

NAVAL POSTGRADUATE SCHOOL MONTEREY, CALIFORNIA



THESIS

**AN FDDI-BASED SOLUTION
FOR THE
SYSTEMS MANAGEMENT DEPARTMENT
COMPUTER LABORATORY NETWORK**

by

Carlos M. Chavez

March, 1996

Thesis Advisor:
Associate Advisor:

Norman F. Schneidewind
James Emery

Approved for public release; distribution is unlimited.

19960801 085

DTIC QUALITY INSPECTED 1

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing 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 March 1996		3. REPORT TYPE AND DATES COVERED Master's Thesis
4. TITLE AND SUBTITLE AN FDDI-BASED SOLUTION FOR THE SYSTEMS MANAGEMENT DEPARTMENT COMPUTER LABORATORY NETWORK			5. FUNDING NUMBERS	
6. AUTHOR(S) Carlos M. Chavez				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) FDDI is one of the latest evolutions in shared-media technology. Originally intended as a high-speed backbone for interconnecting networks, it has become a viable alternative for organizations that seek better response time and bandwidth capacity from their local area networks (LANs). However, this fiber-based standard is an expensive departure from the more familiar, and perhaps more mature, IEEE 802 token-ring and Ethernet standards. Thus, developing an FDDI-based network may present considerable economic and technical risk to an organization. This study examines the application of FDDI technology as an upgrade to the Systems Management Department's token-ring network. It reviews the protocols that comprise the standard, addresses design considerations for developing an FDDI network, evaluates the existing token-ring LAN, and proposes an FDDI solution. This study concludes that the risks of implementing an FDDI-based upgrade, can be mitigated using an evolutionary design strategy.				
14. SUBJECT TERMS FDDI; Fiber Optic Communications; Token-Ring; Network Design; Network Topology; Local Area Network			15. NUMBER OF PAGES 148	
			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. Z39-18 298-102

Approved for public release; distribution is unlimited.

**AN FDDI-BASED SOLUTION FOR THE
SYSTEMS MANAGEMENT DEPARTMENT
COMPUTER LABORATORY NETWORK**

Carlos M. Chavez
Lieutenant Commander, United States Navy
B.S., United States Naval Academy, 1981

Submitted in partial fulfillment
of the requirements for the degree of

**MASTER OF SCIENCE IN INFORMATION TECHNOLOGY
MANAGEMENT**

from the

NAVAL POSTGRADUATE SCHOOL

March 1996

Author:




Carlos M. Chavez


Approved by:



Norman F. Schneidewind, Thesis Advisor



James Emery, Associate Advisor



Reuben T. Harris, Chairman, Department of Systems Management

ABSTRACT

FDDI is one of the latest evolutions in shared-media technology. Originally intended as a high-speed backbone for interconnecting networks, it has become a viable alternative for organizations that seek better response time and bandwidth capacity from their local area networks (LANs). However, this fiber-based standard is an expensive departure from the more familiar, and perhaps more mature, IEEE 802 token-ring and Ethernet standards. Thus, developing an FDDI-based network may present considerable economic and technical risk to an organization.

This study examines the application of FDDI technology as an upgrade to the Systems Management Department's token-ring network. It reviews the protocols that comprise the standard, addresses design considerations for developing an FDDI network, evaluates the existing token-ring LAN, and proposes an FDDI solution. This study concludes that the risks of implementing an FDDI-based upgrade, can be mitigated using an evolutionary design strategy.

TABLE OF CONTENTS

I. INTRODUCTION	1
A. BACKGROUND	1
B. PURPOSE	2
C. SCOPE	3
D. METHODOLOGY	3
E. ORGANIZATION OF THESIS	4
II. FIBER OPTIC COMMUNICATIONS	5
A. INTRODUCTION	5
B. OPTICAL COMMUNICATIONS	5
1. Optical Communications Link	5
2. Fiber Optic Light Propagation	6
3. Fiber Modes	8
4. Attenuation and Dispersion	9
a. Modal Dispersion	9
b. Chromatic Dispersion	10
C. CHARACTERISTICS OF LINK COMPONENTS	11
1. Optical Transmitters	11
2. Optical Receivers	12
3. Fiber Types	12
D. SUMMARY	13
III. FDDI FUNDAMENTALS	15
A. INTRODUCTION	15
B. INTRODUCTION TO FDDI PROTOCOLS	15
1. OSI Reference Model	15
2. FDDI Protocols	17
C. FDDI TOPOLOGY	18
1. Nodes	18

2.	Protocol Structure of Nodes	21
3.	Ports	21
4.	Connection Rules	23
5.	Ring Reconfiguration	24
6.	Optical Bypass	26
D.	SUMMARY	28
IV. FDDI PROTOCOLS		29
A.	INTRODUCTION	29
B.	PHYSICAL MEDIUM DEPENDENT STANDARDS	29
1.	Multimode Fiber PMD (MMF-PMD)	30
a.	Fiber Link	30
b.	Transmitters and Receivers	31
c.	Media Connectors	32
d.	Mixing Fiber Types and Sizes	33
e.	Optical Bypasses	34
2.	Low-cost Fiber PMD (LCF-PMD)	34
a.	Fiber Link	36
b.	Transmitters and Receivers	36
c.	Media Connectors	37
3.	Twisted Pair PMD (TP-PMD)	37
a.	Multilevel Transmission-3 (MLT-3), Scrambling, and Equalization	37
b.	Media Specifications	38
c.	UTP Installation Considerations	38
C.	PHYSICAL MEDIUM INDEPENDENT STANDARDS	39
1.	4B/5B and NRZI Coding	39
2.	Elasticity Buffer	41
3.	Smoother	42
4.	Repeat Filter	43
D.	MEDIA ACCESS CONTROL STANDARDS	44
1.	Timed Token Access Protocol	45
a.	Claim Process	45
b.	Valid Transmission Timer	46
c.	Beacon Process	46
d.	Management of Asynchronous Traffic	46
e.	Management of Synchronous Traffic	48
2.	Frames	49

a.	Types of Frames	51
b.	Addressing Format	52
3.	Frame Stripping	52
E.	STATION MANAGEMENT STANDARDS	53
1.	Connection Management	53
2.	Ring Management	54
3.	Frame-Based Management	55
a.	Neighbor Information Frame	55
b.	Status Information Frame	56
c.	Echo Frame	57
d.	Resource Allocation Frame	57
e.	Request Denied Frame	57
f.	Status Report Frame	58
g.	Parameter Management Frame	58
h.	Extended Service Frame	58
F.	SUMMARY	58
V.	NETWORK DESIGN	61
A.	INTRODUCTION	61
B.	STRATEGIC PLANNING	61
C.	DEFINING THE PROBLEM AND FINDING A SOLUTION	62
1.	Determining the Feasibility of a Solution	62
2.	Determining System Requirements	63
3.	Defining the Geographical Scope	64
4.	Selecting a Network Standard and Defining Its Topology	65
5.	Documenting Hardware and Software Requirements	65
6.	Calculating Network Costs	66
7.	Determining the Feasibility of a Network Solution	66
D.	METHODOLOGY FOR DESIGNING AN FDDI NETWORK TOPOLOGY	67
1.	Fiber-based Design	67
2.	Copper to the Desktop Level	76
3.	Additional Design Issues	77
E.	SUMMARY	77

VI. FINDING A SOLUTION FOR THE SYSTEMS MANAGEMENT DEPARTMENT	79
A. INTRODUCTION	79
B. INTRODUCTION TO SMD NETWORKS	79
C. JUSTIFICATION FOR AN FDDI SOLUTION	80
D. EVOLUTIONARY VERSUS REVOLUTIONARY DESIGN	81
1. Technical Risk	82
2. Limited Staffing	83
3. Limited Budget	83
4. Preferred Approach	83
E. BASELINE REVIEW OF THE TOKEN-RING NETWORK	84
1. Topology	84
2. Integrated Network	84
3. Token-Ring Segment 8TR	86
4. Token-Ring Segment 4TR	88
5. Token-Ring Segment 0TR	88
6. Cable Routing	91
7. Physical Security	92
F. SUMMARY	92
VII. THE TARGET SOLUTION	93
A. INTRODUCTION	93
B. PRELIMINARY LINK DESIGN DECISIONS	93
C. TARGET SOLUTION	95
1. Trunk Design	95
2. IN-158 Tree Design	96
3. IN-224 Tree Design	99
4. IN-250 Tree Design	101
5. Hardware Components	101
6. Software Components	103
D. COMPATIBILITY ISSUES	104
1. Optical Cables and Connectors	105
2. Adapter Cards	106

3.	Concentrators	108
E.	EVOLUTIONARY DEVELOPMENT OF THE TARGET SOLUTION	110
1.	Stage 1: Two-Node Network	110
2.	Stage 2: Introduction of a Dual-Attachment Concentrator	112
3.	Stage 3: Installation of the Fiber Trunk	113
4.	Stage 4: Migration of the Computer Labs	117
5.	Stage 5: Target Implementation	118
F.	SUMMARY	121
VIII.	CONCLUSION	123
	APPENDIX: COST COMPARISON BETWEEN FIBER, STP, AND UTP	125
	LIST OF REFERENCES	127
	INITIAL DISTRIBUTION LIST	129

LIST OF FIGURES

2-1. Fiber Optical Communications Link	5
2-2. Maximum Entrance Angle	7
2-3. Dispersion Effect on Signal Pulses	10
3-1. OSI Reference Model	16
3-2. FDDI Standards Comparison to the OSI Reference Model	17
3-3. FDDI Topology	19
3-4. Types of Concentrators	20
3-5. Example of an FDDI Network	20
3-6. Port Designations	22
3-7. Normal Ring Operation	25
3-8. Reconfiguration After Cable Fault	25
3-9. Reconfiguration After Second Fault	25
3-10. Reconfiguration After Station Power-down	26
3-11. Optical Bypass Analysis	27
4-1. Media Interface Connector	33
4-2. MMF-PMD Connector Keying	34
4-3. Simplex Connectors	34
4-4. Standard FDDI Encoding versus MLT-3 Coding	38
4-5. Elasticity Buffer	42
4-6. Repeat Filter Operation	43
4-7. FDDI Frame	49
4-8. Frame Status Field	50
4-9. Frame Control Field	51
4-10. Address Fields	52
5-1. Conceptual Design Example	69
5-2. Logical Topology Example	70
5-3. Alternative Tree Configurations	71
5-4. Interconnecting Stations using Patch Panels	71
5-5. Physical Topology Diagram	73
6-1. Systems Management Department's Token-Ring LAN	85
6-2. Token-Ring Segment 8TR	87
6-3. Token-Ring Segment 4TR	89
6-4. Token-Ring Segment 0TR	90
6-5. Token-Ring Cable Layout	91
7-1. Trunk Design	97
7-2. IN-158 Tree Design	98
7-3. Dual-Homed Configuration for IN-158 Servers	99
7-4. IN-224 Tree Design	100
7-5. IN-250 Tree Design	102
7-6. Stage 1 of Network Evolution	111
7-7. Stage 2 of Network Evolution	113

7-8. Stage 3: Fiber Cable Installation	115
7-9. Sample OTDR Plot	117
7-10. Stage 4 of Network Evolution	119

LIST OF TABLES

2-1. Fiber Optic Wavelength Bands	6
2-2. LEDs versus Laser Diodes	11
2-3. PIN Photodiodes versus APDs	12
3-1. FDDI Connection Rules	23
4-1. Multimode Fiber PMD Specifications	30
4-2. Impact of Using Alternative Fiber Types with a 62.5/125 Transceiver	32
4-3. Losses (dB) Due to Mixing Multimode Fiber Types	35
4-4. Low-Cost Fiber PMD Specifications	35
4-5. Impact of Using Alternative Fiber Types	36
4-6. 4B/5B Encoding	40
4-7. Information Contained in an SIF Response	57
5-1. Network Design Methodology	63
5-2. FDDI Design Methodology	68
5-3. Simple Loss Calculation Method	74
5-4. Determining Losses Using Average and Standard Deviations	75
7-1. Preliminary Link Decisions	94
7-2. Target Solution Hardware Components	103
7-3. Fiber Cable Considerations	106
7-4. Adapter Card Considerations	108
7-5. Concentrator Considerations	109
7-6. Stage 1 Component Requirements	112
7-7. Stage 2 Component Requirements	114
7-8. Stage 4 Component Requirements	120

LIST OF SYMBOLS, ACRONYMS, AND ABBREVIATIONS

$\theta_{\max}$	Maximum Entrance Angle
V	Normalized Frequency
ANSI	American National Standards Institute
APD	Avalanche Photodiode
ATM	Asynchronous Transfer Mode
BN	Backbone Network
bps	Bits per second
CFM	Configuration Management
CMT	Connection Management
CRC	Cyclic Redundancy Check
DAC	Dual-attachment Concentrator
DAS	Dual-attachment Station
DOD	Department of Defense
dB	Decibel
ECF	Echo Frame
ECM	Entity Coordination Management
EIA	Electronic Industries Association
EMI	Electromagnetic Interference
FCC	Federal Communications Commission
FDDI	Fiber Distributed Data Interface
FIFO	First-in-first-out
GUI	Graphical User Interface
IEEE	Institute of Electrical and Electronic Engineers
ILD	Injection Laser Diode
ISO	International Standards Organization
LAN	Local Area Network
LCF	Low-cost Fiber
LED	Light Emitting Diode
LLC	Logical Link Control
MAC	Media Access Control
MAN	Metropolitan Area Network
MAU	Multistation Access Unit
MIB	Management Information Base
MIPS	Million instructions per second
MIC	Media Interface Connector
MLT	Multilevel Transmission
MMF	Multimode Fiber
NA	Numerical Aperture
NAC	Null-attachment Concentrator
NIF	Neighbor Information Frame

NOS	Network Operating System
NPS	Naval Postgraduate School
NRZI	Non-return-to-zero-inverted
OSI	Open Systems Interconnection
OTDR	Optical Time-Domain Reflectometer
PCM	Physical Connection Management
PHY	Physical Medium Independent
PMD	Physical Medium Dependent
PMF	Parameter Management Frame
ppm	Parts per million
R	Reset
RAF	Resource Allocation Frame
RDF	Request Denied Frame
RMT	Ring Management
S	Set
SAC	Single-attachment Concentrator
SAFENET	Survivable Adaptable Fiber Optic Embedded Network
SAS	Single-attachment Station
SIF	Status Information Frame
SMD	Systems Management Department
SMT	Station Management
SNMP	Simple Network Management Protocol
SONET	Synchronous Optical Network
SRF	Status Report Frame
STP	Shielded Twisted-pair
TCP/IP	Transmission Control Protocol/Internet Protocol
THT	Token Holding Time
TP	Twisted-Pair
TRT	Token Rotation Time
TT	Token Rotation Timer
TTRT	Target Token Rotation Time
TVX	Valid Transmission Timer
UTP	Unshielded Twisted-pair
WAN	Wide Area Network
WFWG	Windows for Workgroups

I INTRODUCTION

A. BACKGROUND

Local area networks (LANs) have undergone tremendous development in the past decade. The technology has evolved from simple computer-to-computer interfaces, to networks connecting hundreds of nodes sharing hardware, software, and databases. Both commercial and government agencies have realized the tremendous benefits of LANs, propagating their use beyond the traditional company office or campus network. Computing resources are now interconnected to form backbone networks (BNs), metropolitan area networks (MANs), and wide area networks (WANs).

Two of the biggest challenges to a network's performance is its ability to handle bandwidth intensive applications and a large number of simultaneous users. These problems became evident in WANs that connected dozens of networks and literally hundreds of nodes. The increased demand for capacity quickly consumed all available bandwidth afforded by copper-based media. To alleviate this problem, network designers turned to a new medium--fiber optic.

The use of fiber optics has also undergone a growth explosion in the past decade. Communications companies have realized the advantages of fiber optics and have invested heavily in replacing copper-based circuits with this medium. These advantages include its immunity to electrical and magnetic interference, the added security due to its resistance to unauthorized tapping, and its reduced size and weight--20 times lighter and 5 times smaller than equivalent copper cable. Certainly its biggest advantage is its tremendous bandwidth, enabling it to achieve impressive digital data rates with near error-free transmission.

Until recently, the cost of fiber technology made it unsuitable for all but the highest performance communications systems and networks. Advances in technology have reversed this trend. Improvements in fiber optic quality, light sources, light detectors, and the procedures used to splice circuits have reduced the cost of implementing this medium. Fiber optics is now a feasible solution to the bandwidth limitation problems on LANs.

To meet the demand for optical support in LANs, the Fiber Distributed Data Interface (FDDI) standard was established. The standard was developed by the American National Standards Institute (ANSI), and represents an evolutionary progression in shared-media technology. It is based on a token-passing ring architecture that employs timed-token rotation to exchange information at a rate of 100 Mbps. Coupled with near error-free transmission rates of only 1 in 1,000 million bits, FDDI is particularly attractive for handling emerging bandwidth-intensive applications such as video conferencing and telemedicine.

The logical topology of an FDDI network consists of a primary and secondary ring that interconnect wiring concentrators and stations. Normally, data is transmitted on the primary ring while the secondary remains in a standby condition. In the event of a link or node failure, traffic is automatically routed onto the secondary ring, in a counter rotating direction. This preserves the network's functionality.

This self-healing capability captured the attention of ship-board network designers. FDDI standards were used to develop the Survivable Adaptable Fiber Optic Embedded Network (SAFENET) military standard. FDDI-based networks are now being installed aboard Navy aircraft carriers and Coast Guard buoy tenders. Specifically, the Coast Guard is using SAFENET standards to bridge meteorological, navigation, engineering, and data processing systems onto an FDDI backbone. This enables high speed data communications between a myriad of ship-board systems, and between standard workstations. (Hewell, 1994)

Unfortunately, FDDI standards have failed to revolutionize the LAN market. Despite the reduced cost of optical components, businesses have preferred the less-costly copper-based alternatives over FDDI. Realizing this fact, ANSI recently approved a standard for FDDI transmission across copper. Capable of supporting 100 Mbps capacity at less cost, this new standard may provide the impetus for world-wide growth of the technology.

B. PURPOSE

The purpose of this thesis is to provide a recommendation for replacing the Naval Postgraduate School's Systems Management Department's *Token-Ring-based* network with FDDI standards. Specifically, it will examine the hardware, software, and maintenance requirements necessary to implement FDDI. In addition, it will examine the compatibility between FDDI and

other LAN technologies. The goal is to provide a recommendation for migrating the department's current networks to this mature and proven standard.

It is also intended that this thesis serve as a reference for other organizations that are considering a similar endeavor. It will be written to serve as an example for evaluating an existing LAN technology with the objective of migrating to FDDI standards.

C. SCOPE

The principle goal of this research is to provide a recommendation for implementing FDDI as a supplement or replacement to the Systems Management Department's token-ring network. To support this goal, this research will involve an in-depth review of FDDI standards, a baseline assessment of the current networks, and the preparation of an FDDI-based solution.

D. METHODOLOGY

Designing and implementing a fiber-based network is considerably different from developing one of the more traditional copper-based technologies. For this reason, research began with a focus on fiber-optic communications systems. Indeed, developing and maintaining a fiber-based network requires an understanding of the benefits, limitations, and unique components of fiber-optical communications systems.

Next, research focused on the fundamentals of FDDI to gain a better understanding of the technology's basic topology, configuration, and operation. Furthermore, it exposed the unique terminology used to describe FDDI networks. This was followed by an in-depth review of FDDI protocols to determine the issues that govern design and implementation of these networks. When feasible, the International Standards Organization's Open Systems Interconnection reference model was used to evaluate FDDI protocols in comparative terms to other technologies.

To prepare an FDDI solution, the educational goals of the Systems Management Department were examined and a baseline review of the existing networks was conducted. The baseline assessment entailed a complete review of all network hardware and software in order to identify critical elements, network weaknesses, compatibility conflicts, and the geographical scope of existing technologies. This assessment was crucial to the design of an FDDI solution.

Lastly, an FDDI solution was prepared based on the research findings and baseline assessment results. This recommendation represents an evolutionary approach to developing an FDDI solution. This approach was deemed necessary to minimize the risk to the stakeholders.

E. ORGANIZATION OF THESIS

The remainder of this thesis is organized as follows: Chapter II provides an introduction to fiber-optic communications. It presents the characteristics of light propagation through fiber-media and the components of an optical link. Chapter III is an introduction to FDDI. It provides a fundamental look at the topology, hardware components, and general operation of an FDDI network. Chapter IV is a detailed examination of the various protocols that comprise the FDDI standards. Chapter V contains an overview of network design principles and a methodology for designing FDDI networks. Chapter VI presents justification for designing an FDDI solution for the Systems Management Department and a baseline review of the token-ring network. Chapter VII contains the recommended solution for implementing FDDI standards. Chapter VIII provides concluding remarks and additional recommendations.

II. FIBER OPTIC COMMUNICATIONS

A. INTRODUCTION

FDDI uses fiber optics to achieve high rates of data transfer. To appreciate the advantages and disadvantages of this media, network designers must understand the principles of fiber optic communications and the characteristics of optical components. This chapter is intended to introduce the reader to these topics.

This chapter provides an overview of the components of an optical communications link, the characteristics of light propagation through fiber, and the characteristics of light transmitters, light receivers, and fiber types. In addition, the advantages and disadvantages of various optical components will be discussed.

B. OPTICAL COMMUNICATIONS

1. Optical Communications Link

The basic components of an optical communications link is shown in Figure 2-1. The link receives input signals in the form of electrical pulses representing data. This data is fed into a driver which controls the generation of optical signals by the source. The light source, either a light emitting diode (LED) or an injection laser diode (ILD), generates the light pulses which are launched into the fiber. The light source, either a light emitting diode (LED) or an injection laser diode (ILD), generates the light pulses which are launched into the fiber.




Figure 2-1. Fiber Optical Communications Link.

The light source is carefully joined to the optical-fiber using a connector. Light pulses are launched into the fiber and begin propagating. The light continues to propagate the length of the fiber until it reaches the optical detector or light sensor. This sensor, either a PIN diode or avalanche photodiode (APD), converts the received light back into electrical pulses. These pulses are amplified and fed into the receiving equipment.

The driver-light source components of the block diagram are called a transmitter; the optical detector and output circuit, a receiver. To enable two-way communications, stations are equipped with a combined transmitter and receiver, or transceiver. In addition, two fibers are required to support full duplex transmissions. When fibers are not long enough to extend from one station to another, they are spliced to improve coupling efficiencies.

2. Fiber Optic Light Propagation

Fiber optic transmission systems operate in the infrared band of the frequency spectrum. More specifically, light sources generate pulses within one of the wavelength windows in Table 2-1 (Freeman, 1991).

Wavelength Window	Nominal Wavelength
810-850 nm	820 nm
1220-1340 nm	1330 nm
1540-1610 nm	1550 nm

Table 2-1. Fiber Optic Wavelength Bands.

In general, the goal of any communications system is to launch as much power into the medium as possible. In the case of an optical communications system, this ensures the signal is detectable and discernible when it arrives at the receiver. Critical to this factor is the careful alignment of optical components to the fiber.

