of 100 Mbps over the same distance on Category 5 UTP [Intel 1]. Other higher-capacity schemes use lower frequencies (25-30 MHz) and data rates per pair but use multiple twisted pairs to achieve similar data rates to 100Base-TX. Category 3 UTP can maintain higher-capacity specifications at this lower frequency. Conceivably, the right encoding scheme on a standard 4-pair segment of Category 5 UTP could carry an aggregate data rate in excess of 500 Mbps over 100 meters. Other losses, such as cross-talk interference among twisted pairs and intersymbol interference, increase with frequency, thereby limiting the practical distance and data rate of UTP.

UTP is an important factor to be considered for the NPS network architecture. Its low cost, scalability (especially if additional pairs are already in place), and transmission characteristics make it ideal for star-topology network configurations with radii less than 100 meters.

Coaxial cable has better frequency characteristics than twisted pair and hence can be used effectively at higher frequencies and data rates over greater distances. Because of its shielded concentric construction, coaxial cable is much less susceptible to interference and cross-talk than twisted pair. The principal constraints on performance are

attenuation, thermal noise, and intermodulation noise. Intermodulation noise is not a factor for baseband digital signaling. Experimental systems have achieved data rates of over 800 Mbps using coaxial cable and broadband signaling techniques [Stallings 2, p. 75]. However, for 50 Ohm thicknet coaxial cable, which requires two-way baseband digital signaling on a shared medium, the maximum usable bandwidth, W , is about 25 MHz and the maximum attainable signal-to-noise ratio, S/N , is just over 40 dB. Applying Shannon's theorem, thicknet coaxial cable is limited to just over 134 Mbps in a LAN environment. Thinnet is similarly limited over shorter distances.

Thicknet and thinnet coaxial cables make up a significant portion the NPS network cabling. Coaxial cable's practical bandwidth and data rate are not limiting factors for use as a transmission medium. Other non-technical factors, such as standard network interfaces, economics, and ability to adapt to higher-performance network configurations, limit coaxial cable's usefulness in architecture changes. For example, there are no known 100 Mbps distribution devices that have thinnet connections at 100 Mbps ports.

Optical fiber outperforms both twisted pair and coaxial cable in LAN environments. Because optical fiber transmissions use a different part of the frequency spectrum, they are immune to electromagnetic interference (EMI) and radio frequency interference (RFI). The lowest grade of optical fiber signaling, light emitting diodes over multi-mode optical fiber, has a data capacity roughly five times greater and over a distance ten times greater than that of UTP. This makes it useful for longer distances transmissions and higher capacity channels such as network backbones.

Wireless network technologies have not evolved to the point of being considered for higher capacities than the guided media network technologies being used at NPS. Network managers, however, should not rule out using wireless options in certain situations. Wireless technologies are appropriate for areas where use of other transmission media is cumbersome, such as open spaces in the library, older buildings with solid walls, floors, and ceilings with no place for cable runs, and from building-to-building when cabling is technically or economically infeasible. Wireless alternatives are appropriate for some sections of Hermann Hall.

b. System Bottlenecks

A network or an application running on a network is no faster than its slowest link. [Metcalfe], the "father" of Ethernet, observes that a network is seldom the bottleneck in end-to-end systems. More often, bottlenecks in applications are in disk access, operating systems, and input/output (I/O) buses. Potential bottlenecks in a system are summarized in Figure 6. Network managers and users alike must look at applications and the network end-to-end to identify the bottlenecks in performance before investing in higher-capacity networks as the solution.

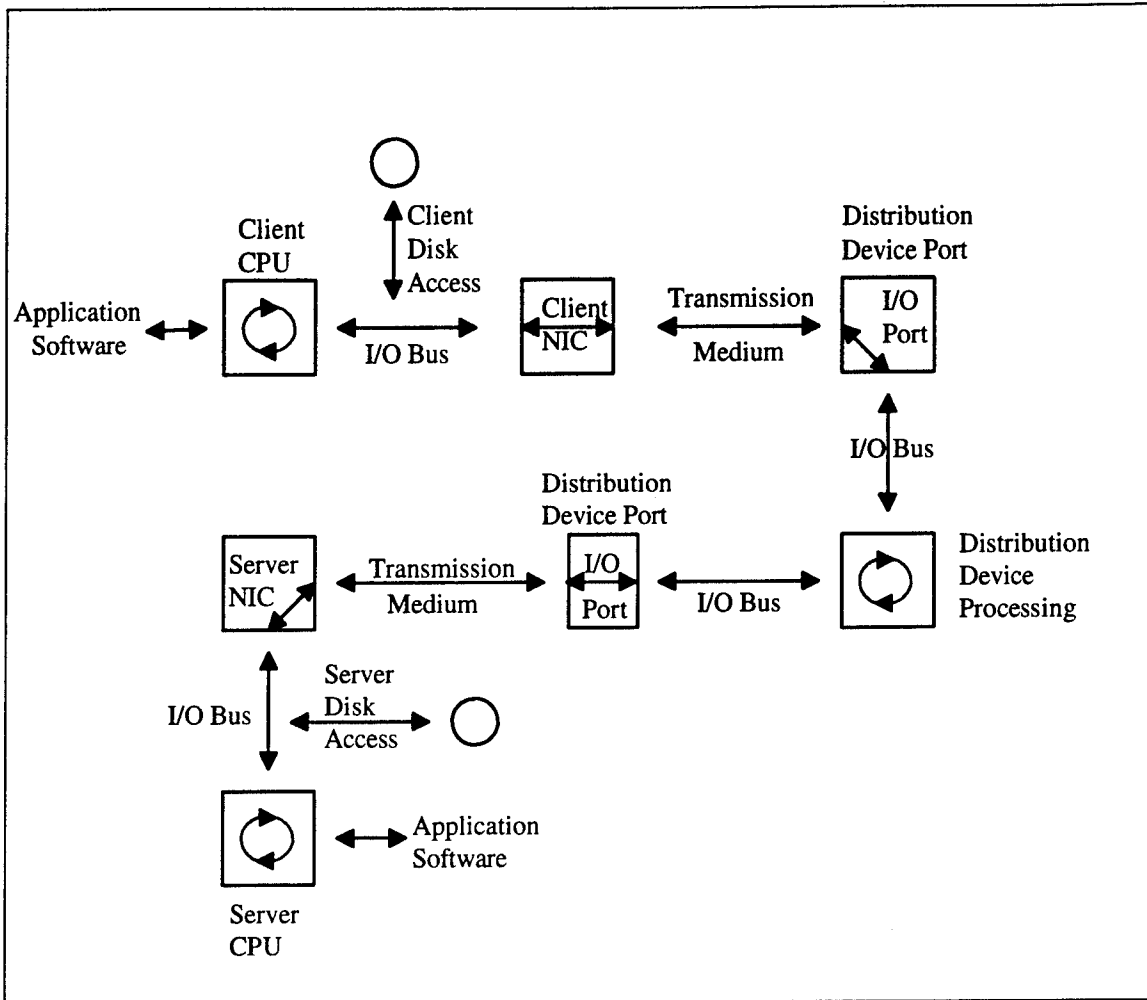


Figure 6. Potential system bottlenecks from end-to-end.

[Bryan 2], in a discussion about system bottlenecks, finds microprocessors like the Pentium, PowerPC, and other RISC chips deliver hundreds of MIPS [million instructions per second] to the desktop. However, associated component subsystems cannot deliver data to those CPUs at anything close to processor capacity. Sophisticated applications demand evermore resources such as access to end system storage or information residing elsewhere on the network. The result is an I/O bottleneck.

The culprit is input/output (I/O) bandwidth caused by peripheral buses within devices at the nodes on the network such as servers and computers running compute-intensive applications. At 100 Mbps, bus capacity is a significant limiting factor in overall network performance. [Hewlett-Packard 1]

[Abrahams] proposes overcoming I/O bottlenecks at several levels. First, he recommends using intelligent NICs with plenty of memory for buffering. Also, NICs should provide speed-matching buffers and both data link and transport layer services. Finally, the network operating system should provide both lower-speed and higher-speed transfer service.

Table 6 summarizes theoretical and practical throughput for some I/O buses in devices connected to the NPS computer network. The vast contrast between theoretical and practical throughput has more to do with performance of the typical NIC connecting the bus and demands by other peripheral devices on the bus. Lower-end NICs rely heavily on the host CPU to do much of the processing of the network data. Higher-end NICs mirror [Abrahams]'s recommendations. In order to maximize the available performance of the network, devices must support intelligent NICs on higher-capacity buses.

Bus	Maximum bus throughput (Mbps)	Practical NIC throughput (Mbps)
ISA (PC)	40	0.2-0.4
EISA (PC)	256	1.3-5.0
PCI (PC)	256/512	3.5

Table 6. I/O bus throughput. After [Ricutti], [Glass], [Abrahams], [Poutain], [Newman], [Infonetics] and [Bryan 2].

2. Network Technologies

Advances in network technologies drive changes in NPS network architecture. It is clear from Table 6 that older technologies supported by ISA I/O buses are significant bottlenecks to higher-bandwidth applications and network technologies [Strauss 2]. Machines on higher-capacity network segments must have greater than ISA I/O bus capacities or the higher-performance network technology is overkill. NPS network managers do not have to discard these older machines and NICs but redistribute network resources to better match overall network performance with user needs.

Increased channel capacity of network topologies and creative ways of using available bandwidth (e.g., compression, micro-segmentation, and switching) increase the number of options available to network managers for improving the performance of their LANs. Reductions in performance/size and performance/price ratios for network technologies are also driving forces in network architecture options. Figure 7 shows trends that make higher-performance network technologies more attractive as time advances.

Higher-performance technologies may be more attractive but there are risks involved in investing too heavily and too early in the life of these technologies. An example is Asynchronous Transfer Mode (ATM). Although this technology promises to provide scalable, higher bandwidth and performance, it is still in its infancy. As "leading-edge" network managers apply this technology to their LANs at great expense, they are rediscovering many problems that were solved long ago in older, more mature

network technologies [Strauss 1]. Older, more proven alternatives present less risk and cost to network managers but at lower levels of performance. Figure 8 provides a relative comparison of maturity of higher-performance network technologies.

NPS network managers may find some advantages in investing in leading-edge network technology. Research projects in computer networking and higher-bandwidth applications could benefit. They must consider available alternatives and remain flexible in deploying these alternative to allow for changes and improvements in network technologies.




Figure 7. Cost/Performance trends of network topologies. After [Pigg].

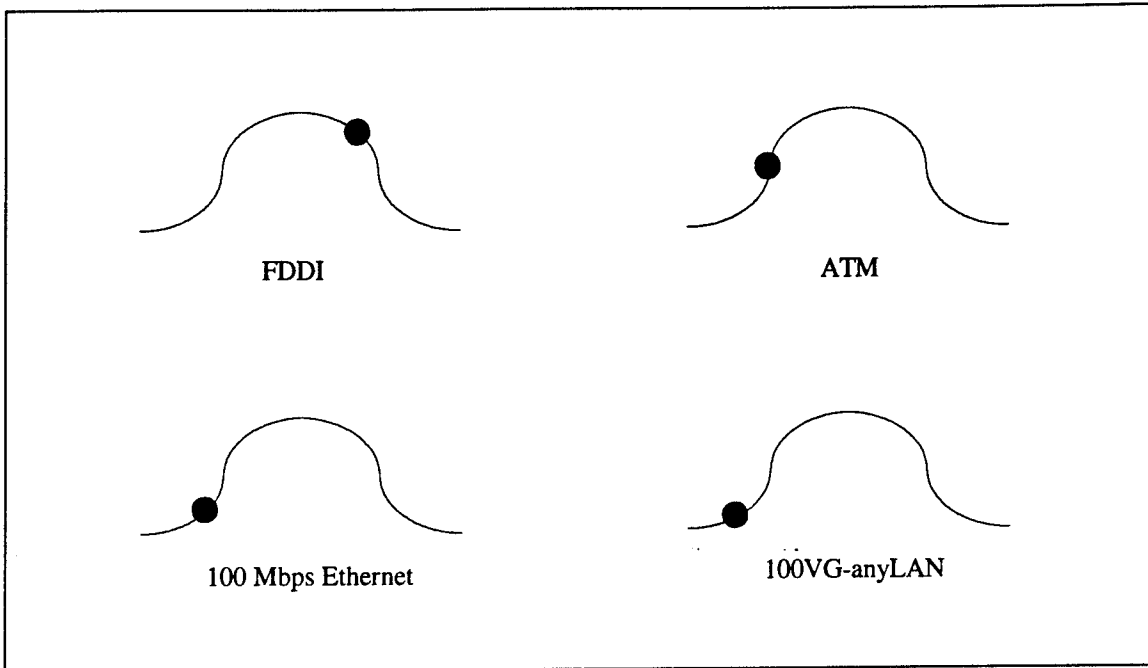


Figure 8. Maturity of higher-performance network technologies. After [Pigg].

E. CHAPTER SUMMARY

Numerous issues drive the network architecture at NPS. The most important issues are organizational. Although a vision for the organization as a whole exists on paper, it is not yet embraced by the entities that comprise the organization. The network is a support function for the organization's wide variety of users on the campus. Each group of users has goals and missions and requires different levels of support from the network. Disparate funding sources and inconsistent distribution of network management and technicians with varying levels of knowledge and expertise across the subnetworks make the network as a whole difficult to manage.

The client-server and distributed computing paradigms have changed the emphasis of the centralized computing resources. Computer technology has distributed the power of computing to the end user. This has resulted in applications that demand more from the network architecture. Few of these applications, by themselves, drive the need for higher-performance networks. However, their aggregate use and the increased number of users do drive the need for increased network performance.

Higher-performance networks enable the use of new technologies and development of newer and larger applications. Network technology continues to advance. Decreasing costs of network technologies and the emergence of newer alternatives present myriad choices to network management.

As a result of these numerous driving forces, NPS network managers must plan, build, and maintain a flexible, maintainable network architecture. This architecture must meet the needs of the organization and its users while integrating and enabling new technologies and applications. Subsequent chapters explore issues related to such a network architecture.

V. MEASURES OF NETWORK BENEFITS

This chapter explores measures of benefits used in decision making regarding network architecture alternatives. The challenge in measuring the benefits of a certain network architecture is determining the appropriate metrics. Some measures focus on efficiency of the network architecture (how well the network performs). The more difficult measure is the effectiveness of the network architecture. [Sprague and McNurlin] contend that the largest payoffs in any information technology system lie in improving effectiveness rather than in efficiency but most measurement techniques tend to focus on efficiency. Thus, many benefits remain unmeasured. Investments in network architecture are difficult to cost-justify because the network itself provides no direct benefit. Only the applications that run on the network can provide measurable benefits to an organization. [Sprague and McNurlin]

Because the network by itself does not provide direct benefits to the organization, but only supports the applications that run on it, its goals must include providing the best support to applications that benefit the organization. Thus, both measures of efficiency and effectiveness are appropriate in determining its worth. Economic and performance measures attempt to determine the efficiency of the network. Other measures attempt to determine the effectiveness of the network.

A. EFFICIENCY MEASURES

A plethora of efficiency measures are available to network managers. Some are simple cost measures while others go into great detail on measured and theoretical performance. This section discusses a few measures that NPS network managers might consider in measuring the benefits of the network in terms of efficiency.

1. Cost per Port on a Distribution Device

The cost per port is a phrase that is often used by vendors in trying to sell their network distribution devices. Some network managers use this as a quick figure for comparison of similar distribution devices. Its advantages are that it is readily quantifiable, available, and able to be forecast based on historic trends. This measure's weaknesses are that it only looks at one small part of the entire network, it does not consider the number of nodes serviced by each port, nor does it consider the demands of each port. This measure also ignores life cycle costs.

2. Cost per Network Adapter

A more telling measure than the cost per port on the distribution device is the cost per network adapter or NIC, especially in shared-medium topologies such as ethernet and token ring. This measure, like the previous measure, is readily quantifiable, available, and able to be forecast based on historic trends. It likewise suffers from not identifying all the costs, demands of the network as a whole, or life cycle costs.

3. Cost per Node on the Network

This measure is a combination of the two measures above plus other network costs. Cost per port of the network distributes the cost of the network among all the nodes on the network. It factors in costs of distribution devices, network cabling, connectors, and network interfaces.

4. Cost per Megabit of Performance

[Pigg] presents three measures of efficiency used by the Yankee Group, a communications industry research, planning, and consulting organization, for comparing networking technologies over a forecast period of several years. The first is cost per connection, which is similar to the cost per node on the network discussed above; the second is cost per Megabit of performance for each LAN technology; and the final measure is cost per Megabit of performance per port on the subnetwork (assuming ten users per shared-media LAN). These measures are used to project the costs of network alternatives three years into the future based on historical data and industry forecasts.

[Pigg]

5. Bandwidth per Port

A simple performance measure of any distribution device is the maximum bandwidth available at each port. This measure looks very narrowly at the potential performance of just one port. It does not consider the aggregate capacity of the device or what effect of bandwidth demands on the device as a whole has on performance of an individual port.

6. Packet Throughput

Network World, *PC Week*, and other testing laboratories have established benchmarks for measuring switch performance on maximum throughput (e.g., ethernet packets/second) without packet loss based on packet size [Haugdahl]. [Choi and Kim] model the expected throughput for the ethernet protocol based on a network load density function, an Internet Protocol (IP) packet density function, a probability of no arrivals, and an equivalent load in time function. [Stallings 2, Chapter 6] provides simple performance models for maximum possible throughput for both ethernet and token ring networks. Packet throughput has the advantage of providing a measure of channel efficiency.

Throughput provides useful data put through the communications channel. It does not just measure the full capacity of the channel, but instead removes the "overhead" data required for a particular scheme below the network layer.

Throughput accurately portrays the data rate of information that is useful at the network layer. Taken alone, throughput is a good measure for distribution devices but, by itself, this measure is inadequate for measuring performance of end-user devices. For those devices, it must be used in conjunction with a measure of CPU utilization to give a more accurate measure of efficiency [Intel 2]. Varying packet sizes and the bursty nature of network traffic also affect throughput.

Packet throughput provides a measure or a group of measures for set of given packet sizes transmitted at fixed intervals. This may not represent actual network traffic,

but it does provide a measure for comparison with other alternatives or with established benchmarks.

Network managers must use caution when comparing alternatives to benchmarks. Some devices are optimized to perform well based on a certain benchmark. A benchmark may not accurately reflect the conditions of their particular network under which the devices might perform substantially different. [Weston], [Yager]

7. Network Simulation

[Hamilton] outlines no-cost network management tools collected or developed by the Texas A&M Computer Simulation Research Group and used for obtaining predictive data about networks. These tools are available from file transfer protocol (ftp) sites on the Internet. These and numerous commercial products allow network designers to model and test network configurations under various conditions on a workstation software package without actually disturbing the network or investing in cabling and distribution devices. Some of the commercial products are Cadence Design Systems, Incorporated's BONES PlanNet; Synetics' LAN SoftBench and LANSim; MIL 3, Incorporated's OPNET; and CACI Products Company's NETWORK II.5 and COMNET III.

If the model or algorithm used to simulate the network fails to account for some actual occurrences on the network, then problems will only be discovered after changes to the network are actually implemented.

8. Capacity Assessment

[Vis] presents a simple formula for the assessment of LAN performance in lieu of sophisticated analytical models and tools. It is based only on the most essential performance parameters. Vis contends his "rule of thumb" is a fairly accurate conservative measure when compared to a closed queuing model of a time sharing system. Appendix D discusses this measure in more detail.

Capacity assessment provides a quick, no-cost estimate of performance of some networks under certain conditions. It is not appropriate for networks more complex than a single shared-medium topology such as the NPS computer network.

9. Network Monitoring

The most practical approach to measuring the performance of network architecture is monitoring the "vital signs" of the network over time and applications use. Various commercially-available, integrated or single-use, network management tools monitor and analyze network performance. They have varying capabilities, which include such things as protocol analysis, network availability, network and application utilization, traffic analysis, cable testing, and other features. Network management applications also monitor maintenance history, repair costs, and time required for management of the network. When used effectively, these applications provide a means of documenting life cycle costs of network architecture. [Miller 1]

Specialized equipment also measure performance of a network architecture. An example is a Hewlett-Packard LANalyzer series network analyzer, which remotely

monitors packet counts, collisions, and ethernet network anomalies such as jabbers and runts. Network technicians use hand-held, battery-powered analyzers for mobile trouble-shooting on network segments.

Network monitoring provides real-life efficiency data. When combined with network maintenance data, monitoring provides a clear picture of network performance. Unfortunately, the only way of truly determining how a change in application or network configuration will effect network performance is making the change and then monitoring the performance. If the changes cause problems on the network, then users of the network suffer through decreased performance or loss of service while the problems are rectified.

B. EFFECTIVENESS MEASURES

Effectiveness of a network is more difficult to measure because it must be linked to the (often intangible) goals of the organization. This section provides network managers with a sampling of available measures of effectiveness to be considered when determining the benefits of the network.

1. Business-Value Approach

A business-value approach relates network technology strategic issues to economics vs. productivity concerns in an organization [Feldman]. [Spada] concurs and recommends that investment in network technology be measured not only by the price of the products and services alone but also by the strategic benefit of the technology for an organization. [Spada] makes this case in support of investments in ATM but a

business-value approach can apply more broadly to measure the effectiveness of computer network architecture in general. Network managers must consider how investment in network architecture will further the business goals of the organization by helping them manage information more effectively and competitively. Investments in network architecture are more justifiable when quantifiable measures of the contributions of network investments are applied towards issues such as decreased business time cycles, higher quality products and services, and increased customer satisfaction. Technology must link to business goals to increase the perceived business value of the investment in order to offset the costs of the investment. The network architecture must deliver a critical commodity, information, to people who need it when they need it in order to improve the effectiveness of the organization in meeting its goals. [Spada]

Along a similar tack, [Capetanakis] suggests matching performance and cost factors to the needs of an organization; Table 7, below, summarizes the factors to be considered. He also recommends considering product maturity, installed cabling plant and support for specific application types such as multimedia when determining which alternative is best for an organization.