Optical systems rely on connectors to align and connect the fiber to the transceivers. Connectors vary in size and coupling efficiency. Better quality connectors use keying features and gluing compounds to provide high coupling efficiencies. This ensures the maximum amount of light is launched into the fiber. However, quality connectors alone do

not guarantee the best coupling efficiency. Equally important is the quality of the fiber.

Light is launched into an optical fiber as a series of rays. These rays propagate along the fiber due to differences in the index of refraction between the fiber core and cladding. The number of rays accepted by the fiber is directly related to this difference in refraction indices.

As shown in Figure 2-2, rays enter the fiber at different angles and begin propagating. There is a maximum angle, however, that if exceeded causes the ray to enter the fiber's cladding. This angle, called the maximum entrance angle ($\theta_{\max}$), is calculated using the following formula:

$$2.1 \quad n \sin \theta_{\max} = \sqrt{n_1^2 - n_2^2},$$

where n is the index of refraction of air (~ 1), n_1 is the index of refraction of the core, and n_2 is the index of refraction of the cladding. Thus, the difference in refraction indices of the cladding and core directly affect the number of rays, and therefore power, entering the fiber.

The right side of equation 2.1 is referred to as the numerical aperture (NA) of the fiber,

$$2.2 \quad NA = \sqrt{n_1^2 - n_2^2}.$$

It is a direct indication of the quality of the fiber--the amount of optical power accepted by




Figure 2-2. Maximum Entrance Angle. After Jain, 1994.

the fiber varies as the square of its numerical aperture (Freeman, 1991). Vendors use NA to specify the quality of their fiber.

Though many rays may fail to enter the fiber due to poor alignment or a low numerical aperture, a considerable number still enter and begin propagating. However, the majority of these rays are reflected in a manner such that they will cancel each other. Only a small number will actually reach the receiver. The rays that continue to propagate are called modes. The number of modes supported in a fiber is dependent on the type of fiber, the wavelength of the light source, and the numerical aperture of the fiber.

To calculate the number of modes propagated, the normalized frequency (V) of the fiber is determined first, as shown in the equation below (Jain, 1994).

$$2.3 \quad V = \frac{2\pi a}{\lambda} NA.$$

In this equation, a is the radius of the core and λ is the wavelength of the light.

Once the normalized frequency has been determined, the number of modes is calculated by using one of the following formulas: $V^2/2$ for step index fibers; $V^2/4$ for graded-index fibers. Step index fibers contain cores whose index of refraction remain constant; the abrupt differences between the core and cladding cause the signal to propagate. Graded-index fibers are manufactured such that the index of refraction of the core decreases radially from its center. This gradual change causes the light pulse to bend and refract along the fiber's axis.

3. Fiber Modes

Fibers are categorized as single-mode or multimode. The distinction is dependent on the number of modes propagated--a distinction directly related to the size and construction of the core as discussed above. In general, fibers with a core diameter less than 10 μm are single-mode; fibers with a core diameter between 50 and 100 μm are multimode (Jain, 1994).

Given the small diameter of single-mode fiber cores, only one ray enters and propagates down the core axis. The large diameter of multimode fibers on the other hand,

can propagate hundreds of rays. Although this means multimode fibers can accept more light power, it does not imply they are the preferred media for all applications.

Multimode fibers contain more impurities than single-mode fibers, which causes a linear degradation of the signal as it propagates the fiber. Moreover, because the light pulse consists of a large number of rays traveling different paths along the fiber, these rays arrive at the receiver at slightly different times. This phenomenon changes the characteristic of the signal.

4. Attenuation and Dispersion

There are two phenomenon associated with fiber optic transmission that affect performance of systems: attenuation and dispersion. Attenuation is the absorption and scattering of light energy due to impurities in the fiber. It causes a loss in signal amplitude that affects the receiver's capability to detect and recognize a valid signal from noise.

Vendors specify the attenuation of their fiber in terms of decibel (dB) loss per kilometer--the longer the fiber, the greater the loss. Network designers account for the affects of attenuation by adding these losses to other link losses. The total loss cannot exceed the difference between the power launched into the fiber by the transmitter and the power required at the receiver to discern the signal.

In addition to suffering a loss of signal strength, light pulses experience a widening effect as they travel the length of the fiber. This widening effect is called dispersion. After a certain distance, the pulses become so wide that succeeding pulses interfere with each other and the signals are no longer discernable at the receiver. Thus, if the dispersion is high, the data rate must be slowed in order to separate the signals. This causes a reduction in the bits per second or bandwidth. There are two types of dispersion: modal and chromatic.

a. Modal Dispersion

In a multimode fiber, numerous rays that constitute the light pulse propagate along the fiber, traveling different paths and hence different distances--the rays that continuously reflect off the cladding travel longer distances than the ray that travels along the axis of the fiber. Since all the rays travel at the same speed, different rays arrive at the receiver at different times. This causes the pulse width to increase with the length of the




Figure 2-3. Dispersion Effect on Signal Pulses. From Jain, 1994.

fiber, as shown in Figure 2-3. This effect is called modal dispersion and occurs only in multimode fibers.

Modal dispersion ultimately results in a limit called modal bandwidth. This limit is specified as a bandwidth-distance product and is used by vendors to grade their fibers. A fiber with 500 MHz-km bandwidth-distance product will have a bandwidth of 250 MHz at two kilometers. (Jain, 1994)

To improve the bandwidth-distance product, the number of modes propagated must be reduced. This may be accomplished by using a fiber with a smaller core, using a fiber with a higher numerical aperture, or using a light source that generates a signal with a higher wavelength (1500 nm versus 1300 nm). The end-result is that the normalized frequency (V) of the fiber is reduced, causing a corresponding reduction in the modes propagated. (Jain, 1994)

b. Chromatic Dispersion

Although single-mode fibers are not subjected to modal dispersion, they are bandwidth limited by chromatic dispersion. This type of dispersion is a function of velocity differences associated with the range of wavelengths produced by a light source.

Light sources are not capable of producing single wavelength, instantaneous emissions. Instead, they produce a light pulse of varying wavelengths. These different wavelength components travel at different speeds and therefore arrive at the receiver at different times. This in turn causes a pulse-widening effect similar to that produced by modal dispersion.

C. CHARACTERISTICS OF LINK COMPONENTS

1. Optical Transmitters

Fiber optic communication systems use either LED or laser diode light sources. These sources vary in terms of performance, reliability, and cost.

Overall, the performance of lasers are far better than LEDs, as shown in Table 2-2 (Jain, 1994; Freeman, 1991). They provide greater power, narrower spectral width, faster rise and fall times, and better coupling efficiency. Such characteristics allow them to achieve transmission over tens of kilometers, compared to just a few kilometers for LEDs.

Characteristic	LED	Laser Diode
Power launched	100 μ W (fiber w/.2 NA)	5-7 mW (fiber w/.2 NA)
Spectral width	25-50 nm	2-5 nm
Rise/Fall time	3-20 ns	.5-2 ns
Coupling efficiency	2%	50%
Lifetime	200,000 hours	50,000-100,000 hours

Table 2-2. LEDs versus Laser Diodes. After Jain, 1994.

Lasers, however, have complicated drive circuits and require more power than LEDs. This makes them less reliable and considerably more expensive. Moreover, they pose a potential safety hazard.

The very focused and high power beam of invisible light produced by a laser can cause eye damage. To prevent such accidents, government safety regulations require controlled access to equipment and media that use high-powered lasers. Access is limited to trained and certified personnel only. (Jain, 1994)

To summarize, laser diodes are the preferred light source when long link distances are involved. In short link applications, network designers prefer the lower cost, higher reliability, and safety of LEDs.

2. Optical Receivers

Optical receivers use light detectors called photodiodes to convert the received light signal back into an electrical current. There are two types of photodiodes used in optical communications: PINs and APDs. The characteristics of each are presented in Table 2-3.

Characteristic	PIN	APD
Responsivity	0.5-0.7 $\mu\text{A}/\mu\text{W}$	30-80 $\mu\text{A}/\mu\text{W}$
Bias voltage	10V	100+ Volts
Temperature sensitivity	Less	More
Availability	1300 nm available	Mostly 850 nm
Cost	Less	More

Table 2-3. PIN Photodiodes versus APDs. After Jain, 1994.

Of the two types of photodiodes available, APDs provide the greatest responsivity--amount of power generated by the photodiode for a given light input. However, they require considerably more power to operate. Moreover, they generate their greatest responsivity gain at a wavelength of 850 nm; though readily available for this wavelength, they may be difficult to find for 1300 nm and 1500 nm applications. (Jain, 1994)

PINs, in contrast, require less voltage, are less temperature sensitive, cost less, and are readily available for 1300 nm wavelengths--the wavelength used in FDDI standards. Therefore, with the exception of their lower responsivity, PIN photodiodes are the preferred detector. (Jain, 1994)

3. Fiber Types

As indicated above, fibers are classified as either single-mode or multimode. Multimode fibers are further categorized as either step-index or graded-index.

Multimode step-index is the easiest fiber to manufacture. It has a thick core diameter (about 30 to 200 μm) that permits a large amount of power to be launched into the fiber. (Jain, 1994) This characteristic makes it suitable for use with cheaper LED sources. Its drawback, however, is that modal dispersion limits the fiber to a bandwidth-distance product

of 10 to 100 MHz-km (Freeman, 1991).

A multimode graded-index fiber also has a large core making it compatible with LED light sources. Though more expensive than step-index fibers, it reduces the effects of modal dispersion, thus permitting bandwidth-distance products of 300 MHz-km when coupled with LEDs. If coupled to a laser diode, it can achieve bandwidth-distance products from 400 to 1000 MHz-km. (Freeman, 1991)

Because of its small core, a highly focused beam is necessary to launch the light into a single-mode step-index fiber. This requires the use of edge-emitting LEDs or laser sources that are considerably higher in cost than surface emitting LEDs. Its advantage, however, is that bandwidth-distance products up to 1000 GHz-km are possible. (Jain, 1994)

D. SUMMARY

This chapter presented the basic concepts of optical communications. Furthermore, it presented the components used in optical communications, and the tradeoffs between different types of light sources, receivers, and fiber.

Transceivers consist of a light source, either a LED or laser diode, and light detector, either a PIN or APD. Due largely to the high cost of laser diodes, the preferred transmitter is the LED for distances up to a few kilometers. Due to limited availability of APDs in the 1300 nm wavelength—a specified FDDI standard—and their high power requirements, PINs are the preferred photodetectors.

Fiber types include single-mode, multimode step-index, and multimode graded-index. Single-mode fibers are used with laser diodes to achieve high bandwidth-distance products. These fibers are not compatible with the cheaper surface-emitting LEDs.

Either multimode step-index or multimode graded-index fibers may be used with LEDs. Of the two, multimode graded-index provides the greatest bandwidth-distance products.

Lastly, the quality of a fiber is specified by vendors in terms of attenuation losses, bandwidth-distance products, and numerical apertures. Attenuation is directly related to the degradation of the light signal due to impurities in the fiber; bandwidth-distance is a function of the fiber's dispersion characteristics; and numerical aperture relates to the

amount of power accepted by the fiber and the modes propagated. Each of these factors influence the choice of one fiber over another. Designing a link for a particular application requires choosing the right combination of fiber and transceiver.

III. FDDI FUNDAMENTALS

A. INTRODUCTION

The FDDI topology is based on a dual ring of trees configuration. A cable consisting of a fiber pair interconnects the nodes of the network to form a dual logical ring topology. Normally, only one fiber is used for data exchange between nodes. This fiber is referred to as the primary ring. The second fiber or secondary ring is in a standby mode. It becomes active when a node or link failure occurs along the primary ring. If a failure occurs, the other nodes detect the condition and reroute traffic onto the secondary ring, preserving network functionality.

The purpose of this chapter is to take a closer look at the topology of FDDI networks. More specifically, it will address the different types of nodes used in these networks, their protocol structures, and their interconnection rules. Moreover, the ring reconfiguration capability of FDDI networks will be discussed.

This chapter begins with a overview of FDDI protocols and their relationship to the International Standards Organization's (ISO) Open Systems Interconnection (OSI) model. This overview is intended to serve as an introduction to support other topic discussions throughout the chapter. FDDI protocol specifications will be examined in detail in Chapter IV.

B. INTRODUCTION TO FDDI PROTOCOLS

1. OSI Reference Model

Protocols are communication rules used by network computers to communicate with each other. These protocols are implemented through software and hardware products. To promote the compatibility of protocols, and ultimately the design of open systems, the ISO approved the OSI reference model as a standard in 1983. (Schivley, 1994) This model serves as a logical framework for defining and developing communication rules across seven layers of functionality: application, presentation, session, transport, network, data link, and physical layers.

Each layer of the model represents a different level of functionality required to support communications between two network computers. The model works in a top-down fashion beginning with the application layer and ending with the physical layer, as shown in Figure 3-1. During a communications exchange, the user enters data at the application level. The data undergoes a transformation process as it passes through each layer, until it is finally transmitted across the physical medium to a receiving node. At the receiving node, the data undergoes a reverse transformation back up the model, where it is viewed by the second-party user at the application layer.




Figure 3-1. OSI Reference Model. From Schivley, 1994.

To effectively communicate, nodes must have compatible protocols at their required layers of interface; a node does not have to implement all the layers of the model. For example, routers use only the network, data link, and physical layers to manage traffic routing functions. Likewise, standards need not address all layers of the model. IEEE 802 standards, for example, only implement protocols that correspond to the two lowest layers of the model. These standards rely on other higher-layer protocols to complete communications to the application layer.

In addition to promoting the development of standardized protocols, the OSI model serves as a useful tool for analyzing the compatibility of different networks. By referencing

the protocols of different standards employed on a network to the OSI model, a network designer can identify potential compatibility problems. This approach is used in the analysis of FDDI protocols.

2. FDDI Protocols

FDDI standards specify protocols that map into the physical and data link layers of the OSI model as shown in Figure 3-2. The standards include a physical medium dependent (PMD) sub-layer, a physical medium independent (PHY) sub-layer, a media access control (MAC) sub-layer, and a station management (SMT) entity. Though the figure shows a logical link control (LLC) sub-layer, this protocol is not an FDDI standard. FDDI uses the LLC specified in the IEEE 802 protocols. (Jain, 1994)

The PMD sub-layer specifies the characteristics of the media, interconnection with that media, and the characteristics of the transceivers. Transceiver specifications include transmitter power, frequencies, receiver sensitivities, and repeater-less-distances between nodes. (Shah, 1994)

There are separate PMD standards for the different types of media. In the case of optical-based networks, PMD specifications describe the characteristics of the fibers, connectors, and optical transceivers. In the case of copper-based networks, the characteristics of the copper medium and electrical transceivers are prescribed.




Figure 3-2. FDDI Standards Comparison to the OSI Reference Model.

The PHY sub-layer interacts with the PMD sub-layer to perform functions that include 4B/5B and non-return-to zero invert-ones (NRZI) encoding and decoding, buffering, smoothing, repeat filtering, and initialization of the medium. This sub-layer specifies protocols regardless of the media used—fiber or copper. It exchanges data with the MAC sub-layer, completing the link between the OSI data link and physical layers. (Shah, 1994)

The MAC standard specifies the data, token and management frame structures, frame check sequence generation and verification, the timed token access protocol, and addressing. It also specifies a set of control frames called MAC frames that are used for ring initialization and fault isolation. (Shah, 1994)

SMT is an entity that overlaps the media access control, and physical medium dependent and independent sub-layers. It provides high-level and low-level monitoring and management of these sub-layers. It further specifies the ring initialization, error monitoring, and ring fault recovery procedures. (Mills, 1995)

C. FDDI TOPOLOGY

The network topology consists of a cable that interconnects various network nodes, as shown in Figure 3-3. The cable is physically referred to as the trunk and contains two optical fibers to support data communications. One fiber serves as the network's primary ring and handles all traffic during normal operation. The second fiber serves as a backup or secondary ring. Traffic is redirected to this ring only in the event of a link or node failure on the primary ring. When both rings are active, data is transmitted in counter-rotating directions.

1. Nodes

FDDI network nodes are categorized as either stations or concentrators. Stations are active nodes that transmit and receive information in the form of optical signals; concentrators serve as distribution points that provide the necessary components to connect multiple stations to the network.

Nodes are further categorized by the manner in which they connect to the trunk. A station or concentrator may connect to both rings of the trunk, referred to as a dual-attachment node, or to just the primary ring, a single-attachment node.




Figure 3-3. FDDI Topology.

Stations that connect to the primary ring only are referred to as single-attachment stations (SASs). Stations that connect to both rings are called dual-attachment stations (DASs). Besides the obvious difference in physical connections, the two types of stations differ in their network functionality.

Due to their physical connection to the trunk and software programming, DASs are capable of reconfiguring the logical flow of data between nodes. Should a link or station failure occur, the DAS redirects traffic onto the secondary ring. This capability is normally reserved for stations that are critical to network operation, such as minicomputers or file servers. SASs, in contrast, only connect to the primary ring. They cannot reconfigure the flow of traffic. They are nodes that are not critical to network functionality, and hence do not require the fault tolerance of a DAS. An example of a SAS is a desktop computer.

Concentrators are classified as either single-attachment, dual-attachment, or null-attachment, as shown in Figure 3-4. Single-attachment concentrators (SACs) are connected to the primary ring of the trunk generally through a dual-attachment concentrator. They are used to interconnect a series of SASs or other SACs. Dual-attachment concentrators (DACs) connect to both rings of the network. They provide the fault tolerance of a dual-attachment




Figure 3-4. Types of Concentrators. From Jain, 1994.

station and are used to establish tree topologies that cascade into SACs and SASs. Null-attachment concentrators (NACs) are not connected to the trunk per se, but serve as a trunk root. They may be used to interconnect a series of stations.

Figure 3-5 illustrates how the various types of nodes are connected to form an FDDI network. The primary and secondary rings are shown to illustrate the flow of data between the nodes. In a physical representation of this network, the data paths would be replaced by a single connection representing the fiber-pair cable. (Jain, 1994)

The network in the figure expands through three levels of hierarchy. This expansion could be increased if needed by cascading additional concentrators. Refinement of the tree to greater detail is allowed, provided the total length of the network does not exceed 200 km. In addition, the total number of nodes cannot exceed 500.




Figure 3-5. Example of an FDDI Network.

2. Protocol Structure of Nodes

Connection between a node and fiber is accomplished through a port. Each port consists of a PMD and PHY protocol pair. The remaining node protocol structure depends on the type of station or concentrator used. (Jain, 1994)

In addition to the PMD-PHY pair, single-attachment station protocols contain a MAC and an SMT. Dual-attachment stations require two PMD-PHY pairs, an SMT, and at least one MAC to access the media. A second MAC is not required since the FDDI secondary ring is in a standby mode. When a fault occurs, the ring reconfigures and signal paths within the station can be rerouted through its single MAC. (Jain, 1994)

In some applications, dual-attachment stations may be configured with two MAC protocols. This configuration is used when the network designer intends to transmit and receive data on both the primary and secondary rings simultaneously. Such a configuration can achieve a 200 Mbps transfer rate. In the event of a primary ring failure, the network can still reconfigure traffic onto the secondary ring. This causes the network to revert back to 100 Mbps. (Mills, 1995)

Concentrators require a PMD-PHY pair for each trunk ring connection and for each single-attachment node connection. They also require an SMT protocol. The installation of MACs in concentrators is optional. Theoretically, since concentrators do not need to transmit messages, they do not require a MAC. However, in order to enable remote station management features, a MAC is necessary. This will enable communications between stations and concentrators to support SMT frame based protocols, which will be discussed in Chapter IV.

3. Ports

The number of active ports on a node depends on the type of node and its network configuration. A port is required for each connection to a logical ring and, as in the case of a concentrator, for connecting to additional nodes that form the tree topology. The following paragraphs discuss the various FDDI port designations, as illustrated in Figure 3-6.

Single-attachment nodes use a single port to connect to the primary ring of the trunk.




Figure 3-6. Port Designations.

This port, called a slave port or S-port, accepts a connection that consists of two fibers. One fiber handles signals traveling from the ring to the node; the second returns signals from the node back to the ring.

Dual-attachment stations use two ports, designated A- and B-ports, to connect to the logical rings. Each port accepts a cable that consists of two fibers that serve as the rings of the trunk. Traffic is routed from the primary ring into the node, through the A-port. Within the node, traffic is routed back onto the primary ring through the B-port. Conversely, the secondary ring is connected in a manner such that traffic enters the node through the B-port, and returns to the ring through the A-port.

Concentrators are configured with either an S-port or combination A- and B-port. The configuration depends on their connection to the trunk--single-attachment or dual-attachment. Moreover, they also contain master ports or M-ports that interconnect other stations. M-ports on a dual-attachment concentrator may connect to single-attachment stations or single-attachment concentrators. M-ports on a single-attachment concentrator will interconnect single-attachment stations only. A concentrator can have a varying number of active M-ports. (Jain, 1994)

4. Connection Rules

Understanding port configurations is important to the installation of a network. Incorrectly connected ports result in ring perturbations that degrade functionality or inhibit network operation.

For example, the normal port connections for two dual-attachment nodes is A-to-B as shown in Figure 3-6. If the stations are incorrectly connected in an A-to-A or B-to-B configuration, the two logical rings become twisted. Likewise, an A-to-S or B-to-S connection causes a wrapped condition that reduces the two logical rings to just one. Although both of these ring conditions are legal, they inhibit full network functionality.

Moreover, A-to-M connections result in the formation of multiple rings if a B-to-M connection is also active. This connection would prevent a signal entering an A port from exiting through the B port, a condition referred to as prevent through. Some vendors may implement these connections by disabling the A-port. In addition, M-to-M connections also result in multiple rings that interrupt the normal flow of data between stations.

To preclude the inadvertent misconnection of ports, FDDI standards require vendors to label node ports with their appropriate designation. Moreover, special keyed connectors are recommended to prevent some of the common configuration mistakes. Nonetheless, users may still make errors that affect network operation.

A summary of FDDI connection rules is contained in Table 3-1. Connections that are identified as valid but undesired are treated differently by the station's SMT, depending on the policy set by the network manager. (Jain, 1994)

Port Connection	A	B	S	M
A	Valid; Undesired	Valid	Valid; Undesired	Valid; Prevent Through
B	Valid	Valid; Undesired	Valid; Undesired	Valid; Prevent Through
S	Valid; Undesired	Valid; Undesired	Valid	Valid
M	Valid	Valid	Valid	Valid; Undesired

Table 3-1. FDDI Connection Rules. After Jain, 1994.

5. FDDI Ring Reconfiguration

FDDI networks have the ability to reconfigure in the event of a link or node failure. When a failure occurs, stations on either side of the failure detect the condition and automatically reroute traffic. Reconfiguration will also occur as a result of a node power-down. This feature allows FDDI to preserve functionality.

An FDDI network can tolerate a single failure at the trunk level without seriously impacting ring operation. Two or more failures will segment the ring. When a fault occurs within a station or along a link, it is detected and the ring automatically reconfigured to reroute traffic onto the secondary ring. Reconfiguration occurs within milliseconds and is controlled by station management. If the failure that caused the reconfiguration is corrected, the restored condition is subsequently detected and the ring reconfigured for normal operation. (Joshi, 1986)

Since single-attachment stations are normally connected to the network through a concentrator, securing power to a SAS will be detected by protocols within the concentrator and the signals routed around the port serving that station. Securing power to a dual-attachment station has significantly different consequences; a loss of power would be interpreted by other stations as a failure, causing the network to reroute traffic onto the secondary ring.

The following figures illustrate FDDI's ring reconfiguration capability. During normal operation, traffic flows along the primary ring of the network while the secondary ring remains idle, as shown in Figure 3-7.

In Figure 3-8, a cable fault has occurred between DAS 2 and DAC 2. Station management detected the failure and reconfigured the logical flow of data as shown. At DAS 2, traffic flow was wrapped-around at the incoming port back onto the secondary ring. Within DAC 2, traffic was routed to the port serving station 6.

In Figure 3-12, a second link failure has occurred. Station management detected the new failure and caused a wrap-around condition at DAS 1 and a rerouting of traffic between the single-attachment stations in DAC 2. The result is DAC 2 and its single-attachment stations are segregated from the trunk ring and thus, the rest of the network. The resulting




Figure 3-7. Normal Ring Operation.




Figure 3-8. Reconfiguration After Cable Fault.




Figure 3-9. Reconfiguration After Second Fault.

functionality of each sub-network depends on the capabilities of the various nodes.

Lastly, Figure 3-10 illustrates network response to a single-attachment station power-down. In this case, Station 3 has been secured causing the wiring concentrator to reroute internal traffic in a manner that bypasses the affected port, preserving ring functionality.




Figure 3-10. Reconfiguration After Station Power-down.

When single-attachment stations are connected through a concentrator, securing power to them will not cause a ring reconfiguration. The protocols within the concentrator will simply bypass the port serving the secured stations. If however, power is secured to a dual-attachment station, the logical flow of data around the ring is disrupted. Other stations detect the problem and would reconfigure the ring to maintain functionality. This is an undesirable condition since a second dual-attachment node failure or power-down will result in a segmented ring. To prevent this undesired reconfiguration, dual-attachment stations may be equipped with an optical bypass.

6. Optical Bypass

An optical bypass is a relay that maintains ring continuity when a failure or power-down condition occurs in a dual-attachment node. The bypass routes incoming light signals directly to the next station by switching the light path from the incoming port to the outgoing port. The signals are switched either optically or through prisms. As a result, the signals are not regenerated electronically. (Jain, 1994)

The use of optical bypasses is optional. When used, they enhance the functionality of the network by eliminating ring reconfigurations that would otherwise occur if a dual-attachment node were taken off-line. Unfortunately, they also cause considerable loss of signal strength and compound link loss calculations.

For example, the network in Figure 3-11 consists of three dual-attachment nodes, each configured with an optical bypass switch. The transceivers used in these nodes meet the FDDI standards for power levels, providing 11 dB difference in minimum power levels between transmitter and receiver (Jain, 1994). The network fiber has an attenuation rating of 1.5 dB per kilometer, the connectors a loss of 1.2 dB each, and the splices a loss of .5 dB each. Thus, the overall loss between DAS 1 and DAC 1 is 4.4 dB, and between DAC 1 and DAS 2, 5.9 dB--well within the 11 dB budget between active nodes.

If, however, DAC 1 were powered-down, the optical bypass in this node would redirect traffic to DAS 2, while causing a 1.9 dB loss in signal strength. Moreover, since the link has been extended from DAS 1 to DAS 2, the signal incurs a total loss of 12.2 dB ($4.4 \text{ dB} + 1.9 \text{ dB} + 5.9 \text{ dB}$). This exceeds the 11 dB budget limit. Thus, the network has become attenuation limited. This example illustrates the compounding problems presented by the installation of bypass switches. (Jain, 1994)




Figure 3-11. Optical Bypass Analysis.

D. SUMMARY

FDDI standards do not implement protocols across all seven layers of the OSI reference model. In fact, the FDDI protocols only implement the physical and data link layers of the model. Thus, to enable application-to-application communication between two stations, a second protocol stack such as the Transmission Control Protocol/Internet Protocol (TCP/IP) is required.

The FDDI topology consists of stations and concentrators interconnected by fiber-pair cables. These nodes are categorized as either single-attachment or dual-attachment nodes. Single-attachment nodes connect to the primary ring only, and are not capable of reconfiguring the trunk should a failure occur. Dual-attachment nodes are connected to both logical rings and can support reconfiguration. When a reconfiguration of the ring occurs, stations wrap the flow of traffic from the primary ring onto the secondary ring.

When power to a node is secured, the network responds to the change in condition. If the node is a single-attachment node connected through a concentrator, the concentrator internally reroutes signals around the affected node. If the node is a dual-attachment node, the power-down will result in a ring reconfiguration. More than one dual-attachment node power-down will segment the trunk into sub-networks. This undesired reconfiguration can be prevented by installing optical bypasses in the dual-attachment nodes.

Interconnection between stations is critical to network functionality. To preserve the dual logical ring topology, network builders must ensure valid port connections are used. Incorrectly connected ports may result in twisted rings, wrapped rings, or segmented rings. Such perturbations may be handled differently depending on the vendor. Others may be controlled by the network manager using station management protocols.

IV. FDDI PROTOCOLS

A. INTRODUCTION

This chapter examines the various FDDI protocols in detail. It begins with a discussion of the physical media dependent standards for multimode fiber, low-cost fiber, and twisted pair media. Since the scope of this project is limited to a single building, the use of single-mode fiber media is not required and has therefore been omitted from the discussions.

Following the discussion of media dependent standards, information regarding each successive protocol--physical medium independent and media access control--is presented. The chapter ends with an overview of station management.

B. PHYSICAL MEDIUM DEPENDENT STANDARDS

During development of FDDI standards, fiber was chosen as the principle medium because of its high capacity and low attenuation. These characteristics were deemed essential to providing the high data and low error rates necessary to support emerging technologies. Fiber is also lightweight, immune to radio frequency interference, chemically resistant, and temperature variation resistant. In addition, it does not generate electromagnetic interference, and provides greater security against deliberate taps to gain unauthorized access to data. Moreover, the cost of installed fiber has fallen rapidly and is currently within 10-15 percent of that of Category 5 copper (Raynovich, 1995).

The biggest drawback to fiber optic communications is the cost of optical transceivers. For this reason, ANSI developed standards for single-mode fiber, multimode-fiber, and low-cost fiber. These separate standards use transceivers that differ in performance and cost. Network designers can choose the standard that fits their needs.

Moreover, considering the investment businesses have made in copper-based media, ANSI developed standards for both shielded twisted pair and unshielded twisted pair. The obvious advantage to using copper-based media is its affordability; the drawback is its susceptibility to electromagnetic influence, crosstalk, and noise. Regardless, Category 5 unshielded twisted pair can achieve the bit and error rates required by FDDI standards, for

distances up to 100 meters (Shah, 1994). This represents a cost tradeoff that weighs heavily in a decision to use fiber optics as the infrastructure for a network.

1. Multimode Fiber PMD (MMF-PMD)

The MMF-PMD is based on a 1300 nm wavelength optical signal. This wavelength was selected due to: (1) its lower attenuation and dispersion compared to the 850 nm wavelength; and (2) its compatibility with the cheaper LED light sources--1550 nm wavelength signals are generated with laser sources. The PMD specifications for multimode fiber are summarized in the following table. (Jain, 1994)

MMF-PMD Characteristic	PMD Specification
Optical Spectrum	1300 nm
Fiber Link	
Fiber type	50/125, 62.5/125, 85/125, 100/140 μm
Modal bandwidth	500 MHz-km (62.5/125 @ 1300 nm)
Maximum link loss	11 dB
Maximum length	2 km
Transmitters	LEDs/Lasers
Center wavelength	1270-1380 nm
Average power	-20 dBm to -14 dBm
Receivers	PINs/APDs
Wavelength detection range	1270-1380 nm
Detectable power range	-31 to -14 dBm
Connectors	Duplex-FDDI (port and polarity keying)

Table 4-1. Multimode Fiber PMD Specifications. After Jain, 1994.

a. Fiber Link

MMF-PMD standards permit the use of four different fiber types; the default type is 62.5/125 μm graded-index fiber. This fiber provides the best performance for the 1300 nm optical spectrum. When using other types, careful bandwidth and loss analysis is required to ensure the link meets minimum performance requirements. These analyses are not required when using 62.5/125 fiber, provided the other link and component

specifications in Table 4-1 are met. (Jain, 1994)

The maximum distance allowed between nodes in an FDDI network is 2 km. This distance limit serves as a basis for setting bandwidth and attenuation standards for the fiber. Extending the distance beyond 2 km will require higher quality light sources and better quality fiber. Conversely, using less quality transceivers or fiber will require shortening the distance between nodes.

To effectively support the transmission of FDDI signals, the fiber must provide a minimum bandwidth of 250 MHz. With a maximum allowed distance of 2 km between nodes, this equates to a fiber bandwidth-distance product of 500 MHz-km. (Jain, 1994)

To ensure sufficient optical power is launched into the fiber, the numerical aperture for 62.5/125 fiber must be .275. This equates to a maximum entrance angle of 16 degrees--a requirement easily achieved by LED light sources. (Jain, 1994)

The total signal loss in a link cannot exceed 11 dB. This includes loss due to fiber attenuation, connectors, and splices. Connectors cause losses that range from .5 to 1 dB, while splices cause losses that range from .1 to .5 dB (Baker, 1986). In the worst case, a link with two connectors and one splice will incur a loss of 2.5 dB, leaving 8.5 dB remaining for fiber attenuation. If the fiber has an attenuation of 4 dB/km, the link will meet the loss budget of 11 dB. Fiber attenuation is typically 2 dB/km (Jain, 1994).

b. Transmitters and Receivers

Either LED or laser light sources may be used with multimode fibers. Although laser sources provide better performance characteristics than LEDs, their high cost make LEDs the preferred source. The minimum average output power of these sources must fall between -20 dBm and -14 dBm. (Jain, 1994)

Receivers may use PINs or APD detectors to convert received optical signals back into electrical pulses. Although APDs are more sensitive than PINs, they also require more power, are more temperature sensitive, and cost more than PIN detectors. Thus, PIN detectors are the preferred diode for receivers. (Jain, 1994)

Receivers must be capable of discerning signals in the power range between

-31 dBm and -14 dBm. The minimum limit ensures that a signal transmitted at the minimum power level of -20 dBm, can incur a loss up to 11 dB and still be discernable at the receiver. The maximum limit ensures the receiver will not be overloaded by the transmitter on short links.

Transceivers are designed for a particular fiber type. The 11 dB margin provided by the standard is based on the fact that the transceiver and fiber have been matched. Connecting a transmitter designed for use with 62.5/125 fiber to a 50/125 fiber will result in less light launched into the fiber. Conversely, connecting that same transmitter to a 100/140 fiber may result in a signal gain. The result depends on the design of the transceiver, the diameter of the fiber core, and the fiber's numerical aperture. Table 4-2 shows the impact on link design caused by gains and losses associated with using a 62.5/125 transceiver with different fiber types. (Jain, 1994)

Fiber Type	Transmitter Loss/Gain	Receiver Loss/Gain	Remaining Margin
50/125 μm (NA=0.20)	5.0 dB Loss	0.0 to 1.0 dB Gain	6.0 to 7.0 dB
50/125 μm (NA=0.21)	4.5 dB Loss	0.0 to 1.0 dB Gain	6.5 to 7.5 dB
50/125 μm (NA=.22)	4.0 dB Loss	0.0 to 1.0 dB Gain	7.0 to 8.0 dB
85/125 μm (NA=0.26)	2.0 dB Gain	0.0 to 2.6 dB Loss	10.4 to 13.0 dB
100/140 μm (NA=0.29)	2.0 dB Gain	0.0 to 4.0 dB Loss	9.0 to 13.0 dB

Table 4-2. Impact of Using Alternative Fiber Types with a 62.5/125 Transceiver.
After Jain, 1994.

c. Media Connectors

FDDI standards prescribe the use of specially designed duplex media connectors. These connectors are used to interconnect the optical fibers into the transceivers of the nodes, and to interconnect fiber cables.

For interconnecting nodes, a single duplex-connector plug called a media interface connector (MIC) is used at each end of a fiber pair. The MIC connects into a matching receptacle on the node's adapter card. The design of the plug and receptacle allows inexperienced network builders to interconnect nodes without fear of making incorrect

connections. The geometry of the MIC and receptacle prevent the inadvertent flip-flop or polarity reversal of fibers—a polarity reversal would connect two transmitters to the same fiber. Moreover, the plug and receptacles are designed to provide adequate alignment of the adapter card and fiber, to minimize losses due to poor coupling efficiency. An example of a MIC is shown in Figure 4-1.




Figure 4-1. Media Interface Connector. From Mills, 1995.

Moreover, MICs and fiber cable connectors have special keying features. These keys, which differ in size and position on the duplex plugs and receptacles, are intended to prevent the misconnection of node ports. Incorrectly connected ports, such as an A-to-A port connection, affects network functionality as discussed in Chapter III. Figure 4-2 shows the keying geometry used for multimode fibers; single-mode fiber keying is different.

In addition to duplex connectors, commercial vendors offer several types of simplex connectors. These connectors, such as the ST and SC connectors shown in Figure 4-3, may be used in place of duplex connectors; two simplex connectors are required to replace each duplex connector. Although these connectors are cheaper than duplex connectors, they do not provide keying features that prevent the misconnection of fibers and nodes. They do not meet the duplex specifications of FDDI standards.

d. Mixing Fiber Types and Sizes

FDDI permits the mixing of different fibers between stations. Doing so however, may cause additional losses due to core diameter and numerical aperture




Figure 4-2. MMF-PMD Connector Keying. From Jain, 1994.

differences, as shown in Table 4-3. Due to the significant difference in core diameters between single-mode and multimode fibers, these fibers cannot be interconnected.

e. Optical Bypasses

Optical bypasses are optional for use in dual-attachment nodes. The performance characteristics include a maximum attenuation level of 2.5 dB, a maximum switching time of 25 ms, a maximum media interruption time of 15 ms, and a minimum interchannel isolation level of 40 dB. The interchannel isolation limit is necessary to prevent crosstalk between light signals traveling on the primary and secondary rings. Both rings are active during fault conditions that cause ring reconfiguration. (Jain, 1994)

2. Low-Cost Fiber PMD (LCF-PMD)

Low-Cost fiber standards were developed to offset the high cost of optical components used in the MMF-PMD standard. Although the name implies low-cost fiber, the




Figure 4-3. Simplex Connectors. From Shah, 1994.

Receiving Fiber	Transmitting Fiber					
	50 μm NA=0.20	50 μm NA=0.21	50 μm NA=0.22	62.5 μm NA=0.275	80 μm NA=0.26	100 μm NA=0.29
50 μm (NA=0.20)	0.0	0.2	0.4	2.2	3.8	5.7
50 μm (NA=0.21)	0.0	0.0	0.2	1.9	3.5	5.3
50 μm (NA=0.22)	0.0	0.0	0.0	1.6	3.2	4.9
62.5 μm (NA=0.275)	0.0	0.0	0.0	0.0	1.0	2.3
80 μm (NA=0.26)	0.0	0.0	0.0	0.1	0.0	0.8
100 μm (NA=0.29)	0.0	0.0	0.0	0.0	0.0	0.0

Table 4-3. Losses (dB) Due to Mixing Multimode Fiber Types. From Jain, 1994.

cost savings are not really a result of using cheaper fiber, but a result of relaxing transceiver power requirements. To accommodate the use of lower-powered transceivers, the maximum link distance allowed is reduced to 500 meters. Table 4-4 summarizes the specifications of this standard.

LCF-PMD Characteristic	PMD Specification
Optical Spectrum	1300 nm
Fiber Link	
Fiber type	50/125, 62.5/125, 85/125, 100/140, 200/230 μm
Maximum link loss	7 dB
Maximum length	500 meters
Transmitters	
Center wavelength	1270-1380 nm
Average power	-22 dBm to -14 dBm
Receivers	
Wavelength detection range	1270-1380 nm
Detectable power range	-29 to -14 dBm
Connectors	Duplex-SC or Duplex-ST (polarity keying only)

Table 4-4. Low-Cost Fiber PMD Specifications. After Jain, 1994.

a. Fiber Link

LCF-PMD standards recommend the use of 62.5/125 multimode graded-index fiber. Other fiber types are permitted; however, the network designer must again match the fiber type and transceiver.

As was the case with multimode fiber, connecting a 62.5/125 transceiver to other fiber types, will impact the link margin. The gain or loss will depend on the diameter of the fiber's core and the fiber's numerical aperture. Table 4-5 summarizes the impact on link budget as a result of mixing different low-cost fiber types with a 62.5/125 transceiver. These impacts can be easily avoided by using 62.5/125 fiber. (Jain, 1994)

The standard limits the maximum link length to 500 meters. As a result, potential problems with dispersion are avoided. Thus, the standard does not specify a bandwidth requirement. (Jain, 1994)

b. Transmitters and Receivers

The cost savings associated with LCF-PMD is a result of the lower-powered, less-sensitive transceivers allowed by the standard. The LCF transmitters are required to generate a minimum signal level of -22 dBm, compared to -20 dBm for MMF-PMD transmitters. LCF receivers must be able to detect a minimum signal level of -29 dBm, compared to -31 dBm for MMF-PMD receivers. These differences in performance translate into optical component savings.

The drawback to reducing transceiver performance specifications is the

Fiber Type	Transmitter Loss/Gain	Receiver Loss/Gain	Remaining Budget
50/125 μm (NA=0.20)	5.0 dB Loss	0.0 to 1.0 dB Gain	2.0 to 3.0 dB
62.5/125 μm (NA=0.275)	0.0 dB	0.0 dB	7.0 dB
85/125 μm (NA=.26)	1.0 dB Loss	0.0 to 2.6 dB Loss	5.5 to 9.0 dB
100/140 μm (NA=0.29)	2.0 dB Gain	0.0 to 4.0 dB Loss	5.0 to 9.0 dB
200/230 μm (NA=0.40)	2.0 dB Gain	0.0 to 8.0 dB Loss	1.0 to 9.0 dB

Table 4-5. Impact of Using Alternative Fiber Types. After Jain, 1994.

corresponding reduction in margin available for other link losses. The difference between the minimum transmitter power and the minimum detectible receiver signal yields a margin of only 7 dB. Careful link analysis is necessary to prevent exceeding this margin. Moreover, the use of optical bypasses is restricted.

c. Media Connectors

Like the MMF-PMD standard, LCF-PMD specifies the use of duplex connectors. These connectors are patterned after the SC and ST simplex connectors and provide polarity keying. Unlike MMF-PMD duplex connectors however, they do not provide port keying to prevent incorrect port connections. They are labeled with port designations to aid network installers in making link connections. (Jain, 1994)

3. Twisted Pair PMD (TP-PMD)

ANSI developed the TP-PMD standard as an alternative to the high cost of optical components. The standard specifies the use of shielded twisted pair (STP) or unshielded twisted pair (UTP) media. The standard also specifies the use of a special coding scheme, signal scrambling, and signal equalization. These specifications are necessary to overcome the drawbacks of high speed transmission across copper-based media.

a. Multilevel Transmission-3 (MLT-3), Scrambling, and Equalization

The standard bit rate for an FDDI coded signal is 125 Mbps. The standard employs a two-level encoding scheme that results in two bits transmitted for each complete cycle. Thus, the maximum frequency for a series of 1 bits is 62.5 MHz. Though this bit rate and frequency is easily managed by fiber optics, attempting to transmit the same signal across copper-based media presents two serious problems.

At 62.5 MHz, signals produce considerable electromagnetic interference (EMI). Though this EMI problem can be averted by using STP cabling, the levels of radiated energy for UTP exceeds the restrictions imposed by the Federal Communications Commission (FCC). These restrictions apply to all electromagnetic emissions generated by frequencies beyond 30 MHz (Shah, 1994).

To reduce EMI, the standard specifies the use of a multilevel transmission scheme, MLT-3, and scrambling. MLT-3 encodes the FDDI two-level signal into a three-

level signal. This creates a 4 bit per cycle rate as shown in Figure 4-4, reducing the maximum frequency to 31.25 MHz. The scrambler then spreads the signal's energy across its frequency spectrum, as it is transmitted onto the media. These two techniques together, reduce the EMI produced in UTP to within restrictions imposed by the FCC. (Jain, 1994)




Figure 4-4. Standard FDDI Encoding versus MLT-3 Coding. After Jain, 1994.

Moreover, at 62.5 MHz, signals are severely attenuated when transmitted across copper-based media. The degree of attenuation varies as a function of signal frequency--higher frequency components are attenuated more than the lower frequency components, causing a distorted signal.

To overcome this problem, an equalization filter is used at the receiver. This filter is designed to emphasize the high frequency components of the received signal while leaving low frequency components unchanged. Although this reduces the impact of severely-attenuated high frequency components, it also emphasizes high frequency noise. (Jain, 1994)

b. Media Specifications

The TP-PMD standard specifies the use of Type 1/2, 150 Ohm STP or Category 5, 100 Ohm UTP cabling. These cables support data transmissions for distances up to 100 meters. Cable interconnection with nodes is accomplished using either RJ-45 connectors for UTP or 9-pin D-connectors for STP. (Mills, 1995)

c. UTP Installation Considerations

Although the effects of EMI in UTP are reduced by MLT-3 coding and signal scrambling, the unshielded characteristic of this cable makes it particularly vulnerable to

other sources of EMI. Jain (1994) provides the following guidance regarding installation of UTP.

(1) Electronic Industries Association (EIA) Category 5 UTP cabling consists of four unshielded twisted pairs of 24 gauge copper wires, enclosed in a thermoplastic jacket. Although this cabling is often used to support telephone communications, data and voice traffic should not be transmitted simultaneously.

(2) Cabling should be routed around florescent lamps and high-powered equipment; cables should remain a minimum of one foot from such devices.

C. PHYSICAL MEDIUM INDEPENDENT STANDARDS

The PHY standards interact with the FDDI physical media dependent standards and media access control standards, to enable data exchange between the data link and physical layers of the OSI reference model. The PHY standards specify the use of 4B/5B and NRZI encoding and decoding, buffers, smoothers, and repeat filters.

1. 4B/5B and NRZI Coding

Fiber-based FDDI networks use a combination of 4B/5B group encoding and NRZI bit encoding, to transmit data across the media. The 4B/5B coding takes 4 data bits and encodes them into a 5 bit code symbol that enhances error protection and control. This results in a symbol rate of 125 Mbaud with an efficiency of 80%--10 symbol bits transmitted for every 8 data bits. NRZI then converts each logical 1 of the code group into a transition between positive and negative optical power levels. Each logical 0 is represented by an absence of a transition. (Shah, 1994)

Each valid code symbol consists of no more than three consecutive 0's, as shown in Table 4-6. This feature ensures there is sufficient power level transitions to enable stations to extract clocking information. Stations use clock information to synchronize with incoming data frames. (Shah, 1994)

Another important advantage to using 4B/5B encoding is that the maximum DC component of a symbol is within 10% of the transmitted AC signal. This reduces the baseline wander of the AC signal, limiting the impact on bandwidth capacity. (Jain, 1994) It also simplifies the design of interface and circuit components (Hammar, 1992).