Technology Issue	switched ethernet	fast ethernet	FDDI	CDDI	ATM
LAN type	switched	broadcast	broadcast	broadcast	switched
Mbps to desktop	10	100	100	100	155
distance	100m	100m-2km	2km	100m	100m-2km
status	stable	emerging	stable	stable	emerging
standardized	yes	9 months	yes	yes	partially
cabling	UTP-3	UTP-3,5, fiber	fiber	UTP-5	UTP-5, STP, fiber
multimedia	maybe	maybe	no	no	yes
price of hub and interface card	\$700	\$500	\$3000	\$2000	\$4000

Table 7. LAN upgrade technology tradeoffs, April 1994. After [Capetanakis].

2. Quality of Service

Another measure of effectiveness of a computer network architecture is users' perception of quality of service provided by the network. If users' needs are satisfied and remain satisfied for the life cycle of the network architecture, then investment in changes or upgrades to the network have no added benefit according to [Gibbs]. [Gibbs] points out that networking is about supporting personal productivity and not an end in itself. He states, "What matters most is what happens at the desktop...If users are not adequately supported then the network is failing." User satisfaction surveys attempt to capture how well a network is supporting its users.

[Seitz] discusses user-oriented objective measures of quality that are primarily designed for voice and video transmissions over telecommunications networks. These measures attempt to correlate end-user perceptions of satisfaction of a system with

objective measurements of system parameters. End-users are not interested in the internal architecture of the network, but only the quality of the system's end product. Users perceive differences in quality of network alternatives independent of the technology behind the alternatives. Thus, measures can objectively and subjectively compare network alternatives. The objective measures include speed, accuracy, and dependability of access; information transfer time; and transmission disconnect time. The subjective measures include user satisfaction with a transmission, considering the same parameters as the objective measures, and as collected by opinion surveys. [Seitz]

Criterion Function	Speed	Accuracy	Dependability
Access	-Access time	-Incorrect access probability	-Access denial probability
Information transfer	-Throughput -Block transfer time	-Block error probability -Block misdelivery probability	-Block loss probability
Disconnect	-Disconnect time	-Disconnect failure probability	

Table 8. Parameters for user-oriented measures of quality. After [Seitz].

The NPS network administration can develop surveys to measure user satisfaction and perception of quality. The perception-of-quality measures described by [Seitz] are especially appropriate for videoteleconferencing and distance-learning applications. The advantage of these measures is that the users of the network are involved in determining the effectiveness of network performance. The user-oriented objective measures of

quality are in their infancy and have few correlating studies outside of Seitz's work. This may provide an opportunity for further research at NPS.

Prestige closely relates to quality of service. An organization gains intangible benefits among its peers when its network provides the highest quality of service or implements the latest networking technologies in an effort to provide the highest quality of service. The respect and admiration drawn from similar organizations and the attention paid by others may increase an organization's ability to attract more customers.

Quality-of-service measures capture user feedback on the benefits of the network. This feedback may offset indications by measures of efficiency that might lead network managers to believe otherwise. If performance measures indicate the network is doing fine but the users are not satisfied, then that issue must be addressed. Similarly, if efficiency measures indicate the network is performing below par but the users are satisfied, then network managers must reconsider their priorities for changes to network architecture.

C. NPS MEASURES

NPS network management should use cost per node on the network when comparing costs of network alternatives. Network simulation and network monitoring should be combined to assess the performance of network alternatives. Network management could then combine the economic and performance measures for a single metric based on the ratio of cost to performance. Economic measures that consider total network costs and performance measures that encompass the entire network are

appropriate for use when measuring the efficiency of the NPS computer network.

Network managers can compare costs of existing configurations with the costs of proposed configurations. Similarly, they can compare the performance benefits of existing configurations with proposed changes in configuration by monitoring, modeling, and simulations.

NPS network management should measure the effectiveness of network architecture through studies based on user satisfaction with test-bed network configurations. Network managers could set up alternative architecture configurations, measure user satisfaction opinions and compare results to determine appropriate network alternatives.

As a diverse academic institution, NPS may have difficulty in applying business-value approaches. Private sector industries and public sector organizations with well-defined goals and customers can apply these measures to a "bottom line" figure. Disparate funding sources and expenditure accounting at the school make it difficult to quantify a "bottom line." Therefore, the benefits of the network cannot clearly relate to a financial figure for the organization or for any particular department within the organization.

The network helps different departments, tenant organizations, and different users in varying ways that may be difficult to quantify. Because the network provides support to a variety of users with differing goals and functions, it is difficult relating the network benefits to definitive business goals. However, as a research institution, NPS users need

technology near the leading edge which may require support by the network. To remain near the leading edge, the network must be able to adapt to changing technologies quickly and at minimal incremental cost.

Opinion surveys and/or direct observation could measure quality of service to each department or user before and after changes to the architecture both in test-beds and on the entire network. Network management can use these measures to gauge the effectiveness of the network in keeping it users near the leading edge of technology.

Network management can proactively manage the network using these measures. Planners should incorporate user satisfaction into network objectives by ensuring a planned architecture will satisfy user needs and respond to growing and changing needs in the future. Similar instruments could measure prestige among peer organizations, customers (program and research sponsors) and potential customers.

The problem with this approach is that true strategic benefits and quality of service cannot be seen until *after* significant investments are made. These investments change the existing architecture and enable applications that take advantage of the change. Both the changes and the enabled applications require significant capital investments.

As a follow-on to the comparison of network configurations, network management could compare the costs of incremental changes in performance. For instance, if a particular architecture is in place, how much does it cost to increase the performance of the network by replacing components? Do the distribution devices and NICs need to be changed? Do the devices at each node need to be changed? Does the transmission media

need to be replaced? All these incremental costs could be considered when comparing network architecture alternatives discussed in the following chapter.

VI. REASONABLE ARCHITECTURE ALTERNATIVES AND TRENDS

This chapter discusses alternatives of network architecture, considers relative advantages and disadvantages of each, presents trends of local area network architecture, and relates the alternatives to NPS.

A. ARCHITECTURE ALTERNATIVES

As local area networks become interconnected and the number of users and applications grow, network administrators seek alternatives to network architecture that meet the needs of the users and that allow some means of managing the network.

Network administrators face myriad combinations of cabling and distribution devices. The following discussion addresses the main components of network architecture, shown in Figure 9. The discussion then presents alternatives for each component.

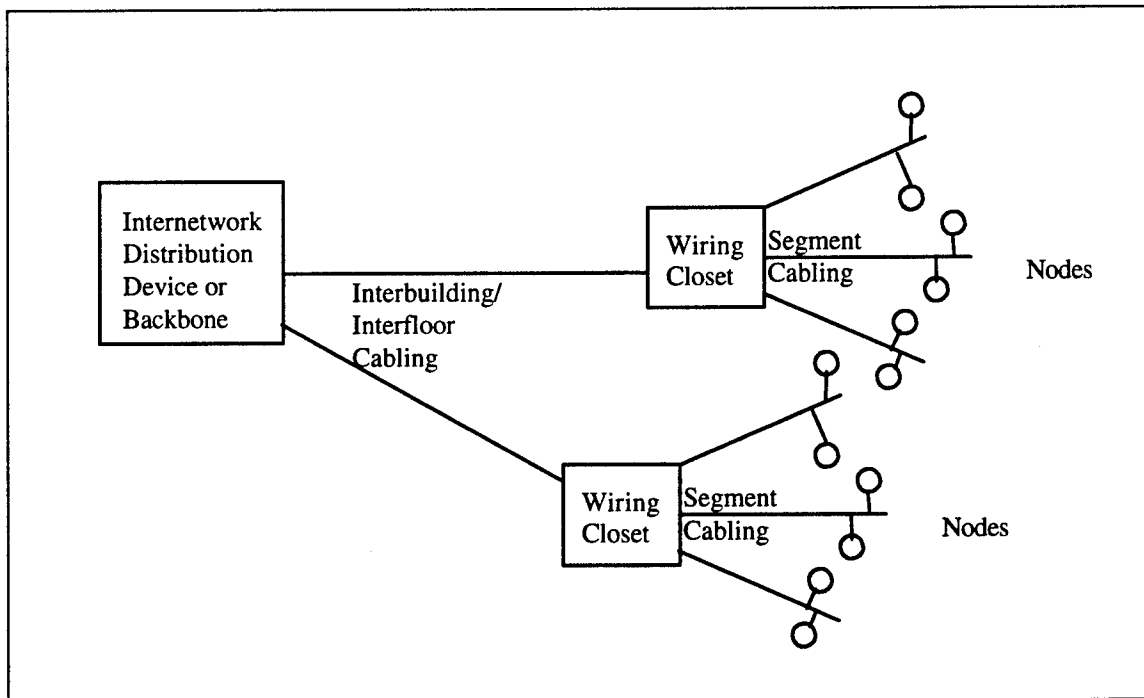


Figure 9. Typical NPS subnetwork architecture.

1. Cabling

The three main categories of cabling are cabling between buildings on campus ("interbuilding") and between wiring closets; cabling within each wiring closet; and cabling from the wiring closets to each node on a LAN segment. A wiring closet is considered any central distribution or servicing point for cables in a network. Figure 10 shows the categories of cabling.

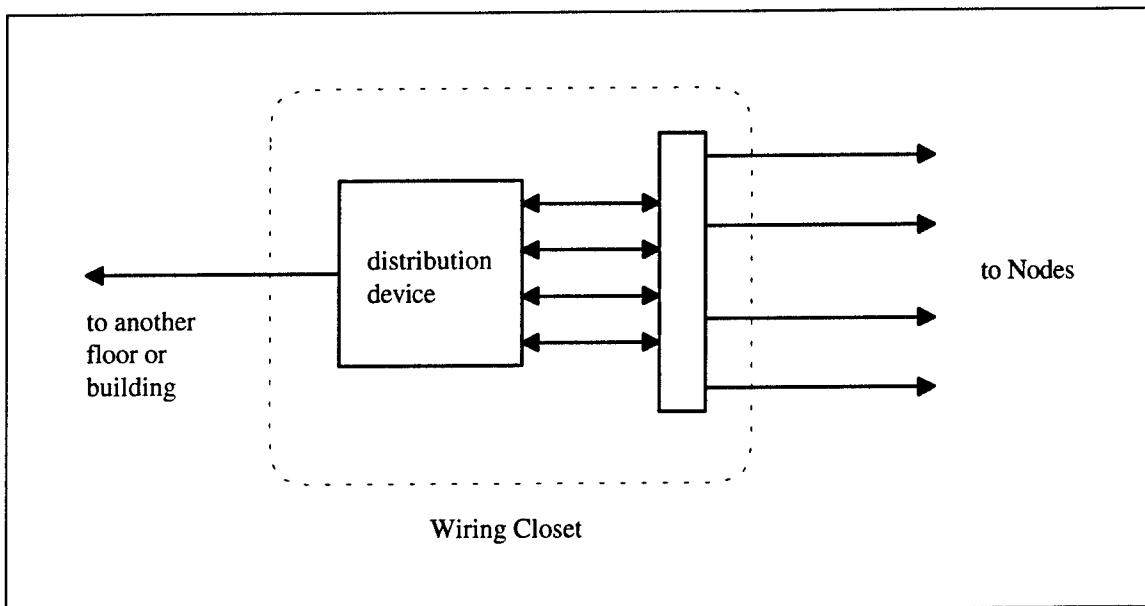


Figure 10. Three main categories of cabling in a LAN.

a. Cabling Between Buildings and Between Wiring Closets

This area of cabling addresses internetwork connectivity. The primary requirement for this cabling is capacity of the transmission medium over the distance between two distribution devices. The financial costs of time and effort related to the installation of this cabling normally outweigh the cost of the medium itself. [Marks] notes

that labor charges typically comprise 40 to 80 percent of new cable installation costs and 60 to 90 percent of the costs involved in changes to an existing cable system. [Marks] also recommends that "the transmission medium chosen should meet the network's anticipated needs for at least 10 years; 20 would be even better." The Naval Postgraduate School's Connectivity and Network Committee recommended that any cable installed in support of a campus networking project meet the networks needs for 20 years [Leahy]. Therefore, installation of a new cabling plant should allow for growth, including higher-bandwidth networks and additional users that might be added in the next two decades. The reasonable alternatives for campus LAN internetwork transmission media are multi-mode optical fiber and coaxial cable.

Coaxial cable, often called "backbone cable," is the less expensive of the two options. It is easier to pull and more tolerant to bending than optical fiber. The IEEE 10Base5 standard limits the length of this 50 Ohm coaxial cable to 500 meters between repeaters and a capacity of 10 Mbps using ethernet signaling techniques. The standard allows up to four repeaters between any two nodes which effectively extends the cable to 2500 meters. End connectors are less expensive and easier to install than optical fiber connectors.

Multi-mode optical fiber has advantages over coaxial cable in data capacity, and cable length between repeaters. In network systems, 62.5/125 multi-mode optical fiber can operate between 112 and 420 Mbps over 1000 meters depending on system operating wavelength [Botelho]. The standards for FOIRL ethernet limit the

length of multi-mode optical fiber to 1000 meters between repeaters. The newer IEEE 10BaseF (10 Mbps) ethernet standard limits the length of the optical fiber segments to 2000 meters between repeaters.

The physical topology among the buildings on the campus or between wiring closets can be mesh, star, bus or ring as discussed in Chapter II. Mesh and star topologies are the more fault tolerant options. Alternate paths can be established or troublesome paths between distribution devices can be isolated using these two options. Ring and bus topologies normally require less cable and time to install.

b. Cabling within Wiring Closet :

Cabling within a wiring closet links the distribution devices within a wiring closet and the cable that connects to the nodes on the subnetworks. In its simplest form, the cable from each node or LAN segment connects directly to the distribution device. This form minimizes transmission losses due to intermediate connections at the expense of configuration flexibility. Each connector or termination of a cable introduces additional transmission losses. Each connector on a segment cannot maintain the same connectivity quality as the transmission medium could by itself nor can each connector exactly match the impedance of the transmission medium. Quality of the transmission signal remains better with fewer connections between transmitter and receiver. The trade-off is that direct connections do not have flexibility that allows combining, separating, and rerouting segments of cabling to alternate paths. An option for UTP terminates cables from nodes or segments at a connector block and subsequently connects segments to the distribution

device via 25 pair telco cabling. Other types of cables use patch panels and patch cords with AUI, BNC or optical fiber connectors. Pre-wired connector blocks, which connect to the distribution device, provide for a relatively neat wiring closet at the expense of flexibility. The most flexible option uses patch panels to terminate the cable from the node and patch cable with modular connectors to the distribution device. This flexibility comes at the expense of increased transmission losses, which are introduced by additional modular connections. This option also makes it difficult to trace the cabling as patch cables hang down over the patch panel, the distribution device, and other patch cables.

c. Cabling from Wiring Closet to Network Nodes

Alternatives for cabling from the wiring closet to the network nodes are: unstructured, structured, or wireless.

An unstructured alternative implies that connecting computers is an afterthought to construction of a building or office space. The transmission medium lies along whatever path possible to network nodes. It avoids significant intrusion into the building construction. Unstructured cabling winds around obstacles, gets coiled behind furniture, and lies out in the open. Some typical unstructured cabling schemes hang from ceilings, are shoved under carpet or left lying openly or taped to the floor.

The simplest and typically least expensive cabling choice for an unstructured alternative is the 50 Ohm coaxial cable used for 10Base2 wiring (RG-58). The disadvantages of this medium are several. As a bus topology, 10Base2 cable is easy to connect but not easy to *properly* connect and maintain. BNC connectors tend to come

loose at splices and NICs when the cable is moved. Loose connectors introduce additional transmission losses to the medium. Isolating these problems without a cabling diagram is time-consuming--technicians must trace the coaxial cable that meanders about the building connecting nodes on the subnetwork. Another problem introduced by the ease of connection is that network managers and users themselves may easily exceed standards for topology-specified segment length and number nodes. More users gain connectivity through the network without knowing the effect of their actions on network performance.

Coaxial cable used for 10Base5 installations is more expensive but more sturdy than 10Base2 cabling. The number nodes allowed on a segment is over three times that allowed for 10Base2. Disadvantages include the necessity and cost of external transceivers for nodes and the potential for damage to the cable core at the connection of each transceiver. Isolating problems is difficult for the same reasons as for 10Base2 coaxial cable.

Unstructured cabling schemes also use multi-mode optical fiber from the wiring closet to network nodes. Optical fiber's main advantages are greater distance allowed between the wiring closet and the node, greater data-carrying capacity, and its improved signal quality and immunity to electromagnetic interference and electronic eavesdropping. The disadvantages are the cost of the fiber and, especially, the associated connectors on NICs and distribution devices. Fiber cannot be physically tapped into like coaxial cable without interrupting the transmission. This characteristic enhances security

but makes the transmission medium much less adaptable to change. Because of the cost disadvantages compared to copper-based solutions and lack of adaptability to network configuration change, network management should not consider optical fiber for the majority of unstructured cabling schemes.

A structured approach implies the use of a complete solution for wiring all communication devices including telephones and computer network devices. This solution should use defined cabling standards that include design, layout and logic as key factors [Newton]. The main components of a structured wiring system are:

- Drop cable--the cable that runs from the computer to a network outlet
- Cable run--the cable that runs from the outlet to the wiring closet
- Patch panel within the wiring closet
- Distribution device within the wiring closet (such as a concentrator).

The components of concern for this category of cabling are the drop cable and the cable run. The patch panel and the distribution device were discussed in the previous section. A typical structured cabling scheme runs in cableways, connects devices to the network via patch cables and wall-mounted jacks, and is routed through hollow walls and false ceilings. A structured cabling system has advantages over unstructured systems in five ways, according to Anixter, a leading supplier of structured wiring systems [Newton]:

- It eases network segmentation--the job of dividing the network into pieces to isolate and minimize traffic--and thus congestion.
- It ensures that proper physical requirements are met, such as distance, capacitance, and attenuation specifications.
- It means adds, moves, and changes are easy to make without expensive and cumbersome rewiring, thus a more scalable and flexible physical topology.
- It radically eases problem detection and isolation.
- It allows for intelligent, easy and computerized tracking and documentation.

This configuration uses nearly any data communications transmission medium. Unshielded twisted pair (UTP) in any category is a reasonable alternative because of ease of installation and performance in star topologies with segments of less than 100 meters. Presently, Category 5 UTP is the most scalable from current to foreseeable future network needs [Marks]. It can support at least 100 Mbps using current technologies. Tests promise a boost in capacity to over 600 Mbps over distances of less than 100 meters in the near future. The disadvantage of a structured cabling scheme is its greater installation cost due to additional labor charges.

The third approach involves the use of wireless connectivity to reach from the wiring closet to the network nodes. Chapter III discusses wireless technologies in some detail. [Stallings 2], [Freeman], and others discuss the advantages and disadvantages to each of the three main types (infrared, microwave, spread spectrum). Wireless schemes are not constrained by cable location. They are limited by mutual antenna visibility between devices, free space propagation losses, and/or electromagnetic interference. Wireless options may be appropriate when the costs of structured cabling

schemes are prohibitive because of the existing building structure. At NPS, these options should be considered in areas where asbestos-filled walls prevent a low-cost copper or optical fiber cabling solution, such as Hermann Hall.

2. Network Connectivity

This portion of the network architecture consists of distribution devices on the subnetworks. Subnetworks are those part of the campus network logically and/or physically separated from other segments of the network. Network-layer addressing schemes achieve logical separation among subnetworks. Individual cabling schemes physically separate the subnetworks by cable route, floor location, or building.

a. Network Interface Card

The type of network interface card (NIC) used by a machine on a network node depends on the topology of the subnetwork and the machine on the node. NICs are available for all types of topologies and the standard cabling that supports each specific topology. A change in subnetwork logical (e.g., 10Base5 to FDDI) or physical (e.g., 10Base2 physical bus to 10Base-T physical star) topology frequently demands a change of every NIC on the subnetwork. There are exceptions to this limitation that allow some transition without replacing NICs. Some NIC manufacturers provide terminal connections on the cards to a combination of the cabling that support the topology (e.g. a combination of 10Base-T RJ-45, 10Base2 BNC, and 10Base5 AUI connectors on one ethernet card). Manufacturers now offer NICs with capabilities built-in that allow reconfiguration from

lower-bandwidth ethernet network connections to one of the emerging 100 Mbps channel capacity standards without removing the card. [Schnaidt 1]

b. Repeater

Repeaters allow further segmentation of a subnetwork beyond that which is provided by other distribution devices. Repeaters also extend the physical length of cable allowed to be used in a shared-media topology. The number of repeaters used on a subnetwork is limited by the topology. For example, in IEEE 802.3 networks, the maximum number repeaters on a segment is four and the maximum length of the cable is 2.5 km for 10Base5 and 925 m for 10Base2. For IEEE 802.5 token ring networks, the maximum number of repeaters is 250.

c. Concentrator Hub

A concentrator hub allows wiring in a IEEE 802.3 10Base-T subnetwork to be arranged in a physical star topology similar to common telephone wiring closet configurations. This configuration allows easier isolation of a troublesome node on the subnetwork than found with 10Base5 or 10Base2 configurations. The wiring of the subnetwork remains logically connected as a bus in keeping with the IEEE 802.3 standard. The IEEE standard specifies each port on the hub is limited to two devices sharing a common UTP cable. The maximum length of a UTP cable extending from the hub is 100 meters. Configurations typically use only 90 meters to allow for drop cables, patch cables, and connector transmission losses. There is no limit specified as to how many ports can

be connected to a hub but adding more active nodes to a hub on a CSMA/CD subnetwork increases the probability of collisions and reduces throughput.

d. Switching Hub

Switching hubs build on the advantages found in using concentrators discussed above. [Olsen] and [Schnaidt 2] describes various switching hubs. All ethernet switching hubs improve upon standard 10 Mbps ethernet in higher data throughput, better administrative security, and centralized management of complex internetworks.

Higher-layer software can dynamically manage bandwidth allocation, transmission reliability, and traffic prioritization. Most switching hubs provide a high-bandwidth backplane inside the device. In these high-end distribution devices, the aggregate throughput of the internal backplane of the device limits switching between connected LANs. Typical backplane throughput exceeds 150 Mbps.

[Durr] describes an extension of this alternative, the "hub-of hubs" network. As the subnetwork grows, network managers stack and cascade switching hubs to increase the number of nodes serviced, while providing segmentation without using a router. The disadvantage of this configuration is that it assumes a homogeneous MAC-layer topology among the LAN segments for non-routing hubs.

3. Internetwork Connectivity

This portion of the network architecture consists of distribution devices that interconnect the subnetwork segments described above and that provide connection to wide area networks.

a. Shared Backbone

In this configuration, LANs share a backbone with bridges and transceivers for internetwork connections. Figure 11 shows a basic configuration that provides connectivity among physically separated LANs. A shared-backbone configuration is inexpensive and simple. A shared-backbone configuration does not provide for network addressing above the MAC layer, nor is it tolerant to faults created by a transceiver at a node on the backbone.

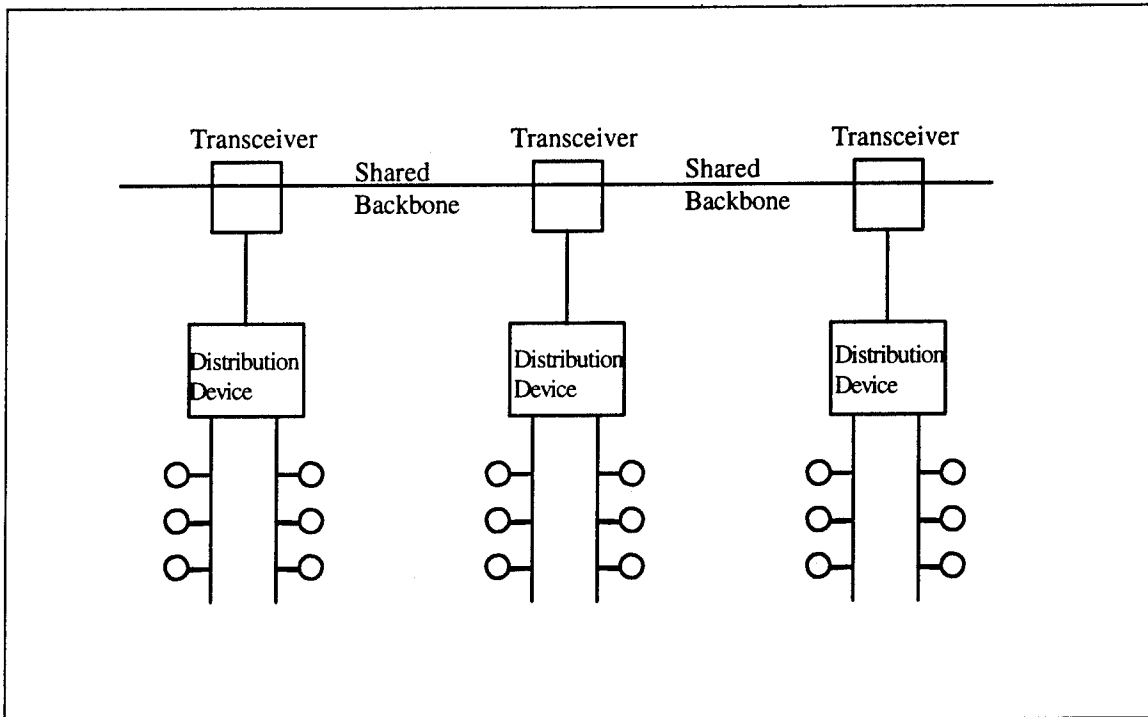


Figure 11. Shared backbone configuration.

b. Collapsed Backbone

In this configuration, LANs are segmented and interconnected using routers and/or bridges. As shown in Figure 12, the backplane of a router acts as the backbone of the network, thus this configuration is typically referred as a "collapsed backbone." This alternative takes advantage of all the properties of the router to segment LANs and to connect heterogeneous LANs. It also has the disadvantages of a router--expensive and complicated. The number of ports in a router and the aggregate capacity of its backplane limit the collapsed-backbone configuration.

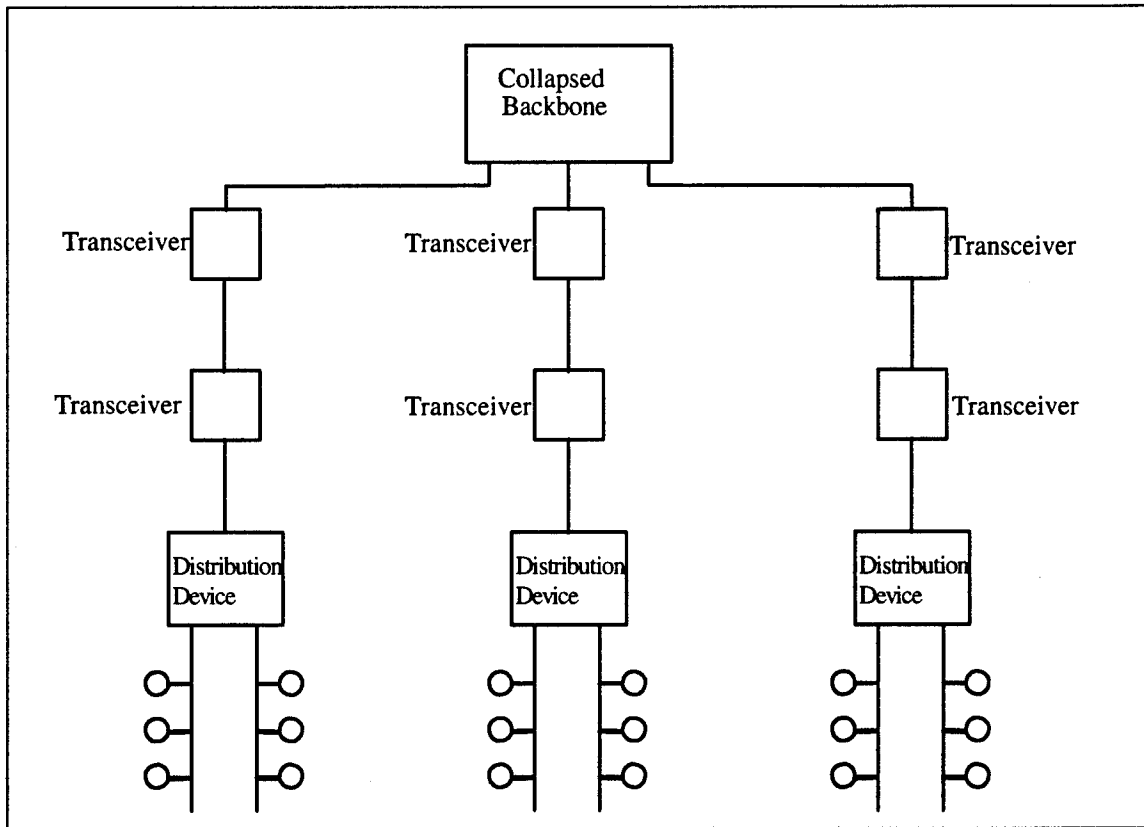


Figure 12. Collapsed backbone.

c. Backbone between Routers

As the network grows and fills the router ports, a logical extension connects routers together. As shown in Figure 13, additional routers, connected via a backbone, allow the addition of LAN segments. Each segment retains its full channel capacity. This configuration provides for fault tolerance at the routers. If a router malfunctions, its attached segments can be redirected to ports on another router. The disadvantage is increased complexity within each router.

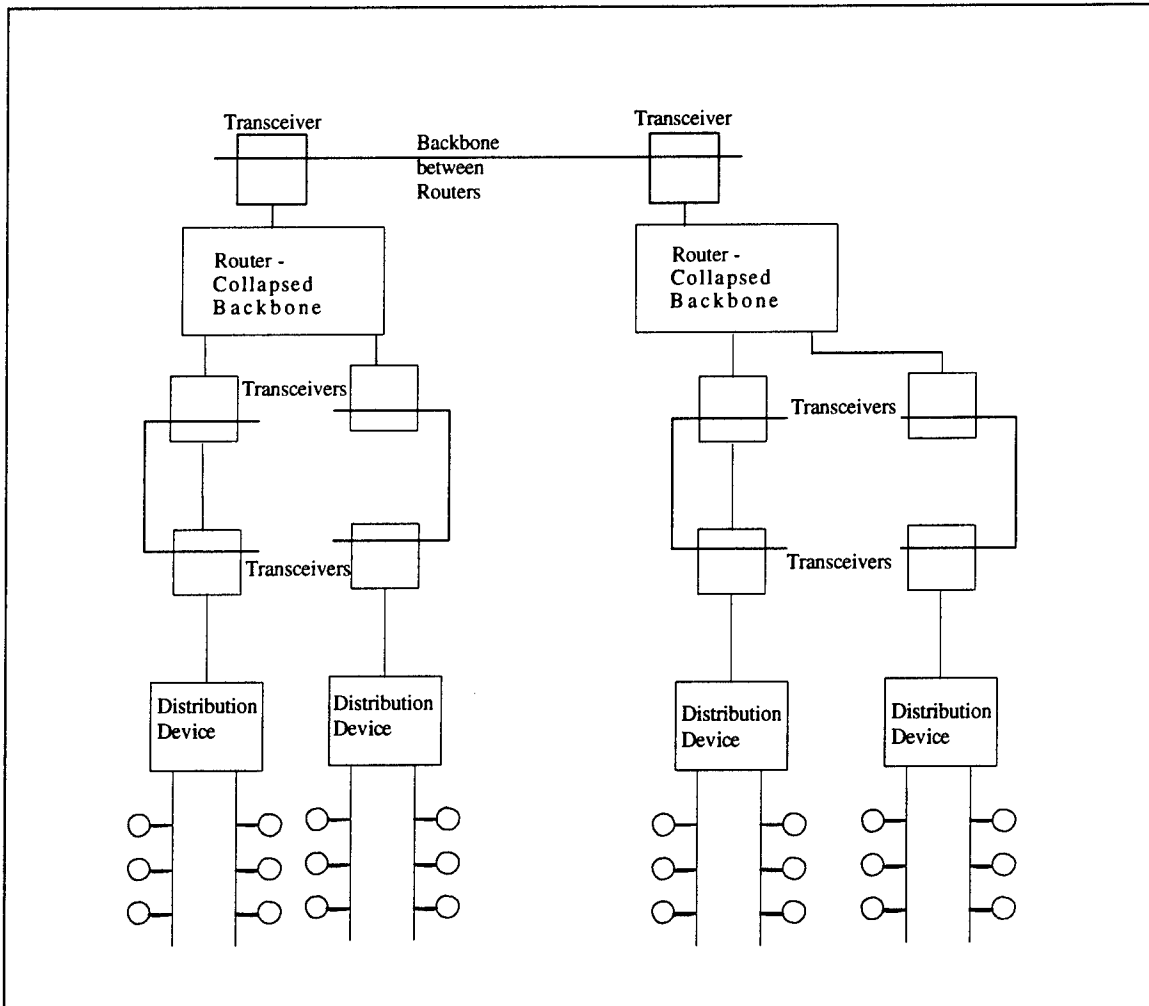


Figure 13. Backbone between routers.

B. ARCHITECTURE TRENDS

Computer networking-related periodicals are significant sources of information on trends in network architecture. Rather than discussing the technical merits of the each of the alternatives, as presented in more professional journals, the periodicals focus on what alternatives are available and being sold, installed, and used. This section summarizes network architecture trends as discussed in periodicals such as *BYTE*, *Cabling Business Magazine*, *Communications Week*, *Computer Shopper*, *Computerworld*, *Government Computer News*, *Infoworld*, *LAN Computing*, *LAN Times*, *Network Computing*, *Network World*, and *Networking Management*.

1. Meeting Immediate Needs

Shared-media alternatives, with low (10 Mbps or less) channel capacity to each node, meet the immediate needs of many networks. They have advantages over newer technologies of proven performance and lower initial costs in cabling, NICs, and distribution devices. Their biggest disadvantages are the risk of not being able to meet the future needs of their users and the probability of costly replacement of cable, NICs, and distribution devices in the near future.

2. Scalability and Integration

Another network architecture trend focuses on scalability and integration. User demands and industry manufacturers, seeking market niches, drive network architecture technologies towards divergent higher-bandwidth topologies such as ATM, 100VG-AnyLAN, and FDDI [Nolle]. At the same time, the need for interconnection

among computer networks still exists. Network architecture designs consider scalability as an important factor especially after basic network connectivity needs have been met. Networks are capable of supporting current user needs and of allowing quick, easy, and inexpensive upgrades in distribution devices and network capacity. The upgrades within existing distribution devices anticipate the need to support future demands for higher bandwidth applications, greater flexibility, simplified network management, and increased distributed processing. Switching technologies and structured cabling provide steps towards a scalable switch-based network architecture. Their drawbacks are high initial costs for distribution devices and costs of any necessary recabling.

Cabling, installed with the future of the network and the organization in mind, allows for increases in channel capacity, additional users, and recovery from link or node failure. For example, when pulling multi-mode optical fiber, installers pull additional fiber pairs and leave them "dark." Likewise, network builders pull and terminate additional segments of UTP using EIA/TIA-568 standards when star-topology wiring is being installed. These additional cables allow the network to grow, accommodate higher bandwidth requirements, and allow different encoding schemes from emerging technologies during the expected life of the cabling plant.

3. Backbones

As the demand for higher-bandwidth connectivity among LANs increases, LANs need interconnection through higher-bandwidth backbones and internetwork distribution devices. FDDI is the leading installed alternative meeting this need. This technology has

the advantage of more than a decade of experience in providing shared-media 100 Mbps connectivity among devices. The future of higher-bandwidth backbones appears to be switch-based alternatives, such as ATM, that use star-topology wiring. Leading-edge sites are installing ATM switches capable of filling the role of higher-bandwidth collapsed backbones. Other sites are waiting for reasons including clarification of ATM standards, further proof of the technology in operational environments, introduction of ATM applications, and reduction in prices of ATM switches and interfaces to prices comparable with established technologies. [Masud]

4. Virtual LANs

A trend in network architecture brought about by business organizational changes is virtual LANs. This network management alternative is based in switching technologies and star-topology wiring. A virtual LAN logically groups nodes into networks, which are independent of their physical wiring. Workgroup applications and horizontal organizational hierarchies drive this trend. As the trend of flatter organizations continues, users require connection in logical workgroups regardless of their physical locations in the organization. The challenge of the network providing sufficient throughput among the workgroup nodes without sacrificing performance of the network for other users.

[Schnaidt 2]

The network administrator can assign workgroups by network layer address or by port number rather than by MAC address. For example, IP subnetting can be used to provide a virtual LAN. [Schnaidt 2]

High-end ethernet switches provide the means of virtual networking with existing technology at ethernet speeds. ATM switches are beginning to support virtual LANs at higher speeds. [Durr]

C. NPS APPLICABILITY

NPS could benefit from a reasonable combination of cabling, network connectivity and internetwork connectivity presented in this chapter. Network administrators could gauge network benefits using measures such as those discussed in the previous chapter to compare network architecture alternatives. Quantifiable measures of efficiency alone may not justify changes to all or part of the existing architecture. Measures of effectiveness may be the deciding factors in any significant changes to the network in order to meet the anticipated needs or satisfaction of the users. [Spada]

New network installations should use Category 5 UTP in a star-topology structured-wiring configuration (e.g., EIA/TIA-568 standard) and switch-based technologies whenever possible on subnetworks from wiring closets to network nodes. Cabling within wiring closets should include patch panels and patch cords to the distribution devices for maximum flexibility with acceptable losses compared to the use of punch-down blocks. Cabling between buildings and between wiring closets should be multi-mode optical fiber to provide higher bandwidth over longer distances than copper-based solutions. Redundant paths between distribution devices and buildings should be installed whenever possible in a physical mesh configuration to allow rapid recovery from damage to the primary transmission medium. Distribution devices should

be scalable, survivable, and capable of supporting all existing network traffic. If distribution devices cannot meet future needs of the network, then network technicians should be able to replace them with minimal disruption to the rest of the network. Likewise, if a node demands higher bandwidth, network technicians should only have to replace the NIC and reconfigure the distribution device in the wiring closet with minimal effect on the rest of the network.

A computer network architecture alternative that meets the goal of maximizing scalability, sustainability, survivability, and security, is EIA/TIA-568-compliant structured wiring and switch-based distribution devices. Deliberate transition to this alternative does not come without cost. The cost of campus-wide cable installation and investment in new or upgraded distribution devices may drive any such changes out of reach. These costs may also bring the network to the attention of financial planners for the school. The competitive position of network transition funding may improve when senior organizational management understand the benefits of a switch-based architecture and when management is committed to bringing this alternative to NPS. Transition strategies are discussed in the subsequent chapter.

VII. REASONABLE TRANSITION STRATEGIES

This chapter discusses reasonable strategies for transition of the network architecture from its current configuration to the switch-based architecture presented in the previous chapter. The chapter discusses several strategies for migration and then presents some examples of how other networks have been changed or are being changed.

Each of the strategies presented considers meeting the present and future needs of the network users with the ultimate goal of a switch-based architecture. Each strategy is driven not necessarily by availability of technology, but by how best to match users' needs with technology in support of the overall goals and vision of the organization.

[Feldman] ties strategic issues associated with any network architecture to economics vs. productivity. He suggests network management ponder the following questions.

- When is network traffic load going to require migrating from existing solutions?
- If ATM is the "next generation" of technology will the network be "left behind" waiting for prices to drop?
- On glutted LANs, can a less-expensive solution be used to buy time? If so, how long?
- Are the routers compatible with ATM or will they have to be changed, too?
- What is the bottom line the organization is willing to invest in new technology?

Another question could be added: "If ATM doesn't turn out to be all it is promised to be in the LAN environment, can we minimize our losses?" [Schoenstadt 2]

A. STRATEGIES

This section discusses strategies primarily derived from [Serjak] and modified to address some of the specific network architecture on the NPS campus.

1. Higher-Bandwidth Switch-Based Architecture

This "forklift" strategy recognizes the inadequacy of the current network architecture in meeting the current needs of the users, and in providing scalability, sustainability, security, and survivability. It recognizes that a complete rebuild of the network from the bottom up is necessary to meet all the demands of the network. All existing network connectivity and internetwork connectivity devices and cabling are "hailed out by a forklift." Another forklift load brings in the new infrastructure. This strategy links the backbone, wiring closets, and all connections to servers and other higher-demand devices together via higher-bandwidth switches in a mesh configuration. Existing lower-bandwidth subnetworks still play a role, but only on the periphery through interfaces on the higher-bandwidth wiring closet switches. Eventually, the subnetworks are rewired for incorporation into the switch-based architecture. ATM switches are appropriate for use as the higher-bandwidth switches in this strategy, as shown in Figure 14.

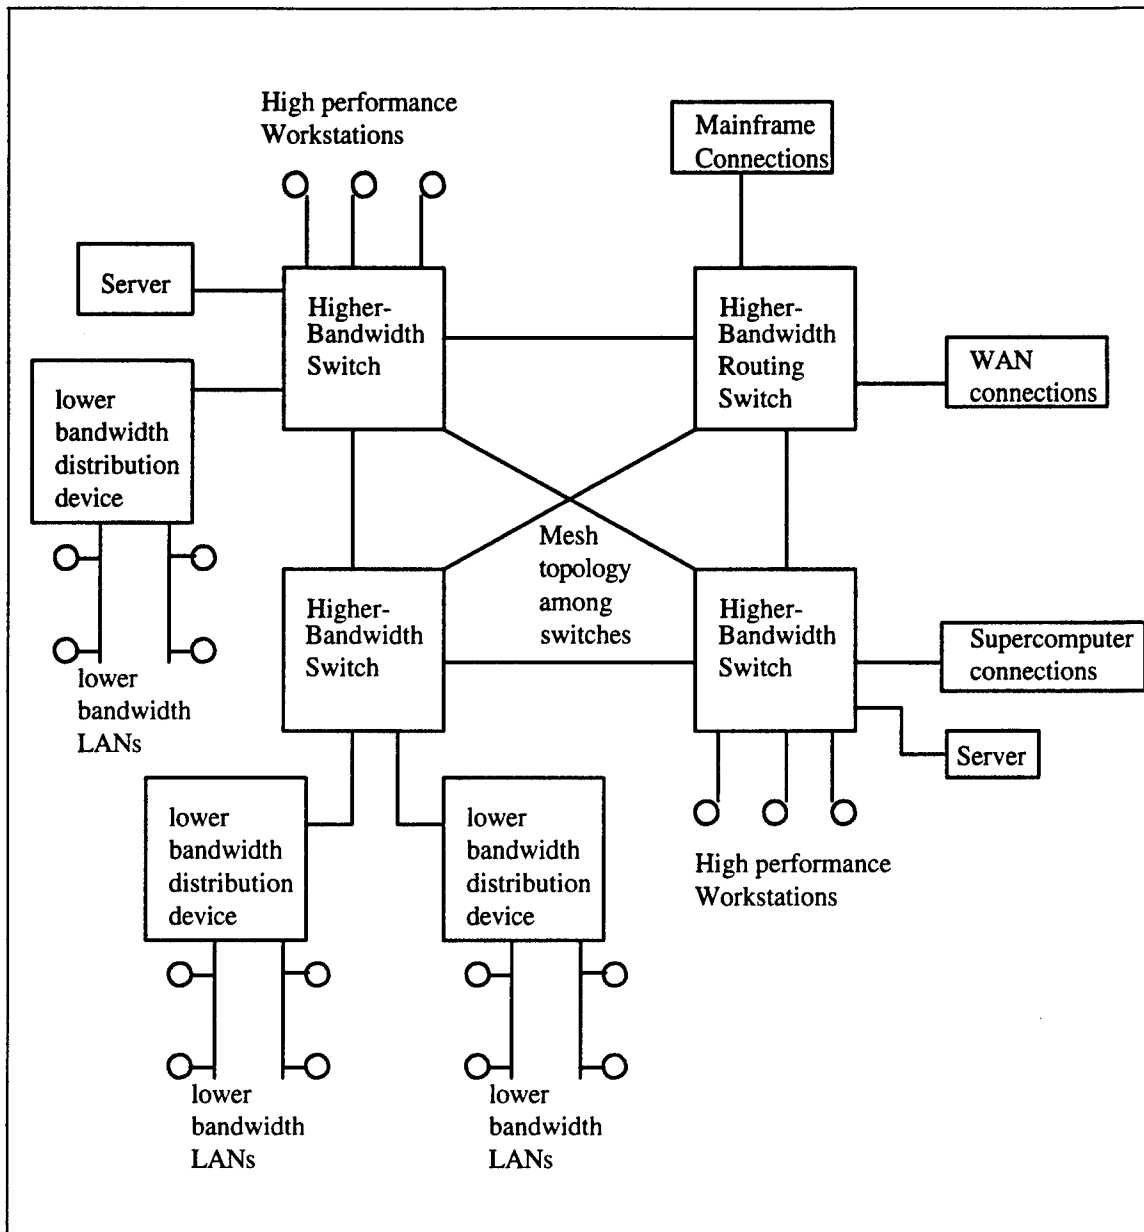


Figure 14. Higher-bandwidth switch architecture.