Symbol	4B/5B Code	FDDI Purpose	Symbol	4B/5B Code	FDDI Purpose
0	11110	Data symbol 0	Q	00000	Quiet line state
1	01001	Data symbol 1	I	11111	Idle line state
2	10100	Data symbol 2	H	00100	Halt line state
3	10101	Data symbol 3	J	11000	Frame start delimiter
4	01010	Data symbol 4	K	10001	Frame start delimiter
5	01011	Data symbol 5	T	01101	Terminates data stream
6	01110	Data symbol 6	R	00111	Reset (status indicator)
7	01111	Data symbol 7	S	11001	Set (status indicator)
8	10010	Data symbol 8	L	00101	Embedded delimiter
9	10011	Data symbol 9	V	00011	Violation (invalid code)
A	10110	Data symbol 10	V	00110	Violation (invalid code)
B	10111	Data symbol 11	V	01100	Violation (invalid code)
C	11010	Data symbol 12	V or H	00001	Halt when received
D	11011	Data symbol 13	V or H	00010	Halt when received
E	11100	Data symbol 14	V or H	01000	Halt when received
F	11101	Data symbol 15	V or H	10000	Halt when received

Table 4-6. 4B/5B Encoding. After Mills, 1995.

As shown in Table 4-6, there are 32 (2^5) possible symbols; 16 symbols represent data, nine symbols are used for control purposes, and seven are interpreted as violations. The violation symbols are not transmitted by stations; they occur as a result of noise on the network.

The symbols Q, I, and H are used to monitor the state of the medium. The Q or quiet symbol indicates the absence of any frame transmissions. A series of 16 consecutive Q symbols may indicate a break in the fiber. The I or idle symbol is transmitted between tokens and frames as a frame preamble. It is the most common symbol and is used extensively for frame synchronization. It consists exclusively of logical 1 bits providing the highest clock content at a frequency of 62.5 MHz. The H or halt symbol is used in conjunction with quiet symbols for communication between two nodes during connection

setup. The V symbol, when received by a node, is also interpreted as a halt symbol. (Jain, 1994)

A key feature of the PHY standard is its ability to communicate the status of the media through line state indications. The standard uses the 4B/5B control symbols to define the following states: quiet line state, idle line state, active line state, master line state, halt line state, and noise line state. Two states in particular, the master line and halt line states, are used between stations during connection line setup. (Jain, 1994)

2. Elasticity Buffer

Each node in a FDDI network has its own clock for transmitting and receiving data. Although these clocks are required to have a frequency accuracy of 50 parts per million (ppm), slight differences in clock speeds between two stations presents a problem.

If a station transmits bits at a faster rate than another station can receive them, there is a potential for loss of data. The same may occur if the transmitting station operates at a slower speed than the receiving station. The purpose of the elasticity buffer is to overcome this problem of differential station timing.

The buffer uses a first-in-first-out (FIFO) storage concept and pointers to track input and output. As shown in Figure 4-5, a receiving station's incoming data stream is loaded into a buffer at the transmitting station's clock rate. The input pointer starts at mid-position of the receive buffer and may drift up or down depending on the difference in speed between the receiving and transmitting stations' clock rates. Once the buffer is half-full, the receive station begins to pull bits from the elasticity buffer. This data pull is tracked by the output pointer. Ideally, if the clock rates of both stations are equal, the pointers will remain at a minimum separation of just half the buffer size. (Mills, 1995)

The buffer size is a function of the maximum allowable difference between the receiving and transmitting clocks and the maximum frame size. A maximum frame size of 4500 bytes equates to 45,000 code bits. With a maximum clock difference of 100 ppm (50 ppm per station), the number of bits gained or lost during a frame transmission is 45000×10^{-4} or 4.5 bits. To handle either a slower or faster incoming clock rate, the buffer needs at least a 9 bit-storage capacity. (Jain, 1994)




Figure 4-5. Elasticity Buffer. After Mills, 1995.

If the clock rate exceeds the 50 ppm requirement, an underflow or overflow condition will result. To manage this problem, elasticity buffers are generally larger than the minimum 9 bit size. Unfortunately, the larger the buffer size, the more the impact on ring performance.

Since buffer output only commences when the buffer is half-full, there is a delay in movement of data. For a 9-bit buffer, a 4.5 bit code delay results in a 36 nanosecond (ns) station delay ($4.5 \text{ bits} \times 8 \text{ ns/bit}$). The larger the buffer, the longer the station delay and ring latency. (Jain, 1994)

3. Smoother

Between each frame transmission there exists an interframe gap consisting of idle symbols. FDDI standards require at least 16 idle symbols between frames to ensure stations have sufficient time to reset and correctly receive the next frame. Transmitting nodes ensure this requirement is met by generating the 16 idle symbols, called a preamble, prior to transmitting each frame.

Although frames are generated with the requisite 16-symbol preamble, clock differences between network stations may cause elasticity buffers to erode the number of idle symbols. The purpose of the smoother is to monitor the length of the preamble, and correct for any deficiencies. It adds symbols back to the preamble by borrowing them from a source. This source may be a gobbler mechanism that strips and stores idle symbols

attached to partial frames circulating the network. When the smoother encounters a short preamble, it borrows from the gobbler. Eventually, however, the bank of idle symbols may be depleted, causing frame synchronization problems. (Jain, 1994)

4. Repeat Filter

As a frame travels the network, it may become corrupted with noise. FDDI MAC protocols are designed to handle this problem by converting invalid symbols of the frame into idle symbols. Unfortunately, not all stations are configured with a MAC. MAC-less stations depend on a repeat filter to eliminate invalid symbols. The operation of a repeat filter is shown in the figure below.




Figure 4-6. Repeat Filter Operation. From Jain, 1994.

When the filter detects invalid symbols within a frame, it begins changing the frame into a series of four halt symbols followed by a continuous stream of idle symbols. The idle symbols are transmitted until the beginning of the next frame is detected. This results in a truncated frame that travels to the next node. (Mills, 1995)

When the frame reaches the next MAC-configured node, the four halt symbols are replaced with idle symbols. At subsequent stations the last two non-idle symbols in the figure are converted into idle symbols as it is retransmitted. After passing through a series of stations, the frame is completely converted to idle symbols. (Jain, 1994)

D. MEDIA ACCESS CONTROL STANDARDS

FDDI uses a timed-token rotation protocol to manage station access to the media. The token is a small control frame that circulates the ring. As each station receives the token, it either repeats the token and passes it on, or it captures the token to transmit data. If the token is repeated, it continues to circulate the ring passing from one station to the next, each station regenerating the token as it passes through its physical and MAC protocols.

When a station captures a token, the station transmits its queued data frames. When it has completed transmission, the station regenerates and releases the token. To prevent any single station from monopolizing the network, timers are used to manage the amount of time each station may hold the token.

FDDI segregates traffic into delay-sensitive synchronous data and asynchronous data. To prevent the disruption of time-critical data, synchronous traffic has priority over asynchronous traffic. Thus, when a station captures a token, it will transmit its synchronous traffic first, then transmit any asynchronous traffic until its timer expires.

Within the asynchronous class, network administrators can assign eight priority levels to data. Using a load-level-based priority scheme, they can control the transmission of traffic under various network load conditions. For example, an administrator may implement a policy that permits all priority-level traffic to be transmitted until the network load reaches 15%. At a 15% load, only priority levels 3 through 8 will be transmitted. Regardless of the policy implemented, when the load reaches 100%, all priority levels are restricted—only synchronous traffic is transmitted. (Jain, 1994)

In certain situations, a station may have a considerable amount of asynchronous data to transmit. To improve the efficiency of data transfer in this situation, network administrators may permit stations to generate restricted tokens. Thus, when the station's timer expires and it is forced to release the token, it marks the token as restricted and transmits it to the next station. Other stations on the ring cannot capture the token unless they are involved in the restricted data transfer or have synchronous traffic to transmit. These restricted tokens are not normally used. (Jain, 1994)

1. Timed Token Access Protocol

A station can only transmit data while holding the token. To prevent a single station from monopolizing the network, stations are limited in the time they can hold a token, called the token holding time (THT). After capturing the token, a station must release the token when the THT expires. It may continue its transmission only after recapturing the token. The time required for the token to circulate the ring and return is called the token rotation time (TRT) and is calculated as follows (Jain, 1994):

$$4.1 \quad \text{TRT} = \text{Number active stations} * \text{THT} + \text{Token travel time}$$

The equation reveals a potential problem between a large THT and network response time. The longer each station holds the token (large THT), the more efficient its transfer of data. Unfortunately, a longer token holding time causes an increase in the token rotation time. This translates into degraded response time for the network as a whole. (Jain, 1994)

a. Claim Process

To reach a compromise between THT and TRT, FDDI MAC protocols use a claim process in which each station bids for a desired token rotation time. The lowest bid becomes the target token rotation time (TTRT) for the network. The THT is the difference between the TTRT and the token travel time. (Jain, 1994)

The claim process is executed anytime the network is initialized, a station joins or leaves the network, or whenever a station detects the token has been lost. It begins when one station starts transmitting special MAC frames called claim frames. Once a station begins transmitting these frames, it does so until the claim process is complete.

The claim frame contains the station's desired TTRT or bid, and is transmitted to the next station on the network. The receiving station compares the bid to its own target time. If its target is lower, it begins its own bidding by negating the bids from the previous station and transmitting its own continuous stream of claim frames. If its target is higher, it ceases its bid and repeats the received claim frames.

The process continues around the ring until a station receives its own bid

signifying it has won the bid process. The station then generates and releases a new token. In the event a station receives an equal bid from another station, a comparison between addresses is made to determine the winner. After completing the process, each station knows and stores the new TTRT. This time is used to control the length of time a station may hold the token for transmission. (Jain, 1994)

b. Valid Transmission Timer

Station MACs use valid transmission timers (TVXs) to control the claim process. When a station detects the absence of valid transmissions, the timer is initiated. When it expires, the claim process is started. It is programmed into each station and can be modified by the network administrator. The value set should provide sufficient buffer to preclude ring reinitialization each time a fault occurs. It should be greater than the sum of the ring latency and the time to transmit a maximum-sized frame. (Jain, 1994)

c. Beacon Process

If the claim process fails to correct a fault condition, the station's MAC initiates the beacon process in an attempt to isolate the problem. Beacons may also be initiated by station management protocols.

The process uses special MAC beacon frames. Transmission of these frames begins anytime a station suspects there is a break in the ring. The frames are generated continuously until the problem is isolated. As each station receives the beacon frames, it repeats the frames to the next station. If a station receives its own beacon frames, the ring is intact and the claim process started to regenerate a new token.

Eventually, the station down-line from the break will begin transmission of beacon frames. Since the break occurs between itself and the previous station, it will not receive its own or any other station's beacons. After a period of time, the station concludes the break is located prior to itself. (Jain, 1994)

d. Management of Asynchronous Traffic

To manage its network access, each station uses a series of timers and counters. These include the following (Shah, 1994):

- (1) T_Opr. The time result established by the claim process and

loaded in the T_Opr register; represents the TTRT.

- (2) Token Rotation Timer (TT). The timer used to measure successive arrivals of the token at a station.
- (3) Token Holding Timer (THT). The timer used to control the period a station can hold the token for transmission of asynchronous traffic.
- (4) Late_Ct. A counter that increments when the station's TT expires.

Following ring initialization, the winner of the claim process generates and releases the token. Meanwhile, each station sets its TT to T_Opr and begins decrementing the timer. When a station captures the token, it loads the time remaining in its TRT into the THT, resets the TT to T_Opr, and begins transmitting data.

When the station completes its transmission, the token is released and the THT is reset to zero. Should the THT expire before the station has transmitted all of its traffic, transmission of the current frame is completed and then the token is released.

If a station's TT expires before the token arrives, it increments its Late_Ct counter to one and resets the TT to T_Opr. The Late_Ct flags the station to prevent it from capturing the token for asynchronous transmissions. When the token finally arrives, the Late_Ct counter is reset to zero and the token retransmitted. The TT is not reset.

Since the TT continues to decrement as the token passes the station, the station will have less time to transmit when the token is captured on the subsequent round. Thus, the station does not cause an undue delay of the token for other stations when it eventually captures the token. If the TT expires two consecutive times, the station reinitializes the ring.

It is important to note that since the station's TT is reset when the token is received, the longer the station transmits on the current rotation, the less time remaining in the TT for transmission on the next token capture. In fact, if the TT expires, the station is prevented from capturing the token altogether.

Moreover, a late token condition at one station often means the TTs at other

stations have also expired. Since these stations are not allowed to capture the token, the speed of the token around the ring increases. This varying speed of the token makes the network self-regulating. It provides an average token rotation time equal to the TTRT, and ensures a maximum station access time of $(n - 1) * TTRT + 2 * \text{Ring Latency}$, where n is the number of stations (Shah, 1994).

e. Management of Synchronous Traffic

Synchronous transmission capability is an optional feature for an FDDI station. To enable this service, the station must be equipped with the requisite hardware that contains preprogrammed timer functions. (Mills, 1995)

Synchronous data transmissions take priority over asynchronous traffic. Thus, when a station captures the token, it always transmits its synchronous traffic first. Even if the TRT has expired and the Late_Ct incremented—a late token condition—a station may still capture the token if it has synchronous data to transmit.

The duration of synchronous transmission is controlled using a pre-allocated value stored in a synchronous timer. When a station is turned on, it requests a synchronous bandwidth allocation from a central manager. Based on the overall bandwidth and T_Opr, the manager assigns a percentage of the bandwidth to the station. This percentage is based on 64 Kbps units or 125 μ s time slots. Each station may have a different allocation. (Mills, 1995)

When the token arrives at a station with queued synchronous traffic, the station begins transmitting the traffic until the queue is emptied or the synchronous timer expires. When either occurs, the THT is loaded with the time remaining in the TT and asynchronous traffic is transmitted. If the TT has already expired, then the token is released. (Shah, 1994)

During synchronous transmissions, the TT continues to decrement. In conditions of high synchronous traffic flow, the station TTs may expire before asynchronous transmissions are commenced or completed. For this reason, synchronous transmissions are limited to less than $2 * TTRT$ (Jain, 1994).

2. Frames

Figure 4-7 shows the structure of an FDDI frame. The frame has a maximum length of 4,500 bytes, providing an information content of 4,486 bytes when using 16-bit addressing, or 4,478 bytes when using 48-bit addressing. The figure indicates the size of each field in terms of the number of symbols that comprise the field. The reader will recall that a symbol is a grouping of four bits of data that will be encoded into five bits for transmission. Thus, the maximum length of a frame is 36 Kbits or 9,000 symbols. (Shah, 1994)




Figure 4-7. FDDI Frame.

The *preamble* consists of a minimum of 16 idle symbols transmitted before the beginning of each frame. It provides the necessary separation between frames to allow stations to delineate one frame from another. It is also used to recover clocking information. The *starting delimiter* field marks the beginning of the frame with a unique J-K symbol pair.

The *frame control* field identifies the class of frame, the type of frame, and the length of the *address* fields. The *address* fields identify the transmitting and receiving stations, and can accept either 16-bit or 48-bit addresses.

The frame's *information* field is a variable length field containing data from the higher layers of the OSI reference model. The *frame check sequence* field contains a 32-bit cyclic redundancy check (CRC) polynomial value used for error detection--the same CRC

used in IEEE 802 protocols (Jain, 1994). The *ending delimiter* field contains a T symbol to identify the end of frame data.

The last field of the frame, the *frame status* field, consists of three indicators as shown in Figure 4-8. These indicators can consist of either an R symbol (Reset) or an S symbol (Set). When the frame is transmitted, the transmitting station sets the indicators to R.




Figure 4-8. Frame Status Field.

When a station detects the incoming data frame, it immediately begins repeating the frame to the next station. When the station receives the destination address within the frame, it compares that address to its own. If it makes an equal comparison, the frame is then copied into its receive buffers. If the frame is copied without errors, the station changes the *address recognized* and *copied* indicators within the frame, to an S symbol. The frame eventually returns to the originating station to signal satisfactory receipt of traffic. (Mills, 1995)

As each station receives the frame and transmits it back onto the network, it checks the frame for errors using 32-bit CRC. If an error is detected, the *error detected* indicator is changed to an S symbol by the station. Thus, any station on the network may detect an error. Once a station has changed the indicator to an S, subsequent stations cannot change the value even if their CRCs indicate the frame is correct. When the frame returns to the

transmitting station, it discovers the indicator has been set and retransmits the frame. (Mills, 1995)

a. Types of Frames

The contents of the *frame control* field is shown in Figure 4-9. It indicates the class of service (synchronous or asynchronous), the length of the *address* fields, and the type of frame. If the class of service is asynchronous, the priority level will also be indicated. The types of frames include tokens, LLC frames, MAC frames, SMT frames, and void frames. (Jain, 1994)




Figure 4-9. Frame Control Field. After Mills, 1995.

A token is a special frame that permits stations to transmit onto the media. It consists of only three fields, a *starting delimiter*, a *frame control* field, and an *ending delimiter*. Each field is two symbols long: the *starting delimiter* field contains the unique J-K symbol pair used to mark all frames and tokens; the *frame control* field indicates whether the token is restricted or unrestricted; and the *ending delimiter* field contains a unique T-T symbol pair to mark the end of the token.

The LLC frame is the most common frame and is used to carry higher layer information between stations. MAC frames are frames defined for use at the MAC layer only, and include claim and beacon frames. SMT frames are reserved for carrying station management data between network nodes. Lastly, void frames are used as markers for

various functions by the MAC protocols. In particular, they are used to reset the valid transmission timers at all stations. (Jain, 1994)

b. Addressing Format

FDDI supports broadcast, multi-cast, and individual addressing. The structure of the *address* field is shown in Figure 4-10.




Figure 4-10. Address Fields. After Mills, 1995.

3. Frame Stripping

FDDI traffic consists of data frames, frame fragments, idle symbols, and tokens. To prevent the buildup of frames and fragments, the MAC protocol employs several techniques designed to remove obsolete traffic.

In general, the transmitting station is responsible for removing its own frames. As stations receive incoming frames, they automatically retransmit the data back onto the media through their output port. When the source address field of the frame is received and recognized as its own, the station begins stripping the frame, and transmits idle symbols in its place. The result is a fragmented frame consisting of a starting delimiter, frame control field, destination address, and a series of idle symbols.

A similar situation leads to truncated tokens. Stations desiring to capture the token must repeat the incoming frame until the frame control field is received and recognized. At that moment, the token is captured and idle symbols transmitted in place of the rest of the

token. The result is a truncated token consisting of a starting delimiter followed by a series of idles.

Frame fragments and truncated tokens continue to circulate the ring from station to station until reaching the active station holding the token. This station's MAC then receives and discards all traffic. This makes the network self-regulating. (Mills, 1995)

Stations may employ a gobble feature. The gobble strips and stores idle symbols from the end of frame fragments. These symbols are used by the PHY smoother to ensure proper length of frame preambles as discussed earlier. (Jain, 1994)

E. STATION MANAGEMENT STANDARDS

SMT standards enable the fully distributed management of an FDDI network--any network station can monitor and manage the operation of the network. It also supports the automatic detection, isolation, and recovery of faults. SMT features apply to the MAC, PHY, and PMD sub-layers of network stations. (Jain, 1994)

SMT consists of three components: connection management, ring management, and frame-based services. Connection management manages functions related to the physical connection and communication between two stations. Ring management is used to monitor the ring for faults and subsequently, to isolate them. Frame-based services apply to the general operation and monitoring of the network. (Jain, 1994)

1. Connection Management

Connection management (CMT) manages the optical media interface, the physical connection between ports, and the configuration of PHY and MAC elements of network nodes. These functions are accomplished by its three components, entity coordination management (ECM), physical connection management (PCM), and configuration management (CFM). (Jain, 1994)

Each node uses one ECM component to manage its media interface and ports. Its ECM also interfaces with its PCM to coordinate fault tracing and path testing within the node. Fault tracing is initiated by ECM should the media become unavailable. If an optical bypass is installed in the node, it is activated by the node's ECM.

PCM is used to manage the connection of a node's port to a neighbor's port. It

exchanges bits with adjacent ports using line state indications. These bits allow the PCM to coordinate its port configurations with the connected ports. After coordinating the configuration, it signals CFM to complete the configuration. A PCM is required for each active port in a node.

As its name implies, CFM manages the configuration of a node's MAC and ports. After receiving instructions from the PCM, it makes the port connection or disconnection as appropriate. It tracks active connections using paths that represent the logical flow of data on the primary and secondary rings between nodes. It also tracks the data paths within the node.

Furthermore, CFM monitors port connections during normal and wrapped conditions to detect illegal port connections. These illegal connections (see Chapter III for connection rules) are detected and reported by the node's CFM. The actual management of the condition varies depending on the condition, and in some cases on the vendor's implemented protocols.

The collection of CMT components provide a number of useful services that maintain network operation. These functions include link testing, link error monitoring, fault tracing, ring scrubbing, ring purging, and wrapping avoidance. (Jain, 1994)

2. Ring Management

While CMT is used to resolve problems at the physical layer between two stations, ring management (RMT) is used to handle MAC-level problems across the network. RMT functions include announcing the presence of faults, monitoring the restricted tokens, and detecting and resolving duplicate addresses. (Jain, 1994)

The reader will recall that if the claim process fails to resolve a communications problem, the station that initiated the process begins transmitting beacon frames. This process continues until the station downstream of a break concludes the fault lies somewhere between itself and its upstream neighbor. Once this has been determined, the station sends a directed beacon to the network's management station.

If nodes are permitted to generate restricted tokens, it is necessary to monitor this operation for malfunctions. This ensures other stations are not restricted from gaining access

to the media due to a station's failure to remove its own restricted token. Each station monitors the restricted token operation with a timer. When the station sees a restricted token, it starts the timer; when the timer expires, RMT initiates the claim process which eventually results in the generation of a new unrestricted token.

Lastly, RMT is used to detect and resolve duplicate addresses. This is necessary to prevent stations from stripping frames intended for other stations, due to address conflicts. There are several possible algorithms used to detect and resolve this problem.

3. Frame-Based Management

Frame-based management involves the use of special SMT frames to initialize, monitor, and manage normal ring operations. These frames use the standard format of FDDI frames, with some modifications to the information field to enable the management functions. The exchange of these frames follows the normal convention for asynchronous transmissions.

There are numerous types of SMT frames that enable the frame-based functions. These frames are separated into three classes: announcement frames, request frames, and response frames. Announcement frames are generated periodically by the network stations to report connection, control, and management related data about themselves. Request frames are generated by stations to request information from other stations. Response frames are generated in response to another station's query. The following is a summary of SMT frame types from Jain (1994).

a. Neighbor Information Frame

Neighbor information frames (NIF) are used by MACs to determine their upstream and downstream neighbors. Periodically, a MAC will transmit a NIF announcement or a NIF request to downstream stations. This frame contains information about the requesting MAC and its upstream neighbor. The broadcast address of this frame is set to *next station addressing*. (Jain, 1994)

The first station to receive the frame retains the sending MAC's address as its upstream neighbor. As it continues to retransmit the signal, it sets the *address recognized* field. When the station captures the token, it transmits a NIF response frame that tells the

requesting station its own address. The requesting station then records this address as its down stream neighbor.

Any other station on the ring can monitor the transmission of these request and response frames. It uses information about the senders to construct logical or physical ring maps of the network, by generating lists of successive MACs or PHYs respectively. Information in the NIF indicates the type of node, the number of MACs, number and types of ports, state of the ports, and the internal configuration of MACs and PHYs.

In addition to monitoring NIFs, a station's MAC may demand information to construct the maps. It does so by generating a series of NIFs to successive upstream stations or transmitting one NIF with a broadcast address. Such features are used by network managers as an aid to trouble shooting faults or implementing a new network.

The periodic interval between transmission of NIF announcements can be set by network managers between 2 and 30 seconds. Jain (1994) recommends the individual station intervals be selected randomly using a uniform distribution. This prevents an excessive number of NIFs on the network at the same time. The default value is 30 seconds.

b. Status Information Frame

Status information frames (SIF) are used between nodes to exchange information about themselves. The information exchanged may be in reference to a station's configuration or a station's operation. The exchange of information is initiated by a request for data from one station to another. The information is used for configuration, performance, and fault monitoring. (Jain, 1994)

The response to an SIF configuration request contains detailed information about the responding station's software and hardware configurations. It also contains information regarding its upstream and downstream neighbors. Information that may be included in the response is show in Table 4-7.

The response to an SIF operation request contains information about the station's various timers, the TRT thresholds for asynchronous priority levels, number of frames lost, number of frames in error, number of frames not copied due to buffer overflow, and the number of times parameters in the station have been changed by local or remote

- | |
|--|
| <ol style="list-style-type: none">1. Type of node2. Number of MACs and their addresses3. Number of ports, type of port, and state of each port4. Type of remote port to which each port is connected5. SMT versions supported by the station6. Wrapped or unwrapped condition of the station7. Port connection policies supported by the station8. Station delay due to primary and secondary ring latencies; reported in bytes9. Upstream and downstream neighbors' addresses10. Internal arrangement of MACs and ports11. Time stamp for the frame |
|--|

Table 4-7. Information Contained in an SIF Configuration Response. From Jain, 1994.
management commands. This information is used to monitor the performance of the station,
and for fault isolation.

c. Echo Frames

Echo frames (ECFs) are used by SMTs for testing between two stations. One station initiates the testing by transmitting an ECF containing a maximum of 4454 bytes of data to another station. The receiving station responds by transmitting the same information back. These frames are used to determine if the receiving station's port, MAC, and SMT are functioning. (Jain, 1994)

d. Resource Allocation Frame

Resource allocation frames (RAFs) are used by stations to announce, request, and receive network resources. Currently, they are only used to allocate synchronous bandwidth to stations configured for and requiring this service. These frames may also be used by the network manager to monitor allocations. (Jain, 1994)

e. Request Denied Frame

Request denied frames (RDFs) are transmitted anytime a station cannot respond to another station's request. Such a problem may occur, for example, if the stations' SMT versions are different. A reason for denial is included in the RDF. (Jain, 1994)

f. Status Report Frame

Status report frames (SRFs) are used by stations to report certain events and abnormal conditions. Periodically, stations will verify their status and check for changes in events such as a neighbor change, a configuration change, or illegal connection attempt. If a change or abnormal condition has occurred, the station reports it using an SRF. (Jain, 1994)

g. Parameter Management Frame

Any station on a network can monitor and manage the operation and parameters of other stations remotely. This management is accomplished using parameter management frames (PMFs). (Jain, 1994)

Two types of PMF frames are used to manage network operation. PMF *get request* frames are used to query stations for their operating parameters. PMF *set request* frames are used to modify the parameters. PMF *get request* frames may be broadcast, while PMF *set request* and all PMF responses to requests are individually addressed. Only the PMF *get response* frame must be implemented; other frames may be suppressed to prevent remote control of network operations.

Though not specified by SMT standards, the PMF may contain authorization data to control which stations may execute network management features. The implementation of authorization functions is vendor specific.

h. Extended Service Frame

Extended service frames are used by vendors to implement a new frame and its supporting protocol. It may be rejected by a station that does not recognize the implementation of the new protocol. (Jain, 1994)

F. SUMMARY

ANSI developed four different PMD standards of which three were presented: MMF-PMD, LCF-PMD, and the TP-PMD. The MMF-PMD specifies the use of 62.5/125 multimode fiber. Mixing of different fiber types is permitted, however, additional link losses may result due to fiber size and numerical aperture differences. The fiber must provide a minimum bandwidth of 250 MHz. The maximum link distance allowed is 2 km.

The standard also prescribes the use of transceivers that provide a minimum power difference of 11 dB between the transmitter and receiver. This difference provides considerable margin for signal losses caused by attenuation, connectors, and splices.

Vendors manufacture two types of connectors: duplex and simplex. The MMF-PMD prescribes the use of MIC duplex connectors that provide station and polarity keying. These connectors provide good coupling efficiencies and are easy to connect and disconnect.

The LCF-PMD standard relaxes the transceiver specifications--the difference between transmitter and receiver power levels is only 7 dB. Although these relaxed specifications translate into cost savings, they also reduce the available margin for other link losses. For this reason, it is necessary to limit the maximum link distance to 500 meters.

Like the MMF-PMD standard, the low-cost standard also recommends the use of 62.5/125 fiber and duplex connectors. The connectors are based on the SC and ST simplex connectors. They do not provide station keying features like the MIC connectors.

The TP-PMD standard prescribes the use of STP and UTP media. To enable communications across this media, it is necessary to use a special coding scheme, signal scrambling, and signal equalization. The standard can achieve 100 Mbps transfer rates up to a distance of 100 meters.

The PHY standard specifies the use of 4B/5B and NRZI encoding to transmit data across the medium. This encoding scheme ensures there is sufficient power level transitions for station clocking purposes. It also reduces the baseline wander of the signal.

Of the 32 possible encoding symbols, 16 symbols represent data, nine are used for control signaling, and seven indicate the presence of noise. Combinations of certain symbols are used to indicate line states.

Furthermore, the PHY specifies the use of elasticity buffers, smoothers, and repeat filters. Elasticity buffers are used to prevent the loss of data due to clock differences between stations. Smoothers monitor the preambles of frames to ensure that at least 16 idle symbols are present for station-to-frame synchronization. Lastly, repeat filters are installed in MAC-less stations to help eliminate invalid symbols that are circulating the network.

MAC standards specify the protocols that govern access to the media to include the

timed token protocol, lost token procedures, and the structure and control of frames. The timed token protocol supports both synchronous and asynchronous transmissions. If a station has both types of traffic to transmit, the synchronous traffic takes priority and is transmitted first. To manage transmission of asynchronous traffic, eight priority levels may be designated.

To prevent any station from monopolizing the network, timers and counters are used. One timer in particular, the THT, limits the period that a station may hold a token for asynchronous transmissions. A separate timer is used to monitor synchronous transmissions.

If a station detects a lost token or other ring problem, it initiates the claim process to regenerate a new token. If the claim process fails, the station initiates the beacon process to locate the fault.

The MAC standard also prescribes the structure of frames. There are several types of frames to include tokens, LLC frames, MAC frames, SMT frames, and void frames. These frames are generated and subsequently removed by a transmitting station. However, due to the manner in which frames are repeated as they travel through a station's protocols, parts of the frame are retransmitted before the station recognizes a frame as its own. These fragmented frames continue to circle the ring until they are stripped by downstream stations.

SMT consists of three components: CMT, RMT, and frame-based management. These components are used to initialize, control, and monitor MAC, PHY, and PMD components of the network. Moreover, they are crucial to fault isolation and recovery.

V. NETWORK DESIGN

A. INTRODUCTION

Organizations develop networks to improve their efficiency, coordination, responsiveness, and reach. The need to develop a network may be based on a particular problem that impairs the organization's ability to operate or an opportunity to improve its current operations. If a network solution is deemed necessary, achievable, and desired, measures are taken to find and implement the right solution.

The purpose of this chapter is to examine the procedures involved in finding and developing a network solution. It will address the design of networks from a systems approach, present a methodology for designing an FDDI topology, and discuss the decisions and procedures for implementing an FDDI infrastructure.

B. STRATEGIC PLANNING

For decades, many organizations have made the common mistake of applying technology to resolve problems and automate processes. Often, the end-result was a stovepipe system that failed to meet the expectations of its stakeholders and users. Today, organizations are moving away from this approach to building information systems.

Businesses are learning the value of information resources planning. Now, top executives and their information resource officers develop information resource plans that govern their acquisition of information systems. The goal is no longer just to automate, but to develop and design systems that provide a strategic advantage. This entails re-engineering old processes to improve efficiency and responsiveness, and rethinking strategies in a global sense to improve the organization's range and reach. Ironically, this change has caused a paradigm shift within many organizations--a direct result of advances made in networking and telecommunications technology itself.

Federal agencies are also learning the value of information resources planning. Though their acquisition of information systems is controlled by federal regulations and policies, they are developing and annually revising five-year information resource plans. The goal of these plans is to increase the efficiency and effectiveness of acquisition strategies

while simultaneously minimizing the cost of controlling information. Achieving this goal is facilitated by acquiring the right systems based on the functions and missions of the organization. Failing to develop and maintain an information resource plan may result in fraud, waste, and abuse.

C. DEFINING THE PROBLEM AND FINDING A SOLUTION

Strategic planning often identifies a need for a solution without accurately defining and bounding the problem. Before committing to the design of a network, it is necessary to expend time and perhaps money to properly define the problem. Only then can the stakeholders understand the scope of the problem and the business opportunities afforded by a solution.

Defining the problem is more complex than just noting an inadequacy in a computing or networking process. It involves defining the problem in terms of the missions, functions, and objectives of the organization. Network designers agree this step is necessary to find the best solution to meet the needs of the organization.

Designers also concur the approach to network design should be based on a systems perspective. That is, the solution should extend beyond the physical design and implementation of a network, and encompass the other important aspects that make a solution successful--security, management, maintenance, compatibility, user training, and documentation. Failing to address these important issues will only lead to an inadequate solution.

There are countless approaches to designing and developing a network. At a minimum, however, the design process will involve the general steps outlined in Table 5-1. These steps are an adaptation of a more in-depth network design approach recommended by Fitzgerald (1993).

1. Determining the Feasibility of a Solution

Once the problem has been defined, it may be necessary to determine the feasibility of pursuing a solution in order to satisfy the skepticism of stakeholders. This will likely involve an economic risk analysis to determine if the costs and benefits of a solution justify its acquisition. The goal is not to justify a particular design, but to illustrate the benefits of

Step	Description
1	Determine feasibility of implementing a solution.
2	Determine new system requirements based on the needs of the organization.
3	Define the geographical scope of the new network.
4	Select a network standard or standards that can achieve the network requirements in terms of reach, range, and responsiveness. Define its topology.
5	Document hardware and software requirements.
6	Calculate network costs.
7	Determine feasibility of implementing the network.

Table 5-1. Network Design Methodology.

a solution against the energy and resources expended in developing and implementing that solution.

It may also be necessary to analyze some of the risks to the development of a successful solution. Such an analysis may include a technical risk assessment to determine if the technology and expertise can be acquired to implement and maintain the system, an operational risk assessment to determine if the recommended solution will be accepted and used by operators, a schedule risk assessment to determine if the solution can be implemented within a given time-frame, and a political risk assessment to determine if management is committed to developing the system.

Though largely subjective in nature, these assessments nonetheless provide useful information that can be used to prepare training and implementation plans that mitigate the risks. At a minimum, the organization's management will be able to gage the potential resistance to the new system and any resulting change in the culture of their people. Recognizing and anticipating these factors may mean the difference between a successful solution and an expensive disaster.

2. Determining System Requirements

Once a network solution is deemed necessary, achievable, and desired, the next stage is to identify the requirements of the new system. Determining the system requirements begins with identifying the networking and processing needs of the organization, based on its mission and functions. These needs are then translated into networking goals. Evaluation criteria are established for verifying successful achievement of these goals in terms of

transaction time, processing time, accuracy, reliability, documentation, and training.

Next, a functional analysis is conducted on the organization's current operations to identify and document business processes, software application requirements, database requirements, and network interconnectivity requirements. In support of this analysis, a baseline assessment of existing computer systems is conducted. This baseline assessment provides a functional understanding of current networks, an opportunity to identify networking problems, and insight into the hardware and software components of existing systems that will influence the design of a new network. Furthermore, a message traffic analysis may be completed to document transaction and processing requirements.

The results of these analyses and the networking goals defined earlier are used to prepare new network requirements. These requirements are expressed in terms of network capacity, response time, interconnectivity, management, and future growth. These requirements are then prioritized to differentiate between must-have versus nice-to-have requirements. Finally, the requirements are recorded, and then added to the network documentation.

3. Defining the Geographical Scope

Defining the geographical scope of a project is necessary to prevent designers from limiting their focus to a particular solution. Indeed, the geographical scope of a project may span hundreds or even thousands of miles. A solution for such a network may entail developing a fiber-based backbone that interconnects several networks. Or perhaps a better alternative is to lease private lines from a telecommunications provider.

On a smaller scale, selecting Ethernet to meet the processing needs of 20 users may be a practical solution if these users are confined to a single room or even building. If, however, these users were spread out across a college campus, Ethernet may prove to be inadequate.

The point to emphasize is that the geographical scope of a project affects the selection of a network standard, the design of its topology, the selection of hardware and software components, and the evaluation of alternatives. It is necessary to define the scope before selecting a network solution.

4. Selecting a Network Standard and Defining Its Topology

After the geographical scope of the project has been evaluated, the next step is to select a network standard that achieves the requirements documented in earlier steps. In many situations, more than one standard may be a viable solution. If such is the case, the organization should consider and evaluate each standard as a potential alternative. Selecting the best alternative will depend on the results of the final feasibility analysis.

After selecting a standard, the next task is to define its topology and select the supporting media. The topology actually consists of three different presentations: a logical topology that illustrates the signal paths between nodes, a physical topology that illustrates the physical connection of media components and network nodes, and a layout that shows the actual media paths through the conduits and wiring closets of a building, or through the trenches between buildings. Each presentation may take on a ring, bus, star, or mesh configuration depending on the design decision. (Jain, 1994)

The logical topology depends largely on the network standard. In the case of Ethernet, the logical topology is based on a bus configuration. In the case of token-ring and FDDI, the logical topology uses a ring configuration. Although, the logical topology is easy to conceptualize, developing the physical topology and layout is more difficult. The presentation depends on the relative location of network components and the manner in which they are interconnected.

Network designers use the logical topology of the selected standard to guide them in the design of the physical topology. Indeed, a network designed with the same physical and logical topologies requires the least amount of media to interconnect the various components. This translates into cost savings for the organization. (Jain, 1994)

Choosing an appropriate media depends on the specific application and the network standard chosen. At a minimum, the network designer must weigh the tradeoffs between unshielded twisted pair, shielded twisted pair, coaxial, and fiber. Influencing factors include cost, capacity, internode distances, signal attenuation, and transceiver requirements.

5. Documenting Hardware and Software Requirements

The purpose of this step is to document the hardware requirements identified during

the design of the physical topology and layout of the new network. Moreover, it involves identifying and documenting the software required to support the networking and processing needs of the organization. The goal is to produce hardware and software purchase lists to be used for cost analysis.

The hardware purchase list should document all the components required to support network operations such as computers or work stations, servers, hard disks, adapter cards, media, connectors, and concentrators. It should also include those components required to support internetworking, telecommunications, security, backup, recovery, and maintenance.

Software purchase lists should identify the software required to support networking, internetworking, and general processing. This includes computer operating systems, network operating systems, security software, virus protection software, application programs, internetworking software, and network management programs for diagnostic testing and maintenance.

6. Calculating Network Costs

Determining network costs for the life of the system is a necessary step prior to weighing the costs against the benefits of the system. These life cycle costs include the cost of hardware, software, telecommunications, network management, testing, and maintenance. They may be directly related to the network implementation and operation, or indirectly related as in the case of personnel training.

7. Determining the Feasibility of a Network Solution

Once the network has been defined, it may be necessary to conduct a second feasibility analysis. This analysis will be more accurate than the first, now that the network costs have been determined. The results of this analysis are presented to the stakeholders to get final approval for implementing the recommended solution.

It is important to note the final decision to implement a design plan rests with the stakeholders. Though the costs and benefits analysis may provide indisputable results that the tangible and intangible benefits will exceed the costs of the system, the final decision may be less objective than most stakeholders would like. This is particularly true if there is considerable technical, operational, schedule, and political risk.

If the stakeholders are dissatisfied with the solution, the network designers may present a more feasible alternative based on feedback. Although this implies finding a satisfactory solution rather than a requirements-driven solution, it is a reality that all organizations face. Thus, network designers may review and re-prioritize the system requirements, survey other network solutions, and prepare another proposal for evaluation. The goal is to find the right solution to meet the business functions of the organization, within economic, operational, and schedule constraints.

D. METHODOLOGY FOR DESIGNING AN FDDI NETWORK TOPOLOGY

In general, the steps leading to the selection of any networking solution parallel the processes described in the previous sections. The actual design of the networks, however, may follow different conventions. The design process is influenced by the geographical scope of the problem and the standard selected. This is particularly true in the design of an FDDI network.

An FDDI network consists of peculiar hardware and software configurations that affect its design. Moreover, the design of a solution requires careful analysis of each link to ensure adequate performance of the network. Therefore, the purpose of this section is to introduce one approach to designing an FDDI network that captures these peculiarities and special link considerations. The focus is on the design of a single-building network.

1. Fiber-Based Design

Designing an intra-building fiber-based network involves defining the network topology and selecting an appropriate fiber type. One approach to defining the topology is presented in Table 5-2.

This methodology begins with a study of the internetworking requirements, the baseline assessment, and building diagrams. The intent is to develop a conceptual representation of the network before actually designating and connecting nodes. Several configurations are possible and each should be considered. For example, one configuration may be to connect all nodes onto the trunk as DASs. Another alternative is to interconnect the stations through concentrators onto a backbone. Each has its advantages and drawbacks, and should be evaluated. Figure 5-1 is an example of how such a conceptual representation

Step	Procedure
1	Determine the location of all computer equipment to be networked.
2	Determine the location of existing wiring closets, patch panels, conduits, and installed cable.
3	Draw a preliminary logical topology diagram. a. Determine which computer equipment are critical nodes of the network and designate them as potential DASs; identify groupings of potential SASs for interconnection through DACs or SACs. b. Interconnect the DASs and DACs to form the dual-ring. c. Construct logical tree diagrams representing the connection of SASs to parent DACs or SACs. d. Determine which DASs are candidates for optical bypass switches and mark them accordingly.
4	Using the logical topology diagram, draw a cable layout plan using a building blueprint as a guideline. The layout plan should take advantage of existing wiring closets, patch panels, conduits, and installed cable.
5	Identify potential location of additional patch panels, wiring closets, and splices. Mark them on the layout plan accordingly. Due attention should be given to providing adequate security to network components when selecting these locations.
6	Determine if established networks will be connected to the trunk. Identify the location of bridges and routers necessary to support these interconnections on the layout plan.
7	Calculate the number of nodes that comprise the network including bridges and routers. If the number exceeds 80, consider segmenting the network into two or more sub-networks, interconnected through bridges and routers (Shah, 1994). Adjust the logical and layout plans accordingly.
8	Conduct a traffic analysis of the trunk. If the estimated traffic approaches 50% of the FDDI bandwidth, again consider segmenting the network into two or more sub-networks (Shah, 1994). Allow a 10% growth factor in the calculations. Adjust the logical and layout plans accordingly.
9	Conduct a link analysis of the layout plan on a link-by-link basis: a. Trace the link's path to determine fiber length, and number of connectors, adapters, and splices. Links that terminate at stations containing an optical bypass should be extended to successive links. This permits link analysis of bypassed nodes in a worst case situation. b. Conduct a loss analysis of the link to determine if the optical power loss exceeds FDDI limits. This is necessary to preclude the installation of a link that fails to provide the performance standards due to excessive loss. Calculate losses using simple loss computation or 99-percentile loss computation techniques. Component specifications should be obtained from the manufacturer. c. Conduct a bandwidth analysis of each link.

Table 5-2. FDDI Design Methodology.

might appear.

Once the conceptual representation is completed, the next step is to develop the logical topology of the network. First, the critical elements of the network are identified and designated potential dual-attachment nodes. These elements may be concentrators or stations. Other components are designated as potential single-attachment nodes.

Next, the nodes are logically grouped to form the trunk and trees of the topology. Beginning with DACs and DASs, the nodes are interconnected as illustrated in Figure 5-2. It is helpful at this point to label all nodes, ports, and signal paths. This allows the designer




Figure 5-1. Conceptual Design Example.

to verify port connections, as well as the reconfiguration capability of his design. It is also useful during network installation, and thus becomes a valuable part of documentation.

The logical tree diagrams are prepared in the same manner. Beginning with one tree, the network is expanded to interconnect the single-attachment stations. The diagrams should indicate the same detail as the trunk diagram. Each additional tree is completed in the same manner.

The network designer will discover a myriad of alternatives to interconnect stations that comprise the trees. For example, Figure 5-3 shows two options for interconnecting six stations using different configurations of concentrators. Each option provides a different level of survivability. Moreover, using three concentrators as opposed to one will increase




Figure 5-2. Logical Topology Example.




Figure 5-3. Alternative Tree Configurations.

the network cost. The use of concentrators can be eliminated completely by interconnecting the stations as DASs through patch panels, as shown in Figure 5-4.

The designer should consider these various alternatives. The final design will depend on the number of stations, their relative location, the desired level of reconfiguration, and




Figure 5-4. Interconnecting Stations using Patch Panels.

cost. In general, cost is a function of the number of concentrators, the number of ports, and adapter cables required to interconnect the stations.

After all the logical tree diagrams have been drawn, the designer may consolidate the individual trunk and tree diagrams into a comprehensive diagram. This may, however, prove to be overwhelming. If the interconnection points between the trunk and tree diagrams are clearly marked, constructing a comprehensive diagram is not necessary.

After the logical topology has been diagramed, the next step is to construct the physical topology diagram. To accomplish this task, it is necessary to determine the location of existing conduits, wiring closets, and pre-installed cabling. Every effort should be made to take advantage of this existing infrastructure.

The physical topology should include all the hardware components necessary to complete the interconnection of the nodes as depicted in the logical diagrams. Moreover, it should include components such as encryption devices and secure cabling that are necessary to provide the degree of security required. Figure 5-5 is an example of a physical topology diagram.