This strategy requires much financial commitment by an organization towards improvements in services provided by the network. In this "forklift" strategy, where a significant portion of the architecture is replaced *en masse*, the network requires complete redesign. This redesign brings together representative users, network planners, and network maintainers to maximize the use of in-place cabling and equipment while making the most of newer higher-bandwidth switch-based technologies.

The most difficult challenge in this strategy is rewiring of subnetworks that are shared-medium bus or ring topologies that are connected to a collapsed backbone by coaxial cabling. These subnetworks require conversion to star topologies that support a single node per segment, which is connected to a centralized distribution device. A distribution device in this configuration must then connect to other distribution devices by a transmission medium such as multi-mode optical fiber, which has high capacity over longer distance .

2. Evolution

Network builders may have difficulty acquiring adequate funding and organizational commitment necessary for a forklift transition to a higher-bandwidth switch-based topology. They need a more incremental approach, portrayed in Figure 15, when the essentials for a forklift transition are not available and yet the users still demand improved network performance. The (lack of) organizational commitment and funding in this case preserves the life of the existing network architecture as long as possible by

supplanting the infrastructure with switch-based solutions in one of the following evolutionary strategies.

The major transitions in the early part of one evolutionary strategy include micro-segmentation of each LAN to reduce the number of users and bandwidth needed on each segment. Network performance is then improved by increasing the performance of the backbone and finally gradually upgrading the distribution devices. Separating higher bandwidth-demanding applications and services, such as file servers, onto their own segments reduces the bandwidth needs of the other segments. Segmenting along organizational boundaries, e.g. work groups or divisions, also reduces the bandwidth demands on other segments. As this "micro-segmentation" extends to its limits, a star topology forms with dedicated capacity, through the repeater or switch, to each node on a subnetwork. Scalable switches then replace ethernet repeaters. These switches could contain replaceable modules for greater flexibility and improved fault tolerance. An investment in structured cabling systems aids this star topology by concentrating each segment into centrally-located wiring closets, which contain the new switches.

As the demand for sharing of information among segments and switches increases, the backbone that connects them becomes congested. In the case of NPS, the subnetworks are themselves backbones which are further concentrated into the collapsed backbone of the routers. Network administrators could upgrade these subnetwork backbones to FDDI or other 100 Mbps alternatives as long as the routers and the switches on the subnetworks each have the capability to integrate into an FDDI configuration. The

current NPS routers (Cisco AGS+) are limited to four or less FDDI interfaces, which operate at full capacity while the remaining 14 interfaces are limited to lower-capacity LAN topologies. [Howard], [Cisco 1]

If upgrading a subnetwork requires installing new multi-mode optical fiber, then installation of additional "dark" fiber should be considered to allow for future upgrades. As the aggregate data rate on each subnetwork begins to exceed the capacity of the subnetwork backbone, higher-bandwidth switching modules can replace the switch modules in the wiring closet and in the router. These upgraded modules could use technology such as found in ATM modules which can transmit data over the multi-mode optical fiber that was installed for the FDDI network. (As of August 1994, Cisco Systems, Inc. is committed to delivering an ATM interface for the AGS+ within a year [Howard].) Similarly, as aggregate network traffic overburdens the collapsed backbone on the backplane of the routers, scalable high-speed switches can replace the routers using switch-based technology such as ATM.

At this point, network management should consider the trade-off between installing an intermediate switch at each remote building and installing additional pairs optical fiber between the centrally-located collapsed-backbone switch and the remote buildings. If the cost of installing additional optical fiber (one pair for each wiring closet plus "dark" fiber) is greater than that of an additional intermediate switch, then install switches in the remote buildings. These switches can further distribute to the switches in the wiring closets. If sufficient "dark" fiber is already in place between the router and the

remote building then each wiring closet switch can link directly to the collapsed backbone switch with the addition of short fiber patch cables.

After the collapsed backbone has been replaced, newly incorporated higher-bandwidth wide area network (WAN) cabling and switch connections can take advantage of emerging high-speed WAN technologies and distribution of high-bandwidth applications to sites off the NPS campus.

Network management can now consider replacing or supplementing the switches in the wiring closets with ATM switches. The ATM switches are capable of providing dedicated higher-bandwidth service to the nodes that are not restricted by internal limitations of the desktop system. The ATM switch in the wiring closet could also support the remaining lower-bandwidth nodes through a port to the previously installed switch. The need for the lower-bandwidth switch diminishes as more nodes transition to ATM interfaces. This lower-bandwidth switch would remain needed as long as individual nodes cannot justify dedicated access to a port on the ATM switch.

The advantage this strategy is it allows network managers to come into switch-based architecture gradually while learning the best utilization for switching technology in their environment on small scales at the segment level before integrating it campus-wide. It may also solve isolated bottleneck problems without additional investment in higher-bandwidth solutions.

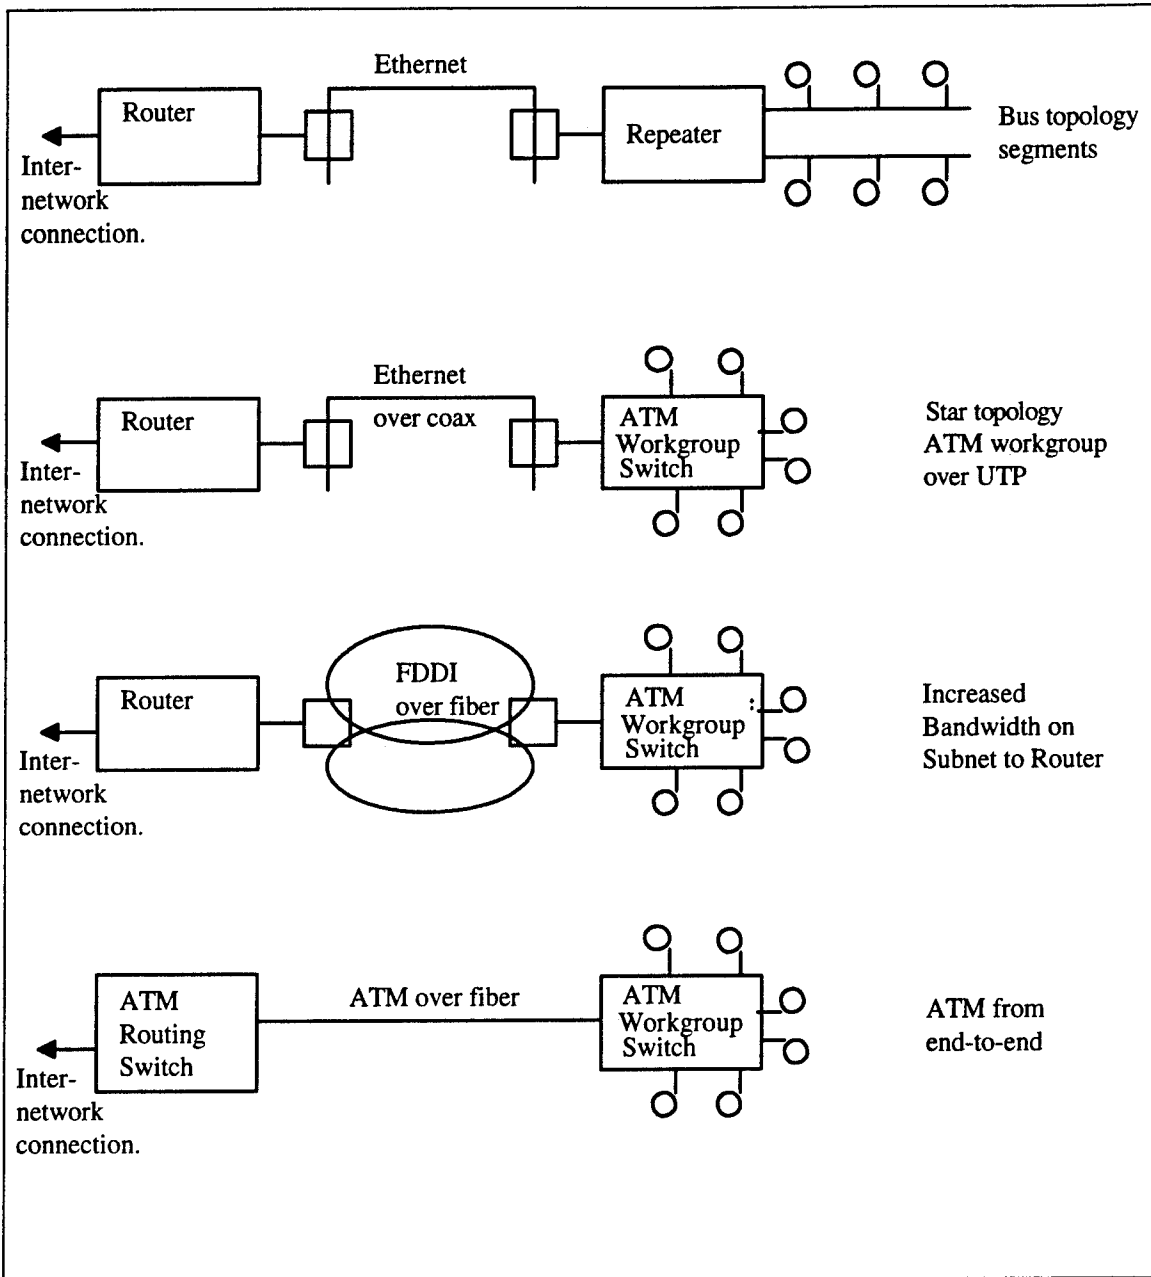


Figure 15. An evolutionary transition to ATM.

A second evolutionary process is nearly the opposite approach from the first evolutionary strategy. "A productive strategy is to apply ATM first to the weakest area--campus and building backbones, second to the wide area network, and third to the client interface [Serjak]."

The first step in this process is replacing the centralized routers with scalable high-speed routing switches. This alleviates the backbone problem of aggregate channel capacity, because the internal switch fabric of ATM-based switches are capable of a much higher sustained throughput than a conventional software-based router [Cisco 2].

The second step in this strategy is upgrading wide area network (WAN) connections. Higher-bandwidth connections through common carriers are expensive, so network management must consider using this connection to its fullest, most economical potential, lest the advantages of having a higher-bandwidth WAN connection be lost. If network management cannot justify the higher-bandwidth connection, then they should consider delaying the upgrade of WAN connections. In the NPS network environment, upgraded WAN connections may be necessary to meet the needs of applications such as distance learning, videoteleconferencing, and more efficient sharing of large data sets with external computing resources.

The final step in this strategy involves gradually spreading higher-bandwidth switches such as ATM across the campus network starting at the collapsed backbone switch and step-by-step upgrading of the network until workstations at each node are brought into a higher-bandwidth switch-based network architecture. Replaced

collapsed-backbone routers could be moved out to the subnetworks most in need of greater capacity through segmentation. (In this case, the router acts as an intermediate distribution device at a network entry point into a building, similar to the intermediate switch discussed in the first evolutionary strategy.) That same router could then be moved to the wiring closet when it no longer meets the network demands as an intermediate switching device. This preserves the end-node investment in lower-bandwidth technologies for as long as necessary, improves overall network performance, and does not force higher-bandwidth switching down to users that do not need it or cannot afford it. It does, however, remove the highly complex legacy router away from centralized management and expertise--one of the reasons for a collapsed backbone configuration in the first place.

The advantage of this second evolutionary process is that it provides a high-bandwidth solution to the wiring closet, while delaying the investment in cabling from the wiring closet to the nodes that is required to moved the entire network to a switch-based architecture. This strategy retains the investment in lower-bandwidth technologies by continually pushing the distribution devices towards the periphery--closer to the desktop, without changing the nodes. This strategy also takes advantage of the maturity and pricing of switch-based technology user-network interfaces. As time progresses, prices come down and standards improve and stabilize.

This second evolutionary alternative's disadvantage is the risk taken in replacing a familiar, marginally-performing technology with an unknown and untested

high-performing technology. While the bugs are worked out on the integration of ATM on the collapsed backbone, users and applications that require routing through the backbone are in danger of not having that service available.

Evolutionary strategies in general allow the migration to take place at a slower rate with less up-front expenses than with the "forklift" strategy. Each step can be deliberate, well-defined, and involve a small work force that builds on the lessons learned from changing of each segment.

3. Parallel Backbones

The parallel-backbone strategy solves the problem of meeting the high bandwidth needs of a few user and subnetworks, while retaining the current network architecture. This strategy is appropriate if there are highly-polarized performance demands on the network. Some subnetworks and users require high performance because of such applications as rapid access to data bases, on-demand videoteleconferencing, and rapid transfer of large data files. The remaining subnetworks and users continue using the existing architecture because their network use is limited to primarily lower-bandwidth demands such as office automation applications and basic electronic mail functions.

This strategy implements the higher-bandwidth solution from the WAN connection to the desktop or supercomputer for a few applications and users. It employs an ATM collapsed-backbone switch and links to the lower-bandwidth backbone (Figure 16).

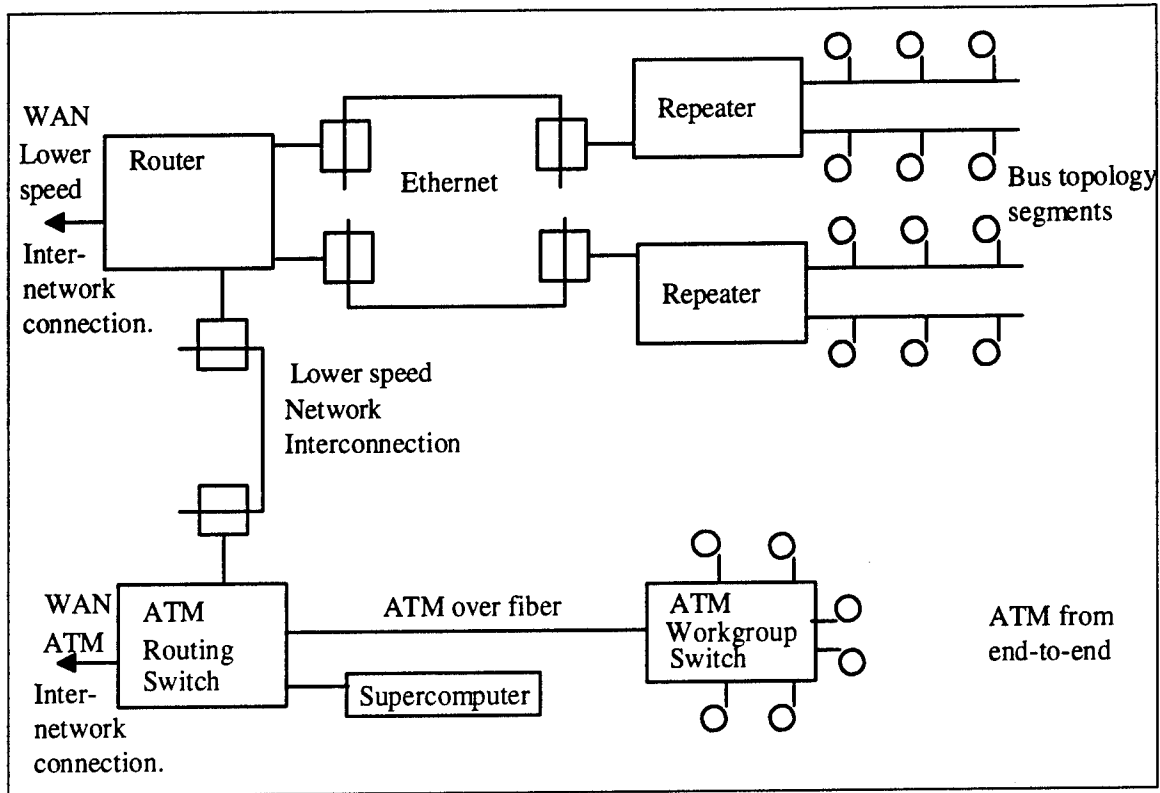


Figure 16. Parallel backbones strategy.

This strategy serves two purposes--it relieves the congestion on the current network and it provides the bandwidth where it is needed most. [Serjak] believes this strategy is inevitable for organizations that need to implement an end-to-end ATM solution to support selected bandwidth-hungry applications. [Bolles], [Masud], and others in [MacAskill 4] believe this strategy is too expensive to maintain and recommend evolutionary strategies using combinations of less expensive and more mature technological solutions.

The disadvantages of this strategy include added installation costs without any reuse of the existing infrastructure and added complexities and responsibilities for network

management . It is also unlikely that ATM switches for this strategy will have the interfaces necessary to handle a migration of the rest of the network onto them without additional expense.

Despite its disadvantages, this strategy is appropriate for NPS. There are applications awaiting the use of ATM-like performance on the campus within a few departments. These applications include videoteleconferencing and distance learning research, meteorological and oceanographic data transfers, and interactive simulations. Many other departments and subnetworks require only relatively minor upgrades such as additional segmentation, improved file servers, and small recabling projects to improve their performance. The parallel backbone strategy allows NPS network administration the opportunity to become comfortable with the technology without totally committing the organization to an immature LAN technology. The rest of the campus can continue along an evolutionary path towards a switch-based architecture that may integrate with the ATM backbone sometime in the future when users require high-bandwidth solutions.

B. EXAMPLES

1. Fairfield University

Fairfield University is a small religion and liberal arts academic institution in Fairfield, Connecticut. Applications on its computer network are primarily administrative and require only lower-bandwidth connections. In 1993, the network consisted of three hundred "dumb" terminals connected to a Digital Equipment Corporation VAX and 400 PCs--many of them stand-alone. Ethernet networks did not permeate the campus. Their

infrastructure and networking capabilities had not been upgraded in several years.

[Mulqueen]

Fairfield University had a growing data transfer requirements including PC connections to all dormitory rooms and multimedia applications being developed by the university staff. Faculty members were developing interactive, multimedia teaching aids that supplemented a newly installed community antenna television (CATV) network that provided foreign language programming and serviced four cable television stations run by the university. Research initiatives were bogged down by slow network data transfer rates.

Fairfield's new network architecture will consist of a multi-mode optical fiber backbone among all the buildings and multi-mode optical fiber from switching hubs to nodes in a star-topology configuration. When completed in late 1994, the new network architecture will integrate all voice, data, and video resources--including CATV. The university will soon require that students have PCs in their dormitory rooms. A minimum of 1,100 rooms will have a PC connection to the network. There will be 3,000 to 3,500 connections to the network, when the library, administrative, and faculty offices are included. Fairfield University's director of communications and technology services feels that with star topology and multi-mode optical fiber cabling, he will have the bandwidth to accommodate more than just traditional data and the network will be well-positioned for providing network services demanding the bandwidth available with ATM. Already, with only early phases of the project complete, reduction in data transfer times have aided

research initiatives. One Fairfield University researcher noted reduced transfer times for data files, particularly video, from "hours to minutes and minutes to seconds." The network management's goal is to make available all data types--voice, data, and video resources, to every user as standard resources on the backbone.

Fairfield University is using a forklift strategy costing approximately one million dollars. It is taking this opportunity to rip out the existing infrastructure and become a showcase of multimedia networking. Network administration compared the costs of fiber connections with the cost of more traditional copper wire connections. Although the cost of the installation was higher, they felt the difference was insignificant when compared with the increased potential for high bandwidth applications with no electronic interference. They also feel that they will not have to recable every few years, as they felt they would have if they had committed to doing the project with copper cable options. Prestige has played a factor, too. Network managers from other academic institutions and from corporations have inundated Fairfield University with requests for information on how and what is being done with the optical fiber installation.

Fairfield chose a reasonable strategy for areas where network services are severely lacking or non-existent. Fairfield's existing network consisted of outdated technologies that could easily be scrapped. This is an approach that could be used at NPS when offices or departments move to new or different facilities. It is not as critical, in late 1994, for network planners to choose optical fiber as the transmission medium as it was in early 1993, when Fairfield University committed to multi-mode optical fiber. Category 5 UTP

and STP now provide bandwidth comparable to that of multi-mode optical fiber over reasonable distances such as from the hubs to the dormitory rooms or desktops. Optical fiber remains the reasonable choice for the backbone. When moving departing offices and equipment, network technicians could remove any old network cabling and network devices and install a new architecture before the new occupants arrive. Network management could coordinate this effort prior to, and during, the move-in with the arriving occupants. This coordination would ensure maximum flexibility and service once the move-in is complete.

2. Johns Hopkins School of Medicine/University of Maryland Medical Center

Johns Hopkins University (JHU) School of Medicine and the University of Maryland Medical Center (UMMC) are involved in medical research. One particular project focuses on improving efficiency in treating children born with abnormal skulls. Computer applications draw from a data base to predict the effect normal growth will have after surgery is performed to correct skull abnormalities. The data base is a massive store of magnetic resonance imaging (MRI) and computed tomography (CT) scans for about 300 patients. Each patient file contains an average of 50M bytes. [MacAskill 2], [Klett]

The late 1993 network configuration was as shown in Figure 17. At JHU, the network consisted of high-end desktop workstations connected via an ethernet hub. An ethernet backbone connected the hub to a router. The router was connected to a WAN ATM switch over T-1(1.5 Mbps) using frame relay. The WAN ATM switch was

connected to another WAN ATM switch at UMMC over a leased T-1 line. At UMMC, the ATM switch was connected to another router over T-1. This router was then connected to imaging and scanning equipment over shared ethernet.