To construct the topology diagram, the designer must be familiar with the characteristics of FDDI products available on the market. These products include patch cable lengths, adapter cable lengths, types of cables, and concentrator port configurations. Thus, it is beneficial for the designer to conduct a preliminary market survey to determine what products are available and their cost, before preparing the physical topology diagram.

After the physical topology has been prepared, the next step is to ensure the links will meet FDDI performance standards. This entails evaluating each link to determine its length, signal loss, and expected bandwidth. These analyses are necessary to prevent the design of links that later fail to perform adequately.

There are two methods for evaluating link losses: the simple loss computation method and the 99-percentile loss computation method. The first method is less precise; it presumes a maximum loss value for each connector, splice, and fiber link. The second method uses the loss variance of each component to derive a more accurate calculation of the signal losses. (Jain, 1994)




Figure 5-5. Physical Topology Diagram.

Both techniques begin with tracing the proposed path of the link and identifying the number of connectors, splices, and the fiber length along the cable that extends from one active node to the next. If an optical bypass is to be installed at the receiving node, the node is considered failed and tracing continues to the next bypass-less node. This presumes a worst case situation. If the extended link fails to meet loss limitations, an intermediate node should be installed or the optical bypass removed.

The simple loss computation method uses the maximum loss value for each component to determine the total link loss. The total link loss is then compared to the difference in power levels between the transmitter and receiver to determine the remaining margin. These steps are illustrated in Table 5-3.

The remaining margin should be at least 1 dB to allow for transceiver aging, temperature variations, and component specification variations. Furthermore, if using single-mode fiber, an additional dB margin is needed for reflection and dispersion losses. If the link fails to provide the necessary margin, another active node should be added to the link, at an

From Path Analysis	
1. Total length of fiber	200 meters
2. Number of connectors	6
3. Number of splices	1
4. Number of bypasses	1
Component Specifications	
5. Fiber attenuation	1.5 dB/km
6. Loss per connector	0.8 dB
7. Loss per splice	0.5 dB
8. Loss per bypass	1.5 dB
9. Minimum output power of transmitter	-20 dBm
10. Minimum detectable power of receiver	-29 dBm
Calculations	
11. Maximum allowable loss (subtract line 10 from line 8)	9.0 dB
12. Total attenuation loss (multiply lines 1 and 5)	.3 dB
13. Total loss for connectors (multiply lines 2 and 6)	4.8 dB
14. Total loss for splices (multiply lines 3 and 7)	0.5 dB
15. Total loss for bypasses (multiply lines 4 and 8)	1.5 dB
16. Total loss (add lines 12 through 15)	7.1 dB
17. Remaining margin (subtract line 15 from line 11)	1.9 dB

Table 5-3. Simple Loss Calculation Method. After Jain, 1994.

intermediate location. (Jain, 1994)

The 99-percentile calculation method uses the average and standard deviations of component losses to compute a less risky calculation of total loss. First, the mean and standard deviation specifications for individual components are obtained from the manufacturers. Next, these values are used to determine the mean loss and variance for the different components and then the link, as shown in Table 5-4.

Once the mean loss and variance for the link have been determined, the 99-percentile link loss can be calculated. From statistics tables, the 99-percentile calculation is equal to $\mu + 2.326\sigma$, or 8.78 dB loss for this example. Thus, if using the transceivers described in Table 5-3, the remaining margin would be less than .3 dB--a risky design to say the least. (Jain, 1994)

Once the loss analysis is complete, the network designer should ensure each link meets the FDDI bandwidth specifications. This step is easily accomplished by ensuring the fiber's bandwidth capacity is at least 500 MHz-km, and the link length does not exceed 2 km. If these specifications are not met, an in-depth bandwidth analysis is necessary. An example of this analysis is available in Jain (1994).

To conduct these loss and bandwidth analyses, the network designer must make some preliminary choices regarding the type of transceivers, fibers, connectors, splices, and bypasses that will be used to construct the network. If the analyses indicate a particular link is questionable, he should reevaluate the link using components with better performance

Component	Component Loss			Total Loss	
	Mean μ_i	Std. Dev. σ_i	Quantity n_i	Mean $n_i \times \mu_i$	Variance $n_i \times \sigma_i^2$
Fiber (per km)	1.50 dB	0.20 dB	200 m	$0.2 \times 1.50 = 0.3$	$0.2 \times 0.20^2 = 0.01$
Connectors	0.80 dB	0.20 dB	6	$6 \times 0.8 = 4.8$	$6 \times 0.2^2 = 0.24$
Splices	0.50 dB	0.15 dB	1	$1 \times 0.5 = 0.5$	$1 \times 0.15^2 = 0.02$
Bypasses	1.50 dB	0.50 dB	1	$1 \times 1.50 = 1.5$	$1 \times 0.50^2 = 0.25$
Link				$\mu = 7.1$	$\sigma^2 = 0.52$

Table 5-4. Determining Losses Using Average and Standard Deviations. After Jain, 1994.

characteristics. This may be a cheaper alternative than installing another active node.

After the network designer is satisfied with his design, he prepares a hardware purchase list. Using his physical topology as a guide, he records the type and quantity of all components required to construct the network. This purchase list, along with copies of the logical, physical, and building blueprints, are retained as part of the system's documentation.

2. Copper to the Desktop Level

In some situations, an organization may find it an expensive proposition to extend fiber to the desktop. The increase in cost is not related to the fiber itself, but to the cost of optical transceivers. A less expensive alternative is to use copper.

Recall that the TP-PMD permits the use of UTP or STP to support FDDI communications between nodes. Although the maximum distance between nodes is more restrictive than fiber, it is a viable option in a workstation environment. A designer may consider building a fiber-based dual-ring to the concentrator level, then use UTP or STP to extend the network to the desktop. This hybrid design eliminates the need to purchase tens or perhaps hundreds of optical transceivers.

The drawback to building a hybrid solution is the complexity added to the network. Combining TP-PMD and MMF-PMD standards requires additional levels of hardware and software that complicates network maintenance, testing, and troubleshooting.

Network design considerations differ slightly from those presented in Table 5-2, when using copper media. First, due to the higher attenuation of copper media, the maximum distances between nodes is limited to 100 meters. This factor may change the physical topology of the network considerably.

Furthermore, the network designer must be cognizant of the susceptibility of UTP to EMI. EMI generated by electrical equipment such as florescent lighting can corrupt data transmissions. To avoid this problem, UTP cabling should be routed around EMI-rich environments; a better alternative is to use STP media.

Lastly, copper media is not as secure as fiber-optic media. It is easier to tap and monitor for electrical emissions. Therefore, it may be necessary to implement additional security measures to protect data transmissions along copper links. This may include the use

of secure cabling, secure wiring closets, and encryption devices, depending on the degree of security required.

3. Additional Design Issues

It is important to note this methodology is intended to serve as a guideline for designing a network; it is not a comprehensive solution to network design. There are other factors not included in these design steps that will influence the final layout. For example, the network designer must consider future growth when designing the infrastructure--failing to do so will only lead to an inadequate solution.

Designing a scalable network may be as simple as installing a patch panel mid-way along a link. Connectors at the patch panel interconnect the fibers to complete the connection. The advantage of such a design is that the cable has been prepared to accept additional nodes. Otherwise, the link would have to be severed, connectors installed, and the link tested before the nodes are added.

Furthermore, the analyses described earlier presumed that the same fiber type was used throughout the network. If this were not the case, additional losses due to adapters, as well as size and numerical aperture differences between fiber types, would have to be included in the loss analysis (see Chapter IV). Moreover, an in-depth bandwidth analysis would be necessary to ensure the link meets bandwidth specifications.

Finally, to emphasize a point made earlier, the design must give due consideration to providing the degree of security commensurate with the importance of the network and its data. Recognizing the fact that an organization's data may be its most important asset, failing to enforce adequate security measures may spell disaster. These security measures must go beyond merely protecting the data from unauthorized users; they must consider physical threats, both man-made and natural, that could render a network unusable.

E. SUMMARY

Organizations have learned the value of information resources planning. Designing and implementing network solutions no longer involves applying technology to simply automate business functions. It involves defining the problem in terms of the organization's mission, and developing a solution that achieves some strategic advantage.

There are numerous approaches to designing a networking solution. A systems approach includes a feasibility study, a requirements analysis, an evaluation of the project's geographical scope, the selection of a network standard or standards, and the design of the network. The decision to move beyond design and into implementation rests with the stakeholders.

From a practical standpoint, the most ambiguous step in the systems design approach is the feasibility analysis. It is often difficult to evaluate the intangible benefits and opportunity costs of a system in order to justify its acquisition. Likewise, the technical, operational, schedule, and political risks are equally difficult to assess. In many instances, however, the need for an in-depth feasibility analysis is a moot point. In the case of educational and military systems, the need for computers and networks is a given. The question is which network should be used.

Designing an FDDI network is similar to designing most other networks. There are, however, certain peculiarities that must be considered. One approach is to develop the logical topology of the network first. Documentation that is useful in completing this task include the requirements data, building blueprints, and a copy of the FDDI port connection rules.

Next, the physical topology diagram is prepared. To complete this task, it may be necessary for the designer to conduct a market survey to determine the types and configuration of FDDI products available. The physical topology includes all hardware components necessary to complete design of the network. This should include components required to support future growth and security requirements.

Once a preliminary physical design has been completed, a link analysis is conducted to ensure the links meet FDDI loss and bandwidth limitations. If the link fails to meet these restrictions, different components with better specifications may be used. Any changes are then incorporated into the logical and physical diagrams.

Upon completion of the diagrams, the hardware purchase list is prepared. This list documents all hardware requirements necessary to build the network. It is retained with the topology diagrams, as part of the system's documentation.

VI. FINDING A SOLUTION FOR THE SYSTEMS MANAGEMENT DEPARTMENT

A. INTRODUCTION

The Systems Management Department (SMD) is one of more than a dozen academic departments at the Naval Postgraduate School (NPS). The department consists of approximately 40 staff personnel and 75 professors and lecturers, responsible for managing the administrative functions and masters education for more than 450 students. In addition, the department maintains three LANs that are essential to the education and research endeavors of students and instructors alike.

This chapter takes a closer look at the importance of these LANs to the training and development of students. It provides an overview of these networks in terms of functionality and purpose, to build a case for justifying the introduction of an FDDI network. This justification is based largely on the training obligations of the SMD.

This chapter also sets the groundwork for developing a recommended solution. It addresses the importance of developing an evolutionary rather than a revolutionary solution, and presents the results of the baseline review. The actual solution will be addressed in Chapter VII.

B. INTRODUCTION TO SMD NETWORKS

The SMD's LANs include an integrated 16 Mbps token-ring network, a 10 Mbps Ethernet network, and a 230.4 Kbps AppleTalk network. The token-ring LAN is the mainstay of the SMD's networks. It supports 50 simultaneous users and spans two floors of Ingersoll Hall, with a reach extending to three computer labs, the Church Computer Center, and the campus backbone. This network is used extensively by students and instructors to support study and research efforts.

The token-ring network consists of user computers located in labs IN-224 and IN-250, and file servers located in lab IN-158. The user computers in IN-250 are used by students for application processing. Available applications include popular word processing, spreadsheet, and database programs. In addition, the computers can access specialized

programs used to support teaching objectives for a number of different courses.

The segment of computers in IN-224 is used to instruct students in the design, implementation, management, and maintenance of networks, through direct hands-on exposure. This lab is used strictly as a dedicated networks lab for instructing a myriad of network-related courses.

The Ethernet LAN is also located in lab IN-224. It is a small network consisting of five user computers, a server, and printer. Configured to support a number of word processing, spreadsheet, and database applications, it is also used primarily as a network design, maintenance, and management tool.

The AppleTalk network is located in lab IN-158. In conjunction with a number of IBM-compatible computers, it is used largely for testing and research purposes.

C. JUSTIFICATION FOR AN FDDI SOLUTION

One of the principal goals of the SMD is to expose students to the technologies they will encounter in follow-on tours. This is particularly important for students engaged in curricula focusing on computing and networking technologies. It is for this reason that the current mix of LAN technologies is maintained by the department.

However, these networks do not provide a comprehensive exposure to all technologies. In particular, they do not provide the exposure to fiber optic-based networks that are becoming more commonplace.

Fiber optic communications is becoming the media of choice. For example, telecommunications providers have realized the tremendous capabilities of fiber and are committed to replacing traditional long-haul, copper-based networks with synchronous optical network (SONET) standards.

SONET adds a new dimension to the speed and capacity of communications systems. The basic SONET building block, OC-1, can be multiplexed to an OC-48 level to achieve a transfer rate of 2.488 Gbps. The traditional pulse code modulation hierarchy used in copper-based communications can only achieve a maximum line rate of approximately 274 Mbps. (McClain, 1994; Freeman, 1991)

Furthermore, SONET can be designed with self-healing capabilities that permit ring

reconfiguration. Sprint, which is in the process of building 39 SONET four-fiber rings across the United States, demonstrated this capability when it cut the fiber lines of one of its SONET links. The failure was detected and corrected within 8 milliseconds. (Menefree, 1994)

These features of self-healing and tremendous capacity captured the attention of the Army, and led to the recent installation of a SONET-based backbone between Fort Bragg and Fort Hood, Texas. The backbone uses an OC-3 link capable of 155.52 Mbps. It was selected to carry asynchronous traffic mode (ATM) transmissions at a rate of 155 Mbps. The backbone interconnects two dozen routers that expand into LANs servicing 600 users. (Masud, 1994)

The intent is not to shift focus to SONET, telecommunications, or to emerging technologies like ATM; the intent is to focus on the underlying technology--fiber optics. The Army's installation of a SONET-based backbone is a clear example that there is a pressing need for students and faculty to keep abreast of this technology. Installing an FDDI network would enhance hands-on experience and exposure to fiber-based communications.

Moreover, the recent installation of FDDI aboard Navy and Coast Guard ships increases the likelihood that students will become responsible for managing these assets during future assignments. Therefore, it is essential they gain experience in the design, implementation, management, and maintenance of this technology. Such exposure is not possible through textbooks alone. Thus, the growing use of fiber in general communications fields and the Department of Defense's (DOD) adoption of FDDI standards presents two strong arguments in favor of finding an FDDI solution for the SMD's token-ring network.

D. EVOLUTIONARY VERSUS REVOLUTIONARY DESIGN

There are two approaches to network design. The first is based on a revolutionary solution: the target is defined and implemented as a wide-spread replacement for existing systems. It is used when a solution represents a significant departure from current business practices, such that delaying its development would only result in lost business opportunities. The second is an evolutionary design: the target solution is defined in stages that are implemented over time. This approach allows the organization to experiment with

new technologies through the various stages of project development.

The goal of a revolutionary design is to eliminate any unnecessary delays. It is commonly used when the business practices of an organization dictate the need for an effective, timely solution. Indeed, time itself is often the culprit behind failed systems. Over time, an organization's business functions, management strategies, and top-level personnel change. These changes cause project changes, delays, and even cancellations. Furthermore, in an era where network technology continues to grow at an accelerated rate, there is the potential that today's solutions will be overrun run by tomorrow's innovations. Under these circumstances, a revolutionary system may be the best answer.

There are however, situations where an evolutionary design is the preferred approach. In circumstances of limited staffing, limited budget, and limited experience with a particular technology, an evolutionary approach may be a better alternative. These are the conditions facing the SMD.

1. Technical Risk

An organization faces high technical risk when it lacks the in-house expertise to implement a technology. To offset this risk, the organization may hire skilled specialists, seek specialized training for its personnel, or perhaps purchase a microsystem to train and develop its in-house expertise. If the technical risk is deemed excessive, the organization may consider outsourcing rather than face the prospects of a failed system.

Within the SMD, there are only three individuals assigned to manage and maintain the department's LANs. Although these individuals have extensive experience in network design, implementation, management, and maintenance, they have not implemented nor managed an FDDI network. Regardless, their combined experience in dealing with network issues mitigate the technical risk of implementing an FDDI solution.

This risk, however, increases with network complexity. The larger the network the more difficult it is to resolve hardware and software compatibility issues. Thus, implementing a revolutionary FDDI network that represents a radical change in current technology requires experience working with the standard--experience that is currently lacking in the SMD.

This experience can be gained, however, by developing and experimenting with a small FDDI network. Over time, the network can be evolved into the final target solution. This approach allows network builders to gain valuable hands-on experience with the new standard--experience that is then applied to subsequent stages of the evolution.

2. Limited Staffing

Attempting to implement a revolutionary design for the SMD's token-ring network, would require taking the current network off-line. With only two people available to work on the target network, it could take weeks or perhaps months to install the hardware, program the software, and resolve compatibility problems. Bringing the network down for just one week has serious repercussions on student training.

Again, an evolutionary approach would mitigate the impact on student training. The current token-ring network can be segmented and incrementally evolved to FDDI standards. As each segment is implemented, students can be trained in the use of the new standard to minimize the affect on educational and research goals.

3. Limited Budget

Implementing an FDDI solution is an expensive proposition. Indeed, the cost of a 50 node network will run tens of thousands of dollars. Such an investment may not be possible on a limited budget.

In an evolutionary design, however, the solution is developed over time. Thus, the cost element is spread across several budget cycles easing the financial burden of an FDDI solution. Furthermore, this approach reduces the risk of purchasing components that are marginally compatible--a fact that will likely be discovered while experimenting with the technology during evolution.

Another advantage is that the stakeholders are not committed to the target solution. The evolution may be modified or even canceled if desired, minimizing the sunk costs. Purchasing all the requisite components up-front, however, generates considerable sunk costs.

4. Preferred Approach

Due to limited staffing, a limited budget, and a lack of experience with FDDI, the

recommended solution will be based on an evolutionary design. The recommended solution will consist of a series of stages intended to achieve a target solution. The implementation of each stage, as well as the final solution, will be left to the discretion of the stakeholders.

E. BASELINE REVIEW OF THE TOKEN-RING NETWORK

The token-ring LAN is an IEEE 802.5-based network that uses a series of multistation access units (MAUs) and shielded twisted-pair wiring to establish the ring topology shown in Figure 6-1. For network management purposes, the network has been divided into three segments: 0TR, 4TR, and 8TR. The numerical value of the designation corresponds to the last digit of the lab's respective room number.

1. Topology

The token-ring network is based on a logical ring topology that uses a token to control access to the media. The token travels the logical ring from one lab to the next, across Type 1 STP cabling that interconnects the MAUs. The MAUs, which contain a ring-in and ring-out port for accepting the STP cabling, interconnect the intra-lab computers.

The MAUs serve as multi-port hubs. They accept adapter cables that are connected to the adapter cards installed on client computers and servers. The MAUs change the logical topology of the network into a physical star configuration. They contain circuitry that detects the presence of operating computers, in order to connect them into the network ring. When computers are secured, the MAU reconfigures in order to segregate the computer from the rest of the network. This preserves the flow of data around the logical ring.

The maximum distance between computers and an installed MAU is approximately 20 feet. In circumstances where the adapter cable is of insufficient length to extend between the computer and MAU, a patch cable is used to complete the connection. This cable is also Type 1 STP media.

2. Integrated Network

The computers connected into the token-ring form an integrated network--any computer can communicate with any other computer, regardless of its location on the network. Network communications is managed by the Windows for Workgroups (WFWG) network operating system (NOS). This peer-to-peer NOS allows two dedicated Pentium




Figure 6-1. Systems Management Department's Token-Ring LAN.

servers to share their resources with the client computers, and any client to share its resources with any other client.

There are actually three Pentium servers connected to the ring through the 8TR segment. Only two of these servers, designated PN3 and PN6 in Figure 6-2, are configured with WFWG. They contain application files that are accessed by the clients throughout the network. The third server, designated PN9, is currently configured with the Windows NT NOS. Currently, it does not provide any functionality to the integrated network. The intent, however, is to eventually migrate the network from WFWG to Windows NT. This server provides the network managers an opportunity to become familiar with the peculiarities of this NOS.

To connect to the network, client computers must load certain communication drivers and NOS files. This task is accomplished at power-on through modifications made to their CONFIG.SYS and AUTOEXEC.BAT routines. These routines contain the necessary instructions to load the drivers, set default drives, load the NOS, and establish the computer as a user with the servers. In particular, there are two instructions within these files that designate the WFWG as a network version of Windows.

The computers in lab IN-250 are used by students for application processing. Their virtual drive designations describe paths that map to application files loaded on the servers. The computers in lab IN-224 are dedicated to networks training. These computers are used to illustrate network operating and management principles. They are configured to load default drives which can be changed through network connect and disconnect functions.

There are two printers connected to the network, one located in each of the student labs. These printers are managed by their individual servers. The client printer managers are normally configured such that each printer serves the computers located within the same lab. They may, however, be accessed by any network computer using network connect functions.

3. Token-Ring Segment 8TR

The token-ring segment in lab IN-158 interconnects three Pentium servers, a 486 server, and six user computers. Two Pentium servers, PN3 and PN6, are configured as WFWG file servers. The third server, PN9, is configured as a Windows NT server. Although




Figure 6-2. Token-Ring Segment 8TR.

it is connected to the network, client computers do not access this server. It is used as a test bed by the network staff to learn the Windows NT NOS prior to implementing it on the network.

The fourth server, TN4, is a carry-over from the network's recent upgrade to WFWG. This server used to function as a file server for DOS and Windows-based applications. It was accessed by the six user computers that are also interconnected to this segment. Five of these 486 MHz computers are still configured with old software that used TCP/IP to communicate

between user computers and servers. The remaining computer, TN10, has been configured with WFWG.

Furthermore, each user computer is fitted with an IBM 3270 terminal emulation card. This enables connection to the mainframe in the Church Computer Center, via coaxial cable, for mainframe terminal emulation. Though rarely used, this function can be demonstrated for instructional purposes. Two user computers are also equipped with modems.

4. Token-Ring Segment 4TR

As shown in Figure 6-3, the 4TR segment of the integrated network includes 17 client computers, a print server, and a printer. These computers are either 486 or Pentium machines that have been configured with WFWG NOS. They may be shared with other network resources.

Some client computers are configured with add-on devices. These peripheral devices include a scanner, an external CD-ROM, and modems. These devices are not shared across the network. They are not invoked by the WFWG NOS.

Nine computers are configured with user-version emulation software. This software is another carry-over from previous network configurations. The software was used by the computers to access the Gateway servers that are also connected to this segment. These servers and their 3270 Gateway-version software enable terminal emulation with the Church Computer Center's mainframe. This emulation is not currently used in the WFWG environment but may be invoked using DOS commands.

Lastly, MAU 3 is connected via STP to a Cisco router located in the Church Computer Center. This router is used to interconnect the token ring with the campus backbone. It allows communication with other campus computer systems, as well as off-campus systems through the Internet.

5. Token-Ring Segment 0TR

This segment consists of 24 client computers, a print server, and a printer, as shown in Figure 6-4. The clients and server are 486 computers configured with the WFWG NOS. These computers are used mainly for application processing. These applications are stored on the Pentium servers in lab IN-158. Several computers are configured with modems.




Figure 6-3. Token-Ring Segment 4TR.




Figure 6-4. Token-Ring Segment 0TR.

6. Cable Routing

The data grade STP cable begins in lab IN-158. It is routed into the false overhead where it rests on the suspended ceiling. From IN-158, it is routed along the overhead and through wall partitions, to a position over the computer center. It then passes through a vertical conduit to the second deck.

On the second deck, the cable is routed back into the false overhead where it extends to IN-224. Within this lab, the cable is routed to the first of three MAUs. From the third MAU, it is routed back into the overhead, across wall partitions, and into lab IN-250.

Within IN-250, the cable descends the wall to interconnect with the first of four MAUs. From the fourth MAU, the cable is routed back into the overhead, across the second deck false ceiling, and down through the same vertical conduit used to route the first cable. The cable extends back across the false ceiling to IN-158. Figure 6-5 illustrates this cable layout. The cable lengths between IN-158 and IN-224, IN-224 and IN-250, and IN-250 and IN-158, are 225 feet, 225 feet, and 300 feet, respectively.




Figure 6-5. Token-Ring Cable Layout.

7. Physical Security

Network security is enforced largely through physical access to the client computers and servers. IN-224 and IN-250 have cipher-locked doors that restrict access to these computer rooms. Access is limited to students and staff who have signed for the combination. A security log is maintained to record the distribution of combinations.

Within the student labs, additional physical security measures are in place. Computers, monitors, and servers have been physically secured to desks and walls using a myriad of cable-based devices. These measures are intended to discourage pilferage.

Access to lab IN-158 is restricted to a small number of staff personnel and research students. Class sessions are held in this lab but only under the direction of a professor. Thus, this lab provides a greater degree of access security than the other two labs.

F. SUMMARY

The introduction of FDDI-based networks aboard Navy and Coast Guard vessels increases the likelihood that NPS students will one day be responsible for managing these assets. Coupled with the growing trend of using fiber-based media in a myriad of telecommunications fields, these two points suggest a need for establishing an FDDI network in order to broaden students' exposure to LAN technologies.

As presented in this chapter, the baseline system is an integrated token-ring network that interconnects computing assets in three different labs. The computers on the network operate in a peer-to-peer configuration based on the WFWG NOS. One computer lab is dedicated to network training, another for student application processing, and the third as a test and research lab. The goal of this thesis is to prepare a solution for implementing FDDI protocols in place of this architecture, without impacting the current functionality of the baseline system.

VII. THE TARGET SOLUTION

A. INTRODUCTION

This chapter presents a solution for replacing the Systems Management Department's token-ring network. It begins with an overview of preliminary link design decisions followed by a recommended FDDI solution for interconnecting the token-ring stations. Logical topology diagrams are used to facilitate this discussion. Though consideration was given to constructing physical topology diagrams, it was deemed such diagrams would not be a true representation of the final infrastructure. To construct physical representations of the network, some decisions are necessary regarding which specific vendor products will be used—a proposition deemed premature at this stage.

Next, a number of compatibility issues, both hardware and software, are discussed. These issues can be mitigated through careful selection of compatible vendor products that achieve the degree of functionality and survivability desired in the network. Thus, to assist the network designer in his selection of vendor products, the numerous characteristics of various network components are discussed.

Finally, an evolutionary approach to achieve that solution is described. This approach is presented in a series of stages that increase in functionality and complexity with each evolution. Each stage is intended to mitigate the technological risk and cost of developing the network.

As is the case with any project, there are a number of different alternatives to consider. These alternatives provide different levels of network functionality and survivability; the tradeoff is usually cost. Several of these design alternatives are presented for consideration by the stakeholders.

B. PRELIMINARY LINK DESIGN DECISIONS

Table 7-1 outlines the preliminary link decisions for the target solution. These decisions were based on the results of the baseline review, FDDI restrictions, and desired network performance. Supporting justification for these decisions are presented next.

Inter-lab Links	
Media	62.5/125 μ m multimode fiber, NA = .275, 2.5 dB/km or better attenuation, 500 MHz•km bandwidth-distance product
Connectors	SC connectors at link ends; MIC connectors into nodes
Maximum length	2 km
Intra-lab Links	
Media	STP Type 1/2
Maximum length	500 m
Optical bypass	Not required

Table 7-1. Preliminary Link Decisions.

From the baseline review, it was determined that the maximum link distance between the three computer labs was less than 300 feet. At these short distances, many of the concerns associated with attenuation and link losses are mitigated. In fact, any of the FDDI PMD standards may be used in the target solution. In light of this fact, the decision was made to use a hybrid design for the link infrastructure. The network will use a combination of fiber for the trunk and twisted pair media to the desktop level.

The decision to incorporate twisted pair media is based on the high cost associated with extending fiber to the desktop. To illustrate this point, a cost comparison between fiber, STP, and UTP to the desktop level was conducted. This comparison used the products of a single vendor to interconnect 20 stations into a fiber trunk, via concentrators. The results revealed an 18% cost savings using STP compared to using fiber alone. Moreover, a 48% cost savings was possible using an UTP solution. This cost comparison is presented in the attached Appendix.

Although the UTP alternative can achieve greater cost savings, the tradeoff is its susceptibility to EMI. For this reason, STP is selected as the preferred media. Both media support 100 Mbps throughput, provided the distance limitations are not exceeded.

One argument against using either twisted pair options is the greater potential for unauthorized tapping compared to fiber. In this application, however, this issue is not a real concern. Since the STP media will only be used at the desktop level, it is protected within the secure confines of the lab. Although this does not prevent sabotage or deliberate tapping

by authorized users, it is deemed to be adequate. Meanwhile, the more vulnerable links between labs will consist of fiber.

Inter-lab links will use 62.5/125 multimode fiber as prescribed by FDDI standards. Although less quality fiber could be used given the short link distances, using this default type will ensure sufficient link budgets and bandwidth-distance products are available to support future growth possibilities. Furthermore, it eliminates the need to do in-depth loss and bandwidth analysis.

Nodes will be interconnected to the fiber trunk using MIC connectors. These connectors were chosen for their low insertion loss, ease of use, and station and polarity keying. These characteristics will help facilitate the numerous network changes that will occur during the target solution's development stages.

Since the inter-lab links will have to be routed through firewalls and conduits, it would be impractical to use cables with pre-fitted connectors. Therefore, SC type connectors will be installed after the cable has been routed. Extending the link into the trunk nodes will be accomplished using an SC-to-MIC patch cable.

An alternative to using the SC connector at the link ends is to splice a MIC pigtail cable directly to the link. However, considering the limited flexibility of making changes to this configuration, and the fact that considerable changes will occur during the evolution of the target solution, the use of SC-to-MIC patch cables is deemed to be the better option.

Lastly, considering the relatively small size of the network, only a handful of concentrators are required at the trunk-level to construct the tree topologies that represent the individual labs. Since it is unlikely that more than one concentrator will fail at any given time, the threat of a segmented network is extremely low. Thus, the use of optical bypasses is not deemed necessary.

C. TARGET SOLUTION

1. Trunk Design

The target solution uses a fiber-based trunk as shown in Figure 7-1. The fiber cables will contain a single fiber-pair to support the FDDI dual-ring topology. The fiber media will

connect a series of dual-attachment concentrators that expand into the tree topologies within each lab.

Though the figure only shows two concentrators within labs IN-224 and IN-250, the actual number will depend on the level of survivability desired. For example, if all the lab stations were connected into a single DAC, a failure in that DAC would bring the whole lab down. By increasing the number of DACs, the impact of a single DAC failure is mitigated.

However, there is a tradeoff associated with using a large number of DACs. First, adding more DACs to the trunk increases the cost of the network. Furthermore, a trunk constructed with a large number of nodes is more susceptible to node failure. Though ring reconfiguration preserves the trunk's functionality, there is still the increased potential for segmentation if other nodes should also fail. Two DACs are considered adequate for each of the student labs.

Figure 7-1 also shows that fiber cable will be used to connect stations to the DAC located in lab IN-158. The reason for using this media instead of STP is a result of the evolutionary design of the system. This reason will become more evident later.

2. IN-158 Tree Design

Figure 7-2 is an exploded view of the logical tree design for IN-158. The tree will consist of single-attachment stations representing the network servers, and a computer station, TN10. These stations are interconnected through DAC 1 onto the fiber trunk.

The reader will note that not all of the computers presently located in this lab are shown (see Figure 6-2). Only those stations that currently access the network have been included in the figure. If the remaining computers are added to the network, the logical topology may change considerably. Depending on the number of ports supported by the concentrator, an additional DAC or cascading SACs may be required to connect the remaining stations.

One concern regarding IN-158's design is the network's susceptibility to a failure in DAC 1. If this concentrator fails, the servers are isolated from the network. To reduce this susceptibility, the servers could be connected in a dual-homed configuration.




Figure 7-1. Trunk Design.




Figure 7-2. IN-158 Tree Design.

Recall that both the A and B ports of a dual-attachment station may be connected to the M-ports of a concentrator. This configuration, though valid according to FDDI protocols, results in a condition called prevent through—the flow of data between the ports is disrupted to prevent the creation of multiple rings. Some vendors implement this condition by placing the A-port in a stand-by mode while the B-port is active. Using special protocols, if the B-port fails the A-port is activated.

Dual homing takes advantage of these protocols by connecting the ports of a dual-attachment station into the M-ports on two different concentrators. During normal operations the B-port handles all data flow on the logical ring. Should the concentrator servicing the B-port fail, the A-port is activated and the station maintains logical connection through the second concentrator.

Figure 7-3 is an illustration of how the servers may be interconnected in a dual-homed configuration. The obvious tradeoff to this improved survivability is the cost of the




Figure 7-3. Dual-Homed Configuration for IN-158 Servers.

additional components required to implement this feature. It is interesting to note, however, that adding a second DAC in lab IN-158 may be necessary anyway, if the other computers in the lab are to be connected to the network.

3. IN-224 Tree Design

Figure 7-4 shows a logical topology solution for lab IN-224. The actual number of concentrators required will depend on their port configurations. Some concentrators can only connect eight stations, while others may connect up to 24 stations. The number of ports depends on the vendor and the type of media supported.

The concentrators in this lab use fiber-based ports to connect to the trunk, and STP-based ports to interconnect the stations. Such configurations are possible using concentrators that are based on a modular design. These concentrators contain a management module that controls the functions of the concentrator, as well as a number of empty slots that accept port modules. These port modules may be configured to accept fiber,




Figure 7-4. IN-224 Tree Design.

STP, UTP, or any myriad of permutations. Thus, the designer selects the appropriate media given his particular application, and builds the concentrator shell as required.

To maintain connectivity with the campus backbone, the Cisco router is shown connected to DAC 3. Though the logical connection is simple to illustrate, making the physical connection will prove to be more difficult. The STP cable that currently connects to the router uses a special connector tailored for MAU ports. An adapter cable will be required to connect this cable to the standard 9-pin D-connector used for FDDI STP ports. If an adapter cable is not available, the STP cable will have to be rerouted. If such is the case, it would be easier to connect the router into the concentrator located in lab IN-158 than to route the cable between floors and into IN-224. Moreover, the router will have to be reconfigured with the correct ports and reprogrammed to support FDDI protocols.

A potential flaw with this design is that a failure in DAC 3 will segregate the router from the network. This problem can be circumvented by configuring the router as a dual-attachment station. As a DAS it can be connected directly into the trunk as another node, or connected into DAC 2 as a dual-homed station. The drawback to these alternatives, however, is that a second cable is required to configure the router as a DAS.

Furthermore, a failure in DAC 3 will isolate the print server, PN1. To improve the availability of this printer, it could be configured as a dual-homed station by connecting it into DAC 2. However, considering the close proximity between DAC 2 and DAC 3, it is easier and cheaper to simply disconnect PN1 from DAC 3, and reconnect it into DAC 2.

4. IN-250 Tree Design

Figure 7-5 is a logical representation of the tree design for IN-250. The number of stations in this lab probably exceeds the STP port capacity for any two-concentrator configuration. Therefore, it was prepared with a single-attachment concentrator to illustrate how this segment may appear. Once again, the actual topology will depend on the products used in constructing the network.

5. Hardware Components

The different hardware components required to construct the network are outlined in Table 7-2. The quantity of each component required will depend on the capabilities of




Figure 7-5. IN-250 Tree Design.

the vendor products selected by the network designer, as well as the network's final design. Only FDDI-compatible products will be used.

Components

1. 62.5/125 μm multimode fiber
2. SC connectors
3. 62.5/125 MIC-to-SC patch cables
4. 62.5/125 MIC-to-MIC adapter cables
5. Type 1/2 STP cable with 9-pin D-connectors
6. Dual-Attachment concentrators capable of supporting both fiber and STP media
7. Single-Attachment Concentrators capable of supporting STP media
8. Adapter or network interface cards

Table 7-2. Target Solution Hardware Components.

As previously discussed, the fiber trunk will be constructed using 62.5/125 multimode fiber. These fibers will connect to the patch cables using SC connectors. The patch cables will connect into the trunk-level nodes using MIC connectors.

Four of the DACs will use 9-pin D-connectors to transmit signals from their M-ports into the STP media. The fifth concentrator, located in lab IN-158, will use fiber media to exchange data with the stations. Thus, some concentrators will be configured to handle two types of media.

Since the four computers in lab IN-158 will use fiber media to exchange data with their concentrator, fiber-based adapter cards will be required. Furthermore, in order to complete the physical connection into the concentrator, MIC-to-MIC adapter cables will be required. The remaining stations, including the Cisco router, will use STP-compatible adapter cards and media.

6. Software Components

The software elements required to support the operation of this target solution include software drivers, station management and access control software, and network management software. Software drivers are used to facilitate data exchange between the NOS and the FDDI protocols embedded in the adapter cards. These drivers are written to support specific NOSs in either a DOS or OS environment. They support most of the popular NOSs to include Novell Netware Workstation, Windows for Workgroups, and Windows NT.

They are normally included with the adapter cards as a package purchase.

Station management and access control software is used within the nodes to enable and manage SMT and MAC protocols. As was the case with software drivers, they are provided as a standard package when purchasing FDDI components. When purchasing software for concentrators, the network builder will discover that many vendors produce generic concentrators as opposed to dual- or single-attachment concentrators. The difference is a matter of programming. The vendor will supply the software required to program the concentrator as either a DAC or SAC. They will also discover that vendors use different versions of SMT software.

Network management software is used to support higher level processing of SMT's connection, ring, and frame base management features. This software allows network managers to monitor network counters and ring conditions, as well as control the operation of the network. Most support full remote operations from any station on the network using graphical user interfaces (GUIs). This software is available through FDDI product vendors, and is usually based on simple network management protocols (SNMP). It is particularly useful in monitoring network performance and troubleshooting problems.

D. COMPATIBILITY ISSUES

Perhaps the greatest challenge to implementing an effective FDDI solution is ensuring compatibility between the various software and hardware components that comprise the network. Although the purpose of developing the FDDI standards was to promote the manufacture of "open" products, these standards may be implemented in the form of hardware, firmware, or programmable software, at the discretion of the vendor. Thus, compatibility between products from different vendors is not always assured.

Furthermore, it is critical to ensure that FDDI hardware and software components are compatible with existing computer software and hardware configurations. For example, when purchasing adapter cards, the designer must ensure the card is compatible with the computer's bus configuration. Likewise, he must ensure the correct software drivers are available for the NOS of the target solution.

It is not unusual for a change in one software protocol to generate a rippling problem

through other software protocol stacks. Take for example, the communications protocols that were used on the baseline token-ring network, prior to installing the Windows for Workgroups NOS. The network used TCP/IP protocols to communicate between nodes within and external to the network. However, when the NOS was installed, it was discovered to be incompatible with the version of TCP/IP protocols in use. To enable the NOS, the TCP/IP protocols were disabled. Communications between computers are now accomplished using network station names instead of IP addresses.

A lack of TCP/IP protocols may present a problem in the target solution. Although the software drivers are available to enable data exchange between computers using the Windows for Workgroups or Windows NT NOS, many network management software products that use SNMP rely on TCP/IP protocols to coordinate management functions. Thus, the TCP/IP version will need to be updated if this software is to be enabled.

An alternative to updating the TCP/IP protocol stack is to load another set of protocols that are compatible with the network management software. Once again, however, the designer must examine the compatibility of this software with other protocols used throughout the network. This problem is further exacerbated when mixing products from different vendors.

Some of the problems of compatibility can be mitigated by carefully evaluating the market products available, prior to purchasing them. In fact, Jain (1994) suggests that prior to mixing products network designers should ask vendors if their products have been tested for compatibility with other vendors' products. There are important characteristics of each component in the network that must be examined before settling on a particular product.

1. Optical Cables and Connectors

As indicated earlier, the target solution will use 62.5/125 multimode dual-fiber cable to form the trunk of the network. This requirement, however, only identifies the minimum characteristics of the cable. There are other important characteristics that must be considered. These characteristics affect the cost of the cable, its installation, and its operation. Table 7-3 provides a summary of cable characteristics that must be considered.

1.	Manufacturer _____	Model/Part number _____
2.	Cable type: ☐ Plenum ☐ Riser ☐ Adapter ☐ Breakout	
	Tight buffered: ☐ Yes ☐ No ☐ Not applicable	
	Gel filled: ☐ Yes ☐ No ☐ Not applicable	
	Number of optical fibers: ☐ One ☐ Two ☐ Other _____	
3.	Fiber type:	
	Multimode 1300 nm: ☐ Yes ☐ No	
	Size: ☐ 50/125 ☐ 62.5/125 ☐ Other _____	
	Maximum attenuation _____ (dB/km)	
	Typical attenuation _____ (dB/km)	
	Minimum bandwidth distance product: ☐ 500 MHz•km ☐ Other _____	
4.	Preconnectorized: ☐ Yes ☐ No	
	Type of connectors: ☐ MIC ☐ ST ☐ SC	
5.	Outer diameter of the cable _____ (inches/millimeters); ☐ Not applicable	
6.	Weight of the cable _____ (lbs/ft or kg/km); ☐ Not applicable	
7.	Minimum bend radius of the cable _____ (inches/centimeters); ☐ Not applicable	
8.	Maximum tensile load of the cable _____ (lbs/Newtons); ☐ Not applicable	
9.	Spool quantity _____ (ft/spool or km/spool); ☐ Not applicable	
	Cost:	
	Price per spool _____; ☐ Not applicable	
	Shipping cost _____	

Table 7-3. Fiber Cable Considerations.

Similar cable considerations should be applied to the purchase of STP Type 1/2 cable. The important point is to ensure the product meets the performance requirements for the network and is compatible with other components. In the case of STP media, FDDI standards specify the use of 9-pin D-connectors. It is incumbent on the network designer to ensure the selected STP cabling is fitted with these connectors.

Likewise, the designer must ensure the correct connectors are used with the fiber media. For example, when purchasing SC connectors, the designer must specify the size of cable that will be used. This is necessary to ensure the connector's ferrule will accept the fiber-strand within tolerance limits, and that the connector's backshell is able to clamp onto the outer jacket of the cable. (Jain, 1994)

2. Adapter Cards

There are a number of adapter card characteristics that must be considered. The card must be compatible with the host computer's bus architecture, the media type, and the connector type. Furthermore, the card's software drivers must be compatible with the NOS,

its SMT version compatible with other component SMT versions, and its programming compatible with network management protocols. (Jain, 1994; Espiritu, 1991)

Adapter cards are available in a number of different configurations. Some cards are designed for single-attachment applications only, while others are capable of supporting dual-attachment functions. All cards are capable of asynchronous data transmissions; synchronous transmission is an optional feature that must be specified if desired. Likewise, dual-homing is a special feature that is not supported by all vendors.

Performance is another issue that must be considered. Even though the standard specifies a throughput of 100 Mbps, the actual throughput may fall considerably short. The performance varies with the card's processor speed and memory. It is also a function of the host computer's speed and the speed of the software driver. (Jain, 1994)

Jain (1994) cautions against comparing throughput rates based on higher layer protocol measurements. For example, a touted 50 Mbps TCP/IP throughput rate is a function of the processor used in the host as well as the adapter card. If the vendor measured that rate using a 50-million-instructions-per-second (MIPS) processor, that same rate will be considerably different using a 25 MIPS processor.

It is also important to determine the transceiver characteristics of the adapter card. This is necessary to ensure that there is sufficient loss budget to support the link distances specified in the network design. Although FDDI standards specify certain minimum performance requirements, the actual components used in a card may not provide the necessary levels of signal power and sensitivity. Thus, despite the vendor's claim his product meets ANSI X3T9.5 standards, the designer may discover that an STP adapter card will only support a link distance of 300 meters, instead of the specified 500 meters. (Jain, 1994)

In summary, there are numerous adapter card characteristics that must be considered before settling on a particular vendor's product. Most of these characteristics are outlined in the Table 7-4. At a minimum, the network designer must ensure the product is compatible with the hardware and software configurations used in the host computers and across the network. Furthermore, he must consider the capabilities of the card in terms of desired functionality and performance.

3. Concentrators

Like their adapter card counterparts, concentrators are also available in a number of different configurations. To ensure compatibility with other system components, the network designer must verify that the concentrator supports both the network media and software protocols. Moreover, he must ensure the concentrator can be configured to perform its intended functions. Table 7-5 is a compilation of the minimum characteristics that should be considered when comparing one product against another.

1.	Manufacturer _____	Model/Part number _____
2.	System bus interface: ☐ 16-bit AT/ISA ☐ 32-bit EISA ☐ Other _____	
3.	Number of attachments per adapter: ☐ One (SAS) ☐ Two (DAS)	
4.	Number of MACs: ☐ One ☐ Two	
5.	Media Type _____	Connector type _____ Maximum distance _____ Loss budget _____
6.	Driver software: Processor _____ Operating system _____ NOS _____	
7.	DMA options _____ IRQ options _____	
8.	Dual homing: ☐ Supported ☐ Not supported ☐ Not required	
9.	Optical Bypass: ☐ Standard ☐ Optional ☐ Not available ☐ Not required	
10.	Multicast table size _____ addresses	
11.	Buffer size _____ Mbytes	
12.	Synchronous transmission supported: ☐ Yes ☐ No	
13.	Network management: ☐ Remotely manageable via: ☐ In-band signaling ☐ Out-of-band signaling	
	SMT versions supported _____	
	Management protocol: ☐ SNMP ☐ Other _____	
	Management Information Bases (MIBs) supported _____	
14.	Upgradability: ☐ Down line-loadable software ☐ Down line-loadable firmware	
15.	Performance:	
	Transmit throughput:	Peak _____ packets/sec _____ packets/sec
		Sustained _____ packets/sec _____ packets/sec
	Receive throughput	Peak _____ packets/sec _____ packets/sec
		Sustained _____ packets/sec _____ packets/sec
	Delay _____ μ sec	
	CPU utilization at peak throughput _____ %	
	Protocol layer at which the performance was measured: ☐ Data link ☐ UDP/IP ☐ TCP/IP	
	☐ Other _____	
	System configuration used for performance measurement: CPU MIPS _____	
	or CPU Systems Performance Evaluation Cooperative (SPEC) mark _____	
16.	Card size _____	
17.	Power requirements _____	
18.	Price _____	
19.	Warranty period _____	

Table 7-4. Adapter Card Considerations. After Jain, 1994.