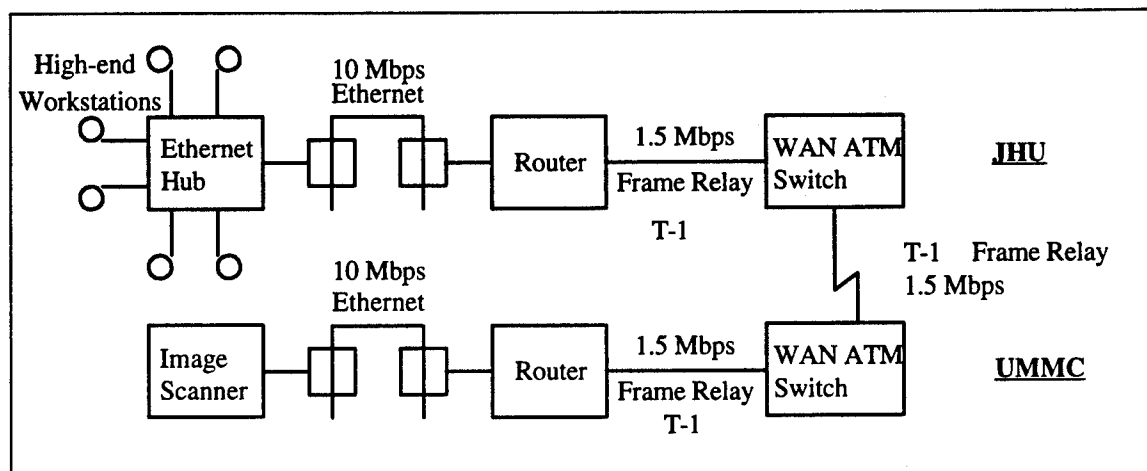


Figure 17. JHU/UMMC late 1993 configuration. After [MacAskill 2].

The need for instant geographic access to huge imaging and graphic files is the most apparent driving force in this example. Specifically, researchers need high bandwidth to access and share the large data base, and develop advanced modeling applications that will help surgeons predict future growth patterns of cranial defects. More researchers using the network to transfer images to develop computerized growth reference models. The long term goal of the organizations is to be a national archive, accessible to researchers around the country who will read and modify the data.

This project is progressing in an evolutionary manner. First, the JHU backbone between the hubs and the router was upgraded to FDDI. This relieved the immediate stress of transferring the data files across the backbone and established optical fiber as the

backbone medium. Second, the 10Base5 shared ethernet at JHU was converted to 10Base-T. The 10Base-T segments were then converted to switched ethernet by upgrading the 10Base-T repeater hub to a switching hub. This change provided dedicated 10 Mbps links to archive workstations and introduced the network to switch-based technologies. Within the next year, the links between the routers and the WAN ATM switches at both sites will be upgraded to FDDI while the WAN ATM link will be upgraded to T-3 (45 Mbps), as shown in Figure 18. Within the next five years, both sites will convert their backbones to ATM. The ultimate goal is ATM to the desktop especially for those areas that manipulate radiological images.




Figure 18. JHU/UMMC network configuration (projected for late 1994). After [MacAskill 2].

Several key issues are being addressed during this migration. Using switched ethernet has retained the investment in NICs and Category 5 UTP cabling on the JHU network. To achieve the goal of being a national archive, network developers are using and maintaining network standards in order to allow relatively easy access from other

medical and research facilities. Cost is not much of a factor in this case because the switched ethernet hubs, ATM equipment, funding, and training were all donated by network equipment vendors. The vendors are using this project as a trial network for their equipment.

The projected late 1994 network configuration, Figure 18, allows quicker transfer of the minimum 24G bytes of data that is moved daily in support of research efforts.

Already, visible improvements include accelerated research initiatives. For example, the network allows videoteleconferencing over the network between physicians during remote consultations while both parties view the same images.

3. Concurrent Technologies Corporation

Concurrent Technologies Corporation (CTC) is a nonprofit subsidiary of the University of Pittsburgh Trust that operates four National Centers of Excellence for the Department of Defense. CTC, located in Johnstown, Pennsylvania, specializes in metallurgy, environmental studies, factory automation, and manufacturing logistics research. The production side of the corporation runs distributed parallel computing applications that require the movement of large amounts of integrated data, graphics and images among engineering staff workstations. These applications were taxing the existing ethernet. [MacAskill 1]

CTC considered several measures when comparing network alternatives but justified their selection based on primarily on one measure. The measures considered were capacity-per-node, stability-of-standards, scalability, price/performance

(cost-per-Mbps-per-node). CTC's principal technical manager found the measures that were the most telling for CTC were scalability and cost-per-Mbps-per-node. Ethernet's cost-per-Mbps was \$50 and not scalable. FDDI's cost-per-Mbps-per-node was \$40 and not scalable. ATM's cost-per-Mbps-per-node was \$32 and scalable.

CTC's new network architecture consists of two networks, an ATM network for the production applications and an ethernet network for the administrative applications. Each ATM switch connects to a Cisco Systems, Inc. AGS+ router via a 10 Mbps link.

The ATM network, shown in Figure 19, consists of four ATM switches that deliver 155 Mbps directly to more than 50 high-end Unix workstations via multi-mode optical fiber connected between ports on the ATM switch and the NICs on the workstations. A switch is located in each of four buildings. The four switches are linked to each other via fiber at 155 Mbps in a mesh topology. The ATM network has significantly increased productivity and reduced time to market. According to CTC's principal technical manager, "Processes that took weeks to complete on the ethernet now only take a day or two." CTC engineers collaborate using desktop videoteleconferencing and a three-dimensional model simulator concurrently. These collaborations do not slow the network, thus other ATM network users continue on the network without noticeable degradation in performance. Users are satisfied with the performance of their applications, as made possible by the ATM network. The network continues evolving as standards for ATM are resolved and more users and switches are added to the network.

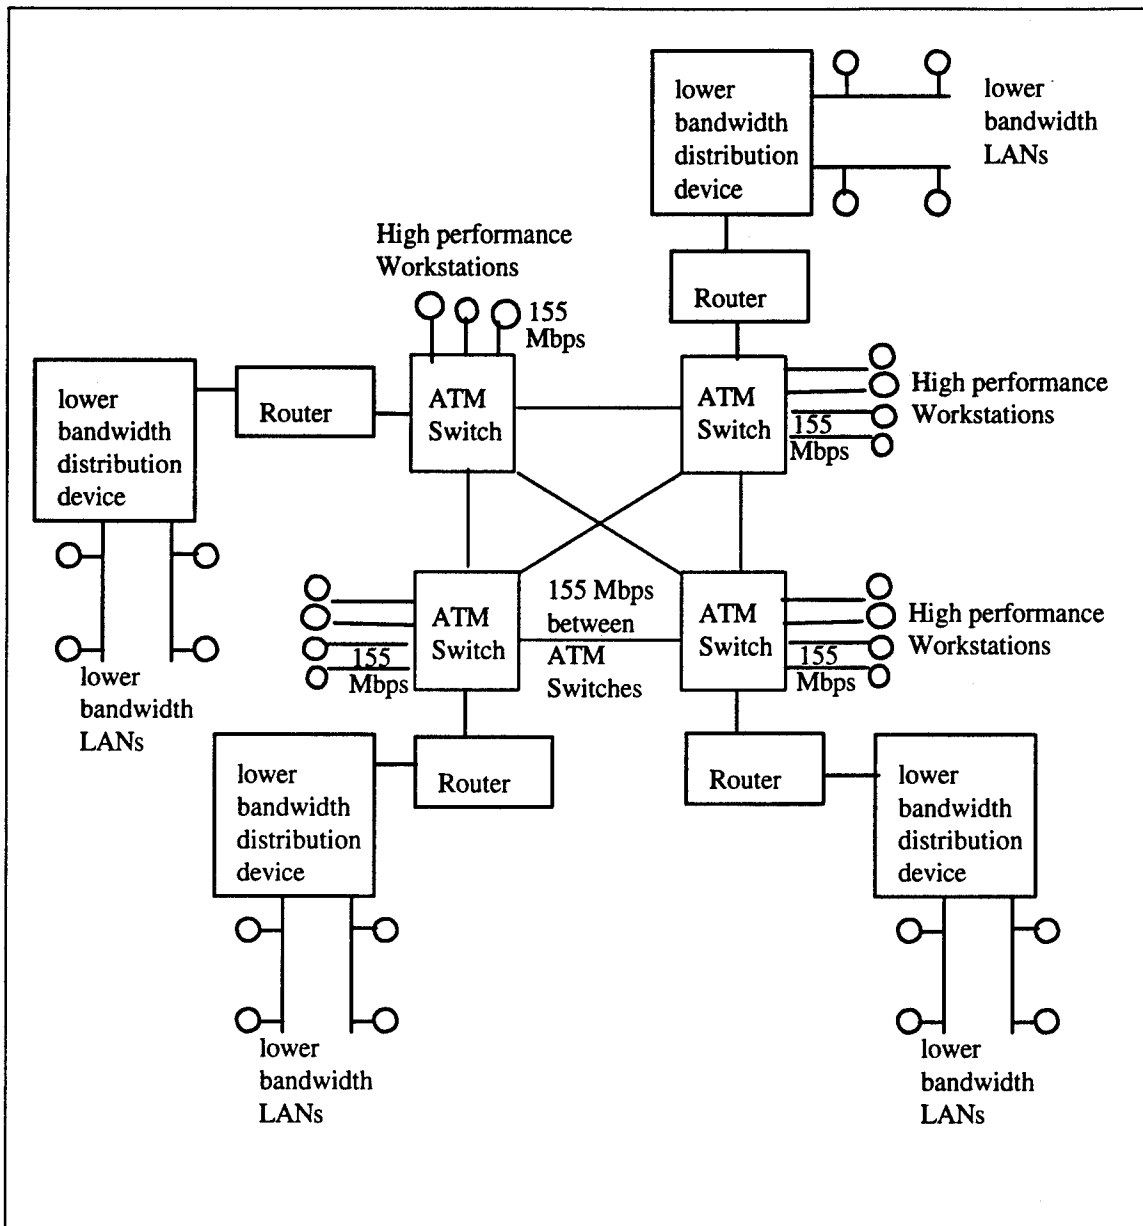


Figure 19. CTC parallel backbone network.

The ethernet network consists of approximately 200 microcomputers.

Applications running on this network consist primarily of electronic mail, accounting, and other lower-bandwidth administrative applications.

CTC used a parallel backbone strategy. The users and applications requiring the movement of large amounts of data, graphics files, and images on the network, justified the move to the ATM network. These applications run more effectively on the ATM network and have been further enhanced by the added capability of the users to collaborate face-to-face through videoteleconferencing. Other users and applications remain on the ethernet network. These other users have benefited from the move as well-- their network is no longer degraded by the higher-bandwidth applications.

This parallel-backbones strategy is a reasonable strategy for NPS for research applications on a few of the subnetworks and to relieve the strain on the existing network by these applications. The existing lower-bandwidth subnetworks could continue evolutionary improvements while network managers gain experience in using higher-bandwidth technologies to meet the needs of applications that greatly benefit from improved performance. Those applications include videoteleconferencing , simulations, and large data file transfers.

VIII. CONCLUSIONS AND RECOMMENDATIONS

A. CONCLUSIONS

The present configuration of the Naval Postgraduate School computer network architecture is primarily a mixture of lower-bandwidth topology subnetworks. If network standards are followed and the subnetworks are properly managed, this configuration could meet the present aggregate needs of most users and applications. However, it cannot effectively support changes in the organization or emerging network applications and network technologies. Any strategic planning regarding network architecture must address these shortcomings.

Organizational goals should dominate the driving forces, but because of the unique nature of the NPS organization as a military, academic, and research institution, the goals remain unclear. Network management, funding, and the architecture itself reflect the nature of the organization--overlapping and diffuse. Thus, as a support service, the network architecture must be flexible enough to respond effectively to a broad spectrum of organizational demands yet still be manageable with no increase in human resources.

Technical driving forces, including practical limitations of transmission media, capabilities of network topologies, and bandwidth demands of applications are more quantifiable than the organizational forces. Because network technologies are always advancing, network managers should choose an architecture that can accommodate change at minimal cost.

Several technological alternatives are available to meet the projected future technical and organizational demands on the network architecture. These alternatives may also allow the network to be more responsive to the organizational demands and improve network management. A switch-based high performance network, using star- topology wiring to network nodes and mesh-topology wiring among distribution devices, provides the most scalable, sustainable, survivable, and secure network architecture.

There are several strategies for migrating to a switch-based network. The most reasonable for NPS is a "parallel backbones" approach. In this approach, the small number of users and applications that can justifiably benefit from the higher performance network move directly to an end-to-end switch-based network. The remaining users and applications remain on the lower-bandwidth network until such time as their move to the higher-bandwidth network becomes practical. The following section discusses this strategy in greater detail as well as other recommendations.

B. RECOMMENDATIONS

To obtain the type of architecture discussed above, this study recommends several improvements to the existing network management and network architecture.

Recommendations vary from general awareness and support of the network to specifics regarding network architecture.

1. General

The ultimate goal of the network is for it never to be the limiting factor to its users or the organization--it should be a reliable support service similar to telephone and

electricity services. The highest levels of management within the NPS organization should recognize the importance of the network architecture in the day-to-day business at the Naval Postgraduate School and support it accordingly. Just as it would be difficult to function for an extended period of time without telephones or electrical power, loss of the computer network for any length of time would become a severe detriment to the operation of the organization.

Towards the goal of never being a limiting factor, the network architecture should migrate to a highly responsive, flexible switch-based architecture using a parallel backbones approach. Management at all levels in the organization, including network managers, should ensure this approach meets the goals and has the necessary support and oversight of the organization. Any upgrades or changes to the network or subnetworks should be towards the goal of a dynamic network--an architecture capable of adapting to changing needs and missions of users and the organization and capable of incorporating new technologies with minimal changes in cabling.

Organizational management must support the computer network architecture as it enables NPS to achieve its *Vision 2000*. In turn, the network can better support the NPS *Guiding Principles*. Properly supported and managed, the network enables--not hinders--the NPS strategic plan.

Network management must take several steps to alleviate the concerns caused by present unclear overall organizational goals and support. Recommendations in [Report] address many of these issues.

- First, Network administration should develop, document, announce, and disseminate a clear strategy and vision of their own within the constraints of the overall NPS organization.
- Second, network management should review their processes and resources to identify their own strengths and weakness within the constraints placed by the organization.
- Third, network personnel should sell the importance of the network to its users, organizational management, and potential customers outside the organization. If these groups are aware of the capabilities and potential of the network, they may be more inclined to support it and upgrades to the architecture.
- Fourth, high-visibility actions could strengthen the position of the network, such as advertising and improving end-user support in the form a centralized network operations center, computing resource trouble desk, and user help desk.
- Fifth, network technologies have advanced sufficiently that network management should consider consolidating network personnel and management into a single, cohesive department.
- Finally, NPS administration should appoint a full-time, qualified "chief information officer" (CIO) to the position of Dean of Computer and Information Services, Code 05. The CIO should have authority, discretion, and organizational support necessary to strengthen the role of the network and NPS computing in general.

2. Network Architecture Specifics

Network managers should invest in a parallel backbone network for those users and applications that demand network capacity that exceeds the practical limitations of the current architecture. Network managers should study the present network's efficiency and effectiveness in depth. Managers do not have enough data about network and subnetwork performance and utilization to position strategic investments in architecture specifics.

Network managers should continue their analysis of applications and tools that will help

monitor and evaluate the network. Similarly, network managers also must develop a plan to test and measure the benefits of any new architecture.

a. Subnetwork Cabling

Network managers should map all network cabling and distribution devices throughout the campus. Having a schematic of the actual cable runs will help managers identify potential bottlenecks and areas where topologies require conformance to networking standards. Network managers should consolidate such mapping in electronic format so they can readily modify and display cabling maps for trouble-shooting and planning. This is part of proactive network management.

Any new cable installations should also consider other organizational requirements for signal distribution such as telephone and television signals to classrooms. To support all potential end-user needs, installations should be compliant with EIA/TIA-568 structured wiring standards. These installations should include at least eight pairs of Category 5 UTP between each node and wiring closet. Network managers should invest in optical fiber cabling for higher capacity data distribution among wiring closets and between buildings.

b. Distribution and End Devices

To improve management and survivability of the network, managers should install uninterruptible power supplies and hot-swappable modules on all critical network components including routers, hubs, and file servers. Any new device on the network

should allow remote management (e.g., using Simple Network Management Protocol) or at least remote notification of network problems that arise at the device.

3. Further Research

As with many studies, research for this thesis uncovered more questions than could be answered within the scope of the thesis. Some opportunities for further research follow.

a. Measures of Benefits

The measures of network benefits discussed in Chapter V require tailoring to the specifics of the NPS network. Study in this area could help network managers better grasp the efficiency and effectiveness of the network. Network performance efficiency requires technical study. Survey and analysis of user satisfaction would enlighten and/or reinforce network administration and policies with respect to network architecture effectiveness.

b. Organizational Issues

A managerial study of the NPS organization and its relation to the network would assist network managers in better supporting the goals and missions of the organization. A greater understanding of NPS organizational dynamics would assist network managers in developing strategies for the network. Study in this area would enhance the understanding of the role of computing at NPS.

c. Network Management Applications

There are numerous tools available on the market that aid in managing computer networks. Evaluation of the options would assist network managers in choosing the applications necessary to build and operate a reasonable network operations center. Computer center personnel have begun research in this area, but in-depth analysis of alternatives would provide them with the opportunity to make a better informed investment in network management tools.

C. SUMMARY

The computer network on the Naval Postgraduate School campus has become an integral part of the operations of the Naval Postgraduate School organization. This study describes the Naval Postgraduate School computer network architecture, driving forces, limitations, and possible measures of network benefits. It considers network alternatives and reasonable transition strategies. The analysis offers recommendations for improvements to the existing network configuration. The analysis of the network architecture provides information, alternatives, and recommendations to assist management in formulating strategic plans that could support the network and NPS to the end of the century.

APPENDIX A. NETWORK ARCHITECTURE QUESTIONNAIRE

The purpose of this questionnaire is to take a "snapshot" view of the computer networks architecture on NPS campus. This questionnaire supports thesis research by Paul Wiedenhoef (PM-31, pewieden@nps).

General/contact info:

1. (sub)network designation:
2. (sub)network general location (building, floor, etc.):
3. Interviewee's name:
4. Interviewee's phone number:
5. Interviewee's e-mail address:
6. Date of interview:

N/W admin.:

1. How did your (sub)network reach its current configuration? (What is the history of your LAN?)
2. What type of medium do you use to connect your nodes/workstations together?
 - ◆ copper wire
 - UTP cat 3
 - UTP cat 5
 - STP 150 ohm
 - STP 100 ohm

- thick coax with AUI
 - thin coax
 - ◆ optical fiber
 - multi-mode
 - ◆ other (?)
1. What type of logical topology do you use in your network?

- ◆ bus CSMA/CD (IEEE 802.3)
 - ethernet 10 Mbps

Thick--10Base5

Thin ("thinnet", "cheapernet")- 10Base2

10Base-T

- "starlan" 1 Mbps--1Base5
- broadband--10Broad36
- ◆ bus token passing (IEEE 802.4)
- ◆ ring token passing (IEEE 802.5)
 - 1 Mbps
 - 4 Mbps
 - 16 Mbps
- ◆ ANSI FDDI
- ◆ proprietary
 - localtalk
 - other (?)

1. Do you have a schematic of your (sub)network? If not, sketch the physical connections of your network.

2. Do you have any repeaters in your network? if so, how many, brand/model, capabilities...
3. Do you have any hubs in your network? if so, how many, brand/model, capabilities...
4. Do you have any routers connected to your network? if so, brand/model, protocols supported, capabilities...
5. Do you have any other type of distribution devices on your (sub)network? if so, brand/model, protocols supported, capabilities...
6. Are you connected to the campus backbone?
7. How are you connected to the campus backbone? (vampire tap, router, bridge, other(?))
8. Do you have any "back door" connections to other LANs or WANs other than a direct connection through the campus backbone? (please list/describe)
9. What is the physical length of your cabling runs? each segment and overall (?)
10. If bus topology, what is the maximum number of collisions observed?
11. If bus topology, what is the average number of collisions observed?
12. Do you have a strategic plan for your (sub)network?
13. What changes do you anticipate making in the near future (less than 5 years) in the architecture of your network?
14. What would you like to change in your (sub)network within the next five years?
15. What protocols are supported on your (sub)network
16. What applications run across/on your (sub)network?
17. What applications do you anticipate running on your network in the next five years?
18. How many users/nodes are on your network?

19. How many logical addresses are on your network? (IP addresses)
20. How many addresses do you anticipate adding in the next five years?
21. If you have a server (or servers) on your (sub)network, what is its IP address?
22. How many man-hours does administration take per week?
23. What are your responsibilities as (sub)network administrator?
24. What physical security measures do you practice to protect your (sub)network?
25. What software security measures do you practice to protect your (sub)network?
26. Do you know of any other (sub)networks nearby or sharing the same cabling that I may not have already identified?
27. Who are the primary users of your (sub)network? (what faculty, what staff, what students, what tenant commands?)
28. What are the names of the users on your (sub)network that run applications that demand the most bandwidth on the (sub)network?

User:

1. What applications do you run [that use your (sub)network]?
2. What applications do you anticipate adding to your (sub)network or needing on your (sub)network within the next five years?
3. What would you like to add to your (sub)network to improve your work or use of the (sub)network that is not included in the response to the previous question?

APPENDIX B. SUBNETWORK SUMMARY

Connection	Pairs	"dark" or damaged pairs	Subnetworks serviced	Notes
(Ingersoll to) Bullard Hall	1	0	131.120.021 131.120.025	aerial; crosses over top of Halligan Hall
Halligan Hall	1	0	(131.120.148) 131.120.149	aerial; .148 routed off .149
Root Hall	9	6	(131.120.062) 131.120.140 131.120.146	underground .062 routed off .140; one pair used to directly connect to IDEA lab computer to VAX in Spanagel - not on the campus network.
Spanagel Hall	12	5	131.120.001 131.120.007 131.120.020 131.120.060 (131.120.061) 131.120.101 131.120.254	aerial; crosses over top of Root Hall then underground between Root Hall and Spanagel Hall; .061 routed off .060; one pair used directly connect VAX in Spanagel to computer in IDEA lab in Root Hall - not on the campus network.
Dudley Knox Library	1	0	131.120.051	underground
Glasgow Hall	12	6	131.120.141 131.120.142 131.120.143 131.120.144 131.120.145 131.120.147	underground
Bldg 223/224	1	0	131.120.056	underground; thinnet across Bldg 223 to 224
Bldg 203/200	1	0	131.120.057	underground; thinnet underground between Bldg 203 and 200
Hermann Hall to Bldg 427	1	0	131.120.080 131.120.081 131.120.130 131.120.131 131.120.132 131.120.133	underground

62.5/125 micron multi-mode optical fiber interbuilding cabling.

Connection	Number of segments	"dark" or damaged segments	Subnetworks serviced	Notes
Ingersoll Hall to Hermann Hall	1	0	131.120.080 131.120.081 131.120.130 131.120.131 131.120.132 131.120.133	underground

Thicknet coaxial cable interbuilding cabling.

Subnetwork designation	Internetwork protocol address	General location(s)	Transmission media	Distribution devices	MAC layer topology	Network layer protocols	Primary users	Primary applications	Devices on nodes	Network Admin point of contact and extension (1)
	(1)(3)	(3)	(1)(3)	(1)(3)	(2)	(4)	(1)(3)	(1)(3)	(1)(3)	(1)
CS Primary	131.120.001	Spanagel Hall 5th floor	thicknet		FOIRL IEEE 802.3	IP APP VIP	CS Dept faculty, staff, students	admin, research		Mike Williams x2550
CS Graphics	131.120.007	Spanagel Hall 5th floor			FOIRL	IP	Computer Science Dept faculty, staff, students	research		Mike Williams x2550
ECE Spanagel	131.120.020	Spanagel Hall 3rd floor	thicknet		FOIRL IEEE 802.3	IP APP	ECE Dept faculty, staff, students	admin, research	Sun	Bob Limes x3216
ECE Bullard	131.120.021	Bullard Hall			FOIRL IEEE 802.3	IP	ECE Dept faculty, staff, students	admin, research		Bob Limes x3216
Space Systems	131.120.025	Bullard Hall			FOIRL IEEE 802.3	IP XNS IPX APP VIP		admin, research		Jim Horning x3199
SM Labs	131.120.030	Ingersoll Hall 2nd floor rms 224, 250		multi-station access units	IEEE 802.5: Token ring	IP XNS IPX VIP APP	students	admin, research: email, FTP, spreadsheets, RDBMS, word processing	PC Mac	Leon Sahlman x3574
SM Labs (secondary to 131.120.030)	131.120.038	Ingersoll Hall 2nd floor rms 224, 250		multi-station access units	IEEE 802.5: Token ring	IP XNS IPX VIP APP	students	admin, research: email, FTP, spreadsheets, RDBMS, word processing	PC Mac	Leon Sahlman x3574
Systems Management (secondary to 131.120.040)	131.120.039	Ingersoll Hall 3rd floor	thicknet thinnet	multi-port 10Base5 transceivers, multi-port 10Base2 repeaters	IEEE 802.3	IP	Systems Management Dept faculty and staff, TQL Office	admin, research	Sun, PC, Mac	Leon Sahlman x3574
Systems Management	131.120.040	Ingersoll Hall 3rd floor	thicknet thinnet	multi-port 10Base5 transceivers, multi-port 10Base2 repeaters	IEEE 802.3	IP XNS IPX VIP	Systems Management Dept faculty and staff, TQL Office	admin, research	Sun, PC, Mac	Leon Sahlman x3574

Subnetwork designation	Internetwork protocol address	General location(s)	Transmission media (1)(3)	Distribution devices (1)(3)	MAC layer topology (2)	Network layer protocols (4)	Primary users (1)(3)	Primary applications (1)(3)	Devices on nodes (1)(3)	Network Admin point of contact and extension (1)
Computer Center Main	131.120.050	Ingersoll Hall 1st floor	thicknet, thinnet cat 5 UTP	single and multi-port 10Base5 transceivers, multi-port 10Base2 repeaters,	IEEE 802.3: 10Base5, 10Base2, 10Base-T	IP XNS IPX VIP APP	computer center administration and operations, students	computer center admin, email, distributed processing	mainframe, Sun PC	Raul Romo x 2004 Terry Gentry x3432
Library	131.120.051	Dudley Knox Library	multi-mode optical fiber thinnet, thinnet, cat 5 UTP	fiber-optic transceiver, single and multi-port 10Base5 transceivers, 10Base-T hub multi-port 10Base2 repeater	FOIRL IEEE 802.3: 10Base5 10Base2 10Base-T	IP IPX VIP APP	Library staff	admin, library research	PC, Mac	Raul Romo x2004 Terry Gentry x3432 Diane Crankshaw x3342
Microcomputer Lab	131.120.052	Ingersoll Hall 1st floor rm 104, 151, 371	Cat 3 UTP, (AUI from router to 151 and fm 104 to 151)	Multi-port 10Base5 transceiver, 10Base-T hub	IEEE 802.3: 10Base2	IP XNS IPX VIP APP	students	word processing	PC,	Joe Rogers x3660
Visualization Lab	131.120.053	Ingersoll Hall 1st floor rm 102A, 102B, 148, 135	multi-mode optical fiber, thicknet, thinnet		IEEE 802.3 10Base2 (*might be ANSI FDDI by now...)	IP APP FDDI	computer center staff, students	Visualization, distance learning (M-bone),	Super-computer, SGI, HP, Sun, Mac,	Mike McCann x2752
Cray/DEC	131.120.054	Ingersoll Hall rm 135, 148	multi-mode optical fiber		ANSI FDDI	IP	computer center staff	Visualization	Super-computer, DEC	Mike McCann x2752
Study Barn/224	131.120.056	Bldgs 223, 223A, 224	multi-mode optical fiber, thinnet	fiber-optic transceiver,	FOIRL IEEE 802.3: 10Base2	IP VIP	meteorology staff	research	Sun, PC	Raul Romo x2004 Terry Gentry x3432

Subnetwork designation	Internetwork protocol address	General location(s)	Transmission media (1)(3)	Distribution devices (1)(3)	MAC layer topology (2)	Network layer protocols (4)	Primary users (1)(3)	Primary applications (1)(3)	Devices on nodes (1)(3)	Network Admin point of contact and extension (1)
TRAC Monterey	131.120.057	Bldgs 203, 200	multi-mode optical fiber, thinnet Cat 5 UTP	fiber-optic transceiver, 10Base-T hub, 8-port 10Base2 repeaters,	FOIRL IEEE 802.3: 10Base2, 10Base-T	IP APP VIP	TRAC Monterey, Security, Contracting (PWC), OA, Command Evaluation,	Simulations email, PV Wave, Frame, FOCAS, CCASS, WP Office	VAX, Sun, SGI, HP, PC, Mac	Jeff Ingram x3087
Oceanography	131.120.060	Spanagel Hall 3rd floor	multi-mode optical fiber thicknet	fiber-optic transceiver	FOIRL IEEE 802.3	IP	Oceanography faculty staff and students	admin and research		
Oceanography Research	131.120.061	Spanagel Hall 3rd floor off 131.120.060	thicknet		IEEE 802.3	IP VIP	Oceanography faculty, staff and students	research		Larry Blalock x2567
COAC Lab	131.120.062	Root Hall 1st floor rms 106-107 off 131.120.140	thicknet thinnet	Sparc 1+ acts as router off .140	IEEE 802.3: 10Base2	IP	Oceanography, USW, ECE	file transfers, NFS	Sun, HP, PC	Stephen Hudson x
MIS Vines / TCP gateway (sub-divided from 131.120.130)	131.120.080	Hermann Hall, PWC Bldgs	thicknet thinnet Cat 3 UTP multi-mode optical fiber	multi-port 10Base5 transceiver, multi-port 10Base2 repeaters, 10Base-T hubs, fiber-optic transceiver,	FOIRL IEEE 802.3: 10Base5 10Base2 10Base-T	IP XNS IPX VIP APP	NPS Administration	email, office suites, C/S apps RDBMS, spreadsheets, CAD/CAE		Joe LoPiccolo x2994
MIS (sub-divided from 131.120.130) (secondary to 131.120.130)	131.120.081	Hermann Hall, PWC Bldgs	thicknet thinnet Cat 3 UTP multi-mode optical fiber	multi-port 10Base5 transceiver, multi-port 10Base2 repeaters, 10Base-T hubs, fiber-optic transceiver,	FOIRL IEEE 802.3: 10Base5 10Base2 10Base-T	IP XNS IPX VIP APP	NPS Administration	email, office suites, C/S apps RDBMS, spreadsheets, CAD/CAE		Joe LoPiccolo x2994

Subnetwork designation	Internetwork protocol address	General location(s)	Transmission media	Distribution devices	MAC layer topology	Network layer protocols	Primary users	Primary applications	Devices on nodes	Network Admin point of contact and extension (1)
	(1)(3)	(3)	(1)(3)	(1)(3)	(2)	(4)	(1)(3)	(1)(3)	(1)(3)	
Physics	131.120.101	Spanagel Hall Physics Dept	thicknet thinnet	multi-port 10Base5 transceiver, multi-port 10Base2 repeaters	FOIRL IEEE 802.3: 10Base5 10Base2	IP DECnet IPX VIP	Physics Admin	email, office suites, C/S apps RDBMS, spreadsheets		Joe Blau x2685
MIS	131.120.130	Herrmann Hall, PWC Bldgs	thicknet, thinnet, Cat 3 UTP, multi-mode optical fiber,	single and multi-port 10Base5 transceivers multi-port 10Base2 repeaters 10Base-T hub fiber-optic transceiver	FOIRL IEEE 802.3: 10Base5 10Base2 10Base-T	IP XNS IPX VIP APP	NPS Administration	email, office suites, C/S apps RDBMS, spreadsheets, CAD/CAE ,		Joe LoPiccolo x2994
Computer Science (HE 5--)	131.120.131	Herrmann Hall	thicknet		IEEE 802.3	IP				
DRMI	131.120.132	Herrmann Hall West Wing	thicknet		IEEE 802.3	IP	DRMI			
DRMI	131.120.133	Herrmann Hall West Wing	thicknet		IEEE 802.3	IP	DRMI			
DMDC	131.120.135	DMDC Monterey			IEEE 802.3	IP	DMDC	personnel record data transfer		
DMDC	131.120.136	DMDC Monterey			IEEE 802.3	IP	DMDC	personnel record data transfer		
DMDC	131.120.137	DMDC Monterey			IEEE 802.3	IP	DMDC	personnel record data transfer		

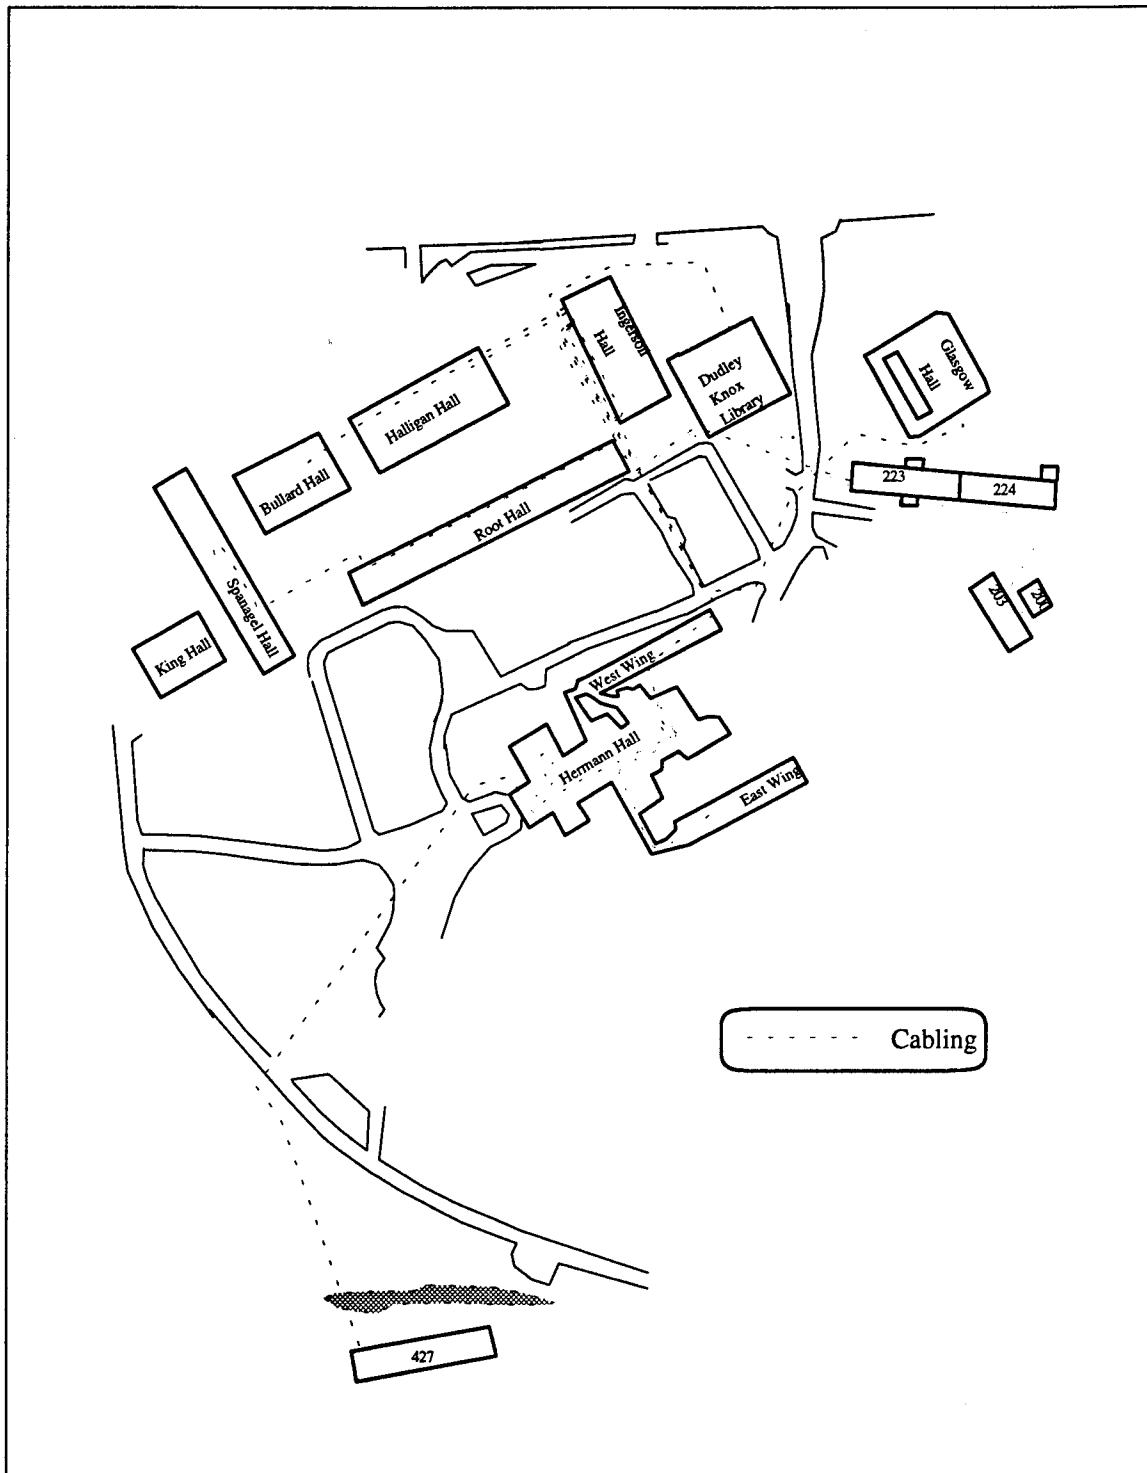
Subnetwork designation	Internetwork protocol address	General location(s)	Transmission media	Distribution devices	MAC layer topology	Network layer protocols	Primary users	Primary applications	Devices on nodes	Network Admin point of contact and extension (1)
Root Hall	131.120.140	Root Hall 1st floor: rms 100-103, 108 2nd floor: rms 200-205, 210-234, 251-277	thicknet thinnet cat 5 UTP cat 3 UTP	4-port 10Base2 repeaters 8-port 10Base-T hubs	IEEE 802.3: 10Base5 10Base2 10Base-T	IP XNS IPX APP VIP	DHRSC, CIFO, IAC faculty and staff, students, Rm 262 LRC, Registrar, Research Services, SCIF Lab,	admin, research	Sun, SGI, HP, PC, Mac,	Raul Romo x2004 Terry Gentry x3432 Chuck Taylor (rm 262 LRC) x2539 Gary Porter (rm 204 Lab) x Joe LoPiccolo (rms 100-103) x2994
Glasgow 1st Floor	131.120.141	Glasgow Hall 1st floor	multi-mode optical fiber cat 5 UTP	10Base-T hub (U-B)	FOIRL IEEE 802.3: 10Base-T	IP XNS IPX APP VIP				Lary Moore x3170
Glasgow 2nd Floor	131.120.142	Glasgow Hall 2nd floor	multi-mode optical fiber cat 5 UTP	10Base-T hub (U-B)	FOIRL IEEE 802.3: 10Base-T	IP IPX APP VIP	Math Dept faculty and staff	admin, research		Lary Moore x3170
Glasgow Math	131.120.143	Glasgow Hall 3rd floor	multi-mode optical fiber cat 5 UTP	10Base-T hub (U-B)	FOIRL IEEE 802.3: 10Base-T	IP IPX APP VIP				
Glasgow LRC 1	131.120.144	Glasgow Hall 1st floor rm 128	multi-mode optical fiber cat 5 UTP	10Base-T hub (U-B)	FOIRL IEEE 802.3: 10Base-T	IP IPX APP VIP			Mac, PC	
Glasgow LRC 2	131.120.145	Glasgow Hall 2nd, 3rd floors rm 203, 309, 318	multi-mode optical fiber cat 5 UTP	10Base-T hub (U-B)	FOIRL IEEE 802.3: 10Base-T	IP IPX APP VIP				Lary Moore x3170
IDEA Lab	131.120.146	Root Hall 1st floor rms 117-123 2nd floor rms 229-255	multi-mode optical fiber, thicknet thinnet	fiber-optic transceiver 8-port 10Base2 repeater	FOIRL IEEE 802.3: 10Base5 10Base2 10BaseF	IP DECnet IPX APP VIP	Meteorology Dept faculty, staff, and students Boundary Layer Studies Group	research, admin	HP, Mac, PC	Russ Schwanz x3177

Subnetwork designation	Internetwork protocol address	General location(s)	Transmission media	Distribution devices	MAC layer topology	Network layer protocols	Primary users	Primary applications	Devices on nodes	Network Admin point of contact and extension
	(1)(3)	(3)	(1)(3)	(1)(3)	(2)	(4)	(1)(3)	(1)(3)	(1)(3)	(1)
NSA	131.120.147	Glasgow Hall 3rd floor	multi-mode optical fiber cat 5 UTP	10Base-T hub (U-B)	FOIRL IEEE 802.3: 10Base-T	IP IPX APP VIP	NSA faculty, staff	admin		Lary Moore x3170
Mechanical Engineering	131.120.148	Halligan Hall (routed off 131.120.149)	multi-mode optical fiber thicknet		FOIRL IEEE 802.3	IP	Mechanical Engineering Dept faculty, staff, students	admin, research		Dave Marco x2809
Aeronautics	131.120.149	Halligan Hall	multi-mode optical fiber thicknet		FOIRL IEEE 802.3	IP IPX APP VIP	Aero Dept faculty, staff, students	admin, research		Tony Cricelli x2910
Public Cluster/Router Backbone (old campus backbone)	131.120.254	Ingersoll Hall/ Spanagel Hall	multi-mode optical fiber thicknet, cat 5 UTP	Cisco routers multi-port 10Base5 transceivers 10Base-T hub	FOIRL IEEE 802.3: 10Base5 10Base-T	IP XNS DECnet IPX APP VIP	Cisco routers, all network users	admin, research, backbone internetwork connections	Sun, AGS+, CGS	Raul Romo x2004 Terry Gentry x3432

Notes:

- (1) Blank cell indicates data not collected or data not available.
- (2) Where MAC topology shown as IEEE 802.3 but specific topology not indicated, data not collected or data not available; FOIRL-Fiber Optical Inter Repeater Link
- (3) This table is not meant to be exhaustive, but representative of the NPS campus computer network.
- (4) IP- DoD Internetwork Protocol; IPX-Novell Internetwork Packet Exchange; APP- Apple Appletalk; VIP-Vines Internetwork Protocol; XNS-Xerox Network Services; DECnet-Digital Equipment Corporation network protocol