Some vendors build concentrator shells that may be programmed as NACs, SACs, or DACs. Most concentrators use a modular design--the concentrator accepts port modules that are configured to support different types of media and their specific connectors. The designer must ensure these concentrators can be configured with the appropriate SMT and MAC protocols to achieve the degree of functionality required. The SMT version, in particular, must be compatible with other SMT versions used throughout the network.

The use of MAC protocols in concentrators is optional. These protocols are required in stations that transmit and receive onto the media. Since a concentrator does not transmit data onto the network per se, a MAC is not required--it can still perform all of its port connection functions using its PHY protocols. However, a MAC-less concentrator is not

1.	Manufacturer _____		Model/Part number _____		
2.	Number of peer attachments: ☐ NAC ☐ SAC ☐ DAC				
3.	Number of port modules: _____				
4.	Total number of M-ports: _____				
5.	Types of ports:				
	Media Type	Ports per Module	Media Size or Category	Connector Type	Maximum Distance
	Multimode Fiber	_____	_____	_____	_____
	STP	_____	_____	_____	_____
	UTP	_____	_____	_____	_____
6.	Number of MACs _____				
7.	Dual Homing: ☐ Supported ☐ Not supported ☐ Not required				
8.	Optical Bypass: ☐ Standard ☐ Optional ☐ Not available ☐ Not required				
9.	Network management: ☐ Remotely manageable via: ☐ In-band signaling ☐ Out-of-band signaling				
	SMT version _____				
	Management Protocol: ☐ SNMP ☐ Other _____ MIBs supported _____				
10.	Upgradability: ☐ Down line-loadable software ☐ Down line-loadable firmware				
11.	Availability: MTBF _____ MTTR _____				
12.	Connection rules: _____				
13.	Graceful insertion: ☐ Yes ☐ No				
14.	MAC-level link confidence testing (LCT): ☐ Yes ☐ No				
16.	Physical dimensions: Height _____ Width _____ Depth _____ Weight _____				
17.	Enclosure: ☐ Desktop ☐ Rack mounted				
18.	Base price _____ Price per expansion module: Fiber _____ STP _____ UTP _____				
19.	Total price _____				
20.	Power requirements _____				
21.	Warranty period _____				

Table 7-5. Concentrator Considerations. After Jain, 1994.

capable of supporting SMT frame-based management since its ability to transmit and receive information frames is inhibited. Thus, to ensure full functionality of SMT functions, the concentrators should be programmed with at least one MAC. This is particularly true if managing numerous concentrators from a remote management station. (Jain, 1994)

To ensure compatibility between different vendor products, the network designer must determine how vendors handle undesirable, yet valid, port connections. The reader will recall that certain port connections result in the creation of twisted or multiple rings. Although FDDI protocols identify these connections as questionable, the manner in which they are handled is left to the vendors. Thus, one vendor's solution to handling these misconnections may be incompatible with another vendor's solution. (Jain, 1994)

Vendors often provide higher-level management programming with their concentrators. Designers should evaluate the capabilities of these network monitoring and management tools before deciding on a particular product. Moreover, they should ensure the software protocols required to support the tools are compatible with existing and proposed protocols.

E. EVOLUTIONARY DEVELOPMENT OF THE TARGET SOLUTION

The implementation of the target network has been separated into five stages. The target solution begins as a simple two-station network, then evolves into progressively more complex configurations until the target solution is achieved. This evolution may progress over a period of one to two years if necessary. The actual period required to implement the solution will depend on the workload of the lab staff, their experienced gained working with the standards through each stage of development, budgetary restrictions, and the final degree of target solution desired.

1. Stage 1: Two-Node Network

This first stage will expose the lab staff to some of the hardware and software components of an FDDI network. At a minimum, the lab staff will gain experience in configuring dual-attachment stations, have an opportunity to test the compatibility of software drivers with the other software components, and gain exposure to fiber-based adapter cards and media. The micro-network will be constructed in lab IN-158 as opposed

to the other labs, to minimize the disruption to student learning.

To build the network, two computers configured with Windows for Workgroups or Windows NT NOS will be required. Other software, such as TCP/IP protocols, should be loaded to test the interoperability of these components with the software drivers. The computers may be 486 or Pentium processor machines. Two of the idle computers currently connected to the token-ring segment within the lab would be ideal for this purpose.

The minimum FDDI hardware components required to connect the computers include four adapter cards and two MIC-to-MIC adapter cables, as shown in Figure 7-6. Although adapter cards are available in either single-port or dual port configurations, a single port will suffice for this stage. In fact, using four single-port cards is preferred over two dual-port cards, since these cards will be used later to connect the pentium servers to the network as optical-based SASs. Therefore, if a cost-savings is to be realized, the adapter cards used to construct this stage of the network should be single-port and bus-compatible with the Pentium servers.

Furthermore, only one MAC per card is required to control access to the media, and dual-homing and synchronous capabilities are not necessary. The software drivers should support both Windows for Workgroups and Windows NT. The cards should contain the




Figure 7-6. Stage 1 of Network Evolution.

latest version of SMT software to preclude the purchase of an outdated product. Finally, the adapter card should support SNMP protocols. These protocols are commonly used by network management programs. Table 7-6 summarizes the network components required to implement this stage of the network.

2. Stage 2: Introduction of a Dual-Attachment Concentrator

The two-node network described in Stage 1 represents a simple dual-ring topology network. In Stage 2, the dual-ring will be expanded into a ring with one tree by introducing a dual-attachment concentrator and single-attachment station as shown in Figure 7-7. The staff will gain experience in configuring a dual-attachment concentrator, have an opportunity to implement and test a network management product, and gain exposure to copper-based adapter cards and media. Once again, this network will remain within IN-158 to reduce the impact on student learning.

The additional components required to complete this stage of evolution is summarized in Table 7-7. A major goal of this stage is provide the lab staff an opportunity to gain experience in configuring a DAC. This experience will prove useful when multiple

<u>Component</u>	<u>Characteristics</u>
4 - Adapter cards	Hardware: - Media - single-port, STP compatible - Connector - 9-pin D-connector - Bus configuration - 16-bit ISA or 32-bit EISA as appropriate Software: 1 MAC required - Dual-homing not required - Synchronous transmission not required - SMT version compatible with other versions - SNMP compatible - Software drivers for Windows for Workgroups and Windows NT
2 - Client computers	Hardware: 486 or Pentium processor Software: - Windows for Workgroups or Windows NT NOS - Application files as required to exercise data transmission
2 - Fiber adapter cable	62.5/125 multimode with MIC connectors

Table 7-6. Stage 1 Component Requirements.




Figure 7-7. Stage 2 of Network Evolution.

DACs will have to be implemented to form the tree topologies within the labs.

Another major goal is for the staff to have an opportunity to implement and use a network management tool. Learning to monitor and control the network will also be useful in later stages when the network takes on more complex configurations. To implement this tool, the staff will need to ensure the nodes are programmed with the proper software protocols. Since SNMP-based tools usually rely on TCP/IP protocols, the version of TCP/IP must be updated in all computers. The new version must be compatible with the NOS, the software drivers, and the network management software.

Table 7-7 shows that only one copper-based M-port is required in the concentrator. For upgradability purposes, however, the full port-capacity of this node must be evaluated before selecting a product. This concentrator will be used later to connect a dozen or more client computers.

3. Stage 3: Installation of the Fiber Trunk

This stage of the evolution will involve the installation of the fiber trunk between the labs. The fiber cable will be routed along the false overheads, through firewalls, and through conduits in the same fashion as is currently used in routing the token-ring media. Figure 7-8

<u>Component</u>	<u>Characteristics</u>
1 - Dual-attachment concentrator	Hardware: Media - fiber-based A- and B-ports; 1 copper-based M-port Connectors - MIC and 9-pin D-connector Optical-bypass not required Software/Firmware: 1 MAC required Dual-Homing not required SMT version compatible with adapter card version SNMP compatible for network management software Network management software
1 - Adapter card	Hardware: Media - single-port, STP compatible Connector - 9-pin D-connector Bus configuration - 16-bit ISA or 32-bit EISA as appropriate Software: 1 MAC required Dual-homing not required Synchronous transmission not required SMT version compatible with other versions SNMP compatible Software drivers for Windows for Workgroups and Windows NT
1 - Client computer	Hardware: 486 or Pentium processor Software: Windows for Workgroups or Windows NT NOS TCP/IP version compatible with network management software and NOS Application files as required to exercise data transmission
1 - STP adapter cable	STP Type 1/2 with 9-pin D-connectors
1 - Fiber adapter cable	62.5/125 multimode with MIC connectors

Table 7-7. Stage 2 Component Requirements.

illustrates this physical layout of the cabling between the labs.

This stage will begin with the selection of an appropriate cable type. As shown in Table 7-3, there are several types to consider. For this application, a plenum cable consisting of two fibers will be adequate. The cable should contain 62.5/125 multimode fibers. Its attenuation characteristic should be at least 2.5 dB/km or better and provide a 500 MHz•km bandwidth-distance product.

Care must be exercised when installing the cable. If the pulling tension exceeds the cable's tensile strength while pulling it through conduits, the fibers will likely be damaged.




Figure 7-8. Stage 3: Fiber Cable Installation.

To prevent this, a tensiometer may be used during cable installation. Furthermore, Jain (1994) recommends against pulling fiber cables through conduits whose cross-sectional area is more than 50% full.

The current condition of the vertical riser between floors in this application presents a problem. At nearly 80% capacity, it is unlikely that the cable can be pulled without damaging the fibers. Fortunately, 70% of the current cabling is old coaxial cable due for removal. This removal will have to be completed before attempting to install the new cable.

Given the difficulty of attempting to route a cable fitted with connectors, the cable will be installed first, then the connectors installed on the fiber ends. It is possible to splice MIC pigtails to the cables' bitter ends rather than using connectors. In fact, in many cases this method is preferred due to the reduced insertion losses compared with using connectors. However, splicing requires special tools and a greater level of expertise to complete. Considering the short link distances in this application, the use of SC connectors will be more than adequate.

Once the cable has been installed, a test of the optical fibers should be conducted. One approach to completing this test is to simply connect two optical-based stations to the link and determine if they are able to exchange data. Unfortunately, this operational test does not reveal any damage to the optical fibers that may have occurred during installation. In an operational test, the stations may be able to exchange data despite damage, provided the additional losses do not exceed the required budget margin. Over time, however, the damage may degrade even further, rendering the link unusable.

To effectively test the fiber, an optical time-domain reflectometer (OTDR) should be used. This tool launches light into a fiber, and then measures the reflection of light energy caused by fiber imperfections and breaks. These measurements are used to determine the fiber's length, attenuation losses, splice losses, and connector losses. It produces a signal strength versus distance plot of the link similar to that shown in Figure 7-9. (Jain, 1994)

The network builder uses the plot to evaluate the performance of the link. The attenuation of the fiber can be determined by measuring the slope of the plot; other losses are evaluated by measuring the vertical height changes in signal strength, where ever spikes

occur. These spikes are the result of back-scattered energy. Their location along the fiber is determined by referencing the horizontal axis of the plot.

The network builder should remember the plot is a function of the transceiver used in the test set. The actual signal strengths and losses for the link can only be determined if the same transceiver to be used on the link is also installed in the test set. Nonetheless, the test results will identify potential problems along the link.

4. Stage 4: Migration of the Computer Labs

Once the fiber cable has been tested, software compatibility issues resolved, and the




Figure 7-9. Sample OTDR Plot.

lab staff feels comfortable with implementing dual-attachment concentrators, the next stage is to build the tree roots for the labs. This will entail configuring another concentrator and moving some of the components from Stage 2 into labs IN-224 and IN-250. The goal is to establish the fiber trunk and trees of the target solution before attempting to connect the client computers into the concentrators.

Figure 7-10 illustrates this stage of the evolution. The DAC, along with its SAS, will be removed from lab IN-158 and connected to the fiber trunk in lab IN-224. The connection will be completed using MIC-to-SC patch cables. In addition, DAS 2 will also be moved to IN-224. This will facilitate the testing and troubleshooting of the network from the second deck, rather than relying solely on DAS 1 in lab IN-158. Once the target solution has been achieved, the remote networking functions should be restricted to stations within IN-158. This will prevent unauthorized personnel from changing the network operation.

Figure 7-10 also shows the introduction of a second concentrator, DAC 2, and another station, SAS 2. This concentrator, like the first, will be connected to the trunk using MIC-to-SC patch cables. It will support STP cabling to the station using 9-pin-compatible M-ports. Table 7-8 summarizes the components required for this stage of development.

5. Stage 5: Target Implementation

This final stage of development will achieve the target solution described earlier. This stage is perhaps the most critical in the sense that multiple concentrators must be configured and the computer stations in each lab must be migrated to FDDI standards. It is the one stage with the greatest potential for interrupting student training. Therefore, it must be planned and executed accordingly.

There are three steps to this evolution that must be accomplished: the servers must be configured and connected via a DAC onto the ring, the correct configuration of DACs and SACs must be implemented for each lab, and the individual computers must be connected into their respective concentrators. To reduce the impact on student training, the recommended approach is to keep the token-ring network operational until these steps are completed. This provides students with some, if not all, the available computing resources to continue their information processing and research endeavors.




Figure 7-10. Stage 4 of Network Evolution.

<u>Component</u>	<u>Characteristics</u>
1 - Dual-attachment concentrator	Hardware: - Media - fiber-based A- and B-ports; 1 copper-based M-port - Connectors - MIC and 9-pin D-connector - Optical-bypass not required Software/Firmware: 1 MAC required - Dual-Homing not required - SMT version compatible with adapter card version - SNMP compatible for network management software - Network management software
1 - Adapter card	Hardware: - Media - single-port, STP compatible - Connector - 9-pin D-connector - Bus configuration - 16-bit ISA or 32-bit EISA as appropriate Software: 1 MAC required - Dual-homing not required - Synchronous transmission not required - SMT version compatible with other versions - SNMP compatible - Software drivers for Windows for Workgroups and Windows NT
1 - Client computer	Hardware: 486 or Pentium processor Software: - Windows for Workgroups or Windows NT NOS - TCP/IP version compatible with network management software and NOS - Application files as required to exercise data transmission
1 - STP adapter cable	STP Type 1/2 with 9-pin D-connectors
6 - MIC-to-SC patch cables	62.5/125 multimode with MIC and SC connectors

Table 7-8. Stage 4 Component Requirements.

Beginning in lab IN-158, the servers will be reconfigured so that these resources are split between the two networks. If necessary, a fourth server may be brought on-line until the network transition is complete. To interconnect the servers into the FDDI fiber trunk, a DAC will need to be configured first. Once the servers are operational, they can be tested for functionality.

Next, the concentrators are prepared for both labs. Based on the vendor products selected to construct the physical network, the appropriate configuration of DACs and SACs

are programmed and interconnected into the trunk. Once the concentrators have been configured, the actual migration of computer stations can begin.

To reduce the impact on students, the transition of lab IN-224 should be conducted during an off-quarter when the network courses are not offered to students. During this quarter, all of the computers within the lab can be configured with the appropriate adapter cards and software, and migrated to FDDI. Once this lab has been converted to FDDI, it may be used as the principle lab until lab IN-250 has been upgraded to FDDI.

Once the labs have been upgraded, the servers should be reconfigured to the FDDI trunk as appropriate. Furthermore, the remote station used to monitor and control the network should be logically transferred to a station in lab IN-158. This will prevent unauthorized personnel from tampering with the network operation.

F. SUMMARY

This chapter presented an evolutionary solution for replacing the Systems Management Department's token-ring network with FDDI technology. Key to this design was minimizing the technical risk, financial risk, and impact on student learning. Although there are numerous network configuration alternatives to consider, the recommended solution will achieve the degree of functionality and survivability required for this particular application.

To ensure compatibility between the various FDDI software and hardware components, the network designer should carefully survey market products. This alone, however, does not guarantee the products will be compatible with each other nor with the current baseline system. These incompatibilities may only be discovered as components are implemented. Once discovered, perseverance is necessary to identify and resolve the problem. These incompatibilities must be resolved before attempting to implement subsequent stages of the evolution.

VIII. CONCLUSION

Implementing an FDDI network can be a costly and risky proposition. Indeed, a fiber-to-the-desktop upgrade to the Systems Management Department's token-ring network could cost over two hundred thousand dollars. Furthermore, implementing an FDDI network requires considerable expertise due to peculiar characteristics of this standard.

The economic and technical risk in this particular application can be mitigated through a carefully planned, evolutionary network design. Evolutionary stages of development as suggested in Chapter VII provide a number of advantages over a revolutionary design. First, a full-fiber solution is not necessary to achieve the advertised response time and throughput rate of an FDDI network; copper-to-the-desktop below the concentrator level is more than adequate. However, even a copper-to-the-desktop solution can be expensive. This economic risk can be mitigated by evolving the network in several stages. Such a design plan carries the network cost across several budget cycles, making the proposed solution less of a financial burden.

Furthermore, an evolutionary design does not commit the stakeholders to a final solution. As the network evolves, changes can be made to either enhance or scale back the design as desired. Recall that the justification for implementing an FDDI network was to improve the students' exposure to different networking technologies. This objective could easily be achieved through a micro-network consisting of just a few nodes. The recommended solution supports the development of such a network.

Moreover, an evolutionary design plan provides an opportunity for the SMD staff to gain experience with implementing and maintaining an FDDI network before attempting to install the final solution. Installing hardware components, programming concentrators, and troubleshooting faults can be both challenging and demanding. This poses considerable technical risk, particularly when implementing a complex solution with limited staffing, as in this case.

Perhaps more important, however, is that the proposed evolutionary design provides ample opportunity for the SMD's staff to identify and resolve software incompatibilities.

Although requisite software drivers are available from FDDI vendors, these drivers are not found among the driver lists contained in common network operating systems, such as Windows for Workgroups, Windows 95, or Windows NT. The presence or lack thereof is often a practical test of a software product's availability and maturity. This is not to suggest that FDDI is an immature standard, but only that the absence of the software drivers in these OSs may lead to compatibility problems between existing software, new operating systems, and the software drivers. This problem is exacerbated when mixing vendor products that use different software solutions to invoke their functions.

Furthermore, this evolutionary design provides the staff with ample opportunity to evaluate the existing software protocol stacks with the software drivers provided by vendors. As indicated earlier, incompatibilities between the TCP/IP protocols and the NOS must be resolved if TCP/IP protocols will be used to implement the full station management functions inherent in FDDI protocols. Once the compatibility issues between the current NOS and TCP/IP have been resolved, the software stack must be evaluated against the FDDI drivers. Once the staff is satisfied that their software products achieve the necessary degree of compatibility, the servers must be programmed, network operations enabled, and the network tested for share violations, timing glitches, and performance. This testing requires both time and experienced network builders. An evolutionary network development plan facilitates this testing process and builds the staff's exposure to FDDI standards.

Thus, implementing an FDDI network, such as this upgrade to a token-ring LAN, poses considerable economic and technical risk. These risks can be reduced by evolving the network over a period of time. In fact, this process is similar to prototyping a network before committing to a final solution. The benefits include increased exposure to the standard, ample opportunity to conduct hardware and software compatibility testing, and a flexible final solution. Such benefits can improve the probability of developing a successful LAN that is tailored to the needs of the organization.

APPENDIX. COST COMPARISON BETWEEN FIBER, STP, AND UTP

The following is a cost comparison between implementing a fiber, STP, and UTP solution to the desktop level. This comparison is based on the cost of products required to interconnect 20 workstations into a fiber trunk, using dual-attachment concentrators. The figures were compiled using a single vendor's products. Since the vendor used in this evaluation does not offer STP or UTP media, conservative estimates were used to represent the cost of this media.

Fiber-to-the-Desktop:

<u>Required components</u>	<u>Cost</u>
2 - DACs (each DAC contains three modules that are configured with 4 fiber ports each)	\$26,550
20 - FDDI fiber-based adapter cards	\$34,900
20 - MIC-to-MIC adapter cables	<u>\$4,800</u>
Total:	\$66,250

STP-to-the-Desktop:

<u>Required Components</u>	<u>Cost</u>
2 - DACs (each DAC contains two modules configured with six STP ports each; and one module with four fiber ports)	\$26,150
20 - FDDI STP-based adapter cards	\$25,800
STP cabling	<u>\$2,500</u>
Total:	\$54,450

UTP-to-the-Desktop:

<u>Required Components</u>	<u>Cost</u>
1 - DAC (contains two modules configured with eight UTP ports; and one module configured with two fiber ports and five UTP ports)	\$13,875
20 - FDDI UTP-based adapter cards	\$19,900
UTP cabling	<u>\$500</u>
Total:	\$34,275

LIST OF REFERENCES

- Espiritu, R.V., *Local Area Network (LAN) Compatibility Issues*, Naval Postgraduate School, Monterey, CA, 1991.
- Fitzgerald, J., *Business Data Communications: Basic Concepts, Security, And Design*, John Wiley & Sons, 1993.
- Freeman, R., *Telecommunication Transmission Handbook*, John Wiley & Sons, Inc., 1991.
- Hammar, G.A., *FDDI Installation And Performance Analysis*, Naval Postgraduate School, Monterey, CA, 1992.
- Hewell, J., Lewis, T., *Onboard LAN*, Marine Electronics, April 1994.
- Jain, R., *FDDI Handbook: High-Speed Networking Using Fiber And Other Media*, Addison-Wesley Publishing Company, 1994.
- Joshi, S., *High-Performance Networks: A Focus on Fiber Distributed Data Interface (FDDI) Standard*, IEEE Micro, 1986.
- Masud, S., *ATM Backbone For Army LANs: Fort Bragg And Fort Hood Take Lead In Everyday ATM Networking*, Government Computing News, August 1994.
- McClain, G., *Handbook Of Networking & Connectivity*, Academic Press Inc, 1994.
- Menefree, C., *Sprint's Self-Healing Fiber Optic Technology*, Newsbytes, November 1994.
- Mills, A., *Understanding FDDI: A 100 Mbps Solution For Today's Corporate LANs*, Prentice Hall International (UK) Limited, 1995.
- Raynovich, R., *Is The Time Right For Fiber? Fiber Is Becoming An Affordable Alternative To Category 5 UTP*, LAN Times, April 1995.
- Schivley, M.A., *Throughput Analysis Between High End Workstations Across an FDDI Network*, Naval Postgraduate School, Monterey, CA, 1994.
- Shah, A., Ramakrishnan, G., *FDDI: A High Speed Network*, PTR Prentice-Hall, Inc., 1994.

INITIAL DISTRIBUTION LIST

	No. of copies
1. Defense Technical Information Center 8725 John J. Kingman Rd., STE 0944 Ft. Belvoir, VA 22060-6218	2
2. Dudley Knox Library Naval Postgraduate School 411 Dyer Rd. Monterey, CA 93943-5101	2
3. Professor Norman F. Schneidewind Department of Systems Management, Code SM/Ss Naval Postgraduate School Monterey, CA 93943-5103	1
4. Professor James Emery Department of Systems Management, Code SM/Ey Naval Postgraduate School Monterey, CA 93943-5103	1