APPENDIX C. NETWORK DIAGRAMS

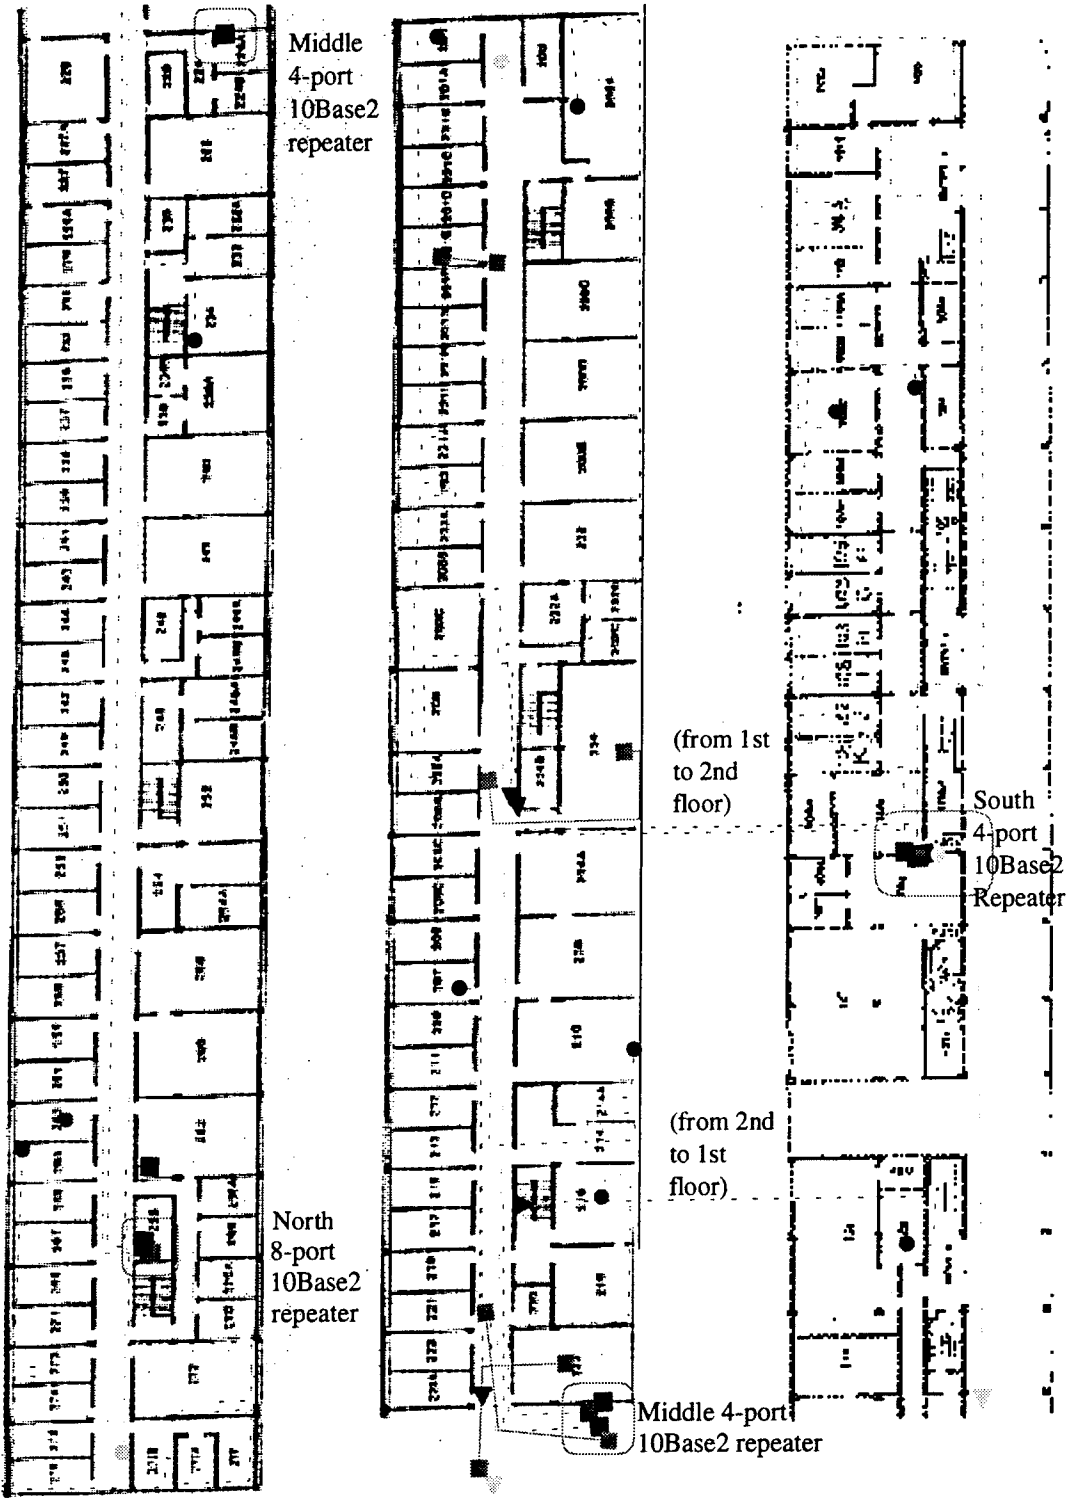


NPS interbuilding computer network cabling.

Root Hall North
(2nd floor)

Root Hall South
(2nd floor)

Root Hall South
(1st floor)



Root Hall 131.120.140 subnetwork cabling.

Notes to accompany 131.120.140 subnetwork diagram:

Optical Fiber

The optical fiber segment terminates from the router in Ingersoll Hall in Root Hall room 268. A FOIRL/10Base5 repeater connects to a segment of (blue) AUI which connects to a 10Base5 transceiver tapped into a thicknet segment in the 2nd floor hall outside room 268.

Thicknet

A thicknet segment runs on the 1st floor primarily in electrical conduit from room 124 until it terminates in room 106B. In room 106B, two transceivers tap into the thicknet. One transceiver connects via AUI to a Sun Sparc 1+ that acts as a router for 131.120.62 subnetwork traffic in the COAC lab, surrounding offices, and room 107. Another transceiver connects to a 4-port 10Base2 repeater in room 106B (hereafter referred to as *South 10Base2 Repeater*).

A thicknet segment runs entire length of 2nd floor in the hall. On the north end, the thicknet is on the west side from room 277 until room 242. Between rooms 242 and 240, it crosses over to the east side of the hall. The thicknet remains on the east side until it terminates at the south end outside room 201.

Several single-port 10Base5 transceivers tap into the cable along its length. Outside room 268, a tap connects via (gray) AUI to a single port 10Base2 repeater in room 268. This single port repeater connects via thinnet to an 8-port 10Base2 repeater (hereafter referred to as *North 10Base2 Repeater*) also in room 268. A (gray) AUI connects *North 10Base2 Repeater* to a transceiver tapped into a segment of thicknet that runs on the 1st floor. Outside room 225, a tap connects via AUI to a 10Base-T hub in room 222 (the public access terminal room). Outside room 223, a tap connects to a 4-port 10Base2 repeater in room 224A (hereafter referred to as *Middle 10Base2 repeater*). Another transceiver, outside 205B/A connects to a 4-port 10Base2 repeater in room 204. the final tap, at room 201E connects to nothing at present.

UTP

In room 222, AUI connects to a 10Base-T hub. Two 10Base-T hubs cascade to connect 15 workstations and a network laser printer.

Thinnet

This note describes the main segments of thinnet on the 131.120.140 subnet by the order of the spaces encountered from repeater to termination.

South 10Base2 Repeater

Segment 1: 106B, 106D, 103K, 103J, 103I, 103H, 103G, 103F, 103E, 103D, 103C
(crosses Segment 2 in room 103C), Hall.

Segment 2: 106B, 106D, Entry 2 (into conduit), (outside), Entry 1 (out of conduit), 100, 100A, 101, 101A, 102A, 102, 103A, 103B, 103C (crosses Segment 1 in 103C).

Segment 3: 106B, 106D, (up to second floor through electrical wiring runs), 2nd floor hall west side outside 204B, cross to east side at 203E/D, 203C, 203D, 203E, 205, 203E, 203D, 203C, 203B, 203A, 201J through 201.

Segment 4: 106B, 106D, (up to second floor through electrical wiring runs), 2nd floor hall west side outside 204B, to short hall opposite 203D, 202C, 202B, 202, 200E, 200D, 200C, 200B, 200A.

AUI segment: to single port 10Base5 transceiver at room 223 on 1st floor thicknet.

Middle 10Base2 Repeater

Segment 1: 224A, 224, hall, 234.

Segment 2: 224A, 224, hall, 220, 218, 216.

Segment 3: 224A, 224, hall, Entry 3 (into conduit), 108 (out of conduit).

Segment 4: 224A, 224, hall, (across hall at 223), 217, 209.

AUI segment: to single port 10Base5 transceiver at room 223 on 2nd floor thicknet.

North 10Base2 Repeater

Segment 1: 268, hall, 272, 277A, 277, across north end of hall, 276, 275, 274, 273, 271, 269, 267, 265, 264.

Segment 2: 268, hall, 262 (LRC 10Base5 bridge).

Segment 3: 268, across hall, 267, 265, 264, 263.

Segment 4: 268, across hall, (skip several room while transiting hall), 229A, 229, 227, 227A, 225, 223A, 223, 221, 219, 217, 215, 213, 212, across hall at 213, 214, 210.

Segment 5: (to single-port 10Base2 repeater in 268)

Segment 6: (dummy load)

Segment 7: (dummy load)

Segment 8: (dummy load)

AUI segment: to 1st floor thicknet segment via 10Base5 transceiver.

APPENDIX D. CAPACITY ASSESSMENT

[Vis] presents a simple formula for the assessment of LAN performance in lieu of sophisticated analytical models and tools. It is based only on the most essential performance parameters. Vis contends his "rule of thumb" is a fairly accurate conservative measure when compared to a closed queuing model of a time sharing system.

The essential performance parameters are LAN transmission speed, C_{LAN} , total capacity required over a period, C , and maximum of the minimal required transfer speed per application, S . These three parameters are associated by an equality that approximates the more general inequality which describes a simple capacity formula for the design of a LAN:

$$S = C_{LAN} - C \quad (\text{Equation 2})$$

The fundamental reasoning behind this, Vis explains, is that the maximum speed at which a certain amount of data can be transmitted over a LAN is, on average, equal to the difference between the total capacity of the network and the average capacity used by other stations on the network. Similarly, when the capacity required, C , on average, by a station requiring transfer speed, S , is small compared to the total capacity required for all stations and applications, C_{LAN} , the formula can be rewritten:

$$C = C_{LAN} - S \quad (\text{Equation 3})$$

Vis provides examples of the application of these formulae used in calculating the effect of adding an application requiring a certain transfer speed to the network and another for calculating the maximum number of users on the network. He also provides theoretical justification for assessment measure.

LIST OF REFERENCES

- [Abrahams] Abrahams, Seth. "What Do You Do When The Network Is 1 to 2 Orders Of Magnitude Faster Than The Backplane Bus?" *High-Speed Fiber Networks And Channels*, Kadiresan Annamalai, Editor, Proc. SPIE 1577, p. 316, 1992.
- [Black Box] Black Box Corporation, *BLACK BOX Catalog*, September 1994.
- [Bolles] Bolles, Gary A., "Fast Networking? Slow Down!" *Network Computing*, v. 5, nr. 3, p. 12, March 1, 1994.
- [Botelho] Botelho, Margot A., "Understanding Fiber Optics," reprint from *LAN Times*, September 1990.
- [Bryan 1] Bryan, John, "Fast Ethernet; No End In Sight," *Computer Shopper*, v. 14, nr. 1, p. 636, January 1994.
- [Bryan 2] Bryan, John, "Fibre Channel Speeds Up," *BYTE*, v. 19, nr. 8, p. 123, August 1994.
- [Capetanakis] Capetanakis, John, "Drive Your LAN Into The Fast Lane," *Government Computer News*, v. 13, nr. 7, p. 37, April 4, 1994.
- [Choi and Kim] Choi, Dug-Kyoo, and Kim, B. G., "The Expected (Not Worst-Case) Throughput Of The Ethernet Protocol," *IEEE Transactions On Computers*, v. 40, nr. 3, p. 245(8), March 1991.
- [Cini] Cini, Al, "Winning The Integration Game," *LAN Computing*, v. 5, nr. 5, p 32(2), May 1994.
- [Cisco 1] Cisco Systems, Incorporated, *Product Information: AGS+*, April 1994.
- [Cisco 2] Cisco Systems, Incorporated, *Technology Brief: Asynchronous Transfer Mode (ATM)*, p. 23, Cisco Systems, Inc., 1993.
- [CNO] Department of the Navy, Office of the Chief of Naval Operations, Letter Serial number 00/4U500182, Subject: Graduate Education Policy, July 27, 1994.
- [Durr] Durr, Michael, "The Virtual Reality Of Switched Networks," *Datamation*, v. 40 nr. 9, p. 34 May 1, 1994.

[Feldman] Feldman, Robert, "Ramping Onto The ATM Superhighway," *LAN Times*, v. 11, nr. 2, p. 62, January 24, 1994.

[Feltman] Feltman, Charles, "ATM: The Grand Unifier," *LAN Magazine*, v. 8, nr. 12, p. 52, November 1994.

[Fraser and Henry] Fraser, A. G., and Henry, P. S., "Transmission Facilities for Computer Communications," *Computer Communication Review*, v. 22, nr. 5, p. 57, October 1992.

[Freeman] Freeman, Roger L., *Telecommunication Transmission Handbook*, 3rd ed., John Wiley & Sons, Inc., 1991.

[Gibbs] Gibbs, Mark, "When Network Computing Is Really Not Work Computing," *Network World*, v. 11, nr. 24, p. 26, June 13, 1994.

[Glass] Glass, L. Brett, "Inside EISA," *BYTE*, v. 14, nr. 12, p. 419, November 1989.

[Hamilton] Hamilton, Drew, "Network Monitoring, Management and Simulation," *Army Research, Development and Acquisition Bulletin*, p. 55, July-August 1994.

[Haugdahl] Haugdahl, J. Scott, "Switch Trio Boosts Bandwidth," *Network World*, v. 11, nr. 30, p. 47, July 25, 1994.

[Hewlett-Packard 1] Hewlett-Packard Company, "Improving Performance With Integrated 32-Bit Networking," *Network World*, v. 11, nr. 28, p. S12, July 11, 1994.

[Hewlett-Packard 2] Hewlett-Packard Company, "Making Fast Ethernet Real," *Network World*, v. 11, nr. 28, p. S16, July 11, 1994.

[Howard] Howard, Matthew D., Manager, Systems Engineering, Cisco Systems, Inc., phone interview with the author, August 3, 1994.

[Infonetics] Infonetics Research, Inc., *Server Bandwidth Congestion--Practical Limits; 10Mbps Versus 100Mbps Solutions*, Infonetics Research, Inc., 1994.

[Intel 1] Intel Corporation, *Introduction To Fast Ethernet And EtherExpress™ PRO/100 Adapters*, white paper, 1994.

[Intel 2] Intel Corporation, *Reducing Bottlenecks For Maximum Performance*, technical report, 1994.

[Katzan] Katzan, Harry, Jr., *Local Area Networks: A Manager's Guide*, Carnegie Press, Inc., 1983.

[Klett] Klett, Stephen P., "Health Care Looks To ATM For Imaging Apps," *Computerworld*, v. 28, nr. 29, p. 60, July 18, 1994.

[Leahy] Leahy, Kevin M., *The Proposed Naval Postgraduate School Campus Network: Computer Communications For The 1990's*, Master's Thesis, Naval Postgraduate School, Monterey, California, p. 5, March 1988.

[MacAskill 1] MacAskill, Skip, "ATM Provides High-Speed Highway To Success For Engineering Company." *Network World*, v. 11, nr. 26, p. 24, June 27, 1994.

[MacAskill 2] MacAskill, Skip, "Of Doctors, Kids And ATM," *Network World*, v. 11, nr. 26, p. 1(2), June 27, 1994.

[MacAskill 3] MacAskill, Skip, "Users Bank On Switching To Cure Bandwidth Woes," *Network World*, v. 11, nr. 14, p. L2, April 4, 1994.

[MacAskill 4] MacAskill, Skip, "Users Warn Of Moving Too Fast To High-Speed LANs," *Network World*, v. 11, nr. 19, p. 5, May 9, 1994.

[Marks] Marks, Howard, "Cable Vision," *LAN Magazine*, v. 8, nr. 4, p. 79, April 1993.

[Masud] Masud, Sam, "Hold Out Just A Little Longer On Those ATM Purchases," *Government Computer News*, p. 41, May 2, 1994.

[McClimans] McClimans, Fred J., "ATM Vs. FDDI: Comparing Functionality, Features, and Benefits," *Telecommunications*, v. 28, nr. 2, p. 28, February 1994.

[Medici] Medici, Mark, *Ethernet Network Questions And Answers*, version 2.11, Internet ftp site: dorm.rutgers.edu, pub/novell/DOCS, April 19, 1993.

[Metcalf] Metcalfe, R. M., "Computer/Network Interface Design," *IEEE Journal On Selected Areas In Communications*, v. 11, nr. 2, p. 174, February 1993.

[Miller 1] Miller, Mark, "LAN Management Tools: A Perplexing Choice," *Network World*, v. 11, nr. 23, p. 48, March 21, 1994.

[Miller 2] Miller, Mark, "Life On The Fast LAN," *Network World*, v. 11, nr. 6, p. 48, March 21, 1994.

- [Minoli] Minoli, Daniel, "Designing Scalable Networks," *Network World*, v. 11, nr. 2, p. 17, January 10, 1994.
- [Molloy] Molloy, Maureen, "Pushing The Speed Limit," *Network World*, v. 11, nr. 1, p. 5(4), January 3, 1994.
- [Molta] Molta, Dave, "Managing LAN Growth," *Network Computing*, v. 4, nr. 10, p. 6, October 15, 1993.
- [Mulqueen] Mulqueen, John T., "A Study In Fiber," *Communications Week*, nr. 500, p. 1, April 11, 1994.
- [Nelson] Nelson, Mark, "The Squeeze On Data," *LAN Magazine*, v. 9, nr. 5, p. 130, May 1994.
- [Newman] Newman, Michael, Newman and Associates, telephone interview with the author, August 4, 1994.
- [Newton] Newton, Harry, *Newton's Telecom Dictionary*, Harry Newton, 1993.
- [Nolle] Nolle, Thomas, "Too Many Choices Cloud LAN Technology Direction," *Network World*, v. 11, nr. 14, p. 90, April 4, 1994.
- [Norman] Norman, David F., Naval Postgraduate School, Interview with the author, July 18, 1994.
- [Olsen] Olsen, Florence, "Switching Hubs Pour On the Speed: Although The Devices Have the Same Purpose, No Two Are Exactly Alike," *Government Computer News*, v. 12, nr. 17, p. 55, August 16, 1993.
- [One Network Place] One Network Place, *Networking Products Catalog*, Fall 1993.
- [Pigg] Pigg, Jennifer, "The Bottom Line: What Price ATM?" presentation given at a conference entitled *Implementing ATM*, Santa Clara, California, May 16, 1994.
- [Poutain] Poutain, Dick, "A Different Kind Of RISC," *BYTE*, v. 19, nr. 8, p. 185, August 1994.
- [Report] *Report of the Committee on the Status of Computing at NPS, July 27, 1994.*
- [Ricciuti] Ricciuti, Mike, "PCI: Server Bus For The Fast LANe," *Datamation*, v. 40, nr. 6, p. 61, March 15, 1994.

[Roberts] Roberts, Erica, "Fast Ethernet Moves Closer To Standard," *LAN Times*, v. 11, nr. 7, p. 7, April 11, 1994.

[Schnaidt 1] Schnaidt, Patricia, "Plug In At 100," *LAN Magazine*, v. 9, nr. 3, p. 71(8), March 1994.

[Schnaidt 2] Schnaidt, Patricia, "Switch Hunt," *LAN Magazine*, v. 9, nr. 6, p. 75, June 1994.

[Schoenstadt] Schoenstadt, Arthur L., Interview with the author, August 4, 1994.

[Seitz] Seitz, Neal B., et al, "User-Oriented Measures of Telecommunication Quality," *IEEE Communications Magazine*, v. 32, nr. 1, p. 56(11), January 1994.

[Serjak] Serjak, Christopher W., "Migrating To ATM: Four Scenarios To Take Users From Today's Networks To Switched Gigabit 'Fabrics'," *Distributed Computing Monitor*, v. 8, nr. 9, September 1993.

[Shannon] Shannon, C. E., "A Mathematical Theory Of Communication," *Bell Systems Technical Journal*, nr. 27, pp. 379 and 623, 1948.

[Siefert] Siefert, R., Networks and Communications Consulting, electronic mail posted to comp.dcom.lans.ethernet news group, July 18, 1994.

[Sirbu] Sirbu, Marvin, *Bandwidth Requirements For Different Applications As A Function Of Time And Size Of Data Transfer*, Carnegie Mellon, 1994.

[Spada] Spada, Mary E., "Does ATM Make Good Business Sense Now?" *LAN TIMES*, v. 11, nr. 13, p. 56, July 11, 1994.

[Sprague and McNurlin] Sprague, Ralph H., Jr., and McNurlin, Barbara C., *Information Systems Management In Practice*, Third Edition, Prentice Hall, Inc., p. 13, 1993.

[Stallings 1] Stallings, William, *Handbook of Computer-Communications Standards, Volume 2, Local Network Standards*, Macmillan Publishing Company, 1987.

[Stallings 2] Stallings, William, *Data and Computer Communications*, 4th edition, Macmillan Publishing Company, 1994.

[Strauss 1] Strauss, Paul, "ATM Backbones: Complex? And How!," *Datamation*, v. 40, nr. 14, p. 48, July 15, 1994.

[Strauss 2] Strauss, Paul, "Plan Your 2001 Network Now," *Datamation*, v. 40, nr. 15, p. 32, August 1, 1994.

[Vis] Vis, Jan, "A Simple LAN Performance Measure," *Computer Communication Review*, v. 24, nr. 1, p. 7, January 1994.

[Wargo] Wargo, Linda, Total Quality Leadership Office, Naval Postgraduate School, Interview with the author, September 6, 1994.

[Weston] Weston, Rusty, "Your Results Will Vary," *Open Computing*, v. 11, nr. 9, p. 43(9), September 1994.

[Wittman] Wittman, Art, "Order Out Of Chaos," *Network Computing*, v. 4, nr. 3, p.70(8), March 1993.

[Yager] Yager, Tom, "How I Learned To Live With Benchmarks," *Open Computing*, v. 11, nr. 9, p. 56(4), September 1994.

BIBLIOGRAPHY

- Abrahams, Seth. "What Do You Do When The Network Is 1 to 2 Orders Of Magnitude Faster Than The Backplane Bus?" *High-Speed Fiber Networks And Channels*, Kadiresan Annamalai, Editor, Proc. SPIE 1577, p. 309(9), 1992.
- Andren, Emily, "Case Study: FMC," *Network Computing*, v. 5, nr. 2, p. 72, February 15, 1994.
- Andren, Emily, "Case Study: Stanford University," *Network Computing*, v. 5, nr. 2, p. 70, February 15, 1994.
- Andren, Emily, "Case Study: Toronto Dominion," *Network Computing*, v. 5, nr. 2, p. 73, February 15, 1994.
- Axner, David H., "Wiring Hubs: Keys To The Network Infrastructure," *Networking Management*, v. 10, nr. 12, p. 48(5), November 1992.
- Bantz, David F., and Bauchot, Frederic J., "Wireless LAN Design Alternatives," *IEEE Network*, v. 8, nr. 2, March/April 1994.
- Biagioni, Edoardo, Cooper, Eric, and Sansom, Robert, "Designing A Practical ATM LAN," *IEEE Network*, v. 7, nr. 2, p. 32(8), March 1993.
- Black Box Corporation, *BLACK BOX Catalog*, September 1994.
- Bochenski, Barbara, "Developers Need Network Know-How; Communications Integral To Architecture," *Software Magazine*, v. 13, nr. 14, p. 96(1), September 15, 1993.
- Bolles, Gary A., "Fast Networking? Slow Down!" *Network Computing*, v. 5, nr. 3, p. 12(1), March 1, 1994.
- Botelho, Margot A., "Fiber Characteristics Impact Network Performance," *LIGHTWAVE, The Journal of Fiber Optics*, September, 1991.
- Botelho, Margot A., "Understanding Fiber Optics," reprint from *LAN Times*, September 1990.
- Bourne, Philip E., "Mish-Mash Management," *DEC Professional*, v. 12, nr. 5, p. 18(6), May 1993.
- Bryan, John, "Fast Ethernet; No End In Sight," *Computer Shopper*, v. 14, nr. 1, p. 636(4), January 1994.
- Bryan, John, "Fibre Channel Speeds Up," *BYTE*, v. 19, nr. 8, p. 123(5), August 1994.
- Bryan, John, "LANs Make The Switch," *BYTE*, v. 19, nr. 9, p. 113(6), September 1994.
- Capetenakis, John, "Drive Your LAN Into The Fast Lane," *Government Computer News*, v. 13, nr. 7, p. 37(2), April 4, 1994.

- Carlson, A. Bruce, *Communication Systems*, Second Edition, McGraw-Hill, 1975.
- Chernicoff, David, "Remedies For Sluggish LANs: How To Boost Performance With A Minimum Of Pain," *PC Week*, v. 10, nr. 41, p. N12(1), October 18, 1993.
- Choi, Dug-Kyoo, and Kim, B. G., "The Expected (Not Worst-Case) Throughput Of The Ethernet Protocol," *IEEE Transactions On Computers*, v. 40, nr. 3, p. 245(8), March 1991.
- Cini, Al, "Winning The Integration Game; An OSI-based Paradigm Helps You Understand Protocol Relationships In Your Internetwork," *LAN Computing*, v. 5, nr. 5, p. 33(2), May 1994.
- Cipra, Barry A., "Wavelet Applications Come To The Fore," *SIAM News*, November 1993.
- Cisco Systems, Incorporated, *Technology Brief: Asynchronous Transfer Mode (ATM)*, Cisco Systems, Inc., 1993.
- Cisco Systems, Incorporated, *Product Information: AGS+*, Cisco Systems, Inc., April 1994.
- Connectivity and Network Committee, *Minutes of Committee Meeting*, Naval Postgraduate School, September 1987.
- Corning Incorporated, *Corning™ Optical Fiber*, March 1992.
- Crawford, Tim, "Super Hubs," *Infoworld*, v. 16, nr. 15, p. 60(9), April 11, 1994.
- Danielle, Diane, "Cable Awareness," *Network Computing*, v. 4, nr. 10, p. 110(4), October 15, 1993.
- Datamation*, "Networking In The '90's: The Infrastructure Emerges," *Datamation*, v. 39 nr. 20 p. N1(6), October 15, 1993.
- Davie, Bruce S., "Architecture And Implementation Of A High-speed Host Interface," *IEEE Journal On Selected Areas In Communications*, v. 11, nr. 2, p. 228(12), February 1993.
- Department of the Navy, Office of the Chief of Naval Operations, Letter Ser. 00/4U500182, Subject: Graduate Education Policy, July 27, 1994.
- Dix, John, "Collapsed Backbone Paves Way For Imaging System," *Network World*, v. 11, nr. 14, p. 16+, April 4, 1994.
- Durr, Michael, "Managing The Switch To Virtual Networks," *Datamation*, v. 40, nr. 11, p. 65(3), June 1, 1994.
- Durr, Michael, "The Virtual Reality Of Switched Networks," *Datamation*, v. 40 nr. 9, p. 34 (3), May 1, 1994..

- Farris, Jeffrey L., *Saber LAN Setup Guide*, 2nd ed., Saber Software Corporation, 1993.
- Feldman, Robert, "Ramping Onto The ATM Superhighway," *LAN Times*, v. 11, nr. 2, p. 61(4), January 24, 1994.
- Feltman, Charles, "ATM: The Grand Unifier," *LAN Magazine*, v. 8, nr. 12, p. 52(12), November 1994.
- Fraser, A. G., and Henry, P. S., "Transmission Facilities for Computer Communications," *Computer Communication Review*, v. 22, nr. 5, p. 53(9), October 1992.
- Freedman, Alan, *Electronic Computer Glossary*, The Computer Language Co. Inc., 1993.
- Freeman, Roger L., *Telecommunication Transmission Handbook*, 3rd ed., John Wiley & Sons, Inc., 1991.
- Gentry, Terry, W. R. Church Computer Center, Naval Postgraduate School, Monterey, California, technical insights and practical applications of network cabling and distribution devices, March-September 1994.
- Gibbs, Mark, "When Network Computing Is Really Not Work Computing," *Network World*, v. 11, nr. 24, p. 26, June 13, 1994.
- Gillooly, Caryn, "ATM A Lifesaver For UVA Medical Center," *Network World*, v. 11, nr. 11, p. 1+, March 14, 1994.
- Haight, Timothy, "Growing Pains," *Network Computing*, v. 4, nr. 3, p. 86(5), March 1993.
- Halwachs, Tom, Naval Postgraduate School, Monterey, California, Interview with the author, April 1994.
- Hamilton, Drew, "Network Monitoring, Management and Simulation," *Army Research, Development and Acquisition Bulletin*, p. 55, July-August 1994..
- Hart, J. A., W. R. Church Computer Center, Naval Postgraduate School, Monterey, California, interviews and demonstrations with the author, July-August 1994.
- Haugdahl, J. Scott, "Switch Trio Boosts Bandwidth," *Network World*, v. 11, nr. 30, p. 47(4), July 25, 1994.
- Hewlett-Packard Company, "A Stackable Architecture For An Open Network Future," *PC Networking's New Dimensions*, a supplement to *Network World*, v. 11, nr. 28, p. S16(2), July 11, 1994.
- Hewlett-Packard Company, "Improving Performance With Integrated 32-Bit Networking," *PC Networking's New Dimensions*, a supplement to *Network World*, v. 11, nr. 28, p. S12(2), July 11, 1994.

- Hewlett-Packard Company, "Making Fast Ethernet Real," *PC Networking's New Dimensions*, a supplement to *Network World*, v. 11, nr. 28, p. S16(1), July 11, 1994.
- Horwitt, Elisabeth, "TCP/IP To Become Leading Protocol." *Computerworld*, v. 28 nr. 21. p. 69+, May 23, 1994.
- Howard, Matthew D., Manager, Systems Engineering, Cisco Systems, Inc., phone interview with the author, August 3, 1994.
- Infonetics Research, Inc., *Server Bandwidth Congestion--Practical Limits; 10Mbps Versus 100Mbps Solutions*, Infonetics Research, Inc., 1994.
- Intel Corporation, *Introduction to Fast Ethernet and EtherExpress™ PRO/100 Adapters*, white paper, 1994.
- Jazwinski, Andy, "Is Emulation Adequate For Traffic Analysis?" *Network World*, v. 11, nr. 14, p. L12, April 4, 1994.
- Jones, Hap, "Data Cabling Infrastructure," *Cabling Business*, v. 2, nr. 9, p. 8(7), November/December 1992.
- Katzan, Harry, Jr., *Local Area Networks: A Manager's Guide*, Carnegie Press, Inc., 1983.
- Klessig, Bob, "Integrating ATM Into Legacy Networks," presentation given at a conference entitled *Implementing ATM*, Santa Clara, California, May 16, 1994.
- Klett, Stephen P., "Health Care Looks To ATM For Imaging Apps," *Computerworld*, v. 28, nr. 29, p. 60, July 18, 1994.
- Klett, Stephen P., Jr., "Router Software Allocates Bandwidth." *Computerworld*, v. 28, nr. 20, p. 16, May 16, 1994.
- LAN Magazine*, "Cal Poly, IBM Create Virtual Classroom," *LAN Magazine*, v. 9, nr. 4, p. 12-14, April 1994. (no author listed)
- Leahy, Kevin M., *The Proposed Naval Postgraduate School Campus Network: Computer Communications For The 1990's*, Master's Thesis, Naval Postgraduate School, Monterey, California, March 1988.
- Leslie, Ian M., McAuley, Derek R., and Tennenhouse, David L., "ATM Everywhere?" *IEEE Network*, v. 7, nr. 2, p. 40(6), March 1993.
- Lippis, Nick and Morency, John, "Reducing Router Network Transmission Costs," *BYTE*, v. 18, nr. 13, p. 70(2), December 1993.
- Loudermilk, Stephen, "Hospital Makes Change To Switched LAN Network," *LAN Times*, v. 11, nr. 11, p. 51(2), June 13, 1994.

- MacAskill, Skip, "ATM Provides High-speed Highway To Success For Engineering Company," *Network World*, v. 11, nr. 26, p. 1(3), June 27, 1994.
- MacAskill, Skip, "Hospital Lays Groundwork For Network," *Network World*, v. 11, nr. 25, p. 25(3), June 20, 1994.
- MacAskill, Skip, "New ATM Options Move To The Fore-Front," *Network World*, v. 11, nr. 14, p. 1(3), April 4, 1994.
- MacAskill, Skip, "Of Doctors, Kids And ATM," *Network World*, v. 11, nr. 26, p. 1(2), June 27, 1994.
- MacAskill, Skip, "Users Bank On Switching To Cure Bandwidth Woes," *Network World*, v. 11, nr. 14, p. L2, April 4, 1994.
- MacAskill, Skip, "Users Warn Of Moving Too Fast To High-Speed LANs," *Network World*, v. 11, nr. 19, p. 5, May 9, 1994.
- Marks, Howard, "Cable Vision," *LAN Magazine*, v. 8, nr. 4, p. 77(13), April 1993.
- Masud, S. A., "Navy Multimedia Flies On Fiber," *Government Computer News*, p. 37(3), February 7, 1994.
- Masud, S. A. and Menke, Susan M., "Give ATM Time To Mature, NASA Engineer Advises," *Government Computer News*, p. 8, February 7, 1994.
- Masud, Sam, "Hold Out Just a Little Longer On Those ATM Purchases," *Government Computer News*, p. 41(2), May 2, 1994.
- McCabe, Ron, "Tech Talk: Can Your Hard Disk Subsystem Ride The Local Bus?" *Computer Shopper*, v. 13, nr. 12, p. 621(6), December 1993.
- McClimans, Fred J., "ATM vs. FDDI: Comparing Functionality, Features, and Benefits," *Telecommunications*, v. 28, nr. 2, p. 27(5), February 1994.
- McNutt, Barbara, "SCSI And Beyond," *BYTE*, v. 19, nr. 8, p. 111(4), August 1994.
- McQuay, D. S., Stone, L. L., and Trepanier, D. M., "Checking Aboard," paper presented for IS 2000 class, Naval Postgraduate School, Monterey, California, Fall 1993.
- Medici, Mark, *Ethernet Network Questions And Answers*, version 2.11, Internet ftp site: dorm.rutgers.edu, pub/novell/DOCS, April 19, 1993.
- Metcalf, R. M., "Computer/Network Interface Design," *IEEE Journal On Selected Areas In Communications*, v. 11, nr. 2, p. 173(7), February 1993.
- Miller, Mark, "LAN Management Tools: A Perplexing Choice," *Network World*, v. 11, nr. 23, p. 48, March 21, 1994.
- Miller, Mark, "Life On The Fast LAN," *Network World*, v. 11, nr. 6, p. 41(6), March 21, 1994.

- Minoli, Daniel, "Designing Scalable Networks," *Network World*, v. 11, nr. 2, p. 11(8), January 10, 1994.
- Molta, Dave, "Managing LAN Growth," *Network Computing*, v. 4, nr. 10, p. 6(9), October 15, 1993.
- Moore, Steve, "Network Management: How Much Is Too Much?," *Computerworld*, v. 28, nr. 27, p. 56, July 4, 1994.
- Moore, Steve, "Software-Based Virtual LANs Raise Questions," *Computerworld*, v. 28, nr. 24, p. 85, June 13, 1994.
- Morency, John, Lippis, Nick, and Hindin, Eric, "Virtual LANs Promise Big Gains, Pose New Risks," *Network World*, v. 11, nr. 18, p. 1(3), May 2, 1994.
- Mulqueen, John T., "A Study In Fiber," *Communications Week*, nr. 500, p. 1(2), April 11, 1994.
- Nelson, Mark, "The Squeeze On Data," *LAN Magazine*, v. 9, nr. 5, p. 129(4), May 1994.
- Newman, Michael, Newman and Associates, interview with the author, July 7, 1994.
- Newman, Peter, "ATM Local Area Networks," *IEEE Communications Magazine*, v. 32, nr. 3, p. 86(13), March 1994.
- Newton, Harry, *Newton's Telecom Dictionary*, Harry Newton, 1993.
- Nicholson, Andy, "HIPPI And High Performance LANs; Supercomputer Technologies For High Performance," *Dr. Dobbs Journal*, v. 18, nr. 6, p. 58(7), June 1993.
- Nolle, Thomas, "Too Many Choices Cloud LAN Technology Direction," *Network World*, v. 11, nr. 14, p. 90, April 4, 1994.
- Norman, Dave F., Naval Postgraduate School, Director of Academic and Administrative Computing, interview with the author, July 18, 1994.
- Olsen, Florence, "Switching Hubs Pour On the Speed; Although The Devices Have the Same Purpose, No Two Are Exactly Alike," *Government Computer News*, v. 12, nr. 17, p. 55(2), August 16, 1993.
- One Network Place, *Networking Products Catalog*, Fall 1993.
- Panettieri, Joseph, "When Fast Is Not Fast Enough," *Informationweek*, p. 26(6), April 18, 1994.
- Paris, Fred, "Wiring The Schools: K-12 To The University," *Cabling Business*, v. 4, nr. 6, p. 48(11), June, 1994.
- Parnell, Tere, "Capacity Planning: Rules And Reality," *Netware Technical Journal*, v. 2, nr. 2, p. 7(5), March-April 1994.

- Poutain, Dick, "A Different Kind Of RISC," *BYTE*, v. 19, nr. 8, p. 185(2), August 1994.
Report of the Committee on the Status of Computing at NPS, July 27, 1994.
- Ricciuti, Mike, "PCI: Server Bus For The Fast LANe," *Datamation*, v. 40, nr. 6, p. 61(4), March 15, 1994.
- Rohde, David, "Time May Be Right For ISDN," *Network World*, v. 11, nr. 14, p. 1(3), April 4, 1994.
- Romo, Raul, Computer Operations and Network Services, W. R. Church Computer Center, Naval Postgraduate School, Monterey, California, technical insights and practical applications of network cabling and distribution devices, March-September 1994.
- Sant'Angelo, Rick, "Understanding LAN Bandwidth," *Netware Technical Journal*, v. 2, nr. 2, p. 14(4), March-April 1994.
- Schnaidt, Patricia, "Plug In At 100," *LAN Magazine*, v. 9, nr. 3, p. 71(8), March 1994.
- Schnaidt, Patricia, "Switch Hunt," *LAN Magazine*, v. 9, nr. 6, p. 75(6), June 1994.
- Schoenstadt, Arthur L., numerous interviews, advice and discussions with the author, March-September 1994.
- Seitz, Neal B., et al, "User-Oriented Measures of Telecommunication Quality," *IEEE Communications Magazine*, v. 32, nr. 1, p. 56(11), January 1994.
- Serjak, Christopher W., "Migrating To ATM: Four Scenarios To Take Users From Today's Networks To Switched Gigabit 'Fabrics'," *Distributed Computing Monitor*, v. 8, nr. 9, p. 3(16), September 1993.
- Shannon, C. E., "A Mathematical Theory Of Communication," *Bell Systems Technical Journal*, nr. 27, pp. 379 and 623, 1948.
- Sirbu, Marvin, *Bandwidth Requirements For Different Applications As A Function Of Time And Size Of Data Transfer*, Carnegie Mellon, 1994.
- Skorupa, Joseph, "Tough Choices Ahead In Hub Selection," *Network World*, v. 11, nr. 18, p. 80(8), May 2, 1994.
- Smith, Ben, and Udell, Jon, "Linking LANs," *BYTE*, v. 18, nr. 13, p. 66(24), December 1993.
- Smith, J. M., and others, "Editorial: High-Speed Computer/Network Interfaces," *IEEE Journal On Selected Areas In Communications*, v. 11, nr. 2, p. 169(3), February 1993.
- Spada, Mary E., "Does ATM Make Good Business Sense Now?" *LAN TIMES*, v. 11, nr. 13, p. 56(1), July 11, 1994.
- Sprague, Ralph H., Jr., and McNurlin, Barbara C., *Information Systems Management In Practice*, Third Edition, Prentice-Hall, Inc., 1993.

Stallings, William, *Handbook of Computer-Communications Standards, Volume 2, Local Network Standards*, Macmillan Publishing Company, 1987.

Stallings, William, *Data and Computer Communications*, 4th ed., Macmillan Publishing Company, 1994.

Strassman, Paul, "Centralization's Payback," *Computerworld*, v. 28, nr. 23, p. 35(1), June 6, 1994.

Strauss, Paul, "Plan Your 2001 Network Now," *Datamation*, v. 40, nr. 15, p. 32(6), August 1, 1994.

Strauss, Paul, "Seize Control Of Your 2001 Network," *Datamation*, v. 40, nr. 15, p. 35(3), August 1, 1994.

Stuck, B. W., and Arthurs, E., *A Computer Communication Network Performance Analysis Primer*, Prentice-Hall, 1985.

Stuckey, John F., "Link Loss Budget For UTP," *Cabling Business*, v. 2, nr. 9, p. 46(4), November/December 1992.

Tibbets, John, and Bernstein, Barbara, "Political Primer For Enterprise Networks," *BYTE*, v. 18, nr. 13, p. 78(2), December 1993.

Traw, C. B. and Smith, J. M., "Hardware/Software Organization Of A High-Performance ATM Host Interface," *IEEE Journal On Selected Areas In Communications*, v. 11, nr. 2, p. 240(14), February 1993.

Tuck, Edward F., et al, *The Calling™ Network: A Global Telephone Utility*, Calling Communications Corporation, 1993.

Varhol, Peter D., "Applications Influence Network Design," *Information World Direct*, p. 51(1), November 1992.

Varhol, Peter D., "The Hardware Mix," *Information World Direct*, p. 50(6), November 1992.

Vis, Jan, "A Simple LAN Performance Measure," *Computer Communication Review*, v. 24, nr. 1, p. 7(5), January 1994.

Wallace, Bob, "Fast Ethernet's Luster Is Fading Fast," *Infoworld*, v. 16, nr. 26, p. 53(2), June 27, 1994.

Wargo, Linda, Total Quality Leadership Office, Naval Postgraduate School, Monterey, California, Interview with the author, September 6, 1994.

Weston, Rusty, "Your Results Will Vary," *Open Computing*, v. 11, nr. 9, p. 43(9).

Wittman, Art, "Can We Run Just TCP/IP?" *Network Computing*, v. 5, nr. 2, p. 78(2), February 15, 1994.

Wittman, Art, "Don't Let Protocol Madness Overwhelm You," *Network Computing*, v. 5, nr. 2, p. 66(6), February 15, 1994.

Wittman, Art, "Order Out Of Chaos," *Network Computing*, v. 4, nr. 3, p. 70(8), March 15, 1993.

Wittman, Art, "The ATM Experience: Still Smoke And Mirrors," *Network Computing*, v. 5, nr. 3, p. 44(1), March 1, 1994.

Yager, Tom, "How I Learned To Live With Benchmarks," *Open Computing*, v. 11, nr. 9, p. 56(4), September 1994.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center 2
Cameron Station
Alexandria, VA 22304-6145
2. Library, Code 52 2
Naval Postgraduate School
Monterey, CA 93943-5002
3. Dean of Computer and 2
Information Services, Code 05
Naval Postgraduate School
Monterey, CA 93943
4. Director of Academic and 2
Administrative Computing Services, Code 051
Naval Postgraduate School
Monterey, CA 93943
5. Director of Military Operations, Code 04 1
Naval Postgraduate School
Monterey, CA 93943
6. Professor Arthur Schoenstadt 5
Department of Mathematics, Code MA/Zh
Naval Postgraduate School
Monterey, CA 93943
7. Professor James C. Emery 1
Department of Systems Management, Code SM/Ey
Naval Postgraduate School
Monterey, CA 93943-5103
8. Professor Barry Frew 1
Department of Systems Management, Code SM/Fw
Naval Postgraduate School
Monterey, CA 93943-5103

9. Dr. Ted Lewis 1
 Department of Computer Science
 Naval Postgraduate School
 Monterey, CA 93943

10. Professor Rex Buddenburg 1
 Department of Systems Management
 Naval Postgraduate School
 Monterey, CA 93943-5103

11. Mr. Raul Romo 1
 Network Operations, Code 51
 Naval Postgraduate School
 Monterey, CA 93943

12. Ms. Terry Gentry 1
 Network Operations, Code 51
 Naval Postgraduate School
 Monterey, CA 93943

13. Commandant (G-L) 1
 U. S. Coast Guard
 2100 2nd St., S.W.
 Washington, DC 20593-0001

14. Commandant (G-TP-2) 1
 U. S. Coast Guard
 2100 2nd St., S.W.
 Washington, DC 20593-0001

15. United States Military Academy 1
 Directorate of Information Management
 Business Operations and Plans Division
 Attn: CPT Paul R. Logan
 West Point, NY 10996

16. LCDR Paul Wiedenhoeft, USCG 2
 c/o Commandant (G-CCS-2)
 U. S. Coast Guard
 2100 2nd St., S.W.
 Washington, DC 20593-0001